Paper 1076-2017

Auditing in SAS® Visual Analytics

Elena Muriel, Amadeus Software Ltd

ABSTRACT

Many organizations are using SAS® Visual Analytics for their daily reporting. But as more users gain access to the visual tool, it is easy to lose track of what data is being used, what reports are being accessed, and what elements of the system are classified as critical. With SAS® Visual Analytics comes a governance exercise that all organizations should provision for, as otherwise it jeopardizes its maintenance and performance.

This paper explores the three different auditing areas that can be configured with SAS Visual Analytics and the different metrics that are associated with them. It presents how to configure the auditing, the data sources that are being populated on the background, and how to exploit them to expand your reports beyond the pre-created audit reports.

Consideration is also given to the IT and infrastructure side of enabling auditing mechanisms, with data volumes and archiving practices being at the heart of the discussion.

INTRODUCTION

Many organizations are adopting SAS Visual Analytics as their preferred reporting tool. Your organization could have a simple non-distributed Visual Analytics environment with just a few data sources, or it could be a distributed environment with very complex ETL processes. However, one of the problems that organizations of any scale will face, is the management of the content they produce and its administration.

In Visual Analytics, users are able to create different types of content, each with their corresponding management challenges:

- Reports. Consider if all the existing reports are being used and when reports should be archived or removed altogether from the system. Also analyze what data dependencies exist for the Visual Analytics reports.
- Explorations. Consider if there are explorations that were one-off queries and have not been accessed for some time, and the data dependencies.
- LASR tables. Consider who is consuming loaded tables, what are their refresh rate and if there are tables that could be loaded on demand rather than being available all of the time.
- Queries. When using Visual Analytics queries, consider the output produced by them and if they are being scheduled. Consider the impact that can introduce a change on a Visual Analytics query to the table that supports a report.

Most of these questions can be answered by looking into the activity related to Visual Analytics that is recorded in the SAS audit tables. This paper has been written using SAS Visual Analytics 7.3 and SAS 9.4 (TS1M3).

AUDIT TYPES

With all instances of SAS Visual Analytics there are three administration reports that access the different auditing areas that can be configured with SAS. Consider these available reports as an entry point into the auditing data, as you can improve on them to give you a higher level of detail regarding how users interact with Visual Analytics.

The Administration reports can be found under the following path in SAS metadata:

SAS Folders\Product\SAS Visual Analytics Administrator\Reports\Usage

The table below includes a list of the reports and the data sources that are used in its construction:

Report Name	Description	Data Sources
Administrator overview	Uses data from the middle-tier **audit** service.	AUDIT_VISUALANALYTICS
Relationship report	Uses data from the middle-tier **relationship** service.	RELATIONSHIPS_VISUALANALYTICS
Midtier performance report	Uses agent collected metrics (**ACM**) data.	FILEMOUNTS HOSTPLATFORMS HTTPCHECKS IOMSERVERS METADATASVRS NETWORKINTERFACE TCSSERVERMGRS WEBAPPSERVER WIPDATADB

Table 1 - Visual Analytics administrator reports and tables

Although the definition of these reports is present in all installations of SAS Visual Analytics, the data feeding these reports is not automatically available, and extra configuration is needed to enable each of them.

The mechanism for loading usage data into Visual Analytics is the same for the three different areas. There are some agents that collect the audit metrics, these metrics are stored in a database or set of tables, and extracted at set intervals. The extracted data is dropped into the Append area inside the EVDMLA directory that is located inside the Autoload folders for the Visual Analytics Administrator. Once in that location, the autoload facility will load it into the LASR Server at set intervals ready to be consumed.

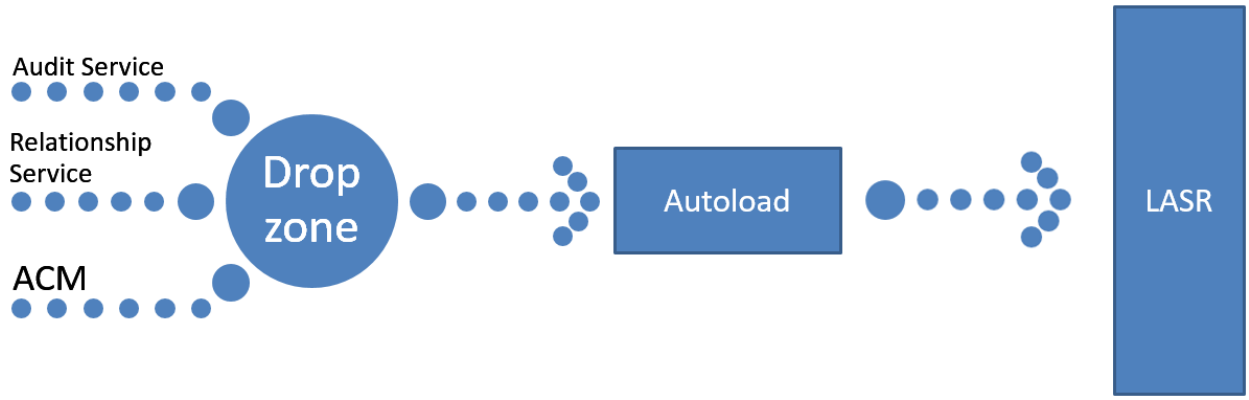

Figure 1 - Audit load process

AUDIT SERVICE

The data provided by the SAS Middle Tier Audit Service is probably the one that most SAS administrators will be interested in. The data captured by the Audit Service includes user, tables and report level information, enabling identification of the most utilized LASR tables, users that are more active and reports most commonly opened.

Capturing audit metrics is switched on/off through SAS Management Console, by accessing the Configuration Management plugin and the following navigation path:

→Application Management

→Configuration Manager

→SAS Application Infrastructure

→Visual Analytics 7.3

and setting the option va.AuditEnabled to true:

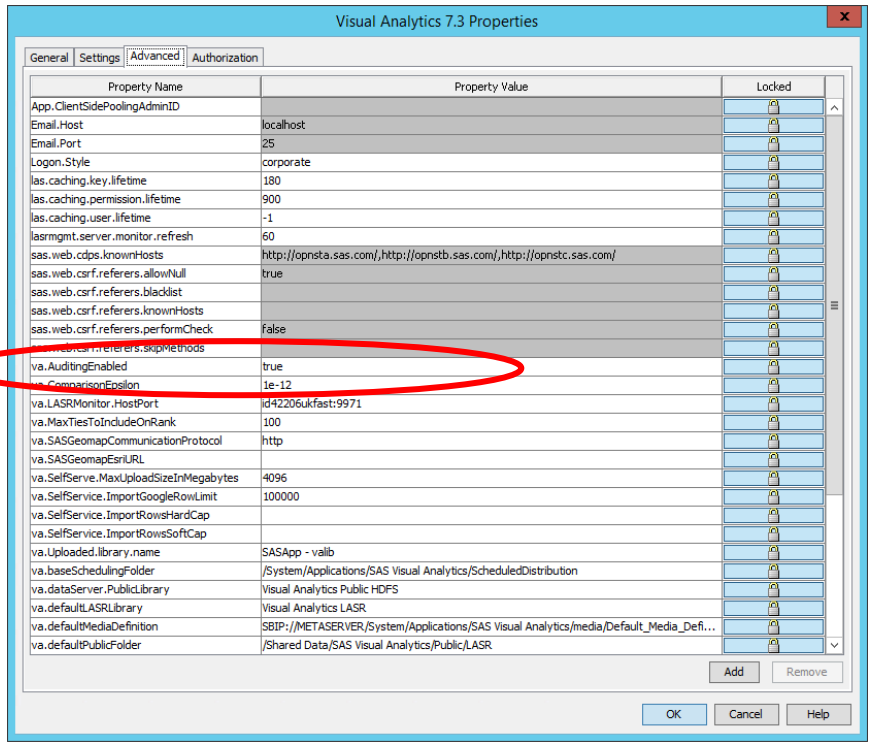

Caption 1 - Enabling the audit service

Once the auditing has been enabled, the system will start collecting data, which is stored in the Web Infrastructure Platform (WIP) SharedServices database. The WIP database is a PostgreSQL (Postgres) database that can be managed by using some of the Postgres admin tools (e.g. pgAdmin III) and from this location data can be extracted. Particularly, the audit data is stored in the sas_audit and sas_audit_entry tables, where the sas_audit table contains one entry per event and sas_audit_entry can contain multiple entries per event (storing different bits of information related to that event).

The auditing process extracts the data from the SharedServices database and moves it into the Append location under:

LevX\AppData\SASVisualAnalytics\VisualAnalyticsAdministrator\AutoLoad\EVDMLA\Append

where the autoload process loads it into LASR. More information on how to configure the auditing and autoload processes for Visual Analytics can be found on the SAS Middle-Tier Administration Guide [Ref. 1].

