SAS6660-2016

Using Metadata Queries To Build Row-Level Audit Reports in SAS® Visual Analytics

Brandon Kirk and Jason Shoffner, SAS Institute Inc., Cary, NC

ABSTRACT

Sensitive data requires elevated security requirements and the flexibility to apply logic that subsets data based on user privileges. Following the *SAS® Visual Analytics: Administration Guide* gives you the ability to apply row-level permission conditions. After you have set the permissions, you have to prove through audits who has access and row-level security. This paper provides you with the ability to easily apply, validate, report, and audit all tables that have row-level permissions, along with the groups, users, and conditions that will be applied. Take the hours of maintenance and lack of visibility out of row-level secure data, and build confidence in the data and analytics provided to the enterprise.

INTRODUCTION

SAS® Visual Analytics Administrator provides an easy way to apply data filter expressions, giving you the ability to control access to particular rows of data. As an administrator overseeing the row-level security, you can quickly get overwhelmed as you scale to more tables, adding more filter expressions required by the business. In this paper, we share tools and code that enable you to monitor conditional grants. There are different levels of knowledge on applying row-level security. If you have been exposed to applying conditional grants, the “Apply row-level security” section will most likely be a review. Also, see the “Recommended Reading” section for additional information about applying row-level security. New techniques for insights into your applied row-level security are in the “Visibility into the Access Control” and “

Reporting on the Security Controls” sections. Once you understand the principles and apply the tools on your environment, you will have the ability to confidently secure sensitive data for analytics while providing audit reports in *SAS® Visual Analytics* on the applied conditions.

APPLY ROW-LEVEL SECURITY

Applying row-level security is fairly easy using SAS Visual Analytics Administrator. From this application, you can find the table on the left-hand navigation window and select the Authorization properties by right-clicking the table. See Figure 1.

Figure 1 - SAS Visual Analytics Administrator Table Authorization

The Authorization properties show the selected table and all of its applied permissions. Find the user or group and click in the Read permission on that row. See Figure 2.

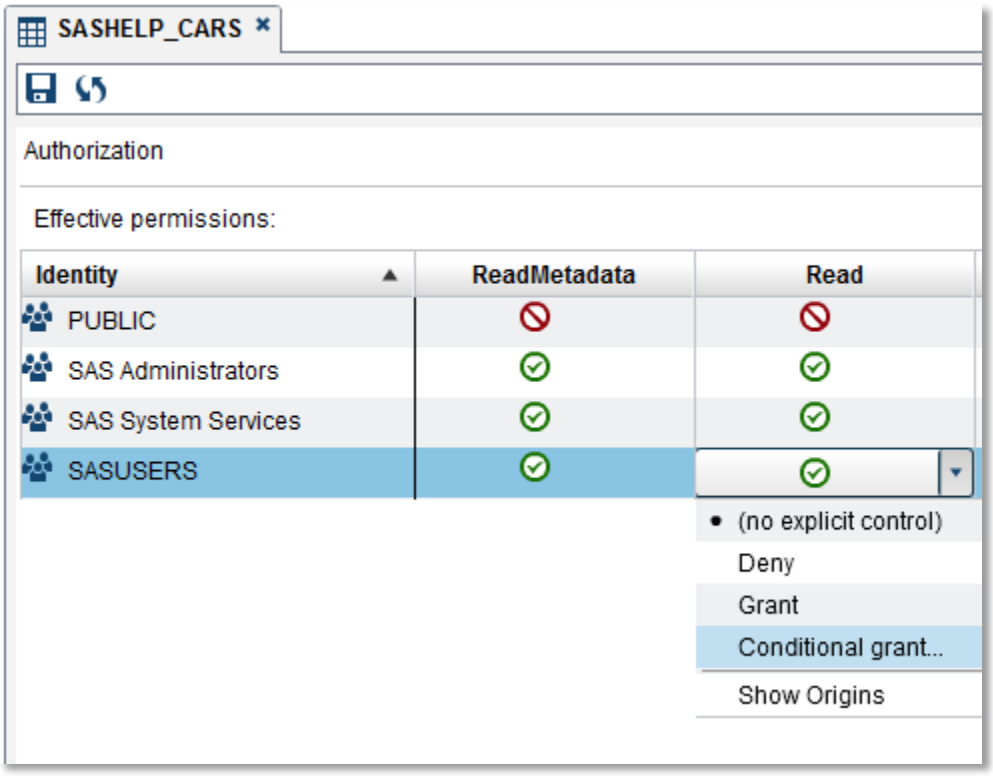

Figure 2 - Conditional Grant Prompt

These conditional grants can be very simple, but they can also be more complex and use the Identity Properties from the logged on user. Here are some examples.

