

SAS® and SAP Business Warehouse on SAP HANA – What's in the Handshake?

Tatyana Petrova, SAS Institute Inc.

ABSTRACT

Is your company using or considering using SAP Business Warehouse (BW) powered by SAP HANA? SAS® provides various levels of integration with SAP BW in an SAP HANA environment. This integration enables you to not only access SAP BW components from SAS, but also to push portions of SAS analysis directly into SAP HANA, accelerating predictive modeling and data mining operations. This paper explains the SAS toolset for different integration scenarios, highlights the newest technologies contributing to integration, and walks you through examples of using SAS with SAP BW on SAP HANA. The paper is targeted at SAS and SAP developers and architects interested in building a productive analytical environment with the help of the latest SAS and SAP collaborative advancements.

INTRODUCTION

Handshakes are powerful. A healthy, positive handshake can produce agreement, collaboration, combined efforts, and mutual benefits to the parties. A healthy, positive handshake between leading innovators in the field of information technology can deliver new opportunities to sophisticated enterprises that are looking for information-leveraging advancements.

In 2013, SAS and SAP established a partnership that brought to life a new set of technologies that unite advanced analytics with the power and speed of a modern database.

The SAP in-memory database, SAP HANA, is designed to provide real-time, massively parallel processing. It combines the full set of features of a relational database with in-memory computing and efficient data compression capabilities. It is also built for scalability, since parallelization is incorporated into the platform design. SAP HANA attracts existing SAP BW users as well as other business intelligence and analytics tools consumers who want to get data analysis questions answered faster and with support of significantly larger volumes of data.

For more than a decade, SAS has been integrated with SAP BW and its predecessors. This integration has included the ability for SAS to retrieve SAP metadata and content for complex statistical and analytical calculations.

The newest development originated from the strengthened handshake is focused on direct communication with the SAP HANA layer and movement of computing tasks inside the SAP HANA database where the data sources reside. Although you can still work with SAP BW application layer, the new technology gives you a faster route into the SAP HANA database.

This paper explains four integration options and how each can be useful, depending on the analytical task, the load, and where you are in your adoption of SAP HANA as an underlying data platform.

Each integration option is reflected in one of the following product packages:

1. SAS® Data Surveyor for SAP
2. SAS/ACCESS® Interface to SAP HANA
3. SAS® Predictive Modeling Workbench for SAP HANA
4. SAS® High Performance Predictive Modeling Workbench for SAP HANA

This paper also includes information about SAS solutions for SAP HANA.

Co-innovations between SAS and SAP are ongoing; new products and changes to the existing ones are in development. This paper is based on the second maintenance release for SAS 9.4 and SAP NetWeaver BW 7.3 on SAP HANA 1.0 SPS8.

SAP TERMINOLOGY

Before jumping into the integration specifics, it is useful to review some basic SAP terminology. The SAP BW modeling layer is robust and consists of various building blocks. For the purpose of this paper, we'll provide simplified definitions for a few types of objects:

- **DataStore Objects (DSOs):**
objects that store transactional data at its most granular level. Depending on the purpose of a DSO, it can consist of a table or a table with extra objects that support its operations.
- **InfoObjects:**
smallest SAP BW Business Intelligence building blocks. Other objects such as InfoCubes and DSOs consist of InfoObjects. Examples of InfoObjects are Characteristics (categorical values) or Key Figures (numeric measures).
- **InfoProviders:**
objects that enable data for analysis purposes. InfoCube is a type of InfoProvider.
- **InfoCubes:**
SAP BW representation of an OLAP entity. Consists of a set of dimension tables and fact tables.
- **HANA-optimized InfoCubes:**
an SAP BW InfoCube converted into an SAP HANA version. In a HANA-optimized InfoCube, dimension tables are logical definitions (as opposed to physical tables) that remove a level of complexity from the modeling perspective and that speed up data loading. As suggested by the term, HANA-optimized InfoCubes are optimized for queries execution in SAP HANA.

SAS DATA SURVEYOR FOR SAP

The first level of integration gives you the ability to access SAP BW objects from SAS.

This functionality has been available for years. It was established long before SAP HANA entered the database arena. SAS Data Surveyor for SAP enables you to work with BW entities. It supplies plug-ins and transformations that enable other SAS products (for example, SAS® Data Integration Studio and SAS® Management Console) to integrate with SAP BW.

In the SAP BW on HANA architecture, SAS Data Surveyor for SAP communicates with the application layer. It provides the capability to navigate SAP BW metadata from SAS and extract data and objects definitions from SAP into SAS for further analysis.

Here are some typical tasks that you can accomplish with the help of SAS Data Surveyor for SAP:

- **Exploring SAP BW metadata:**
From SAS Data Integration Studio, you can browse SAP BW metadata that is extracted from SAP into SAS, or you can use the SAP BW Metadata search feature in SAS Data Integration Studio to search pre-extracted SAP metadata. For example, you can search which InfoCubes contain data on 'Customers'.
- **Reading SAP BW InfoCubes:**
Using SAS Data Integration Studio as a client, you can import all or a portion of an SAP BW InfoCube into a star-schema-based collection of SAS data sets. You can then use SAS to query these SAS data sets. The InfoCube importing job is auto-generated by SAS Data Surveyor for SAP.
- **Analyzing and reporting on SAP data:**
For example, using SAS® Enterprise Guide® and its graphical interface (Query Builder), submit a query on SAP data (DSOs) and produce a report.
- **Writing data back to SAP BW:**

SAS Data Surveyor enables you to transform SAS data sets into SAP BW InfoObjects. SAS data sets should conform to the particular structure. Both DSOs and Infocubes can be written from SAS to SAP BW. This functionality is available from SAS Data Integration Studio via SAP transformations.

Figure 1 highlights the roles of SAS and SAP for some typical client requests.

