

Managing the SAS® Development Life Cycle across Environments and within a Single Production Environment

Andrew Howell, ANJ Solutions

ABSTRACT

How many environments does your organization have - three (Dev, Test, Prod), five (Dev, SIT, UAT, Pre-Prod, Prod), or maybe only one?

Once you've built your SAS® process - an ETL job, a model, an exploration, or a report - how should you promote it across these environments?

If you have only one environment, is a development life cycle still possible? (Yes, it is.)

Historically, the traditional systems development life cycle (SDLC) spans multiple environments (for example, Dev/Test/Prod). This approach has benefits-primarily to ensure that change in one environment does not adversely impact others, but costs and release time-frames mean this is not always practicable.

Some sites now adopt a two-platform implementation: Non-Production and Production. Non-Prod exists for technology change, such as new software, hot fixes, database connections, and so on. At these sites, the business runs wholly within the Production environment, yet still requires a business-specific life-cycle management within the Production environment.

And, of course, all promotion must include thorough testing.

Other questions to consider are:

- Can this promotion process be automated?
- Can this process extend beyond business content to include configuration settings?

This paper investigates the SAS tools available to promote content between environments or between functional areas of a single environment, and how to automate and test the promotion process.

Just imagine: a weekly automated and tested promotion process? Let's see

INTRODUCTION

A TRADITIONAL SOFTWARE DEVELOPMENT LIFECYCLE

In a traditional Software Development LifeCycle (SDLC), the various phases were well segregated, taking place in physically (or at least logically) separate environments.




Figure 1. Traditional Software Development Lifecycle (SDLC)

Development takes place within a Development environment; changes are promoted to a Test environment for testing & user sign-off, and then promoted to a Production environment, where the content is available for consumption. All change takes place within the Development environment, whereas NO change takes place on the Production environment.

Note: For the purpose of this paper, this scenario has been simplified to 3 environments - Development, Testing and Production. It is feasible there could be many more, depending on the level of testing required; for example Development, Systems Integration Testing, User Acceptance Testing, Pre-Production, Production.

In this methodology, changes are rarely implemented "on-demand". Typically, changes are bundled into "release cycles" - possibly daily, but more likely weekly, monthly or quarterly.

Similarly, content can be **demoted** – for example, if any levels of testing fail, content (including test cases & test data) can be demoted to a prior environment for debugging.

However this methodology presents issues:

- By the very nature of the release sequence, each "downstream" platform is a (*major or minor*) release behind its predecessor.
- Different changes have different implementation cycles.
- Often the non-Production platforms are not of the same technical specifications as the Production platform, often compromising the testing phase (in particular, stress testing).
- Non-Production data is often not of the same quality as Production data.
- One critical flaw of the traditional SDLC release cycles is that the frequency of release cycles was generally not sufficiently responsive to ever changing Business Requirements.

THE MODERN SDLC

More commonly today, the traditional SDLC methodology is still in place for "technology" changes (patches, hot fixes, new software), which are still promoted to the Production platform as "Operational" components, as per traditional SDLC processes.

However, business activities now more commonly take place directly on the Production platform in a "Discovery" or "Sandpit" workspace, with steps (hopefully) taken by Platform & System administrators to ensure Business processes cannot impact Operational processes.

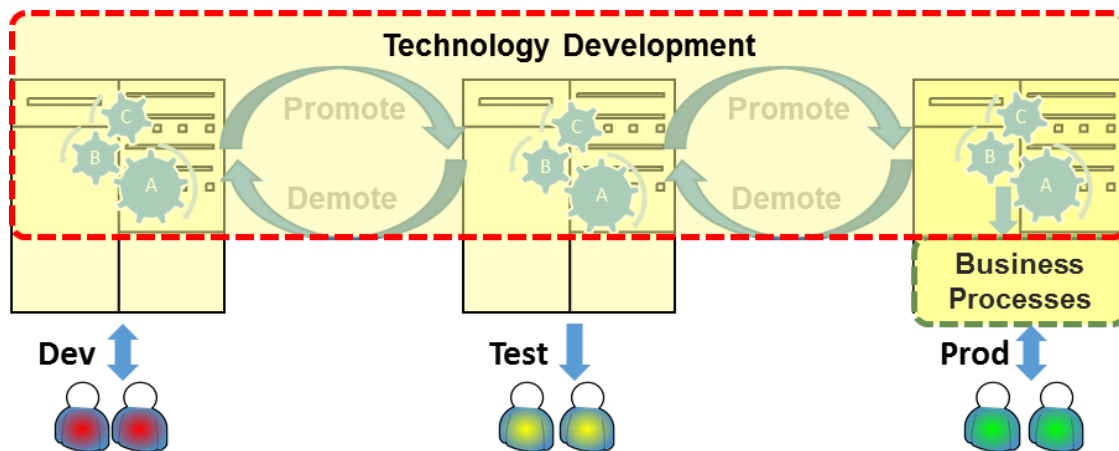


Figure 2. Modern Software Development Lifecycle (SDLC)

So the question must be asked: Can a Business “self-manage” its processes within a single environment?

