

Metadata Promotion in SAS® 9.2

Eric Bourn, SAS, Cary, NC

Diane Hatcher, SAS, Cary NC

ABSTRACT

Promotion of metadata content is typically used to support movement across Development, Test, and Production environments. In SAS® 9.1.3, we provided tools for full repository promotion and partial promotion. In SAS® 9.2, these tools have been enhanced to streamline functionality and provide batch interfaces. SAS® Management Console will continue to provide interfaces for full repository replication and partial promotion via the export/import framework that exists in SAS 9.1.3 with the BI Manager plug-in. In SAS 9.2, we have implemented a batch interface for the partial promotion framework that will allow you to create a schedulable and repeatable process for moving a set of metadata content across your environments.

INTRODUCTION

IT infrastructure Library (ITIL) Service Support processes define “change control” as the process of controlling change, including requests for change, approvals, execution, and validation. “Change management,” on the other hand, is the process responsible for handling changes to IT infrastructure or IT services (Stickel 2005).

Change control is the process to coordinate changes to content in a given environment in a rational manner. Typically, organizations set up three environments: a playpen environment (Development), a quality assurance environment (Test), and a published environment (Production). Changes can be made in an ad hoc fashion in Development (DEV) and promoted to Test (TEST) for quality assurance. After the content is validated, the changes can be promoted to Production (PROD) for broad end-user consumption. Promoted content is usually not editable by end users, and any changes must go back through the DEV/TEST/PROD process.

With SAS® 9, the SAS® Metadata Server is at the core of the intelligence platform. The SAS Metadata Server facilitates the sharing of content across SAS applications to provide a consistent representation of the business to all constituents in your organization. As a result, it is critical that the information that is surfaced to end users in a production environment has been rigorously protected from the impacts of change. This paper focuses on “change control” for the SAS Metadata Server. This is intended for SAS system administrators who are familiar with the SAS® 9 architecture and the role that the metadata server plays. We focus on the tools available for promoting metadata.

TYPES OF PROMOTION

There are two main methods for handling promotion of metadata content: full repository promotion and partial promotion. Full repository promotion is typically used to stage a baseline system or promote a wide range of changes over an existing deployment on a scheduled basis. Partial promotion is used in an ongoing manner, in support of project-level changes over time.

FULL REPOSITORY PROMOTION

Full repository promotion is more of an environment management tool. It involves overwriting the target environment with all the metadata content from the source environment. Any changes that were made in the target environment are lost. Therefore, full repository promotion should be done in only the following specific situations:

- establishing a baseline system infrastructure
- promoting an initial Portal environment (such as a baseline set of portal pages)
- duplicating an existing environment for another application
- scheduling a regular promotion to resynchronize the target environment with a wide set of metadata changes

Full repository promotions can occur across only separate metadata environments. This means that you cannot use full repository promotion to copy content from one custom repository to another custom repository in the same metadata server. Object ID references are copied during the promotion, so you cannot have duplicate ID references in the same metadata environment.

As a best practice, you should always promote all repositories in the metadata environment. Because custom repositories contain associations to metadata content that exist in its parent repository dependency tree, promotion of the entire metadata environment will maintain the integrity of those associations.

To execute a full repository promotion, you can use the Metadata Manager plug-in in SAS® Management Console or create a SAS program using macros provided by SAS. Using the Metadata Manager plug-in, you can define two types of jobs: a replication job or a promotion job. The replication job creates an identical copy of the metadata

repository. The promotion job creates a copy of the metadata repository, but it allows you to substitute values for key environmental variables, such as server name and port numbers that match your target environment. In SAS 9.2, we have combined the two types of jobs into a single job type. Replication will provide all the features of full repository promotion, including the substitution of values.

These tools handle promotion of metadata only; they do not deal with any associated physical files. You should set up a complementary process to handle physical content that is executed at the same time as the metadata promotion. This could be a combination of regenerating content (such as rerunning jobs and reports) and copying physical content from one location to another. The physical content files that you should consider are:

- webDAV content for reports and other portal content
- source code directories for stored processes and data integration jobs
- physical data tables and databases
- custom portal themes and custom portlets

Documentation for these tools can be found in the *SAS 9.1.3 Intelligence Platform: System Administration Guide*.