The final table containing all of the audit data is a SAS data set that can be found under:

\LevX\AppData\SASVisualAnalytics\VisualAnalyticsAdministrator\AutoLoad\EVDMLA\audit_visual analytics.sas7bdat

From all the variables included in this table, it is interesting to look into:

- Action_type: contains a description of what has happened;
- Object_type: describes the type of object that was involved;
- Newtable_name: contains the name of the LASR table that was affected (if any).

Other important variables included are userid and timestamp_dttm for when a particular action took place.

Alphabetic List of Variables and Attributes				
#	Variable	Type	Len	Format
7	action_success_flg	Char	1	
4	action_type	Char	20	
1	audit_id	Num	8	
8	audit_info	Char	200	
6	executor_nm	Char	60	
13	newclient_id	Char	20	
16	newelapsed_time	Num	8	
21	newemail_recipients	Char	200	
20	newemail_sender	Char	200	
19	newexport_object	Char	200	
17	newexport_output	Char	200	
18	newexport_rows	Char	200	
11	newlasr_server_name	Char	200	
9	newlocation	Char	200	
14	newreport_elements	Char	200	
15	newserver_app	Char	60	
12	newtable_name	Char	200	
5	object_type	Char	100	
10	oldlocation	Char	200	
2	timestamp_dttm	Num	8	NLDATMS27.
3	userid	Char	60	

The values held under Action_type and Object_type are defined in lookup tables present in the SharedServices database. There are many types of objects available, however there are some key ones that are more relevant when identifying Visual Analytics reporting activity.

For Object_Type look for entries containing:

- Table. This will be used to represent a LASR table;
- VisualDataQuery. For Visual Analytics Queries;
- VisualExploration. For Visual Analytics Explorations;
- Report.BI. For Visual Analytics reports.

Caption 2 - Variables in Audit_VisualAnalytics table

For Action_type, it is more relevant to look at the available actions for each object type. The most common combinations are included in the following table:

Object_type	Action_Type	Comments
Report.BI	Copy	Visual Analytics reports have a lot of actions associated with them. Interesting values to analyze are Open, to see how many reports are being used, or Create, to see the rate at which the business might be adopting Visual Analytics.
	Create	
	Delete	
	Execute	
	Export	
	Move	
	Open	
	Print	
	Save	
	SendEmail	
Server.LASR	Cancel	To provide information on when was the last time a LASR server was started.
	Start	
Table	Add	The Read action is useful in order to see which LASR tables are being used.
	Delete	
	Read	
	Release	
	Update	
VisualDataQuery	Create	For Visual Analytics queries, it is useful to see how many queries are being developed in order to load the data. If there is complex data manipulation, then it might be beneficial to consider other methods for performing it.
	Delete	
	Execute	
	Open	
	Save	
VisualExploration	Create	Similarly to the Visual Analytics reports, it is useful to report on the rate that explorations are being created in order to analyze adoption of the tool, and the Open event to see
	Delete	
	Export	

Object_type	Action_Type	Comments
	Open	how many of those created are being used again.
	Print	
	Save	
	SendEmail	

Table 2 - List of useful Object types and Action types

Below is an extract of the Audit_VisualAnalytics data set with some of the information it records:

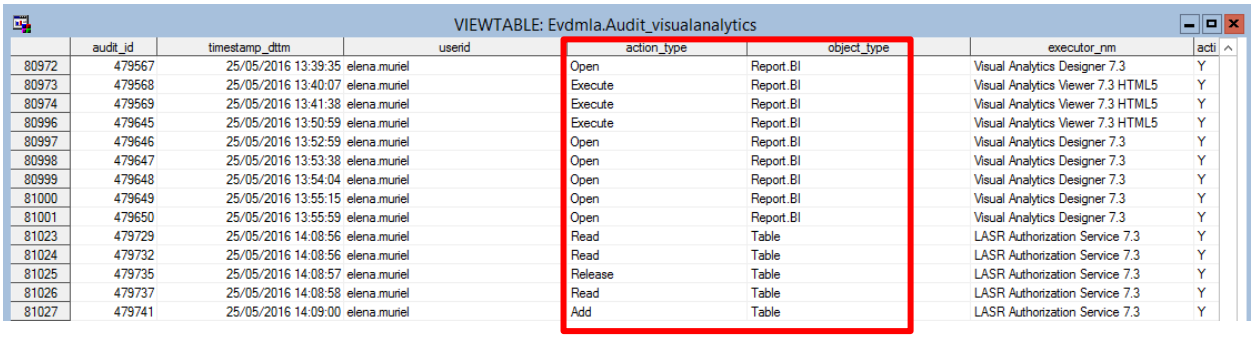

VIEWTABLE: Evdmla.Audit_visualanalytics

	audit_id	timestamp_dttm	userid	action_type	object_type	executor_nm	acti
80972	479567	25/05/2016 13:39:35	elena.muriel	Open	Report.BI	Visual Analytics Designer 7.3	Y
80973	479568	25/05/2016 13:40:07	elena.muriel	Execute	Report.BI	Visual Analytics Viewer 7.3 HTML5	Y
80974	479569	25/05/2016 13:41:38	elena.muriel	Execute	Report.BI	Visual Analytics Viewer 7.3 HTML5	Y
80996	479645	25/05/2016 13:50:59	elena.muriel	Execute	Report.BI	Visual Analytics Viewer 7.3 HTML5	Y
80997	479646	25/05/2016 13:52:59	elena.muriel	Open	Report.BI	Visual Analytics Designer 7.3	Y
80998	479647	25/05/2016 13:53:38	elena.muriel	Open	Report.BI	Visual Analytics Designer 7.3	Y
80999	479648	25/05/2016 13:54:04	elena.muriel	Open	Report.BI	Visual Analytics Designer 7.3	Y
81000	479649	25/05/2016 13:55:15	elena.muriel	Open	Report.BI	Visual Analytics Designer 7.3	Y
81001	479650	25/05/2016 13:55:59	elena.muriel	Open	Report.BI	Visual Analytics Designer 7.3	Y
81023	479729	25/05/2016 14:08:56	elena.muriel	Read	Table	LASR Authorization Service 7.3	Y
81024	479732	25/05/2016 14:08:56	elena.muriel	Read	Table	LASR Authorization Service 7.3	Y
81025	479735	25/05/2016 14:08:57	elena.muriel	Release	Table	LASR Authorization Service 7.3	Y
81026	479737	25/05/2016 14:08:58	elena.muriel	Read	Table	LASR Authorization Service 7.3	Y
81027	479741	25/05/2016 14:09:00	elena.muriel	Add	Table	LASR Authorization Service 7.3	Y

Caption 3 - Audit_VisualAnalytics table

REPORTS

Let's have a look at the reports that can be constructed using the audit data. By default, SAS provides the Administrator Overview report, which is a good starting point for monitoring Visual Analytics activity.

This report contains several sections to do basic reporting on the most used reports, explorations, users and tables, by simply providing frequency counts. It also gives you a bit more detail on existing data sources.

