

Paper 4662-2020

From Silos to Ecosystem - Evolution of SAS Analytical Platforms at Elisa Corporation

Prasanna Pandian, Elisa Corporation; Jarno Lindqvist, SAS Institute Oy

ABSTRACT

SAS in Elisa was originally a standalone product where metadata, data and users were siloed in three separate environments for SAS Enterprise Guide, SAS Enterprise Miner and SAS Visual Analytics. Data tasks were manually executed, and the data produced by SAS was not available to other environments. Open source team relying on R and Python were not able to utilize the SAS datasets directly. In late 2018, as part of the enterprise BI architectural strategy, Elisa ventured to build a new SAS Analytics & Data Ecosystem - a collection of infrastructure, analytics and applications used to capture and analyze data. By deploying modern SAS Viya and SAS 9.4 platforms Elisa implemented standardized approaches, automation, tools, and processes.

Current architecture enables them to effortlessly share data artifacts and collaborate on projects across corporate functions and platforms. The produced data is now shareable among **all SAS users while the heavier analytical processing is pushed to SAS Viya. Elisa's open source team has access to SAS datasets in Jupyter Notebook and R Studio via SAS Viya integration.** All SAS-produced data is exposed via REST APIs and corporate reporting users utilize them for continued analysis in various tools. Modern Elisa nurtures a sustainable common analytics ecosystem and the goal is now set on realizing the promise of analytics and making the most **of Elisa's data assets with right set of architecture and tools.**

INTRODUCTION

Elisa is a Finnish pioneer in telecommunications and digital services since 1882. They serve approximately 2.8 million consumer, corporate and public administration organization customers, and have over 6.2 million subscriptions in their network operating across Finland, Estonia and internationally. **Elisa's** core values in all their business are Customer orientation, Responsibility, Renewal, Results orientation and Collaboration. **In 2018 Elisa's revenue was EUR 1.83 billion euros, and they employed 4,800 people.** Elisa has been a SAS customer since 1997 and is currently using a wide range of SAS solutions serving the needs of several departments.

Prasanna Pandian works for Elisa as BI Architect and SAS platform owner being responsible for selecting and administering the platforms for managing data, analytics and business intelligence. Jarno Lindqvist works for SAS Finland since 2005 and has been **Elisa's** Technical Account Manager since 2019 working closely with Prasanna to ensure Elisa gets value from their investment in SAS.

The idea for writing this paper came from our **experiences of working together on Elisa's SAS 9.4 and Viya platform.** Since the SAS Viya platform was new in the house, there was much need for proving the value and solving some technical obstacles. It started in 2018 when Elisa opted for SAS Viya reference customer program with SAS Viya. In this paper we tell a story of how things have evolved to where they are and what steps have been taken and decisions made along the way to build this new SAS ecosystem.

PARTNERSHIP OVER THE YEARS

Elisa first opted SAS in 1997 and the first use case was credit scoring - trying to predict which customers are likely to end up as defaulters paying their bills. After that, next use case was churn prediction and it was carried out in both Finland and Estonia. These were successful use cases and Elisa was able to adopt success already during those days with SAS analytics.

Base SAS and Enterprise Miner were the technologies used during the initial days and soon afterwards the wider SAS adoption in analytics started. In the next phases SAS Visual Analytics was added to the stack to visualize the SAS data created from the analytical use cases. Later on, Elisa opted for SAS Cost and Profitability Management solution in their finance department to allocate costs more precisely.

As Elisa grew through acquisitions into the IT service business there was need to improve process efficiency in customer support. This led to adopting SAS Enterprise Content Categorization to help automate the support ticket process.

To bring the story to current times, when SAS Viya was introduced Elisa became one of the first organizations to bring in this latest SAS generation in the form of SAS Visual Analytics and Visual Statistics.

Elisa is a heavy user of reporting tools and recently the need arose to share SAS datasets with users on other reporting platforms. This need was met with SAS Federation Server that can among other capabilities be used to provide a standard ODBC API on SAS datasets.