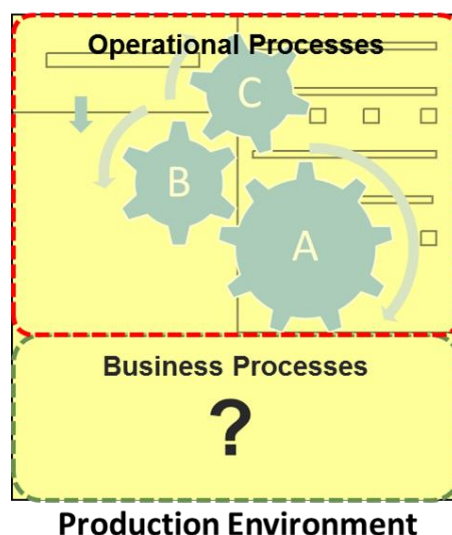


Figure 3. Business Processes within a Production environment

Without defined business processes, a single all-purpose Business Workspace risks quickly becoming unmanageable.

INTENDED AUDIENCE

This paper is intended primarily for those responsible for the promotion of SAS metadata, either from one environment to another, or within a single environment; these may include:

- Platform Administrators who perform the promotion of business content;
- Platform Administrators who design a platform security model to allow the Business to manage their own business content;
- Business Unit Administrators who perform promotion of business content.

SCOPE

"Promotion is the process of copying selected metadata and associated physical content within or between planned deployments of SAS software." - **SAS 9.4 Intelligence Platform, System Administration Guide, Fourth Edition.**

The scope of this paper is limited to promotion of metadata content (and where relevant, associated physical files) between existing SAS platforms of the same software version, or within an existing SAS platform. It does not extend to platform migration from earlier versions of SAS, nor deployment of SAS software.

Similarly, this paper does not discuss of lifecycle management of SAS models developed with SAS Enterprise Miner & managed with the SAS Model Manager.

TERMINOLOGIES USED IN THIS PAPER

- Business Unit - an organisational group on the SAS platform, typically within a Production environment.
- Business Unit Administrator - a senior SAS user within the Business Unit, who has been delegated Administrative privileges by the SAS Platform Administrator
- BusDev - a workspace (both metadata & physical disk) for Business Unit development, typically within a Production environment.
- BusOps - a workspace (both metadata & physical disk) for consumption of protected Business Unit content, typically within a Production environment.

SAS SOFTWARE DISCUSSED

Fortunately, the tools used by Platform Administrators to promote content between environments, can also be used by managers of business metadata to promote content within a single environment:

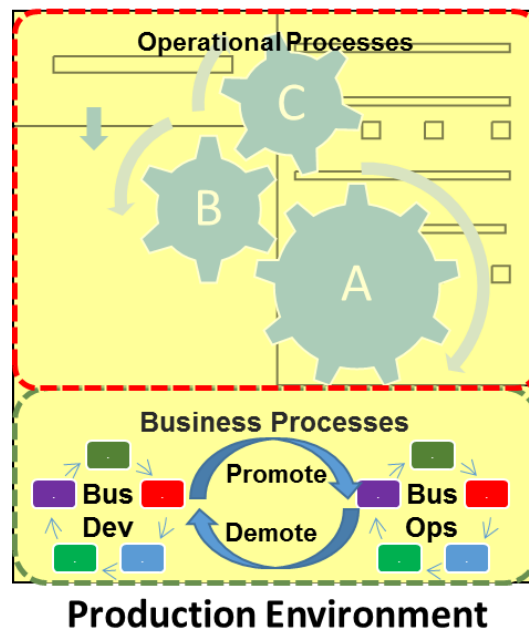


Figure 4. Business Software Development Lifecycle (SDLC) within a Production environment

The focus of SAS tools discussed tools available for managed promotion of SAS metadata, available in SAS Management Console (and with lesser functionality in SAS Data Integration Studio & SAS OLAP Cube Studio), and equivalent command line tools.

- SAS Export & Import Wizard (for SAS metadata & associated physical content)
- SAS Batch Export Tool & SAS Batch Import Tool (for command-line export/import of SAS metadata).
- Other supporting command-line tools (for example, **sas-make-folder**)

For promotion of non-metadata content, this paper discusses:

