Paper A01-2009

High Performance Analytics with In-Database Processing

Stephen Brobst, Chief Technology Officer, Teradata Corporation, San Diego, CA
Keith Collins, Senior Vice President & Chief Technology Officer, SAS Institute, Cary, NC
Paul Kent, Vice President, R&D, Platform Research and Development, SAS Institute, Cary, NC
Michael Watzke, Teradata Architect, Teradata Corporation, Fitchburg, WI

ABSTRACT
A new era of high performance analytics has emerged. In-database processing for deep analytics enables much faster turnaround for developing new analytic models and dramatically reduces the cost associated with data replication into "shadow" file systems. This talk examines these trends and quantifies the impacts. We also provide best practices for execution of a phased deployment strategy of these capabilities.

INTRODUCTION
It is estimated that the time spent in data management tasks using traditional methods to build and deploy a predictive model is approximately 70% of the timeframe from project initiation to realization. In the typical scenario, an experienced SAS® programmer is expending inordinate effort locating, extracting, integrating, and preparing data for model development. Data assets are often gathered from across multiple sources that exist in a variety of repository technologies. In competitive times when speed of execution is critical, this lack of agility is unacceptable. Moreover, the resource costs associated with inefficiencies in manually constructing data sets via gather and integrate tactics is prohibitive when considering the lost value squandered with inefficient use of precious human (and hardware) resources.

The solution to this predicament is to leverage reusable data assets in an enterprise data warehouse. In this framework, analysts provision data for constructing predictive models from a pre-integrated repository of data. This repository must allow for retention of historical, detailed data to be used for constructing training sets to drive model development. In addition, once a model is constructed the repository must support production scoring processes. The tricky part is providing a data warehouse implementation that allows efficient access for SAS programs. Historically, SAS access to relational databases has had a relatively high overhead.

OLD STYLE DEVELOPMENT FRAMEWORK: DISTRIBUTED DATA MARTS
In traditional implementations of SAS with a relational data warehouse the model development and scoring is typically performed outside of the RDBMS using a shadow file system designed for efficient access by SAS programs. These shadow file systems are typically non-relational data mart structures and are usually managed on distributed servers outside of the mainstream data warehouse environment.

There are multiple problems with this approach. First of all, the performance penalty for extracting data from a relational data warehouse is quite high. Even with efficient implementation using SAS® Access and PROC-SQL to prepare and retrieve data from the data warehouse, the performance cost of extracting data through a single threaded interface and (usually) across a network is very high.

The second problem is that the replication of data from a relational repository into a shadow file system is costly in terms of infrastructure. The storage cost of duplicating detailed data to support advanced analytic processes can be significant. Moreover, the processing, extraction, and network cost for duplicating data from a data warehouse environment to the analytic environment is also significant in terms of hardware utilization.

There is also a risk factor, from a security perspective, associated with data replication. The data needed to build sophisticated prediction models must be detailed. Very often, this detailed data is extracted into working sets that are not full secured. Given the increasingly close attention being paid to protecting personally identifying information (e.g., HIPAA in the health care industry), it is an unacceptable vulnerability to allow unmanaged replication of data within an enterprise.

Lastly, and most importantly, the human cost for managing the duplicate data sets is quite high. This effort is usually forced upon the very resources that are highly skilled analytic modelers. The result is that significant human resources are dedicated to data management when they would be better leveraged performing analytic modeling. Imagine if that all the time spent by analytic experts managing data was, instead, spent building more models with greater accuracy. The business value proposition for transforming data management time into analytic modeling time is significant.

NEW STYLE DEVELOPMENT FRAMEWORK: IN-DATABASE PROCESSING

In the new style of development, analytics are performed in-place within the relational data warehouse repository [1]. By integrating SAS processing inside the data warehouse RDBMS it is possible to avoid the costs of data duplication. Rather than creating separate SAS data sets that are manipulated using PROC SORT, PROC RANK, and Data Step Language constructs such as MERGE, SET and other SAS processing capabilities, the idea is for the SAS program to be executed inside the RDBMS using SQL primitives.

An added advantage of the in-database processing approach is that SAS will be able to leverage the parallel processing capability of the data warehouse RDBMS. In the case of a Teradata implementation of in-database processing with SAS, the RDBMS becomes a parallel harness for executing SAS programs. This architecture allows SAS to exploit the massively parallel processing (MPP) execution capability of the Teradata engine across tens, hundreds, or even thousands of Intel Standard High Volume (SHV) Servers.

