

## **TIME TO HARVEST: OPERATIONALIZING SAS® ANALYTICS ON THE SAP HANA® PLATFORM**

Christoph Morgen, SAP SE, Germany

### **ABSTRACT**

The ultimate value of farming analytics efforts in organizations with a maximum harvest is only achieved if analytics can also be operationalized at the level of core business applications. Mapped to the use of SAS® Analytics, how can the fruits of SAS be shared with Enterprise Business Applications by SAP? Learn how your SAS® environment – including the latest of SAS® In-Memory Analytics – can be integrated with SAP® applications based on SAP's In-Memory Platform SAP HANA®. We'll explore how a SAS® Predictive Modeling environment can be embedded inside SAP HANA and how native SAP HANA data management capabilities like SAP HANA views, smart data access and more can be leveraged by SAS applications and thus contribute to an end-to-end in-memory data management and analytics platform. Come and see how you can extend the reach of your SAS Analytics efforts with the SAP HANA integration.

### **INTRODUCTION**

With the SAS® Predictive Modeling Workbench for SAP HANA, organizations can now leverage tight integration between a SAS Advanced Analytics environment and the SAP HANA platform, the backbone of SAP core business applications. Moreover SAP HANA can service as a lightning fast in-memory database for non-SAP applications like SAS.

SAS and SAP have integrated SAS Advanced Analytics with the SAP HANA platform to speed up the entire analytical life cycle – from preparing and exploring data to developing, deploying and monitoring models. By enabling SAS predictive model scoring inside the SAP HANA platform provides faster valuable insights from SAP HANA's unique in-memory data processing capabilities and furthermore opens up the use of SAS predictive capabilities embedded into core business applications running on SAP HANA.

The purpose of this paper is to provide an overview of the current integration capabilities between SAS and SAP HANA, and to provide some details on relevant SAP HANA enhancements for SAS users and architects seeking to integrate their SAS and SAP HANA environments. A specific focus will be given to the topic of how SAS predictive models can be utilized by SAP HANA-based applications, providing insights into how the fruits of SAS analytic efforts can be easily shared and consumed by audiences and applications beyond the classic SAS applications.

### **GENERAL INTEGRATION AND SAP HANA ENHANCEMENTS**

#### **SAS AND SAP HANA INTEGRATION ENHANCEMENTS**

The key and foundation for integrating a SAS application with SAP HANA are the capabilities of SAS/ACCESS® to SAP HANA. It enables implicit push-down of SQL-based application logic into the SAP HANA platform as well as the use of HANA-specific SQL capabilities and extensions via explicit PROC SQL pass-through. A similar approach to push down application logic to SAP HANA can be found in the SAP ABAP® on HANA-based optimizations.

SAS applications like SAS® Marketing Automation or SAS® Asset Performance Analytics have been optimized to make use of this approach by delivering a specifically designed application data model for SAP HANA, to leverage the power of SAP HANA and the integration capabilities with related SAP applications.

These applications provide excellent examples of how you can optimize and accelerate your SAS application with SAP HANA by optimizing your data access patterns using PROC SQL or SQL push-down enabled SAS procedures, pushing the heavy data lifting or aggregation tasks to SAP HANA. Further optimization steps may involve moving data models from SAS to SAP HANA and leveraging more SAP HANA-specific capabilities like calculation views in order to leverage on-the-fly calculation of derived or aggregated information using virtualized data flows or virtualized multi-dimensional structures.

Enabling the core strength of SAS around Advanced Analytics with SAP HANA is a central focus of integrating SAS and SAP HANA. Hence the SAS Predictive Modeling Workbench for SAP HANA and its components are of key importance, bridging the path for enabling the SAS High-Performance Data Mining and Statistic extensions, interconnecting with SAP HANA for fastest SAS predictive modeling. Building on these fundamental integration paths, SAS® Visual Analytics and SAS® Visual Statistics have as well been enabled to work with SAP HANA to complete the family of SAS in-memory applications.

## SAP HANA DATA MANAGEMENT ENHANCEMENTS FOR SAS APPLICATIONS

### SAP HANA capabilities for SAS Analytics

If you store your application data in or consume your analytic base from SAP HANA, you may benefit from a variety of specific SAP HANA capabilities. Overall, you'll be able to utilize a fully SQL enabled in-memory data store. The in-memory columnar data store and its highly optimized data operators provide best performance specifically for analytic scenarios. Furthermore, the columnar store allows more agile columnar changes to tables, e.g. adding a new column without a complete table re-organization. A specific flavor of SAP HANA tables called flexible tables even supports up to 64,000 columns.

Reducing the in-memory data footprint is a key approach taken in most SAP HANA-based applications. This starts with highest in-memory data compression of your columnar store data and by reducing the need for indexes. An additional important step is to store only a single atomic copy of any data point, i.e. storing no aggregates at all, but calculating any derived information on the fly by completely virtualized data- models and calculation flows.

### SAP HANA enhancements for SAS Applications

While the in-memory core of SAP HANA provides maximum speed and performance, including support for scale-out- and largest scale-up in-memory systems, more flexibility and options in architecting an SAP HANA system beyond the in-memory borders have become more important, especially in Big Data- and Internet-of-Things-applications.

The SAP HANA **Dynamic Tiering** option adds smart, disk-based extended storage to your SAP HANA system. SAP HANA Dynamic Tiering enhances SAP HANA with large volume, warm data management capabilities. Data of highest value remains in-memory, whereas cooler, less-valuable data can be saved to the extended store. Tables within the extended store are an integral part of SAP HANA and thus are located within the SAP HANA database catalog and can be used as if they were persistent SAP HANA in-memory tables. Users and applications see the entire system as one single database. SAP HANA Dynamic Tiering allows adopting your SAP HANA environment with a multi-temperature data management approach according to your scenario and performance requirements, e.g. not necessarily all the data relevant to your SAS analysis requires to reside in-memory, but may be stored in the extended warm store. It can still be accessed from SAS as well as the SAS predictive models can be applied.

