Paper 11769-2016

Analytics and Data Management in a Box: Dramatically Increase Performance

Greg Otto, Paul Segal, and Tho Nguyen, Teradata Corporation

ABSTRACT

Organizations are collecting more structured and unstructured data than ever before, and it is presenting great opportunities and challenges to analyze all of that complex data. In this volatile and competitive economy, there has never been a bigger need for proactive and agile strategies to overcome these challenges by applying the analytics directly to the data rather than shuffling data around. Teradata and SAS® have joined forces to revolutionize your business by providing enterprise analytics in a harmonious data management platform to deliver critical strategic insights by applying advanced analytics "inside" the database or data warehouse where the vast volume of data is fed and resides. This paper highlights how Teradata and SAS strategically integrate in-memory, in-database, and the Teradata relational database management system (RDBMS) in a box, giving IT departments a simplified footprint and reporting infrastructure, as well as lower total cost of ownership (TCO), while offering SAS users increased analytic performance by eliminating the shuffling of data over external network connections.

INTRODUCTION – THE TERADATA UNIFIED DATA ARCHITECTURE

The Teradata Unified Data Architecture (UDA) brings diverse data sources and processing capabilities into an environment where data is stored in a "best fit" data management system, and then processed in the analytical engine that is best suited to the work at hand.

Figure 1 illustrates the components of an analytical environment based on the Teradata UDA, with all components interconnected by the high speed Teradata BYNET over InfiniBand:

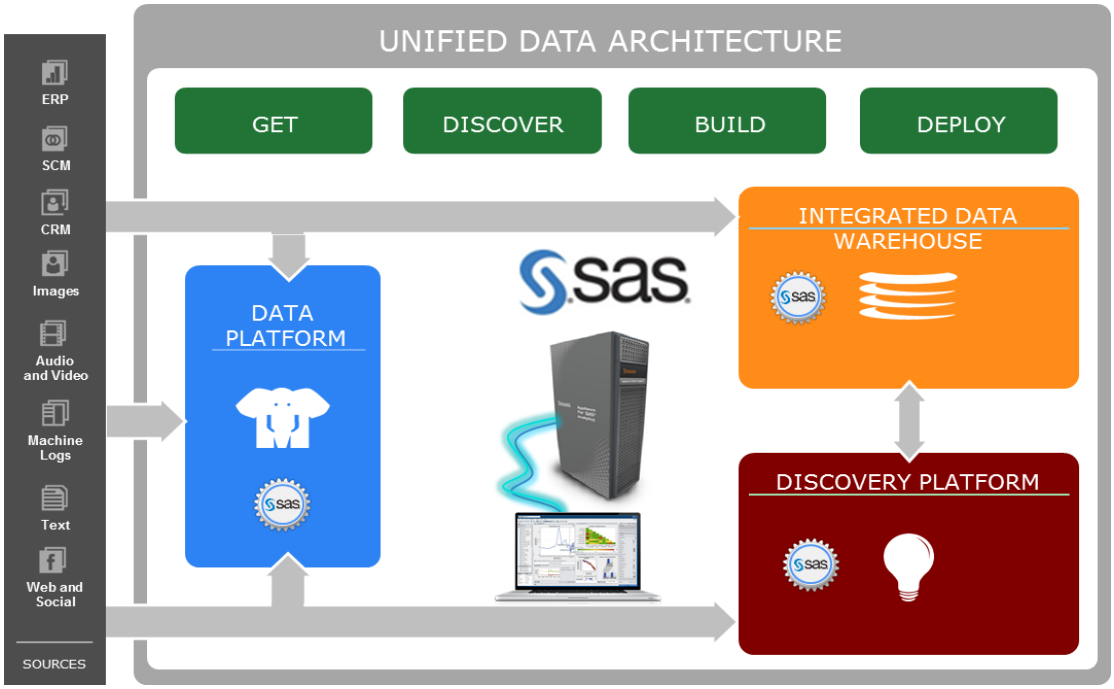

Figure 1. Components of the Teradata Unified Data Architecture (UDA)

There are three primary design objectives when architecting an analytical environment based on the Teradata UDA:

1. Store the data in a system that is best suited to the characteristics of the data (volume, complexity/structure, frequency of access), along with the most critical and frequent processing requirements. Best fit data solutions provide the analytical results that are needed at the appropriate scale and with the best balance of price/performance.
2. Move the processing to the data. Analyze the data in place to minimize data movement and replication. Minimizing movement improves performance, reduces data management overhead, and improves data governance.
3. When the data does need to move from data store to the analytical processing engine, make the movement seamless (transparent to the end user) and fast. Fast, transparent data movement lets the users focus on delivering value from data vs. gluing data together and waiting for it to show up.

