



Come closer

Denmark Telecom Leader Cuts Costs with 6 Terabyte Data warehouse on SQL Server 2005

Overview

Country or Region: Denmark
Industry: Telecommunications

Customer Profile

Based in Copenhagen, Denmark, TDC provides telecommunications services to customers in 12 countries. TDC generates annual revenue of more than U.S.\$7 billion.

Business Situation

TDC needed a data warehouse to facilitate analysis of data from more than 60 disparate systems, accessed by 2,000 business analysts and other internal users.

Solution

TDC created a 6-terabyte data warehouse using Microsoft® SQL Server™ 2005 Enterprise Edition (64-bit) and SQL Server Analysis Services running on the Microsoft Windows Server® 2003 64-bit Enterprise Edition operating system.

Benefits

- “Single Truth” with multiple views
- Ability to support more analysts
- 80 percent reduction in processing time
- Integrated development environment
- Easier data analysis

“We have 6 terabytes of data now, and anticipate data growth of about 30 percent per year. One of the reasons we chose SQL Server 2005 is that it has the scalability we need to match our growth.”

Jan Andersen, Systems Architect, TDC

From its headquarters in Copenhagen, Denmark, TDC provides land-line, mobile, and data communications services to customers in 12 countries. The company, which reported 2005 revenue of U.S.\$8 billion, needed to gather data from more than 60 disparate systems into a central data warehouse to support some 2,000 business analysts and other internal users. TDC deployed its 6-terabyte data warehouse using Microsoft® SQL Server™ 2005 Enterprise Edition (64-bit) database and SQL Server Analysis Services hosted on HP ProLiant DL585 servers with AMD Opteron™ Dual-Core 64-bit processors. The data warehouse provides the company with a centralized, authoritative, source of information across the organization. The cost of doing complex analysis, which previously required importing historical data from tape and working across disparate systems, has been slashed.

“By using SQL Server and Analysis Services to create our data warehouse we have eliminated the need to argue over the validity of data. We now have a 'single version of the truth.'”

Peter Møllebjerg Andersen, Head of Analysis, TDC

Fast Facts	
Total data	6 terabytes
Data growth	30 percent per year
Largest table	1 billion rows
Multidimensional cubes	100+
Largest cube	32 dimensions, 450 billion aggregations
Database	Microsoft SQL Server 2005 Enterprise Edition (64-bit)
Operating system	Microsoft Windows Server 2003 64-bit Enterprise Edition
Processors	AMD Opteron™ Dual-Core 64-bit processors 875 series running at 2.2GHz on HP ProLiant 585 servers

Situation

TDC is Denmark's leading provider of land-line, mobile, and data communications services. The company's 22,000 employees provide service to 3 million land-line customers, 2.5 million mobile customers, 1 million Internet subscribers, and 924,000 cable television customers in 11 European countries. It also serves some 6 million customers in Oman in the Middle East. Some 55 percent of its revenues come from international activities. The company, which had 2005 revenue of U.S.\$8 billion, prides itself in using its international experience to benefit its customers, as it realizes significant economies of scale and transfers competencies achieved in one market to its other markets.

The TDC IT infrastructure is complex, including mainframe computers, UNIX systems, and computers running Microsoft® Windows Server® operating system and servers. As the company has grown, often through acquisitions, the diversity of its IT infrastructure has increased. This complexity resulted in disparate methods of data harvesting and the consequent formation of

islands of data that couldn't be easily accessed or merged with other data.

TDC had a collection of third-party reporting applications, including SAS Essbase, and Brio. For day-to-day operations, the disparate reporting systems and types of data meant that TDC analysts had to spend too much of their time becoming experts at extracting data, when they needed to instead concentrate on their core competencies of analyzing data for business planning and control. The problems of disparate information systems were complicated further by differences in how different business units approached data quality, handling, and transformation. Variations in how data was integrated for reporting resulted in different organizations creating different “versions of the truth.”

