Paper 891-2017

Enterprise Data Access Management in a Multi-Tenant SAS BI environment

Chun-Yian Liew, ING Bank N.V.

ABSTRACT

Sometimes it might be beneficial to share a SAS® Business Intelligence environment with multiple tenants within an enterprise, but at the same time this might also introduce additional complexity with regard to the administration of data access. This paper shows one possible setup by sharing a high level overview of such an environment within the ING Bank in the Netherlands for the Risk organization.

INTRODUCTION

There are many aspects to be taken into consideration when deploying a SAS Business Intelligence environment within a large enterprise, i.e. System Architecture, Security, Sizing, Non-Functional Requirements, End-User Training, etc. But for the Risk COO organization within ING Bank there was one aspect that was especially challenging for multiple reasons: data access management.

As the Risk COO organization supports many risk-related processes for the global headquarters and many local units worldwide, the decision was made to deploy a centralized Risk COO SAS BI environment for all these different user groups. One of the data access management challenges was to provision the appropriate access rights from ING Bank Enterprise Identity Access Management (IAM) Service into the Risk COO SAS BI environment.

After many meetings and sessions with a wide variety of stakeholders we managed to come up with a balanced solution. This paper will provide insight how this solution looks like.

ABOUT ING

ING is a global financial institution with a strong European base, offering banking services through its operating company ING Bank. The purpose of ING Bank is empowering people to stay a step ahead in life and in business. ING Bank's more than 51,000 employees offer retail and wholesale banking services to customers in over 40 countries.

Sustainability forms an integral part of ING's corporate strategy, which is evidenced by ING Group shares being included in the FTSE4Good index and in the Dow Jones Sustainability Index (Europe and World), where ING is among the leaders in the Banks industry group.

As at end-2016, ING serves more than 35 million customers.

ABOUT RISK COO

Risk COO is a corporate staff department that integrates risk-related activities on change, support, operations, reporting as well as data and systems. It drives the information architecture and analytics capabilities of the Risk organization. It concentrates on change, data management, reporting, servicing the Risk community and optimizing processes. It also enables Business Intelligence inside the Risk organization.

DATA ACCESS MANAGEMENT

The purpose of data access management is that you are able grant the correct person or system the appropriate access to certain data objects within an IT environment. So how can this be achieved within a SAS Business Intelligence environment? SAS does provide the SAS Management Console which will allow SAS administrators to assign a person to a user group with specific access rights, but how can you control this process from an audit perspective? In other words, how can you validate that the right persons or systems have been correctly granted access to data objects? The process to control this can be provided by Identity and Access Management.

Identity and Access Management (IAM) is the governance, processes and technology that ensures that the right persons and systems have the correct accesses to assets - on a need to know basis. An asset can be an application or a system and also BYOD's. So with IAM you can control who has access to your SAS BI environment.

IDENTITY ACCESS MANAGEMENT

There are different approaches available to implement access control to support Identity Access Management (IAM), but ING Bank has chosen for the Role Based Access Control (RBAC) approach.

RBAC is used in the majority of large enterprise, because it is a convenient and cost-effective way of controlling employee access to resources (i.e. Buildings, IT Systems).

ROLE BASED ACCESS CONTROL

Within ING Bank we try to follow the following 3 recommended principles with regard to RBAC:

- An employee has a minimum number of RBAC Roles
- The total number of RBAC Roles must be minimized
- RBAC Roles are based on business processes and activities

The last principle is sometimes very difficult to follow when a business department supports many different business processes and activities but where everybody in the department is assigned to the same Business Role (BR). For this reason ING Bank had to introduce a new role type: Additional Role (AR), to support application specific requirements.

Figure 1 represents a simplified model of the principle of RBAC.