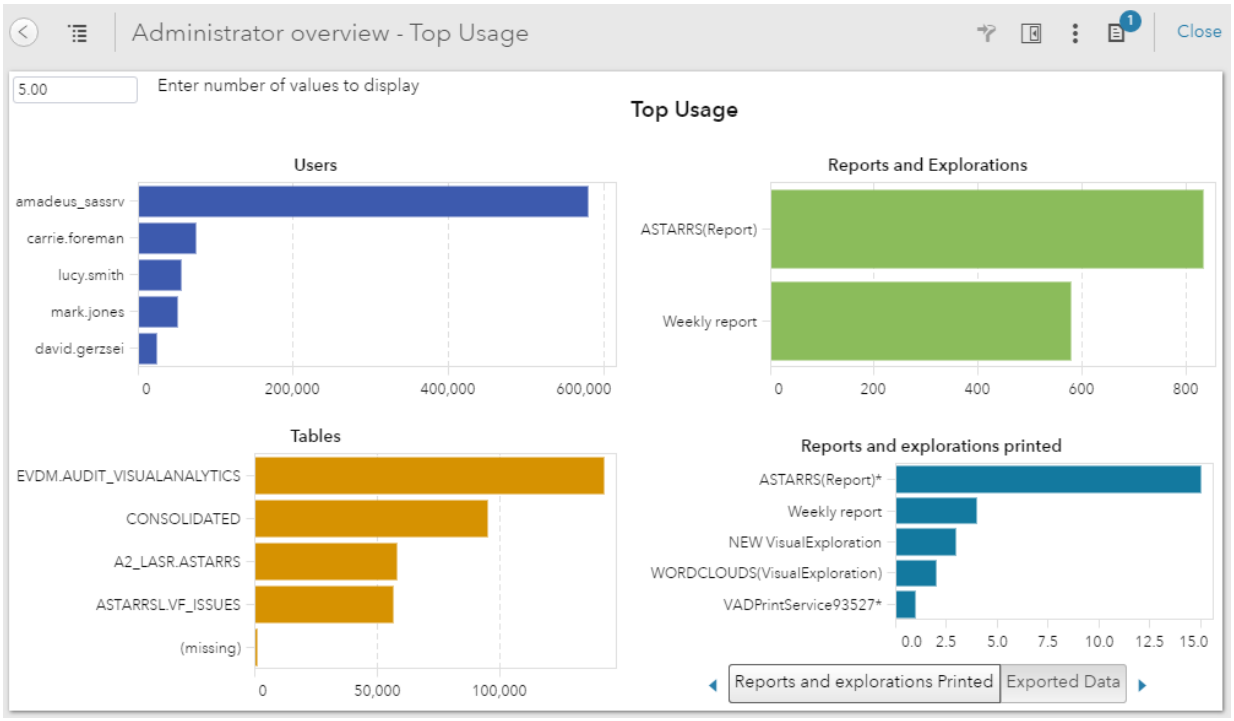

Caption 4 - Administrator Overview report

By using simple manipulation techniques, the information presented in this report can be expanded. Some of the questions we might want to answer are related to what reports are being used, by whom and when they were being accessed.

This can be achieved by requesting information on object_type of Report.BI and action_type of Open from the audit table.

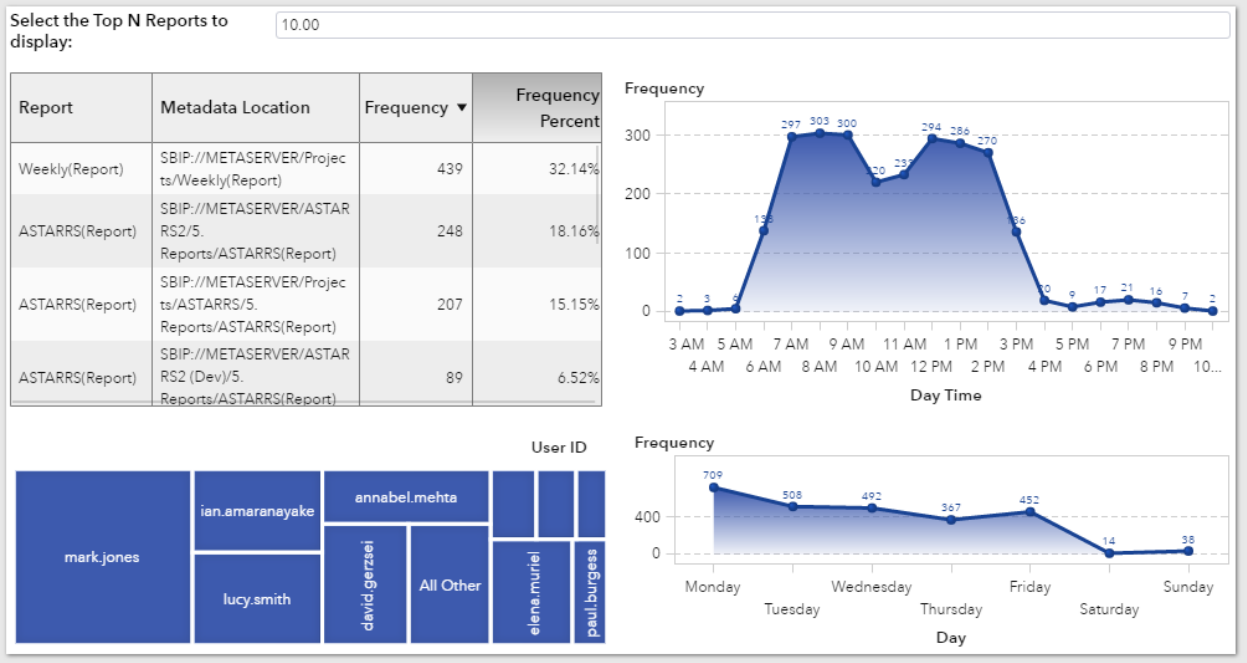

Caption 5 - Top N report

By adding a bit more information, it is easy to see the most utilized reports, where they are located and the number of times that they have been accessed. It can also be important to see the time of the day and day of the week that users are accessing these in order to flag particular access patterns. We can clearly see from the information displayed, that these reports are mainly accessed during office hours and there is a dip around lunch time when people might be having a break. The top 10 reports are accessed more frequently on a Monday, suggesting that this is the peak access time for reports.

Information related to the users consuming these reports is also presented. Using metadata queries it is possible to link users to metadata groups and potentially see which departments are more heavily utilizing the reports. It would also be possible to obtain email addresses from metadata enabling Administrators to contact report users to communicate any changes happening to the reports they use.

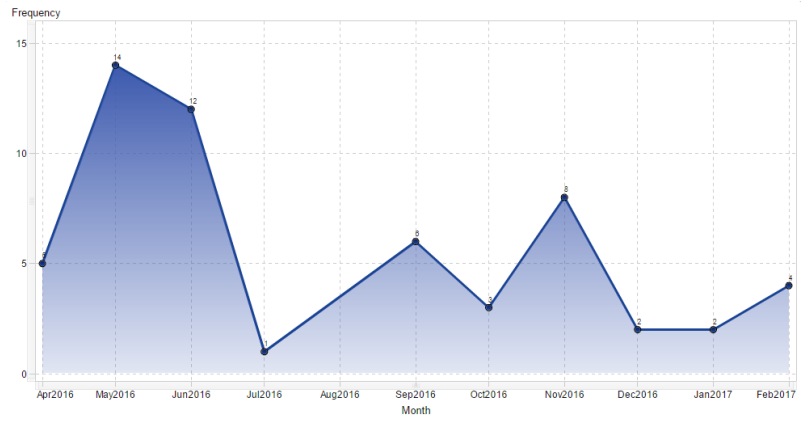

Another useful measure when introducing a new reporting tool is the rate at which the business is generating new content.

This can be done by analyzing the Create action type.

In this graph we can see an initial peak for creating reports when the tool was introduced and some quiet time around the UK summer holiday period.

Caption 6 - New reports created

For maintenance reasons, it is important to identify those reports that are being underutilized. For these reports Administrators could check the time that has passed since the last visit and determine if the data that was feeding that report is still needed.

It is also good to check with the business, so a useful step could be to identify the last person that accessed the report and check if it is still needed before decommissioning.

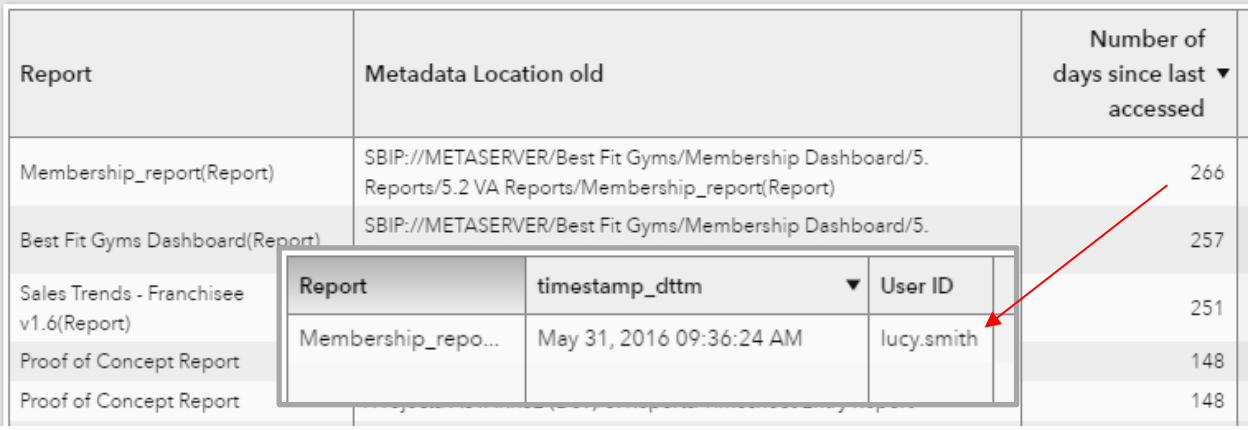

Report	Metadata Location old	Number of days since last accessed
Membership_report(Report)	SBIP://METASERVER/Best Fit Gyms/Membership Dashboard/5. Reports/5.2 VA Reports/Membership_report(Report)	266
Best Fit Gyms Dashboard(Report)	SBIP://METASERVER/Best Fit Gyms/Membership Dashboard/5.	257
Sales Trends - Franchisee v1.6(Report)		251
Proof of Concept Report		148
Proof of Concept Report		148

Report	timestamp_dttm	User ID
Membership_repo...	May 31, 2016 09:36:24 AM	lucy.smith

Caption 7 - Underutilized report

Another useful indicator is to explore the frequency at which the reports are being consumed, such as daily, weekly, etc., as this can determine refresh rates or even the necessity of having data sources available all the time. If reports are not frequently used, consider implementing a mechanism for loading the data for them on-demand.