Simple Conditional Grants

Simple conditional grants are just that, simple. They are nothing more than a Boolean condition that limits or filters the rows of data for a given user or user group. For example, if you need to limit the data in a data set that contains global data for a particular user group to only US information. The conditional grant looks like this:

Group: *US Users*

Conditional Grant: *country = 'US'*

When a user is a member of the *US Users* group, the only rows that are returned are the ones that meet the condition of country='US'. All other rows are not available to that user. Likewise, if a user is not a member of the group *US Users,* this condition is not evaluated.

Note: When a user is a member of two or more groups that have conditional grants associated with them, these conditions are appended together with an OR statement. That way the user can see subsets of data from each applied condition that is associated with each group. For example: country='US' OR country='Mexico'.

Complex Conditional Grants

Complex conditional grants contain more than one conditional grant with AND or OR operators. For example, sometimes the data loaded into SAS® Visual Analytics is not standard. Using the preceding

simple example, think about how to handle it if both 'US' and 'USA' are valid text in the country category. It would look something like this:

Group: *US Users*

Conditional Grant: *country = 'US' or country = 'USA'*

Note: Some standard text operators (for example, IN) are not available through the equation editor that was used to apply conditional grants with SAS Visual Analytics Administrator. These functions can be used if the conditional grant is applied through batch tools. The "Batch" section in this paper provides more information about this topic.

Identity Properties

Several identity properties are available to you when the conditional grants are being evaluated. These are like session variables that enable you to evaluate things such as the user's ID or user's name for that logged on user. Table 1 shows the available identity properties that are available when building conditional grants.

Identity Name	Description	Substitute Value	Example
SAS.PersonName	The user name as defined in the metadata	"SUB::SAS.PersonName"	"SAS Demo User"
SAS.IdentityGroupName	The group name as defined in the metadata	"SUB::SAS.IdentityGroupName"	"Group Name"
SAS.IdentityName	The user or group name as defined in the metadata "SUB::SAS.IdentityName"	"SUB::SAS.IdentityName"	"SAS Demo User"
SAS.IdentityGroups	The groups of which a user is a member	"SUB::SAS.IdentityGroups"	"PUBLIC","Group Name"
SAS.Userid	The authenticated user ID	"SUB::SAS.Userid"	SASDEMO@D14887
SAS.ExternalIdentity	A site-specific external identifier	"SUB::SAS.ExternalIdentity"	"8591"

Table 1 - Identity Properties

One identity property that is used often is the SAS.Userid. This is extremely powerful because at the data prep time a new column can be created in the data with a delimited list of user IDs. This is populated by some other business rules, which imitate the security of the data in the source system. A good example is HR data. Only the individual and possibly the individual's organizational hierarchy (manager, director, and so on) should see his/her records. This delimited list of user IDs is fairly easy to generate based on that HR hierarchy. The conditional grant in that case looks something like this:

Group: *HR Reports Consumers*

Conditional Grant: *UserID_List CONTAINS "SUB::SAS.Userid"*

In this example, you might find that some power users should be able to see the entire data set for analysis purposes. A standard Read permission (no conditional grant) will not work if the user is already in the group *HR Reports Consumers* because the conditional grant that is associated with that group will be applied. Therefore, a TRUE Boolean statement on a conditional grant is required for this power user

group to get all the records, regardless of other groups of which these users might be a member. Here is a good trick to use:

Group: *HR Reports All Access*

Conditional Grant: *1 = 1*

Batch

To simplify a more complex security model that uses many conditional grants, use the batch tools for Metadata Authorization found in the *SAS® Intelligence Platform: Security Administration Guide*. Several command-line tools allow you to apply and view conditional grants in batch. Here is a simple example of using the sas-set-metadata-access tool:

