

Teradata Columnar

DATA WAREHOUSING



Store More Data; Analyze Them Faster

Faced with massive data sets, a growing user population, and performance-driven service level agreements, organizations everywhere are under extreme pressure to deliver analyses faster and to more people than ever before. That means businesses like yours need faster data warehouse performance to support rapid business decisions, added applications, and better system utilization.

And as data volumes continue to increase—driven by everything from longer detailed histories to the need to accommodate big data—companies require a solution that allows their data warehouse to run more applications and to be more responsive to changing business environments. Plus, they need a simple, self-managing system that boosts performance but helps reduce administrative complexities and expenses.

Teradata® Columnar, a transformative advance in in-database technology, delivers new levels of speed for increased performance and higher compression rates that let you store even more data. It's a hybrid row and column database that fully integrates columnar and row-based tables establishing new levels of flexibility, performance, and compression and strengthening your data warehouse analytics capabilities.

Extreme Performance

Teradata Columnar enables the Teradata Database to run faster and deliver higher compression rates. Its advanced architecture delivers the most comprehensive ready-to-use environment available for organizations confronting the challenge of delivering analytics in the world of accelerating big data.

Teradata Columnar is a leading-edge solution that clears performance bottlenecks by storing data in columns—unlike most relational database management systems which only store data in rows. It allows you to

“With Teradata Columnar, Teradata combines the performance advantages of its new hybrid row and column capabilities with Teradata [Database]’s high availability, rich SQL support, optimization, and workload management. Teradata customers can now enjoy the benefits of columnar database capabilities integrated into the Teradata platform.”

– David Menninger, Vice President and Research Director, Ventana Research

mix-and-match columnar and row-based physical storage when it best suits applications. So the applications get the right data at the right time.

With Teradata Columnar, only the data in the columns required for a query are pulled into memory for processing, vastly reducing the time-constraining input/output (I/O) of a row-based approach that would read data from all the columns.

Dramatic Compression

In addition to the significant performance benefits of storing data in columns, Teradata Columnar unblocks I/O congestion in yet another way: compression. Compression reduces the I/O required to read data into memory because the necessary data to answer a query are compressed to a fraction of the original size. Plus, Teradata Columnar automates compression administration, freeing DBAs from the analysis necessary to choose from several complex compression alternatives. The database automatically uses the best compression mechanisms for the situation.

TERADATA.

Teradata Columnar:

- Dramatically compresses data so you can effectively store larger amounts in the same footprint.
- Dynamically evaluates new data automatically to determine appropriate compression algorithms.
- Changes the compression mechanism or dictionary compression values as the data evolve to maintain optimum compression.
- Chooses among several compression mechanisms, including: dictionary, run length encoding, NULL, trim leading or trailing zeros or blanks, Delta from mean, UNICODE to UTF8.
- Works in conjunction with multi-value, algorithmic, and block level compression.

Teradata Columnar offers benefits to the business and technology facets of your organization.

For Business:

- Improves data warehouse performance by only reading the data absolutely required for each query.
- Lets you respond to more queries faster while storing more data.
- Integrated solution can mix and match the storage formats in hybrid row and column storage delivering the highest performance.
- Simple to use; minimal DBA effort required, and no changes necessary to existing applications or queries.

For IT:

- Fully integrates columnar and row-based data orientation at the column level establishing new levels of flexibility, performance, and compression.
- Compresses the data to effectively store larger amounts of data in the same footprint.
- Dynamically adjusts compression mechanisms for optimal storage as data evolve over time.
- Distinct architectural approach creates containers within each column so data can be compressed more efficiently.

Easy to Use

With Teradata Columnar, DBAs have an additional physical database design option that provides the utmost freedom to use column- or row-oriented storage for any portion of the database to produce extreme query performance and dramatic data compression. The Teradata Database will intelligently choose row or column storage formats; and for columnar data it will automatically choose the best compression method or methods.

Integrated

Teradata Columnar is integrated into Teradata Database so that it works with all Teradata Database features. It works with other performance features such as indexes, Teradata Intelligent Memory, and the Teradata optimizer. Highly compressed columnar data increase the leverage of Teradata Intelligent Memory™ by allowing even more very hot data to fit into memory for even higher performance. Also, being part of Teradata Database provides unlimited scalability, high availability, and self-managing administration. This means you no longer have to sacrifice other database attributes to take advantage of column-oriented database tables to reduce I/O and improve performance. Columnar can now be used in a data mart or a large integrated data warehouse.

All-in-the-Family Availability

All Teradata workload-specific platforms can benefit from Teradata Columnar for specific analytical workloads. However, each platform benefits differently based on its unique architectural design. For example:

Teradata Active Enterprise Data Warehouse was built for active workloads. It is enhanced with Teradata Columnar because it can now store even more data in solid state drives (SSD) providing higher performance.

Teradata Data Warehouse Appliance was designed for data scan-intensive workloads. Teradata Columnar improves this capability through executing with pinpoint accuracy on data, streamlining the query, and reducing I/O. The result is increased performance on analytic workloads.

Teradata Integrated Big Data Platform stores massive volumes of data. Teradata Columnar provides dramatic compression on unique data types, which enables the storage of even larger data sets in the same foot print. This reduces the price-per-terabyte for deep history analytical solutions.

Why Teradata?

Teradata helps companies get more value from data than any other company. Our big data analytic solutions, integrated marketing applications, and team of experts can help your company gain a sustainable competitive advantage with data. For More Information

To find out how Teradata Columnar can help you and your business store more data than ever—and analyze them faster—so you can improve your decision-making capabilities and grow a stronger, more productive business, contact your local Teradata representative or visit Teradata.com.

How Can You Take Advantage of Teradata Columnar?

All types of businesses and industries can benefit from Teradata Columnar, including:

Telecommunications—While a consumer is waiting on the phone with a customer service agent, the agent must search the consumer's Call Detail Records (CDR) which may have more than 100 columns. However, to answer the caller's specific questions, the agent may need data from only a few of the columns. Teradata Columnar can decrease, by 90 percent, the amount of data needed for the query. This significantly improves response time to the consumer by reducing I/O and dramatically improving query performance.

Financial Services—When a bank marketing manager uses CRM data to more effectively personalize each sales contact, she often needs fewer than 10 attributes about the consumer, such as customer number, demographic characteristics, the last product purchased, and the last channel used. However, a CRM system captures hundreds of customer attributes. With Teradata Columnar, the amount of data read from the customer record is reduced because only the 10 required attributes are read. This capability supports high performance and the millisecond response time to queries required for inbound marketing.

Retail—A purchasing agent ordering products for his store chain may only be interested in a selected number of items based on specific SKUs. However, retail demand chain management systems store a much broader set of attributes for each item sold, including the date, store locations, and SKU, to meet a variety of reporting requirements. But the business user only wants data relevant to the specific order and doesn't require all the metrics. With columnar, the database only reads the data referenced in the agent's question. The result: higher performance when compared with reading all of the columns in the table.

10000 Innovation Drive, Dayton, OH 45342 Teradata.com

Teradata Intelligent Memory is a trademark, 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 © 2015 by Teradata Corporation All Rights Reserved. Produced in U.S.A.

04.15 EB6443



TERADATA.