

SAS® Metadata Security 101:

A Primer for SAS Administrators and Users Not Familiar with SAS

Charyn Faenza, F.N.B. Corporation

ABSTRACT

The purpose behind this paper is to provide a high-level overview of how SAS® security works in a way that can be communicated to both SAS administrators and users who are not familiar with SAS. It is not uncommon to hear SAS administrators complain that their IT department and users just don't "get it" when it comes to metadata and security. For the administrator or user not familiar with SAS, understanding how SAS interacts with the operating system, the file system, external databases, and users can be confusing. Based on a SAS® Enterprise Office Analytics installation in a Windows environment, this paper walks the reader through all of the basic metadata relationships and how they are created, thus unraveling the mystery of how the host system, external databases, and SAS work together to provide users what they need, while reliably enforcing the appropriate security.


INTRODUCTION

In any company SAS System owners and administrators need to communicate the basics of SAS security. One reason is to introduce security to future system administrators. It is often the first step in training a SAS power user who will be taking on the task of managing SAS users and the SAS metadata environment. Analysts and other general SAS users, however, can also benefit from understanding how security impacts their ability to access, read, and write data. Other business partners that need to understand how the SAS security framework fits into the overall corporate strategy, who may or may not know how to use SAS, include IT, Audit, and Compliance staff.

In order to provide each of these stakeholders a clear overview of how security works, this paper examines a basic use case. It is important to note, however, that SAS metadata security is highly customizable so the configurations and examples in this paper are presented only for illustrative purposes. Individual configurations and recommended settings will vary based upon a site's unique requirements.

SAS CONFIGURATIONS & THE FILE SYSTEM

The environment that is used throughout this paper is a standard installation of SAS 9.4 Office Analytics on a single Windows Server using Windows Active Directory as the authentication provider. The server has four drives with J:/ as the primary location for user created content as seen in **Figure 1**.



There are many settings that are configured when SAS is installed that have an impact on your security choices. In the scope of this white paper we examine three of the configuration choices:

- The file path for the **SASUsers** folders
- The Workspace Server Definitions
- The Logical Workspace Server Security Options

Figure 1 – Server Configuration

File Path for the SASUsers Folders

In this installation, the default location for **SASUsers** has been changed to **J:\SASUsers** instead of the default Windows user profile location, typically **C:\Users\userID\Documents**. As a result, the folders that are created for each must be managed by Active Directory as they are exposed on the file system instead of secured in the users Windows profile by the operating system.

Workspace Server Definitions

As a part of the Workspace Server configuration the server that will be used to create the user connections is identified, in this case as **SASOA.MYDOMAIN.US**. Once the server is selected, using the advanced options, the administrator can specifically identify the file path that is presented to users through the SAS Enterprise Guide application as shown in **Figure 2**. Users who are granted the capability **'Open Files from SAS Server'** and the appropriate security are able to navigate and open or import files from this location from within SAS Enterprise Guide.

In this environment **J:\Files (\SASOA.MYDOMAIN.US\J\$Files)** is identified as the content location.