Example call for setting permission:

/opt/mis/sas94/SASHome/SASPlatformObjectFramework/9.4/tools/sas-set-metadata-access "/Sample Data/LASR/SASHELP_CARS (Table)" -grant “USA Cars(UserGroup)":Read -condition "Origin = USA"

Here are some benefits in using these batch tools:

- *easy promotion of conditional grants across environments*
- *easy copy/paste conditional grants from table to table*
- *more complex operators available than in SAS Visual Analytics Administrator (for example, IN)*
- *ability to apply conditional grants from a spreadsheet or* SAS® *data set*
- *ability to see multiple conditional grants at once*
- *easy creation of restore points in case something goes wrong*

In fact, provided in the appendix are a few SAS® macros and SAS® Data Integration Studio custom transformations that really simplify the application of conditional grants by leveraging these batch tools. You simply point the macro or custom transformation to an Excel spreadsheet and to the LASR library that contains the tables to which you are applying the conditional grants.

This process uses the Artifacts that we introduce in “Visibility into the Access Control” section. Back to our most complex model, 15+ tables with 80+ conditional grants on each table, this method of applying the conditional grants has been simplified. This job runs every day and the business user that maintains the security does not rely on SAS Visual Analytics Administrator since all the conditional grants are contained in one spreadsheet.

Things to Consider

Applying these conditional grants seem easy at first glance. In some simple cases, they will remain fairly straightforward. But keep in mind, the more complex you make the security model, the more complex it is to troubleshoot and to maintain. You should be aware the following things before implementing row-level security.

1. **Performance**: The evaluation of conditional grants impacts performance. No matter how you view the data (for example, SAS® Visual Analytics Viewer, SAS® Visual Analytics Designer, or SAS® Visual Analytics Explorer), the statement that is generated for the conditional grants will be applied and evaluated on every object associated to that table. Think of it as a filter statement that uses the WHERE syntax. If, for example, you have 10 or so conditional grants on different

groups (one per country) on one table, the WHERE statement will be fairly complex. If the user is a member of all those groups, the conditional grant turns into this:

Country = “US” OR Country = “Canada” OR Country = “Mexico”……etc.

The performance impact is specific to the environment’s capacity, the table size, and the operators that are used in the conditions (IN, =, CONTAINS). Typically the more explicit the condition the better the performance. It is a good idea to perform some benchmarks on how performance is impacted and try to keep this to a minimum.

2. **Space**: This really comes into play when you are preparing your data in order to make the conditional grants work. The great example of this is when using the identity properties (for example, SAS.Userid). What if you have 100 or even a 1,000 users and can see a row of data? If you are using the user ID, that could be 5-6 characters in length plus a delimiter for each user. This is not a huge deal with a small data set, but in SAS Visual Analytics we often are working with tables with millions of rows. Adding a column that is 4,000+ characters in length quickly increases the overall size of a table with millions of rows.
3. **Maintenance**: This is the one area that will be improved significantly by using the batch tools to apply the conditional grants, especially when applying similar or identical conditional grants on multiple tables. See the “Batch” section.
4. **Troubleshooting**: The SAS Visual Analytics Administrator is decent doing some simple troubleshooting. You can easily look at a conditional grant for a particular table. This can get very tedious when you have many conditional grants on a table since you can view only one at a time. The “Visibility into the Access Control” section introduces a way to gather this information and provide it in consumable reports.
5. **Permission Precedence**: You need to understand that security is evaluated in a hierarchal fashion. Being a member of a group directly and indirectly can have very different results. For a deeper dive into identity hierarchies, see the *SAS® Intelligence Platform: Security Administration Guide.*

VISIBILITY INTO THE ACCESS CONTROL

Now that you have applied your conditional grants, looking at all of them at a once is not an option within SAS Visual Analytics Administrator. The following valuable process enables you to gather the necessary information from metadata in a way that is easy to consume.

In detail, a conditional grant creates an access control entry on a metadata table. The filter is an associated condition to the access control entry. You need to understand this level in order to query the metadata to build out a table with all the conditions. Figure 3 shows an example when viewing this hierarchy (**Table>Access Control>Associated Condition**) and the identities with that association in the SAS® Metadata Browser.