Key technologies that enable SAS integration in the Teradata UDA include:

- Teradata DBMS – Massively Parallel Processing (MPP) relational database supporting complex SQL processing and high concurrency analytic workloads.
- Other data stores (including Hadoop) – Systems that support specialized analytical functions, or provide more cost-effective storage for massive volumes of data.
- Teradata Appliance for SAS In-Memory Model 750 – A platform for hosting the entire suite of SAS In-memory analytics products, pre-staged ready to install SAS and fully integrated with the Teradata UDA.
- SAS Managed Server – Orchestrates In-Database and In-Memory operations on the data where it sits in the data stores. SAS applications and programs hosted on SAS Application Servers benefit from the high-speed interconnect between all UDA components when data movement is required.
- Teradata BYNET V5 interconnect – Based on InfiniBand technology, BYNET V5 provides load balanced and highly available high-speed data transfer between all processing nodes and clusters in the UDA.
- Single system operational view – Teradata integrated server management and Teradata Viewpoint monitoring for all components in the UDA.
- Teradata® QueryGrid™ -- Provides transparent, high performance data integration across multiple data sources on the fly.

For each of these key technologies this paper will describe how the SAS and Teradata jointly provide a harmonious data management platform that delivers critical strategic insights by applying advanced analytics “inside” the database or data warehouse where the data resides.

TERADATA DATA WAREHOUSE

TERADATA DATA WAREHOUSE APPLIANCE

Teradata Data Warehouse Appliances are fully-integrated systems, purpose-built for data warehousing. They feature the Teradata Database, SUSE® Linux operating system, enterprise-class storage and BYNET interconnect all pre-installed and ready to run.

Teradata delivers a family of platforms to support a wide range of data management and decision support workloads from high volume data repositories, to departmental and application specific applications, and high concurrency / high performance enterprise data warehouse applications.

The Teradata Database is a proven, robust data store with a mature parallel-aware optimizer, advanced compression, sophisticated storage and partitioning strategies, and advanced workload management. Hundreds of SAS customers interact with Teradata daily to extract valuable insights using SAS advanced analytics.

SAS/ACCESS TO TERADATA

SAS/ACCESS to Teradata on a SAS Managed Server delivers comprehensive and efficient integration with the Teradata database, with capabilities that include:

- Implicit SQL pass through – Data access by SAS procedures is translated to Teradata SQL, with push down of filters and data transformations including SAS formats
- SQL push down for In-Database enabled procedures
- Bulk load, update, and export to support large data integration tasks – Integrated with Teradata Parallel Transporter (TPT)
- SAS SQL procedure:
    - PROC SQL Implicit pass through – SAS SQL is converted to Teradata SQL syntax
    - PROC SQL Explicit pass through – Supports sending Teradata SQL directly to the database for specially optimized processing

IN-DATABASE PROCESSING

SAS and Teradata have worked together to enhance many SAS procedures and solutions to leverage the power of Teradata's parallel SQL for In-database processing. In many cases In-database processing is automatic and transparent to the user.

All of the steps in the analytic life cycle can benefit from In-database processing:

Data Exploration

- SAS/ACCESS to Teradata
- Base SAS procedures

Data Preparation

- SAS DS2 Code Accelerator for Teradata
- SAS Data Quality Accelerator for Teradata

Model Development

- SAS Analytics Accelerator for Teradata
- SAS Enterprise Miner
- SAS/STAT

Model Deployment

- SAS Model Manager

Model Deployment

- SAS Model Manager
- SAS Scoring Accelerator for Teradata

IN-DATABASE ENABLED SAS BASE PROCEDURES

Built on the foundation of SAS/ACCESS to Teradata, these Base SAS Procedures use Teradata SQL for In-Database processing of all, or in some cases partial results. See the SAS Reference for each procedure for details on how the procedure is In-database enabled

- PROC APPEND
- PROC CONTENTS

- PROC COPY
- PROC DATASETS
- PROC DELETE
- PROC FORMAT
- PROC FREQ
- PROC MEANS
- PROC PRINT
- PROC RANK
- PROC REPORT
- PROC SORT
- PROC SQL
- PROC SUMMARY
- PROC TABULATE

THE SAS EMBEDDED PROCESS: FOUNDATION TECHNOLOGY FOR IN-DATABASE PROCESSING

The SAS® Embedded Process (EP) on Teradata is a SAS-developed table function server that processes data using the SAS DS2 Language in cooperation with Teradata database units of parallelism (Teradata AMPs). The SAS EP is fully integrated with Teradata for database security, workload management, process priority, resource controls, accounting, exception handling and query logging.

