

How Teradata Makes the Smart Grid Smarter

The Active Smart Grid Analytics Platform

By: Casey Higgins,
Senior Solutions Architect,
Teradata Corporation

TERADATA[®]

How Teradata Makes the Smart Grid Smarter

Table of Contents

<i>Executive Overview</i>	2
<i>Introduction</i>	3
<i>Drivers of the Active Smart Grid Analytics Platform from Teradata</i>	3
<i>The Teradata Active Smart Grid Analytics Platform Enablers</i>	4
<i>Teradata Platform</i>	6
<i>Teradata Tools and Utilities</i>	8
<i>Teradata High Availability Components</i>	13

Executive Overview

This white paper describes how Teradata leverages 30+ years of data warehousing experience to provide the most powerful solution to utility companies that are transforming their business with Smart Grid technology. Teradata is the only company that can provide the information technology infrastructure and services to support the requirements for a true Active Smart Grid Analytics (ASA) platform.

How Teradata Makes the Smart Grid Smarter

Introduction

The advent of smart grid and smart meter technology is changing the utility industry's business model. As a result, utility companies face new requirements to provide advanced analytical and decision-making capabilities. The need to integrate data from many systems such as the Meter Data Management System, the Customer Information System (CIS), the Outage Management System and others, increases the value of the grid while supporting data delivery to a wider range of applications and data consumers. Utility companies will need to provide this information for many business needs such as the corporate portal, outage management, dynamic pricing, and meter-to-cash analysis.

Teradata Corporation leverages 30 years of experience in data warehousing to enable the most powerful solution for utility companies to help them transform to accommodate this new business model. Teradata makes available the infrastructure to support the requirements of the Active Smart Grid Analytics (ASA) platform as demand increases for growing data volumes, new business users, and new data subject areas. The active aspect of the ASA platform means that the requirements must support the ability to quickly move data into the analytical environment to perform analytics and drive decision making based on detailed granular data throughout the utility. The following discussion provides an overview of an active smart grid analytics platform from Teradata.

Drivers of the Active Smart Grid Analytics Platform from Teradata

Like the telecommunications, retail, and airline industries, the utilities industry has the ability to drive business processes based on detailed customer behavior. These drivers include being able to:

- > Determine the true effectiveness of energy efficiency programs through the analysis of AMI data and provide more accurate credits.
- > Identify financial issues in the meter-to-cash process, such as incorrect re-bill processing or improper reading adjustments.
- > More accurately identify transformer issues and load increases based on individual customer usage patterns, rather than simply on peak day estimates.
- > Improve customer communications and program enrollment by targeting customers across multiple dynamic segments, rather than a fixed set of predetermined segments.
- > Use Geospatial capabilities to visualize customer usage and event trends.

Many other examples of using detailed meter data exist. The critical aspect is providing the utility company with the ability to support these many applications of the data without creating an unmanageable clutter of data repositories.

A second set of drivers for an ASA platform focuses on architectural necessities to support the many data consumers who need the valuable information. Four necessary drivers must be considered within an ASA platform:

- > Provide the performance and scalability to support:
 - Large data volumes of interval and event data from smart meters.
 - An increasing number of users accessing the data in different ways, such as very tactical Web services and deep historic trend analysis.
 - Improved data latency from the time events have occurred to the time decisions are driven by the data.
- > Manage costs by simplifying the development and administration of the integration of differing data sources into a common platform, while maintaining data granularity and consistency.
- > Ensure data security through the proper control and auditing of data across the environment to support the different user requirements, from the external customer to the financial analyst.
- > Ensure data accuracy through the use of standardized and agreed upon business rules applied to the data.

How Teradata Makes the Smart Grid Smarter

The Teradata Active Smart Grid Analytics Platform Enablers

Active Smart Grid Analytics Platform Architecture

The architectural approach to providing an information platform for smart grid analytics focuses on centrally integrating data to support the varying requirements of a utility. As mentioned above, the platform must support both business and infrastructure requirements such as security, performance, scalability, and easy enterprise integration. The smart grid platform incorporates the same Teradata architecture that many industries, such as

telecommunications, financial services, retail, travel, and transportation, are using for Active Data Warehousing™.

Figure 1 illustrates the flow of data from the meter to the ASA Platform for a mature data environment. Teradata provides the ability to start small and add new data sources as new business requirements need to be fulfilled. As the data are sent by the meter to the MDMS and head-end via the necessary communication channels, it is gathered and loaded into a staging area within the data warehouse. The most common method for gathering the data is to perform change data capture (CDC) on the MDMS. Data are also

loaded from other sources to enrich the value of the meter data interval and event data. The most critical source is the CIS system, which provides critical customer master data and customer financial transaction data. Once the data are loaded into the staging area, they are moved to the core layer using any one of the ETL tools with in-database capabilities. Transformations of the data and referential integrity are performed while the data are moved to the core layer. Data users access the data through different semantic views, depending on the need of the users.