PARTIAL PROMOTION

Partial promotion provides more granular promotion capabilities that allow fine-grained control over which metadata objects are promoted. As such, promotion is not restricted to multiple metadata environments but can be executed within a single repository, across custom repositories in a single environment, or across repositories in different environments, even across hosts. This flexibility allows partial promotion to be the primary mechanism for executing changes within the SAS environment.

Essentially, the act of partial promotion in the SAS environment is exporting content from the source repository and importing that same content into the target repository. When exporting content, SAS creates a SAS package file (.spk file extension) that is essentially a .zip archive. The package file is stored on the file system of your environment. The SAS package file provides the flexibility to store external files in addition to the metadata objects. Therefore, you can also export associated physical content with the metadata in the export package file.

In SAS 9.1.3, there are two main tools supporting partial promotion.

- BI Manager plug-in version 1.4 for SAS Management Console

The BI Manager plug-in provides a user interface for handling the export and import of content, including re-mapping properties to match the properties in the target repository in order to account for different paths, library names, and server names.

The BI Manager 1.4 plug-in provides partial promotion for the following metadata objects:

BI-Related Content	DI-Related Content	Other Content
Reports	Data Integration Jobs	Mining results
Stored Processes	Libraries	Metadata Tree folders
Information Maps	Tables	
Data Explorations	Generated Transformations	
	Notes	
	Documents	
	External Files	
	Deployed Jobs	
	Job Flows	

- SAS® Data Integration Studio 3.4

SAS Data Integration Studio 3.4 has integrated the same underlying framework as the BI Manager plug-in as it relates to partial promotion, including copy/paste-special, which provides the ability to copy and paste metadata objects within a single repository as well as across repositories. The legacy copy/paste support in earlier versions of SAS Data Integration Studio is still supported, but we recommend that the newer paste-special function be used for promotion activities.

For detailed documentation of partial promotion capabilities in the BI Manager plug-in and SAS Data Integration Studio, see the *SAS 9.1.3 Intelligence Platform: System Administration Guide*.

PARTIAL PROMOTION IN SAS 9.2

In SAS 9.2, we have enhanced the partial promotion framework to provide more flexibility to incorporate existing change management processes. We recognize that change management must be treated as a repeatable process and allow multiple people to make changes to the environment. It must be easy to organize and select a subset of objects based on specific criteria; for example, only new objects created since a certain date, only a specific type of object, such reports or SAS Data Integration Studio jobs, or perhaps only objects that have been modified since the last promotion process.

The partial promotion framework in SAS 9.2 adds support for a wider range of metadata objects. Additional objects supported include:

BI-Related Content	DI-Related Content	Other Content
Enterprise Guide Projects	OLAP Schemas	Channels and Subscribers
Add-In for MS Office Documents	OLAP cubes	Message Queues
Prompts		

In addition, when you import scheduled job flows, you can attempt to reschedule them as part of the import process.

SAS MANAGEMENT CONSOLE IN SAS 9.2

As mentioned above, in SAS 9.1.3, the BI Manager plug-in in SAS Management Console provides an interface for the partial promotion framework. In SAS 9.2, we have removed the BI Manager plug-in, but incorporated the functionality into the metadata content tree view. As you can see below in Figure 1, there are two main tabs: the **Plug-ins** tab, and the **Folders** tab. The **Folders** tab surfaces the metadata content tree.

