

Paper 373-2012

The Top Five Migration Considerations for Success

Andrew Mott, SAS Institute Inc., United Kingdom

ABSTRACT

There are different ways to move house; you can empty the contents of your house into a truck and transport your belongings to a new house. Or, put your house with all its belongings on the back of a truck and relocate the house. Many people have performed the former, and some the latter; everyone has a different experience but a similar set of considerations. The same is true of migrating a SAS[®] environment. There is a large number of SAS users who have migrated to SAS[®] 9.2, and some to SAS[®] 9.3, so there is a wealth of experience upon which to draw.

This paper reviews a number of customer experiences to draw out the top five considerations for a successful migration.

INTRODUCTION

There are many things to consider when migrating an existing SAS environment to a new SAS environment. For example, questions like “How should we approach the migration?”, “How long will it take?”, and “What should we do to prepare?” are some good open questions that have no immediately obvious or simple answers.

These questions are focused around the narrow viewpoint of the SAS administrator, but the answers impact the whole SAS user community. From the perspective of the IT function within an organization, a migration might seem like the ideal time to perform some form of environmental upgrade. After all, why not migrate that SAS environment that is constantly using all the available resources to a bigger platform or to one that complies with the latest IT standards of the organization?

When an organization considers a migration, there are typically many stakeholders that can influence the process, each influencing the technical approach in their own way.

Finally, it's worth noting that for the purposes of this paper, the term migration is defined as the act of moving content from a SAS environment that is based on a specific release of SAS software to a more recent release of SAS software.

MIGRATION DEFINITIONS

From a technical perspective, there are two migration pathways one can follow when migrating a SAS environment. One is to use the SAS[®] Migration Utility, and the other is to use the import and export functionality that exists in SAS[®] Management Console. As the choice of these two pathways is key to a successful migration, it's worth documenting the different pathways and the reasons for choosing each approach. For the purposes of clarity within this paper, I am calling these the “migration pathway” and the “promotion pathway.”

It's worth noting that both of these migration paths focus on the migration of the SAS metadata and environment configuration information between environments. The migration of data between environments is discussed later in this paper.

MIGRATION PATHWAY

As stated above, this pathway is based on using the SAS Migration Utility; in essence, this is a like-for-like migration, where the source and target environment are very similar in terms of architecture and infrastructure. The process flow can be seen in Figure 1.

