

# Teradata Database's Ease of Management

**A Total Cost of Ownership Advantage**

By:

Rob Armstrong

Carrie Ballinger

Teradata Corporation



# Teradata Database's Ease of Management

## Table of Contents

<i>Executive Overview</i> .....	2
<i>Introduction</i> .....	3
<i>31 Tasks You Never Have to Do with a Teradata Database</i> ....	3-4
<i>Teradata Benefits that Contribute to Lower Total Cost of Ownership</i> .....	4-9
<i>Conclusion</i> .....	9

## **Executive Overview**

There has been an interesting change in the data warehouse arena in the past few years. As more and more vendors target their products toward business analysis, customers are starting to see the databases and other tools as commodity purchases. This has led to data warehouse platform decisions being centered on hard dollar costs and savings of acquisition. Unfortunately, the perimeter of this costs and savings analysis is often unduly narrow, not taking into account the cost of maintaining the system, its growth, or the work associated with the end users of the warehouse.

This paper addresses the savings that accrue when constructing a data warehouse with a relatively self-managing database, such as Teradata® Database. These savings are significant enough to affect Total Cost of Ownership (TCO).

# Teradata Database's Ease of Management

## Introduction

The real cost of the data warehouse is not in acquiring the hardware and software (operating system or the DBMS), but rather in the application development and overall management effort. Hardware and software prices are decreasing while people costs are increasing, so addressing the recurring people costs is a logical place to begin in significantly reducing TCO. A closer look will also provide some compelling reasons why migrating to Teradata Database is very cost-effective.

This paper will help you to begin to question and understand your real costs. By using the points in this paper as guideposts to assist in that process, you will gain a deeper understanding of just how Teradata Database can dramatically reduce the total cost of ownership.

The first section is a simple list of tasks that are no longer required when a Teradata Database is installed.

The second section contains charts that list benefits of the Teradata Database, along with the potential contribution made by each benefit toward a reduction in TCO. As the impact of these benefits varies across different applications, we have not attempted to quantify the dollar or percent savings demonstrated by each contributing benefit. We leave that exercise to you — based on your understanding of your own environment.

## 31 Tasks You Never Have to Do with a Teradata Database

To understand how Teradata Database can reduce your TCO, you must first understand that it is fundamentally different from other databases. One way to understand this difference is to consider some of the tasks that no longer need to be done — by anyone. The following are some of the most obvious tasks that are eliminated when a conversion to Teradata Database is made.

### Implementation

1. Understand and set extensive operating system tuning parameters to achieve optimal database operations.
2. Create and format data files to hold the data and the indices.
3. Write programs to determine how to divide data into partitions.
4. Determine the size and physical location of each table and index partition or simple tablespace.
5. Code the space allocation for each partition or underlying file structures.
6. Embed partitioning assignments into CREATE TABLE statements.
7. Determine the level of parallelism to be assigned to tables or indices.
8. Assign and manage special buffer pools for parallel processing.

9. Associate tables and/or queries with degrees of parallelism.
10. Code the definition and allocation for temporary workspace.
11. Create rollback segments or log files.
12. Ensure that the data are spread evenly across disks and controllers.
13. Design database defensively to prevent arduous operations, such as table scans, joins, and aggregation.
14. Build summary tables before end users can access the data warehouse.
15. Build and partition materialized view logs.
16. Build and partition indices on top of materialized views.
17. Determine how materialized views are updated, asynchronously or synchronously.
18. Add hints or otherwise rewrite SQL.

### Support

19. Monitor partition size.
20. Monitor and tune temporary work and sort spaces.
21. Monitor and tune buffer pool assignments.
22. Monitor and tune parameters and control blocks that enable parallel execution.

# Teradata Database's Ease of Management

## Planning

23. Perform periodic table and index re-orgs, such as unloads and reloads, dropping and rebuilding.
24. Convert data types of mainframe data sets prior to a data warehouse load.
25. Sort and/or split the data before a load job.
26. Set up multiple load jobs from a mainframe to the data warehouse to load a single table in parallel.
27. Manually restart the multi-step load process when failure occurs.
28. Calculate and configure fail-over plans for DBMS, such as HACMP.