The SAS Embedded Process:

- Runs SAS DS2 language natively under the SAS Threaded Kernel (TK) runtime environment
- Runs on all Teradata nodes
- Is fully multi-threaded (number of EP threads = number of Teradata AMPs)
- Uses high speed in-memory buffering between the AMPs and EP threads

The block diagram of the SAS EP on Teradata in figure 2 illustrates how the EP is tightly integrated with the Teradata database system:

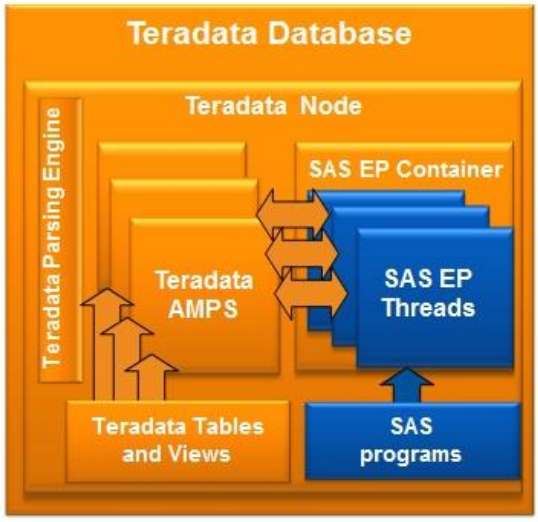

Figure 2. SAS Embedded Process on Teradata

When a Teradata view is used as input to the SAS EP, the SQL in the Teradata view is evaluated and materialized in Teradata spool before it is passed to the EP. That view could easily include multiple data sources combined by Teradata QueryGrid, for seamless data integration across multiple data sources, on the fly, with In-Database processing in Teradata. Teradata QueryGrid is explained in more detail in a subsequent section of this paper.

IN-DATABASE ACCELERATORS FOR TERADATA

Beyond using the basic Teradata SQL language for In-database processing, SAS and Teradata have jointly developed advanced In-database accelerators using purpose-built integration technologies like the SAS Embedded Process.

SAS Scoring Accelerator for Teradata

The SAS Scoring Accelerator for Teradata turns models created by multiple SAS analytical tools into SAS DS2 language for scoring inside the database using the SAS EP. Models are "published" to Teradata using SAS macros provided with the Scoring Accelerator, or they can be published using SAS Model Manager. SAS Model Manager also maintains metadata about the models to assist with organizing and deploying the models to production processes.

Models generated by these SAS products can be published to Teradata for scoring:

SAS/STAT

SAS/EM

Visual Statistics

SAS DS2 Code Accelerator for Teradata

The SAS DS2 language can be used to express data transformation and computational operations that used to require multiple SAS program steps involving SAS procedures and Data Steps linked together by data sets stored in SASWORK. Now these operations can be converted to run in fewer steps (in some cases just a single step), reducing data movement. Response time is improved by taking advantage of massively scalable processing inside the database.

The DS2 procedure automatically sends the program to the database when it meets requirements for the code accelerator:

- The SAS Embedded Process is installed on the parallel data source
- The DS2 program contains a "thread" code block
- Code accelerator is licensed, and a program option requests acceleration (DS2ACCEL=YES)

In one case study, a computationally intensive scoring program used by a retailer was converted to DS2. The new process reduced processing steps in SAS from six down to one PROC DS2 step which ran inside the database. Time taken to process one billion transactions was reduced by 20X compared to the original process, with elapsed time reduced from minutes to seconds (see Secosky, 2014).

Results like this are achieved by moving the processing to the data and by leveraging parallel execution threads in the database. The SAS DS2 Code Accelerator leverages parallel database processing for partitioning and sorting of BY groups prior to parallel DS2 thread execution in the SAS EP, as illustrated in figure 3:

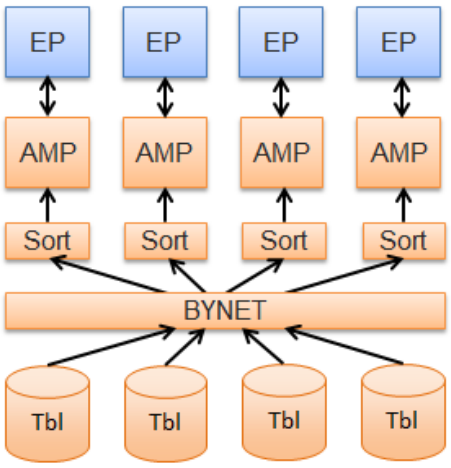

Figure 3. Teradata partitions and sorts BY groups for each DS2 threads in the SAS EP

SAS In-Database Decision Management for Teradata is an example of a solution that extends In-Database capabilities using DS2 running in the Teradata database. The Decision Management solution uses proven methods to prepare and cleanse data, develop and execute analytical models, and score those models according to the conditions under which the models and business rules are valid, all resulting in more accurate, consistent, data-driven decisions which lead to effective actions.