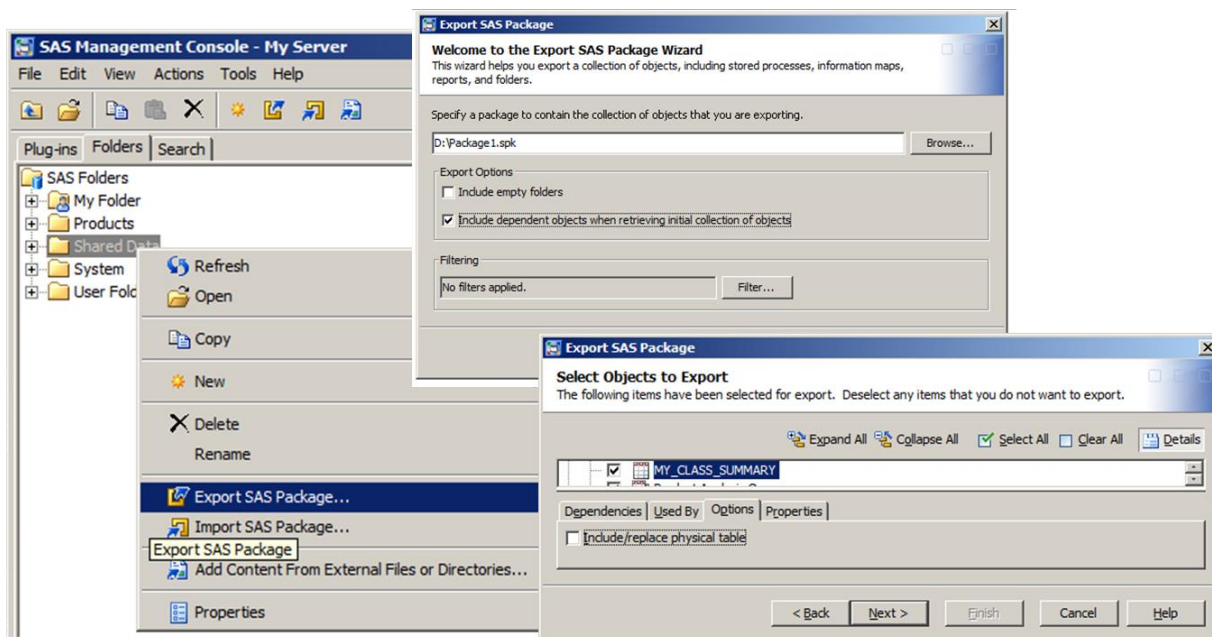
- SAS Enterprise Guide Migration Wizard
- SAS Enterprise Guide Project Maintenance
- SAS Enterprise Guide Search Utility

SAS EXPORT & IMPORT WIZARDS

The SAS Export & Import Wizards are an interactive series of steps to export metadata content (and possibly associated physical content) from one environment to a SAS package file (*.spk), which can then be imported into another environment, or to an alternate location of the original environment.

The wizards can be used in SAS Management Console, SAS Data Integration Studio & SAS OLAP Cube Studio, although SAS Management Console allows the broadest range of metadata content to be exported & imported. The Export & Imports Wizards within SAS Data Integration Studio and SAS OLAP have a narrower scope of metadata objects, specific to ETL flows and OLAP cubes respectfully.

Note: SAS Data Integration Studio does enable archiving content into versioning systems, for example Apache Subversion (SVN) server & Concurrent Versions System (CVS) server. SAS Management Console does not have this feature, although it would be a useful inclusion in future releases. However, versioning is not included in the scope of this paper.