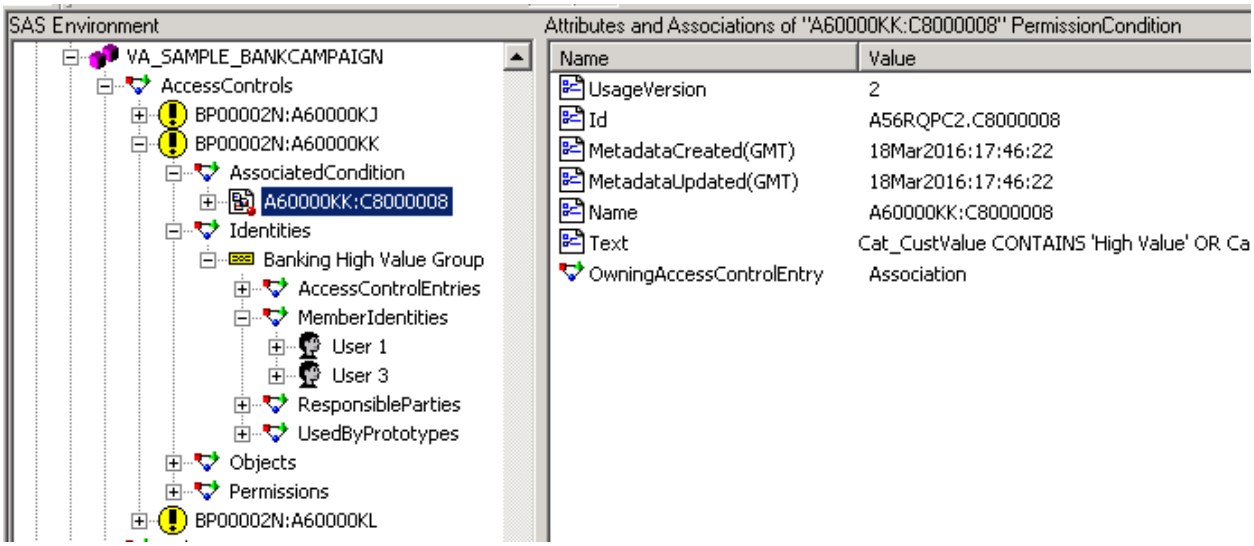

Figure 3 - Metadata Browser on Access Controls

Note: SAS Metadata Browser can be found in SAS® Display Manager under Solution>Accessories>Metadata Browser.

To accomplish this, first pull the needed hierarchy of metadata information together into a collection of artifacts. These artifacts are very beneficial for reporting and monitoring the health of your SAS Visual Analytics environment. As you traverse the metadata to build the Artifacts collection, you will gain intelligence and reporting capabilities on the following objects:

- *Libraries*
- *Tables*
- *Access Controls (conditional grants)*
- *Identities (groups and users)*

Table 2 list the names of the macros that are needed to extract the information for the Artifacts collection.

Program name	Description
MetadataQueryArtifacts.sas	Output: ARTIFACTS and ARTIFACTS_RELATIONSHIPS This is a macro from the SAS® Audit, Permorfance and Management package that has been integrated into SAS® Environment Manager. ARTIFACTS contains table, library and other objects stored in metadata. ARTIFACTS_RELATIONSHIPS contains the linkage between objects. For example, and table is a member of a library.
MetadataQueryLibrary.sas	Output: ARTIFACT_LIBRARIES A data set of all libraries in metadata with identifiable attributes such as engine, host, and port. The data can be used to pragmatically allocate libraries and leveraged for processes that need library information.