The in-database processing framework also eliminates a significant amount of the data management overhead associated with the old style approach. However, it is important to note that model development and model scoring are most effectively implemented with "flattened" data sets rather than the normalized (or star schema) forms typically found in a relational data warehouses. To perform the transformations from relational form to an analytic data set structure, tools are required to automate the pivoting of data from a relational form into observation data sets to feed into model building and scoring processes.

Best practice is to create knowledge worker databases to serve as analytic sandbox repositories to facilitate model development and scoring. The sandbox repositories should have working space to create training sets and scoring data sets. Analytic knowledge workers need the ability to create tables and manipulate data directly in the sandbox repositories. The sandboxes provide RDBMS storage and manipulation facilities, and should be on the same platform as the data warehouse to allow for cost effective integration between the production warehouse data and the analytic data sets.

ENABLING TECHNOLOGY IN THE SAS SYSTEM

The enabling technology for the SAS in-database processing emerged from joint R&D work between SAS Institute and Teradata Corporation. The simplest form of in-database processing is the introduction of SAS capability to transform native SAS operations into standard Teradata SQL. For example, SAS 9.2 is able to generate SQL to execute PROC FREQ directly inside the Teradata® RDBMS rather than using the old style method that extracts data out of an RDBMS and then calculates the frequencies on a separate SAS server. The in-database method allows all data access and data aggregation to take place completely inside the database without data movement and data duplication. See Figures 1 and 2 to illustrate the difference between old world and new world SAS execution with in-database processing. Many SAS features have been implemented automatically generated SQL as a result of the joint R&D between SAS Institute and Teradata Corporation.

However, some SAS functions are more advanced than what can be easily expressed in the SQL language. The advanced analytic and mathematical functionality of SAS goes beyond the data manipulation capabilities of ANSI standard SQL. As a result, SAS Institute and Teradata Corporation have worked together to extend the functionality of the Teradata® RDBMS by embedding SAS code directly inside of it. This was achieved by architecting a "data parallel" implementation of SAS functions by embedding SAS libraries within the Teradata® RDBMS using compiled C functions within ANSI compliant User Defined Function (UDF) constructs. These relational extensions can be used to embed C, C++, or Java programs inside of the Teradata® RDBMS.

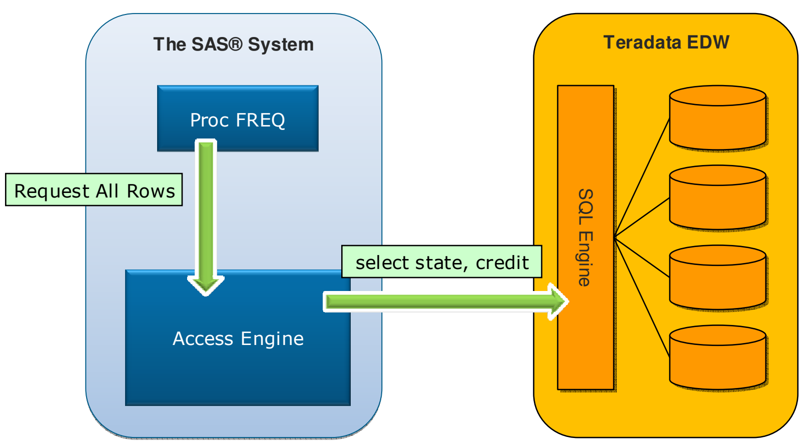

Figure 1: SAS Execution in the Old World

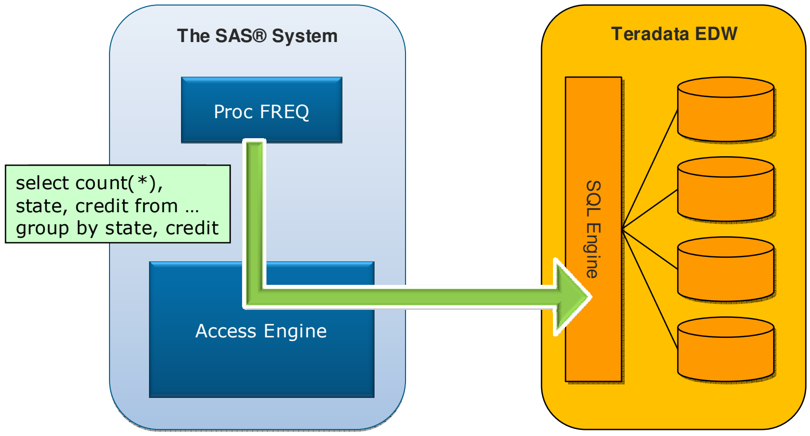

Figure 2: SAS Execution in the New World

SCORING ACCELERATOR™ WITH IN-DATABASE PROCESSING

Another use of the in-database processing technique is the acceleration of scoring algorithms via parallelization inside the Teradata® RDBMS. The SAS® Scoring Accelerator™ enables translation of scoring models created in SAS® Enterprise Miner™ into Teradata specific functions to be executed directly within the Teradata® RDBMS. These functions are translated into optimized C code and are fully parallelized across all servers in the massively parallel processing (MPP) architecture.