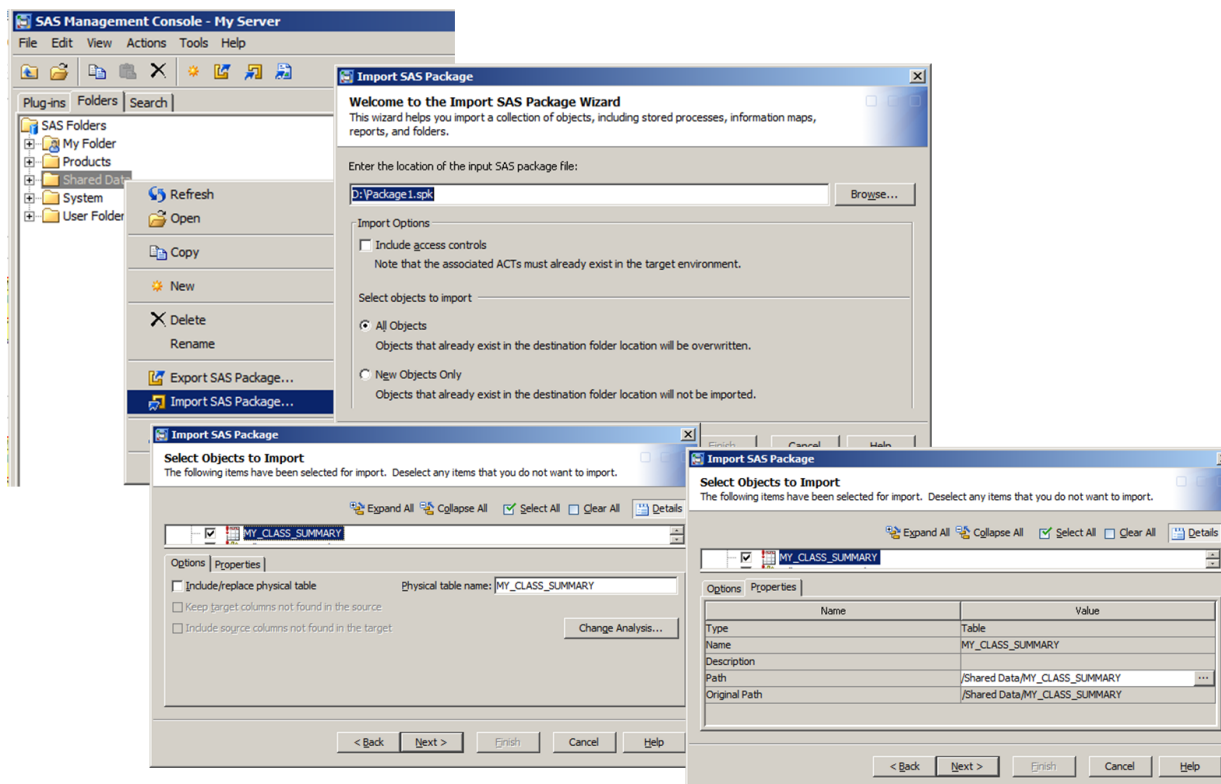


Display 1. SAS Management Console – Export Wizard screenshots

By default, the Import Wizard will import a package to the same folder location from which the metadata was exported. (In the example above, the **Shared Data** metadata folder.

However the Import Wizard allows metadata content to be modified during the import process. For example:

- Relocation of metadata objects. If the object is a folder, then all items within the folder, including sub-folders are also relocated;
- Reassigning of tables to alternate libraries, or libraries to alternate SAS Application Contexts;
- Reassign paths to alternate physical data.



Display 2. SAS Management Console – Import Wizard screenshots

This also means metadata exported from a “BusDev” metadata folder within a Production environment can be imported to an alternate “BusOps” metadata folder **within the same environment** – ideal for business units trying to manage their own metadata content.

When importing, it is generally advisable to perform a phased import rather than an “all at one” import process. This is particularly important if importing security objects (Access Control Templates, users, groups, roles) prior to importing subsequent objects (libraries, tables, jobs, reports) potentially affected by security authorizations of the imported security objects.

General recommendations for the sequence of importing metadata objects:

1. Security objects (Access Control Templates, users, groups, roles)
2. Libraries
3. Tables
4. ETL flows
5. Mid-tier content (reports, etc)

COMMAND-LINE TOOLS FOR SAS PROMOTION

Whilst the interactive tools are a useful means of better understanding the promotion process, in an ideal world I would advocate an “unattended promotion” process, especially to the Production platform. Being unattended it would most likely execute more quickly, plus the promotion process could be scheduled to occur overnight, at a time of minimal (if any) business usage of the platform.

The Batch Export & Import Tools (ExportPackage & ImportPackage) are located in the following folder:
\$SASHOME/SASPlatformObjectFramework/9.4

These tools offer (almost) the same functionality as the Export/Import Wizard; the primary exception being that although table metadata & external file metadata are exported, the underlying physical tables & external file are not included in the exported package (whereas this is offered as an option in the Export/Import Wizard).

If external data IS required to be promoted (typically reference data), an additional step is required to retain a copy of the data to be promoted.

SAS BATCH EXPORT TOOL

The following is an example of a command to export all objects created or modified since 01Jan2017 (and any dependant objects) in the /Business/Released folder & subfolders.

ExportPackage

```
-profile "MyProfile"           (connection profile)
-package "package_pathname.spk" (destination package)
- objects "/Business/Released" (will export everything in each of these
                               "/Other/Folders" folders & subfolders)
-log "log_file_pathname.log"   (log file of the export process)
-since 01Jan2017:12:00:00      (everything created or modified
                               since midday 01Jan2017)
-includeDep                    (include dependent objects, recommended)
-includeEmptyFolders           (recommended)
-subprop                        (create an external substitutions
                               properties file, important if altering
                               metadata properties when importing)
```

A full syntax list is available in Chapter 26 ("Using the Batch Export and Import Tools") of the SAS® 9.4 Intelligence Platform: System Administration Guide.

With options such as **-since <<datetime>>** it is possible to perform scheduled exports of any changed metadata since a given point in time, along with any dependant metadata objects (libraries, tables, etc).

I would generally recommend the **-includeDep** option be used, to ensure all associated objects are also included in the exported package. These items can always be selectively included/excluded during the importing process (for example, if target libraries or tables already existed in the target environment).

SUBSTITUTION PROPERTIES FILE

The Substitution Properties File can be used to modify objects' properties and/or associations prior to performing an import process. The file can be manually created during the Export process, or can be extracted from the exported package file.

If the Import process requires certain metadata properties and/or associations, this file can be modified to apply the necessary changes during the import process. This is especially important if promoting metadata within a single environment; without modifying the substitution properties file, imported object would simply overwrite the source objects "in place".

Below are some examples which would require modifying the substitution properties file:

- Object properties:
 - Metadata folder location
 - Directory & file paths of associated physical content
- Associations:
 - Assign an imported table to an alternate library (for example, from the OraDev library to the OraOps library in the same environment)
 - Libraries, ETL jobs, stored processes to an alternate SAS Application Server Context

Following one of the above examples, any tables requiring reassignment to an alternate library, the following properties would need to be modified for each table. (Note: Modify only "Target" properties, not "Source" properties.)

```
[Connections: Library]
Library[1].SourcePath=/Business/Development/Libraries/OraDev(Library)
Library[1].TargetPath=/Business/Operations/Libraries/OraOps(Library)
```

SAS BATCH IMPORT TOOL

As mentioned earlier when discussing the SAS Import Wizard, it is generally advisable to import content in a logical sequence (Security objects first, then libraries, then tables, ETL flows, mid-tier content, etc). When using the ImportPackage tool, this is achieved using the **-types** option. For example:

```
-types ACT, User, UserGroup, Role
-types Server, Library, Table
-types Job
-types InformationMap, Report
- etc..
```

A full list of object types is available in Chapter 24 ("Using the SAS Intelligence Platform Batch Tools") of the SAS® 9.4 Intelligence Platform: System Administration Guide.