“The problem was that our analysts needed to become experts on so many different systems just to get at the data they needed to study,” explains Peter Møllebjerg Andersen, Head of Analysis at TDC. “This problem was compounded because we work with data from more than 60 sources. No one analyst could possibly have a full systems overview.”

TDC needed a business intelligence (BI) solution that would enable the company to:

- Perform extract, transform, and load (ETL) processes on data from all sources and consolidate it into a data warehouse to provide a single “version of the truth.”
- Create multidimensional cubes to support data analysis.
- Reduce the cost of analyzing disparate data.

Solution

TDC replaced its collection of third-party reporting applications with a unified BI data warehouse created using the Microsoft SQL Server™ 2005 Enterprise Edition (64-bit)

“Our analysts are extremely happy with our SQL Server 2005 data warehouse because now they can spend their time doing analysis instead of digging for data.”

Jacob Moesgaard, General Manager, Finance Department, TDC

Complete Solution - SQL Server 2005 helps TDC gather and analyze data to generate knowledge that enhances its decision making.

database running on the Microsoft Windows Server 2003 64-bit Enterprise Edition operating system.

Called CUBUS, the data warehouse includes 6 terabytes of data hosted on two instances of SQL Server, with the largest instance holding about 3 terabytes. The data warehouse currently supports more than 2,000 users in some 332 groups. The number of users is expected to more than double over the next two years.

The solution was originally deployed using SQL Server 2000, and then upgraded to SQL Server 2005 in a migration process that TDC praised for simplicity. “The migration only required minimal code change,” says Andersen. “We were glad to see that the DTS [Data Transformation Service] packages we created with SQL Server 2000 ran fine with SQL Server 2005, meaning that we don’t have to rewrite them until we have a reason to, at which point we do so using SQL Server 2005 Integration Services. That really helped cut the time required for migration.”

The data warehouse is hosted on a three-node cluster using HP ProLiant DL585 servers, each with four AMD Opteron™ Dual-

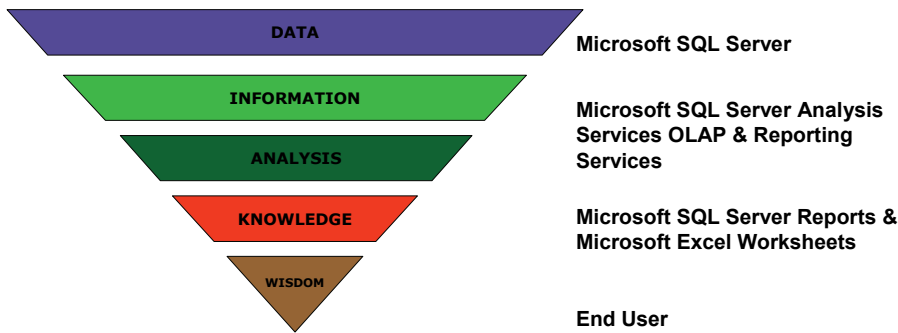
Core 64-bit processors running at 2.2 GHz, and with 32GB of RAM.

The data warehouse includes information from a wide variety of sources, including the Microsoft Business Solutions–Navision® software (now part of Microsoft Dynamics™), Oracle, SAS running on VMS, UNIX, and Windows® operating systems, IBM DB2, Tandem, VAX/VMS, and SQL Server. The data also comes from a range of groups within the organization, including human resources, financial (accounting, budgeting, capital expenditures, and operational expenditures), production, marketing, and customer service.