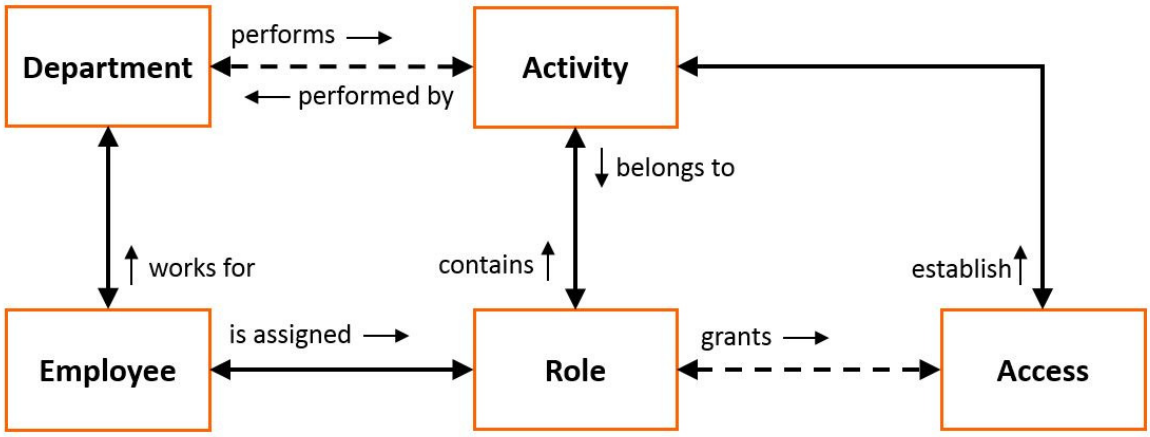

Figure 1 - Simplified RBAC Model

SAS SECURITY FUNDAMENTALS

The SAS Business Intelligence platform provides a metadata-based authorization layer which supplements the security protection defined on the host environment and other systems. You can use the metadata authorization layer to manage access to almost any metadata object.

The SAS Business Intelligence platform supports three different types of access controls:

- Direct Control: Explicit, based on individual grants and denials
- Direct Control: Access Control Template (ACT), pattern based grants and denials

- Indirect setting: based on inheriting access controls from someone else, somewhere else or from a special status.

For simplicity, SAS recommends to set access control permissions on containers (such as folders) instead of on individual content objects, whenever possible.

RISK COO SAS BUSINESS INTELLIGENCE IAM MODEL

When you combine the ING Bank IAM RBAC model with the recommendation from SAS to define access controls on containers it becomes clear why ING Bank has chosen for Access Control Templates to define the permissions within the Risk COO SAS BI environment. But by defining the access controls on containers you are implicitly also relying on the indirect setting type of access control as all objects within a container will inherit the permissions of the parent container.

In our setup each functional user group will get their own SAS Application Server context containing a Workspace Server, Stored Process server and Batch server. Within this Application server context, all users belonging to the corresponding user group are in principle in full control (create, update and delete) of all data objects contained within this context. The biggest advantage of this model is that there is little to none additional configuration required to define specific permissions, i.e. at folder level.

Figure 2 shows a more detailed overview of how the access rights for our Risk COO SAS BI environment are provisioned via ING Bank's Enterprise IAM service.