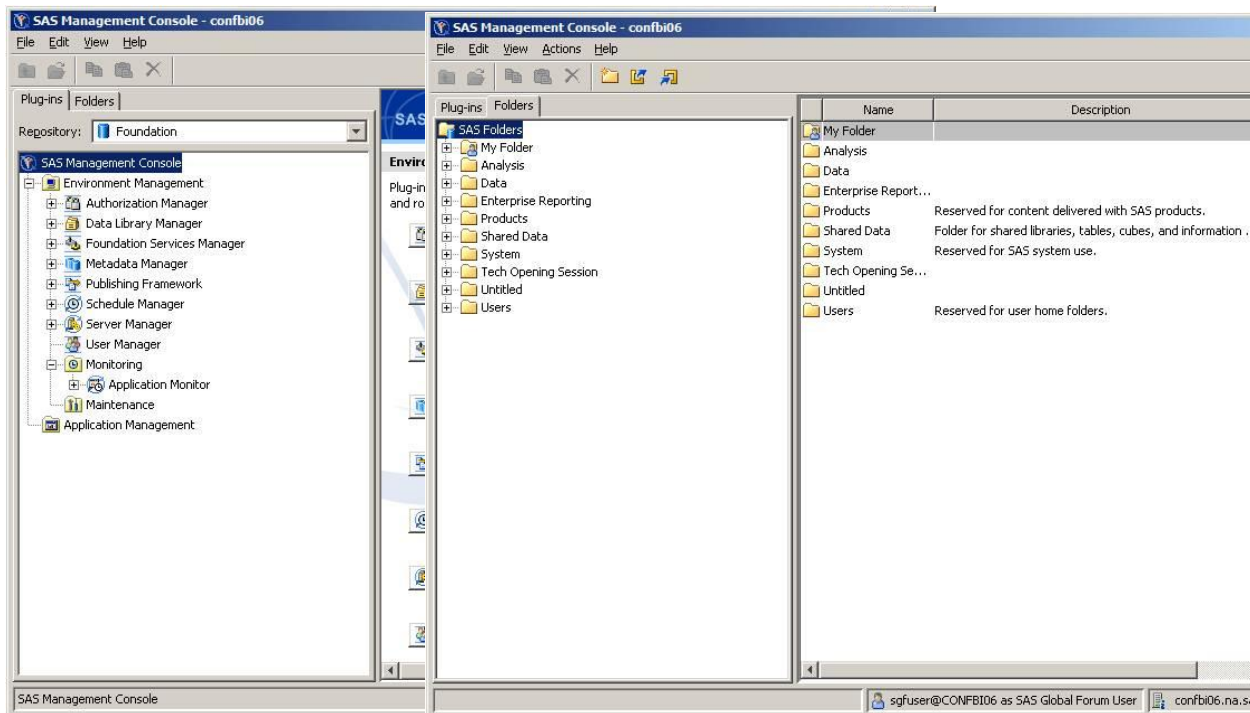


Figure 1. SAS Management Console in SAS 9.2

Right-clicking on a folder in the SAS Folders tree will surface a menu, where you can select to Export or Import a SAS Package. This is the partial promotion framework, and selecting either option will launch a wizard interface.

EXPORT FILTERING

The Export Wizard in SAS 9.2 has been enhanced to support filtering capability (Figure 2). Filtering allows you to subset which objects should be included in the export process. For example, if you are promoting a collection of objects from a Development environment to a Test environment, you can select only those objects that have been modified in a given timeframe.



Figure 2. Options Available on the Export Wizard

A new Filter dialog box has been added in SAS 9.2 (Figure 3). The three types of filters that can be applied when creating the SAS package file are: object name, type, and date filters.