CURRENT STATUS

Currently Elisa utilizes SAS in several business areas, such as B2B and B2C marketing, sales **operations, customer support and corporate finance**. **Elisa's current SAS software stack** includes:

SAS 9.4

- SAS Office Analytics
 - Enterprise Guide
 - Add-In for Microsoft Office
- SAS Enterprise Miner
- SAS Federation Server
- SAS Cost and Profitability Management
- SAS Enterprise Content Categorization

SAS Viya 3.5

- Visual Analytics and Visual Statistics

The image below represents the current Data and Analytics architecture stack at Elisa.

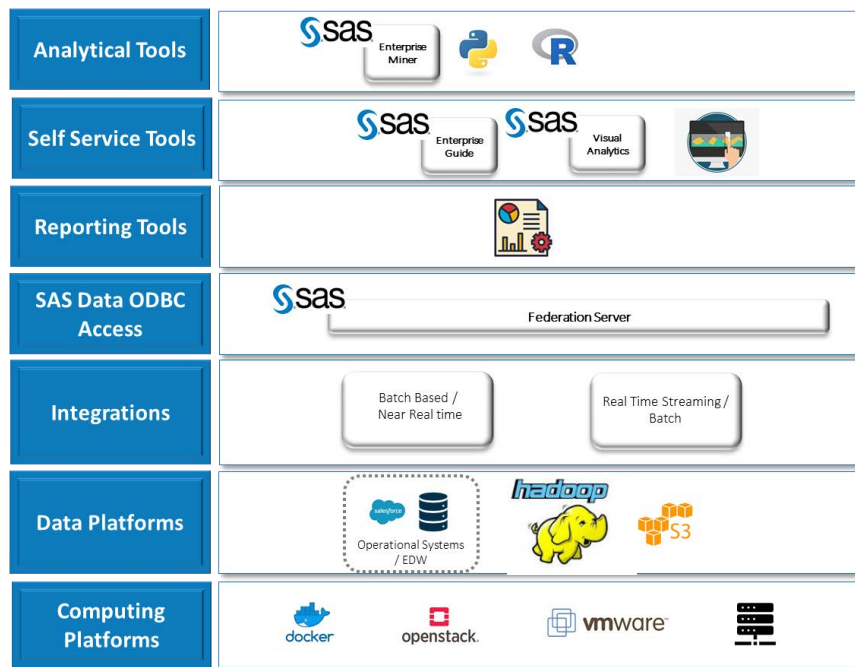


Figure 1 Data and Analytics architecture stack at Elisa

Elisa utilizes many different databases and storage platforms for managing data. Open S3 type of storage serves as an on-premise data for data science purposes. A data warehouse has been used for storing structured data for a long time to support business reporting. There are also conventional transactional databases used as operational datastores. These data sources are available to SAS and other standard reporting tools for traditional reporting and self-service data analysis and **Elisa's Hadoop** data lake data for advanced AI & machine learning use cases. Image below explains the current SAS reference architecture at Elisa.