## Growth and Leverage

29. Alter the parallelism assignments as the number of users or data volume increases.
30. Expand partition boundaries or relocate partition data sets.
31. Add or delete table or index partitions as tables grow.

## Teradata Benefits that Contribute to Lower Total Cost of Ownership

Teradata Database offers potential financial savings in a multitude of ways from the 31 tasks noted previously and

from other distinctions of Teradata Database versus other RDBMS products. Benefits will be listed in the following four data warehouse application development life cycle phases: Planning, Implementation, Support, and Growth and Leverage.

## Planning (See figure 1)

The Planning phase is used to determine the business need; the criteria and metrics to measure success; the logical architecture and data necessary to meet those needs; and finally, a gap analysis to determine how ready the customer is to move forward.

Step	Teradata Benefits	Contribution to Lower TCO
<b>Business discovery</b>	Proven methodology	Reduces time from process reuse
	Experienced facilitators	Increases value by showing what outcomes need to be driven
<b>Logical database design</b>	Experienced consultants	Doesn't require analysis time to pre-guess database compromises
<b>Database assessment/ data mapping</b>	Proven methodology	Reduces total time for assessment
	Mainframe experienced consultants	Enables higher quality and faster results from people who have experienced your platform
<b>Architectural design</b>	"Glass House" understanding	Removes time later in process due to otherwise unforeseen areas of concern
<b>Solution readiness</b>	Ability to use simplified architectures	Reduces time to solution as fewer gaps to close with fewer structures in the architecture
	Reduced education and training needs	Helps lower cost and shorten time to implementation

Figure 1. Planning

# Teradata Database's Ease of Management

## Implementation (See figure 2)

This phase consists of putting the findings of the planning phase into practice. Data are transformed and loaded into the database. Data are checked for integrity and consistency across the model, and where necessary, some preliminary changes may be performed. Front-end applications are developed or installed, and controlled end users are starting to

access the system. Also in this step, the incremental load (or delete) processes are tested and routine data management protocols are developed for turn over to production. At the end of this phase, you have a live data warehouse environment.

Teradata methodology calls for 90-120 days to implement the initial subject area and lay the proper foundation, which can then

be built upon as expansion is required. Contrast this to other vendors who often take longer to implement initial systems which often end up with reduced functionality and little ability to grow. The man-hour difference alone can be hundreds of thousands of dollars.

Step	Teradata Benefits	Contribution to Lower TCO
<b>Configuration planning</b>	Smaller DASD ratio	Reduces hardware cost
	Full parallelism thus lower overall CPU needs	Reduces hardware cost
	Logical space management	Eliminates time spent slicing and assigning DASD segments
	Same database and same approach to SMP or MPP	Eliminates the need for software conversions with growth
	Simple expansion process	Allows true starting with a small system and upgrading later, reducing initial hardware costs
<b>Physical database design and modeling</b>	Over 90% of LDM usable directly in the physical model	Enables easier design effort, days versus weeks
	Modest need for indexing or denormalization	Simplifies design effort further, reduced DASD cost and DBA setup effort
	Ability to support a single data store	Removes duplicated work, eliminates consolidation and synchronization efforts
<b>Configuration planning</b>	Host-based utilities and direct host connection	Eliminates work required to build multi-step job streams for mainframe to UNIX® conversions
	No need to split, order, or convert data	Further reduces time required to implement load processes, no need to re-work with system growth
	Can use parallel processing to accomplish many transformation steps	Reduces required MIPS on host platform, enables quicker execution with parallelism
	Can run with atomic level data and perform aggregation and joins at execution time	Lowers development cost, improves flexibility, makes future change cheaper; ends need to write pre-join or pre-aggregate code for transformation processes