MetadataQueryTable.sas	Output: ARTIFACT_TABLES A data set with information about all metadata tables registered in the environment including the metadata ID (the key for joining to other Artifact data sets). This is a logical place to extend the information gathered on each table. Some examples include table size, record length, usage, and in-memory load date.
MetadataQueryConditionalGrants.sas	Output: ARTIFACT_CONDITIONAL_GRANTS Loops through the MetadataQueryTable.sas output data set to build a consolidated list of row filters associated with tables that have access controls. In a SAS Visual Analytics environment, this output data set will most likely contain only LASR tables.
MetadataQueryIdentities.sas	Output: ARTIFACTS_IDENTITES Queries the metadata to build a comprehensive list of groups, roles, and users. This data set is then joined to see the identities (groups and users) associated with a row-level filter.
CreateVAPermissionAudit.sas	Output: VA_PERMISSIONS_AUDIT Joins all the Artifacts data sets into a single data set. Enables you to create custom reports through SAS Visual Analytics Designer or leverage the reports that are provided in the appendix and introduced in the **Reporting on the Security Controls** section.

Table 2 - Core Macros to Gather Information about Access Controls

Note: The code for these macros is contained on the proceedings download site.

The Artifacts collection enables you to use them as input into a single data set to analyze the LASR tables' security and usage. By joining the artifacts together then loading the data into SAS® LASR™, you can use the power of SAS Visual Analytics to visualize and share the content (See Table 2 - Core Macros to Gather Information about Access Controls– CreateVAPermissionAudit.sas). In the reporting section you will find example of SAS Visual Analytics reports using the data collected. Figure 4 is a sample SAS Data Integration Studio job that joins the Artifacts data sets.