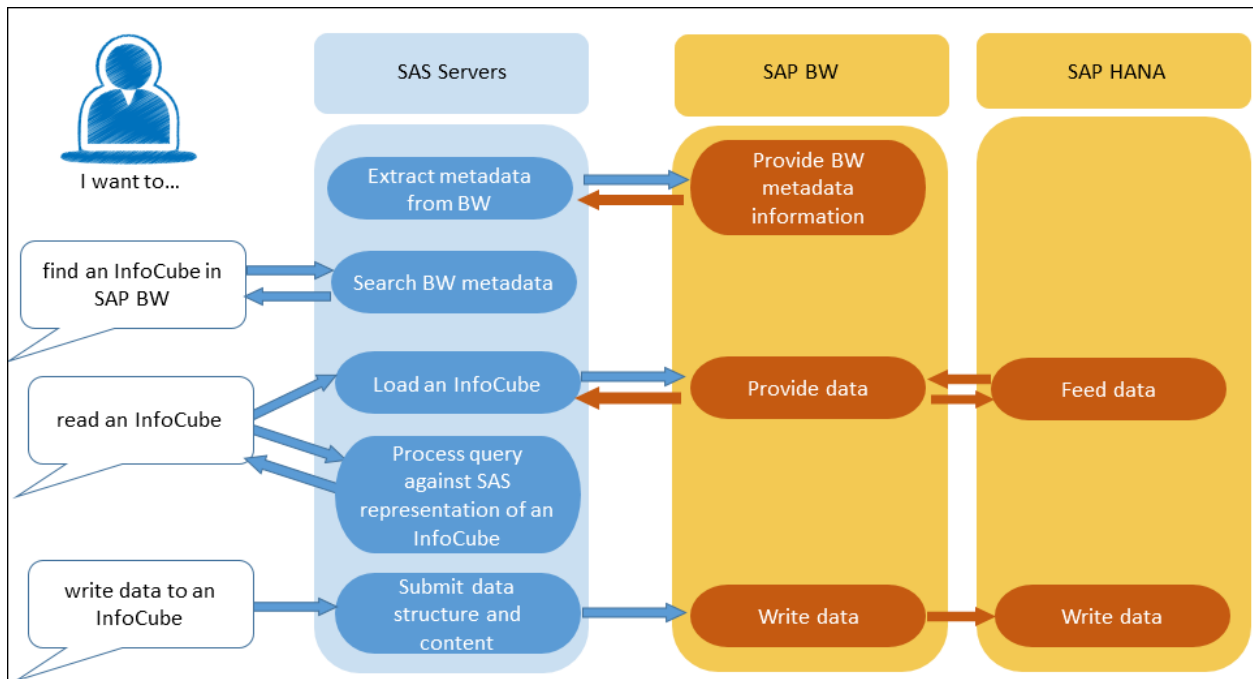


Figure 1. SAS Data Surveyor for SAP: Sample Process Flow for Typical Requests

Does it matter for SAS Data Surveyor for SAP what data platform SAP BW is based on: SAP HANA or some other platform?

SAS Data Surveyor for SAP works with SAP BW regardless of the data platform SAP BW is hosted on. It relies on SAP BW to deliver content from data sources. But starting with the first maintenance release for SAS 9.4, SAS Data Surveyor for SAP is extended with a feature to recognize HANA-optimized InfoCubes. The feature adapts SAS Data Surveyor for SAP wizards to accommodate the simplified data model that HANA-optimized InfoCubes introduce. See Figures 2 and 3 for examples of the SAS Data Integration Studio interface that is used to select SAP InfoCube components to be populated in SAS. Note that extra dimension tables are not presented in the list of available InfoObjects. Rather, a single fact table, master data tables, and the SID tables appear reflecting the simplified HANA-optimized InfoCube structure.

Figure 2 shows the InfoObjects that are available for extraction. The InfoObjects in the list are associated with a particular HANA-optimized InfoCube.

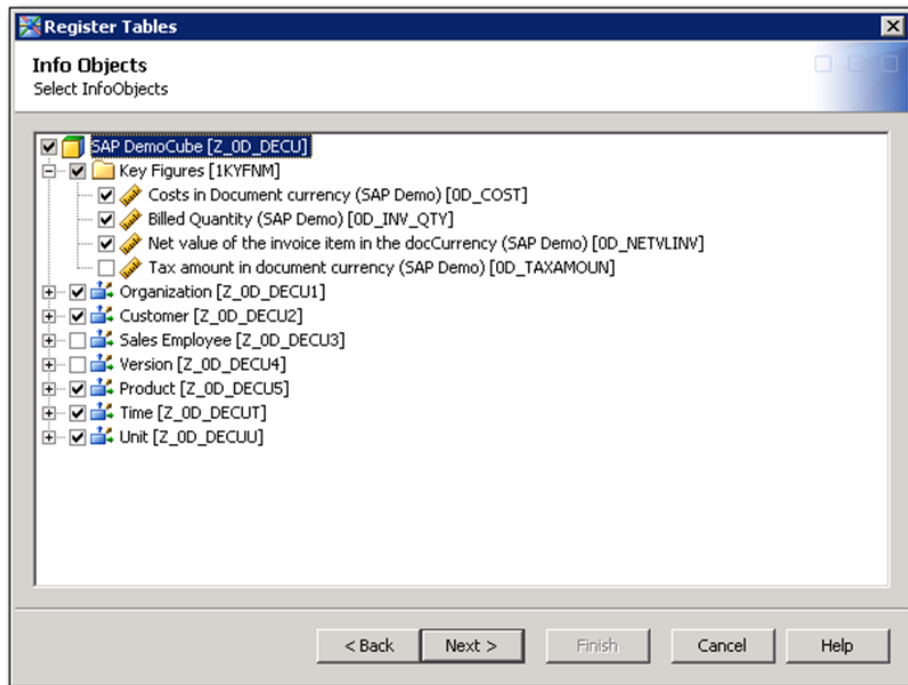


Figure 2. SAS Data Integration Studio: SAP InfoObjects Available for Extraction to SAS

Figure 3 shows the list of associated with HANA-optimized InfoCube objects selected for extraction to SAS.

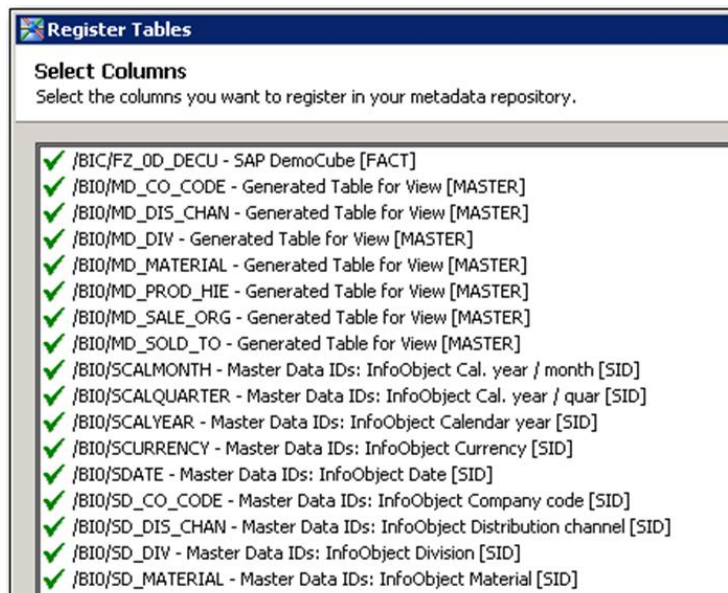


Figure 3. SAS Data Integration Studio: List of Objects Selected for Extraction from SAP BW Hana-Optimized InfoCube

The lack of extra dimension tables comes in especially handy when populating corresponding SAS data sets during the loading of a SAP HANA-optimized InfoCube in SAS Data Integration Studio. You no longer see steps for loading supporting dimension tables in your job flow. Omitting extra dimension tables results in a faster loading transaction.