The SAP HANA **Data Warehousing Foundation** delivers specialized add-on applications, like a Data Distribution Optimizer to plan, adjust and analyze landscape reorganizations for SAP HANA scale-out systems, the Data Lifecycle Manager to support in displacing data between the SAP HANA in-memory persistency and the SAP HANA Dynamic Tiering extended store.

SAP HANA **Smart Data Access** is the extension and federation technology allowing customers to combine SAP HANA data with other heterogeneous data sources like **Hadoop**, SAP Sybase ASE/IQ and various other database systems. Data from remote sources is exposed as a virtual table, which can be combined with other data for example using SAP HANA views. A key capability of SAP HANA Smart Data Access is to support SQL functional compensation, which allows push down of filters, aggregates etc. to the remote sources so that most of the processing is done at the remote source and only smaller result sets are transferred across the network into SAP HANA. Hence tables from remote sources to your SAS environment integrated with SAS/ACCESS interfaces, can similarly be touched via SAP HANA Smart Data Access from your SAP HANA environment.

The SAP HANA **Hadoop integration** is of specific interest and value, as the combination allows to leverage both, fastest in-memory processing with the power of SAP HANA and hadoop's ability to store and process huge amounts of data, regardless of structure. The SAP HANA Smart Data Access adapter for hadoop support Hive SQL as well as Spark. Further within your SAP HANA development environment you can manage hadoop map-reduce routines, wrap them in as virtual functions in SAP HANA and execute them dynamically on the hadoop cluster, providing the ability to access and process data out of your SAP HANA based calculation flow. Specifically with respect to Internet-of-Things use cases, where large amounts of e.g. sensor-, log- and other machine data may require to be persisted, in reach for analysis and ready to be processed in place the hadoop integration will become much more relevant going forward.

Multiple isolated databases within a single SAP HANA system are now supported and referred to as SAP HANA **Multitenant Database Containers**. Database clients connect to specific container. All the databases in a multiple-container system share the same installation of database system software, the same computing resources, and the same system administration. However, each database is self-contained and fully isolated with its own set of database users, database catalog, repository, etc. Although database objects such as schemas, tables, views, procedures, and so on are local to the database, cross-database SELECT queries are possible. This supports in particular cross-application reporting scenarios. A SAS data mart or SAS application data model may be deployed within the same HANA system environment, however shielded within a separate container.

SAP HANA **Smart Data Streaming** adds high-performance event stream processing and complex event processing to the SAP HANA platform, making it easy for developers to incorporate smart stream capture and active event monitoring, alerting and event-driven response to their HANA applications. It allows you to filter, aggregate and enrich raw data before committing it to your database. You can capture data input from a variety of sources including data feeds, business applications, sensors, IT monitoring infrastructure and so on, apply business logic and analysis to the streaming data, store your results directly in SAP HANA and generate real-time alerts and responses.

Finally to conclude with this update on latest SAP HANA Data Management enhancements (based on the SAP HANA Platform 1.0 SPS09) customers can now leverage the SAP HANA **Enterprise Information Management** option, a more natively embedded framework to enhance, cleanse and transform incoming data from any source, provision and real-time- or batch-load the data into your SAP HANA system. SAP HANA EIM consists of two main areas, smart data integration and smart data quality. The adapter framework for remote source integration and real-time replication provides an SDK framework (similar to the smart data access adapter SDK), which allows to build custom adapters and extensions for non-standard sources.

## SAS PREDICTIVE MODELING INTEGRATION WITH SAP HANA

The **SAS Predictive Modeling Workbench for SAP HANA** integrates SAS Advanced Analytics with the SAP HANA Platform to speed up the entire analytical life cycle – from preparing and exploring data to developing, deploying and monitoring predictive models. Pushing the SAS analytic prediction closer to the data helps you extract valuable insights fast and within the core SAP business applications of your organization.

The SAS and SAP HANA integrated analytical life cycle includes

- Data access. Native integration between SAS and SAP HANA to read, write and update data via the SAS/ACCESS Interface to SAP HANA. Including consumption of SAP HANA specific data structures like calculation views.
- Data preparation and exploration. A powerful set of data preparation tools for greater ease and efficiency.
- Predictive model development. Streamlined data mining processes using SAS® Enterprise Miner™ to create highly accurate predictive and descriptive models.
- High-performance predictive model development. Tight integration with SAP HANA's massive parallel processing in-memory environment for analyzing big data with sophisticated modeling techniques available in SAS® High-Performance Data Mining.
- Distributed model deployment and execution. Automated deployment of predictive models to run directly on the data within SAP HANA.
- Model management and monitoring. Integration of analytical model management, deployment and monitoring with your SAP HANA data store.

The SAS Predictive Modeling Workbench for SAP HANA bundles Base SAS®, SAS/STAT®, SAS® Enterprise Miner™, SAS® Model Manager, SAS® Scoring Accelerator for SAP HANA, SAS/ACCESS® to SAP HANA®. With the SAS Scoring Accelerator for SAP HANA enabling SAS's industry-proven advanced analytics algorithms generated predictive models to run on SAP HANA by a SAS® Embedded Process.