In addition to the leverage of the MPP database architecture for high-performance parallel processing, using the SAS® Scoring Accelerator™ avoids data duplication and pulling large data sets across the corporate network. Benchmark tests with and without the SAS® Scoring Accelerator™ demonstrated speed-ups by nearly a factor of fifteen for risk scoring at a large financial services institution.

The SAS® Model Manager 2.2 for SAS 9.2 integrates with the SAS® Scoring Accelerator™ to support model publishing and validation. This integration facilitates seamless use of the work flow designs embodies in SAS® Enterprise Miner™. The use of the SAS® Scoring Accelerator™ for Teradata requires Teradata® RDBMS version V2R6.x, V12, or V13 on a Linux operating system environment. Figure 3 illustrates the approach.

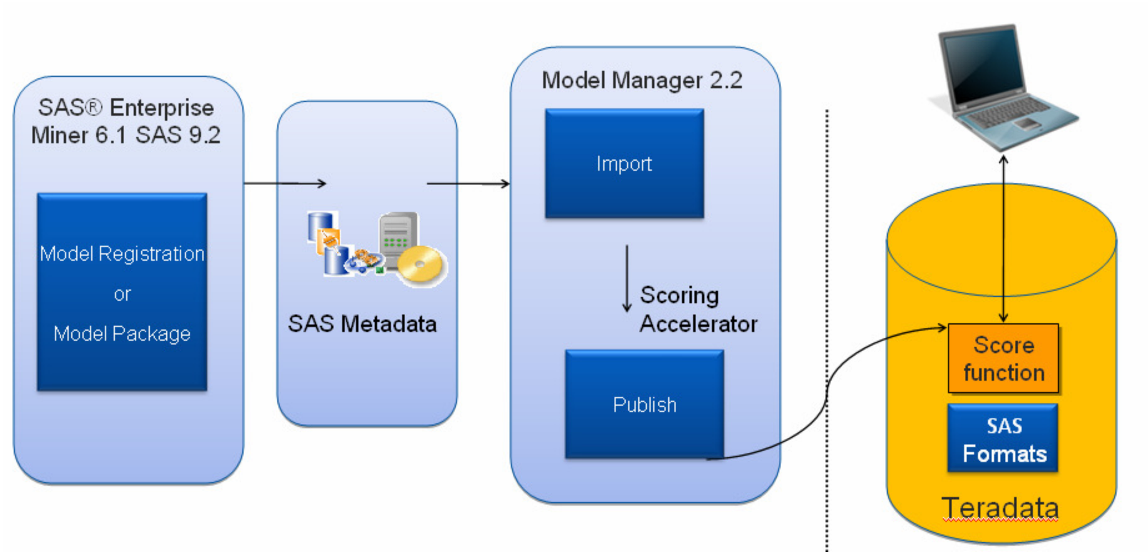

Figure 3: Integration of Model Manager and Scoring Accelerator

ANALYTIC SANDBOX DATABASES

Power users in the SAS environment require the ability to create new tables to support construction of analytic data sets and to import new data sets to facilitate analysis with data that may not yet be populated in a corporate data warehouse. Historically, operational database administrators (DBAs) have been hesitant to allow the access rights to SAS power users that would enable creating tables and loading data from external sources. However, when governed properly, an analytic sandbox can add huge value in facilitating advanced analytics.

It is often necessary to construct intermediate data sets (or to import data sets from external sources) to enable efficiency in deploying advanced analytics. The "old style" DBA rules that prohibit power users from building such data sets has caused a black market in data management whereby information critical to constructing analytic data sets is replicated into servers managed directly by power users. PROC SORT, MERGE Data Steps, and other SAS tools are then used to construct analytic data sets. In many corporate environments, this is the only method available to power users for organizing data in a manner amenable to advanced analytics.