SAS Data Surveyor for SAP uses SAP Remote Function Calls (RFC) as a connection and communication mechanism. RFC is a standard SAP interface that allows SAP native systems and other vendors'

applications to communicate with SAP and exchange information about SAP objects. This mechanism differs from methods used by other SAS-to-SAP integration levels explained further in this paper and requires a separate configuration.

It is important to reiterate that SAS Data Surveyor for SAP does not directly connect to the database platform and does not support in-database processing. The analytics tasks to be performed on this data are processed in SAS preceded with a step of data retrieval from SAP.

To enable you to connect from SAS directly to SAP HANA and move data processing into its powerful in-memory environment, a new set of technologies has emerged.

With these technologies, new levels of integration have become possible enabling faster and broader analytical calculations to be added to customers' scenarios.

SAS/ACCESS INTERFACE TO SAP HANA

The SAS/ACCESS Interface to SAP HANA uses the SAP HANA ODBC driver as its connecting mechanism and enables SAS to reach out directly to SAP HANA.

As with other SAS/ACCESS interfaces, SAS/ACCESS Interface to SAP HANA provides capabilities to read, write, and update data in the SAP HANA database, perform data definition manipulations, and submit database-native function calls. It also opens the door for other SAS data management and analytics products to take advantage of in-database processing in SAP HANA. We cover additional toolsets in the following sections, but first, here is an overview of SAS/ACCESS Interface to SAP HANA.

SAS/ACCESS Interface to SAP HANA provides support for both types of code pushdown:

- SQL Explicit Pass-through:

You write code that can be interpreted by SAP HANA (for example, DML and DDL statements and SAP HANA function calls) and submit it from SAS to run completely in the database.

- SQL Implicit Pass-through:

You write code in SAS syntax and allow SAS to convert it into code that can be interpreted by SAP HANA. As much as possible, this code runs in the database. SQL implicit pass-through comes to play for a typical Base SAS® program when you point the corresponding LIBNAME to SAP HANA data sources.

Figure 4 demonstrates a flow for a typical implicit pass-through request.

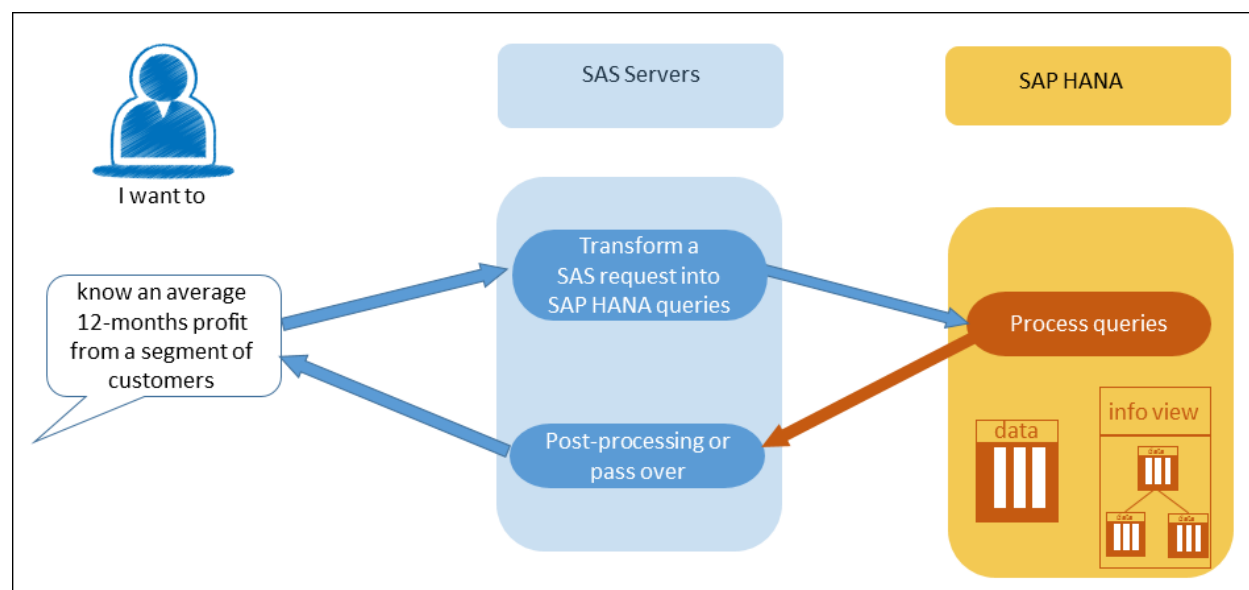


Figure 4. SAS/ACCESS for SAP HANA: Sample Process Flow for a Typical Request

Additional helpful capabilities of SAS/ACCESS Interface to SAP HANA include buffered reads and bulk loading (speeding up reads and writes of data accordingly), extended data types mapping, and a variety of pushdown functions, including aggregation functions. SAS/ACCESS Interface to SAP HANA also supports working with SAP HANA information views. You can submit a query to an information view by referring to it by a LIBNAME and an information view name, as you would when working with a database table. All corresponding subqueries to access and process data sources for the information view are implicitly generated and submitted to SAP HANA on your behalf.

To illustrate the implicit code pushdown in action, we look at two examples:

- The SQL procedure, your gateway into anything RDBMS
- The MEANS procedure, one of the typical statistical procedures

First of all, set the tracing option to log the exact SQL code that arrives at SAP HANA:

```
OPTIONS SASTRACE=',,,d' SASTRACELOC=saslog NOSTSUFFIX details;
```

Establish a connection:

```
LIBNAME hanalib SAPHANA SERVER=myHANAServer INSTANCE=00 SCHEMA=mySchema
USER=myUser PASSWORD=myPassword;
```

EXAMPLE 1. PROC SQL CODE PUSHDOWN

```
PROC SQL;
  SELECT AVG(profit_12_months) AS Profit, IncomeGroup
  FROM hanalib.bankc_in_hana
  WHERE UPCASE(AgeGroupID)="MIDDLE AGED"
  GROUP BY IncomeGroup
  ORDER BY Profit DESC;
QUIT;
```

The statement in the log confirms that execution occurred in the database. Due to the tracing option set, the log also contains the exact SQL received by SAP HANA. Notice that the SAS UPCASE function gets converted to the SAP HANA corresponding UPPER function.

```
SAPHANA_8449: Prepared: on connection 0
  select AVG(TXT_1."PROFIT_12_MONTHS") as "Profit", TXT_1."INCOMEGROUP" from
tapetr.BANKC_IN_HANA
TXT_1 where UPPER(TXT_1."AGEGROUPID") = 'MIDDLE AGED' group by
TXT_1."INCOMEGROUP" order by "Profit"
desc

ACCESS ENGINE:  SQL statement was passed to the DBMS for fetching data.
```