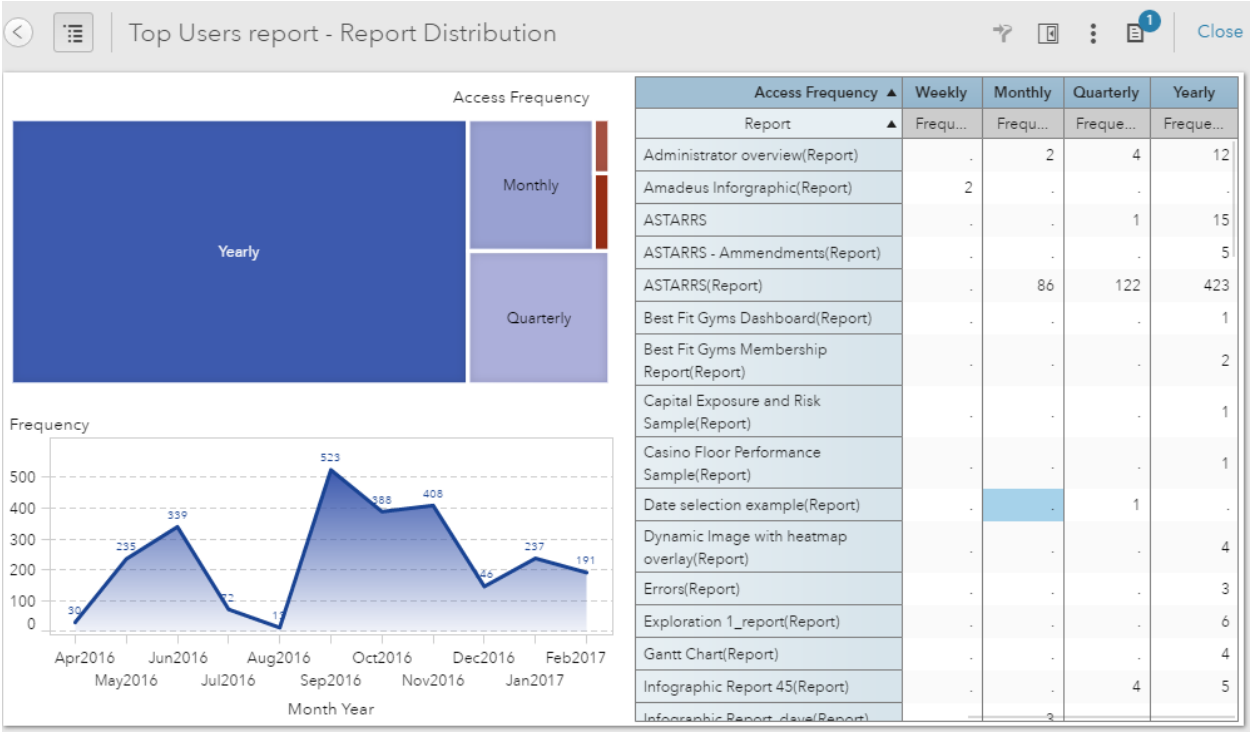

Access Frequency	Weekly	Monthly	Quarterly	Yearly
Report	Frequ...	Frequ...	Freque...	Freque...
Administrator overview(Report)	.	2	4	12
Amadeus Inforgraphic(Report)	2	.	.	.
ASTARRS	.	.	1	15
ASTARRS - Ammendments(Report)	.	.	.	5
ASTARRS(Report)	.	86	122	423
Best Fit Gyms Dashboard(Report)	.	.	.	1
Best Fit Gyms Membership Report(Report)	.	.	.	2
Capital Exposure and Risk Sample(Report)	.	.	.	1
Casino Floor Performance Sample(Report)	.	.	.	1
Date selection example(Report)	.	.	1	.
Dynamic Image with heatmap overlay(Report)	.	.	.	4
Errors(Report)	.	.	.	3
Exploration 1_report(Report)	.	.	.	6
Gantt Chart(Report)	.	.	.	4
Infographic Report 45(Report)	.	.	4	5

Caption 8 - Reports access rate

In the previous graph, we can see a similar pattern to the report creation, where there is a dip when it comes to access reports during the UK summer holidays.

Visual Explorations can be treated very similarly to Visual Analytics Reports. The object type to use in this case is VisualExploration.

When it comes to data loaded in memory, the main metric to obtain is its consumption through reports and explorations. This can be achieved by filtering the audit table by object_type of Table and action_type of Read. It is also useful to include a filter to remove any system tables such as the AUDIT_VISUALANALYTICS from the analysis, as it might skew some of the results.

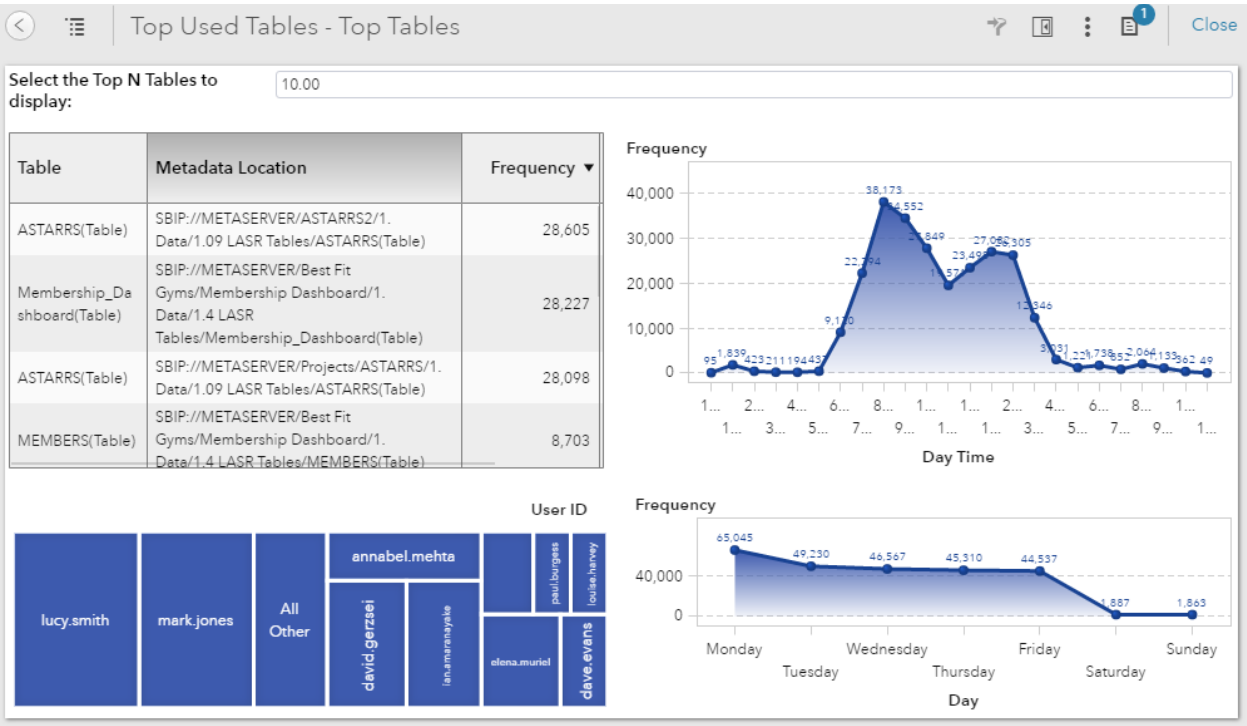

Caption 9 - Top Tables

We can see a similar access pattern when reading LASR tables as for accessing reports. Any underutilized table will be a candidate to be unloaded. But in order to see what reports/explorations it affects, we will need to use the Relationship data, which is the second type of auditing data available and described in the next section.

The age of the data loaded is also a useful indicator. This data is not recorded on the audit tables but Proc Contents provides one technique to extract that information for you:

```
libname lasrlib SASIOLA TAG=your-tag PORT=your-port-number
SIGNER="http://hostname:port/SASLASRAuthorisation HOST="hostname";

proc contents data=lasrlib._all_ out=lasrtables noprint;
run;
```

The table created will contain (among other information) the date and time each table was loaded into memory. By registering this table in metadata we can analyze the age of the data loaded in Visual Analytics.

