

Increase Insights while Reducing Costs and Complexity with Teradata's Unified Data Architecture



DATA PLATFORM

New analytical capabilities have emerged that enable insights beyond traditional analytics. Data volumes have exploded, fueled by the web and Internet of Things, that offer new details on customer behaviors and operational activities. Organizations are struggling to work within limited budgets and to sort through the myriad technology choices. These trends make it increasingly challenging for enterprises to connect the dots for business insights and actions.

In the quest to do things quickly and to be agile, organizations are solving use cases one at a time. This application-centric approach quickly leads to data and analytic silos that diminish in value over time. The silos need to be minimized wherever possible because they increase data redundancy, reduce quality of decisions, increase complexity, and can ultimately slow the organization down.

Data-driven companies realize more value from their data by taking a data and analytic-centric approach to business insights. The data and analytic-centric approach, as opposed to the application-centric approach, allows you to *connect the dots* across your business. This enables breakthrough insights and provides the agility to answer new business questions—all while reducing overall costs and complexity.

The Only Fully Unified Analytics Solution

The Teradata Unified Data Architecture™ (UDA) is the only truly unified data management solution on the market, orchestrating best-of-breed and complementary technologies to align with a multitude of analytical needs.

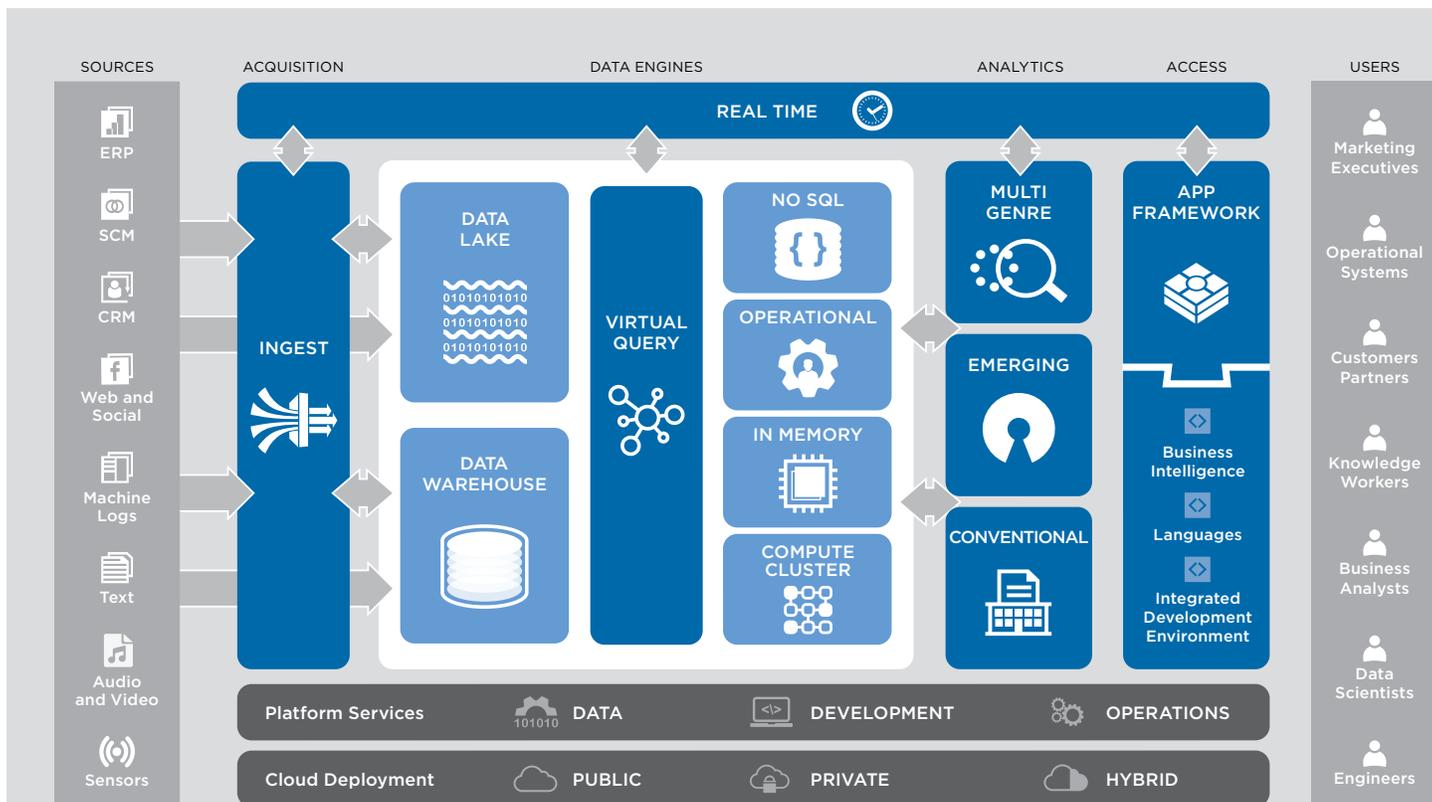


Figure 1. Teradata Unified Data Architecture

Core technologies by themselves are not enough; they have to work together to provide business value. Teradata extends and aligns its broad portfolio of robust products and services with key partnerships to help customers get the highest value from their analytics, providing easy data access for users and simplifying data movement and systems management for administrators.