EXAMPLE 2. PROC MEANS CODE PUSHDOWN

```
PROC MEANS DATA=hanalib.bankc_in_hana;
  VAR Profit_12_Months;
  CLASS AgeGroupID IncomeGroup;
  TYPES () AgeGroupID*IncomeGroup;
RUN;
```

The log confirms successful pushdown.

```
SAPHANA_8453: Prepared: on connection 0
select COUNT(*) as "ZSQL1", MIN(TXT_1."AGEGROUPID") as "ZSQL2",
MIN(TXT_1."INCOMEGROUP") as "ZSQL3",
COUNT(*) as "ZSQL4", COUNT(TXT_1."PROFIT_12_MONTHS") as "ZSQL5",
MIN(TXT_1."PROFIT_12_MONTHS") as
"ZSQL6", MAX(TXT_1."PROFIT_12_MONTHS") as "ZSQL7",
SUM(TXT_1."PROFIT_12_MONTHS") as "ZSQL8",
COALESCE(VAR(TXT_1."PROFIT_12_MONTHS")*(COUNT(TXT_1."PROFIT_12_MONTHS")-
1),0) as "ZSQL9" from
tapetr.BANKC_IN_HANA TXT_1 group by TXT_1."AGEGROUPID", TXT_1."INCOMEGROUP"
```

ACCESS ENGINE: **SQL statement was passed to the DBMS for fetching data**

These examples are over-simplified to highlight the point, but more complex data manipulations or statistical calculations are handled in a similar manner. Some queries are broken down into a set of individual queries for a pushdown. In some situations where not all code conversion is possible or is meaningful, you might see some post-processing take place in SAS.

You can perform powerful calculations with just SAS/ACCESS Interface to SAP HANA combined with Base SAS. In addition, SAS/ACCESS Interface to SAP HANA lays the foundation for other layers of SAS-to-SAP integration and serves as a baseline for several other SAS product packages that are enabled for in-database processing.

If your analytics tasks call for more involved manipulations on SAP HANA data, you might benefit from strengthening the handshake with further integrated toolsets.

SAS PREDICTIVE MODELING WORKBENCH FOR SAP HANA

Does your business application require predictive modeling?

Do you plan to develop predictive models and apply them to target data to identify the most responsive customer segments for your marketing campaign or the likelihood of insurance claims? Do you plan to forecast trends to support your analytics-driven business decision?

If you use or plan to use predictive analytics to perform similar tasks, consider using the SAS Predictive Modeling Workbench for SAP HANA, which helps combine predictive modeling with fast, in-database scoring.

SAS Predictive Modeling Workbench for SAP HANA is a package of products that includes the following main players (as of the second maintenance release for SAS 9.4):

- Base SAS
- SAS/ACCESS Interface to SAP HANA
- SAS Enterprise Guide
- SAS® Scoring Accelerator for SAP HANA
- SAS® Enterprise Miner™
- SAS® Model Manager

These products, while supporting your predictive modeling lifecycle, all play a role in the SAS-to-SAP integration.

Base SAS and SAS/ACCESS Interface to SAP HANA are the foundation.

SAS Scoring Accelerator for SAP HANA is the product that enables in-database scoring.

SAS Enterprise Miner is a powerful, predictive analytics instrument that supports various stages of predictive modeling development from the data preparation and exploration stages all the way through model validation and generation of the scoring code.

SAS Model Manager supports the later stages of the analytical model lifecycle: models deployment and management. Model deployment includes models publishing, while models management includes models evaluation and monitoring. With the help of SAS Model Manager you can operationalize your models and be informed when it is time to retrain your model on an updated set of data, or create a new one.

SAS Enterprise Miner and SAS Model Manager together support the full model lifecycle. The corresponding model definitions to be shared across these tools are communicated via the SAS metadata repository.

Combined with the SAS Scoring Accelerator for SAP HANA, SAS Model Manager provides a user-friendly interface for models publishing to SAP HANA. After a model is published, data stored in SAP HANA can be scored without leaving the database premises.

Alternatively, if you prefer to publish models in batch, SAS Enterprise Miner enables you to export the scoring code as a set of model definition files, which, in turn, can be manually published to SAP HANA from a SAS session or scheduled as a publishing job.

Look at the following usage example, which walks you through the complete process flow using SAS. There are multiple ways the flow can evolve, and for some steps, different tools can be used. But the following example shows a representative sequence of steps.

EXAMPLE

A bank needs to target a new credit card campaign to its segment of customers who previously showed the most return on former credit card campaigns. The bank has data accumulated from the previous campaign that can be used to build an analytical model and predict the behavior of a larger set of customers.

SUMMARY OF THE TASKS IN THE EXAMPLE

Use a **SAS data set** as a data source for model creation. Develop the model. (Processing of model calculations takes place in **SAS® Workspace Server**.) Select the champion model. Register the champion model in SAS metadata. Publish the model in **SAP HANA**. Run the model scoring code to score new data in **SAP HANA** with all scoring taking place in **SAP HANA**.

DEVELOPMENT FLOW IN DETAIL

1. Develop and register the model in SAS Enterprise Miner using the following steps, as shown in Figure 5.
 - a. Prepare the data. In our example, these are nodes Impute and Data Partition.
 - b. Develop alternative models by using various modeling methods. In our example, these are nodes Regression, Decision Tree, and Neural Networks.
 - c. Compare the models and select the champion model.
 - d. Generate the scoring code.
 - e. Register the model in SAS metadata or export the score code as a set of model definition files.
2. Create the model table in SAP HANA using SAS code. The model table holds models scoring code. Perform this step once per an environment setup.
3. Operationalize the model using SAS Model Manager:
 - a. Create a SAS Model Manager project to hold associated objects.
 - b. Import the model from SAS metadata.
 - c. Publish the model to SAP HANA.
4. Run the model scoring code in SAP HANA to score the new data.