Top Five Considerations for Migration Success

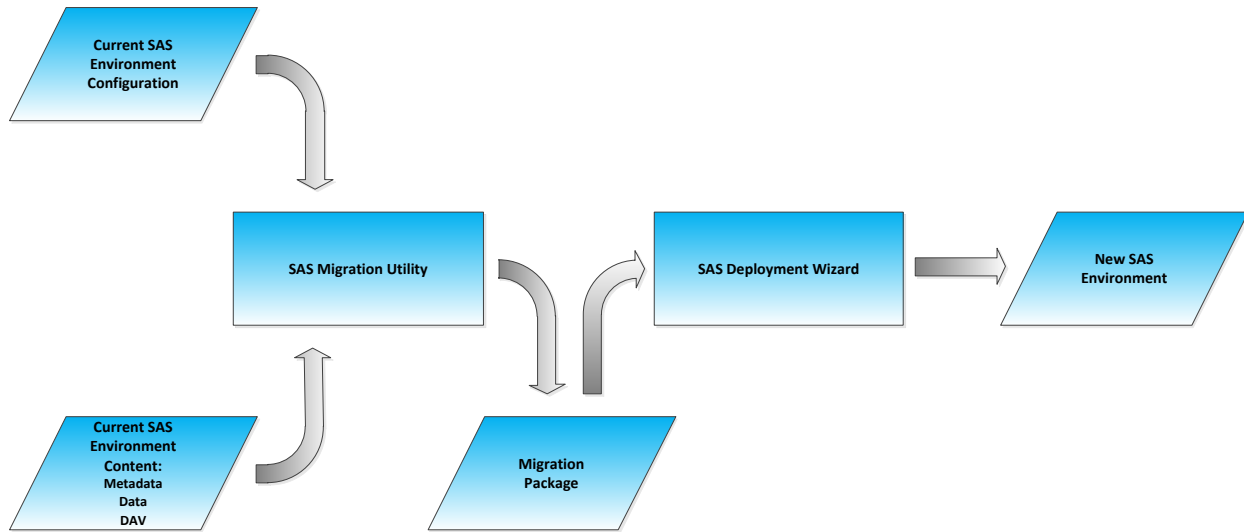


Figure 1: Migration Pathway Process Flow

Essentially, the SAS Migration Utility is run on the source environment, which creates a migration package that contains content and configuration information from the source environment. This package is then used by the SAS[®] Deployment Wizard to configure the target environment. The mechanics of this process are beyond the scope of this paper, but this process results in a target SAS environment that contains all of the content of the source environment.

PROMOTION PATHWAY

The second pathway is the promotion pathway, based on the import and export functionality of SAS Management Console. This pathway is more flexible than the migration pathway in two ways:

1. The source SAS environment is migrated piece-by-piece by the administrator, which means that only the items that are required in the target SAS environment must be migrated.
2. There is no restriction on the infrastructure or architecture of the target and the source environments.

However there are some considerations. The flexibility provided by this pathway requires manual activities and, hence, a greater level of governance (for example, "Have we migrated the table that this report is based on?"), and although there are no restrictions on the architecture and infrastructure of the environments, there are limits on the different types of metadata objects that can be exported from different versions of SAS[®]9. The process flow for this pathway can be seen in Figure 2.

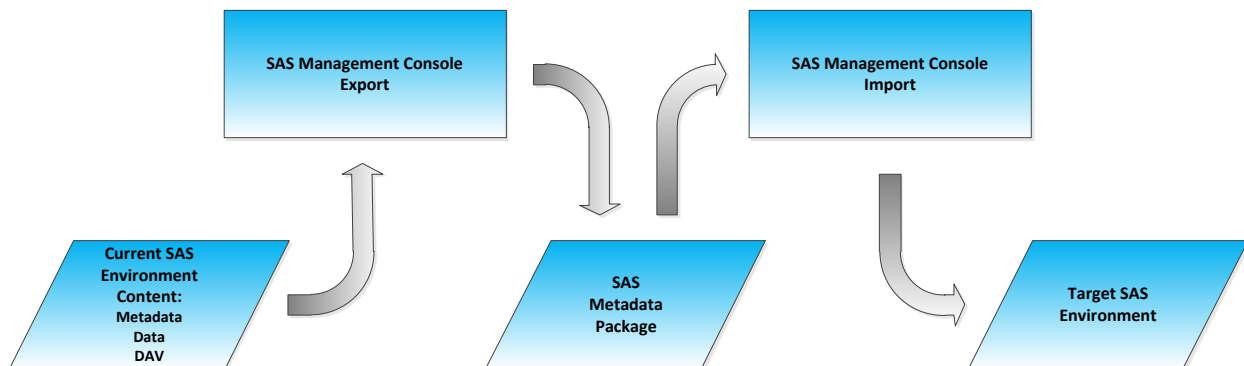


Figure 2: Promotion Pathway Process Flow

Top Five Considerations for Migration Success

Consideration 1: The tools to support SAS migrations have different capabilities, and each tool provides different functionality with the different SAS[®] 9.1.3, SAS[®] 9.2, and SAS[®] 9.3 releases.

RESEARCH METHOD

As can be seen above, there are numerous factors that affect how one might migrate a SAS environment, and there are potentially numerous approaches to migrating a SAS environment. To understand how customers have migrated their environments and to learn from the experiences of those that have gone before, the global SAS Professional Services community have been surveyed.

This survey took the form of 23 questions that were piloted to a small group initially, with feedback confirming the relevance and quality of the questions. The data was then collected during Q4 2011 and Q1 2012.

RESEARCH RESULTS

In all, SAS Consultants responded relating to 19 migration projects, with the breakdown of industries represented in Figure 3.

