

Oracle to Teradata Migration Case Studies

Realizing Value and Improving Capabilities



Oracle to Teradata Migration Case Studies

Table of Contents

<i>Introduction</i>	2
<i>Background</i>	2
<i>Case #1 – A Travel Company</i>	4
<i>Case #2 – An Online Retailer</i>	5
<i>Case #3 – A Financial Data Warehouse</i>	6
<i>Summary</i>	7

Introduction

Many organizations understand the value of data warehousing and have sought to implement a warehouse based on their existing database management system (DBMS). These efforts are usually adequate and sometimes even successful for static reporting and limited queries. These data warehouses are built, however, on a platform designed initially for on-line transaction processing (OLTP) and not decision support. What many of these organizations eventually discover is that these platforms can't handle the demand and workload in a growing and dynamic data warehousing environment. These initial data warehouses eventually limit business improvement and lock companies out of the real value that data warehousing can provide.

More than ever, Oracle customers fall into this category. They are realizing the limits of the Oracle platform for decision support. Although they're likely satisfied with Oracle as their OLTP backbone, they're considering moving to the Teradata® Database for decision support. For those customers, the important questions are:

- > How does the migration take place?
- > What tasks are required to successfully complete the migration?
- > What value will I realize at the end of the process?

This paper is designed to help answer these key questions by profiling three of the more than 200 Oracle customers who have migrated to a Teradata solution. The companies profiled in this paper have

achieved performance improvements, additional business benefits that justify the migration cost, and successful integration of the Teradata Database with the Oracle OLTP environment. Their stories go beyond theory and discuss the reality of their successful migrations from Oracle to Teradata.

Background

Over the past few years, many organizations that have begun implementing data warehousing are pausing to take stock of what they've built. They've cited many reasons for reevaluating their existing architecture. For some, it's because what was built for a specific set of end users doesn't provide utility to other parts of the business. Others discuss serious performance degradation as the system, or systems, are forced to handle

Oracle to Teradata Migration Case Studies

“The legacy implementation had grown so large because of PayPal’s success that a lot of the legacy processes were starting to take more than 24 hours to run and that was just unacceptable to the business. So, we had no choice but to try to optimize and migrate to Teradata. Teradata was absolutely the right platform.”

– Clay Stanley, Director of Information Management Delivery, PayPal, Inc.

not only growing volumes of data, but also larger numbers of users and more complex queries against more complex data sets. Many organizations find that their existing data warehouse can scale along one dimension, i.e., data volumes. But this comes at the cost of other dimensions within their data warehouse environment, and they find they have to limit the number of users or decrease data or query complexity to maintain performance.

In some cases, problems are surfacing because important new business questions can’t be answered in a reasonable length of time, since data must always undergo timely aggregation, indexing, and modeling to ensure performance. There is also the constant concern that as data warehouse requirements continue to grow, these technical limitations will limit business value as well as lead to expanding maintenance costs and lower return on investment (ROI). Regardless of the motivation, many individuals within these organizations who own decision support capabilities now realize that a general-purpose RDBMS

designed initially for OLTP but pressed to also support DSS, doesn’t meet real-world business requirements; requirements to scale along the multiple data warehouse dimensions at the same time (such as data volumes, data complexity, query volumes, and query complexity). Turning to Teradata is the solution for more and more organizations facing this dilemma.

Often, the first foray into data warehousing involved storing summarized data in a data warehouse, but did not let users access it directly. The presentation layer provided access through various data marts designed specifically for a particular set of end users or applications. Frequently, these data marts were only implemented to offload reporting requirements from operational systems. Often, the initial implementation was successful because end users received reports and information they hadn’t seen before. Regardless of the initial success, problems arose as users of the data saw the potential for additional business value and wanted to ask additional questions. The original

infrastructure wasn’t designed to handle the data warehousing workloads that ensued: many users, many concurrent queries, complex queries, much data, new applications, and both large strategic and short tactical queries on the same system at the same time. The original data warehouse was built on a platform that couldn’t manage the unforeseen workload and eventually was viewed as limiting possible improvements to the business.