Data security is maintained by limiting data consumers' access through the

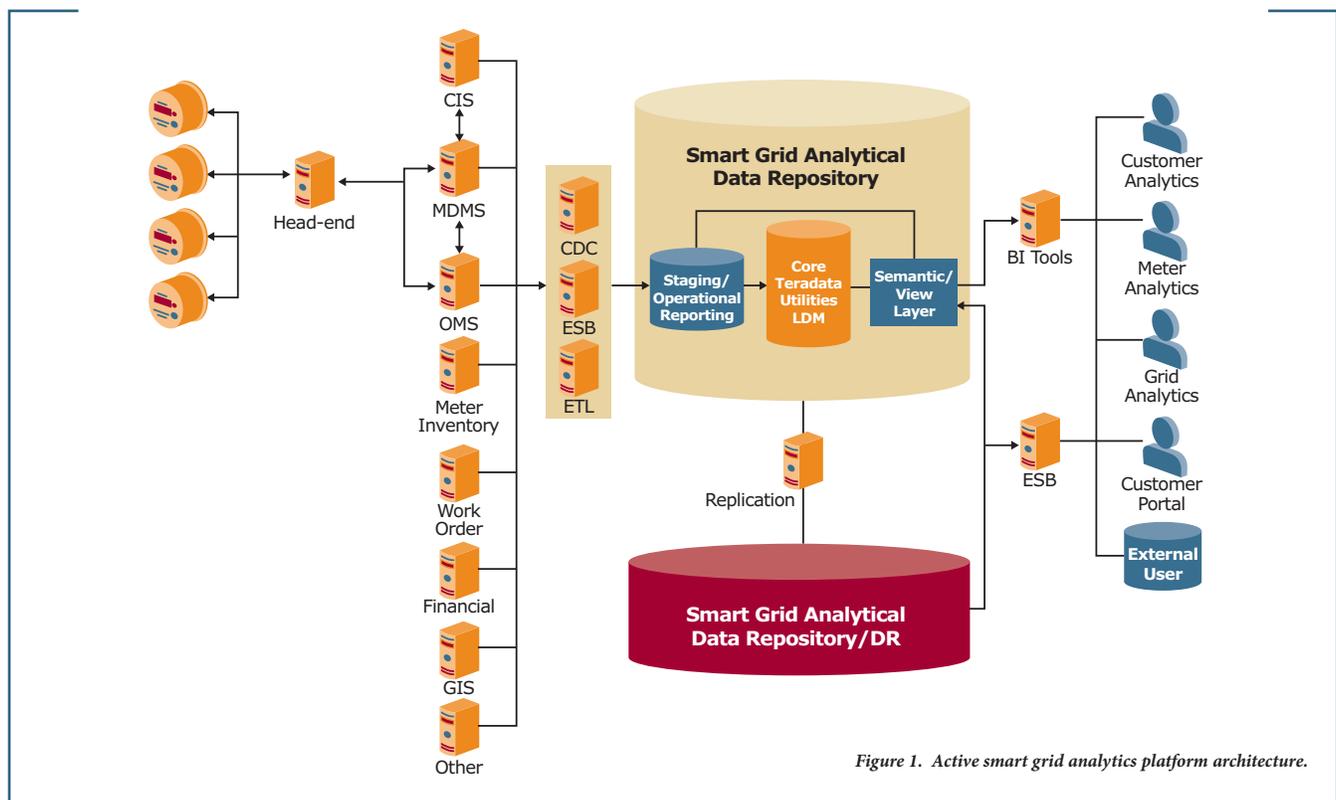


Figure 1. Active smart grid analytics platform architecture.



How Teradata Makes the Smart Grid Smarter

semantic layer, while providing auditing capabilities of all data that are accessed. Further security is managed using capabilities such as LDAP, firewalls, and encryption for hypersensitive data within the platform.

The architectural strength of the ASA Platform is to provide the valuable data resources of the utility to a large data consumer base while providing the ability to maintain data quality, security, consistency, and flexibility, and removing redundancy. Another advantage of the architecture is that it provides cross analysis of subject areas within the smart grid environment. Storing smart meter data alone provides good benefit, but when the meter data are enhanced with data from the CIS, OMS, Work Order, and other systems, the value increases many times.

Since data quality, consistency, and security are maintained at the platform layer, data access can be provided using different methods. Support for a wide range of business intelligence tools can be used to feed dashboards, provide operational reporting, and allow analytics such as ROLAP or Geospatial. Advanced tools can access detailed data for investigation using regression, what-if, correlation, and other analytics. The third aspect of support is the ability to allow tactical access to the data through a service layer integrated with the corporate Service Oriented Architecture strategy.

The final aspect of environment support is the provision of a disaster recovery (DR) infrastructure driven by business needs. Since the environment can grow to many terabytes, simply copying all the data may or may not be the correct approach. The ASA Platform approach allows the utility

to choose only certain data elements or limit the amount of history available within the DR infrastructure. The two platforms can be set up to work in either an active/passive or active/active environment. This approach helps to reduce the complexity and cost of maintaining a standby DR system.

Growing into the Architecture

To get started quickly and simply, Teradata solutions can easily grow as new data needs and applications arise. By providing a *getting started* methodology, referred to as a Phase 0, the utility starts by pulling valuable data from the MDMS and CIS system to begin building out infrastructure. The infrastructure offers business value by utilizing the meter and customer data and delivers the underlying framework for solving many new business problems. See Figure 2 for the simplified architecture.