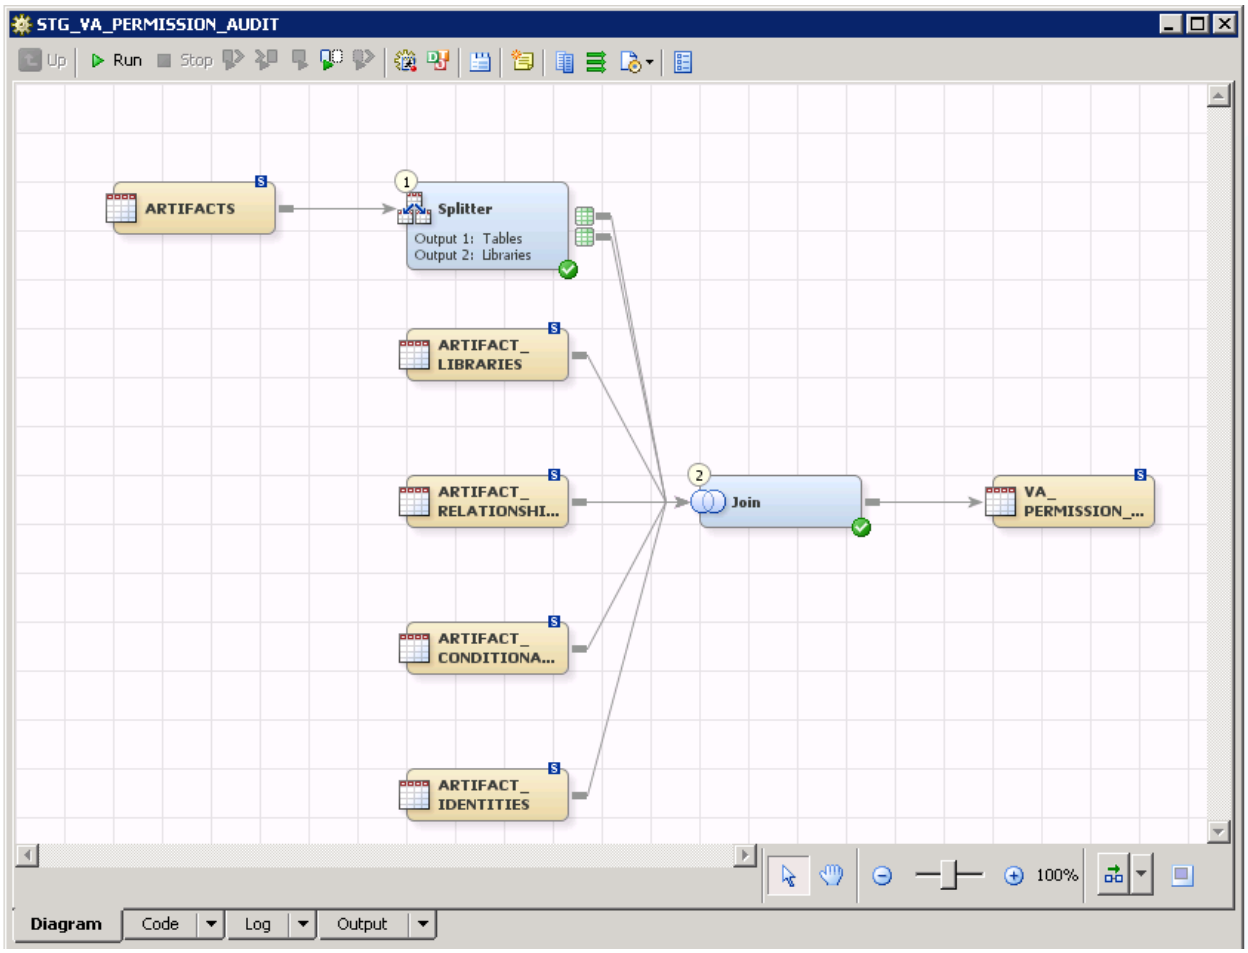

Figure 4 - Create VA Permission Audit Job

REPORTING ON THE SECURITY CONTROLS

Building out custom reports and explorations on the VA_PERMISSIONS_AUDIT allows you to explore the information about libraries, tables, conditional grants, and associated identities. The reporting capabilities benefits you by making it much easier to analyze your environment and assist in troubleshooting your security requirements. This data can answer many high-level questions such as the following:

- *How many tables are collocated in Hadoop as SASHDAT tables?*
- *How many LASR tables do you have in your VA environment?*
- *What conditional grant(s) have been applied to a specific table?*
- *What tables (or rows) can a specific user see?*

Provided are a few custom reports that answer these questions and many others. Figure 5 and Figure 6Figure 6 are screenshots of how you are able to visualize the conditional grants all at once. In SAS Visual Analytics Administrator, this takes many clicks to see all of this information.

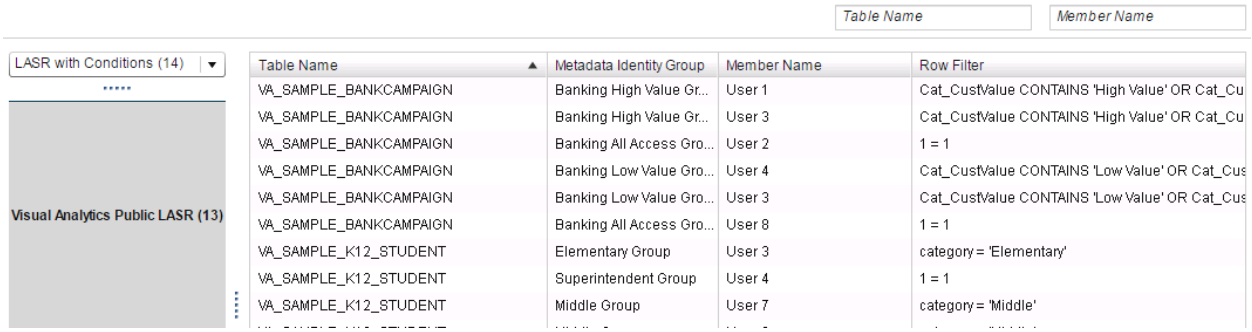

Figure 5 - Row-Level Security Customer Report

Note: Notice how you can see all of the conditional grants for each table at once.

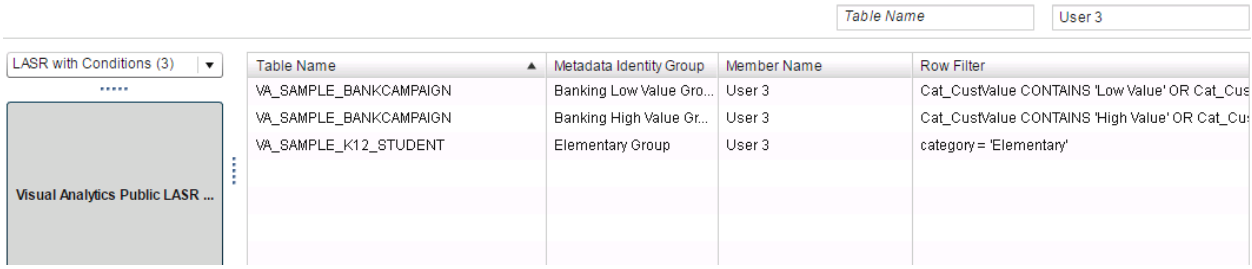

Figure 6 - Row-Level Security Customer Report

Note: Here you can focus on a particular table to see all the conditions applied to a user.