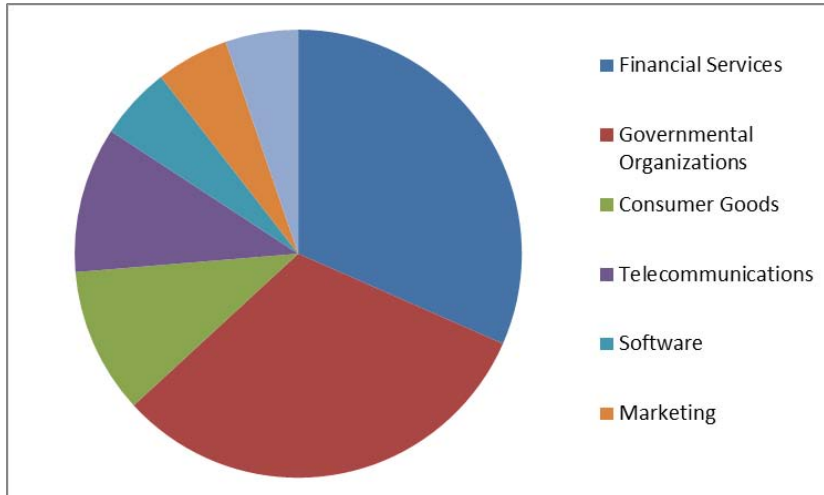


Figure 3: Industry Breakdown

These projects ranged from migrations of small SAS environments, taking two days, to environments with tens of thousands of metadata objects, taking over a year to complete the migration. These SAS environments also cover a number of different industry verticals from the financial services industry to government-related organizations.

The list of SAS technologies that all of these environments supported is too long to list, but it includes SAS[®] Enterprise BI Server, SAS[®] Enterprise Data Integration Server, SAS[®] Enterprise Model Management, and a number of SAS solutions such as SAS[®] Marketing Automation and SAS[®] Financial Management. SAS solution migration specifics are out of scope for this paper.

RESULTS DISCUSSION

MIGRATION TASKS

There are numerous activities and tasks that form part of any SAS migration. For the purposes of simplification, these activities have been grouped into five areas:

1. planning the entire migration activity

Top Five Considerations for Migration Success

2. installation of the target SAS environment
3. migration of SAS metadata from the source environment to the target environment
4. migration of data from the source environment to the target environment
5. testing and validation of the target environment

The responses to these questions have been graphed (Figure 4) by the organization's sector. Clearly this is a small sample, but it's interesting to see some of the correlations. It appears that the Financial Services organizations spend more time on migration and less time on testing, whereas the reverse is true for the time spent on migrating the SAS metadata artifacts.




Figure 4: Time Focused on Different Migration-Related Activities by Industry

Of perhaps more interest is how the size of the SAS environment affects the time spent on different activities during successful migrations. From Figure 5 it appears that larger environments require more time for planning the migration, and more time is required to test and validate the environments than the other migration activities. The SAS environment size was determined by the number of metadata artifacts, with a large environment being an environment with more than 1000 metadata artifacts, a small environment one with less than 100 metadata artifacts, and all environments in between being medium environments. With this in mind, one could have expected that the actual migration of metadata artifacts activity would consume the most time during a migration.

This is supported by the fact that in ten of the large environment migration projects surveyed, eight required a period of parallel running, where both the source and target environments were running at the same time, with the other two operating a switchover approach. From experience, running systems in parallel does enable business continuity and reduces risk of impacting service levels, but it has an overhead in terms of configuration management cycles to keep the environments synchronized, which can elongate the elapsed time for testing a migration. Migration validation and testing like this has become more prevalent in recent times due to regulatory requirements such as Sarbanes Oxley.

Consideration 2: The size of the environment can be correlated with testing and validation time required. One approach to enable adequate testing without impacting the production use of an environment is through parallel running.