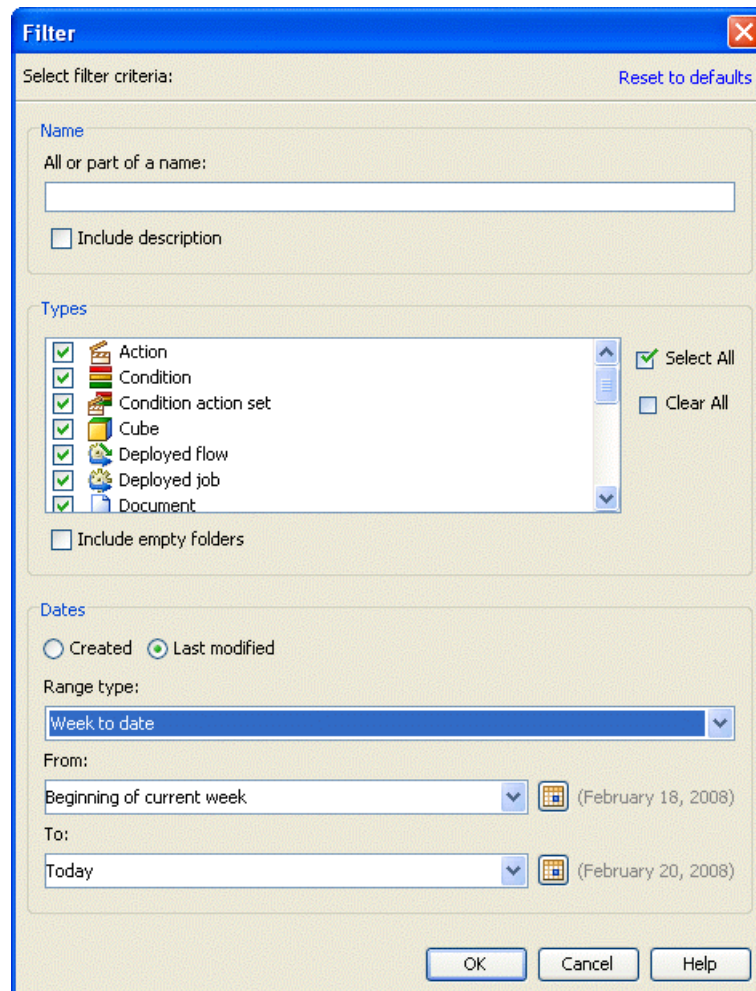


Figure 3. Export Filter Dialog Box

- Name Filtering

Allows you to specify the name, or part of the name, of an object. The “Include description” check box enables you to search the object’s description in addition to the name.

- Object Type Filtering

The filter dialog box contains a list of object types that can be included in an export. By default, all object types will be selected.

When exporting a folder structure, there is usually not a need to include any empty folders in the export process. For example, if a folder doesn't contain any objects or subfolders, it will be excluded from the export. The filter dialog box does allow you to override this setting with the "Include empty folders" check box. This option is not selected by default.

- Date Filtering

Objects can also be filtered based on when they were created or last modified. The Filter dialog box allows you to define your own date range or choose from a list of predefined ranges (Figure 4). Both absolute and relative dates are supported.




Figure 4. Filtering with Custom Date Ranges

OBJECT DEPENDENCIES (IMPACT ANALYSIS)

The dependencies feature within the Export Wizard has been enhanced in SAS 9.2 to support two-directional impact analysis. This feature allows you to view and select the objects on which a specific object depends. The **Used By** tab can display which objects are used by the object selected to be exported. For example, when a library is selected to be exported, the **Used By** tab will discover all tables contained within this library.

In the following example (Figure 5) a set of four information maps were originally selected to be exported. When a single information map is selected, the **Dependencies** tab will retrieve all libraries and tables on which this information map depends. Selecting these objects will include them in the export package. After the library is selected, the **Used By** tab can then be used to discover all tables contained within this library.

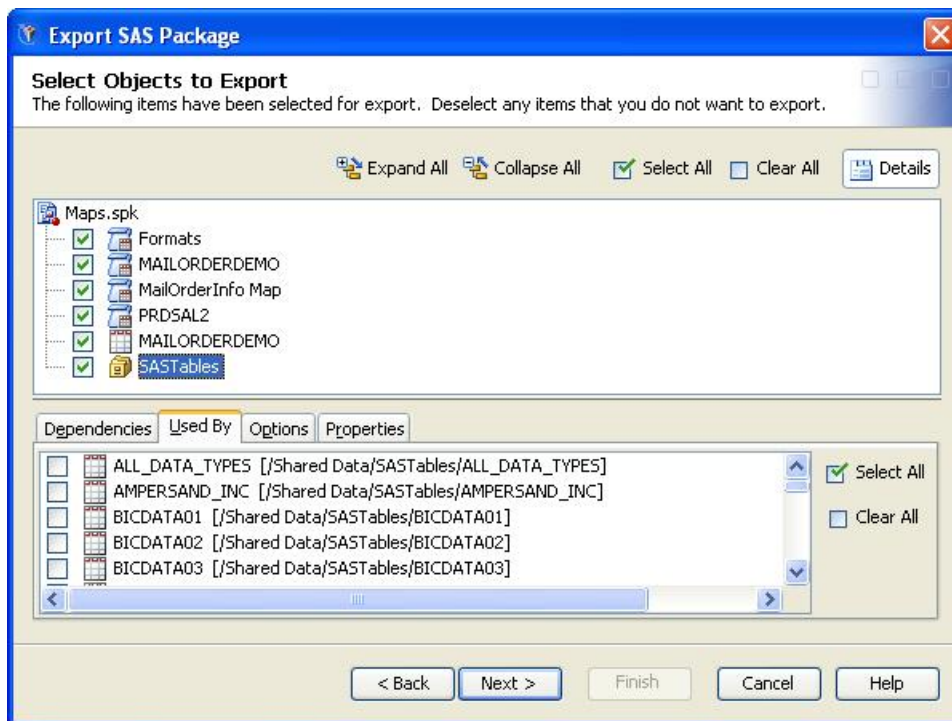


Figure 5. Used By Tab in the Export Wizard

Rather than requiring you to select every object to retrieve its list of dependencies, an additional option was added to allow you to automatically retrieve all dependent objects during the export process. This option, called “Include all dependencies,” is located on the first panel of the export wizard (Figure 2). When selected, this option will recursively retrieve all dependencies on the objects originally selected to be exported. For example, if a report is selected for export, not only will the information map on which the report directly depends be included, but the libraries and tables on which the information map depends will be included as well. When using this feature, the Export Wizard automatically includes all of these dependent objects in the Select Objects to Export window.

