

*Choosing a
DBMS for Data
Warehousing*

*by
William McKnight
McKnight Associates, Inc.*

Outline

Executive Summary	1
The Realities of Data Warehousing Today	2
Applications	2
Data Size	3
Real-Time	3
Architecture Requirements	3
Criteria for a Data Warehouse DBMS selection	4
The Vendor	4
DBMS Features/Functions	5
Sybase IQ	5
Major differences from other RDBMS	5
A review of the IQ Demonstration	9
Queries in the Demonstration	10
Engagement Expectations	12
Compromised Architecture Results	12
Sybase IQ's True Place in Data Warehousing	13
The Proven NCR Solution	13
Parallelism, Optimizer, Scalability	13
Real-Time Capabilities	15
TCO Leadership	15
Continued Relevancy Today	16
Reference Account Review	16
Industry Quotes	17
About the Author	18

Executive Summary

When making product decisions for a data warehousing (DW) environment, the database management system (DBMS) is the most important. Though not the most expensive component, it should be chosen with care and with active discernment over the issues and marketing messages.

You will create a culture around your selected DBMS. You will hire and train your people to support it. It will become the primary driver for hardware and other software selections. Your people will attend user group meetings and interact with others using the DBMS for similar purposes. You will hire consultancy on the DBMS and you will research how to most effectively exploit the technology. You will need vendor support and you will want the vendor to be adding relevant features and capabilities to the DBMS that are needed for data warehousing in the future.

The data warehouse DBMS will certainly have a large enough footprint in the shop to warrant a selection independent of operational technical direction. Data warehouse staff will tend to be dedicated to the data warehouse project for periods of many years. With connectivity and data movement offerings available today, a multiple DBMS enterprise environment is expected.

Data warehouses grow over time. Data volumes will soar over time as history accumulates, third-party data is added, clickstream data is added and new uses require new data. And data warehouses are not just for a handful of marketing analysts anymore. Data shows that over a forward three-year period, the amount of usable warehouse data is likely to increase by a factor of 2.9, and the number of users by a factor of 4.3.¹ With all the investment and value, you'll want and need to leverage your data warehouse for customers, supply chain partners and possibly selectively to the broader Internet. You want to make sure you choose a proven DBMS not just for the initial, known requirements but also for the future, to-be-determined requirements.

Data warehouses over time tend to share other characteristics which include a mixed, unpredictable workload and a strong tendency towards active real-time, not batch, updating. Data warehouses today provide many functions that arguably are operational in nature. It is important to understand this eventuality of the data warehouse environment when selecting the fundamental underpinning DBMS component.

With warehouses, which if successful will become very large databases (VLDB or VLDW), issues become exponentially more difficult than with smaller databases. Some of the consequences of making inappropriate DBMS selection for an active VLDW include:

- Data Quality, a major risk factor for data warehousing, is unable to be addressed adequately due to the impact on data availability
- Compromised architectures which leads to high long-term TCO
- Poor query performance which results in lowered usage
- User confidence in the effort quickly eroding (it's a slippery slope)
- Cost expansion for software, hardware and consulting to "throw at the problems"
- Inability to get to real-time feeds
- Inability to adapt to future data warehousing needs

You want to make sure you choose a proven DBMS not just for the initial, known requirements but also for the future, to-be-determined requirements.

¹ Database Solutions III, survey.com, 1999.

Choosing a DBMS for Data Warehousing

These risks are real with a selection of Sybase IQ, which performs certain queries superbly, but lacks the overall maturity of Teradata. The key differentiators are:

- Teradata's ability to achieve real-time, active data warehousing
- Teradata's proven ability to support large data, large numbers of users and complex queries and grow as your business grows in a production data warehouse environment
- The applicability, disclosure, referencability and repeatability of Teradata's claims
- Teradata's parallel abilities – the key to data warehouse performance
- The design flexibility associated with Teradata implementations and the query latitude afforded to the users
- The low cost of ownership associated with Teradata

The Realities of Data Warehousing Today

No longer an option, the collection, management and use of data is the battleground that many industries engage in today.