Solution Architecture

The new CUBUS data warehouse has a multi-tier architecture that includes:

- **ETL Tier.** Initial ETL is performed, as data from more than 60 sources is brought into a dedicated SQL Server database, to ensure that only valid data is read. Meta data is also selected at this point. ETL is performed using SQL Server 2005 Integration Services and stored procedures, as well as SQL Server 2000 DTS packages created prior to upgrading to SQL Server 2005.
- **Staging Tier.** Data is normalized, cleaned, and adjusted to match dimensions and construction of consistent data models, using more than 200 Integration Services and DTS packages and stored procedures. The ETL and staging work is performed on a dedicated SQL Server database.
- **Data warehouse Tier.** Validated data models are stored in a relational data warehouse on a dedicated SQL Server database.
- **Analytics Tier.** SQL Server 2005 Analysis Services, a middle-tier server for online analytical processing (OLAP), is used to create more than 100 multidimensional data cubes. The most complex cube, the



“SQL Server 2005 Integration Services is a powerful ETL tool.”

Peter Møllebjerg Andersen, Head of Analysis, TDC

financial reports, has 32 dimensions and 450 billion aggregations.

- **Reporting Tier.** Microsoft Office Web Components provide reporting capabilities. TDC conducts ad hoc data mining using SQL Server 2005 Report Builder. TDC is preparing to deploy SQL Server 2005 Reporting Services, a comprehensive, server-based solution for creating, managing, and delivering real-time information to support daily operations and decisions, for creating recurring reports
- **Presentation Tier.** The Microsoft Office 2003 Excel® 2005 spreadsheet software, together with the Microsoft PivotTable® dynamic views provide the primary interface tool through which users access the data warehouse. Any OLAP interface tool that is compatible with Analysis Services also can be used.

TDC's internal developers created applications supporting the data warehouse using Microsoft Visual Studio® 2005 and the Microsoft .NET Framework version 2.0. The .NET Framework is an integral component of the Windows operating system that provides a programming model and runtime for Web services, Web applications, and smart client applications.

Solution Strategies

As with other large organizations, TDC faced the challenge of gaining project acceptance and agreement on product definitions across a large number of organizations within the company. To gain project acceptance, the CUBUS team worked with high-level sponsors within each organization, with special focus on involving each group's finance department. TDC executives viewed gaining support of finance departments across the company as essential. Finance departments are primary users of data, creators of analysis, and are seen as the “keepers and enforcers of the truth”—all of which enabled

people in finance to quickly see the value of the CUBUS data warehouse.

A second challenge was to create standard definitions of data constructs, such as product, segmentation, customer, organization, and other concepts whenever possible, that could be used across the company for the data warehouse. Again, the CUBUS team worked closely with finance departments to determine what data definitions were essential to their own operations. The CUBUS team then worked between finance departments to seek agreement on common definitions that all of the organizations could use.

In cases where a group needed its own specialized definition, the exception was published with the other definitions, with the differences clearly documented to ensure that all could understand what was defined. When common definitions were found across business lines, the responsibility for maintaining these agreed-upon definitions was centralized to ensure continuity through future development and maintenance, and to promote them for use across the company.

Benefits

Uniting data from more than 60 source systems into a central data warehouse has provided TDC with a number of benefits including providing a single set of statistical data with multiple views, the ability to support more analysts with proportionately fewer resources, an 80 percent reduction in processing time, an integrated development environment, and a project return on investment (ROI) of more than 7,000 percent.

"Single Truth" with Multiple Views

The CUBUS data warehouse gives TDC one set of comprehensive statistical data that everyone can agree on. In addition, it provides a central point from which data from

“Our developers like how easily they can create a solution using Visual Studio, Analysis Services, and SQL Server. These products all integrate so easily.”

Jan Andersen, Systems Architect, TDC

dozens of sources can be analyzed through multidimensional cubes and with ad hoc data mining. TDC analysts and others within the company greatly appreciate the consistent data across the company because they have experienced how exasperating it can be when working with disparate data sources and reporting applications.

“Our analysts are extremely happy with our SQL Server 2005 data warehouse because now they can spend their time doing analysis instead of digging for data,” says Jacob Moesgaard, General Manager, Finance Department at TDC.

“We used to spend the first part of every meeting comparing data and trying to determine who had the best information to work with,” recalls Andersen. “Every group had its own analyst with their own ways of looking at data and defining products and business elements. Arguing over who has the best analyst isn’t very productive. By using SQL Server and Analysis Services to create our data warehouse, we have eliminated the need to argue over the validity of data. We now have a ‘single version of the truth.’ We can concentrate on the real task of drawing value from our common set of data.”