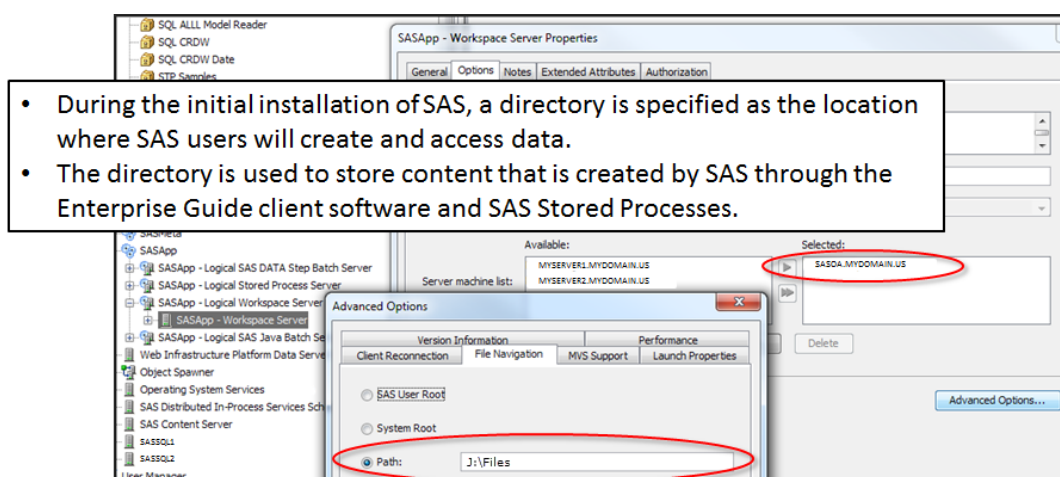
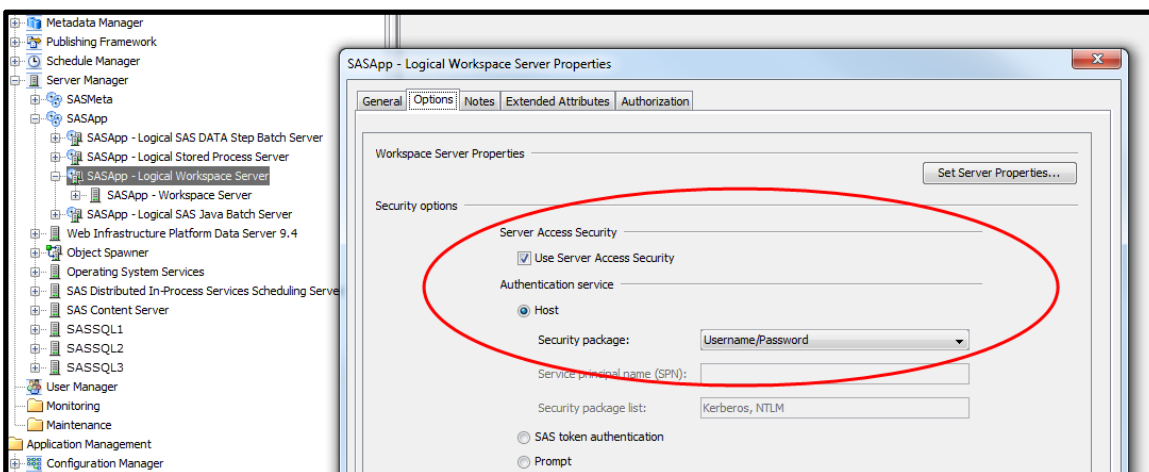


Figure 2 – Workspace Server Definitions

Logical Workspace Server Security Options

By default **'Use Server Access Security'** and the **'Username/ Password'** host authentication security package on the SAS Logical Workspace Server is selected. This configuration tells SAS that users will be authenticated through Active Directory on the host system.



THE THREE A'S OF SECURITY

The first place to start a discussion of security is a review of the three As: Authentication, Authorization and Audit. The *SAS® 9.4 Intelligence Platform: Security Administration Guide, Second Edition* defines the terms as follows:

Authentication:

The process of verifying the identity of a person or process for security purposes.

This answer the question “Who are you?”

Authorization:

The process of determining which users have which permissions for which resources. The outcome of the authorization process is an authorization decision that either permits or denies a specific action on a specific resource, based on the requesting user's identity and group memberships.

This answer the question “What are you allowed to do?”

Audit:

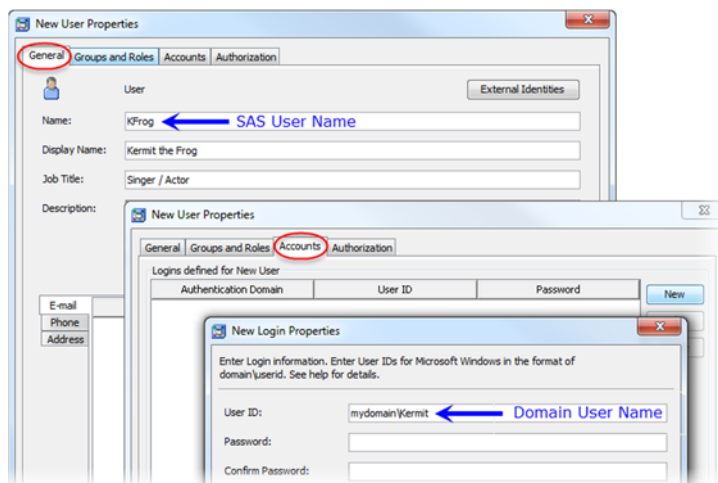
The process of logging of security-related events.