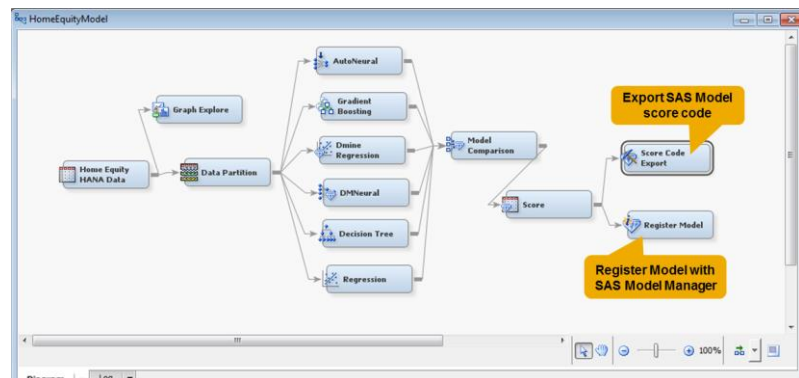
Furthermore the High-Performance Workbench bundle adds additional SAS in-memory speed-up by leveraging SAS® High-Performance Data Mining and SAS® High-Performance Statistics. SAS High-Performance Data Mining makes use of massively parallel in-memory algorithms and enables organizations to derive highly accurate and timely data mining models in minutes, not hours or days, thus massively accelerates the SAS model-building processes.

With SAS High-Performance Data Mining and -Statistics the SAS and SAP HANA integrated environment is typically extended with extra SAS compute nodes for the respective components, dedicated to process the most complex analytic problems running the SAS in-memory algorithms. Data is seamlessly passed between SAP HANA and the SAS High-Performance nodes leveraging the SAS Embedded Process for parallel data feeds to the SAS threads during the run-time of the algorithms.

## PUBLISHING SAS PREDICTIVE MODELS TO SAP HANA

As the focus of this paper is on exploring how to leverage the fruits of your SAS predictive modeling effort within SAP HANA, we will begin with looking at how SAS Enterprise Miner models can be published to SAP HANA.

So let's assume we've built a predictive model based on the home equity sample data and want to score or re-score the new home equity loan applicant data residing within SAP HANA. The SAS Enterprise Miner flow diagram could look like the following:



**Figure 1. Home equity loan default prediction model SAS Enterprise Miner diagram**

As an extension to SAS Enterprise Miner and SAS Model Manager, SAS® Scoring Accelerator for SAP HANA bridges the path to provision trained SAS predictive models inside the SAP HANA environment.

A fast forward approach to publishing the SAS Enterprise Miner generated models would be by making use model of publishing SAS macros. After preparing the SAS score code for in-database execution using the score code export-node in the SAS Enterprise Miner diagram, the score code output directly gets populated. As a next step the `%INDHN CREATE MODELTABLE-` and `%INDHN PUBLISH MODEL (<Options>)` Macros would create an SAP HANA table to store the SAS predictive models and publish the actual model score code from your score code project directory to the HANA table:

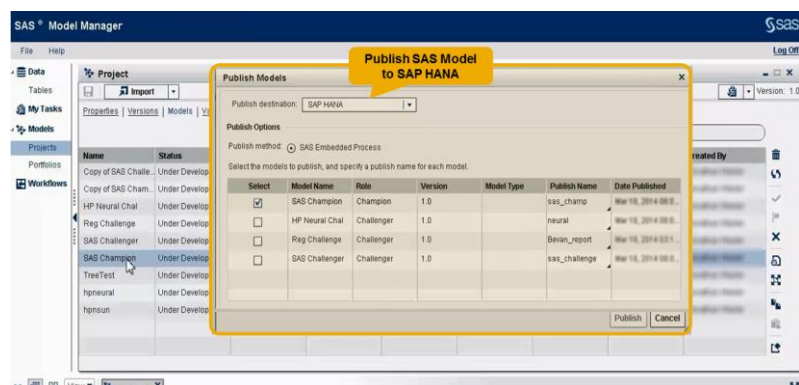
```
/* Provide connection information */
%let INDCONN=%str(server='<HANAServername>' instance=<Instance Number>
                  user=<HANA User> pw=<password>);

/* Create the model table in SAP HANA */
%INDHN_CREATE_MODELTABLE(modeltable=SASMODELTABLE, action=REPLACE);

/* Publish the model */
%INDHN_PUBLISH_MODEL( DIR=/usr/sas/emmodels/hmeq,
                      MODELNAME=hmeq_prod,
                      MODELTABLE=SASMODELTABLE,
                      ACTION=REPLACE);
```

If you require more governance and model management capabilities, then SAS Model Manager would provide the means and framework for a more automated model deployment process. SAS Model Manager allows you to register SAS Enterprise Miner modeling projects, validate, deploy, monitor and retrain SAS predictive models to ensure their sustained business value within your SAP HANA-based business application.

Using a role-based approach, authorized users within SAS Model Manager can quickly put your best models into production to the SAP HANA model repository table using the graphical model publishing user dialog



**Figure 2. SAS Model Manager, publish model to SAP HANA**

Depending on the targeted application and scenario, where the SAS predictive model shall be used and applied, either the scripted or the more managed approach using SAS Model Manager will be most suitable.

The model itself is stored within SAP HANA as SAS language score code and can only be executed by the SAS Embedded Process as detailed in the following section.

**Figure 3.  
SAS Model Table  
in SAP HANA**

MODELNAME	MODELDS2	MODELFORMATS	MODELMETADATA	MODELUUID	NOTES
hmeq_001	ds2_options sas tkmac; dat...	<?xml version="1.0" en...	?		BAD, MODEL=hmeq, TY
hmeq_prod	ds2_options sas tkmac; dat...	<?xml version="1.0" en...	?		Prediction for BAD, MO

**MODELDS2**  
 ds2\_options sas tkmac;  
 data SASEP.out;  
 #local\_TEMP;  
 #local\_LEAF;  
 dcl char(4) \_WARN;  
 dcl double CLAGE;  
 dcl double DEBTINC;  
 dcl double ...