Figure 2, part 1. Implementation

# Teradata Database's Ease of Management

Step	Teradata Benefits	Contribution to Lower TCO
<b>Application development</b>	Lessened need for special indexing, denormalization, or use of temporary tables	Speeds implementation, lowers DBA calculation and involvement
	Optimizer takes SQL from tool and/or user and runs without hints or tuning	Removes DBA time required to capture and analyze tool-constructed SQL and optimize physical model for the queries
	Use of views with no performance degradation	Reduces DASD storage for redundant data, lessens management or physical entities
<b>Data maintenance and management (carries into Support phase)</b>	No table or index re-orgs	Eliminates DBA time to plan and execute these long running processes
	Logical space management	Eliminates tasks associated with calculation and physical placement of data
	Automatic data placement	Reduces analysis time for data demographics, query demographics, and distribution options
	Automatic space management	Removes time necessary to evaluate, allocate, and monitor table and work spaces
	Full parallelism for all queries	Reduces DBA work to analyze and optimize systems workloads and remove or design around serial bottlenecks
<b>System maintenance and management (carries into Support phase)</b>	Minimal tuning parameters	Removes DBA tasks of micro-managing platform, database, and query interactions
	Administrative workstation	Reduces administration of multiple nodes, all seen as single system
	The same for SMP and MPP	Eliminates DBA rework or rebalancing with growth
	Native connects to host	Reduces time for transfer of data, load balances automatically across channels, ends need for UNIX landing areas
	Teradata Database is isolated from the operating system (OS)	Minimal OS tuning, only needs to bring node up, reduces DBA tuning efforts
	Automated DBMS Level Fail-over	Removes set up effort for MPP systems to handle outages or lost nodes
<b>Development staff training</b>	Teradata Database is isolated from the operating system	Reduces the need to understand complex OS tuning issues
	Simplified SQL, database tuning, and utilities	Lowers learning requirements translating to quicker ramp up and lower training costs

Figure 2, part 2. Implementation

# Teradata Database's Ease of Management

## Support (See figure 3)

This phase continues the day-to-day management tasks above, as well as manages the tasks necessary to prove the Teradata system's value with audited return on investment (ROI) and begin the groundwork for additional investment

into the environment. Often times this involves the understanding of the user access and analysis, tuning the system as necessary, and the addition of new data elements (either more detail of existing data or new business data elements).

Again, this is an area that can consume tremendous amounts of resources, and many of these tasks will recur in any well-developed system. In addition to the above advantages in the day-to-day management of the system and data, Teradata also provides significant cost relief in the following areas and tasks.

Step	Teradata Benefits	Contribution to Lower TCO
<b>Logical and physical database reviews and expansion</b>	Able to support third normal form or lightly tuned physical designs	Simplifies planning for and implementing expansion thereby reducing DBA efforts
	Database resource usage and query statistics available	Removes large portion of daily DBA tasks required on other databases
	Logical data management	Reduces DBA effort in planning DASD layout
	Reconfig utility, automatic redistribution of data during growth	Minimizes DBA and support staff planning and execution, no need to rework utilities and queries after utility is completed
<b>Tuning</b>	Optimizer robustness and intelligence	Reduces DBA involvement in collection, analysis, and resolution of performance problems
	Ability to perform well against generic business model designs	Simplifies data model as new elements are added. DBAs do not need to retrofit data into a denormalized schema
	Automatically managed index options	Removes planning and execution of index data management, makes indexed data available in real time with base table
	OLAP and data mining extensions	Optimizes database extensions for OLAP and data mining queries and minimizes DBA tuning and query-writing effort
<b>Assisting business users</b>	Teradata performs well with user-friendly tools	Requires less end-user training and fewer rules and restrictions
	No rewrite or hints required on queries	Demands less time and technical skill, satisfying end users
<b>Capacity planning</b>	Predictable and linear scaling	Simplifies growth calculation reducing the time necessary to perform needs analysis and what-if performance testing
	Account ID to capture usage and query statistics	Benefits from easily accessible and identifiable usage patterns that further reduce analytical phase
<b>ROI Audit process</b>	Proven methodology	Allows quicker and more complete analysis
	Quickly performs new analysis with design changes	Reduces complexity and time needed to run what-if queries (critical in ROI analysis)

Figure 3. Support

# Teradata Database's Ease of Management

## Growth and Leverage (See figure 4)

Any conversation about TCO must include a realistic discussion about how to accommodate growth of the system and the

leveraging of the initial foundation into the longer term strategy. Successful data warehouses grow at a surprising pace.