This answer the question “What did you do?”

Managing SAS logging functionality is covered in detail in numerous other whitepapers. The scope of this discussion focuses on the subjects of Authentication and Authorization.

SAS USERS

The first step in establishing the security program is to identify and set up the users. In this example, the users are already defined in Active Directory in the domain ‘**MYDOMAIN**’; however, all users must also have unique identity set up in SAS. For simplicity, the user names in SAS were created to exactly match the user names in Active Directory; however, this is not required. The identity in SAS is paired with the identity in Active Directory on the Accounts tab of the User Properties in SAS Management Console (**Figure 4**).



While the users exist and must be managed in both realms; the enforcement of corporate password policy resides solely with the authentication provider. SAS pairs the SAS user name with the domain Username / Password credentials provided when the user logs in to the system.

Figure 4 – Linking the Domain User Credentials to the SAS User

AUTHENTICATION

Once a user has been created in SAS and linked to their domain account they are able to log on to the SAS system. Opening up a session in Enterprise Guide results in a cascade of events (**Figure 5**).

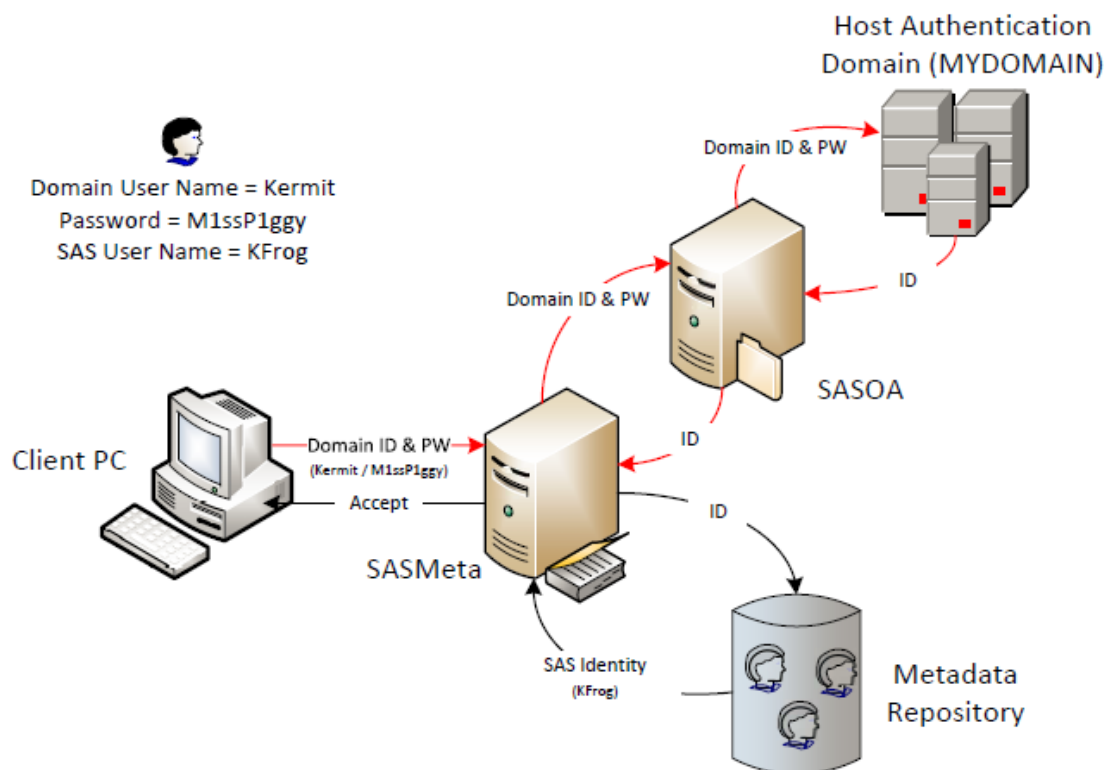


Figure 5 – The SAS Authentication Process

- Users input their domain user name and password through their client program (i.e. Enterprise Guide) which is taken up by the SAS Metadata Server.
- The SAS Metadata Server passes these credentials to the host server.
- The host server authenticates the users in Active Directory according to corporate policy and returns the authenticated ID to the SAS Metadata Server.
- The SAS Metadata Server then resolves the authenticated ID to a specific SAS identity.