According to survey.com, the data warehouse market is estimated to be approximately \$150B for services, software, systems, and in-house spending by 2003. Data warehousing has proven itself over and over again by providing the data to help companies compete in the relevant area of competitive advantage today – business intelligence.

No longer an option, the collection, management and use of data is the battleground that many industries engage in today. Even in a subdued economy, data warehousing and business intelligence remain at the forefront of IT-related spending. This is thanks in large part to the applicability of information directly or indirectly to the income and expense streams within a company and the acceptance of data warehousing as the means to store the data that leads to that information.

To be sure, ROI is demanded and timeframes have shrunk for the delivery of ROI with data warehousing. Projects need to be focused without compromising the ability to leverage the work into support of other areas.

The realities of data warehousing today are:

- Multiple, complex applications will use the data warehouse serving a variety of users
- Data size is exploding and will continue to explode
- Data warehouse data latency is becoming intolerable as real-time data is demanded

Applications

Data warehousing is important to a business to the degree that it can provide either short- or long-term value to the bottom line. Many of the applications from where the return on investment (ROI) is generated today are related to the customer and the ability to attract and retain high value and high potential customers on a marketing budget that must be spent wisely.

Otherwise, data is analyzed for initiatives around fraud detection, marketing automation, workforce management, customer analysis and management, planning and analyzing manufacturing, the supply chain and the contact center and analyzing the key performance indicators of a company's performance and direction. Often, central data for key subject matter like customers, products and sales are applicable to a variety of the applications. This leads to high concurrency requirements.

In summary, data warehousing is applicable across the board to a company. There is scarcely a strategic or tactical company objective that cannot be supported with the information generated from a data warehouse.

Choosing a DBMS for Data Warehousing

Data Size

Data warehouse size is booming due to a variety of factors, but mainly because success begets success with data warehousing. As the initial data generates profitable use and the platform proves able to handle the workload, it is a matter of time until new uses leverage the data and add their different data requirements to the warehouse. And using detailed data in conjunction with summary data is important for effective decision making, further contributing to data overload.

Seldom is it feasible to delete, or otherwise render inaccessible, older data. Plan on data simply accumulating ad infinitum in the warehouse. Plan on loading all the historical data you have to seed the warehouse with as well.

Companies are also realizing the usefulness of data that is generated outside of their confines – so called third-party data. It is no longer difficult or untested to “subscribe” to external data feeds to augment internally generated data. Marketing departments in particular have grown in their sophistication to deal with all kinds of data and the more, the better.

Real-Time

More and sooner is always better. Real-time is becoming in vogue for data warehouses. If overnight loading is good, intra-day loading is better and “real time” loading is best. Data warehouses are absorbing many “operational” functions today that operational systems are unable to provide. Often, it is easier to tackle issues with a modern, supportable data warehouse than it is to change operational systems. But that does not mean the added functionality can impede the system’s operation.

Users demand summarized historical trends but they also demand real-time data. It is critical to offer business users access to critical information as business events occur, enabling them to better focus on driving performance across the business. Enable users to combine real-time data generated by transactional systems with historical time-slice information stored in the data warehouse for serious value.

Do not build non-real-time extraction, transformation and loading routines for data warehouses today that cannot be easily altered for real-time. Your data warehouse must be able to address the real-time data integration challenge.

Architecture Requirements

The trends above are rather obvious to an astute observer of the data warehousing industry. However, acknowledging the industry direction and being able to support it are two different things. Look beyond vendor hype. Look behind the eye-catching headlines for true substance to see if there is real production applicability to you.

Of the major inputs to a data warehouse architecture specification and DBMS selection, eventual data size is one of the most important. At certain levels of data size (in the terabytes), data warehouse programs tend to gravitate towards similar levels of usage and complexity. There are programs in the multiple terabytes today with thousands of users and this represents the eventuality of many data warehouses starting out today.