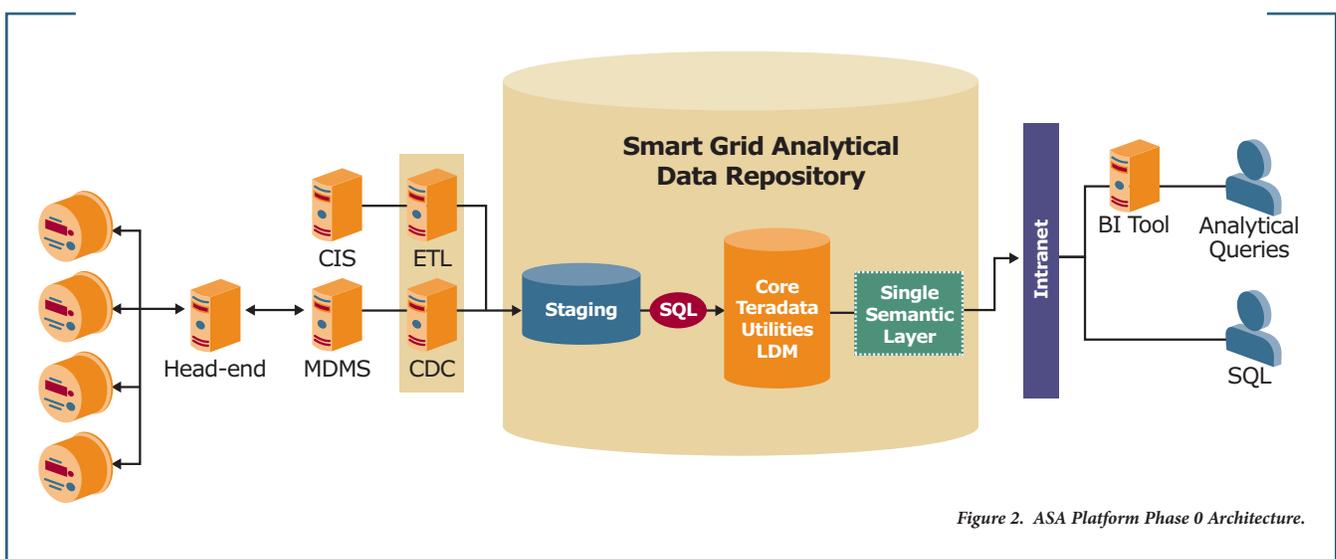


Figure 2. ASA Platform Phase 0 Architecture.

How Teradata Makes the Smart Grid Smarter

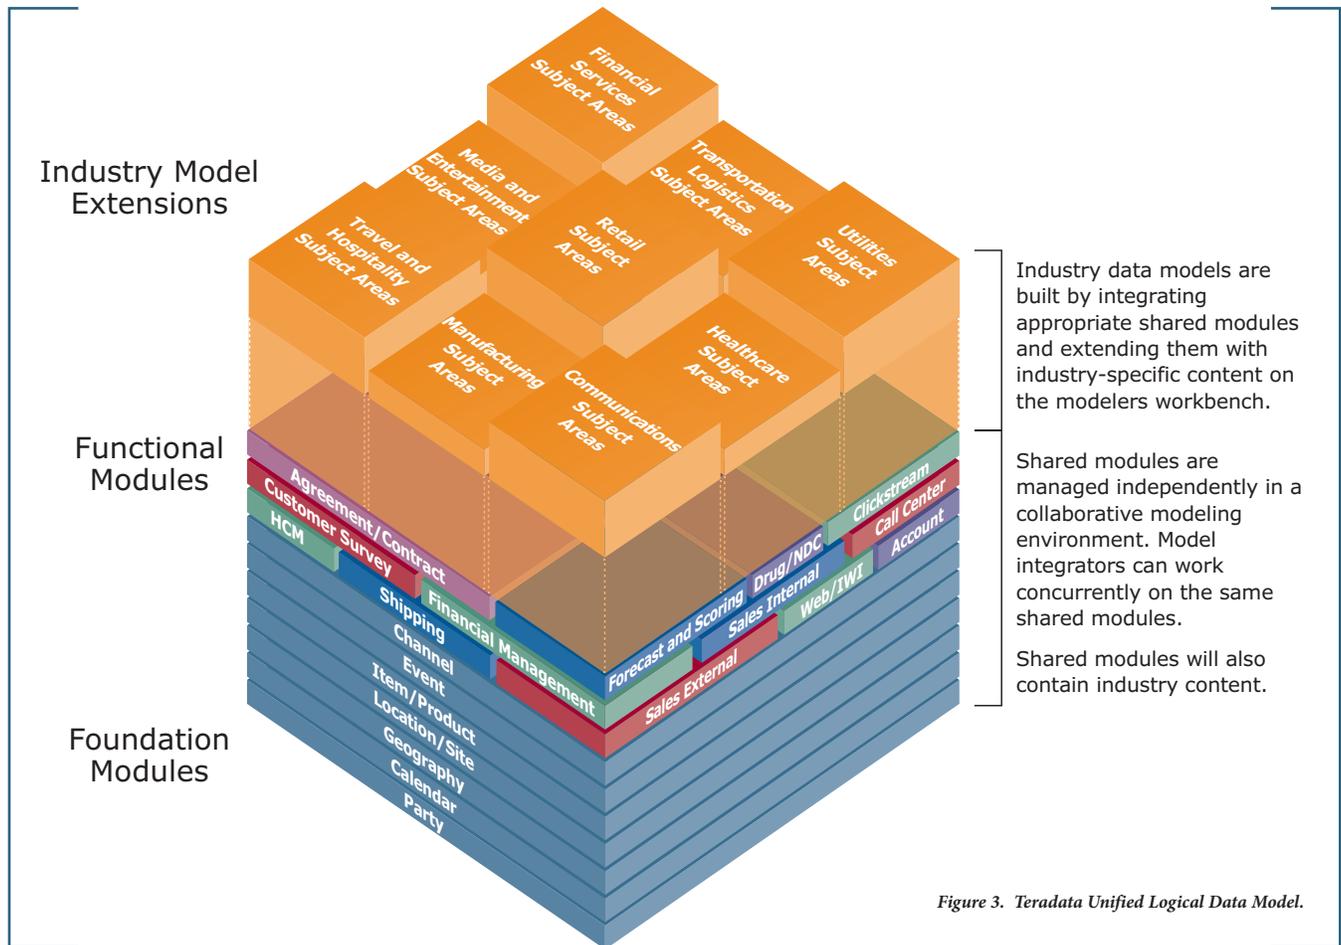


Figure 3. Teradata Unified Logical Data Model.

Teradata Platform

To facilitate the management of data within its ASA Platform, Teradata provides several critical components. These are broken down into three categories: the Teradata Unified Logical Data Model (LDM), the Teradata Database platform, and Teradata enterprise integration.