Data Engines

Throughout the ecosystem, there are a variety of data engines that need to be leveraged. This allows you to store data where it makes sense as well as process data where it makes sense. The main engines are:

Data Warehouse

The data warehouse design remains the fundamental solution for integrating and cleansing heterogeneous data sources across multiple subject areas so that data is consistent, easily reused, and guaranteed to meet SLAs across the enterprise.

Teradata Database runs the world's leading data warehouses. With more than 35 years of cutting-edge development, Teradata Database is the market-leading platform for delivering strategic and operational analytics throughout your organization so users can access a single source of consistent, centralized, integrated data at the lowest total cost of ownership. Coupled with the Teradata Workload-Specific Platform Family, companies have a broad array of choices to meet their requirements for performance, scalability, availability, and budget.

Data Lake

Organizations are looking for cost-effective ways to capture the value of the big data “tsunami”—extraneous data and new data types which in the past were often discarded due to cost, scale or processing limitations. To meet this need, companies are turning to the data lake. A data lake is a collection of long term data containers that capture, refine, and explore any form of raw data at scale, enabled by low cost technologies, from which multiple downstream facilities may draw upon. But the data lake is really a design pattern, and as such there are options on platforms.

Teradata has deep engineering partnerships with Cloudera, Hortonworks, and MapR, offering customers the ability to choose which Hadoop distribution best meets their needs, while ensuring enterprise integration capabilities and the advantage of world-class Teradata service and support.

Virtual Query

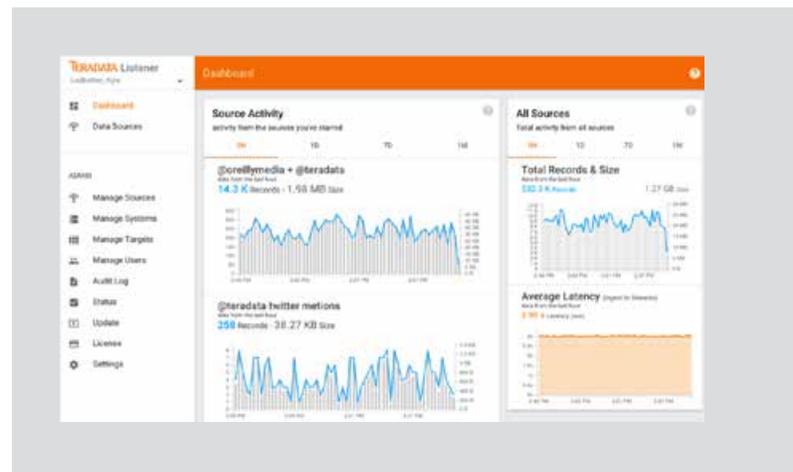
When data is persisted in multiple engines, a virtual query capability is needed to enable users and applications to connect-the-dots, while minimizing data movement.

Teradata QueryGrid™ lets your business work with a seamless data fabric across all of your data and analytical engines for no-hassle analytics. Users get the most value out of all your data by taking advantage of specialized processing engines operating as a cohesive analytic environment. By transparently harnessing the combined power of multiple analytic engines to address a business question, Teradata QueryGrid delivers the most comprehensive analytic solution to businesses.

Ingest

Data ingest is of critical importance as without data there is no analysis. Today's fast moving data streams present new challenges to organizations looking to easily and reliably ingest raw data into a multitude of data engines.

Teradata Listener™ is an intelligent, self-service solution for ingesting and distributing extremely fast moving data streams throughout the analytical ecosystem. Listener is designed to be the primary ingestion framework for organizations with multiple data streams. Listener is built using open source components such as Kafka and Mesos, along with modern software engineering, and Teradata expertise. Listener has comprehensive development and monitoring capabilities to ease the task of identifying, transmitting, and loading data into any of the data engines throughout the UDA.