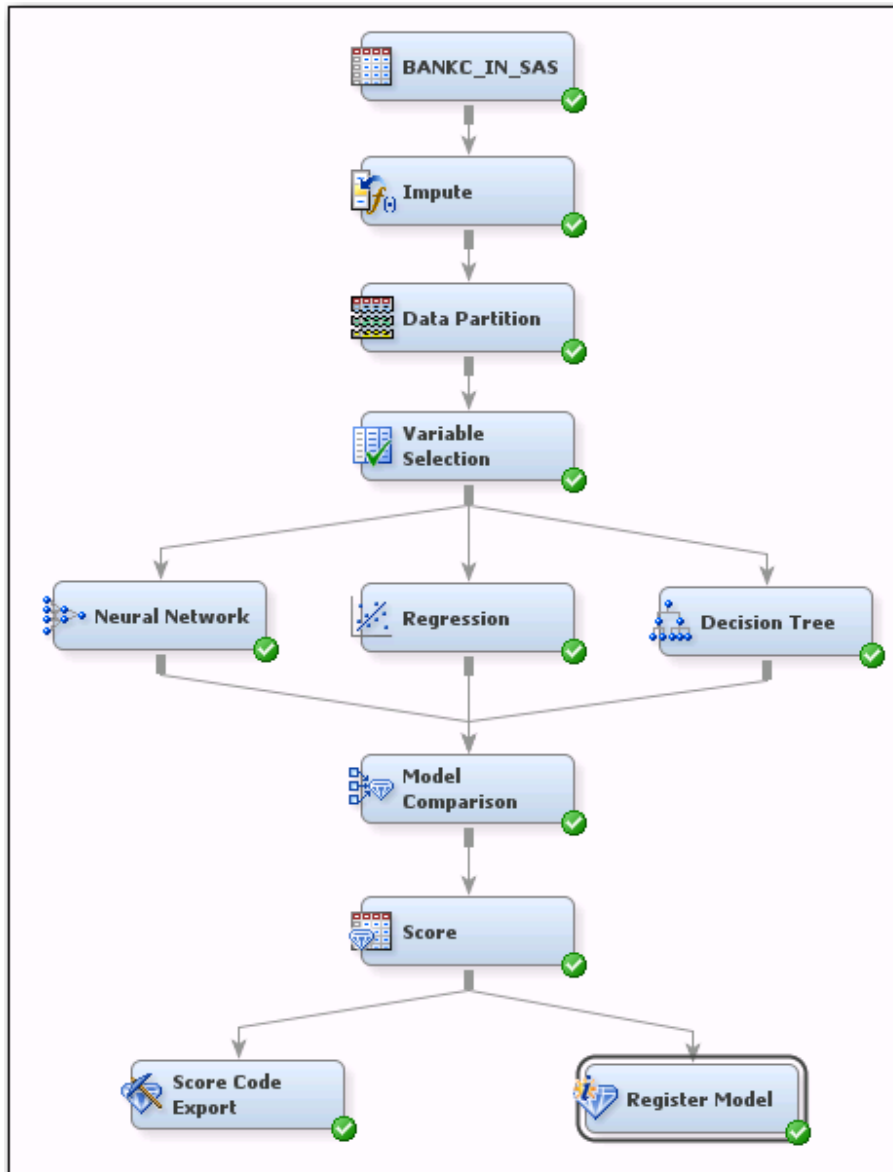


Figure 5. Development Flow for a Credit Card Campaign Model in SAS Enterprise Miner

In this example flow, you see the integration between SAS and SAP HANA during the publishing of the model to SAP HANA and the running of the model scoring code in SAP HANA.

A model is published only once for each model version. Behind the scenes, publishing consists of converting model scoring code to SAS DS2 format and inserting it into a model table in SAP HANA.

Scoring new data, on the other hand, is typically a repetitive and processing-heavy task. Various options are available to execute this step on-demand or on a schedule and can be triggered from either SAS or SAP HANA. Regardless of where the job is submitted from, the scoring process executes fully in SAP HANA.

See Figure 6 for examples of typical tasks and the roles that SAS and SAP play in their processing.

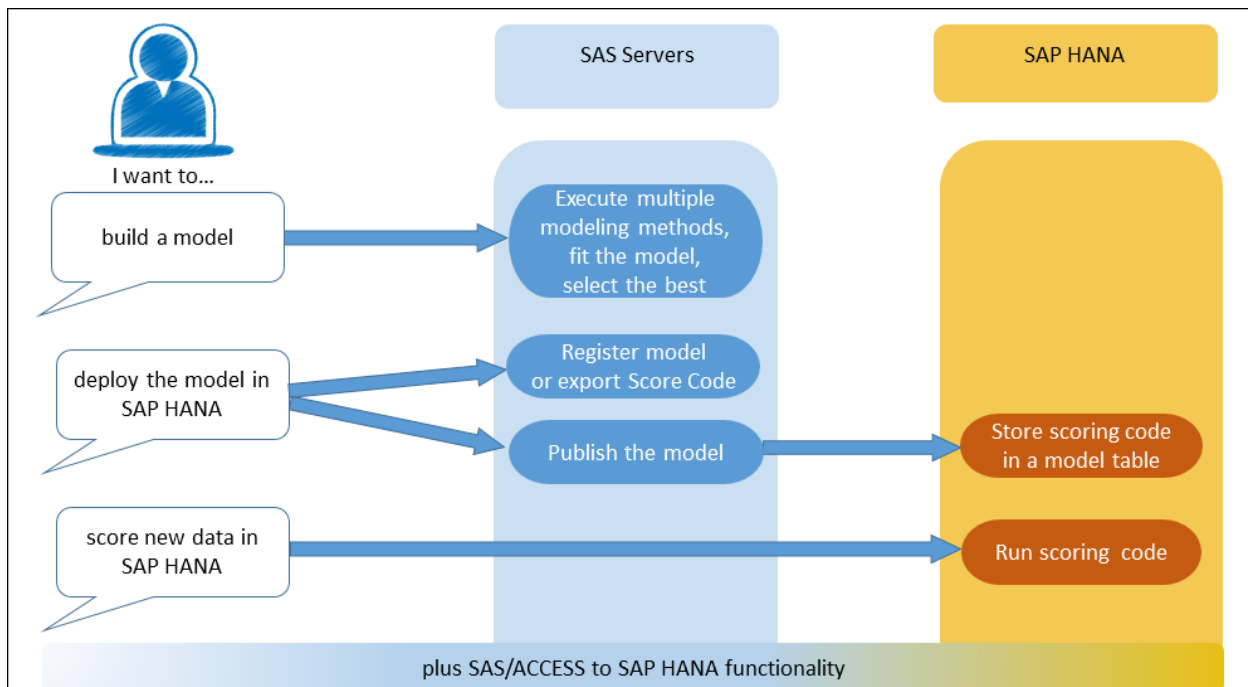


Figure 6. SAS Predictive Modeling Workbench: Sample Process Flow for Typical Requests

CONTRIBUTING TECHNOLOGIES

SAS Predictive Modeling Workbench for SAP HANA builds upon functionality provided with the SAS/ACCESS interface to SAP HANA. In addition, it uses a newly developed, underlying technology called the SAS Embedded Process. The SAS Embedded Process is not a product, but a component that comes with SAS Scoring Accelerator. It engages and maintains communication between SAS and SAP HANA and provides an even stronger collaboration between the systems.