Teradata Unified Logical Data Model

Teradata offers a robust enterprise logical data model (LDM), the LDM, which

supports a number of industries, including utilities, communications, financial services, retail, insurance, travel, and manufacturing. The LDM helps an enterprise cope with abundant data, compressed business cycle times, and constant economic pressure. Teradata's LDM is designed to provide easy-to-follow blueprints for designing an enterprise data warehouse (EDW) that reflects business priorities tailored to an industry's specific needs. Developed through years of experience serving

customers in the industries mentioned, Teradata's industry-specific LDMs clearly organize and structure data, defining which individual information elements are required and how they relate to one another to provide a data model for the entire enterprise. Unifying the LDMs from multiple industries provides the ability to share knowledge from many industries since most corporations cross traditional boundaries incorporated in a single model. By crossing these boundaries,



How Teradata Makes the Smart Grid Smarter

utility companies can incorporate subject areas specific to the utilities industry and enhance the data infrastructures by using subject areas from industries such as the communication network area or the e-business area.

The model uses a three-tier approach. The bottom tier provides foundational subject areas for subjects such as PARTY, GEOGRAPHY, and CHANNEL. The middle tier provides cross-industry application-specific subject areas. The final tier provides industry-specific subject areas.

Specific to the utilities industry, Teradata has extended the LDM to create a utilities-specific Logical Data Model (uLDM). Teradata's uLDM includes critical subject areas around meter infrastructure, customer premise, and other further extensions into the transmission and distribution components of the utility.

The Teradata LDM is a model of the enterprise data, data relationships, topic areas, and business rules relating to enterprise IT. Our LDM offers an integrated, single view of the business that allows business and IT users to communicate about information needs and systems. The LDM presents a picture of all of the pieces of information necessary to run the business. The model facilitates rapid understanding and alignment of business needs, and supports "what if" analyses of different approaches. As the data warehouse evolves, the roadmap assists in identifying the next most valuable data to source.

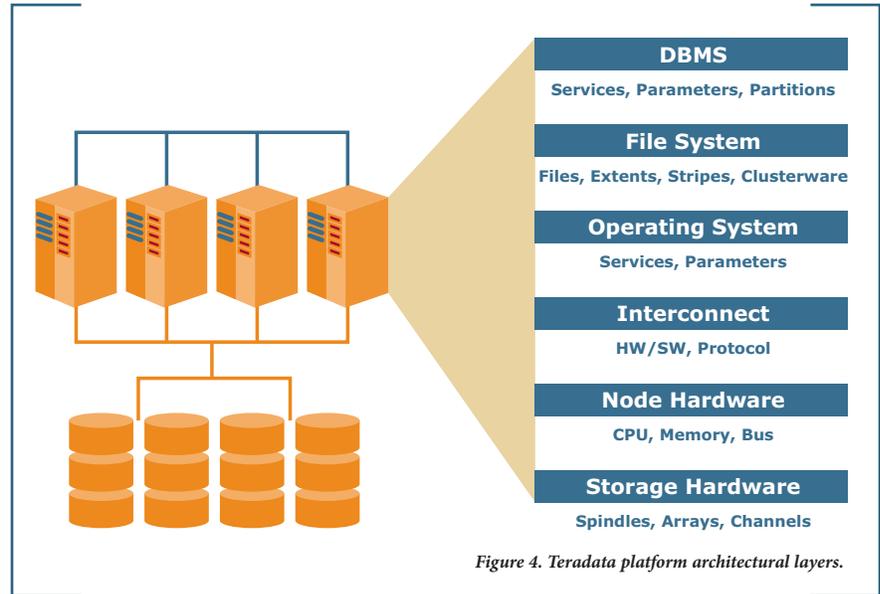


Figure 4. Teradata platform architectural layers.

The LDM approach ensures better data are used for decisions. Data are cleanly stored one time with proper integration, access controls, quality checks, full attribution, and integrity constraints. Not only does this mean better data for fact-based, data-driven decisions, it means greater value for the enterprise as it seeks to turn its information assets into decisions and actions.

Teradata Database Platform

Teradata provides a complete architecture to support the data infrastructure of the data warehouse environment. The architecture and support can be defined as building blocks, as depicted in Figure 4.

At the core of the Teradata solution is the Teradata Database. Teradata Database is an ANSI-standard, relational database management system. Teradata Database's parallel architecture was established from

the beginning as a shared-nothing massively parallel processing (MPP) infrastructure. All activity within the database is done in parallel, and Teradata Database provides linear scalability. This allows some Teradata customers to run data warehouses larger than 1PB, though most run 5-50TBs.

At the lowest level is the hardware infrastructure which is also an MPP shared-nothing architecture. The platform has been optimized to support the needs of a data warehouse environment that includes linear scalability, availability, investment protection, and common components (See Figure 5).

The platform consists of a set of symmetric multi-processing (SMP) nodes interconnected to disk arrays. All nodes are connected using the Teradata BYNET[®]

How Teradata Makes the Smart Grid Smarter

fabric. The BYNET is a specialized interconnect that uses switching technology to allow linear scalability because, as a new node is added, additional capacity is added to the BYNET fabric. In contrast, interconnects based on Ethernet technology or other single-pipe capabilities have a limit. The platform can support one to 4,096 nodes with up to 14TBs of user data space per node.

Teradata delivers powerful, proven solutions that bring together different data sources from across the enterprise to give a single view of the business. Organizations can choose from an entire family of data warehousing platforms. The Teradata Purpose-Built Platform Family has expanded to fit unique business needs. With the latest appliance options including the Extreme Data Appliance, the

Extreme Performance Appliance, the Data Warehouse Appliance, the Data Mart Appliance, and the Active Enterprise Data Warehouse, an organization can start small while trusting that its best-in-class infrastructure will grow with it. Figure 6 shows the Teradata Purpose-Built Platform Family members.