## CONSUMING SAS PREDICTIVE MODELS IN SAP HANA FROM SAS APPLICATIONS

Once deployed, SAS Predictive Models can be called to directly run on the data within the SAP HANA platform, which eliminates the need to move data for scoring purposes out of SAP HANA to the SAS environment and hence further reduce the cost, complexity and latency of the analytic insight. By processing directly in the massively parallel SAP HANA environment, predictive models can provide answers extremely fast, enabling operational business decisions in real time.

Within SAP HANA, the predictive model to be applied is selected from the model table and passed along with the input data to the SAS Embedded Process within the SAP HANA system, in order to generate the predictive results. SAP HANA provides an Application Function Library Framework, which is been leveraged by a SAS provided application function to link from the SAP HANA execution request to the SAS Embedded Process.

If the consuming application of your predictions is a SAS application, you would make use of the SAS provided macro `%indhn run model (<Options>)` which initiates the SAS scoring process within SAP HANA.

The following shows an example of a SAS program code, applying the model “hmeq\_prod” to the data within the SAP HANA table “HMEQ” passing the results to the HANA table “HMEQ\_OUT”, restricted to the columns “CLNO EM\_SEGMENT EM\_PREDICTION”.

```
%let INDCONN=%str(server='<HANAServername>' instance=<Instance Number>
                    user=<HANA User> pw=<password> schema=<hanaschema>);

%INDHN_RUN_MODEL(modelname=hmeq_prod,
                  modeltable=SASMODELTABLE,
                  inputtable=HMEQ,
                  outputtable=HMEQ_OUT,
                  FORCEOVERWRITE=TRUE,
                  keep=CLNO EM_SEGMENT EM_PREDICTION);
```

If SAS Macro and SAS/ACCESS tracing options are switched on, you will find information in the SAS Log like

```
SAPHANA_64: Executed: on connection 2
call "SASLINK"."sas.ep::run"(
  p_model => 'hmeq_prod', b_force => 'Y', default_schema => 'SDEMO',
  input table name => 'SDEMO.HMEQ', output table name => 'SDEMO.HMEQ_OUT',
  sas model table name => 'SDEMO.SASMODELTABLE',
  keep list => ' CLNO EM_SEGMENT EM_PREDICTION ', debug level => 0,
  trace_level => 0, trace_table_name => '#sasep_logfile_ep',
  num_threads => 0, num_data_partitions => 0,
  db max text => 1024, logfile => ?
) with overview
Prepared statement SAPHANA 2
... SASLINK EP starting execution.
... SASLINK EP executing model 'hmeq_prod' with 1 DS2 thread(s) and 1 data partition(s).
... There were 1184844 rows read from the input data set and 1184844 rows written to the
output data set.
... SASLINK EP execution completed with SUCCESS.
```

The first section in the log examples indicates that an SAP HANA stored procedure has been called, then more information about the SAS Embedded Process execution is given.

## APPLYING SAS PREDICTIVE MODELS WITHIN SAP HANA APPLICATIONS

Now you've been asked to extend the usage footprint of SAS predictive models in your organization beyond of SAS applications. It's the SAP HANA application teams who demand to leverage the predictive power of your analytic models. With the SAS Scoring Accelerator for SAP HANA, the reach of a SAS data scientist analytics effort and its fruits in form of consumable predictive models now extends to your organizations core SAP business applications running on SAP HANA.

Valuable insights driven by your SAS predictive models can now be embedded in the machinery of SAP HANA based applications, reaching from real-time-, transactional- to analytic applications, built natively on SAP HANA or being SAP ABAP on HANA-based.

Let's have a look at the various options and flavors to see how your analytic fruits are being harvested in an SAP HANA environment.

### CONSUMING FROM NATIVE SAP HANA APPLICATIONS

#### Using SAP HANA SQLScript to consume SAS Predictive Models

Following up on some details of the example given above, the macro %INDHN\_RUN\_MODEL behind the scenes made use of an SAP HANA SQLScript Stored Procedure "sas.ep::run", which was deployed to your SAP HANA environment during the setup of the SAS Embedded Process.

As a general note, the SAS provided macro %INDHN\_RUN\_MODEL offers a lot of convenience functionality for SAS developers to leverage the SAP HANA embedded SAS scoring capabilities. SAP HANA developers, and potential SAS users with SAP HANA expertise as well, may alternatively make more explicit use of the SAP HANA SQLScript language and development tools. Surely explicit PROC SQL pass-through in SAS provides some of this development experience too.