SAS Data Quality Accelerator for Teradata

Unique to Teradata, the SAS Data Quality Accelerator for Teradata cleanses and transforms data in place in the database using a sophisticated rules database called the Quality Knowledge Base (QKB). With the QKB, SAS EP, and a set of Teradata stored procedures installed on Teradata SAS users can apply the following algorithms to data in place:

- Match code
- Parsing/Casing
- Gender/Pattern/Identification analysis
- Standardization

SAS Analytics Accelerator for Teradata

Unique to Teradata, the SAS Analytics Accelerator contains specialized vendor defined functions in Teradata that enable In-Database processing for a collection of modeling and data mining algorithms.

These SAS procedures run In-Database:

- SAS/STAT Procedures
    - PROC CANCORR
    - PROC CORR
    - PROC FACTOR
    - PROC PRINCOMP
    - PROC REG
    - PROC SCORE
    - PROC TIMESERIES
    - PROC VARCLUS

- SAS/EM Procedures and Operations:
    - PROC DMDB
    - PROC DMINE
    - PROC DMREG (Logistic Regression)
    - SAS/EM nodes for Input, Sample, Partition, Filter, Merge, Expand
- PROC SCORE will score data In-Database using coefficients from these SAS procedures:
    - PROC ACECLUS
    - PROC CALIS
    - PROC CANDISC
    - PROC DISCRIM
    - PROC FACTOR
    - PROC PRINCOMP
    - PROC TCALIS
    - PROC VARCLUS
    - PROC ORTHOREG
    - PROC QUANTREG
    - PROC REG
    - PROC ROBUSTREG

IN-MEMORY INTEGRATION

The SAS EP on Teradata also the enabler for parallel data transfer from the Teradata Data Warehouse to SAS In-memory Appliance Model 750. Teradata BYNET on InfiniBand provides high speed connections from Teradata to the SAS In-memory processes. Teradata BYNET also provides the interconnect among the SAS In-memory worker nodes in the 750. Figure 4 illustrates how the SAS EP copies data on demand from Teradata to the SAS In-memory infrastructure:

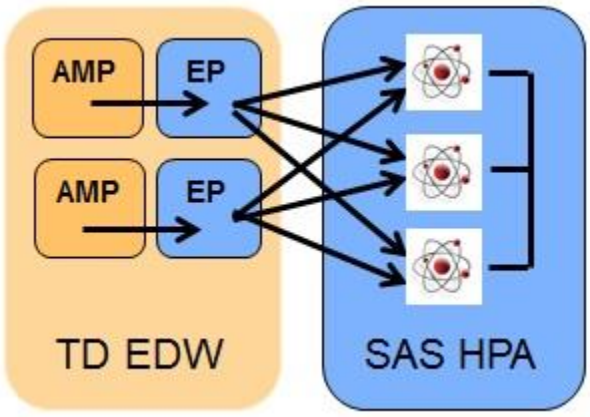

Figure 4. SAS EP on Teradata feeding SAS HPA

The Teradata Appliance for SAS In-Memory will be discussed in more detail in a later section of this paper.

TERADATA APPLIANCE FOR HADOOP

The Teradata Appliance for Hadoop combines optimized hardware, high-speed connectors, software usability enhancements, proactive systems monitoring, intuitive management portals, continuous availability, and linear scalability with Teradata's world-class service and support to provide an integrated system to support Hortonworks or Cloudera Hadoop distributions.

SAS/ACCESS TO HADOOP

SAS Managed Servers with SAS/ACCESS to Hadoop have direct access to Hadoop data and processing capabilities. SAS Input/Output operations are converted automatically into native Hadoop access methods.

IN-DATABASE ACCELERATORS FOR HADOOP

Most of the same SAS In-Database processing capabilities that are available on Teradata are also available on Hadoop, many of them supported by the SAS Embedded Process for Hadoop, including:

- In-Database enabled Base SAS procedures
- Scoring Accelerator
- DS2 Code Accelerator
- High performance parallel data loading from a Teradata Hadoop appliance directly into SAS In-Memory processes on the Teradata Appliance for SAS Model 750

SAS IN-MEMORY ON THE TERADATA APPLIANCE FOR SAS MODEL 750

The Teradata Appliance for SAS High-Performance Analytics Model 750 is fully UDA-integrated platform for hosting SAS In-memory products. SAS High-Performance Analytics products deliver advanced analytic and exploration capabilities at massive scale using In-memory processing.