The following is an example of a command to import objects from a package, optionally selecting specific metadata object types to load, and selecting a Substitutions Properties file.

ImportPackage

-profile "MyProfile"	(connection profile)
-package "package_pathname.spk"	(source package)
-log "log_file_pathname.log"	(log file of the export process)
-subprop "subprop_pathname"	(use an substitutions properties file to alter metadata associations and/or properties when importing)
-types object-type-1, object-type-2, etc	(only load specific metadata object types)

METADATA-MANAGED PROMOTION

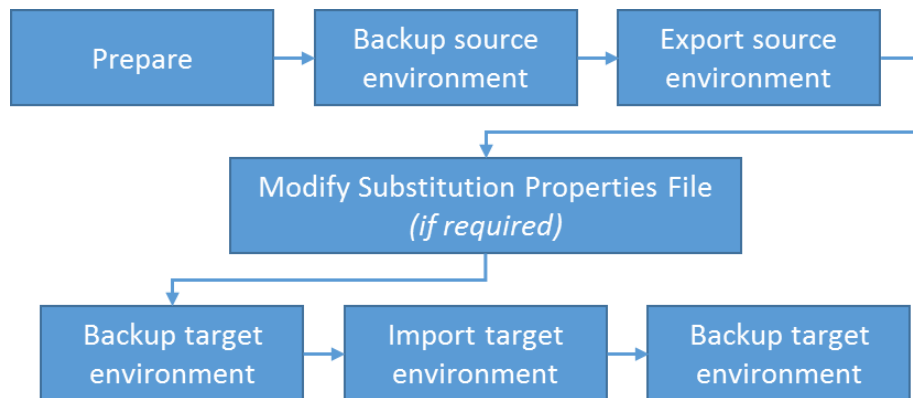


Figure 5. Recommended promotion steps

PREPARATION

Pre-requisites

Critically, ensure the metadata user performing the promotion has the necessary metadata & filesystem permissions. For promotion of metadata within a single environment, a clear metadata security model should be implemented:

- BusDev metadata should be secured, ideally with an Access Control Template (ACT), to ensure the appropriate business users have full metadata permissions to perform the necessary development work within this folder, whilst ensuring non-members of this business unit have ReadOnly or no metadata permissions to the development metadata.
- BusOps metadata should also be secured (ideally with an ACT), to ensure only a member of the BU Admin metadata user group can make changes, whilst ensuring the business users can consume such content with ReadMetadata (and possibly Read) permissions.
- Platform metadata should be secured (ideally with an ACT), to ensure content is available to the business where appropriate, but changes can only be made by a member of the Platform Administrator metadata user group.

A similar security model should be applied to the platform's filesystem.

Promote metadata, physical files or both?

When promoting content from one environment to another, consideration must be given as to whether or not to also promote associated physical files. Often this is not required, as the corresponding files may already exist in the target environment. However, reference data (for example) may be developed in conjunction with metadata, which would require promotion to the production environment.

The Export/Import Wizard has the capability to include associated physical files. However, the Export/Import Batch tools do not include associated physical content. Further, the associated files could make the exported package exceptionally large (and slow the export process), so if physical files are required to be promoted between environments, it will be necessary to promote the physical files separately (albeit simultaneously).

Target metadata folders must exist prior to import

If using the command-line tools to perform the promotion, a useful tool is the **sas-make-folder** command, located in the **\$SASHOME/SASPlatformObjectFramework/9.4/tools** folder. Note: the **-makeFullPath** options ensures any "parent" folders are also created.

BACKUP THE SOURCE ENVIRONMENT

Prior to exporting, perform a (“point in time”) full backup of the source metadata & physical files.

EXPORT SOURCE METADATA

If possible, prior to performing the export, set the source environment's metadata server to Administration mode (to prevent any change during the export process). Ideally, the metadata server should be placed in Administrative mode to ensure only administrators have read/write access to the metadata server during the (hopefully brief) promotion process.

```
proc metaoperate action=pause
    options="<Server STATE='ADMIN' />
        <PauseComment>Metadata promotion in progress..</PauseComment>";
run;
```

Export desired metadata to a package, including all associated content (dependant objects, empty folders, etc)

- If using the Export wizard and physical data is also required, ensure the **“Include/replace physical table”** option is selected for each table and/or file reference.
- If physical data is too large or using the command-line utility, ensure associated physical data is also backed up.

If necessary, resume the source environment's metadata server.

```
proc metaoperate action=resume;
run;
```

UPDATE THE SUBSTITUTIONS PROPERTIES FILE

Should target metadata properties and/or associations require modification, apply the necessary changes to the Substitutions Properties File.