Make sure your DBMS selection can accommodate a true eventual production data warehouse environment.

Your data warehouse must be able to address the real-time data integration challenge.

Ideally the DBMS selection should be the first technology decision made for a data warehouse project.

The vendor's financial stability, the importance of data warehousing to their overall business strategy and their continued research and development in the area of data warehousing towards a well developed and relevant vision are all key components of a vendor's viability in this critical decision.

Criteria for a Data Warehouse DBMS Selection

The data warehouse DBMS selection is critical and acts as a catalyst for all other technology decisions. The technology needs to support both the immediate as well as future, unspecified and unknown requirements. Ideally the DBMS selection should be the first technology decision made for a data warehouse project.

Given the state of the marketplace, described above, the technical architecture should be:

- **Scalable** – In both performance capacity and incremental data volume growth. Make sure the proposed solution scales in a near-linear fashion and behaves consistently with growth in all of database size, number of concurrent users and complexity of queries. Understand additional hardware and software required for each of the incremental uses.
- **Powerful** – Designed for complex decision support activity in a multi-user mixed workload environment. Check on the maturity of the optimizer for supporting every type of query with good performance and determining the best execution plan based on changing data demographics. Check on conditional parallelism and what the causes are of variations in the parallelism deployed. Check on dynamic and controllable prioritization of resources for queries.
- **Manageable** – Through minimal support tasks requiring DBA/System Administrator intervention. It should provide a single point of control to simplify system administration. You should be able to create and implement new tables and indexes at will.
- **Extensible** – Provide flexible database design and system architecture that keeps pace with evolving business requirements and leverages existing investment in hardware and applications. What is required to add and delete columns? What is the impact of repartitioning tables?
- **Available** – Supports mission critical business applications with minimal down time (a fading concept). Check on “hot pluggable” components, understand system down time requirements and any issues that might deny or degrade service to end users. These can include batch load times, software/hardware upgrades, severe system performance issues and system maintenance outages.
- **Interoperable** – Integrated access to the web, internal networks, and corporate mainframes.
- **Affordable** – Proposed solution (hardware, software, services, required customer support) providing a low total cost of ownership (TCO) over multi-year period.
- **Proven** – You don't want to risk a critical decision regarding a fundamental underpinning of the data warehouse environment on an unproven solution.
- **Flexible** – Provides optimal performance across the full range of normalized, star and hybrid data schemas with large numbers of tables. Look for proven ability to support multiple applications from different business units, leveraging data that is integrated across business functions and subject areas.

The Vendor

There are few vendors who understand what it means to build production, mission critical, data warehouse systems.

This decision process should go well beyond the usual feature/function comparisons that are done. The vendor itself should be a major category in these days of vendor consolidation among data warehouse DBMS. The vendor's financial stability, the importance of data warehousing to their overall business strategy and their continued research and development in the area of data warehousing towards a well developed and relevant vision are all key components of a vendor's viability in this critical decision.

Marketplace skill sets and the availability and interest level of resources to be associated with the technology factor in as do the availability of reference accounts that are more

Choosing a DBMS for Data Warehousing

mature than your data warehouse and are actually exploiting, in production, new technical capabilities and new uses for data warehousing.

Vendors put their sales and technical teams through their school. They become used to certain types of data warehouses. Make sure there is enough diversity and number amongst the vendor accounts for there to be one that may be similar to your eventual program. There may not be one single reference that matches your environment exactly but when you see a consistent trend across a wide range of references then there is greater confidence in the capabilities of the vendor and their technology. Also make sure the school and the vision are relevant and backed by R&D.