BATCH PARTIAL PROMOTION IN SAS 9.2

In SAS 9.2, two new tools have been created to allow you to promote a collection of public objects (including folders, stored processes, jobs, tables, etc.) from one metadata server to another without any user interaction. The same functionality in the Export/Import wizards in SAS Management Console is now available with a command line interface. This enables you to create a repeatable and scheduled process for moving metadata and associated content across environments. These two tools, ExportPackage.exe and ImportPackage.exe, can be found within the “commands” directory of your SAS Management Console install.

BATCH EXPORT

The batch export tool enables you to export a collection of objects from a source metadata server to a SAS package file (.spk). It is possible to export individual objects as well as entire folder structures using this tool. This process exports the metadata for the public objects as well as any content files (.sas files, .srx files, etc.) associated to the public objects. This package file can be used later to import to a target metadata server, either with the new batch import tool or the Import Wizard.

Upon running the batch export tool, three files will be created: the SAS package file containing all metadata and associated content, an export log file, and a substitution properties file. This properties file will be explained in more detail below, but essentially it is the file that will allow you to change which target objects to which you want to connect during the import operation. This file also provides you with the ability to modify object specific values, such as an object’s physical path or description. When it is time to run the import command, the substitution file can be used in conjunction with the SAS package file to import its contents into the target metadata server.

EXAMPLES

- To export the contents of an entire folder, such as the top-level folder “Shared Data”:

```
ExportPackage.exe -profile MetadataServerProfile -package SharedData.spk -objects  
"/Shared Data"
```

- To export a specific object within a folder, such as a library definition in the subfolder “Libraries” under the “Shared Data” folder:

```
ExportPackage.exe -profile MetadataServerProfile -package OrionGoldLibrary.spk -  
objects "/Shared Data/Libraries/OrionGold(Library) "
```

- To export all reports (and their dependencies) in the entire metadata content tree that have been updated since the beginning of the year:

```
ExportPackage.exe -profile MetadataServerProfile -package Reports.spk -objects "/"  
-types "Report" -modified -since "Year to date" -includeDep
```

BATCH IMPORT

Any SAS package file created with the batch export tool, the Export Wizards in SAS Management Console, or SAS Data Integration Studio can be used by the batch import tool to import into a target system. As with using the Import Wizard, ensure that your target environment contains the necessary dependencies. For example, all server definitions, directories, etc., need to exist on the target in order for the import process to run as expected. Missing content on the target system could potentially cause the import to be incomplete or fail altogether.

If the source and target environments have the same setup, then a package file can be imported without any additional work. In this case, the package file can be imported using all of its default settings. For example, if a stored process was associated to the “SASApp” server on the source system, by default, it will be connected to the “SASApp” server on the target system.

However, assuming that the source and target systems are different, the batch import tool provides tremendous flexibility in allowing you to manage the mappings between the two environments. The following discusses this process.

The batch import tool uses a substitution properties file as input to assist you in managing these mappings. This file can contain two types of properties: connection properties and substitution properties. Connection properties are used to describe the connection between an object being imported and an existing object on the target metadata server. Substitution properties, on the other hand, allow for modifications to object-specific values, including descriptions, physical paths, port numbers, etc.