DATA ARCHIVE

The MiddleTier audit service data can be very useful when it comes to support the management of your Visual Analytics platform, but it also generates a large amount of data that needs careful consideration. The SharedServices database has a default archive process that is important to manage, as failure to do so can impact performance and stability of the web applications over time.

The data archive process for the audit service data has three steps:

Figure 2 - Archive steps

1. By default, the SharedServices database will archive the audit tables every Monday at the beginning of the day. This process moves records from sas_audit into sas_audit_archive and from sas_audit_entry into sas_audit_entry_archive, to only keep data that is 30 days old in the main tables.

   The scheduling time and frequency of this process is configurable, and it is advisable to change this setting so that the system performs archiving every day.

	audit_archive [PK] serial	object_type_ bigint	action_type_ bigint	frequency_nc bigint
1	1	-1	8	1296000000
2	2	-1	9	1296000000
3	3	31	45	1296000000
4	4	32	1	1296000000
5	5	32	2	1296000000
6	6	32	36	1296000000
7	7	32	45	1296000000
8	8	32	48	1296000000

Caption 10 - Archive frequency

   The amount of data to keep can also be changed by modifying the sas_audit_archive_rule table, where the frequency_no variable holds the number of milliseconds before archiving the combinations of object types.

   The default of 30 days corresponds to 2,592,000,000 milliseconds. In this caption this has been modified to 1,296,000,000 milliseconds or 15 days.

2. The second data archive process is to remove data from the archive tables themselves. This can be done as a one off process by using the Postgres administrator tool or by the scripts available provided by SAS Institute. For details, refer to the Usage Support Note 58599: Managing SAS Audit Tables [Ref. 2].
3. The third step is to remove data from the audit_visualanalytics.sas7dat table. A SAS program can be created to delete records from the SAS data set once they are no longer needed. The LASR table will need to be refreshed in order for the changes to take effect, as by default, data only gets appended.

More information on how to configure archiving of Visual Analytics data can also be found under the Middle-Tier Administration Guide [Ref. 1].

RELATIONSHIP

This is the second area available when collecting auditing data for Visual Analytics. The relationship information focuses on the dependencies between Visual Analytics objects, such as which LASR tables are required for each report.

Once enabled, the load process scans the SAS Folders tree in metadata for items that were created or modified since the last scheduled load operation. The process retrieves these objects relationship information and stores it in the WIP database. The underlying technology used to capture relationship data is the Relationship Reporter, which is open to not just Visual Analytics content but any content stored in SAS Metadata. For more information on the Relationship Reporter refer to the SAS System Administration Guide [Ref 3].

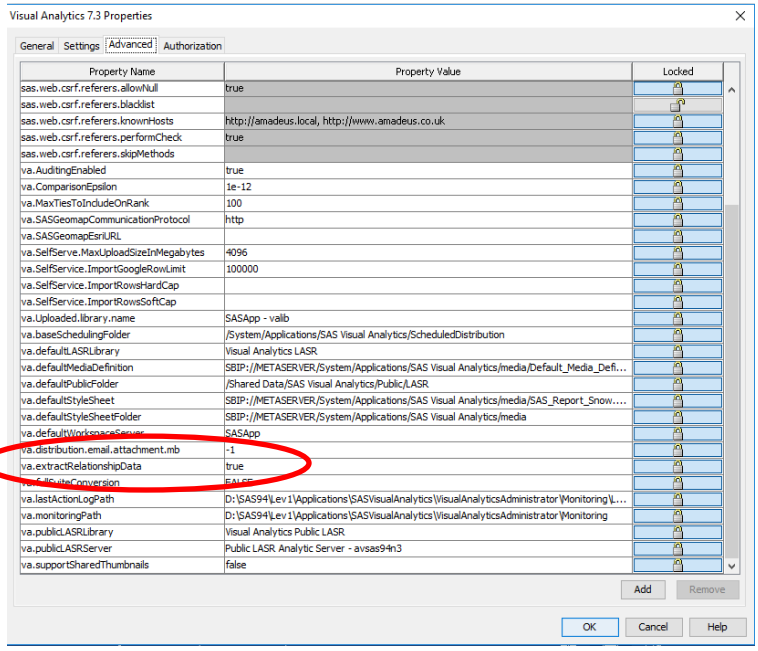

Automatic loading of the relationship data is configured through SAS Management Console.

In order to enable it, navigate to the Visual Analytics properties under:

→Application Management

→Configuration Manager

→SAS Application Infrastructure

→Visual Analytics 7.3

Set va.extractRelationshipData to true and then restart the Visual Analytics Web Application Server.

Caption 11 - Enabling Relationship reporting

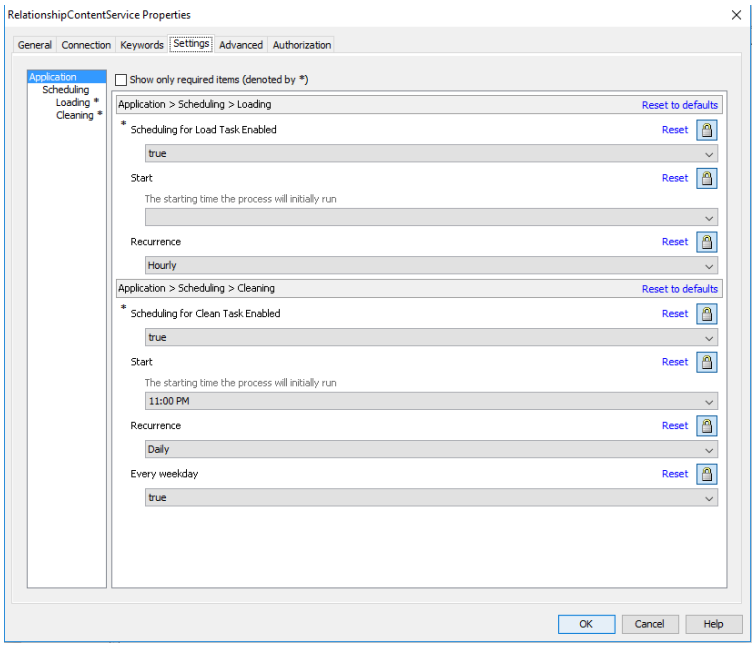

To modify any of its properties, navigate to:

→Application Management

→Configuration Manager

→SAS Application Infrastructure

→Web Infra Platform Services 9.4

→RelationshipContentService

Where scheduling times and frequency can be configured.

By default, the process will be executed every hour daily.

Caption 12 - Properties available for Relationship reports

The data is extracted in a similar way to the Audit Service. From the WIP database it gets extracted to the Append directory in EVDMLA, from where it gets loaded into the LASR Server by the autoload process. The physical SAS data set is stored under:

\LevX\AppData\SASVisualAnalytics\VisualAnalyticsAdministrator\AutoLoad\EVDMLA\relationships_visualanalytics.sas7bdat.

The following table lists the variables that are contained within the Relationship table.

The Relationship_Type can either be “Contains” or “Is Dependent on” and the Direction is always “TO”.

For Visual Analytics governance, the more interesting types of objects present (variable LeftType) are:

- DeployedFlow
- DeployedJob
- Report.BI
- Table
- VisualAnalyticsQuery
- VisualExploration

Where RightType contains:

- Column
- DeployedFlow
- DeployedJob
- Job
- Library
- StoredProcess
- Table
- VisualDataQuery

Alphabetic List of Variables and Attributes					
#	Variable	Type	Len	Format	Informat
2	Depth	Num	8	BEST12.	BEST32.
8	Direction	Char	2	$2.	$2.
3	LeftName	Char	56	$56.	$56.
4	LeftPath	Char	130	$130.	$130.
5	LeftType	Char	17	$17.	$17.
6	LeftTypeLabel	Char	18	$18.	$18.
7	RelationshipType	Char	15	$15.	$15.
9	RightName	Char	54	$54.	$54.
10	RightPath	Char	102	$102.	$102.
11	RightType	Char	15	$15.	$15.
12	RightTypeLabel	Char	17	$17.	$17.
1	SubjectIndex	Num	8	BEST12.	BEST32.

Caption 13 - Relationships_VisualAnalytics

The combination of left and right type, alongside the type of relationship, determines the nature of the relationship between the objects.

For Visual Analytics, the more important combinations are listed in the following table:

LeftTypeLabel	RightTypeLabel
Report.BI	Table
	Stored Process
Table	Column
	Library
Visual Data Query	Table
	Visual Data Query
	Deployed Flow
Visual Exploration	Table
Deployed Flow	Deployed Job
Deployed Job	Job

Table 3 - Types of objects available in relationship interactions

By interrogating the left type and right types of objects, an administrator can gain knowledge of the dependencies that exist between some of the existing objects.