Teradata Enterprise Integration

As mentioned earlier, information delivery needs to be very flexible. Depending on business requirements, different business tools are needed for the support of delivering information to the data users. Teradata provides many types of interfaces into the data warehouse for data consumption. Additionally, Teradata supports the large majority of BI tools, to include MicroStrategy®, Business Objects®, Cognos®, Oracle Hyperion®, Microsoft®, and SAS®.

The following are the Teradata tools specific to interfacing with the data warehouse. Components include tools for delivery and for ingest of data.

Teradata Tools and Utilities

To help make sure an organization can get the most out of its data warehouse investment from the very beginning, Teradata provides a powerful set of tools and utilities that addresses core data warehouse administration and database access needs.

Teradata tools and utilities incorporate industry-standard interfaces to help integrate the Teradata system into an enterprise. Teradata tools and utilities also streamline the data warehouse management tasks it performs every day – such as moving massive volumes of data and accessing multiple data sources in parallel – along with taking crucial steps toward improving overall Teradata system performance. The product suite contains the following tools and utilities.

Teradata Utility Pack

The Teradata Utility Pack uses industry-standard interfaces to help integrate the Teradata system into the enterprise and streamline the data warehouse management tasks that are performed every day. Use this powerful package of tools to:

- > Develop applications according to industry standards.
- > Simplify data analysis and management.
- > Make data available anytime, anywhere, with any query in any language.
- > Leverage the vast amounts of information stored in your data warehouse.

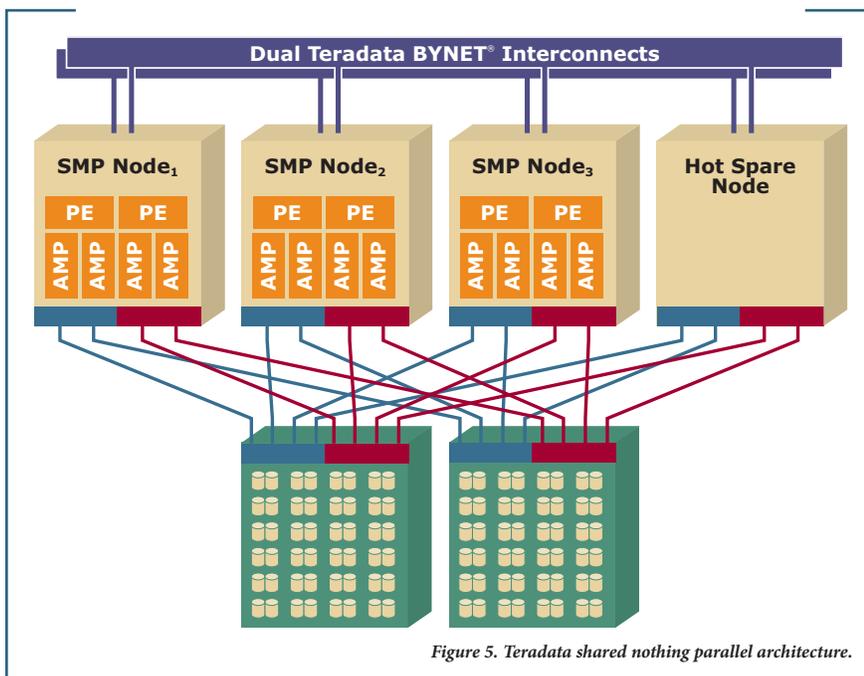


Figure 5. Teradata shared nothing parallel architecture.