The SAS Embedded Process component gets deployed and runs in SAP HANA as a separate process. The data segments are exchanged within a shared memory space in a multithreaded fashion. This approach allows a secure and fast execution of a SAS DATA step (DS2 program) next to the data source.

Looking back at the SAS Predictive Modeling Workbench usage example, the scoring model code generated by SAS Enterprise Miner and published to SAP HANA becomes the DS2 program that can run in database via the SAS Embedded Process.

SAS HIGH PERFORMANCE PREDICTIVE MODELING WORKBENCH FOR SAP HANA

Let's say SAS Predictive Modeling Workbench for SAP HANA explained above appeals to your business scenario.

You might be thinking: *Great, I can publish my scoring code right next to data in SAP HANA and run it right there! My scoring code is processing-intensive, and running it faster would really help me.*

You might also be thinking: *But wait. Scoring data is only a portion of the task. While I develop the model itself, I have to run and rerun my modeling experiments to find the one that best fits the data. This process takes me a really long time. Can you help me?*

SAS High Performance Predictive Modeling Workbench for SAP HANA can help you.

SAS High Performance Predictive Modeling Workbench for SAP HANA is an add-on option for SAS Predictive Modeling Workbench for SAP HANA. It offers distributed high-performance analytics technology to use for models development.

SAS® High-Performance Analytics (HPA) enables you to run modeling algorithms in-memory in a single-machine mode or in distributed multi-nodes environment. Distributed in-memory processing significantly

improves the computing power. Faster execution of algorithms provides opportunity for more experiments, more iterations and model refinements, and more variables to analyze. You can use a greater volume of training data rather than relying on sample subsets. Even better, you can run multiple scenarios in parallel and try new techniques and more advanced statistical methods.

To use SAS HPA in a distributed mode, a separate appliance is required. An appliance in this context is a cluster of computing nodes with SAS High Performance Analytics infrastructure installed. This appliance is an extra player in SAS and SAP BW on HANA architecture and should be added to the ecosystem if you need faster, distributed execution of the model development phase.

Distributed SAS High Performance Predictive Modeling Workbench for SAP HANA is designed to be able to lift data from SAP HANA to the appliance memory quickly, in parallel, and on-demand. The data gets cleared from HPA memory after the analytical calculations are completed. All data lifting happens seamlessly for the user. Here are some of the visible differences for the model creator:

- The data sources feeding into the SAS Enterprise Miner modeling flow reside in SAP HANA.
- High-performance nodes (HPMD) in SAS Enterprise Miner are used instead of traditional modeling nodes.
- By default, properties of high-performance nodes are set at a significantly higher number of iterations.
- As model trains on data on the appliance, sampling of data for the purpose of reducing the training data size becomes unnecessary. HP Data Partition node takes care of splitting the original data source into training and validation subsets and loads them in-memory on the appliance.

Look at the familiar SAS Predictive Modeling Workbench usage example and see how we can improve our model development process and accuracy with High-Performance level of integration.

EXAMPLE

The basic scenario is the same. The bank is using predictive analytics for customer segmentation.

The task and the modeling flow are similar, with the following exceptions:

- the location of data sources for the model development phase,
- the size of the training data,
- characteristics of modeling algorithms,
- an increased number of iterations during the model training phase.

We rephrase the summary of tasks accordingly.

SUMMARY OF TASKS IN THE EXAMPLE

Use the **SAP HANA table** as a data source for model creation. Develop the model. (Processing of model calculations takes place in **SAS HPA**.) Select the champion model. Register the champion model in SAS metadata. Publish the model in **SAP HANA**. Run the model scoring code to score new data in **SAP HANA** with all the scoring taking place in **SAP HANA**.

The process flow sequence does not change except for the new configuration step to point SAS Enterprise Miner to the SAS HPA environment. While the sequence of steps does not change, the nature of data preparation and model development does. High-Performance versions of SAS Enterprise Miner nodes are used for these steps to be run in SAS HPA. High-Performance versions of the methods are not created exactly equal to their non-HP siblings. For example, there are differences in statistics available in train properties for some of the methods. The default settings for HP nodes also differ to reflect the opportunity for the increased number of iterations and depth of calculations.

Due to the increased power of processing, additional analytic algorithms can be tried for model development. In our scenario, an HP Forest node is added to the flow, which can result in a more accurate prediction than the one from a Decision Tree.

DEVELOPMENT FLOW IN DETAIL

1. In SAS Enterprise Miner, add code to be executed at the beginning of the project to set options for the SAS HPA environment.
2. Develop and register the model in SAS Enterprise Miner, as shown in Figure 7.
 - a. Prepare the data. Use High-Performance versions of SAS Enterprise Miner nodes for data preparation: in our example, HP Impute, HP Data Partition, and HP Variable Selection.
 - b. Develop alternative models by using various High-Performance modeling methods: in our example, HP Regression, HP Tree, HP Neural, and HP Forest.
 - c. Compare the models and select the champion model.
 - d. Generate the scoring code.
 - e. Register the model in SAS metadata or export the score code as a set of model definition files.
3. Create the model table in SAP HANA using SAS code. The model table holds models scoring code. Perform this step once per an environment setup.
4. Operationalize the model using SAS Model Manager:
 - a. Create a SAS Model Manager project to hold associated objects.
 - b. Import the model from SAS metadata.
 - c. Publish the model to SAP HANA.
5. Run the model scoring code in SAP HANA to score new data.