However, this black market approach has a significant downside in terms of human and hardware resource consumption to manage duplicate data sets. A better approach, we have found, is to create "personal" databases for power users that allow self-managed data sets to be created. These personal databases are, ideally, integrated into the corporate data warehouse environment in such a way that production data can be directly joined to the power user tables without pulling data across the corporate network.

The main barrier to implementing analytic sandbox databases for most corporate data warehouse environments is the cultural change required among the DBAs to give up some control (to end users) in the area of data management. There is often a fear that allowing end users to create their own tables will result in chaos. However, this need not be the case. Proper governance techniques can (and should) be put into place to manage storage and computing resources. Advanced workload management techniques such as resource partitioning, rules-based workload monitoring, dynamic workload balancing, and workload prioritization according to end user service level goals are critical components of governing computing resources. It is also important to ensure that the analytic sandbox databases do not become data dumping grounds and do not embark upon duplication of data from the corporate data warehouse. Governance is also required to ensure appropriate controls related to data quality in the analytic sandbox databases. A set of recommended governance techniques for production sandboxing is described in [2].

PERFORMANCE IMPACT OF IN-DATABASE PROCESSING

The performance benefits from in-database processing can vary widely, depending on the nature of the analytics brought into the RDBMS. However, there are some basic techniques that can be used to estimate performance impact of the in-database processing. There are two main benefits that can be quantified: (1) reduced data movement, and (2) increased parallelization.

In traditional execution of analytics using SAS ACCESS® and PROC-SQL there is significant data movement from the data warehouse repository into SAS data sets used for computing results (refer to Figures 1 and 2). This data movement is completely eliminated when using the in-database processing technique. Thus, the part of the workload related to data movement, which is often quite substantial, will go to zero.

Increased parallelization comes from use of Teradata's MPP database architecture. Most SAS execution makes use of symmetric multi-processing (SMP) hardware technology. This type of computing architecture allows many CPUs to share processing workloads with coordination through a shared memory infrastructure. Massively parallel processing (MPP) architecture, on the other hand, will use many SMP servers configured with a high performance switching interconnect to allow scaling to a virtually unlimited number of CPUs. Using the SAS in-database processing features will yield linear scalability to huge numbers of CPUs. Thus, the speed-up factor that can be estimated for in-database processing is the number of CPUs available in the MPP system divided by the number of CPUs in the SMP system. A significant sized MPP system will yield a many times factor in speed-up versus an SMP system. Of course, one must also consider other workload running on both the MPP and SMP servers to understand the percent of available resources that the analytic workload will be given on each platform.

As a point-in-case, consider the execution of a PROC FREQ using a traditional methodology versus execution using in-database processing on an MPP environment. In both cases, the data warehouse repository makes use of the Teradata® RDBMS. However, in the first scenario the Teradata® RDBMS is simply a “file server” dishing out data using the SAS ACCESS® application programming interface. In the second case, the Teradata® RDBMS is using SQL code generated by SAS and all data access, joins, aggregations, etc. are performed inside the database using full data parallel execution.

The benchmark results shown are based on SAS® Enterprise Miner™ and the SAS® Scoring Accelerator™ for Teradata scaling from one to twelve Teradata 5550® servers in an MPP configuration - each with two quad core sockets (for a total of 8 cores per server), 32GB of memory per server, and 100 disk spindles per server. Using Teradata’s virtual processor technology, each server is configured with 25 virtual Access Module Processors (vAMPs) for a total of between 25 and 300 virtual CPUs corresponding to the 1 to 12 physical servers in the scaling tests. All servers in the MPP configuration were executing on SUSE® Linux 9 with the

Teradata 12 software. The testing also used the SAS® 9.2c InDB Release 1.3 on top of Windows® Server 2003 SP2. The SAS software was deployed on a Dell® PowerEdge 6250 configured using four 3.66GHz cores with hyper threading (yielding the equivalent of 8 cores) and 8GB of memory. The SAS Server was connected to the Teradata server using a one gigabit Ethernet (1500 MTU, i.e., no jumbo frames).

The difference in execution throughput for this test was nearly a factor of fifty in terms of throughput: 16K records per second versus 750K records per second. To understand the impact of the data movement versus the impact of MPP parallelization we also ran a test that used the traditional SAS ACCESS® method - but over a dedicated high-speed interconnect. With this approach we were able to get a throughput of 77K records per second. Thus, for this scenario, the benefit associated with eliminating data movement time was approximately a factor of five and the performance benefit of the additional CPUs from the MPP processing was approximately a factor of ten (for a total of a nearly 50 times speed up). This makes sense when we consider the relative processing capabilities of the SAS SMP server versus the Teradata MPP server. See Figure 4 for a graphical representation of the performance results.