The planning process that the CUBUS team used in deploying the data warehouse was also helpful, especially in guiding the internal organizations toward agreeing, as much as possible, on using a common set of definitions so all groups across the company can talk about the same things in the same way.

“Product definitions are common knowledge now, so when someone needs a new report or analysis created, the product definition is already part of the cube and we can just build upon it,” notes Andersen. “This helps us create new research products in one tenth the time it used to take.”

With the data warehouse removing the need to debate the validity and accuracy of data, analysts and other users of the data were freed to explore the data looking for new insights.

“We are gaining new knowledge as a company, aggregating information to improve our business knowledge,” explains Andersen. “Employing Analysis Services and performing ad hoc data mining with SQL Server Query Builder helps us to generate information that we didn’t have before. We can deliver knowledge. And wisdom is based on knowledge.”

Ability to Support More Analysts

Building the data warehouse with SQL Server and Analysis Services has greatly simplified the task of supporting TDC’s business analysts and other data-dependent users at TDC. “SQL Server 2005 Integration Services is a powerful ETL tool,” Andersen comments. “And we like the fact that Analysis Services and Integration Services are integrated with the database. We don’t use any third-party applications for our data warehouse infrastructure because everything in SQL Server and the Windows Server operating system integrates so nicely.”

The tightly integrated infrastructure, along with the ease of use of Analysis Services and the SQL Server relational database have helped the CUBUS group to support more users with proportionately fewer staff.

“To a great extent, our analysts are self-supporting because Analysis Services is so intuitive. It even allows us to use Excel as a front end,” Andersen observes. “We used to have five full-time employees supporting 30 users. We now have 12 full-time-equivalent employees supporting more than 2,000 analysts and other users.”

“To a great extent, our analysts are self-supporting because Analysis Services is so intuitive.”

Peter Møllebjerg Andersen, Head of Analysis, TDC

Behind the scenes, analysts benefit from the ease with which new data sources can be brought into the data warehouse. TDC managers have been impressed with SQL Server 2005 Integration Services because it provides an even more robust ETL solution than did DTS.

“SQL Server Integration Services provides a more inclusive approach to working with data than we had with DTS,” comments Andersen. “In designing BI solutions you need to integrate the perspectives and needs of the business analysts. SQL Server Integration Services makes it much easier to do this. It is a very powerful ETL Tool.”

80 Percent Reduction in Processing Time

TDC has enjoyed enhanced processing power with its new data warehouse solution. “We’ve enjoyed a significant increase in our processing power without rewriting code,” says Andersen. “We’ve seen an 80 percent reduction in processing time required for data gathering and cube processing with SQL Server 2005 and the AMD Opteron™ Dual-Core 64-bit processors.”

The increased processing power is enabling the company to run analyses that used to be too time consuming. “One process involves analyzing data from 160,000 profit and loss statements, totaling more than 1.5 terabytes of information,” says Andersen. “This used to be so time consuming that it had to be carefully scheduled. Now it can be done whenever an analyst wants.”

Processing efficiency is enhanced, in part, by TDC’s use of the Table Partitioning feature of SQL Server 2005, which enables fast data loads and simplified maintenance for very large tables. The Table Partitioning gives database administrators the ability to treat multiple tables as a single entity. “We deal with a lot of complex data,” says Andersen. “Our largest table has more than one billion

rows, and when you reach that size you need something like Table Partitioning to facilitate updates.”

The enhanced processing power is especially important to TDC because of its rapidly growing data store. “We have 6 terabytes of data now, and anticipate data growth of about 30 percent per year,” says Jan Andersen, Systems Architect at TDC. “One of the reasons we chose SQL Server 2005 is that it has the scalability we need to match our growth.”