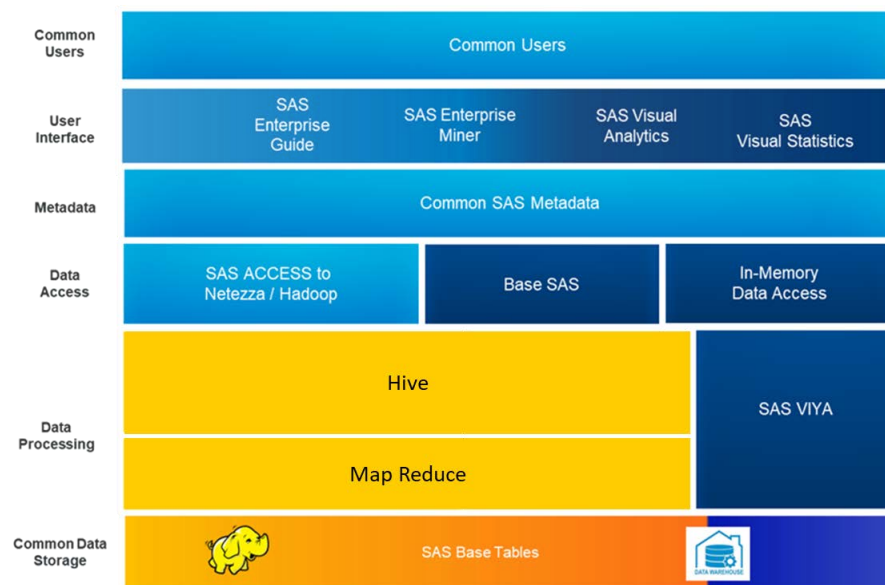


Figure 2 SAS Reference Architecture at Elisa

In addition to SAS software, Elisa utilizes many of the commonly available commercial technologies for data management and reporting, such as Informatica and Microsoft Power BI. The analytics team works with R, Python and Jupyter Notebook for analytics and modelling. The challenge for Prasanna is optimizing this mix and variety of tools and ensuring they meet the needs of business while being cost-effective and productive.

DATA FEDERATION PROVIDES 3RD PARTY ACCESS

The need to make these platforms more integrated and share data paved the way for SAS Federation Server. As Microsoft Power BI is used as common self-service tool at Elisa, they often have the need to access data that was prepared with SAS tools. In earlier days, the only way to make this happen was to export the data to a general format such as CSV and then read the data into Power BI. With SAS Federation Server the constructed SAS data sets can be shared as views via ODBC connection. This makes the process faster and reduces unnecessary work.

The figure below depicts this process, with data sources on the left, which are defined on the Federation Server. Any ODBC compliant application can connect to those data sources provided the user has the necessary credentials and permissions that are controlled on the SAS metadata server of SAS 9.4 platform.

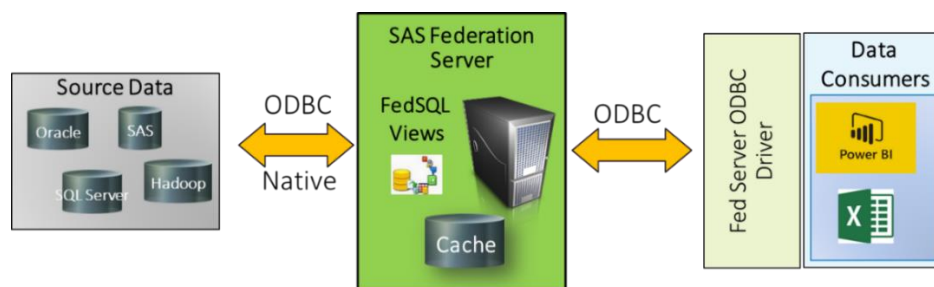


Figure 3 Federation Server data flow at Elisa

OPEN SOURCE INTEGRATION WITH PYTHON AND R

Elisa's data analytics team is actively using Python and R for data processing and modelling. There is often need to use the data sets constructed by SAS users with Enterprise Guide. The previous way to fulfill this need was to produce a CSV file from SAS and then import that to Python or R or write SAS data sets to a database where Python or R could then get connected to.

With SAS Viya there are many more straightforward ways to accomplish this. One simple way to incorporate Python and R and SAS is using SAS SWAT (*Scripting Wrapper for Analytics Transfer*). It enables **Elisa's** data scientists to load and analyze data sets of any size. The value comes from having access to common data and being able to analyze extremely large data sets using the full processing power of CAS, while still retaining the user experience of Python or R on the client side.

The diagram below shows the overall process for a CAS action request from invoking the action to getting the results back. Conversions take place in the client-side to convert Python or R objects to CAS constructs and likewise for returned results.