For administrators, the environment's health and governance is important to monitor. By using the data collected in Artifacts, you can gather insight into this the following types of information and share to users without granting access to SAS Visual Analytics Administrator. The Artifacts collection is extendable by adding additional data to further enhance reporting. For example, you can run queries against data tables loaded in SAS® LASR™ to have metrics such as LASR table size, number of observations, and audit info like LASR table actions.

- What is your table size versus table usage? Are largest tables being used? (See Figure 7.)
- What is your total in-memory footprint?
- Are there any duplicate tables on different SAS®LASR™ Analytic Servers? (See Figure 8.)
- From a governance perspective, how can you verify what credentials are used on source libraries at once? List all SAS Authorization Domains (user ID and password containers) on the all source libraries.

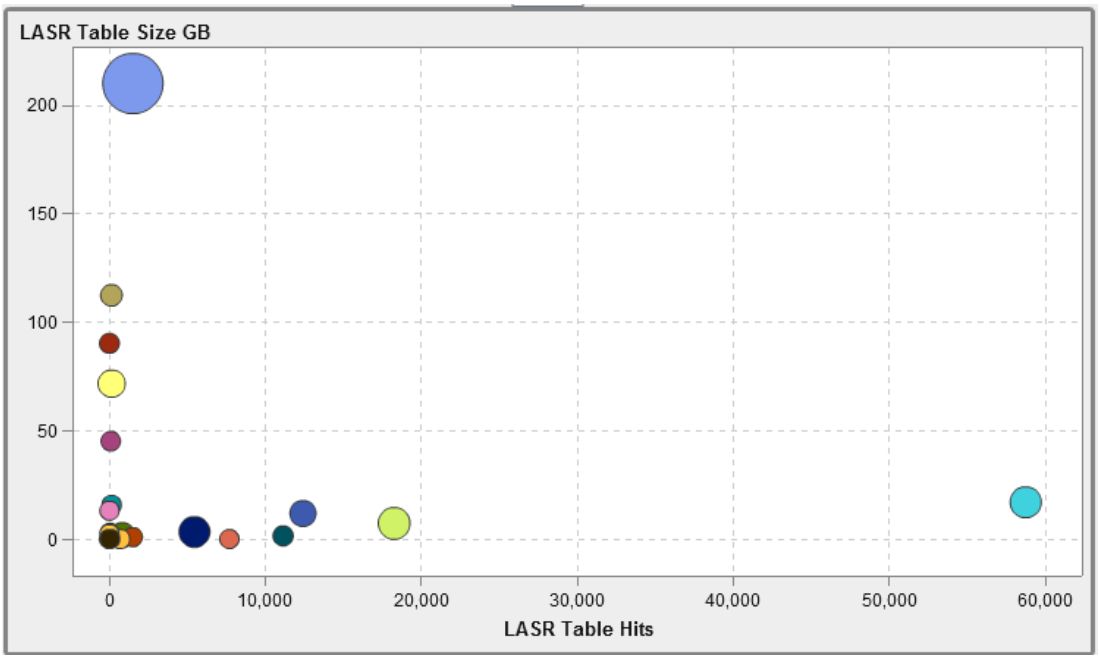

Figure 7 - LASR Table Usage by Size

Note: This table visualizes the size and usage of all LASR tables. When a large table is not being used, it could be a candidate to be dropped from memory to save resources.