This growth must be framed in the proper arena: Is it due to more redundant data

of existing elements, or is it due to new business demands with call for more cross-functional data and deeper detail of existing subjects?

Area	Teradata Benefits	Contribution to Lower TCO
<b>New users, more data and new applications</b>	Easy and automatic re-configuration	Reduces time for planning and execution, shortens downtime caused by growth, eliminates DBA recalculation after growth process
	Support of large number of concurrent users	Negates additional cost of data mart on distributed database platforms to accommodate increased usage
	Cross-functional analysis with access based on a single business modeled data store	Offers less redundant data, less indexing or extracting to data marts, significant savings in set-up and reconciliation efforts
	Security or separation of data via logical views	Reduces DASD requirements, simplifies DBA task or monitoring and managing entities
	Ability to accommodate large data volumes	Single platform negates cost of multiple data mart environment, or breaking data model into separate tables
<b>Data extraction and propagation</b>	Host-based utilities with channel connects	Reduces complexity and planning for loads, reduced MIPS usage on host platform
	Ability to load source data directly from multiple platforms	Eliminates complexity, setup, and error handling of File Transfer Protocol (FTP) usage, eliminates need for cross-platform handling
	Ability to utilize logical views or advanced indexing techniques in lieu of physical data marts	Reduces DASD requirements, reduces processing to maintain marts and summaries, reduces DBA management of data and entities
<b>7x24 requirements</b>	Flexible, parallel data utilities	Decreases time to develop load and extract routines
	Concurrent read/write capability	Reduces DBA/Operations analysis to separate workloads, higher end-user accessibility
	Priority scheduling	Requires less DBA time for scheduling and monitoring of conflicting tasks
	Fault tolerant HW and SW	Increases availability and reduces operations effort to debug and repair system outages
<b>Closing the loop between operational and informational systems</b>	Host connections	Eliminates planning and execution of LAN transfers
	Robust utilities	Reduces time to develop transformation/communication routines
	Supports operational-like, simple data models, requiring few indexes	Lessens planning and execution of data transformation between systems and usage
	Mixed workload environment and large user communities	Reduces complexity and number of platforms as system expands

Figure 4. Growth and leverage

# Teradata Database's Ease of Management

Teradata.com

Teradata provides a foundation that can easily absorb new user communities, allow for atomic levels of data, and provide the access capability to users for new analysis with minimal intervention. Teradata also brings the philosophy of a single store of data to reality. Because no down stream processing is necessary, a closer link can be in place between the feeding operational system and the analysis of the informational system. This also supports the more extended need of feeding cross-functional, consistent data back to operational systems for greater impact to those functions.

The tasks associated with these four life cycle phases do not exhaust all the areas of costs savings associated with a Teradata Database implementation, but they do represent a majority of them. Teradata Database was designed to be self-

managing and flexible in the face of growth and expansion. Taken together, and considered over the life of an application, the cumulative impact of these tasks can offer enormous dollar savings compared against the cost of supporting other popular databases today.

## Conclusion

Compiling a list of benefits that will save you money as you implement a data warehouse is only a part of understanding your total cost of ownership savings with Teradata. It is suggested that you take a long, hard look at what it takes to support your current system, and to quantify what these tasks are costing you today. Many of these costs are not obvious, but some of them can be quite large, such as the time and effort an end user must spend to compile a global view of the business if

there is no centralized data warehouse. Also consider revenue you could have had from lost business, lost because the information needed to identify these opportunities was not in place.

This paper just begins to explore the benefits of the Teradata Database. For additional information, contact your Teradata representative or visit [Teradata.com](http://Teradata.com).

Teradata is a registered trademark of Teradata Corporation and/or its affiliates in the U.S. and worldwide. UNIX is a registered trademark of the Open Group. 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](http://Teradata.com) for more information.

Copyright © 2001-2008 by Teradata Corporation All Rights Reserved. Produced in U.S.A.