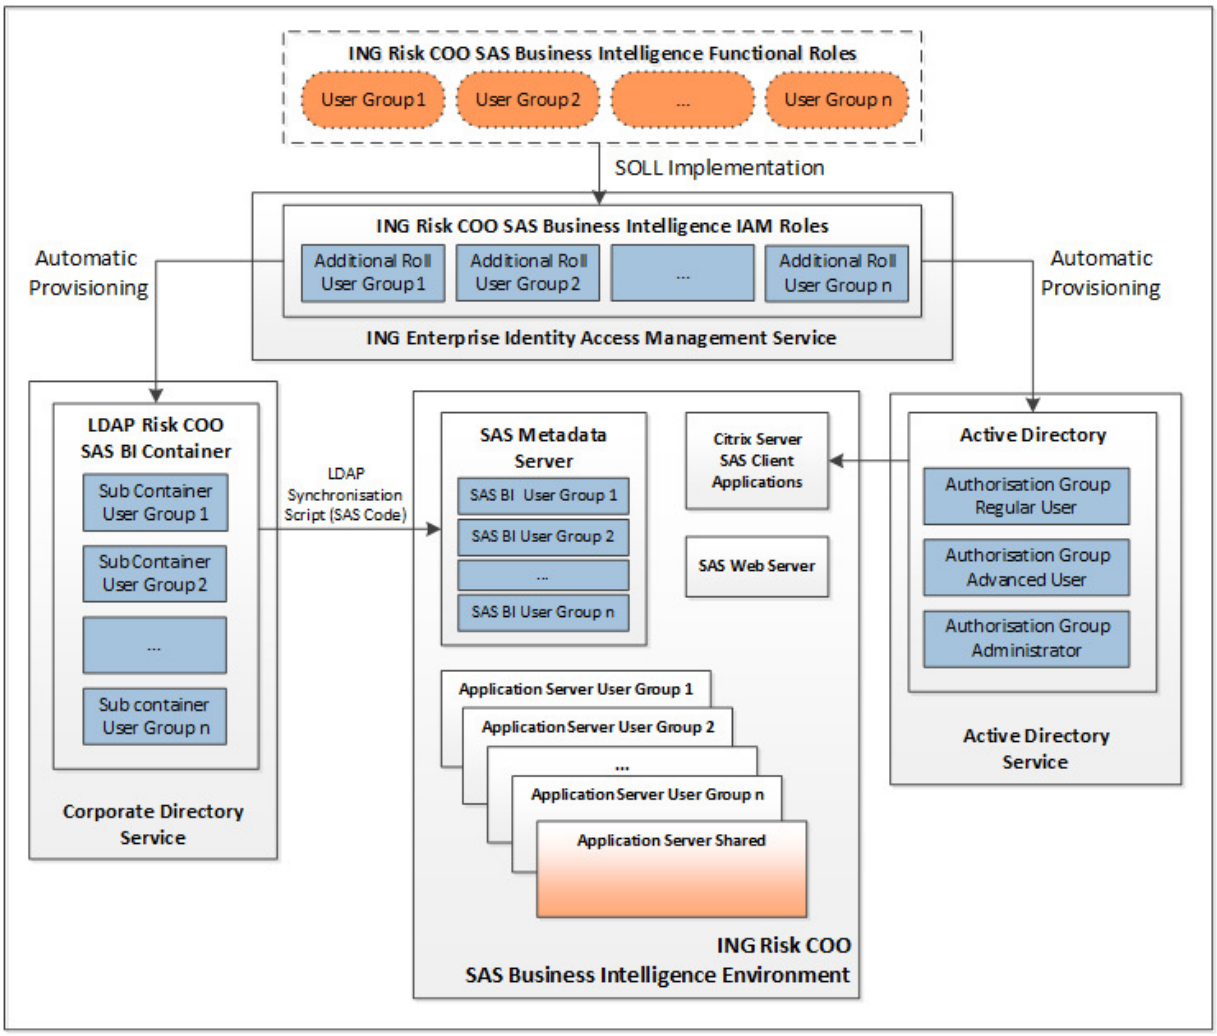

Figure 2 - Risk COO SAS Business Intelligence Identity Access Management Model

The asset owner of the Risk COO SAS BI environment defines the desired functional roles which are documented in a standardized IAM Authorization SOLL document. In the Risk COO SAS BI IAM Authorization SOLL document each functional role is linked to:

- An Additional Role (AR) in the Risk COO SAS BI IAM Solution – within ING Bank a regular Business Role is often still tightly linked to the department of an employee and controlled by the manager of the employee. But as the Risk COO organization required a more fine grained control another type of role was required: Additional Roles are owned by an Asset Owner instead of an HR manager. This will allow the Asset Owner to approve an individual role access requests instead of a manager of a business role. This is necessary as it is not desirable that a complete department is granted access to the Risk COO SAS BI environment by approving a single Business Role.
- A sub container in Corporate Directory service – The LDAP sub container is required to validate the users as member of a group in in the SAS Metadata Server
- An authorization group in Active Directory – The Active Directory group is used to determine which icons of SAS client applications the users are allowed to see on the Citrix Web Desktop

Each iterative change in the IAM Authorization SOLL document has to be approved by the asset owner before the changes are implemented by the responsible IT IAM teams.

NEED FOR A LDAP SYNCHRONIZATION PROCESS

SAS offers a wide range of authentication mechanisms, but as LDAP is one of the leading authentication services within ING Bank the Direct LDAP Authentication mechanism is being used for the Risk COO SAS BI environment. Figure 3 visualizes the SAS Authentication mechanism by way of Direct LDAP Authentication.