BACKUP THE TARGET ENVIRONMENT

Prior to importing, perform a full backup of the target environment, including metadata & physical files.

IMPORT PACKAGE INTO TARGET ENVIRONMENT

If possible, set target environment Metadata Server to Administration mode (to prevent any change during the import process).

Ideally, perform a phased import process, importing multiple times from the same package(s).

If necessary, resume the target environment Metadata Server.

BACKUP THE TARGET ENVIRONMENT

Once importing is complete, perform an additional full backup of the target environment, including metadata & physical files.

POST-PROMOTION TESTING

After any promotion of content, but especially after promoting security changes (Access Control Templates, Users, Groups, etc), it is vital to a full suite of regression testing to ensure the health of the target platform:

- Verify status of target servers & spawners
- Deployment Testing
- Additional testing as required.

DEPLOYMENT TESTING

The Deployment Tester (consisting of a Management Console client, a Deployment Tester server and batch client) can be used to ensure the health of the SAS platform.

The batch client tool, located in **\$SASHOME/SASDeploymentTesterClient/2.2/bin** folder.

Further documentation of the Deployment Tester is available in Chapter 8 ("Using the Deployment Tester") of the SAS® 9.4 Intelligence Platform: System Administration Guide.

NON-METADATA CONTENT ("UNMANAGED" PROMOTION)

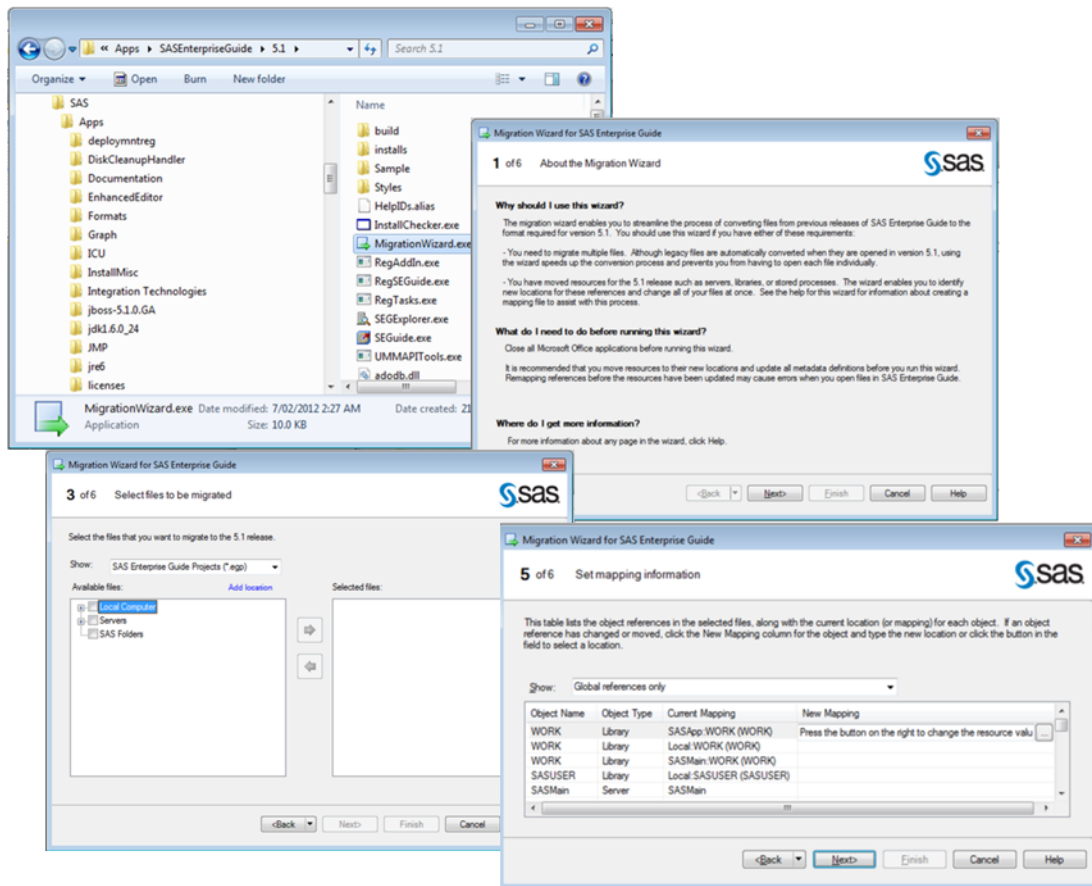
Defining all content in SAS Metadata would naturally greatly assist the promotion process. However, the reality is that a significant amount of business content is stored in Enterprise Guide projects and/or SAS code.

When promoting SAS content, there is a plethora of SAS code and Enterprise Guide projects which do not reside within metadata, and as such require additional steps to promote them to alternate platforms or within an existing platform.