midas

Table Name	Port	LASR Name
CANIT_MIDAS_CAN	10022	Canada MIS LASR Server
CRL_MIDAS_2014_CLEANED	10027	Customer Loyalty LASR
CRL_MIDAS_CLEANED	10027	Customer Loyalty LASR
CRL_MIDAS_TICKETS_MERGED	10027	Customer Loyalty LASR
GIS_MIDAS	10004	SDS LASR - misvap00
ITI_MIDAS_ASSIGN	10046	OneIT LASR Server
ITIP_MIDAS_TEMP	10046	OneIT LASR Server
MIDAS_CUSTOMFIELD	10004	SDS LASR - misvap00
MIDAS_CUSTOMFIELDS_ALL	10001	MIS LASR Server - misvap00
MIDAS_DEPLOYMENT	10004	SDS LASR - misvap00
MIDAS_HWO	10004	SDS LASR - misvap00
MIDAS_INTERNATIONAL_WORKLOAD	10001	MIS LASR Server - misvap00
MIDAS_IPD	10004	SDS LASR - misvap00
MIDAS_MOBILE	10004	SDS LASR - misvap00
MIDAS_MOBILE_ASSETS	10004	SDS LASR - misvap00
MIDAS_OIS_TICKETS	10001	MIS LASR Server - misvap00
MIDAS_OIS_WORKLOAD	10001	MIS LASR Server - misvap00
MIDAS_OPEN_HELPDESK	10004	SDS LASR - misvap00
MIDAS_ROOTCAUSE	10001	MIS LASR Server - misvap00

Figure 8 - Duplicate LASR Tables across Multiple SAS® LASR™ Analytic Servers

Note: See the proceedings download site for a package that contains the custom reports from the screenshots above.

CONCLUSION

The conditional grants for row-level filters are very powerful but can get overwhelming when dealing with many tables and complex security requirements. Implementing the tools in this paper will give you better insight with shareable reports to audit the security.

> "SAS® Visual Analytics has totally ***transformed*** the way I can look at our business," says Farrell. "I spend more time following the data and asking ***informed*** questions."
>
> *Carl Farrell, Executive Vice President and Chief Revenue Officer*

SAS uses SAS Visual Analytics to run the business. The techniques shared in this paper are from our experience in SAS IT managing the SAS Visual Analytics enterprise consumed environment which have provided invaluable reporting and auditing. The current consumer base this SAS IT environment covers over ~6,000 users on ~1800 tables at ~11TB of data. Based on the foundation we shared, a quick report shows we are managing 1,156 conditional grants applied over 161 tables.

RECOMMENDED READING

- *SAS® Intelligence Platform: Security Administration Guide*
- *SAS® Visual Analytics: Administration Guide (https://support.sas.com/documentation/solutions/va/73/en/vaag.pdf)*

  Licensed customers can request the user ID and password from SAS Technical Support.
- Williams, Zuzu. 2015. "Row-Level Security and SAS® Visual Analytics." *Proceedings of the SAS Global Forum 2015 Conference*. Cary, NC: SAS Institute Inc. Available http://support.sas.com/resources/papers/proceedings15/SAS1779-2015.pdf.
- Kirk, Brandon. 2015. "Bust Open That ETL Black Box and Apply Proven Techniques to Successfully Modernize Data Integration" *Proceedings of the SAS Global Forum 2015 Conference*. Cary, NC: SAS Institute Inc. Available http://support.sas.com/resources/papers/proceedings15/SAS1824-2015.pdf.
- Shoffner, Jason. 2015. "SAS® Visual Analytics Environment Stood Up? Check! Data Automatically Loaded and Refreshed? Not Quite" *Proceedings of the SAS Global Forum 2015 Conference*. Cary, NC: SAS Institute Inc. Available http://support.sas.com/resources/papers/proceedings15/SAS1952-2015.pdf.
- Hosking, Jerry. 2015. "Someone Changed My SAS® Visual Analytics Report! How an Automated Version Control Process Can Rescue Your Report and Save Your Sanity" *Proceedings of the SAS Global Forum 2015 Conference*. Cary, NC: SAS Institute Inc. Available http://support.sas.com/resources/papers/proceedings15/SAS1890-2015.pdf.

CONTACT INFORMATION

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

Brandon Kirk	Jason Shoffner
100 SAS Campus Drive	100 SAS Campus Drive
Cary, NC 27513	Cary, NC 27513
SAS Institute Inc.	SAS Institute Inc.

(919) 531-3825
brandon.kirk@sas.com

(919) 531-2110
jason.shoffner@sas.com