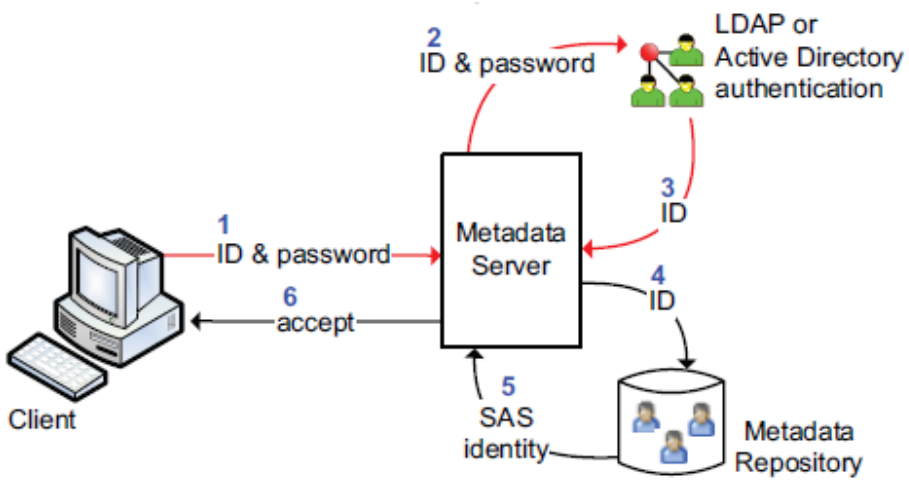

Figure 3 - Direct LDAP Authentication

With the Direct LDAP Authentication mechanism an user can only be granted access to meta data objects when its own identity already exists within the SAS Metadata (User) repository.

We have created a SAS based program to sync all the users assigned to the Risk COO SAS BI environment LDAP sub containers. The program is periodically triggered to synchronize all user identities for each of the LDAP sub containers.

For each LDAP sub container the program will request the list of users and synchronize them with the corresponding SAS User Group within the SAS metadata repository.

NON PERSONAL ACCOUNTS

The most interesting aspect of the IAM Authorization SOLL Document implementation for the Risk COO SAS BI environment is that for each functional user group a separate Workspace Server is created, as part of an Application Server context, with its own data share and Non Personal Account (NPA). The Non Personal Account is required to run SAS processes spawned on its own Linux Workspace Server via the object spawner. This setup will allow users to share data within their own user group and by granting data share access on Linux OS file system level by way of a separate NPA, the data security is implemented on file system instead by SAS meta data configuration. SAS also refers to this concept as mediated access by way of a Privileged Service Account.

If all user groups would have commonly shared the standard Spawned Servers Account on a central stored process, workspace or batch server, each individual user can bypass the metadata security permission defined for a data object. A user then only needs to know the exact file system location of the data object to access the data, even if the data object has no permission defined for the corresponding user group in the SAS Metadata.

SAS illustrates these two different concepts in their SAS 9.4 Security Administration Guide, Third Edition, as shown in Figure 4, where in the left image only metadata security is used whereas in the right image mediated access is used.