*This paper does not include discussion of the changes required to SAS programs stored as *.sas files on a file-system. There are numerous third-party tools which can search and replace text within text files (such as *.sas files), plus I am certain there are others who - like me - have written SAS programs and/or macros with INFILE & FILE statements to load, search, replace & save selected text in SAS programs.*

SAS ENTERPRISE GUIDE - MIGRATION WIZARD

The **SAS Enterprise Guide Migration Wizard** (located in the same folder as the SAS Enterprise Guide executable) is a useful tool for bulk-updating of many Enterprise Guide project simultaneously. It allows for local and server-based searches of Enterprise Guide projects, and then applying changes such as SAS Application Server Contexts (including for example, a BusDev to BusOps within a single platform), as well as SAS Libraries, and more.



Display 3. SAS Migration Wizard screenshots

Capabilities include:

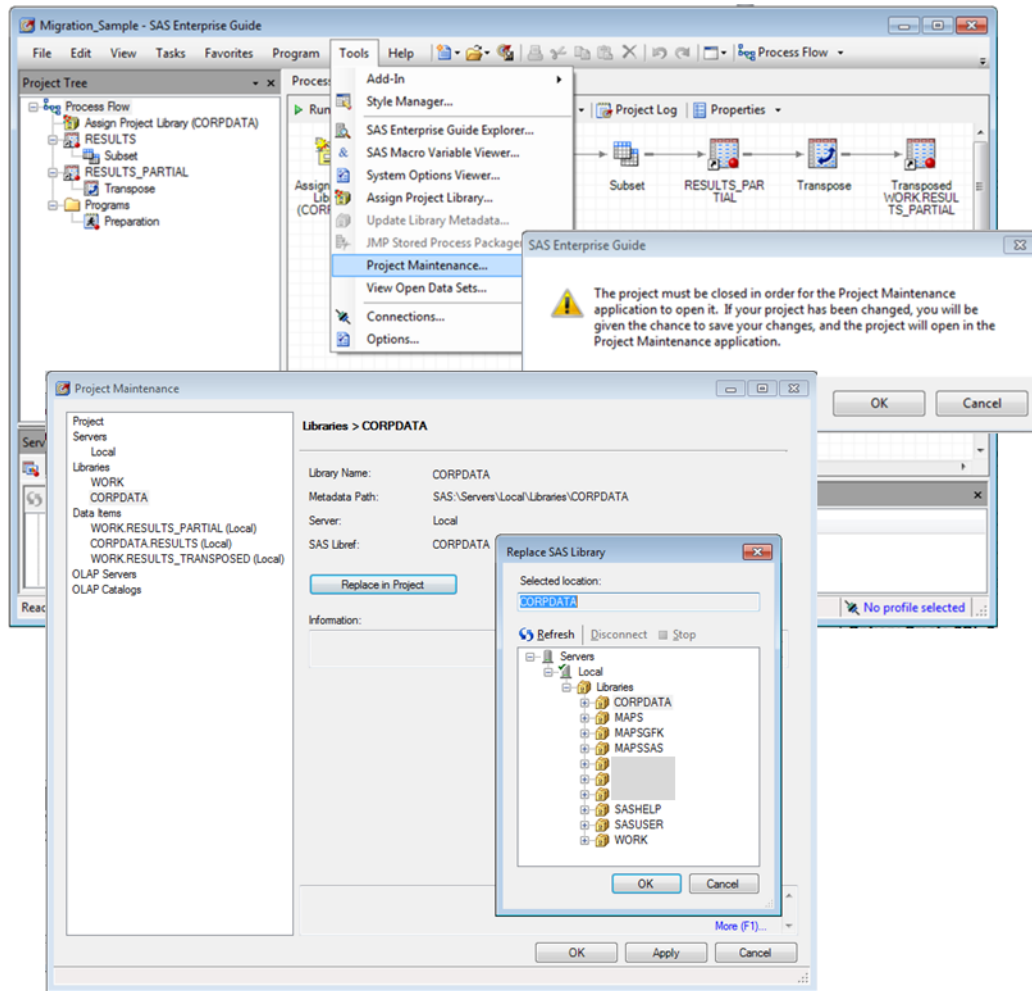
- Upgrading of Enterprise Guide projects to a consistent version
- Updating metadata references (SAS Application Server Context, SAS libraries, etc), which makes it ideal as a tool to promote Enterprise Guide projects from a "Development" to "Operational" status.
- Perform both local and SAS-server searches for Enterprise Guide projects

Note: This utility will not update embedded SAS code within Enterprise Guide projects; this will need to be managed within each Enterprise Guide project.

SAS ENTERPRISE GUIDE – PROJECT MAINTENANCE

Just as the above utility is useful for bulk-modifications of Enterprise Guide projects, within SAS Enterprise Guide is the Project Maintenance utility to enact changes to a single Enterprise Guide project.

It can also apply changes such as SAS Application Server Contexts (including for example, a BusDev to BusOps within a single platform), as well as SAS Libraries, and more.



Display 4. SAS Enterprise Guide Project Maintenance screenshots

Again note: This utility does not update embedded SAS code.

