

Multi-Temperature Data Warehousing with Teradata

Performance Services

Increase Storage Value with Dynamic Performance Management

CIOs across the globe are grappling with volumes of business data exploding at an unprecedented rate. In addition to detailed transaction data, such as call detail records, click-stream logs, and itemized receipts, regulatory and legal retention requirements mandate that years and years of historical data be readily accessible for analysis. The necessity to maintain, analyze, and protect this data is motivating businesses to rethink their approach to storage architecture and evaluate new methods of delivering information to both front-line workers and executive decision makers across the enterprise.

When using a data warehouse to deliver strategic and operational intelligence, the first performance consideration is response time, or 'how much time will it take to complete my queries?' Response

time is largely dependent upon how quickly a disk drive can transfer data to the server, server processing speed, and database architecture.

As disk drive technology has evolved, the speed at which disks retrieve and transfer data to and from a server has remained fairly constant, regardless of the storage capacity of that disk. This means that your response time requirements, not data volume requirements, will determine the number of disks you need for your data warehousing solution.

So while the storage capacity of enterprise disk drives doubles every 12 to 18 months, this trend has not resulted in a reduction in the number of disk drives required for a given performance level. With increasing disk capacities, systems will have far more usable storage than ever before. Unfortunately there is insufficient server input/output (I/O) performance to utilize this



capacity with frequently accessed data. This is causing vendors to recommend that not all of the available disk capacity be used.

Teradata Corporation provides a unique approach to utilizing the unused disk space without impacting performance levels. This is accomplished by adding infrequently accessed yet valuable data to the warehouse which may have been previously stored offline.

Our workload management capabilities and services balance performance and storage requirements based on your needs to maximize the return on your storage investment while maintaining the performance required to support critical functions.

Teradata provides the ability for large volumes of infrequently accessed data to coexist with very frequently accessed data while dynamically balancing query performance resources based on business requirements – we refer to this concept as Multi-Temperature data warehousing.

Benefits of Teradata Multi-Temperature Services:

- > Gain better understanding of historical data to improve offer generation and customer service.
- > Meet data retention requirements for regulatory compliance.
- > Enhance risk assessment of customers, products, and investment portfolios.
- > Achieve greater operational efficiencies.
- > Gain maximum benefit from storage investments.

Multi-Temperature Data Warehousing with Teradata

What is Multi-Temperature?

Temperature is a metaphor used to classify data based on frequency of access – the more frequently the data is accessed, the hotter the temperature of the data. Multi-Temperature data warehousing describes an implementation concept where the Teradata® Database manages both the active (hot) data and the inactive (cool) data, which has not traditionally been kept in the data warehouse due to cost constraints. The larger storage capacities available with today's disk drives combined with Teradata's Multi-Temperature capability create a cost-effective solution for managing both hot and cold data.

Teradata's unique capabilities allow you to maintain lower TCO by managing the workload on the system as opposed to managing the data. By sharing and managing storage across temperature ranges, applications accessing hot data can maintain required performance levels

With Multi-Temperature data warehousing, Teradata customers can manage workloads as opposed to managing data.

while rarely accessed data are still available to users, but with reduced performance priority.

Workload balancing is addressed through a combination of Teradata technologies, such as:

- > Partitioned Primary Index (PPI)
- > Multi-value compression
- > Priority Scheduler
- > Teradata Dynamic Query Manager
- > Teradata Active System Management

Applying Multi-Temperature across Industries


Virtually all industries can benefit from Multi-Temperature data management.

> Retailers often store hot daily and weekly sales data in the data warehouse for comparative analytics (day versus week versus month). For retailers, RFID data will be some of the most frequently accessed data since more retailers are adopting RFID and subsequent event data to make better, faster decisions.

Retailers also store and access cooler data, such as deep history used for market basket analysis, promotional analysis, product movement, or annual sales data comparisons. The determination of whether data is cooling is often defined by business rules (e.g., data more than 90 days old is cold or based on how often the data are accessed.)

> Insurance companies rely on vast amounts of current and historical data for business decisioning. Detailed data and frequently accessed information, or hot data, are necessary to drive response for customer service, such as policy quotes, claims, policy inquiries, underwriting, and claim adjustment assignments.

What is Multi-Temperature Warehousing?



Strategic Intelligence

Temperature determined by:

- Frequency of access
- Frequency of updates
- Data maintenance
- Depth of history

Multi-Temperature Data Warehousing with Teradata

Insurance companies also need to store cold data, allowing access to detailed historical data and information to solve a number of statistical, statutory, and management requirements, including pricing, valuations, underwriting, funding, reserving, audit, regulation, risk management, and corporate governance.

In a multi-temperature warehouse, insurance agents and customer service representatives can access hot data – maintaining required performance levels while rarely accessed data is still available to all users, but with reduced performance priority.