Once users are authenticated they are able to start using SAS; however, the process of authentication only verifies the user's identity, who the user is, not what the user is permitted to do. In any environment, users are often going to be working with data that resides on the SAS server as well as data that resides in external data bases, each with their own security protocols.

SAS Authentication Domains

Regardless of whether the access is requested in code or through a metadata library, access to an external database is governed by the database. In virtually every environment it is undesirable to keep user names and passwords in open code, even if the password is encrypted. To avoid this, database access can be managed using SAS authentication domains and user groups. The *SAS® 9.4 Intelligence Platform: Security Administration Guide* provides the following definition:

Authentication Domain:

A name that facilitates the matching of logins with the servers for which they are valid.

Simply stated, when a user attempts to access a server other than the SAS server they are working on, The SAS Metadata Server reviews the list of authentication domains associated with the user ID to determine if they have a login to the desired server. If a match is found, the username and password stored in the authentication domain is be used to authenticate the user to the server.

In code this:

```
User=SASFinanceUser  
Password="{SAS002}99XX999XXXX9X999999X9XX99X999X9X"
```

Is replaced with this:

```
AuthDomain="Finance Database User";
```

To simplify the process for both users and administrators, the authentication domain can be assigned to a user group instead of a user, provided this is consistent with the security policies of the company. Any user that needs to access that resources is then simply added to the security group. Whether the authentication domain is assigned to a person or a group, it prevents the risk an exposed password presents and also secures the code to prevent unauthorized execution.

AUTHORIZATION

The SAS Environment

Once the user is authenticated, access to data and metadata is provided based upon the individual level of authorization for each user. Authorization, or the ability to perform a specific action on a specific resource, is granted first by the host system (or external database system) and then by the SAS through the capabilities and permissions assigned in metadata. When users access information through a metadata driven interface, such as SAS Enterprise Guide, the combination of host permissions and SAS permissions then determine what the user can see and what actions the user can perform.

The following conditions are applied when accessing data:

In order to access data directly one condition must be met:

1. The UserID must have all of the necessary host-permissions to access the data

To access data through metadata in SAS two more condition are added:

2. The UserID must be able to connect to the SAS metadata server
3. The UserID must have all of the necessary metadata permissions to access the data

As shown in **Figure 6**, all access starts with authentication, shown in blue; however, there are two distinct ways the SAS user can interact with the system once authenticated. Through the first path, metadata authorization, shown in green, the metadata server provides information about the SAS permissions established in the Metadata Repository and the level of access granted to the files and databases through registered libraries, stored processes and other metadata objects.