DBMS Feature/Functions

The DBMS features/functions are indeed a consideration. Some considerations include:

- DBA Productivity Tools that improve the productivity of a DBA in the solution environment
- Monitoring Features and capabilities including APIs supported, interfaces and integration with the hardware and operating system, SQL query tracing and journaling capabilities
- Database utilities with particular attention to their parallel execution abilities, restartability, ability to get granular (like doing a backup at a table, not complete database, level) and impact on availability when running.
- Support for a variety of end user tools (of which there will be several in place in due time in a robust data warehouse environment)
- Performance as indicated by a robust optimizer, indexing schemes, concurrent user resource consumption, DBMS overhead
- Locking and concurrent operation schemes
- Security methodology – standard compliance, row level security
- Intra-query parallel implementation for all possible access paths
- The ability to add on components and have them fully benefit the environment
- Chargeback and accounting features and remote maintenance capabilities
- Support for a variety of schemas and data types
- Etcetera

This list is a starting point for any such DBMS evaluation. A proper evaluation must go beyond the performance of known queries.

Sybase IQ

The Sybase Adaptive Server IQ Multiplex database, commonly referred to as simply Sybase IQ, is addressed here in context of data warehousing since this is the Sybase product referenced in the Gartner Data Warehouse DBMS Magic Quadrant Update dated April 29, 2002 and is the main product sold by the Sybase Business Intelligence Division. Sybase acquired the Expressway DBMS from Expressway Technologies, formerly Henco Software, in October 1994 and renamed it Sybase IQ in February 1996. The Multiplex form of Sybase IQ has been generally available since June of 1999.

Major Differences from other DBMS

Sybase IQ has several major architectural differences from other relational database management systems. The main difference is its physical orientation of data in columns as opposed to rows. This allows it to perform very high selective compression because all of a column's values are physically together. It also provides for excellent performance when you select a small subset of the columns in a table since you do not perform I/O for data that is not needed. See Figure 1.

Sybase IQ has several major architectural differences from other relational database management systems. The main difference is its physical orientation of data in columns as opposed to rows.

Choosing a DBMS for Data Warehousing

Figure 1

Overall Assumptions:

1. Customer record is 500 bytes
2. 10,000,000 customers
3. 8 customer records per 4K page (row-oriented systems)¹
4. Of the 4K bytes on a page, 4000 bytes is usable space
5. 1,250,000 pages containing customer information (row-oriented systems)
6. Date fields are 4 bytes, First Name is 25 bytes, Last Name is 25 bytes, Email Address is 25 bytes

Example 1 - Table Scan

Queries that lightly filter data and need to scan an entire table of values like..

Select first_name, email_addr from customer where month(date_of_birth) = 7

In this example, 3 columns in one table are needed.

Row-wise storage doesn't care about the column count: Table Scan: 1,250,000 pages

Column-wise storage retrieves only the relevant columns:

1. Date of birth column vector (10,000 pages)
2. First_Name column vector (62,500 pages)
3. Email_Address column vector (62,500 pages)

The total pages read is 135,000 pages.

Note that intra-query parallelism could dramatically reduce the runtime of the query either way. With 20 parallel streams, which many row-oriented systems would have, the query would actually run in the same amount of time in row-wise storage if the column-wise storage did not have intra-query parallelism but read all 3 columns in parallel.

Example 2 - Single Record Lookup, 20 columns returned

Select <all information on a customer> from customer

Row-oriented

1. Index lookup – 1-3 pages (although usually the upper levels of an index are cached)²
2. Table Access – 1 page

Total pages read is 4 pages maximum

Column-oriented

1. Index lookup – 1-3 pages
2. Table access – 20 pages (1 page per column)

Total pages read is 23 pages.

Row-oriented systems are better for single row lookup, especially when many columns are needed.

¹ Page sizes can range from 2K to 32K

² It would be 1 I/O in Teradata due to hash lookup.

Choosing a DBMS for Data Warehousing

Example 3 – Simple Join

Get customer first name, last name, address1, address 2, city, state and zip for all customers in the state of Oregon for a mailing.