Integrated Development Environment

TDC internal developers have benefited from how seamlessly their Visual Studio 2005 and .NET Framework development environment integrate with SQL Server 2005. “Our developers like how easily they can create a solution using Visual Studio, Analysis Services, and SQL Server,” says Jan Andersen. “These products all integrate so easily. You can work within Visual Studio to create a solution that will use SQL Server Integration Services, Analysis Services, Reporting Services, Active Directory for security, and custom data types for SQL Server. Our developers love how all of the pieces fit so easily together.”

The company is planning on integrating SQL Server 2005 Notification Services feature into an extensive key performance indicator (KPI) system it is deploying across the organization. “Notification Services will give us an excellent way to deliver very precise and timely information to those who need it,” Andersen says. “KPIs provide vitally important information when tied to strategy, planning, and budgeting. If KPIs aren’t important to a company, then it’s not a data-driven company.”

Easier Data Analysis

Gathering data from more than 60 disparate systems into a unified data warehouse has

For More Information

For more information about Microsoft products and services, call the Microsoft Sales Information Center at (800) 426-9400. In Canada, call the Microsoft Canada Information Centre at (877) 568-2495. Customers who are deaf or hard-of-hearing can reach Microsoft text telephone (TTY/TDD) services at (800) 892-5234 in the United States or (905) 568-9641 in Canada. Outside the 50 United States and Canada, please contact your local Microsoft subsidiary. To access information using the World Wide Web, go to: www.microsoft.com

For more information about AMD products and services, visit the Web site at: <http://enterprise.amd.com>

For more information about TDC products and services, visit the Web site at: www.tdc.dk

slashed the cost of analyzing data at TDC. For example, TDC generates about 960 million billing records a year, while its traffic group generates at least 65 million records a day.

A TDC group needed the granularity of the individual billing records to be able to track problems and predict future development on single products and customers. But the historical data used to be stored on tape after three months, making it difficult to access.

To be able to make its predictions, the group needed two years of data. Further complicating matters, there was no direct link between financial records and billing records prior to the data warehouse. With the data warehouse in place, such data analysis has been greatly simplified.

Microsoft Server Product Portfolio

For more information about the Microsoft server product portfolio, go to: www.microsoft.com/servers/default.mspix

Microsoft SQL Server 2005

Microsoft SQL Server 2005 is comprehensive, integrated data management and analysis software that enables organizations to reliably manage mission-critical information and confidently run today's increasingly complex business applications. By providing high availability, security enhancements, and embedded reporting and data analysis tools, SQL Server 2005 helps companies gain greater insight from their business information and achieve faster results for a competitive advantage. SQL Server 2005 is designed to integrate seamlessly with your other server infrastructure investments.

For more information about SQL Server 2005, go to: www.microsoft.com/sqlserver

Software and Services

- Microsoft Servers
 - Microsoft Windows Server 2003 64-bit Enterprise Edition
 - Microsoft SQL Server 2005 Enterprise Edition (64-bit)
- Microsoft Office System
 - Microsoft Office Excel 2003
- Microsoft Visual Studio 2005
- Technologies
 - Microsoft Active Directory
 - Microsoft .NET Framework version 2.0

- Microsoft SQL Server 2005 Analysis Services
- Microsoft SQL Server 2005 Reporting Services

Hardware

- HP ProLiant DL585 servers with four AMD Opteron Dual-Core 64-bit processors running at 2.2 GHz with 32 G RAM

Partner

- AMD

© 2006 Microsoft Corporation. All rights reserved. This case study is for informational purposes only. MICROSOFT MAKES NO WARRANTIES, EXPRESS OR IMPLIED, IN THIS SUMMARY. Microsoft, Active Directory, Dynamics, Excel, Navision, PivotTable, Visual Studio, Windows, the Windows logo, and Windows Server are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries. All other trademarks are property of their respective owners.

AMD, the AMD Arrow logo, AMD Opteron, and combinations thereof are trademarks of Advanced Micro Devices, Inc.

Document published November 2006