As a first example, as part of your SAP HANA development effort, you can call on the SAP HANA SQLScript Stored Procedure "sas.ep::run" (e.g. in SAP HANA Studio's SQL Editor)

```
drop table HMEQ_OUT ;
truncate table log_output ;
call "SASLINK"."sas.ep::run" (
    p_model=>'hmeq_prod',
    sas_model_table_name=>'SDEMO.SASMODELTABLE',
    input_table_name=>'SDEMO.HMEQ',
    output_table_name=>'SDEMO.HMEQ_OUT',
    logfile=>SDEMO.LOG_OUTPUT,
    keep_list=>'CLNO EM_SEGMENT EM_PREDICTION') with overview;
```

The example shows the procedure call, applying the model "hmeq\_prod" from the model table "SDEMO.SASMODELTABLE" to the data within the SAP HANA table "SDEMO.HMEQ" passing the results to the HANA table "SDEMO.HMEQ\_OUT", restricted to the columns "CLNO EM\_SEGMENT EM\_PREDICTION":

Additionally, the output table must be deleted upfront and you want to truncate SAS Embedded Process Log-Output information. If the Log-Output structure doesn't exist yet, you can create it with a statement like

```
create table "<SCHEMA>".LOG_OUTPUT ("TYPE" VARCHAR(1), "TEXT" NVARCHAR(5000));
```

Still, the procedure "sas.ep::run" provides many convenience functions around the underlying use of a SAS application function for SAP HANA, linking the SAS Embedded Process to the SAP HANA processing via the SAP HANA application function library framework. Through these interfaces, an SAP HANA developer typically requires to create the describing types, parameter tables, a wrapper procedure etc. in order to consume a function for a given scenario.

A shortened example would look like

```
/* Creating and filling a Parameter-Table */
CREATE COLUMN TABLE MYPARMTABLE( "POSITION" INT, "SCHEMA_NAME" VARCHAR(100),
"TYPE_NAME" VARCHAR(100), "PARAMETER_TYPE" VARCHAR(100) );
insert into MYPARMTABLE values(1, 'SDEMO', 'HMEQ_IN_T', 'in');
insert into MYPARMTABLE values(2, 'SDEMO', 'SASMODELTABLE_T', 'in');
insert into MYPARMTABLE values(3, 'SDEMO', 'SAS_OPTIONS_T', 'in');
insert into MYPARMTABLE values(4, 'SDEMO', 'HMEQ_OUT_T', 'out');
insert into MYPARMTABLE values(5, 'SDEMO', 'LOG_OUTPUT_T', 'out');

/* Creating the Wrapper-Procedure for a scenario */
call SYS.AFLLANG_WRAPPER_PROCEDURE_CREATE('SASLINK_AREA', 'SASLINK_XMT', 'SDEMO',
'SASAFL_HMEQ_SCORE', MYPARMTABLE);
```

The example creates a wrapper procedure "SASAFL\_HMEQ\_SCORE" for the given prediction scenario.



Then executing the wrapper procedure and the actual prediction with the SAS Embedded Process looks not too different from the previous examples.

```
/* Calling the Wrapper Procedure */
call SDEMO.SASAFI_HMEQ_SCORE(HMEQ_1M,SASMODELTABLE, SAS_OPTIONS, HMEQ_OUT,
LOG_OUTPUT) with overview;
```

While the effort for an SAP HANA developer going the explicit path of calling the SAS application function for a scenario compared to the more convenient approach of just using the provided procedure “sas.ep::run” is significant, the added benefits of the explicit approach include:

- obtaining SAS HANA development artefacts, which can be deployed along with other application components between environments,
- greater flexibility with respect to the usage context of the SAS prediction call. By using read-only procedures and avoiding to write to permanent output structures, the SAS prediction calls can be leveraged within SAP HANA calculation views, and thus build upon the key concept and approach on the SAP HANA Platform to enable real-time data processing and aggregation, by virtualizing set-based calculation data flows.

In the following, we want to briefly overview the approaches and tools in order to facilitate the efforts of an SAP HANA developer faced with the task of embedding the SAS analytic fruits into an SAP HANA application.

## Using Graphical SAP HANA Editors to consume SAS Predictive Models

The SAP HANA **Application Function Modeler** in SAP HANA Studio is a graphical design editor introduced recently to facilitate development of data flow graphs leveraging SAP HANA application functions as well as batch or real-time data transformation and provisioning scenarios, all in-memory on the SAP HANA Platform.

The extension capabilities of the Application Function Modeler allow one to create and use templates for specific custom functions. The following example shows the use of a template data flow model, embedding the use of the SAS application function, thus enables graphical drag-n-drop use of the SAS prediction function as part of an SAP HANA development effort, by avoiding the manual SQL-Scripting effort indicated in the previous example (creation of describing types, parameter tables, wrapper procedures etc).

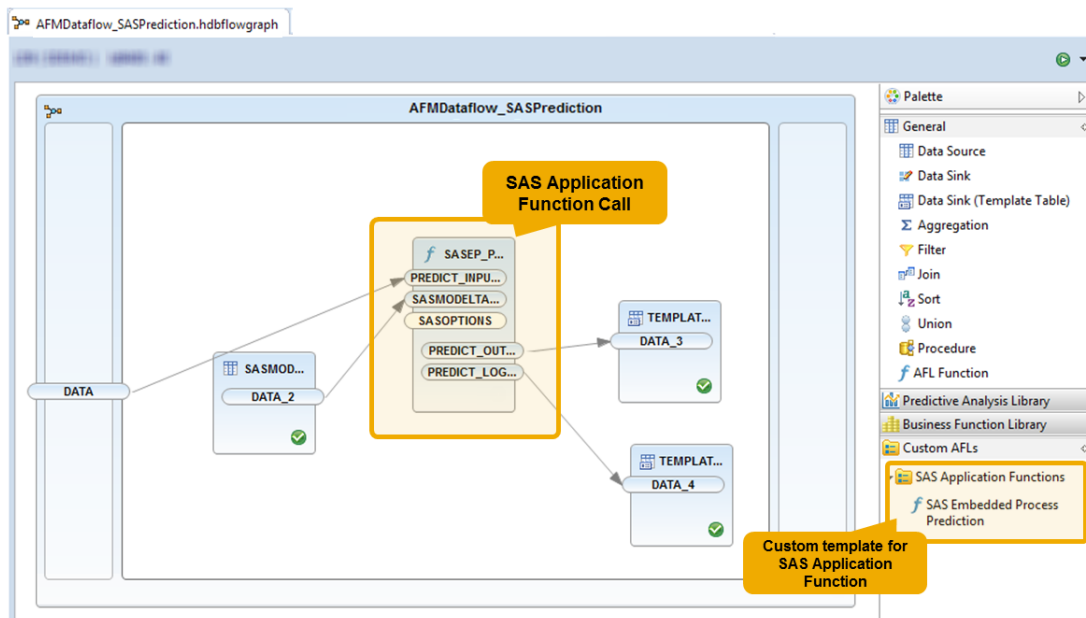


Figure 4. SAS HANA Application Function Modeler flowgraph using the SAS application function