There are two types of analysis that can be performed from the relationship data:

- Impact analysis. Objects that depend on a subject. For example, by using data from an impact analysis on a Table, we can gain information on the reports, explorations and Visual Analytics queries that might depend on it.
- Lineage analysis. Objects that each subject depends on. For example, by using data from a Lineage analysis on a Table, we can gain information on the Libraries and Columns that the table depends on.

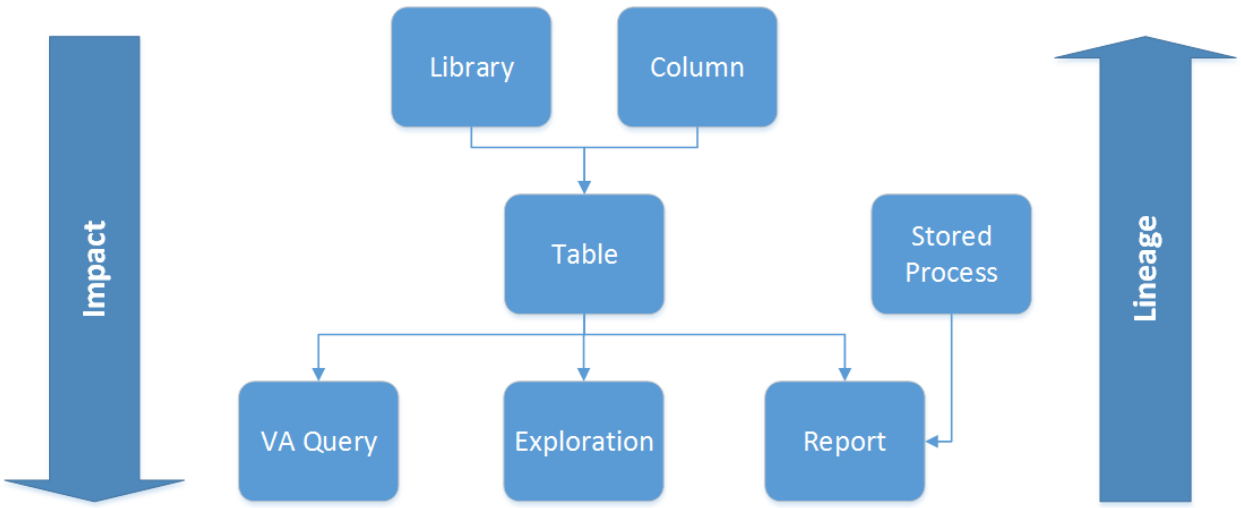

Figure 3 - Impact and Lineage

For Visual Analytics, only direct dependencies are reported but not indirect ones. This means that when looking at a report, the lineage analysis will give us the Tables and Stored Processes that feed into the report, but not the library and column names that feed into the table (indirect dependencies).

REPORTS

The SAS default Relationship Report contains a list of objects and the relations it has with others. When a report is selected, it gives information about what tables contribute to that report.

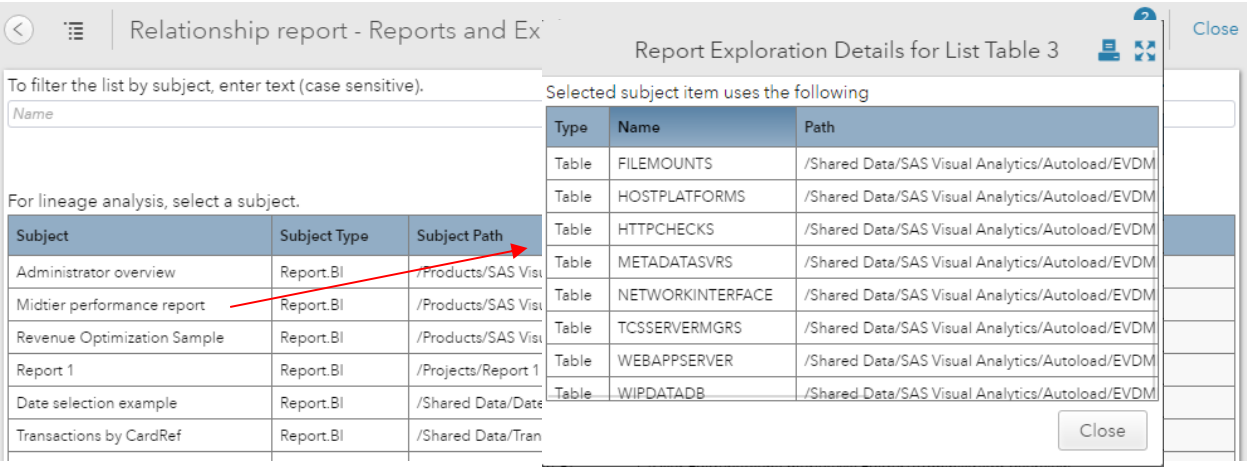

Caption 14 - Relationship report

But there might be a better way of presenting the relationship information by dividing it into the different types of objects.

For example, it might be interesting to see what objects are affected by a particular LASR table, so if I make any changes to the table I know what the impact will be.

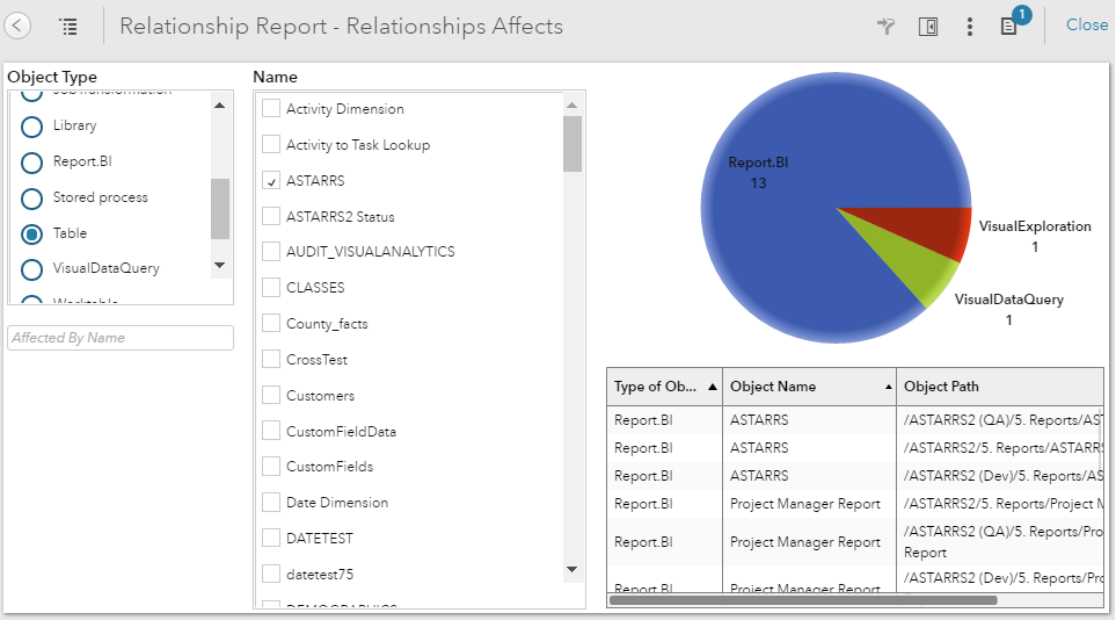

Caption 15 - Relationships Impact

A change to a specific table affects existing reports and explorations, but a change on a report does not affect any other object types, as reports are situated at the end of the dependency tree.

This Impact report is constructed by first listing all of the objects that are affected by something else (RightType), and then displaying the elements that it affects (LeftName, LeftPath and LeftType).

Looking at a Lineage report, this can be constructed by selecting the objects that are at the end of the dependency tree (LeftType) in order to see what other components affect it (RightName, RightPath and RightType).