- > Banks gain access to detailed customer transaction data dating back more than two years for fraud, more than five years for audit, and seven or more years for predictive model development. Knowing your customers and being able to discern patterns against activity requires nearly continuous modeling of detailed historical and current period data. Data warehouse users can simultaneously access hourly, daily, weekly, and monthly financial transactions that can run against behavior models to reduce fraud and criminal activities.

Regulatory and audit requirements are also broader and deeper in scope for banking institutions. Access to a single

source of managed hot and cold data is critical to meeting those demands on time and with a level of accuracy that is now required. Access to consistent and transparent data is vital to the reputation and performance of the organization.

Finally, the institutionalization of enterprise risk management techniques and practices requires large amounts of detailed and historical data. Whether modeling for credit risk and probability of default, or modeling economic series and scenarios, risk management is dependent on both hot and cold data.

- > Communications providers are leveraging near-real-time incoming call detail data to identify best cost routing of roaming traffic. In addition, the most recent 90 days of call detail data are often analyzed to ensure quality of service commitments are satisfied. This frequently accessed hot data coexists in the data warehouse with cooler data, such as 12-24 months of customer data records stored to help resolve a lawsuit or arbitration over intercarrier settlement claims.
- > Manufacturers need to store multiple years of warranty information for easy analysis of trends in defects or quality improvements. For example, due to various federal, state, and local government regulations in such areas as safety (TREAD act in the auto industry) and

emission compliance, detailed data must be stored from 5 to 20 years. With multi-temperature, years of warranty data can coexist in the data warehouse with more frequently accessed data, such as warranty claims, technician network, and early warning systems that give manufacturers early warning signals of potential quality problems in manufactured products.

- > Healthcare organizations find that data and information are vital to the workings of the system. Millions of pieces of raw data pass through various IT and care management systems, and regulations, such as the Health Insurance Portability and Accountability Act (HIPAA) and Sarbanes-Oxley, require that data be managed across the entire organization and not just within a departmental or functional area.

Access to both hot and cold secured data in the data warehouse helps support a number of enterprise needs for product development, pricing, administration, adjudication, and management. Whether for mortality and morbidity, statutory reporting, financial management, or operational performance, a single source of access to both hot and cold data is important so the healthcare organization can meet the demands of customers/patients and regulators.

Multi-Temperature Data Warehousing with Teradata

Teradata.com

Teradata's Multi-Temperature Data Warehousing Services

Teradata multi-temperature services consist of customized analysis and implementation capabilities. First, we perform our Multi-Temperature Analysis Service, where our Professional Services team spends approximately three weeks identifying business processes that can use cold data or a mix of cold and hot data; and defining existing/future application data requirements and existing/future application workload profiles. We also identify storage optimization opportunities and design a solution that potentially includes changes to the data model, load processes, and application access capabilities. All of these elements are incorporated into a comprehensive recommendation that can serve as your data management roadmap.

Then, we perform a Multi-Temperature Implementation Service, where our consultants implement the required changes and take advantage of Teradata's workload management tools and features. After design and implementation of the solution, we conduct thorough testing and monitoring of initial system performance.

Professional Services Support

Teradata Professional Services consultants work directly with you to deploy Teradata Multi-Temperature Services. With more than 2,000 Teradata professionals worldwide, our staff has deep technical expertise in data warehousing, business intelligence tools, applications, and platform technologies. Working with your team, we'll deploy the services to help you perform these key activities:

- > Capacity planning – determine the recommended system configurations that meet your needs for capacity, performance, and price.
- > Physical database design – identify and recommend changes required to enhance system performance.
- > Database maintenance – run performance tests to ensure that loads are faster, examine whether data purge rules have changed, and ensure that a viable recovery strategy is in place.
- > Workload management – achieve workload balance using a number of Teradata tools, including Teradata Dynamic Query Manager, Teradata Dynamic Workload Manager, Priority Scheduler facility, and Teradata Active System Manager.
- > System management – collect data and monitor the data warehouse.

Maximizing Your Data Warehouse Investment

Teradata Multi-Temperature Services is the right solution to maximize your data warehouse investment when you want to achieve greater business value, regulatory compliance, and cost reductions.

With Teradata's differentiated services backed by our global Professional Services staff, we can help you develop a strategy that cost effectively meets your long-term needs. Our goal is simple: to ensure that your enterprise data warehouse can handle large amounts of current and historical data seamlessly, without sacrificing performance as you support your day-to-day operations.

For More Information

To learn more about how Teradata's Multi-Temperature Services offer can assist you with your historical and current data storage requirements, contact your Teradata representative or visit Teradata.com.

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 © 2006-2007 by Teradata Corporation All Rights Reserved. Produced in U.S.A.