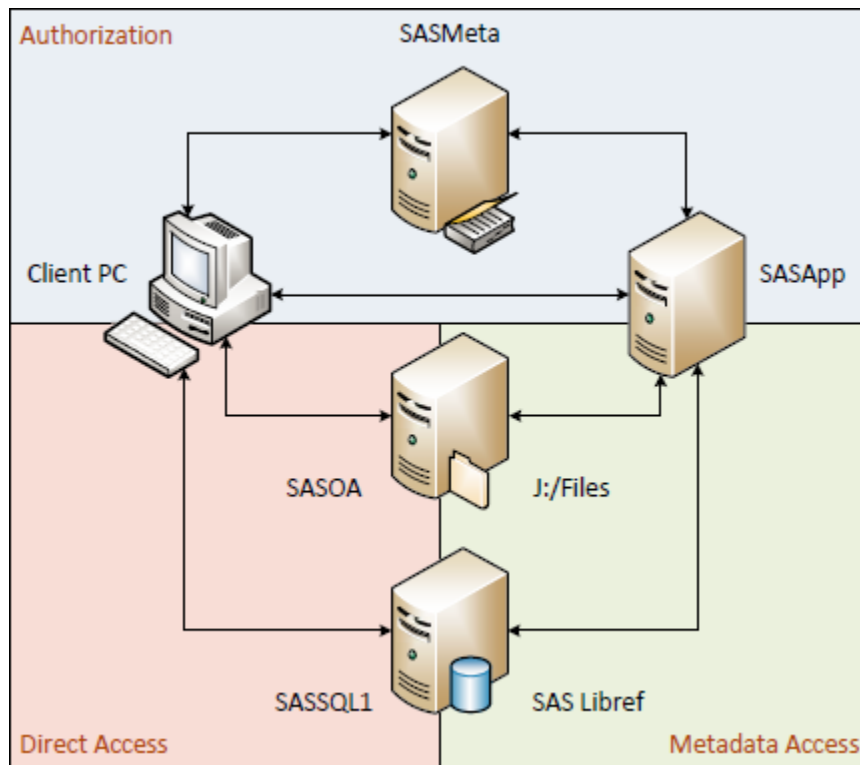


Figure 6 – Data Access Paths

Through the direct access path shown in red, however, the SAS user bypasses the metadata security layer and interacts with the server directly. Users can create libraries directly in their SAS code which then allow them to read or write data based on the permissions granted in Active Directory, regardless of the permissions in the metadata layer. Similarly, users that have the capability **‘Open Files from SAS Server’** enabled in their SAS Enterprise Guide role can view the server in the SAS Application tree and navigate to the server to import or open files directly.

SAS Libraries

In order to understand how SAS grants authorization, it is important to first understand what a SAS library is. By definition a SAS library is a group of SAS files that are stored in the same directory and are accessed by the same engine (a SAS component that can read from or write to a file). SAS has a large number of engines to choose from including Oracle, DB2, Hadoop, Microsoft SQL Server, XLSX, and more. In the example environment there are connections that use the Base SAS engine, which can reads from and writes to SAS datasets and connections that use the SAS/Access Interface to ODBC, which can read from and write to data from applications that are ODBC compliant.

The **LIBNAME** statement is a statement that associates a SAS library with a libref (a shortcut name) and lists the characteristics of the library such as engine type, connection details, and login credentials. For example:

The statement:

```
LIBNAME FIN BASE "J:\Files\Finance\Data";
```

Refers to the following location:

\\saso.mydomain.us\J\$\Files\Finance

The statement:

```
LIBNAME FINDB ODBC CONNECTION=UNIQUE DATASRC=SASSQL1_READER SCHEMA=dbo  
AuthDomain="Finance Database Reader";
```

Refers to the following location:

\\sassql1.mydomain.us\J\$\Data\SAS\FinanceDB.mdf

Writing the libname statement in code establishes a SAS library that uses the identified connection to grant users access to data in the specified location. This is not only true for the DBMS connections, which use the SAS/Access Interface to ODBC but for connections to collections of SAS data sets using the native Base SAS engine as well. Commonly used connections can be stored in metadata as shown in **Figure 7**. To the business users, these libraries look identical, despite having two very different underlying data types.

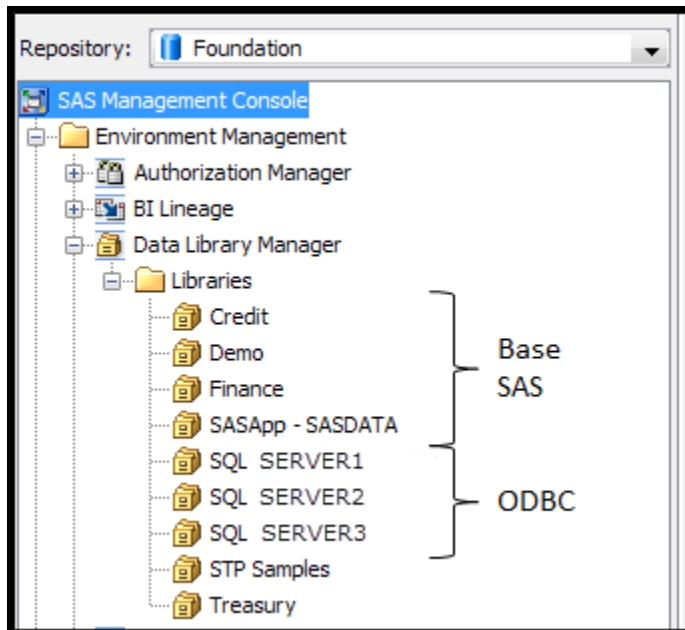



Figure 7 – SAS Libraries in Metadata

DBMS Security

In order for users to access data that is located in an external database, they must have the appropriate permissions on the database. When interacting with the database, SAS has no ability to override the database administrator's security settings. It is the database administrator who controls who can access, update, create and delete objects and records in the database. Access to the database is first created through a standard Windows ODBC connection, which is then referenced in a SAS libname statement or directly in code using a SQL pass-through statement. Regardless of how the access is obtained, it is always subject to the security on the database.