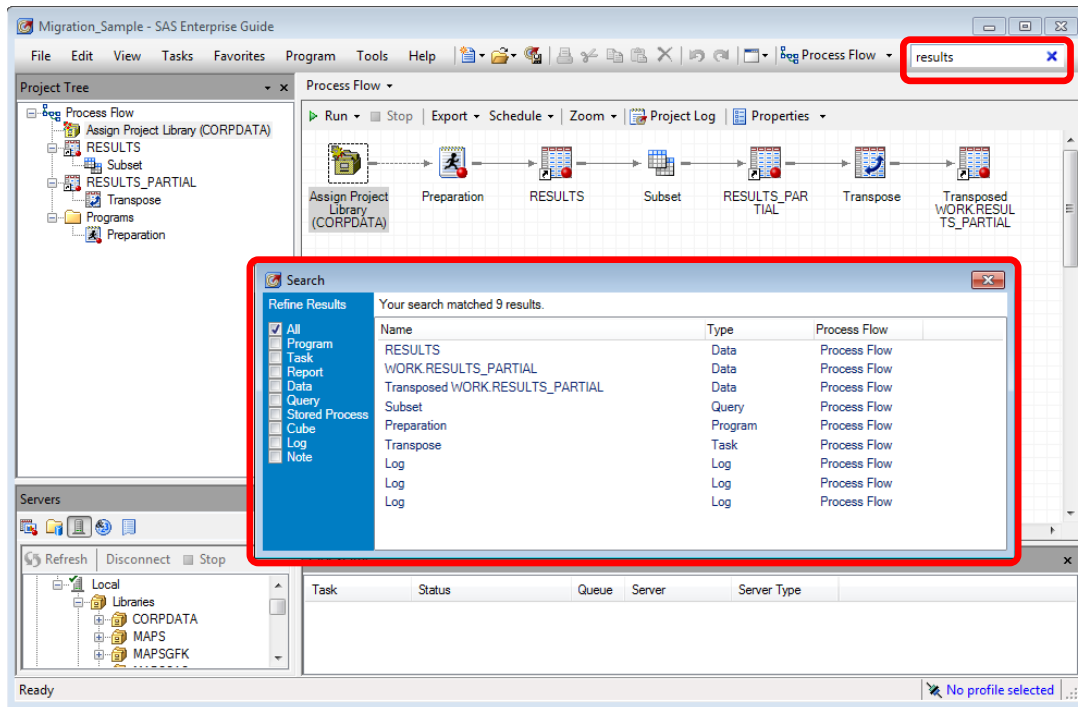
SAS ENTERPRISE GUIDE - SEARCH FACILITY

Enterprise Guide v7.1 now has a useful Search feature to find keywords throughout an Enterprise Guide project.

It is a useful utility for searching for key words within an Enterprise Guide project (including embedded SAS code).

(Prior to SAS Enterprise Guide v7.1, a Search API written by Chris Hemedinger was - and still is - available to download as an Enterprise Guide "add on".)

When promoting Enterprise Guide projects, after running either the Migration Wizard or the Project Maintenance utilities, the search facility is a useful method for searching for keywords in embedded SAS code which may require manual updating.



Display 5. SAS Enterprise Guide search facility

Although this won't perform the modification, it will at least identify which objects require modification.

PROMOTION OF CONFIGURATION FILES

Although this is not typically included in a promotion process, in some cases it may be necessary to promote configuration and/or autoexec.sas files – for example, to append an additional formats catalog to the FMTSEARCH option. Remember to always modify the “**usermods**” files, and never the sasv9.cfg or autoexec.sas files directly.

CONCLUSION

Whilst the SAS Management Console Export & Import Wizards offer the opportunity to better understand how to control the promotion of metadata, the command-line tools available in the **\$SASHOME/SASPlatformObjectFramework/9.4** folder offer the tantalising possibility of a regularly scheduled, fully automated promotion of metadata content – an administrator's dream.

The Enterprise Guide Wizards offer a way to manage the myriad of non-metadata projects & SAS programs.

The combined suite of software means that not only is a Platform Administrator's role simplified when promoting content across environments, but it is also possible to delegate administrative capabilities to the business to “self-manage” the life-cycle of business content within a single Production environment.

REFERENCES

SAS® 9.4 Intelligence Platform: System Administration Guide, Fourth Edition

SAS® 9.4 Intelligence Platform: Web Application Administration Guide, Sixth Edition

SAS Training Course - SAS Platform Administration (Fast Track)

SAS Blogs - <http://blogs.sas.com/content/sasdummy>

ACKNOWLEDGEMENTS

John Amrhein, SAS Global Forum 2017 Chair

Chris Hemedinger, Manager SAS Online Communities, SAS

Scott Bass, SAS New South Wales User Group (SNUG) committee

Deanna Schreiber-Gregory, SAS Global Forum 2017 Presenter Coordinator

CONTACT INFORMATION

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

Andrew Howell

ANJ Solutions Pty Ltd

Melbourne, Australia

Phone: + 61 407 898 513

Email: info@anjsolutions.com.au

Twitter: @AndrewAtANJ

LinkedIn: <http://au.linkedin.com/in/howellandrew/>

SAS Communities: <https://communities.sas.com/people/AndrewHowell>

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.