Top Five Considerations for Migration Success

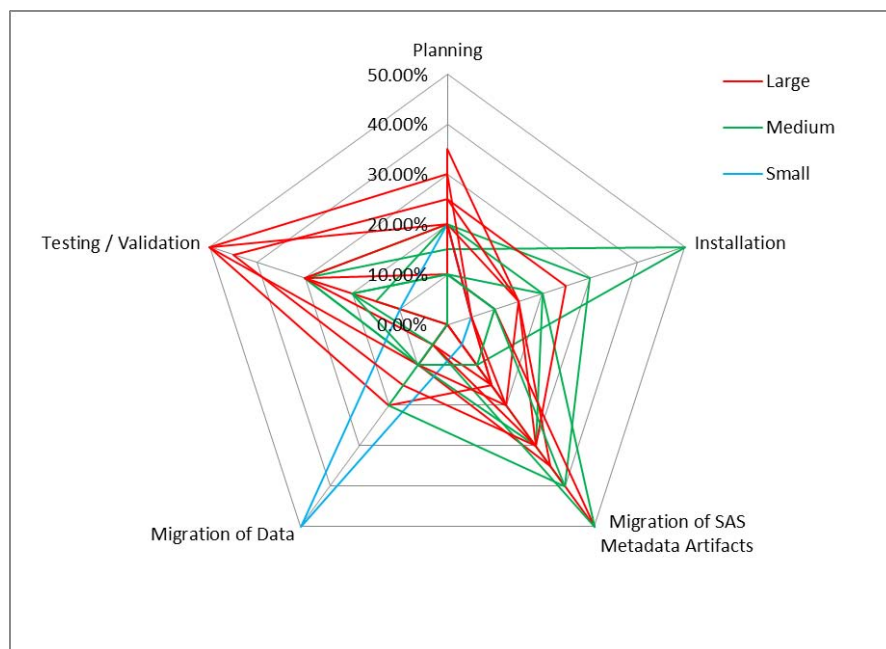


Figure 5: Time Focused on Different Migration-Related Activities by Environment Size

Different industries and different size organizations have different hardware refresh cycles. However, many organizations see a migration as an opportunity to perform two additional activities to their SAS environments, infrastructure changes and environmental “house-keeping.”

INFRASTRUCTURE CHANGES

In the migration projects surveyed, over half of the organizations used a migration as an opportunity to deploy the SAS environment onto a new hardware environment. This is a benefit of supporting parallel running in a cost efficient manner as well as allowing the organization to perform the migration and the deployment to a new environment in a single step. Obviously, from a migration perspective, that has an impact that needs consideration.

The SAS metadata holds significant information about the configuration of the environment (an example of this is application server contexts that can be found in SAS metadata). If the infrastructure supporting the SAS environment alters, then this information needs to be modified. For example, if the organization wants to move from a single server configuration for the SAS Compute Servers (SAS Workspace Server, SAS[®] Stored Process Server, SAS[®] OLAP Server) to a SAS Grid Computing configuration, then the SAS[®] Metadata Server will reflect this through more and differently configured application server contexts in metadata.

Consideration 3: From the two pathways outlined previously, making architectural changes during a migration can indicate that the promotion pathway is more appropriate.

HOUSE KEEPING

It has become clear that even with tight administration of a SAS environment (particularly those that have been in use for an extended period of time, and which support a large and diverse user population), a large amount of SAS metadata can build up that is no longer required by users and provides little value. This belief is supported not only by the primary research performed to support this paper, but also by the author’s experience. This proliferation of SAS metadata could be caused through many uses of the SAS environment, such as new projects replacing old or prototyping new SAS[®] Data Integration Studio jobs or SAS[®] Web Report Studio reports.

So when migrating between SAS environments, you might not want to migrate all of the SAS metadata, but just that

Top Five Considerations for Migration Success

metadata that you consider to be providing value today. Also, it's worth noting that a number of the original metadata definitions and structures that were designed at the inception of the original environment might no longer support the "evolved" use of the environment. The migration project offers the opportunity to leverage the organizational experience of how the environment is used to create a SAS metadata environment that is simpler in terms of administration and maintenance. The requirement to do this is strengthened by the introduction of a new metadata folder structure with SAS 9.2 that significantly simplified the organization of SAS metadata content.

Consideration 4: A new SAS environment might be simpler to manage if redundant metadata is removed, and the new SAS metadata folder structures introduced with SAS 9.2 are implemented.