Teradata solutions were designed for just this kind of dynamic data warehousing workload. For those organizations that have seen the potential of data warehousing and realized the corresponding demand on a DBMS, the Teradata Database was a natural choice. Providing limitless scalability, flexibility, performance, and reliability, Teradata Database consistently outperforms any general purpose DBMS in customer benchmarks. In fact, as more and more companies grow weary of the headaches resulting from supporting a failing infrastructure, many of them are replacing their Oracle decision support infrastructure with one based on a Teradata solution. Fortunately, this is a trend that has been building for years, and the work required to perform a migration is manageable, repeatable, and process driven. Today, customers seeking to migrate can leverage the vast experience gained by Teradata consultants.

More than 200 Oracle data warehouse customers have migrated to the Teradata solution platform. In fact, Teradata has a comprehensive migration program

TERADATA
Raising Intelligence

Oracle to Teradata Migration Case Studies

designed to deliver significant value to those customers who choose to do so.

The migration program provides a specific “Teradata Solution for the Oracle Professional” training curriculum; a set of formal migration processes and procedures designed to ease, shorten, and optimize the migration process; and a consulting and field staff with years of experience performing these migrations.

Three critical steps must occur in a data warehouse migration:

1. Data that currently resides on Oracle must be moved to the Teradata Database. This is done either by transferring the data from one platform to the other or by completely rebuilding the target data sets from the source data. In practice, the data migration is usually a mix of these options.
2. The end users must have their business requirements met, if not exceeded. This is done by migrating the existing application set to work with the

Teradata solution and/or by building a new set of applications to provide new and greater functionality.

3. The Extract, Transform, and Load (ETL) processes must be updated to load the new warehouse from the old and possibly some new sources. This can be a challenge if the existing process uses Oracle’s proprietary stored procedure language, PL/SQL, which only works with Oracle. Teradata’s migration program includes efficient procedures and automated tools that streamline this migration step.

There is no one way to perform these three key steps. There are multiple options for accomplishing these steps, all of which have been tested, tuned, enhanced, and documented over the years. It would be impossible to discuss all of the details of each of these processes, but this paper will examine three particular engagements to highlight the common approaches used in the many Oracle-to-Teradata migrations.

Case #1 – A Travel Company

A large travel retailer had a data warehouse that was built on Oracle. The dilemma they faced was that the existing data warehouse was limited by its scalability, and they were unsure if it could meet future data warehouse growth needs. There were also many problems with performance of both their click stream data loading and analysis as well as with their general data mart and data warehouse functionality. With click stream data, it was taking 10-18 hours per day to load 24 hours worth of data on Oracle, and that was when everything was running correctly. Only one campaign could be run at a time due to this and other limitations. Other data loads were taking four to six hours per day on Oracle in an overnight batch, and this was continuing to grow longer and could not meet the business needs. The company’s overall data warehousing needs demanded a solution that would provide quicker response time to events on the web site, and analytics needed more timely data. The question was whether to abandon their Oracle platform, possibly risking the investment they had already made, or continue to throw money and resources at the existing solution and hope the next Oracle version might solve their problems allowing them to achieve greater results and efficiencies from their data warehousing efforts. They desperately wanted a solution that could handle hourly and near-real-time loads while both tactical queries and deeper

“What we get from the Teradata data warehouse is extremely rapid response. The kinds of queries that we run against the data warehouse, containing over 800 million records, under another platform, if I can say Oracle, would take 24 to 48 hours. We can do that in under an hour.”

- Bert Little, Executive Director of the Center for AgriBusiness Excellence (Tarleton State University), USDA (United States Department of Agriculture)

Oracle to Teradata Migration Case Studies

analytics analysis (larger ad hoc queries) were being performed on the system, but the existing infrastructure couldn't support that vision. They also realized that it was neither practical nor cost effective to believe that the 'next version' of Oracle, that was always just around the bend, would solve their problems. Fortunately, the migration to Teradata allowed them to realize their vision while actually decreasing their overall total cost of ownership (TCO).

This customer was using Oracle version 9iR2, with the warehouse composed of 450 Tables + Views. There was about 1.4 terabytes of data to move as well as hundreds of existing stored procedures. When it came to their existing data load environment, there were more than 200 Informatica® jobs to convert as well as many other custom ETL scripts and load jobs that must be converted over to the new Teradata system. Their existing reporting environment was composed of Cognos®, as well as numerous custom SQL reports. In all, more than twenty thousand lines of PL/SQL had to be converted to ANSI SQL.