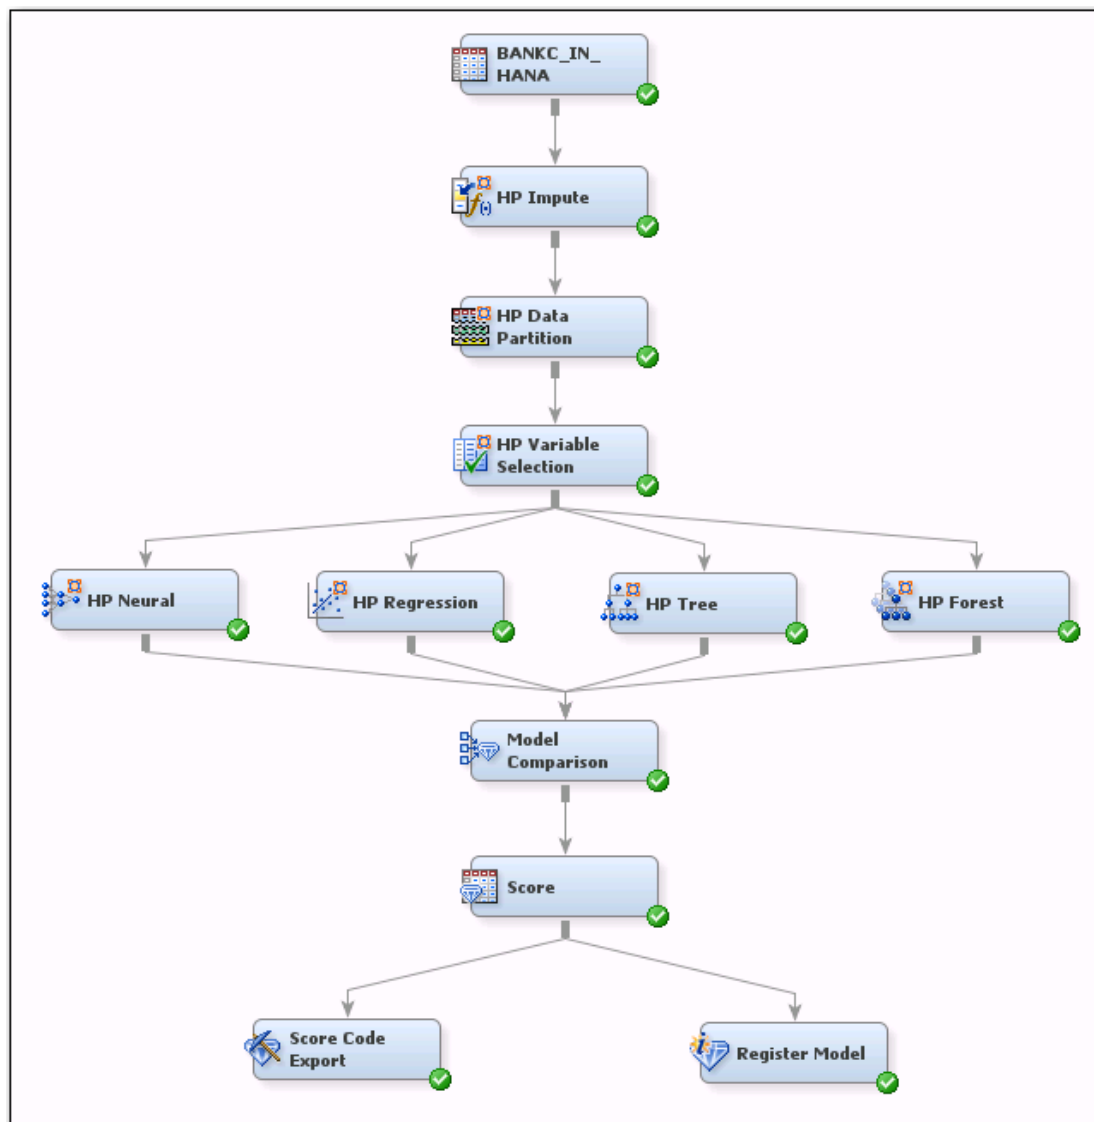


Figure 7. SAS Enterprise Miner Diagram for a Credit Card Campaign Model Development Flow: High Performance Version

As in the Predictive Modeling Workbench example, you see the integration between SAS and SAP HANA during the publishing of the model and the running of the model scoring code. In addition, to establish a high-speed parallel data lifting from SAP HANA into SAS HPA memory, a fast data-exchange mechanism is used. SAS Embedded Process technology, as well as a newly developed Data Feeder for SAP HANA component of SAS HPA, enable an effective communication between the tiers. See Figure 8 for a sample process flow between SAS and SAP HANA for this example.

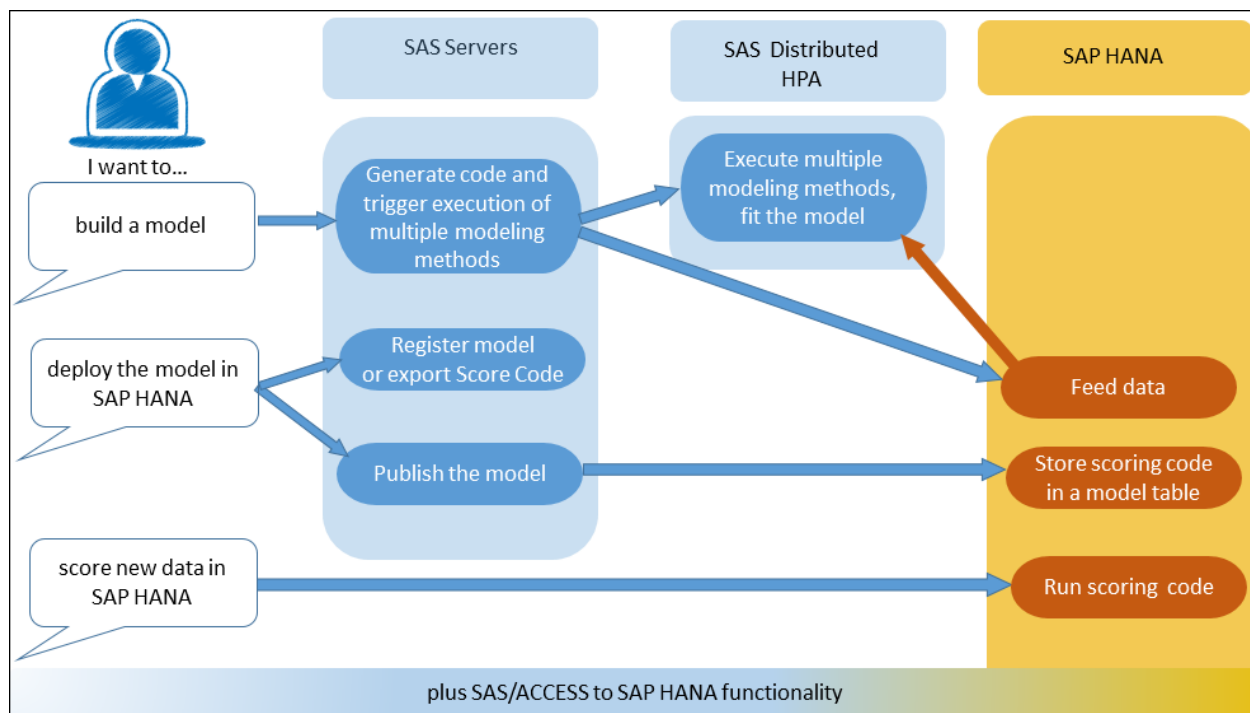


Figure 8. SAS High Performance Predictive Modeling Workbench: Sample Process Flow for Typical Requests

SAS SOLUTIONS FOR SAP HANA

SAS solutions are specialized applications aimed to comprehensively address themed analytics projects. Some of SAS analytics solutions that rely heavily on quick and scalable access to enterprise data, can use the new integration technologies for in-database processing of SAP HANA data sources. These solutions are leveraging SAS/ACCESS Interface to SAP HANA as a connection and a pushdown to database mechanism. Access to SAP HANA data sources is incorporated into the logic of these solutions, and auto-generated database requests are optimized for the best usage of SAP HANA capabilities.