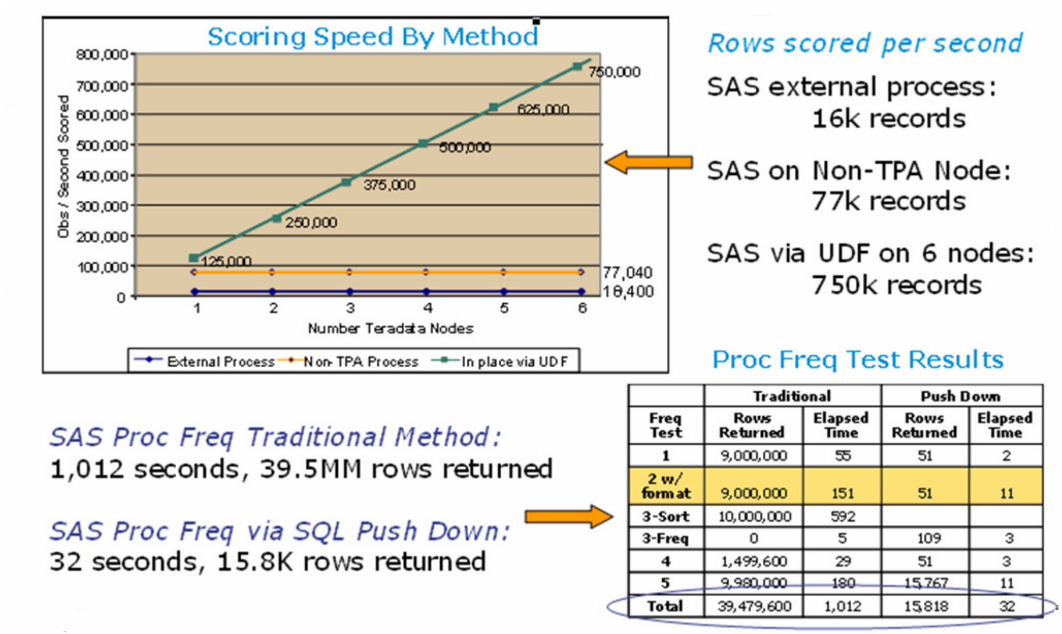

	Traditional		Push Down	
Freq Test	Rows Returned	Elapsed Time	Rows Returned	Elapsed Time
1	9,000,000	55	51	2
2 w/ format	9,000,000	151	51	11
3-Sort	10,000,000	592		
3-Freq	0	5	109	3
4	1,499,600	29	51	3
5	9,980,000	180	15,767	11
Total	39,479,600	1,012	15,818	32

Figure 4: Benchmarking Results

CONCLUSIONS

A new era of high-performance for advanced analytics is upon us. By combining the power of SAS with massively parallel processing enabled RDBMS technology, faster and cheaper analytics can be achieved. The key value proposition, besides cost savings in human and infrastructure resources, is that more advanced models can be developed in shorter time periods with enhanced data security. With quicker turnaround for model construction, more iterations of refinement are possible with a result of more accurate model deployments. The value of the incremental lift delivered by more predictive models resulting from efficient turnaround time using in-database processing is huge for organizations that have begun to exploit this capability.

CONTACT INFORMATION

Stephen Brobst
Chief Technology Officer, Teradata Corporation
17095 Via Del Campo
San Diego, CA 92127
(617) 422-0800
stephen.brobst@teradata.com

Keith Collins
Senior VP & Chief Technology Officer, SAS Institute
100 SAS Campus Drive
Cary, NC 27513
(919) 677-4444
Keith.Collins@SAS.com

Paul Kent
Vice President, R&D, Platform Research and Development, SAS Institute
100 SAS Campus Drive, MC R4452
Cary, NC 27513
(919) 677-8000
Paul.Kent@SAS.com

Mike Watzke
Teradata Architect, Teradata Corporation
2719 Jasmine Drive
Fitchburg, Wisconsin 53711
(608) 277-1292
Michael.Watzke@Teradata.com

REFERENCES

[1] Collins, K. and S. Van Valkenburgh. Use SAS to Get More Power Out of Your Database," SASCOM Magazine, First Quarter, 2008. pp. 4-5.

[2] Brobst, S., M. McIntire, and E. Rado. "Agile Data Warehousing with Production Sandboxing," Business Intelligence Journal, Winter, 2008. pp. 13-22.