To accomplish conversions from Oracle SQL to Teradata SQL, Teradata has developed a series of standard, repeatable processes and automated tools designed specifically for migrating Oracle data warehouses to the Teradata solution. This addresses the need for converting both Oracle SQL to Teradata SQL and Oracle PL/SQL to Teradata SQL. After the initial automated conversions by the tools, minor revisions are usually made to complete the transition and ensure functionality.

Though this amount of work seems daunting for any organization, with Teradata professional services consultants' depth of conversion knowledge, their tried and true methodology, as well as automated conversion tools, the entire migration took eight weeks, including the time that was required for system testing. During the migration process, the existing Teradata Travel Logical Data Model (LDM) was used as the baseline for rehosting the data.

The travel company has been ecstatic with the results post migration. Loads that used to take almost a day, now load in only a few hours, and their new architecture now allows them to utilize trickle loads. Once limited to a single campaign at a time, now another 20 campaigns have been coded and are being run simultaneously against the click stream data. At the same time, their deeper analytics analysis (larger ad hoc queries) performance has increased three to ten times.

Case #2 – An Online Retailer

A large online retailer was utilizing both the Oracle 9i and 10g database platforms to try to meet their analytical business processes. They were quickly finding that with their rapid growth and future analytical business requirements that neither their existing environment nor the Oracle platform could meet their business and technical needs. The company recognized that a new solution was needed to scale and streamline operations to better support its 24x7 customer transaction and reporting loads.

The company wanted to better understand its customers, in order to retain new customers, particularly to harness the large volumes who visit the online storefront during peak holiday shopping. The company also wanted to enable a real-time, single view of the customer to better understand purchasing habits, refine marketing efforts, and more effectively drive business to its website. Their existing Oracle environment was plagued with inconsistent data and reporting problems, thus making these goals impossible to achieve.

Multiple operational reporting environments (Oracle 9i and 10g) moved to Teradata with about one terabyte of data being moved, 850GB from four main Oracle source systems, and 300GB from numerous other source systems. At the same time, new ETL and load code was built out to ensure no data duplication and that data are consistent and accurate before being fed into the data usage systems. In the process, 250 new batch SQL jobs were built, 55 real-time feeds were put in place utilizing Teradata's partner GoldenGate's real-time data movement solution.

At the same time, the reporting environment was moved to Teradata, with more than 70 existing reports being moved to the Teradata environment as well as more than 30,000 lines of custom code replaced with Business Objects®. They had approximately 300 users on the system, including 50 external partners.

Oracle to Teradata Migration Case Studies

The complete migration including re-architecture and putting in place the real-time data movement and analytical capabilities took five months.

The company built out a true enterprise data warehouse (EDW) leveraging Teradata's retail logical data model as a baseline. In that time, they moved six years of history into the model. They also installed and had Teradata CRM with Postmaster operational in that timeframe handling more than five million daily emails. To further increase the value of their EDW, the company installed and is using Teradata Value Analyzer with the reporting system. They are now also leveraging Teradata Warehouse Miner, which runs directly in the warehouse. With the new configuration, one of the first things they were able to do was run a 3+ year customer retention study which they had never been able to do before within their Oracle environment.

The company has experienced huge business benefits post migration. The EDW is now able to be updated in real time, allowing for real-time analytics and event-driven notifications. Marketing and merchandising have greatly increased staff productivity. Where users used to spend 30 to 40 percent of their time gathering data, they have saved approximately one day per week by having all data in a single source – improving their productivity by 20%. In its Books/Video/Music area, it used to take 10 to 12 hours to run analytics. It's now done in fewer than four minutes. Through

customer retention analysis and propensity-to-buy analysis, they have experienced increases in customer retention of 10% to 15%. With Postmaster application online, e-mail campaigns realized a significant decrease in rejected emails – deliverability increased by 10%. The retailer can now determine customer profitability at the individual customer level. This was not possible in the old environment because data were stale and based on averages. The bottom line for this retailer is that almost immediately after migrating to Teradata, the company was seeing huge cost savings and increased revenue. Not only did Teradata solve the retailers' existing business and technical problems, but positioned them well for growth in the future.