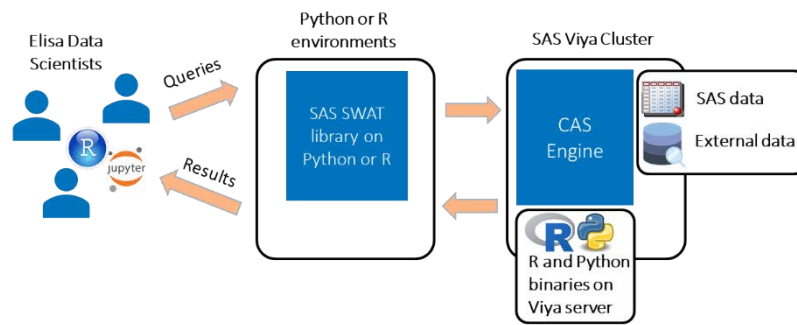


Figure 4 Integration and process flow using Python and R with CAS

BETTER CUSTOMER SERVICE THROUGH CATEGORIZATION

While being a telecommunications provider Elisa offers also IT services to mainly organizational customers. Customer contact Elisa via email and as emails are human created non-structured data, they need to be categorized in order to reach the correct service teams depending on the matter of problem or a type of contract. Previously this resulted in a lot of manual work and Elisa wanted to improve process efficiency through automation. They chose SAS Enterprise Content Categorization to help them automatically categorize the incoming support tickets. To improve their customer satisfaction, Elisa also provides a customer portal on SAS Visual Analytics which their organizational IT service customers can access to see the status of their support tickets.

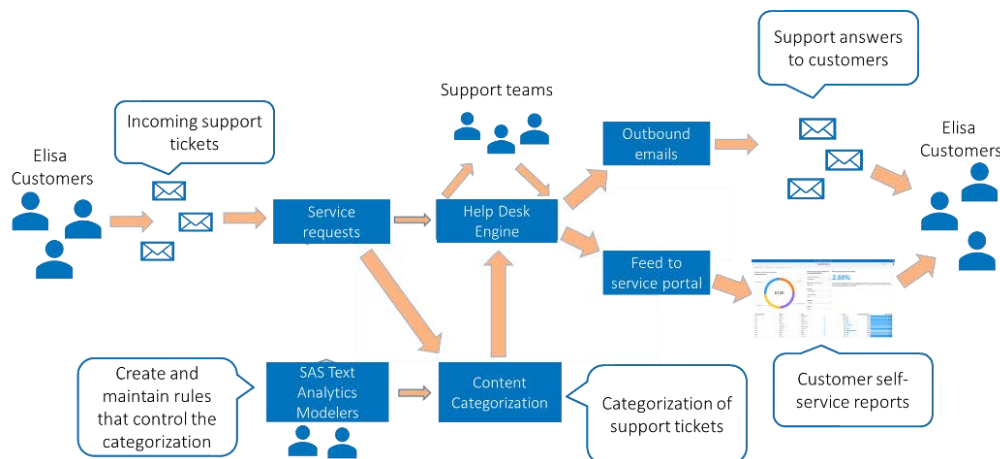


Figure 5 Automatic categorization of support tickets

IMPROVING GOVERNANCE WITH METADATA REPORTING

Every ecosystem built should also be governed properly in order to create a consistent and **proper handling of an organization's data throughout the entire enterprise**. For this purpose, a metadata management tool (see screenshot below) was built in-house at Elisa in order to understand the metadata about SAS Enterprise Guide and SAS Enterprise Miner libraries and licensing information. Elisa also needed to understand user access privileges – which user belongs to which access groups and collect audit information – who has access to what data.

Metadata Libraries Report

This module is built using the SAS metadata and DATA STEP functions to query the metadata repository and returns a list of all libraries and their associated directory or database schema.

License and Product Info

The module identifies which SAS components or products are installed and what the license expiration date is for the system and each component. It also gives the site name and number as well as the system birthday, operating system and any grace or warning period for license expiration.