SAS DIRECTORY STRUCTURES



Two of the directories on the server are locations that are only accessible to the administrators (**Figure 8**):

- **J:/SAS94Admin**
- **J:/SAS94Installs**

The remaining two directories are accessible to users, but in two very different ways.

Figure 8 – Sample Server Directories

J:/SASUsers is the directory that was configured to contain each user's personal folder, accessible through their **SASUSER** library in SAS Enterprise Guide when they log in. When enabled, this library can be used by the SAS Enterprise Guide user to store private SAS content. Behind the scenes, on the server, SAS creates a folder in the designated **SASUsers** directory for each user the first time they log in, which is then used as the file path assigned by that user's **SASUSER** library (**Figure 9**).

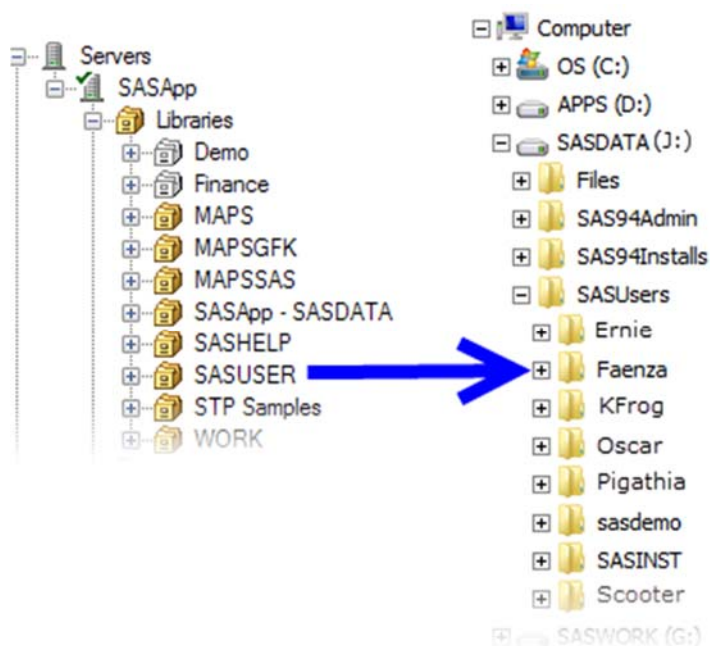


Figure 9 – The SASUSER Library

The default location for this directory in Windows is typically '**C: \Users\user\Documents**' and when the default location is used, the files are secured by the Windows operating system within the user's identity. By selecting an alternate location for the **SASUSER** library, the responsibility for managing permissions shifts from the Windows operating system to Windows Active Directory. If the **SASUsers** directory is left open to all SAS Users, the files that are stored in the individual user folder continue to remain private in metadata; each user only sees the files own **SASUSER** library location; however, they are not secure. An advanced user that is able to authenticate to the server would be able to access files using the direct method.

J:/Files is the content directory for the SAS Users. It is where the general department libraries are stored and is the location that was assigned during configuration as the file path that is made visible in the server tree to SAS Enterprise Guide users that have the '**Open Files from SAS Server**' capability. When the user accesses this location through SAS metadata using a SAS metadata library, both the host security layer and SAS metadata control whether the user can read or write content, or whether the library is even visible. When the user accesses the location directly, it is Active Directory that determines whether the users can read or write content. This can be demonstrated in SAS Enterprise Guide quite easily.

In **Figure 10** the SAS User is assigned the SAS Programmer role in Enterprise Guide which has the '**Open Files from SAS Server**' enabled. The user's metadata security group is 'Finance' which only permits access to the data in the Finance library, which points to **J: /Files/Finance** on the server. The host security, however, does not restrict the user and permits access to all of the subfolders of F:/Files. When viewed through Enterprise Guide, only the Finance library is visible, consistent with the metadata authorization, but all of the subfolders of **J:/Files** are visible, consistent with the host authorization. Should the user choose to he or she can read from or write to a directory they should not have access to, such as the 'Credit' directory.