Case #3 – A Financial Data Warehouse

A diversified financial services holding company that provides services to retail, corporate, and institutional customers first attempted a data warehouse based on an Oracle/SUN platform. Unfortunately, they never brought the system into production. Prior to their migration to Teradata, they had implemented only one application on the Oracle system. All was not lost, though, as they were able to reuse most of the database design and existing processes on the Teradata Database.

When this customer decided to move their Oracle-based system to a Teradata solution, they first assessed what they were going to build on the Teradata solution. Some of the key questions were: Did it

make sense to pick up and move what they had onto the Teradata system? And would the target system be as flexible as they needed for user access?

To answer these questions, they conducted a Business Discovery. This is a process where the customers of the system are questioned about their business requirements so the system can be designed to provide for both their current and future needs.

Relying on established processes and procedures for migration, Teradata professional services consultants easily transferred approximately 200 tables from Oracle to Teradata. Due to the consultants' expertise and repeatable migration processes, the migration of these 200 tables took one week, averaging 15 to 30 minutes per table. The ability to complete such a migration in so little time is the result of conducting a thorough index assessment. This assessment determines how many indices are actually necessary, and almost always results in a dramatic reduction in the number of indices that are needed in the Teradata Database. Typically, the index assessment takes more time than the actual index migration.

The actual movement of data from Oracle to Teradata took about five days. Since the Oracle unload/load utilities couldn't run the whole time, the actual clock time of the migration was approximately 24 hours. Data that were entirely available from the source system were loaded directly from the source system (static tables) and not migrated from Oracle.

Oracle to Teradata Migration Case Studies

The activities to migrate the ETL processes for this customer presented a particular challenge. The existing ETL processes were still being modified for the Oracle environment and were not stable. Source data resided in Sybase operational systems and other smaller third-party data sets.

Ab Initio® was the primary ETL tool being utilized by the Oracle system. The tight integration between Teradata and Ab Initio allowed the consulting team to easily modify the processes that existed so they could point to the new Teradata Database. The remaining ETL processes were created fresh and implemented successfully.

After the migration was complete, end users had immediate access to the data utilizing Teradata's BTEQ and Queryman utilities. Lower data maintenance requirements and performance gains were the biggest benefits this customer realized from the migration. The Teradata solution reduced the run times of the five longest running queries on the Oracle system by a factor of 13. Stability was introduced to the customer's environment, and finally, they could achieve what they had set out to accomplish with their data warehousing initiative.

Summary

Many companies face the question of their current decision support platform's ability to continue growing without drastically

“And the interesting thing is that when we were using an Informix and Oracle data warehouse, we found that we needed tremendous amount more disk space than we did on the Teradata system. And the reason is that most of our tables, and most of our data, on the prior data warehouse was put into smaller summary tables and lots of data marts. We don't need those same things on Teradata, and because of that we've saved a tremendous amount of space.”

– John Conte, IT Director, Wesco

increasing costs and complexity. These companies are also trying to determine whether or not they are satisfied with the ROI provided by their existing infrastructure. We have profiled three companies that determined they were facing diminishing returns, and they were not satisfied with the ROI provided by their existing Oracle-based warehouse system. These companies made the decision to move to the Teradata solution. When they did, they discovered:

- > Teradata has established procedures that simplify the migration process and minimize the migration project.
- > The tasks involved in the migration process are simple, straightforward, and designed to account for existing needs along with future plans.

- > The results of the migration provided a win-win. The cost of administering and maintaining the warehouse environment is significantly reduced, and the capability and performance of the new system is a dramatic improvement.

Customers decide to migrate for many different reasons. This paper has profiled just three of the more than 200 companies who have made the move from Oracle to Teradata. Teradata Corporation would like to add your company to the growing list of firms that have unlocked the power of data warehousing. For more information about migrating from Oracle to Teradata, visit Teradata.com.

Informatica is a registered trademark of Informatica Corporation. Business Objects is a registered trademark of Business Objects SA. Microsoft and Visual Basic are registered trademarks of Microsoft Corporation. SAS is a registered trademark of SAS Institute, Inc. Ab Initio is a registered trademark of Ab Initio Software Corporation. 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.