The resulting dataflow development object can now be combined and transported along with other objects of an SAP HANA application. Upon activation of a dataflow, an SAP HANA SQLScript procedure is created and can be similarly called as shown in the example code below.

```
/* Calling the AFM flowgraph generated procedure, passing in data to process */
CALL "SDEMO"."sas.sasdev::AFMDataflow_SASPrediction"("SDEMO"."HMEQ");
```

Application function models embedding SAS prediction function calls may suit in two ways: firstly, by easing the effort to implement custom prediction scenarios with SAS predictive models on HANA and their re-use in other development contexts; secondly, by facilitating the use of SAS predictive models within a data transformation scenario itself, i.e. through data massaging efforts as the data is provisioned into the SAP HANA environment.

With the SAP HANA **Calculation View Editor** the use of SAS predictive models can be integrated and combined with any other calculation data flow and thus can become a part of the key SAP HANA approach of real-time and virtualized data processing flows on top of in-memory data structures.

There are now two ways to embed the SAS prediction calls we discussed in calculation views.

Firstly, as calculation views support SAP HANA SQLScript table functions as input data structures, wrap the call statement inside a table function and consume that within your graphically modelled calculation view.

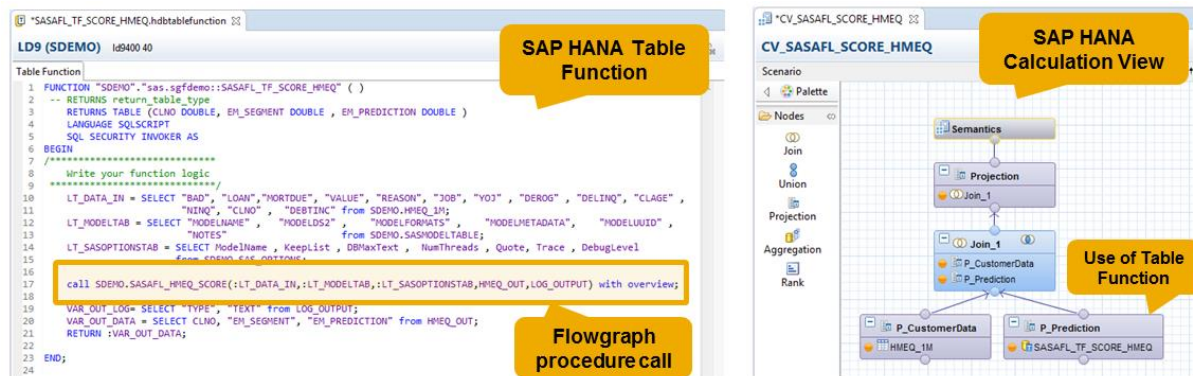


Figure 5. Left: SAS HANA table function referencing the flowgraph generated procedure. Right: SAP HANA Calculation View referencing the table function call.

Secondly, you can wrap the call statement within a script-based SAP HANA calculation view, which allows read-only SQLScript processing. The scripted calculation view itself can again be re-used and referenced within other calculation views.

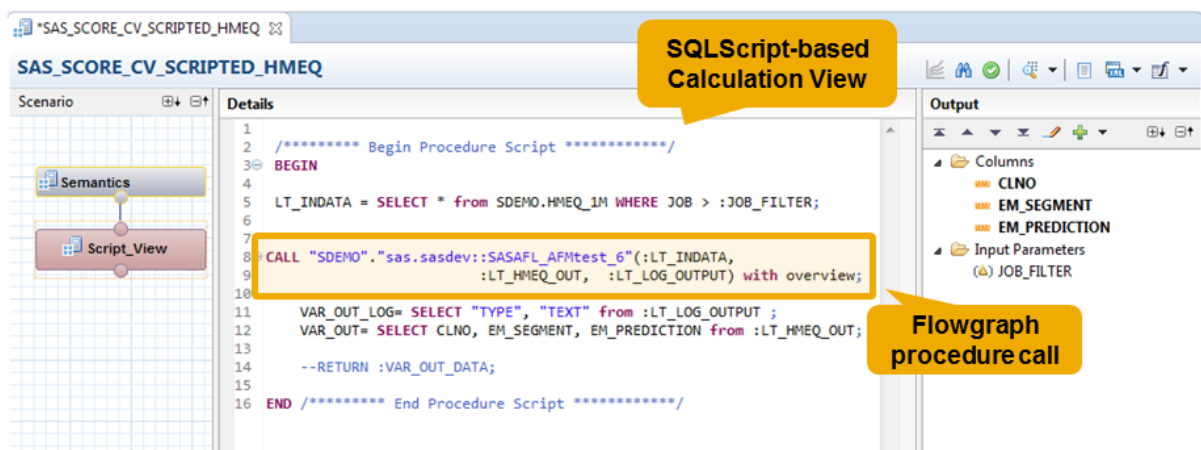
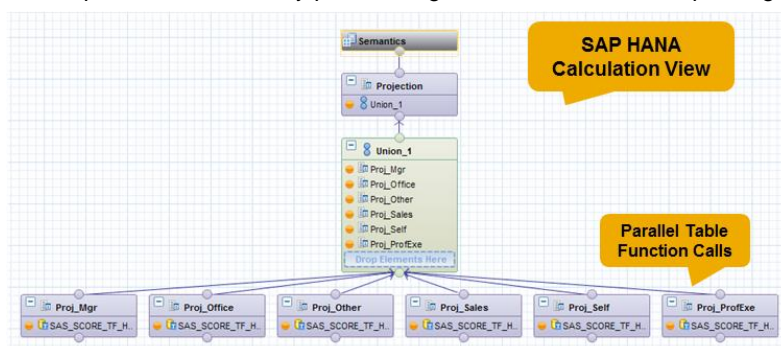