SAS In-Memory utilizes a parallel MPP model, working on the principle of loading and distributing all of the data involved in a particular analytic dataset into the available (combined) RAM of multiple servers in a cluster, and providing complete end to end processing of that data directly from RAM. Exploration and modeling are accelerated because the CPU cores from multiple servers process the analytic jobs in parallel directly against data in memory.

Analytical processes that used to take hours with traditional single-server SAS execution can be reduced to run in minutes, or potentially even seconds using SAS In-Memory analytics. For companies that have very high value analytic processing, this time savings can drive a reduction in cost by supporting many concurrent users while providing new high value insights that would have otherwise been impossible.

To support SAS In-Memory on Teradata, Teradata has developed an appliance that enables SAS In-Memory analytics against any of the Teradata product family, including Teradata and Hadoop.

SAS solutions that utilize the SAS High-Performance Analytics Infrastructure include:

- SAS High-Performance Analytics (HPA) Products
  - SAS® High-Performance Statistics
  - SAS® High-Performance Data Mining
  - SAS® High-Performance Text Mining
  - SAS® High-Performance Forecasting
  - SAS® High-Performance Econometrics
  - SAS® High-Performance Optimization
- SAS Visual Analytics (VA)
- SAS Visual Statistics (VS)

- SAS In-Memory Statistics (IMSTAT)
- SAS High Performance Anti-Money Laundering (HP-AML)

SAS In-memory worker nodes in the Mode 750 appliance support memory configuration options of 256GB, 512GB, or 768GB per node. The CPU to memory ratio is determined by analyzing the total in-memory data footprint, types of processing, and number of users. For compute-intensive workloads the system will perform best with more worker nodes having less memory per node (256GB), but data-intensive workloads can use fewer nodes with more memory per node (512GB or 768GB).

Worker nodes run the SLES11-SP3 operating system and take advantage of the latest Intel Xeon processor technology with options for 16 cores (32 CPUs with hyper-threading) or 24 cores (48 CPUs with hyper-threading) per node.

Figure 5 shows additional specifications of the Model 750 appliance for SAS In-memory analytics:

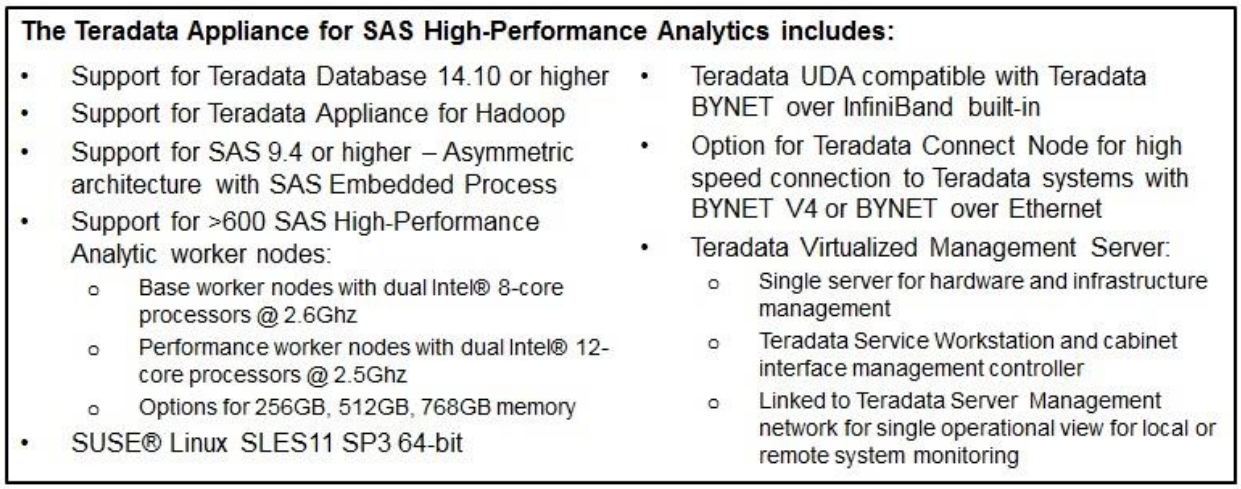

Figure 5. Explore, discover, report and deliver all of your data with the Teradata Appliance for SAS

SAS Managed Servers can also be added to the Model 750 Appliance, providing an opportunity to optimize floor space and data center operations by combining all SAS infrastructure together in a single cabinet.

SAS® MANAGED SERVER FROM TERADATA

INTEGRATED ANALYTICS AND DATA MANAGEMENT IN A BOX

The SAS Managed Server from Teradata enables SAS analytics to run in the Teradata UDA for data exploration, data preparation, model development and model deployment. It extends Teradata systems to provide end-to-end advanced analytics in a single, compact system. The SAS Managed Server delivers the improved performance and enhanced enterprise data management that are required by today's innovative companies.