The following SAS solutions are enabled for in-database integration with SAP HANA in the second maintenance release for SAS 9.4:

- **SAS® Marketing Automation for SAP HANA**

A solution for planning, testing, and executing marketing campaigns. It gives you a comprehensive toolset to build and manage multiple marketing campaigns while your customer data resides in SAP HANA.

- **SAS® Asset Performance Analytics for SAP HANA**

A solution to maximize your asset utilization. It combines sensor data with other critical for systems health monitoring information and provides a framework for maintenance and performance assessment, root-cause analysis of failures, and predictive analytics capabilities to forecast and minimize potential issues in the future.

- **SAS® Field Quality Analytics for SAP HANA**

This solution integrates field quality data from warranty systems, customer service records, and other sources to detect and analyze emerging quality issues to determine the cause and impact on overall performance. Reporting and alerting mechanisms allow for a timely response, while forecasting capabilities are designed for proactive future planning.

WHICH INTEGRATION LEVEL IS THE BEST FOR YOU?

This paper covers four levels of integration, as well as SAS specialized solutions that are available for your engagement from SAS to SAP BW on HANA.

Four levels of integration are reflected in four products packages:

1. SAS Data Surveyor for SAP
2. SAS/ACCESS Interface to SAP HANA
3. SAS Predictive Modeling Workbench for SAP HANA
4. SAS High Performance Predictive Modeling Workbench for SAP HANA

Each level comes with its benefits, costs, and considerations. At which level do you *want* to be engaged? At which level do you *need* to be engaged? Which scenario applies *to your situation*?

The previous sections describe typical functional applications for each integration level.

But mapping features to your analytical tasks is just a part of the equation.

Another part is an overall strategy in advancing your SAP environment and scalability requirements. SAS integration should be consistent with this strategy.

- Where does your company stand in migrating its Business Warehouse data platform to SAP HANA?

Regardless of what your current state of migration is, SAS provides the toolset to accommodate access to SAP data for analytical purposes.

Perhaps you are in the testing mode for one subsystem and in production for another. Or maybe you have replaced the underlying platform but have not started any object conversion to leverage the SAP HANA optimizations, yet. Perhaps you are going through a step-by-step approach with the conversion. Whatever your situation, you can continue to use SAS Data Surveyor for SAP to accommodate for a non-disruptive access to data.

If new data marts are introduced fully in SAP HANA and co-exist with former BW-based applications, and both types of systems' content needs to be called for SAS analytics purposes, you might need to have both types of integration available to you: SAS Data Surveyor-based and SAS/ACCESS Interface to SAP HANA-based.

If full migration is in scope, then connections and technologies based on SAS/ACCESS Interface to SAP HANA should suffice for getting to SAP content.

- Will the data platform be fully replaced with SAP HANA or will portions of data sources feeding into BW be coming from other types of DBMS?

If SAS analytics is expected to be based on data coming from various data platforms at once, in-database capabilities can be affected, depending on the scenario.

- Will you be heavily leveraging recently introduced HANA-optimized InfoProviders? And will you be fully or partially migrating the modeling layer from BW to HANA?

While SAS Data Surveyor supports HANA-optimized InfoCubes and can read and write BW content, in-database processing is not enabled for this approach. If you are investing your resources in moving the modeling logic from SAP BW standard objects into the SAP HANA layer, you will reap more performance benefits when working from SAS directly with SAP HANA entities via SAS/ACCESS Interface to SAP HANA.

- How much data will be used for analytics?

Performance and scalability requirements are at the core of the integration level decision. The stronger the handshake between the systems, the more efficiently the communication and requests exchange occur, and the more optimum resources utilization takes place.

CONCLUSION

For the SAP Business Warehouse consumer, adoption of SAP HANA offers an exciting transformation. If your work relates to SAS integration with SAP and is affected by this change, should you be excited? SAS believes you should.

Use the information in this paper to navigate these expanded integration opportunities with comfort.

You now know which new technologies are available and which SAS product packages support which integration levels. You understand the conceptual differences between your choices. You also know that you can continue to use your current methods of connecting from SAS to SAP. In fact, you can use the current methods to take advantage of some of the modernized SAP BW InfoProviders. But you also know that you can gradually introduce new approaches of integration and plan for faster and more advanced collaboration between SAS and SAP sides of your enterprise information systems.

REFERENCES AND RECOMMENDED READING

Morgen, Christoph. 2014. "Empowering SAS® Users on the SAP HANA Platform." *Proceedings of the SAS Global Forum 2014 Conference*. Washington, DC: SAS Institute Inc. Available <http://support.sas.com/resources/papers/proceedings14/2445-2014.pdf>.

SAS Institute Inc. 2014. "SAS/ACCESS Interface to SAP HANA." *SAS/ACCESS 9.4 for Relational Databases: Reference*, 6th ed. Available at <http://support.sas.com/documentation/cdl/en/acrelldb/67589/HTML/default/viewer.htm#n0wni260o8qeodn18ue7uft6qzhx.htm>

SAS Institute Inc. 2014. "SAS Scoring Accelerator for SAP HANA." *SAS 9.4 In-Database Products: User's Guide*, 5th ed. Available at <http://support.sas.com/documentation/cdl/en/inbug/68170/HTML/default/viewer.htm#p1o0m4ja305mv6n1vj78bsnqcn5k.htm>

SAS Institute Inc. 2014. "SAS Enterprise Miner 13.2: Reference Help", Available at <http://support.sas.com/documentation/onlinedoc/miner/>

SAS Institute Inc. 2014. "SAS Enterprise Miner 13.2: High-Performance Procedures", Available at <http://support.sas.com/documentation/onlinedoc/miner/>

SAS Institute Inc. 2014. "SAS High-Performance Analytics Server 12.2: User's Guide", Available at <http://support.sas.com/documentation/onlinedoc/hpa/index.html>

Palekar, Amol, Bharat Patel, and Shreekant Shiralkar. 2012. *SAP NetWeaver BW 7.3 – Practical Guide*. 2nd edition. SAP PRESS.

Berg, Bjarne, and Penny Silvia. 2013. *SAP HANA: An Introduction*. 2nd edition. SAP PRESS.

ACKNOWLEDGMENTS

The author's sincere gratitude goes to these colleagues for their support and knowledge sharing:

Jeff Bailey
Nancy Bremmer
Heike Czichy
Mark Freskos
Lou Galway
Diane Hatcher
Shirish Javalkar
Dominique Latour
Raghu Mangaraju
Howard Plemmons
Priya Sharma
Jonathan Wexler

CONTACT INFORMATION

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

Tatyana Petrova
SAS Campus Drive, Cary, NC
SAS Institute Inc.
Tatyana.Petrova@sas.com

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.

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