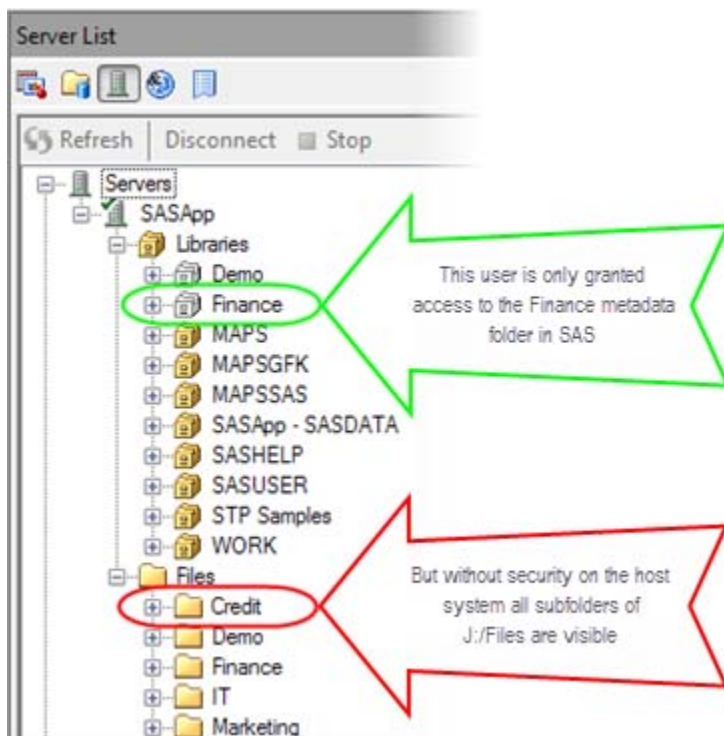


Figure 10 – Metadata Permissions vs. Direct Access Permissions

A closer look at the relationship between the physical file folder structure and the metadata file folder structure is shown in **Figure 11**. To make things less confusing for the users, the metadata tree is set up with the same structure and naming conventions as the physical tree; however, this is not required. For example, the metadata folder 'IT' could be renamed as "Information Technology" in metadata with the metadata library and its tables in the new 'Information Technology' folder pointing to the **F: /Files/IT** location.




Figure 11 – Physical File Folder Structure vs. Metadata File Folder Structure

SAS METADATA BOUND LIBRARIES

As shown in the examples above, there are several ways in which a user without authorization restrictions on the host can use SAS to directly interact with data that they do not have metadata authorization to access. Appropriate host security ensures that each departmental folder is accessible to only individuals within the department; however, they may be instances where an even finer grain of security is required. To accomplish this, administrators can set up a metadata-bound library for the department, which binds all data in the selected library location to a secured SAS metadata object. Once bound, SAS then enforces metadata-layer permissions on the physical data, regardless of how the data is accessed. An overview of the process, from the User Guide "SAS® 9.4 Guide to Metadata-Bound Libraries, Second Edition" is shown in **Figure 12**