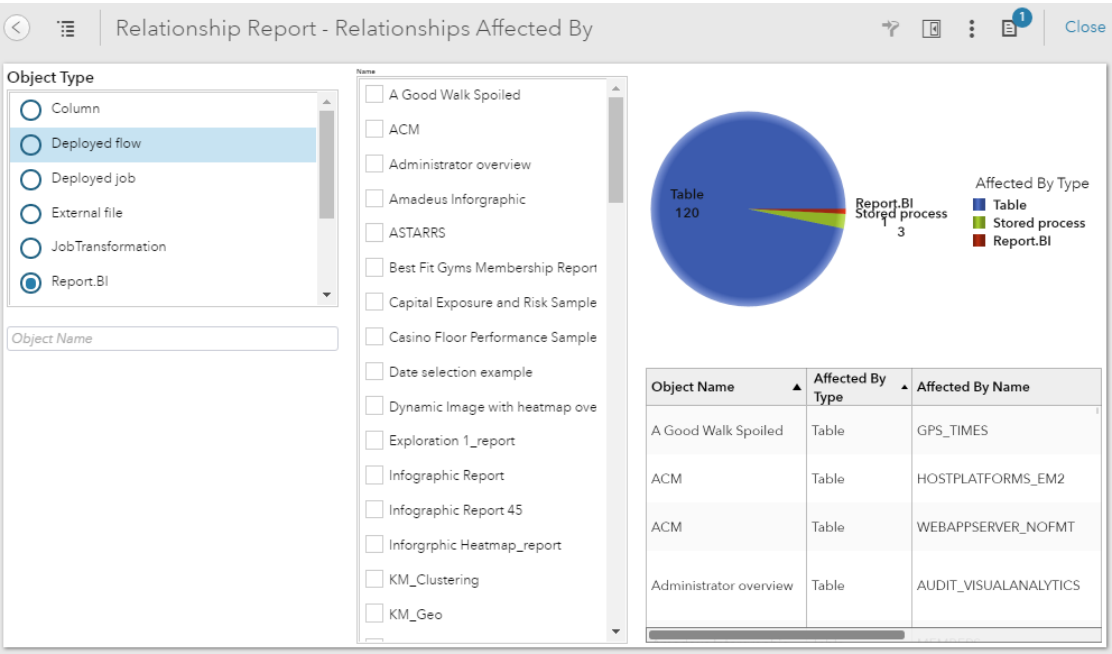

Caption 16 - Relationship Lineage

DATA ARCHIVE

The volumes of data produced by the relationship tables is low compared to the audit service. There are normally no concerns when dealing with this data.

If there is the need to purge old records, simply remove them from the physical SAS data set relationship_visualanalytics.sas7bdat.

ACM

The third reporting area available is related to the Middle Tier performance metrics. These metrics are captured through Environment Manager Extended Monitoring and as the previous cases, it requires extra configuration in order to enable it.

Environment Manager Extended Monitoring is not just used for Visual Analytics, it is an extensive set of metrics that can capture information relating to the actions performed against the SAS servers (from authentication to procedures used in code) and performance metrics captured by reading the log files generated by the SAS platform and Operating System information.

Only the ACM (Agent Collected Metrics) for non-distributed SAS LASR Analytical servers are currently integrated with SAS Visual Analytics and APM (Audit, Performance and Measurements) is not available. Once these metrics are being captured, the data is stored on a series of SAS datasets under SASEnvironmentManager\emiframework\Datamart\acm.

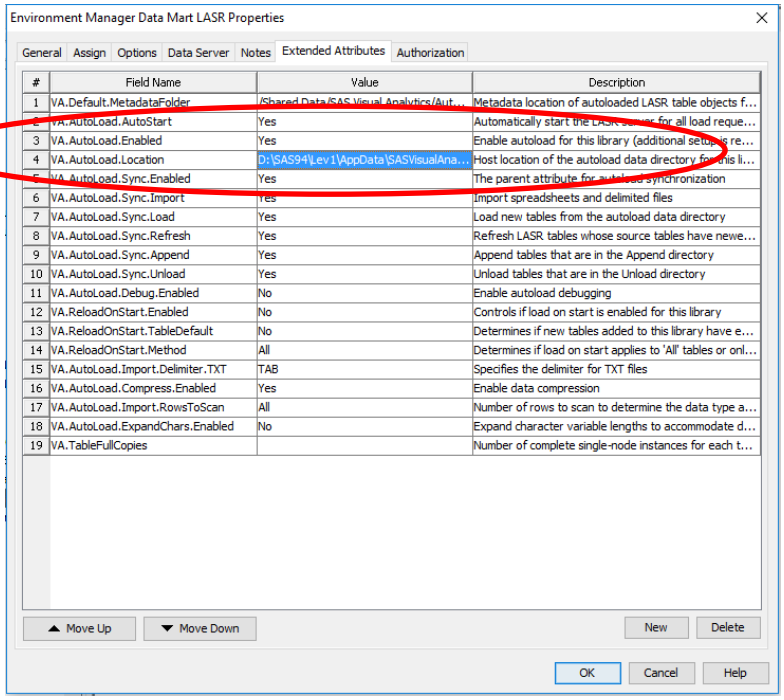

In order to enable to autoload process into LASR, the Environment Manager Data Mart LASR library needs to be configured by enabling the autoload and providing the location of the autoload directory.

This can be done by accessing the Extended Attributes tab on the Library in SAS Management Console.

Figure 4 - Enabling ACM reporting

The extraction process also needs to be enabled, and this has to be done by using some scripts:

emi_init.bat –vafeed for Windows or emi_init.sh –vafeed for Linux

More information can be found in SAS Environment Manager 2.5: User's Guide [Ref. 5].

Enabling the process, copies the data from the Environment Manager tables into the Append folder in EVDMLA in order to be picked up by the Autoload process. As some of these tables have formats applied to their column definitions, it is important to copy the formats catalog from the ACM location to the EVDMLA one.

Not all the available ACM tables will be copied across though, and only the following ones will be included for analysis in Visual Analytics:

Table	Description
FILEMOUNTS	Disks reads and writes data for the file mounts in your environment
HOSTPLATFORMS	Metrics for the platform-level resources such as memory, CPUs, reads/writes per server
HTTPCHECKS	Inbound/outbound and response time for connections to applications
SASAPPSERVERS	Metric data for the SAS Application servers
SASLOGICALSERVERS	Metric data for the SAS logical servers
METADATASVRS	Metric data for the SAS metadata servers
NETWORKINTERFACE	Metric data for the network interfaces
TCSERVERMGRS	Metric data for individual SAS web applications
WEBAPPSERVER	Metric data for SAS Web Application server instances
WIPDATADB	Metric data for the databases within the Web Infrastructure Platform Data Server
EMI_INFO	Metrics on the performance of SAS environment Manager ETL processes

Table 4 – ACM tables available in Visual Analytics

These tables help monitor the performance and overall functioning of the hardware, network and infrastructure on the system.

REPORTS

The reports that are usually created using these metrics are to be consumed by Administrators wishing to look at the performance of the infrastructure and looking to raise alerts to IT departments. These reports can supplement existing Operating System reports.

The existing Midtier performance report in Visual Analytics presents some of the ACM metrics, such as CPU usage, memory usage and connections.

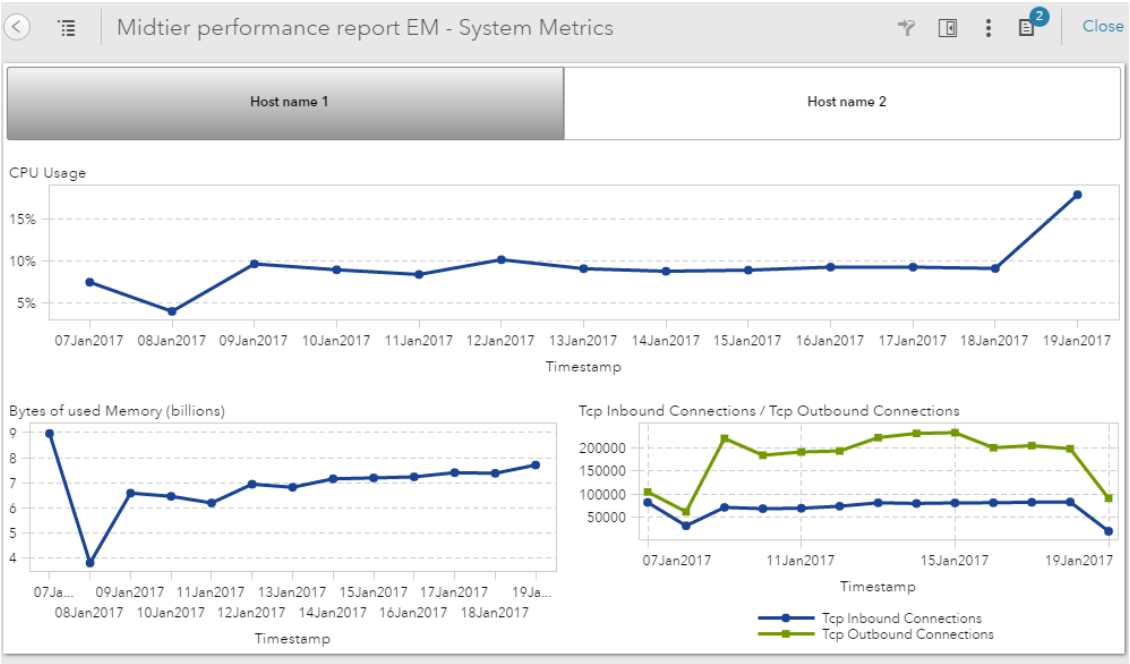

Caption 17 - Midtier performance report System Metrics

Patterns and peaks can be detected and more information can be extracted through the use of date and time dimensions to identify the exact time that an event occurred.