How Teradata Makes the Smart Grid Smarter

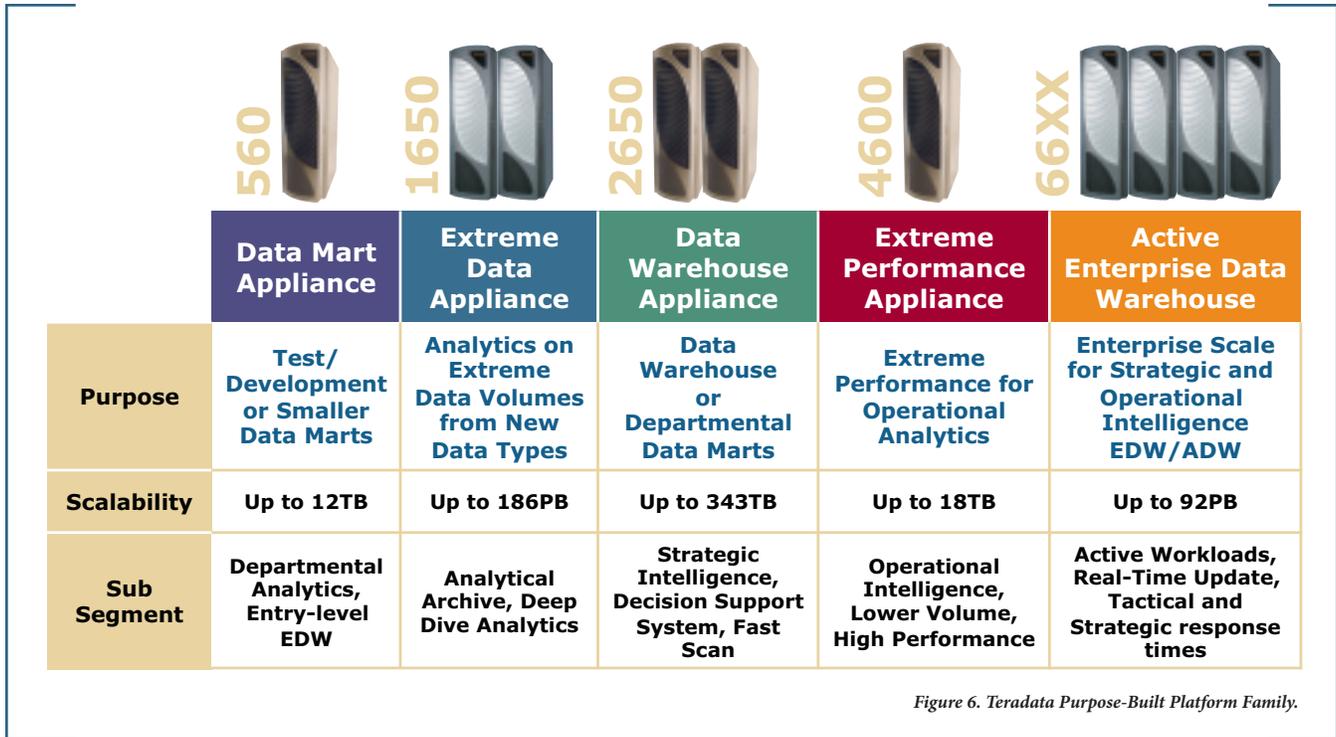


Figure 6. Teradata Purpose-Built Platform Family.

Primary components:

Teradata Open Database Connectivity (ODBC) Driver delivers open standards-based access to the Teradata Database from many popular client applications and tools.

Teradata Call Level Interface (CLiV2) provides an API for sending SQL requests to the Teradata Database.

Teradata JDBC Driver provides Internet and intranet access to Teradata Database via industry-standard solutions.

Basic Teradata Query (BTEQ) Language is a general-purpose, command-based tool that enables communication with one or more Teradata Databases. BTEQ

provides a batch or interactive interface that allows the submission of SQL statements, the import and export of data, and report generation.

OLE DB Provider for Teradata provides open, standards-based access to Teradata Database from many applications on Microsoft Windows®-based clients.

.NET Data Provider for Teradata provides the capability of embedding Teradata connectivity into any application developed using the Microsoft .NET Framework.

Teradata Meta Data Services (MDS) is a comprehensive solution for managing metadata in data warehouse environments

by enabling the location, consolidation, management and navigation of warehouse metadata.

Teradata SQL Assistant, a Microsoft Windows-based SQL query tool, issues queries to any ODBC data source and stores them on the desktop.

Teradata Administrator is an easy-to-use, Windows-based data dictionary management tool.

Teradata Plugin for Eclipse is a plug-in for Teradata to help facilitate development of Web services, Java-stored procedures and Java-based user-defined functions (UDFs).



How Teradata Makes the Smart Grid Smarter

Teradata Load and Unload Utilities

Today's growing data warehouses demand fast, reliable tools to help the enterprise acquire and manage data and flexibility to load large volumes of data from any source at any time. Challenges come from everywhere, including more data sources, growing data volumes, dynamically changing business requirements, and user demands for fresher data. Teradata Load and Unload Utilities offer:

- > Quick access to data for more timely decision making.
- > Solutions for the entire spectrum of load requirements from batch to near real time.
- > Unmatched scalability for large volume loads.
- > Fail-proof loads with checkpoint restart capabilities.
- > Proven technology from the data warehouse technology leader.
- > Integration with industry-leading ETL and ELT tools.

These tools provide a wide array of features that are efficient, scalable, and flexible, allowing an organization to centralize data from disparate sources seamlessly into the Teradata system. They are fully parallel to provide optimal and scalable performance for getting data in and out of the Teradata Database. In addition, Teradata utilities allow you to import and export data to and from host-based and client-resident data sources, including mainframe host databases, enterprise server databases, or departmental data marts.

Teradata Parallel Transporter was designed for increased functionality and customer ease of use for faster, easier, and deeper integration. It provides several methods for data ingest. All of Teradata's ETL partners are certified to integrate with Teradata Parallel Transporter.

Teradata Parallel Transporter Stream is a highly parallel capability designed to continuously move data from data sources into Teradata tables without locking the affected table. Teradata TPump provides near-real-time data into the data warehouse, allowing you to maintain fresh, accurate data for up-to-the-moment decision making.

Teradata Parallel Transporter Load is a highly reliable, parallel-load capability designed to move large volumes of data collected from data sources on channel and network-attached clients into empty tables in the Teradata Database.

Teradata Parallel Transporter Update is a time-tested, highly reliable, parallel-load capability used to create and maintain Teradata Databases. Teradata MultiLoad optimizes operations that rapidly acquire, process, and apply data to tables in a Teradata Database.

Teradata Parallel Transporter Export is a high-speed capability that quickly exports large data sets from Teradata tables or views to a client system for processing, generating large reports or loading data into a smaller database.

Teradata BTEQ, as documented above, supports interactive or batch interface that

allows you to submit SQL statements to import and export data in addition to the Load/Unload data utilities.

Teradata Database Management

A full suite of management tools for the best-of-breed Teradata Database, including:

Teradata Replication Services ensure that the data in transactional data stores and the Teradata Database are synchronized.

Teradata Performance Monitor is an application used for monitoring database sessions, resources usage, and performance statistics.

Teradata Query Scheduler provides a database request scheduling service that schedules queries through client applications, such as the Scheduled Request Viewer and Teradata SQL Assistant.

Teradata Dynamic Workload Manager is a highly reliable, next-generation systems management tool that manages requests to the Teradata Database according to a company's business needs.

Teradata Active System Management is a goal-oriented, automatic management and advisement technology that supports performance tuning, workload management, capacity planning, and configuration and system health management.

Teradata Workload Analyzer helps with the analysis of varying workloads based on historical usage data from the Teradata Database.

The Teradata logo is displayed in a bold, orange, sans-serif font. The letters are closely spaced, and the 'T' is particularly prominent.

How Teradata Makes the Smart Grid Smarter

Teradata Analyst Pack

As application environments expand to include mixed workloads for both decision support and near real-time analytic processing, maximizing the performance of the Teradata Database becomes more challenging. Teradata provides the Teradata Analyst Pack to enable analysis and tuning of database queries for better performance.