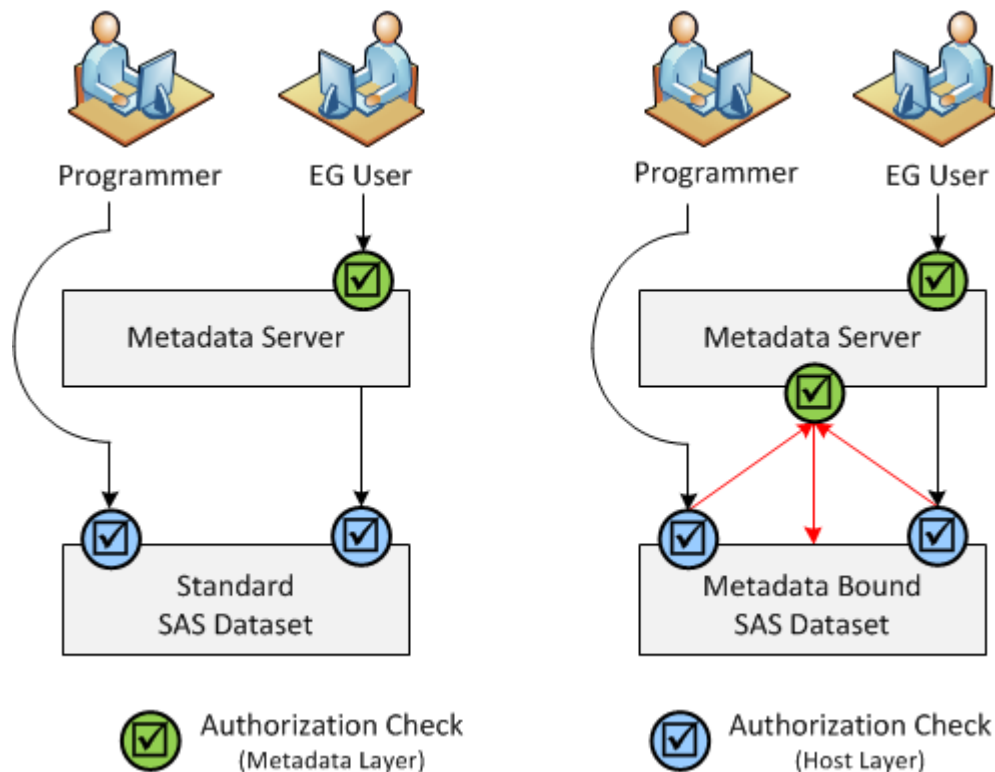


Figure 12 – Metadata Bound Libraries

The obvious benefit to using a SAS metadata-bound library is that it offers much more robust protection than the standard metadata-layer security, including an option to encrypt the data. The intended security is enforced regardless of the access method and the protection is persistent, even if the data is replaced or recreated.

There are, however, limitations to this security feature that should be considered, the most significant being the fact that only base SAS data (data sets and data views) can be bound. Additionally, operating system commands are still able to be used on the data and the metadata-bound library and its contents cannot be located in metadata unless they are also registered as traditional metadata objects.

CONCLUSION

Understanding how SAS interacts with its host server and the host authentication domain is a critical step in building out a strong, effective, and maintainable security program for the SAS environment. Starting with the initial configurations in SAS there are several key elements that must be considered, usually in partnership with non-SAS users, such as the Information Technologists who build and maintain the company's servers and technical infrastructures, as well as SAS business users and administrators. Key points SAS Administrators need to be able to communicate are:

- SAS Users are authenticated by the host and permissions are paired with a SAS identity based upon individual and group permissions
- Access to physical folders and content is controlled by host permissions
- Access to logical folders and SAS data sets in SAS thick client tools (SAS Enterprise Guide / SAS Add-In for Microsoft Office) is controlled by SAS metadata permissions
- Access to external DBMS data is granted through SAS groups and controlled by the DBMS system
- SAS metadata permissions are not applied when a user directly accesses the data in open code
- SAS metadata permissions can be universally enforced on SAS datasets through the use of metadata bound libraries

Providing a clear view of these interactions provides all stake-holders the ability to make better, well-informed decision, and will help the SAS Administrator establish and maintain a highly effective security program.

REFERENCES

- SAS Institute Inc. 2013. *SAS® 9.4 Guide to Metadata-Bound Libraries, Second Edition*. Cary, NC: SAS Institute Inc.
Available at: <http://support.sas.com/documentation/cdl/en/seclibag/66930/PDF/default/seclibag.pdf>
- SAS Institute Inc. 2014. *SAS® 9.4 Intelligence Platform: Desktop Application Administration, Fourth Edition*. Cary, NC: SAS Institute Inc.
Available at: <http://support.sas.com/documentation/cdl/en/bidaag/67972/PDF/default/bidaag.pdf>
- SAS Institute Inc. 2013. *SAS® 9.4 Intelligence Platform: Security Administration Guide, Second Edition*. Cary, NC: SAS Institute Inc.
Available at: <http://support.sas.com/documentation/cdl/en/bisecag/67045/PDF/default/bisecag.pdf>

ACKNOWLEDGMENTS

I would like to extend my gratitude to Paul and Michelle Homes of Metacoda their feedback and guidance in the early stages of this paper as well as their unparalleled support as I learned the ins and outs of SAS Security and metadata management.

CONTACT INFORMATION

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

Charyn Faenza
F.N.B. Corporation
(724) 983-2474
FaenzaS@fnb-corp.com
www.fnbcorporation.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.