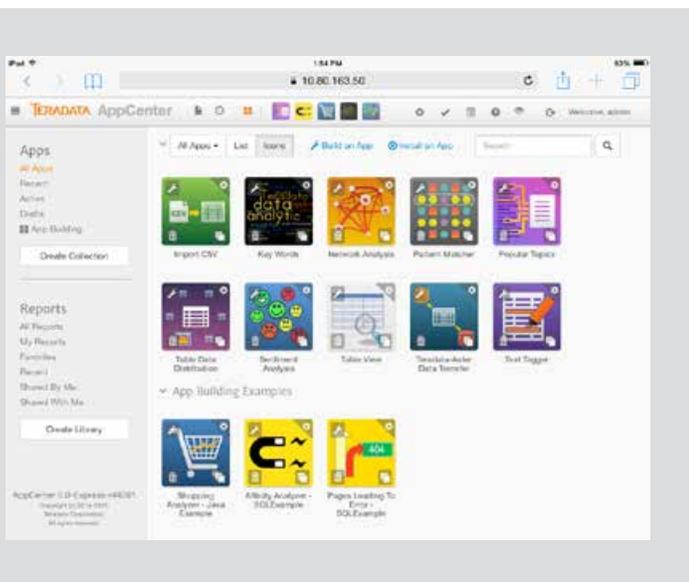
Analytics

Conventional analytics are those that have been tried and true for decades, such as regression, CHAID, decision trees, and countless others. Emerging analytics stem from big data innovation to enable procedural analytical processing such as graph and text analysis. Multi-genre allows for the integration of all types of analytics into a single processing environment.

The Teradata Aster Analytics delivers such Multi-Genre Advanced Analytics™ at scale to empower business users to uncover and operationalize non-intuitive insights. With Aster Analytics, analysts are able to seamlessly combine different analytic techniques against multi-structured data to address any use case within a single solution. This solution includes over 100 prebuilt advanced analytics functions covering multiple analytic genres such as path, pattern, statistics, text, graph and machine learning. It also has business user-friendly interfaces, such as SQL, R and point-and-click applications, that allow analyst to quickly explore, discover, and model data across any source then operationalize business analytics.

App Framework

An App Framework capability fosters collaboration and leads to faster time to value. Teradata's AppCenter provides a user friendly web based environment for the development and deployment of applications across the



business community. Allowing for quick definition and integration of commonly used analytics and then providing a marketplace for reuse and collaborative self-service. Pre-built big data apps are available to solve specific business challenges in industries like consumer financial, entertainment and gaming, healthcare, manufacturing, retail, communications and cable, travel and hospitality.

Real Time

Some data events have very perishable value with respect to time, whereby failure to take action within milliseconds makes all the difference. Examples include a request for payment that may be fraudulent, or sensor reading that may indicate a pending mechanical failure. Real Time is the capability to leverage data in motion in order to take action on events occurring in that moment of time. Real Time capabilities often benefit from tight integration with historical data and analytical tools to improve context and rapidly iterate on enhancing real time algorithms.

Teradata helps organizations leverage the latest open source tools such as Kafka, Spark and others. Additionally, Teradata partners with the leading Complex Event Processing (CEP) vendors in the market.

Platform Services

Teradata is focused on continuing to build out services that enables developers, administrators and analysts to build UDA applications in a way that simplifies the development and operating process. Platform services allows for seamless integration with internal and external ecosystem tools, and fosters collaboration and innovation with the UDA development community.

Teradata's Platform Services are developed with a *cloud-first* mentality, leveraging micro-services and multiple levels of abstraction to ensure scale across the most diverse set of workloads. Examples of commonly used services would be for transforming, caching, logging, authenticating, provisioning, scheduling, and monitoring.

Cloud Deployment

Organizations deploying the UDA have the flexibility to run workloads in on premises, public or managed cloud, and hybrid environment. The Teradata Workload Specific Platform Family can be deployed on premises for maximum performance and peace of mind from deploying behind the firewall.

Tapping into Teradata has never been easier with the choice of purpose-built managed environment or deploy on Amazon Web Services. Both offer superior solutions for data warehousing and advanced analytics. Have confidence knowing that Teradata's cloud products use the same industry-leading software as our engineered systems.

In fact, hybrid deployment is the new name of the game: the right tool for the job, the greatest ROI. Combine on-premises performance with cloud-based flexibility for the absolute best of both worlds.

Services

As a leading data and analytics systems integrator, Teradata can help clients develop business led data strategies, architect and implement diverse analytic ecosystems consisting of open source and commercial software and hardware, and fully manage and support those analytic ecosystems. Teradata has a wide array of skills and expertise, including deep business, industry, and architectural consulting skills to drive business value from Teradata Unified Data Architecture investments. Teradata continues to expand consulting capabilities in support of the Teradata Unified Data Architecture. Think Big, a Teradata Company, brings specialized engineering and data science skills to deliver innovative big data and open source solutions for our clients

A Smart Decision

By providing operational insights that help you increase productivity, streamline costs, and unleash the value of data, Teradata's UDA lets you know more so you can do more—the ultimate competitive advantage. For more information, visit Teradata.com/UDA or call your Teradata representative.

10000 Innovation Drive, Dayton, OH 45342 Teradata.com

Unified Data Architecture, QueryGrid, Teradata Listener, Multi-Genre Advanced Analytics are trademarks, and Teradata and the Teradata logo are registered trademarks of Teradata Corporation and/or its affiliates in the U.S. and worldwide. 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 © 2016 by Teradata Corporation All Rights Reserved. Produced in U.S.A.

04.16 EB6732



TERADATA.