Logins and Associated Identities

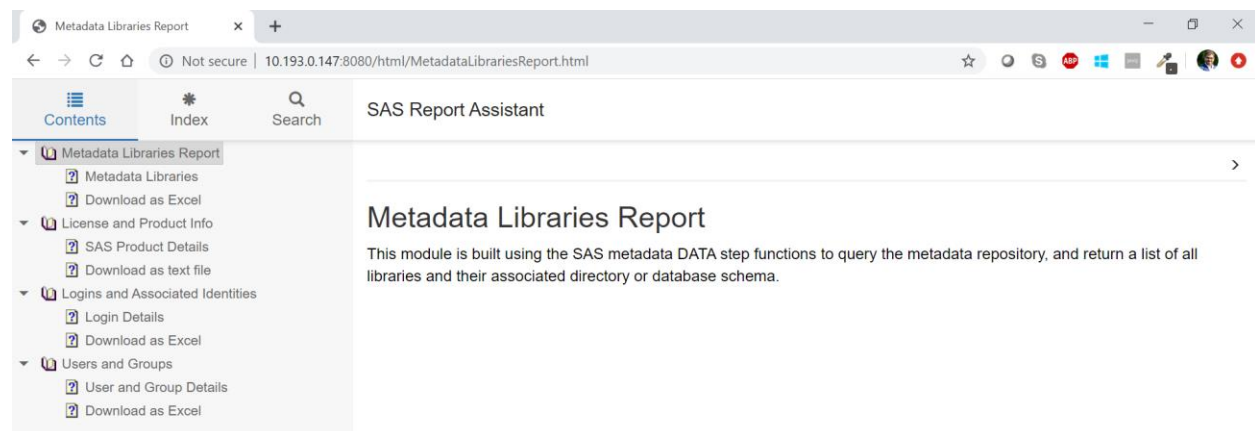
This module is built using the SAS metadata DATA step functions to query the metadata repository and return a list of all logins and the users or groups to which they belong. It returns the authentication domains in which the logins are active.

Users and Groups

This module is built using the SAS metadata and DATA STEP functions to query the metadata repository and return a list of all users and the user groups to which they belong.

Auditing data access

Organizations want to record the user, the table, and the date and time that the table was accessed. This feature to record the audit was built using SAS Logging Facility. This facility is configured to collect, categorize, and filter events and write them to a database table from which then the audit data is then tracked.



The screenshot shows a web browser window displaying the SAS Report Assistant interface. The browser's address bar shows the URL `10.193.0.147:8080/html/MetadataLibrariesReport.html`. The page title is "Metadata Libraries Report". The interface includes a navigation menu on the left with the following items:

- Metadata Libraries Report
 - Metadata Libraries
 - Download as Excel
- License and Product Info
 - SAS Product Details
 - Download as text file
- Logins and Associated Identities
 - Login Details
 - Download as Excel
- Users and Groups
 - User and Group Details
 - Download as Excel

The main content area displays the "Metadata Libraries Report" title and a descriptive paragraph: "This module is built using the SAS metadata DATA step functions to query the metadata repository, and return a list of all libraries and their associated directory or database schema."

CONCLUSION

The value of analytical transformations is an area that is gaining a lot of attention in the upcoming years. It is important to understand first the different users, the data they generate, and how they interact, and second the necessary capabilities that needs to be developed to harness this potential. Architecture discussed here is a step towards this direction. Developing a data-driven culture within organizations, investing in appropriate technology, fostering self-service analytical skills, and promoting a climate of organizational learning are critical factors in realizing value. It is also important to understand that various challenges exist before such a transformation can be achieved, and thus we need to change the existing IT & Business process of how we design information technology and analytical governance & practices.

As a next step, Elisa is planning to bring in more governance on the AI & Machine Learning as the models are created in R, Python and SAS depending on the performance. Thanks to the SAS Model Manager which helps in centralized model registration, versioning, and training and also monitors performance and automates the entire workflow process.

Data stored in diverse platforms generated by different tools also initiated the idea of Data virtualization particularly on the self service where the need is to integrate data from multiple, heterogeneous systems and provide a single access point to the consumers of data. Providing access to data from various data stores to the business teams and allow them to perform data discovery, analysis via SAS and other tools are the future focus areas in Elisa.

Finally, Elisa wants to move more and more towards SAS Viya capabilities as it brings in lot many exciting new features like centralized management of tools, users and data, higher throughput of work, a greater flexibility in terms of options - in memory data, API support, improved model governance and State-of-the-Art Machine Learning (ML) techniques.

CONTACT INFORMATION

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

Prasanna Pandian
Elisa Oyj, Finland
prasanna.pandian@elisa.fi

Jarno Lindqvist
SAS Institute Oy, Finland
jarno.lindqvist@sas.com