DATA MIGRATION

Thus far within this paper, we have focused on the SAS metadata aspects of a SAS environment migration. Another area that is important is the migration of other items, such as SAS data sets, catalogs, and indexes. The impact on these items during a migration is very dependent on the specifics of the data architecture and the source and target environments. Although the average time spent on migrating these items in the migration projects surveyed was only around 10%, there is tremendous variance. Data migration consumed 50% of the elapsed time of one migration project. Therefore, this is an area worthy of consideration.

Neither of the migration pathways discussed previously are aimed at the migration of these items, but there is substantial collateral to support this from a technical perspective on the support.sas.com site including a "migration calculator."

As discussed above, some organizations take the opportunity to change infrastructure during a migration project, which could impact the data architecture, and some organizations might consider altering their data platform. For example, an organization might move to a different database to take advantage of SAS[®] In-Database processing. To assist in this type of data migration, SAS provides numerous tools, including SAS[®] Data Integration Server. (The many details of how to do this are beyond the scope of this paper.) However, this part of the overall migration process alone can take significant time, and the timing of this part of the migration is important in planning an overall SAS migration project.

Consideration 5: Depending on the specific environment, migration of data may be required, and may significantly impact migration timescales.

WORKING PRACTICE

As we can see, there are many factors that influence a migration project. Not all of these factors are based on technical constraints or capabilities. Now that we have discussed the main considerations for a SAS migration, it's valuable to discuss them in context of a working practice migration approach. A derivation of this working practice is currently being performed by one customer on a very large SAS environment, moving from a single server environment to a SAS Grid Computing environment.