Teradata Visual Explain visually displays the query execution plan generated by the Teradata optimizer so that query planners and DBAs can immediately see where improvements can be made and can rewrite certain query sections for analysis.

Teradata Index Wizard automates the process of picking secondary indexes for a given database and workload – a natural grouping of queries – to maximize system performance.

Teradata Statistics Wizard collects usage and performance statistics for analysis by other Teradata performance tools.

Teradata System Emulation Tool allows for the quick projection of a production environment by emulating a larger production system in a smaller test or development environment, reducing the cost of query plan analysis and overall development efforts.

Teradata Viewpoint

Teradata Viewpoint is a product and platform strategy for deploying systems administration tools and end-user self-service capabilities. It is aimed at providing

business users with quick, easy to access status regarding the reports and queries they submitted by enabling them to drill down into current session details when needed. This helps them pace their work habits, make better use of the Teradata data warehouse, and know when or when not to call the help desk for support. For the database administrators and operations staff, Teradata Viewpoint provides an easy way to assess system status to understand resource saturation, trends, and problem queries. This helps them to maintain agreed-to service levels and support business users more efficiently with reduced cost.

Teradata Viewpoint is a rack-mount server appliance that contains everything needed to deploy the systems administration portal. It contains the Intel®-based appliance and software collection that supports the execution of portlets. Each portlet is a software module that displays system or query status within a browser screen. Teradata provides bundles of portlets that serve specific functions, such as system monitoring or business user self service. Over the course of the next few years, Teradata will supply new portlet bundles for a variety of tools, utilities, and applications that the enterprise can leverage for different business needs. Customers can also build their own portlets using a development kit available on the Teradata Download Center.

Portlets available from Teradata include: **Teradata Management Portlets** monitor the health and current status of Teradata servers and in-flight query sessions.

Database administrators, programmers, and operations staff use them as a monitoring dashboard. The portlets included are Capacity Heat Map, Productivity, Query Monitor, Metrics Graph, Today's Statistics, Remote Console, Space Usage, Lock Viewer, Metrics Analysis, Alert Viewer, and Node Resources.

Teradata Self-Service Portlets provide new functionality for business and technical users. They allow quick checking of system status and tracking of specific query sessions. The portlets included are My Queries, System Health, Calendar, Canary Response Times, External Content, and SQL Scratchpad.

Teradata Active System Management Portlets allow DBAs and administrators to fine tune their mixed workload environment by setting up system and workload throttles, filters, classifications, and exemptions. These portlets include Workload Monitor and Workload Health, which are for reporting and monitoring of a mixed workload environment.

Figure 7 depicts five separate portlets appearing in a single browser-based portal page surrounded by red boxes for illustration purposes. Illustrated are the System Health, System Productivity, Capacity Heat Map, Calendar, and My Queries portlets.

Business users of Teradata Viewpoint will quickly be able to determine how busy the Teradata system is. This allows them to pace their daily work for better productivity.