The following example of a properties file contains the different connections referenced by a package file containing information maps, stored processes, and a library.

```
[Connections: Application Server]  
ApplicationServer[1].SourceName=SASApp  
ApplicationServer[1].TargetName=SASApp  
  
ApplicationServer[2].SourceName=Pooled Workspace Server  
ApplicationServer[2].TargetName=Pooled Workspace Server  
  
[Connections: Library]  
Library[1].SourcePath=/Shared Data/AutomationData/AutomationData(Library)  
Library[1].TargetPath=/Shared Data/AutomationData/AutomationData(Library)  
  
[Connections: Table]  
Table[1].SourcePath=/Shared Data/AutomationData/CLASS(Table)  
Table[1].TargetPath=/Shared Data/AutomationData/CLASS(Table)  
  
[Connections: Source Code Repository]  
SourceCodeRepository[1].ApplicationServer=ApplicationServer[1]  
SourceCodeRepository[1].SourceDirectory=C:\\SASServers\\ReportRepository\\Foundatio  
n\\ReportStudio\\Shared\\StoredProcesses
```

```
SourceCodeRepository[1].TargetDirectory=C:\\SASServers\\ReportRepository\\Foundation\\ReportStudio\\Shared\\StoredProcesses
```

```
[/Products/SAS Intelligence Platform/Samples/STP Samples(Library)]  
DirPath=C:\\Program Files\\SAS\\SASFoundation\\9.2\\inttech\\sample
```

When this properties file is created during the export process, all source and target values will match by default. You can make any necessary changes to this file to alter the target mappings. To update the connection properties settings, modify the properties labeled “Target”, such as “TargetName” or “TargetPath”. All of the object-specific settings are included in sections labeled by an object’s path. These settings can be updated by simply modifying the property value. For example, to change which path a library uses, edit the value of the “DirPath” property.

If you are uncertain as to what changes import the package will cause, use the “-noexecute” option on the import command. This option will allow you to view the collection of objects being imported, via the import log, without applying any changes to the target environment.

At the conclusion of the import process, a log is generated containing information on the specific objects that were imported, the content on the target server to which the objects were connected, as well as any warnings or failures that might have occurred.

EXAMPLES

- To import the contents of the “SharedData” package created by the export process:

```
ImportPackage.exe -profile MetadataServerProfile -package SharedDataFolder.spk  
-target "/Shared Data"
```

Since a substitution properties file is not specified, all target mappings will be selected using their default settings.

- The same package can also be imported with an external substitution property file used to manage the source and target connection mappings.

```
ImportPackage.exe -profile MetadataServerProfile -package SharedData.spk -target  
"/Shared Data" -subprop SharedData.subprop
```

- To import only the new libraries contained in the “SharedData” package, along with their access controls:

```
ImportPackage.exe -host "HOST_NAME" -port "PORT" -user "USER_NAME" -password  
"PASSWORD" -package SharedDataFolder.spk -target "/Shared Data" -newOnly -types  
"Library" -includeACL
```

This command allows you to specify the server credentials without needing a server connection profile. Also notice the use of the “-newOnly” and “-types” options. These options can be used to signal the batch import tool to only include new libraries in the import process.

- To verify what will be imported with the “Reports” package, including what changes will be applied based on the substitution properties file:

```
ImportPackage.exe -profile MetadataServerProfile -package Reports.spk -target "/" -  
subprop -noexecute
```

This creates an import log without applying any changes to the target environment.

CONCLUSION

Establishing a formal change control process is essential for managing change within the SAS metadata architecture. SAS provides tools to help you implement the steps in that process. In SAS 9.2, we’ve delivered additional enhancements to make that process more efficient and effective:

- Partial promotion support for more metadata objects
- Filtering capabilities for exporting metadata content for targeted objects or incremental changes
- Batch capabilities to support a repeatable, schedulable process

By combining best practices for process and the promotion tools provided in the SAS® 9 environment, you can establish an effective change control system for metadata-based content.

REFERENCES

SAS Institute Inc. 2006. "Best Practices for SAS®9 Metadata Server Change Control." Available at <http://support.sas.com/resources/papers/MetadataServerchngmgmt.pdf>.

SAS Institute Inc. 2006. *SAS® 9.1.3 Intelligence Platform: System Administration Guide*, Fifth Edition. Available at support.sas.com/documentation/configuration/bisag.pdf.

Stickel, Edward. "Change Control vs. Change Management: Moving Beyond IT." March 14, 2005. Available at www.itsmwatch.com/itil/article.php/3367151.

CONTACT INFORMATION

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

Eric Bourn, Metadata Integration Technologies
SAS Institute Inc.
SAS Campus Drive
Cary, NC, 27513
(919) 531-9186
eric.bourn@sas.com

Diane Hatcher, Global Technology Practice
SAS Institute Inc.
SAS Campus Drive
Cary, NC, 27513
(919) 531-0503
diane.hatcher@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.