Figure 6. SAS HANA SQLScript-based Calculation View.

The table function call or the call of the script-based calculation view allow one to pass in parameters, which enables you to add more dynamic to your SAS prediction call, i.e. by parallelizing the functions calls and passing in different input data, e.g. filtered slices of the same data source.

Figure 7. SAS HANA Calculation View with multiple parameter-driven table function calls





We've looked at the key approaches on how to leverage SAS predictive models within SAP HANA-based native applications, by using SAP HANA SQLScript stored procedures, table functions, flowgraph models using the SAS application function and their embedding into SAP HANA calculation views. Based on those concepts SAP HANA developers in your organization can take action and embed the fruits of your SAS analytic efforts into SAP HANA applications and thus provide higher business value with those applications.

To round this up, surely one can now query the SAP HANA calculation view with any SAP HANA enabled BI client or application. Or we can just simply use a SQL query like the following




Figure 8. SAS HANA Studio SQL Editor query on the Calculation View.

Most surely, within your SAS environment you can make use of SAP HANA calculation views as well, as introduced and discussed in a previous paper (Morgen 2014). If your SAS/ACCESS Libname to HANA is directed to a different schema than the system schema “\_SYS\_BIC” it may be a convenient practice to create a synonym for the respective calculation view within the SAP HANA schema in use by your SAS users.

Example statement creating a synonym for a Calculation View.

```
CREATE SYNONYM SDEMO.CV_SAS_SCORE_PARALLEL_HMEQ
FOR "_SYS_BIC"."sas.sasdev/SAS_SCORE_CV_MULTIPLE_HMEQ";
```

Now, within your SAS application code you can reference the calculation view like a schema table.

```
LIBNAME HANA SAPHANA <options> SCHEMA=SDEMO;
proc SQL;
select count (CLNO)
  from (SELECT CLNO, EM_PREDICTION FROM HANA.CV_SAS_SCORE_PARALLEL_HMEQ
        where EM_PREDICTION > 0.9);
quit;
```

The SAS Log will show the push-down of the query and treat the view synonym as if it were a table.

```
SQL_IP_TRACE: passed down query:      select COUNT(TXT_1."CLNO") from
SDEMO.CV_SAS_SCORE_PARALLEL_HMEQ TXT_1 where TXT_1."EM PREDICTION" > 0.9
SQL_IP_TRACE: The SELECT statement was passed to the DBMS.
```

Note: As a refresher, these are some commonly used SAS tracing options in the context of SAS/Access, which you may want to consider to use here.

```
OPTIONS SQL_IP_TRACE=ALL /* To start with */
SASTRACE=',,,d' SASTRACELOC=SASLOG NOSTSUFFIX /* For more details */
/* To control common implicit SQL pass-through behavior */
DBIDIRECTEXEC SQLCONSTDATETIME SQLREDUCEPUT=ALL SQLGENERATION=DBMS;
```

## CONSUMING FROM SAP ABAP-BASED APPLICATIONS

SAP ABAP® is the programming language core for SAP business applications. Alongside with Java, it builds up to the SAP NetWeaver® Application Server. ABAP will remain the foundation for key SAP applications like SAP Business Suite incl. S/4HANA, Business Warehouse and other new High-Performance Applications build upon and integrated with SAP HANA.

The latest ABAP releases provide numerous SAP HANA integration capabilities. The Code-to-Data paradigm describes the overall recommended approach. It means delegating the processing of data-intensive computation from the application server to the database. Open SQL, a database abstraction layer in SQL processing (comparable to SAS implicit PROC SQL pass-through) and ABAP Managed Stored Procedures, enabling SAP HANA SQLScript source code editing in ABAP (comparable to SAS explicit PROC SQL pass-through), along with other optimizations, enables a tight integration between ABAP and SAP HANA.

The ability to leverage SAP HANA SQLScript from within ABAP programs, hence also enables to leverage SAS predictive models in SAP HANA (the fruits of your SAS analytic effort!) within a huge domain of SAP business applications and customer extensions. A tremendous potential, to add more value to any kind of business process by enrichment with prediction capabilities.