TERADATA

How Teradata Makes the Smart Grid Smarter

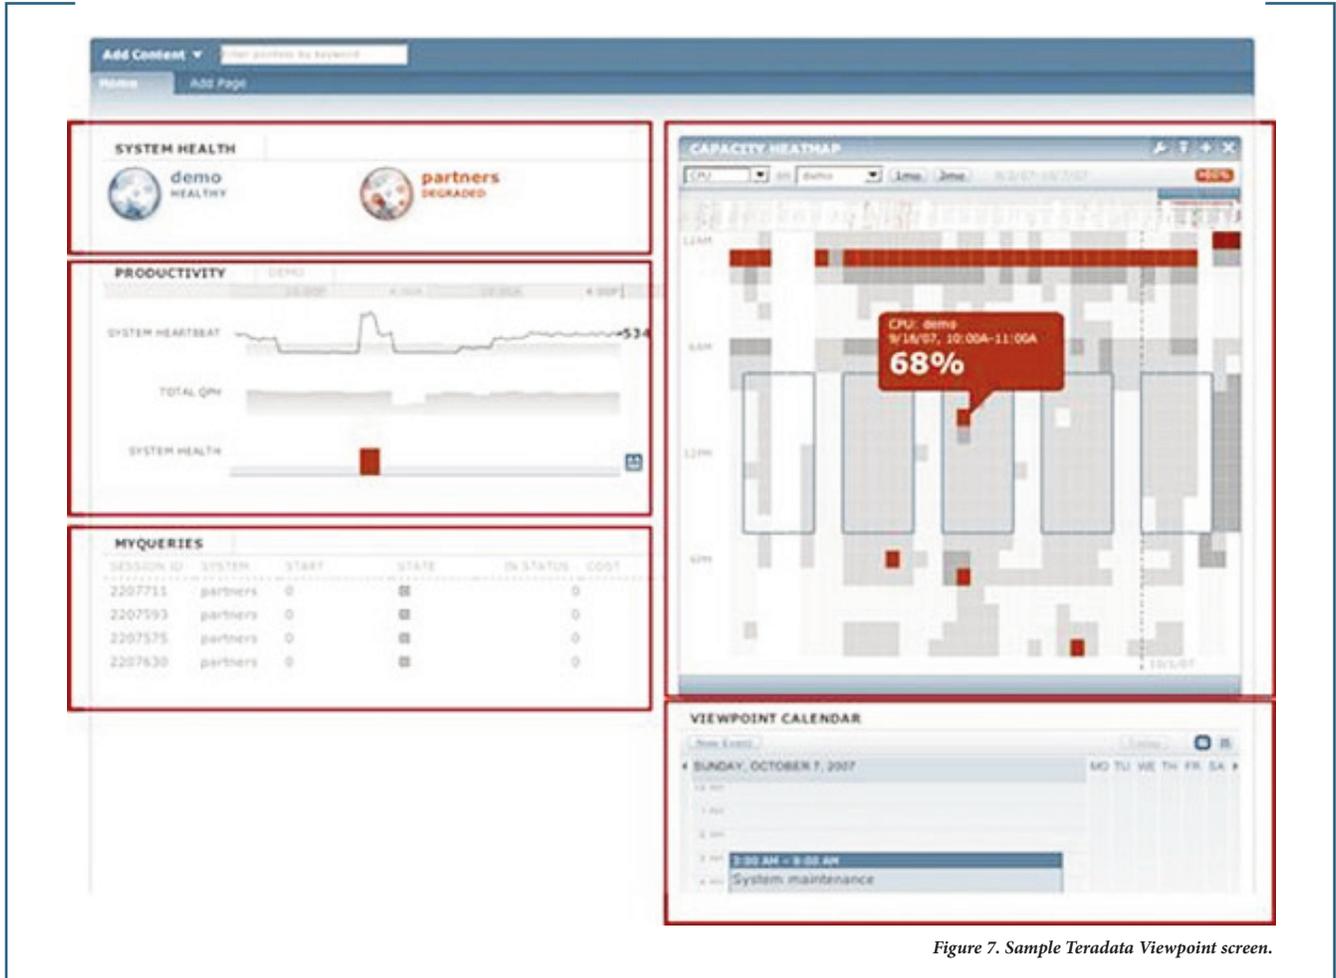


Figure 7. Sample Teradata Viewpoint screen.

They can also use Teradata Viewpoint to directly check on the status of My Queries. Database administrators, IT operations staff, programmers and other technical users will use Teradata Viewpoint to determine system status, trends, and individual query status. By observing trends in system capacity usage, they can better plan new project implementations, batch jobs, and system maintenance to avoid peak periods of use. The technical staff will also

use Viewpoint to quickly troubleshoot developer or business user query problems.

Teradata Backup, Archive, and Restore

Teradata Backup, Archive and Restore (BAR) solutions provide a fully supported and integrated set of industry-leading, matched, and certified components that store and manage data in a way that is compatible with a Teradata data warehouse. Because they are

designed to work within a Teradata system, Teradata BAR solutions minimize the interruption of critical business processes.

The Teradata BAR Framework is the advocated architecture for all BAR solutions. Composed of BAR servers and a Gigabit Ethernet network, the framework positively impacts the performance and availability of a data warehouse.



How Teradata Makes the Smart Grid Smarter

Teradata High Availability Components

This section provides information about other products that are available if businesses need the capabilities.

Teradata Dual Active Environment

The following components are specific to the Teradata dual active environment. The environment is designed to provide maximum flexibility for running a dual active environment or an active/passive environment. The two systems aren't limited by location, and they don't require that both systems are identical.

Teradata Unity enables the Teradata Analytical Ecosystem. It provides a robust solution for routing users between systems based on where their data reside, or based on pre-defined routing rules, or both. Teradata Unity also offers data and

database synchronization through SQL Multicast. Incoming SQL updates are applied to all participating systems within the Teradata Analytical Ecosystem. The updates can be data updates or database updates. In addition, Teradata Unity improves end-user experience through Auto Resubmit. Inflight queries that get aborted following a database restart are automatically and transparently resubmitted to an alternate Teradata system that can satisfy the request.

Teradata Multi-Systems Manager (MSM) facilitates the management of a dual active or passive solution and provides high availability options whether the environment is co-located or across many time zones. Teradata MSM provides the control platform to monitor the activity of data loading, data replication, workload management, and system availability. It receives alerts from all components

involved with the data warehouse and ensures that systems are either healthy or that alerts have been resolved. To monitor the activity of Teradata MSM, monitoring is integrated into a Viewpoint portal. The Viewpoint portal receives real-time status updates of all processes executing as part of the Teradata MSM environment.

Teradata Data Mover provides the ability to perform large-scale replication between two Teradata environments, using highly optimized Teradata Parallel Transporter and archive tools. Teradata created Data Mover, in addition to its data replication capabilities, to improve the impact of large-scale data loading processes. Since many data feeds load data using a batch format, Data Mover extracts and loads to keep up with the high volume. Each Data Mover process is monitored and controlled by the Teradata MSM server.

This document, which includes the information contained herein, is the exclusive property of Teradata Corporation. Any person is hereby authorized to view, copy, print, and distribute this document subject to the following conditions. This document may be used for non-commercial, informational purposes only and is provided on an "AS-IS" basis. Any copy of this document or portion thereof must include this copyright notice and all other restrictive legends appearing in this document. Note that any product, process or technology described in this document may be the subject of other intellectual property rights reserved by Teradata and are not licensed hereunder. No license rights will be implied. Use, duplication, or disclosure by the United States government is subject to the restrictions set forth in DFARS 252.227-7013(c)(1)(ii) and FAR 52.227-19.

Active Data Warehousing is a trademark, and Teradata, BYNET, and the Teradata logo are registered trademarks of Teradata Corporation and/or its affiliates in the U.S. and worldwide. MicroStrategy is a registered trademark of MicroStrategy Incorporated. Cognos is a registered trademark of International Business Machines Corporation in the United States, other countries, or both. Microsoft is a registered trademark of Microsoft Corp. in the United States and/or other countries. Business Objects is the registered trademarks of SAP AG in Germany and in several other countries. 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. Teradata continually improves products as new technologies and components become available. Teradata, therefore, reserves the right to change specifications without prior notice. All features, functions, and operations described herein may not be marketed in all parts of the world. Consult your Teradata representative or Teradata.com for more information.

Copyright © 2011 by Teradata Corporation All Rights Reserved. Produced in U.S.A.

TERADATA®