Designed to Deliver Efficiency and Performance

The SAS Managed Server from Teradata is one or more nodes configured specifically for running traditional SAS Analytics and can reside in any of the Teradata® Workload-Specific Platform Family, as shown in Figure 6:

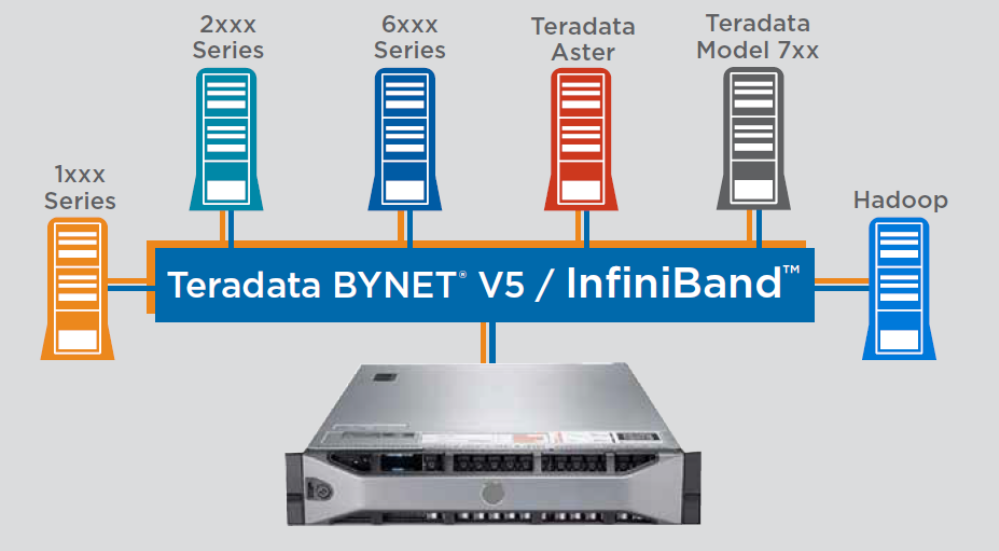

Figure 6. Managed Server can reside in any members of the Teradata Platform Family to run SAS analytics

These servers combine dual Intel® eight core processors, 24-1.2TB data drives, SUSE® Linux operating system, with 256/512/758GB of RAM options, and enterprise-class Mellanox® InfiniBand™ networking infrastructure into a power-efficient system. The servers connect directly to the Teradata Platform Family via InfiniBand providing high-speed local network access.

The SAS Managed Server is enterprise ready:

- Centralized Teradata Viewpoint providing a centralized monitoring point for the entire Teradata ecosystem
- Easy administration, via centralized systems management capabilities

The SAS Managed Server is optimized for performance:

- Fast 8-core Intel processors for multi-user support and for parallel analytic calculations
- Large local file system, configured to support high I/O workload required by traditional SAS processing
- Connected via InfiniBand supporting database transfers and analytic data flow

Long-term benefits of using the SAS Managed Server to support SAS applications include:

- A data management environment supporting enterprise analytics; source your analytic and data warehousing platforms in one box
- Leverage the same SAS data sets and SAS programs that you use today
- High performance platform supporting new SAS applications
- Industry leading support from a single vendor, there's no need for third-party coordination to build out new SAS servers
- Tested and proven support processes, with a collaborative support model spanning Teradata and SAS organizations

TERADATA BYNET V5

Teradata BYNET® V5 is built on InfiniBand technology. Perfected over 30+ years of massively parallel processing experience, BYNET V5 provides low-latency messaging capabilities for maximum data access. This is accomplished by leveraging the inherent scalability and integrity of InfiniBand to load-balance multiple fabrics, seamlessly handling failover in the event of an interconnect failure.

InfiniBand technology delivers performance and resiliency from Mellanox-supplied InfiniBand switches, adapters and cables which are recognized as industry-leading building blocks for high-quality, fully interoperable enterprise class switching systems.

SYSTEMS MANAGEMENT WITH TERADATA VIEWPOINT

Teradata Viewpoint is an advanced web-based system management portal for performance and health management across the Teradata UDA components. It provides a consistent interface via configurable portlets which are easy to use and which allow the user to customize their own systems management dashboard.