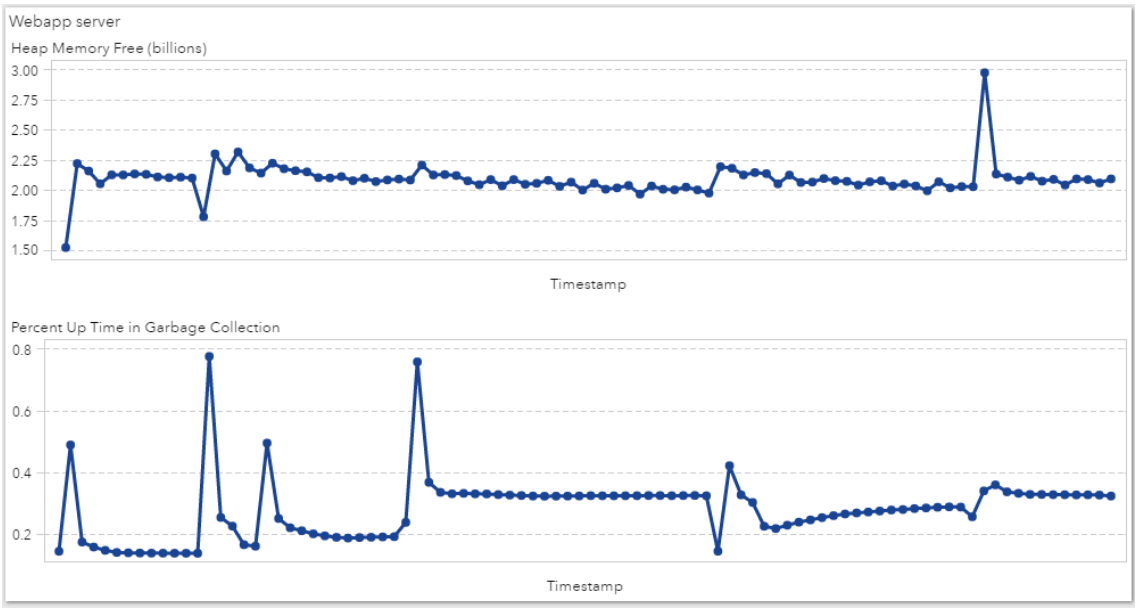

Caption 18 - Midtier performance report Web App server

Once ACM is configured, Environment Manager creates a series of reports using stored processes. These reports contain most of the key metrics for this area and are accessible through the Environment Manager Report Centre.

These reports are a good basis to explore the type of information available.

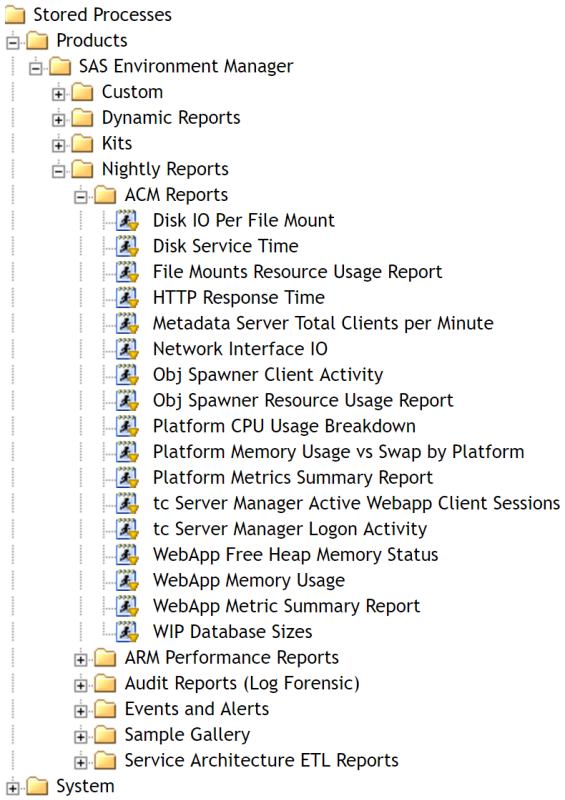

Caption 19 - Environment Manager ACM reports

DATA ARCHIVE

ACM data can get big very quickly, therefore archive routines are very important. Volumes of data kept can be controlled through the EnvMgr Enablement Kit 2.2 as displayed in the next caption.

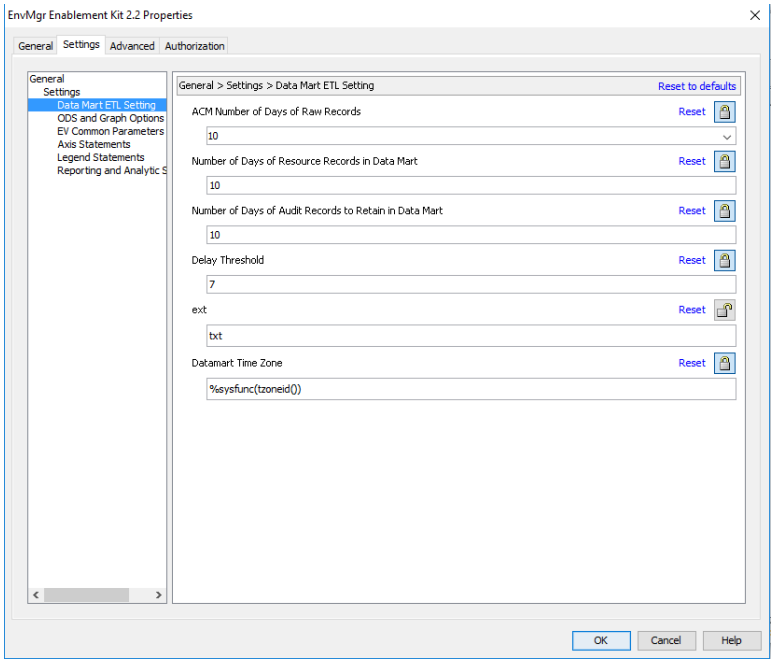

Consider the need for keeping all of the raw data. If the detail is no longer needed, summary tables could be created to save key historical measures.

Caption 20 - EnvMgr Enablement Kit

CONCLUSION

The three auditing mechanisms available for SAS Visual Analytics are an important tool for any SAS Administrator when it comes to manage a Visual Analytics configuration. By using the data available through the Middle Tier Audit Service, Relationship Reporter and ACM, administrators can gain invaluable information regarding the use of the content being developed and how to best maintain it.

By expanding the supplied reports, it is possible to develop an active monitoring system on the utilization of the platform, helping organizations gain the most from their existing software and infrastructure. The paper has shown how to trace utilization/underutilization/dependencies for LASR tables and Visual Analytics reports.

Data archiving is an important factor to consider, as large volumes of data can affect the performance of the platform. Time needs to be spent when defining these processes and identifying the key metrics that are needed.

By understanding the data and the processes than run in the background, SAS administrators can define new reports that help with their maintenance tasks and assess impact on changes to be introduced.

REFERENCES

[1] SAS 9.4 Intelligence Platform: Middle-Tier Administration Guide, Fourth Edition https://support.sas.com/documentation/cdl/en/bimtag/69826/HTML/default/viewer.htm#titlepage.htm

[2] Usage Note 58599: Managing SAS Audit Tables http://support.sas.com/kb/58/599.html

[3] SAS 9.4 Intelligence Platform: System Administration Guide, Fourth Edition, Using the Batch Relationship Reporting Tool

http://support.sas.com/documentation/cdl/en/bisag/68240/HTML/default/viewer.htm#p1dcleibglhr6wn19yzvfz81yw5y.htm

[4] SAS Big Brother, Elena Muriel http://support.sas.com/resources/papers/proceedings16/9680-2016.pdf

[5] SAS Environment Manager 2.5: User's Guide, Second Edition
http://support.sas.com/documentation/cdl/en/evug/69029/HTML/default/viewer.htm#n1gacla7ui0covn1dv65l385uhch.htm

[6] What is going on with my Visual Analytics audit data collection? Gerry Nelson
http://blogs.sas.com/content/sgf/2015/12/29/what-is-going-on-with-my-visual-analytics-audit-data-collection/

CONTACT INFORMATION

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

Elena Muriel
Amadeus Software Ltd.
+44 1993 848010
elena.muriel@amadeus.co.uk
www.amadeus.co.uk