Top Five Considerations for Migration Success

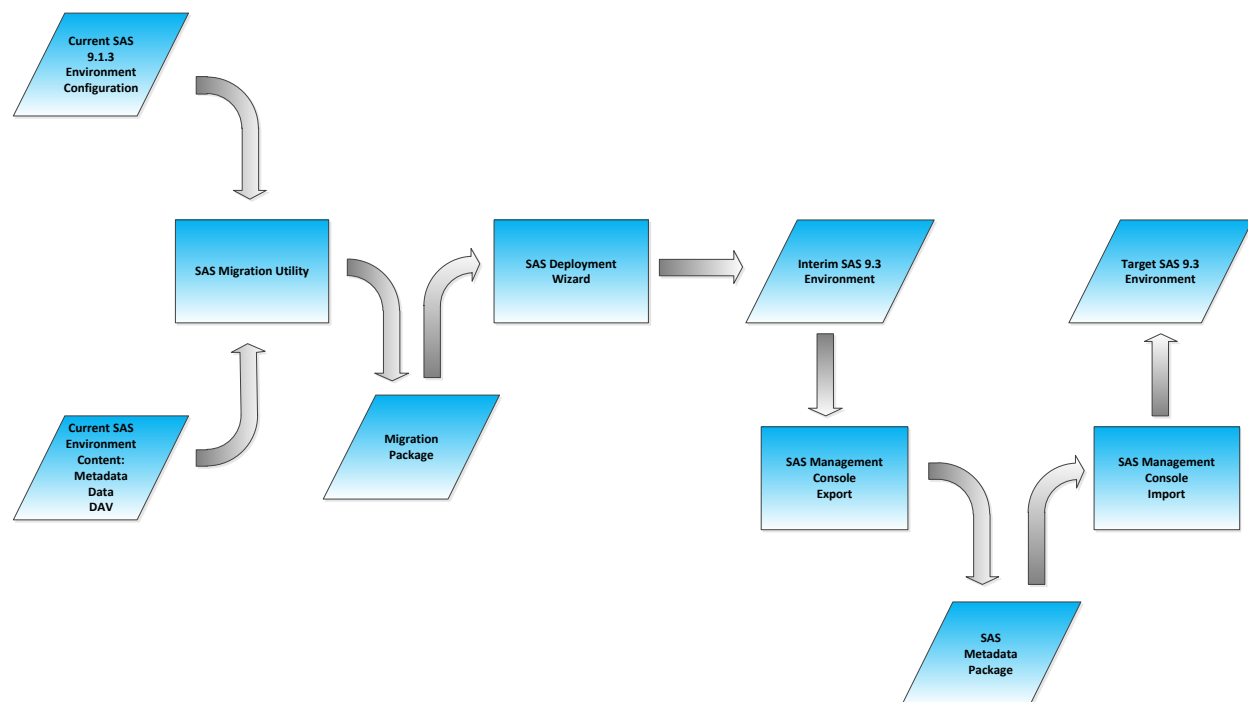


Figure 6: A Working Practice for a Large SAS Environment

In this migration, the organization was moving to a SAS Grid Computing based architecture from an existing single server environment, which meant that we needed to factor in **Consideration 3** and use the promotion pathway. However, the SAS promotion pathway as defined above uses the import and export capabilities provided by SAS Management Console. When we factor in **Consideration 1**, we see that in each of the releases of SAS^{®9}, more capabilities have been added in terms of the types of SAS metadata objects that can be exported and imported. However in SAS 9.1.3, the types of SAS metadata objects that could be migrated was limited, and as the customer had significant amounts of these metadata types in their environment, this was not an option on its own.

One way to meet the output of these two considerations is the use of an interim environment similar to the source environment. In this specific case, the source environment is based on a single Sun Solaris server, so the interim environment is also based on Sun Solaris. However, the SAS version installed on the interim environment is similar to the target environment. The reason for this is that SAS 9.3 supports the import and export of more SAS metadata types.

As this interim environment is similar to the source environment, the SAS Migration Utility can be used to replicate the entire SAS metadata from the source system.

At this point, it's important to note that in this case the interim Solaris environment did not have to replicate the processing capacity of the source or the target environment. In fact, only the SAS Metadata Server and the corresponding WebDAV server need to be available, so virtualization technologies could have been deployed here to ensure that the cost of the migration is minimized.

Once the SAS metadata is replicated to the interim environment and is "operational," the SAS Metadata Server can be started, and the promotion pathway can effectively be followed from the interim environment to the target environment. This approach enables us to comply with **Consideration 4**. Now it's possible to selectively migrate SAS metadata artifacts from the SAS 9.3 interim environment into a well-designed SAS metadata folder structure. One significant benefit of this (which is beyond the scope of this paper) is that this folder structure allows for a simpler security authorization structure to be designed and implemented.

The migration of content other than metadata sits outside this process; however, contemplating **Consideration 5**, migrating these objects to the target environment before the metadata migration from the interim environment to the target environment occurs enables any "plumbing" metadata changes to be tested immediately. For example, this allows metadata libraries to be constructed and tested right away.

Finally, this working practice allows support of parallel running of the source and target environments for the purposes of testing and validation as outlined in **Consideration 2**. There are many details about how to perform this validation that are beyond the scope of this paper. However, it's worth noting that the SAS environment does provide

Top Five Considerations for Migration Success

a platform for creating various validation processes.

CONCLUSION

In the migration of any computer system that provides value to its organization, there are always constraints, requirements, and risks that need to be evaluated. In this paper, I have tried to discuss and detail the major considerations for the migration of SAS environments in the SAS[®]9 era and to propose a logical working practice that can be adapted, modified, or used as is. I understand that this working practice might be aimed at large, more complex environments, but I contend that the considerations discussed are applicable to the migration of any SAS[®]9 environment and their contemplation earlier rather than later should assist in a successful migration.

ACKNOWLEDGMENTS

I would like to acknowledge the peer review and assistance given in the creation of this paper to Steven O'Donoghue and David Stern, both of SAS UK.

RECOMMENDED READING

- SAS(R) 9.3 Intelligence Platform: Migration Guide. Available at <http://support.sas.com/documentation/cdl/en/bimig/62613/HTML/default/viewer.htm#titlepage.htm>
- SAS Migration. Available at <http://support.sas.com/rnd/migration/index.html>
- SAS Migration Compatibility Calculator. Available at <http://support.sas.com/rnd/migration/planning/files/calculator/index.html>

CONTACT INFORMATION

Your comments and questions are valued and encouraged. Contact the author:

Andrew Mott
Marlow, Buckinghamshire
SAS Institute, Inc.
E-mail: Andrew.mott@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.