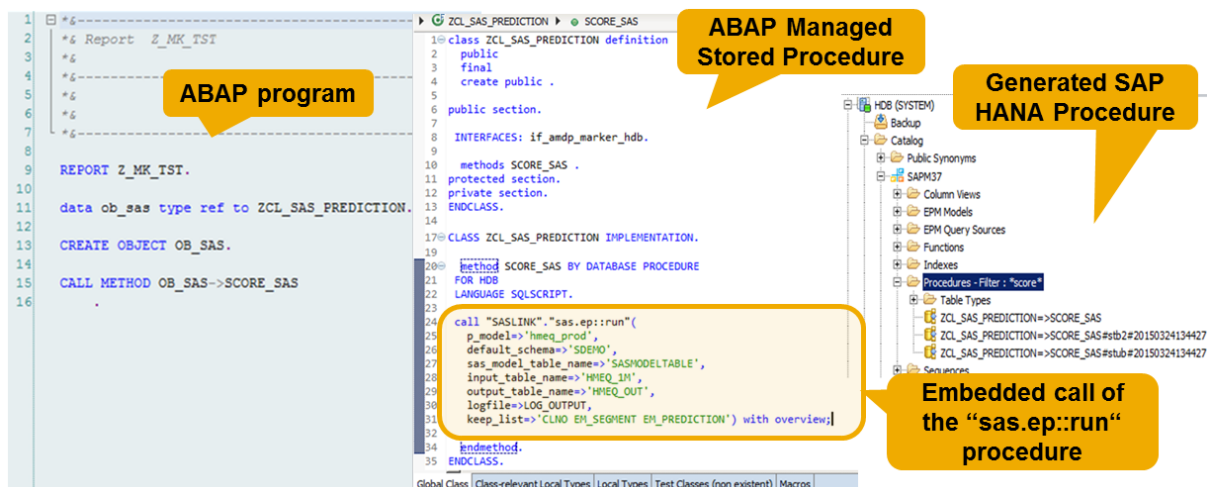


Figure 9. SAP ABAP program calling the stored procedure “sas.ep::run” for prediction calls

## CONCLUSION

With the SAS Predictive Modeling Workbench for SAP HANA, organizations can now leverage tight integration between a SAS Advanced Analytics environment and the SAP HANA Platform, the backbone of SAP core business applications. The entire analytical life cycle and analytic business process is accelerated – from preparing and exploring data, to develop and train predictive models, the automated deployment of trained models into SAP HANA and the overall model management process. SAS predictive models can then be applied entirely embedded on SAP HANA by any type of application.

Ultimately, only by applying predictive models in the context of a business process the business value is finally realized. Hence by extending the reach of your SAS analytics efforts to the domain of SAP business applications running on SAP HANA, organizations are now enabled to harvest and operationalize the analytics fruits to their maximum benefit and impact.

Moreover SAP HANA can service as a lightning fast in-memory database for non-SAP applications like SAS, providing enhanced data management capabilities like SAP HANA dynamic tiering or multitenant database containers providing customers with more flexibility and options to adopt their SAP HANA system landscape to their needs and onboard new applications like SAS with their data stores.

## DISCLAIMER

The contents of this paper are the personal work and view of the author. All details and recommendations in this document have been compiled on a best efforts basis for informational purposes only, without representation or warranty of any kind. The content is considered to be accurate; however neither the author nor SAP SE shall be liable for errors or omissions with respect to the material.

## REFERENCES AND RECOMMENDED READING

Morgen, Christoph (2014): "Empowering SAS® Users on the SAP HANA Platform" , Proceedings of the SAS® Global Forum 2014, available at <http://support.sas.com/resources/papers/proceedings14/2445-2014.pdf>

SAP SE (2014, 2015): For general SAP HANA information explore <http://hana.sap.com/abouthana.html>, documentation is available at [http://help.sap.com/hana\\_platform](http://help.sap.com/hana_platform). For an overview on referenced SAP HANA enhancements introduced with the SAP HANA Platform Support Package Stack 09 see the following blog <https://blogs.saphana.com/2014/10/21/what-is-new-in-sap-hana-sps-09/>

SAS Institute Inc. (2014): "SAS/ACCESS® 9.4 for Relational Databases: Reference, Fourth Edition", SAS/ACCESS Interface to SAP HANA. Available at <http://support.sas.com/documentation/>

SAS Institute Inc. (2014): "SAS® Scoring Accelerator for SAP HANA". The documentation for the SAS Scoring Accelerator for SAP HANA is contained in the following guides: SAS 9.4 In-Database Products: User's Guide, as well as Administrator's Guide, Fifth Edition. Available at <http://support.sas.com/documentation/>

SAS Institute Inc. (2014): "SAS® Predictive Modeling Workbench for SAP HANA". Product Brief, available at [https://www.sas.com/content/dam/SAS/en\\_us/doc/productbrief/sas-predictive-modeling-workbench-sap-hana-hppmw-107314.pdf](https://www.sas.com/content/dam/SAS/en_us/doc/productbrief/sas-predictive-modeling-workbench-sap-hana-hppmw-107314.pdf) and [http://www.sas.com/en\\_us/software/analytics/predictive-modeling-workbench.html](http://www.sas.com/en_us/software/analytics/predictive-modeling-workbench.html)

## CONTACT INFORMATION

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

Name:	Christoph Morgen
Enterprise:	SAP SE
Address:	Dietmar-Hopp-Allee 16
ZIP City:	69190 Walldorf, Germany
Work Phone:	+49 6227 742554
E-mail:	Christoph.Morgen@sap.com
Linkedin:	<a href="http://www.linkedin.com/pub/christoph-morgen/17/57b/543">http://www.linkedin.com/pub/christoph-morgen/17/57b/543</a>

## TRADEMARK INFORMATION

SAS and all other SAS Institute Inc. product or service names are registered trademarks or trademarks of SAS Institute Inc. in the USA and other countries. ® indicates USA registration.

SAP and other SAP products and services mentioned herein as well as their respective logos are trademarks or registered trademarks of SAP AG (or an SAP affiliate company) in Germany and other countries. Please see <http://www.sap.com/corporate-en/legal/copyright/index.epx> for additional SAP trademark information and notices.

Other brand and product names are trademarks of their respective companies.