INTEGRATED DATA MANAGEMENT AND ANALYTICS ENVIRONMENT IN A BOX

Bringing all of these components together into a single box to support integrated data management and analytical processing, Teradata can provide the entire analytical environment in a compact footprint for improved operational efficiency, data governance, and high performance. All of the components of the system shown in figure 7 are linked with the high speed BYNET V5 interconnect:

Figure 7. SAS Managed Servers, Data Store, and SAS In-memory cluster "In a Box"

TERADATA QUERYGRID

Leverage analytic resources, reduce data movement

Teradata QueryGrid delivers transparent data access and local processing across systems within the Unified Data Architecture. It takes advantage of each system's specialized engines by sending parts of queries and data to the other platforms for execution and processing. Data placement and movement are optimized and intelligent query processing is automated for best overall results. This allows organizations to reap the most value from their unified architecture by easily and transparently pushing analytics to the data or moving the data to the right analytic platform, all without IT intervention. Users have more control with the ability to send their query to only one central system (Teradata) as the data hub.

A user will submit a single request to the Teradata Database and receive a completed result back, but some of the work may be performed on one of the other platforms. Other systems get involved because they have the needed data or because they can perform a function that the Teradata Database cannot.

QueryGrid abstracts the complexities of the data location and retrieval mechanism away from the user. All the data administrator needs to know is the name of the data and which FOREIGN SERVER references

that data. Data administrators can add views that make the foreign data access completely transparent to end users.

Teradata QueryGrid:

- Minimizes data movement and data duplication.
- Processes data on the system where it resides, for example a WHERE clause will be pushed down to the foreign system and only the rows required will be surfaced to the coordinating system.
- Optimizes work distribution through "push-down" processing.
- Transparently automate analytic processing between systems.
- Can be combined with SAS® HPA to deliver transformative in-memory capabilities, providing access to all of the data for:
    - Complex modeling and analytics
    - Data visualization and exploration

Figure 8 illustrates how Teradata QueryGrid supports seamless data integration for users of SAS Visual Analytics and SAS Visual Statistics on the Model 750 appliance for SAS High-Performance Analytics:

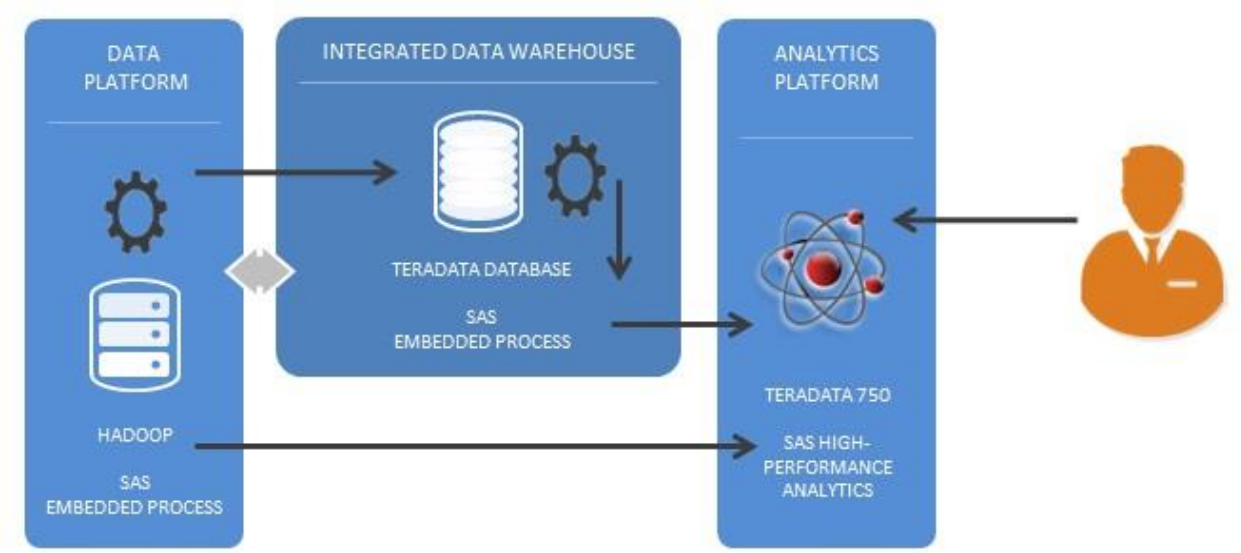

Figure 8. Teradata QueryGrid supports seamless data integration for SAS In-Memory Analytics

STEPS TO CREATING AN INTEGRATED ANAYTICS SOLUTION WITH QUERYGRID

1. Choose the Teradata Database as the primary source system (hub)
2. Define FOREIGN SERVER connections to remote data sources, including Hadoop
3. Create Teradata database views to the tables on the FOREIGN SERVER
4. Run SAS In-database analytics procedures on Teradata views:
    - Views can reference data in Teradata and/or Hadoop on the fly
    - Minimizes data movement with “push-down” processing to the data sources and with In-Database processing in Teradata
5. Run SAS High Performance Analytics or SAS Visual Analytics on 750 appliance:
    - Provides seamless access to remote sources (Hadoop, others available)
    - Utilizes the SAS Embedded Processes on Teradata for fast parallel loading into memory