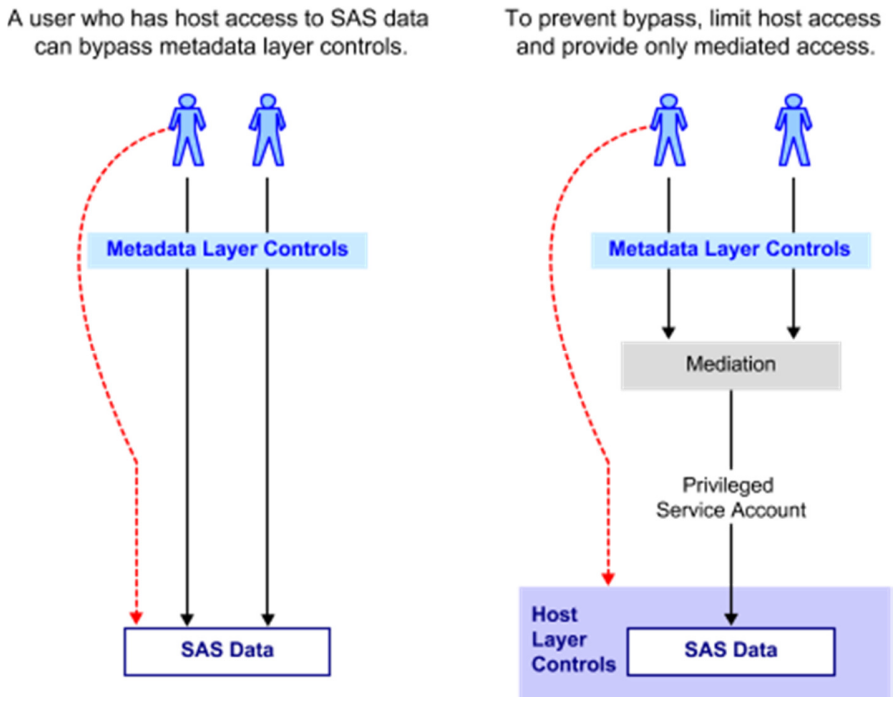

Figure 4 - Host Access to SAS Data

It is important to mention that such setup is not the preferred setup from an auditability perspective, as SAS consultants had recommended to create a separate Linux account for each individual user to enable auditing of individual actions. Unfortunately this does not fit within ING Bank's RBAC model based on roles instead of individual users. So one of the downsides with this setup is that it is somewhat harder to audit individual user actions, compared to a default SAS BI deployment, as the SAS processes are executed by a NPA account.

CONTROLLED VS. AD HOC DATA

As the initial Risk COO SAS BI environment was primarily meant for supporting the credit risk domain, one of the business requirements described the need for one shared data folder containing credit risk related and reference data. This data is sourced from operational data warehouses and is treated as the single source of truth and therefore referred to as controlled data. We achieved this by creating a Shared Application Server context. The SAS Data Integration ETL process extracting the data from the operational data warehouses is granted write access to the Shared Application Server context. All the other User Group Application Server context NPA accounts are granted read-only access on the Shared Application Server context.
As all the User Group NPA accounts are granted full control within their own Application Server context, the users are able to manage their own Ad Hoc data.

But as the scope of usage has extended over time in order to support other risk domains (i.e. Non-Financial Risk, Operational Risk, IT Risk) within the Risk organization, the current setup with one shared data folder is not sufficient anymore.
We are currently still investigating the best approach to restructure our shared and controlled data setup based on the updated business requirements.

CONCLUSION

An Enterprise Identity Access Management solution using Role Based Access Controls can provide a nice starting point for setting up the Data Access Management foundation for a multi-tenant SAS BI environment. But like many other processes, the reality is often much more complex than the model. Our biggest ongoing challenge is that the business users would like to use data access rules which do not always easily fit into the preferred RBAC model.

Some examples in our case:

- An user should have access to multiple user group data shares
- The shared data share is not really shared (i.e. the local units should only be able to see their own local data in the shared data folder)
- Complex data authorization requirements which do not fit the model of defining access permissions on containers

Although different solutions can be defined to solve each individual data access requirement, it is still recommended to dig into the underlying rationale for more complex authorization requirements. Try to look for options to simplify the data access implementation in order to keep your SAS BI environment maintainable.

REFERENCES

SAS Institute Inc. “SAS® 9.4 Intelligence Platform: Security Administration Guide, Third Edition” Available at http://support.sas.com/documentation/cdl/en/bisecag/69827/PDF/default/bisecag.pdf

ACKNOWLEDGMENTS

The author gratefully acknowledges so many of his ING Bank colleagues, SAS consultants and external consultants with whom he has worked. Thank you for your efforts, and thank you that you have been willing to share your knowledge and experience. The author further would like extend his heartfelt thanks and gratitude to the following individuals:

Edgar Vader, ING Bank N.V.
Richard Veenman, ING Bank N.V.
André van de Riet, OCS Consulting B.V.
Mike Vervuren, OCS Consulting B.V.
Therese Ambachtsheer, ING Bank N.V.

RECOMMENDED READING

- Forrest Boozer, SAS Institute Inc. and Christoper Zogby, Zogby Enterprises, Inc. Paper 300-2007 “Case Study in Synchronizing Identities in the SAS® 9 Metadata Server with an Enterprise Security Provider”
  Available at https://support.sas.com/rnd/papers/sgf07/sgf2007-syncmetadata.pdf

CONTACT INFORMATION

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

Chun-Yian Liew
ING Bank N.V.
Chun.Yian.Liew@ingbank.com
http://www.ing.com