- Data movement and consolidation is transparent to SAS users that access the Teradata views

The following sample Teradata SQL creates a FOREIGN SERVER reference to a Hadoop cluster, creates a view that pushes a date filter to Hadoop, and then creates a combined Teradata view that can be used directly by SAS users to analyze history in Teradata data combined on the fly with new transactions in Hadoop:

```
/* Foreign server details (Hadoop in this case) */
CREATE FOREIGN SERVER hadoopserver USING
   server('hadoop.teradata.com')
   hosttype('hadoop')
   port('9083')
   hiveport('10000')
   username('hive')
DO IMPORT WITH SYSLIB.load_from_hcatalog ;

/* View to data in Hadoop, with date filter pushed to Hadoop */
CREATE VIEW hadooop_new_web_events AS
   SELECT user_id, event_id, datestamp, event_details
   FROM hadoop_web_events@hadoopserver
   WHERE datestamp BETWEEN '2016-03-12' AND '2016-03-18' ;

/* View to data in Teradata combined with data from Hadoop */
CREATE VIEW all_web_events AS
   SELECT user_id, event_id, datestamp, event_details
   FROM td_web_event_history
      UNION ALL
   SELECT user_id, event_id, datestamp, event_details
   FROM hadoop_new_web_events ;

/* SAS users can now access the combined data using "all_web_events" */
SELECT * FROM all_web_events ;
```

Teradata QueryGrid supports connections to multiple data sources including Hadoop (Hortonworks and Cloudera), Teradata, Teradata Aster, Oracle, MongoDB and others (contact your Teradata representative for availability).

CONCLUSION

The SAS and Teradata partnership has been revolutionizing the way organizations make critical business decisions with advance analytics since it began in 2007. Helping over 200 joint customers to create integrated, optimized and agile analytics environments, the partnership continues to deliver innovative, best-in-class solutions to make solving your Big Data Analytics challenges easier by:

Embedding SAS within the Teradata Unified Data Architecture, bringing the analytical processing to the data where it resides

- Improving data governance, availability and latency by minimizing permanent copying of data
- Streamlining modeling from days/weeks to minutes/hours with the power of parallel processing
- Maximizing the value of your existing investments while controlling your costs by combining data management and analytical processing in a single complex

Data analysts can focus on developing insights leading to effective actions, rather than worrying over and waiting for with complex data management tasks, supported by technologies like Teradata QueryGrid and the high-speed interconnect within the Teradata UDA.

REFERENCES

Secosky, Jason. 2014. “Parallel Data Preparation with the DS2 Programming Language.” Proceedings of SAS Global Forum 2014. Available at http://support.sas.com/resources/papers/proceedings14/SAS329-2014.pdf

RECOMMENDED READING

SAS and Teradata Partnership
http://www.teradata.com/SAS

Teradata Unified Data Architecture™ (UDA)
http://www.teradata.com/products-and-services/unified-data-architecture

Teradata Data Warehouse Appliance 2800 and the Teradata Database
http://assets.teradata.com/resourceCenter/downloads/Datasheets/EB8381.pdf

Teradata Appliance for Hadoop
http://www.teradata.com/products-and-services/appliance-for-hadoop

SAS Managed Server
http://assets.teradata.com/resourceCenter/downloads/Brochures/EB9117.pdf

Teradata Appliance for SAS® High-Performance Analytics Model 750
http://www.teradata.com/product.aspx?id=18660

Teradata Viewpoint
http://www.teradata.com/products-and-services/viewpoint

BYNET V5 and QueryGrid
http://www.teradatamagazine.com/v14n02/Tech2Tech/Harmonious-Orchestration/
http://www.teradata.com/Teradata-QueryGrid

SAS® Visual Analytics for Teradata and SAS® Visual Statistics for Teradata
http://assets.teradata.com/resourceCenter/downloads/Brochures/EB8742.pdf

SAS and Teradata Advanced Analytics Advantage Program for Teradata
http://www.teradata.com/Resources/Brochures/SAS-and-Teradata-Analytic-Advantage-Program

SAS and Teradata Advanced Analytics Advantage Program for Hadoop
http://www.teradata.com/partners/SAS/SAS-and-Teradata-Advanced-Analytics-Program

CONTACT INFORMATION

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

Greg Otto
Greg.Otto@teradata.com

Paul Segal
Paul.Segal@teradata.com

Tho Nguyen
Tho.Nguyen@teradata.com