Select first_name, last_name, address1, address2, city, state_code, zip from customer a, state b where a.state_code=b.state_code and b.state_name = 'Oregon'

Assumptions

1. There is 1/50 of the total customer base in Oregon

Row-wise storage further assumptions:

2. The index on customer.state is appx. 6 bytes (2 for state ID and 4 for the table location)
3. Appx. 666 index records on a page
4. One index entry per customer record so there would be appx. 15,000 index pages

Column-wise storage further assumptions:

Column Name	Length (Bytes)	Records/Page	# Pages
First_name	25	160	62,500
Last_name	25	160	62,500
Address1	50	80	125,000
Address2	50	80	125,000
City	25	160	62,500
State	2	2000	5,000
Zip	9	444	22,523

Row-wise storage:

1. Read the State table to get the state ID for "Oregon" (small, 2 pages)
2. Take this ID out to the customer.state index and read appx. 300 index pages (15,000 / 50) to get the locations of the Oregon records in the customer table
3. Go to the customer table with this list and read the customer records – appx. 250,000 pages (1,250,000 / 50)

The total pages read is $2+300+250,000 = 250,302$. Alternatively, some RDBMS provide the option of ordering the table ("clustering") according to 1 or more columns. If the table was clustered by state ID, the 300 index pages would not have to be read. And instead of 250,000 table pages, you'd read 31,250.

Column-wise storage:

1. Read the state.state_name column vector to locate the state ID for "Oregon"
2. Read the state.state_id column vector to locate the state ID for Oregon customers (1 page)
3. Go to the customer.state column vector to get the record IDs of all the Oregon customers (up to 5,000 pages)
4. Scan each of the customer first name, last name, address1, address 2, city, state and zip column vectors and retrieve the information on the customers in the record ID set for Oregon customers
5. Put the records together

The total pages read for just step 4 is 465,023.

While column-wise storage may be compressed 50% (with commensurate decompression cost at read-time), the row-wise storage may be compressed 25% if desired.

The differences only get more pronounced the more joins, predicates and columns retrieved get involved.

Choosing a DBMS for Data Warehousing

Column-orientation greatly assists a compression strategy due to the high potential for the existence of similar values in columns of adjacent rows in the table.

The Sybase IQ approach to DBMS data management is not anything new. In fact, the venerable mainframe database Model 204 uses a similar approach. It is an interesting approach to solving the complex, ad-hoc query requirements of data warehousing but let's take a look at the major points of uniqueness.

The Multiplex option gives Sybase IQ the ability to have multiple servers (not end user PCs or thin clients) accessing the disk system. There is one write server for data loading and updating. This server can be combined with one or more read servers in the environment for query processing.

In a Sybase IQ-Multiplex environment, only the single write node can perform updates. You must designate this node when you start the system. Any user or application which needs to perform an update must log on to that particular node. Sybase IQ requires the replication of catalog information from the write operation to reader servers.

Compression is also a key Sybase IQ capability. By most accounts including the demonstration referenced in the next section and a Winter Corp. paper,² Sybase IQ can compress data by over 50%. Column-orientation greatly assists a compression strategy due to the high potential for the existence of similar values in columns of adjacent rows in the table. 50% compression is relatively high by comparison to other DBMS.

Compression is a trade-off that comes with several prices, most notably the time required to decompress and the high use of memory. The disk storage saved by compression has been rapidly lowering in price over the past several years as well. Other DBMS products have compression capabilities as well. For example, there are numerous examples of Teradata customers achieving 25% or higher compression.

Sybase IQ also has many different index types. It is recommended and common practice to index all columns at least once and, for some columns, more than once. A production data warehouse system would therefore have thousands of indexes.

Sybase IQ does not lock a table to ensure data integrity but instead uses an approach called versioning. When a transaction begins, a marker is placed on rows that are inserted, deleted, or changed. A query sees the data as it was when the query was initially submitted. Once the transaction's changes are committed, new queries will see the changes.

In Sybase IQ Multiplex, this situation is complicated because each database server maintains its own catalog or data dictionary. This catalog is stored in a different Sybase database product, either Adaptive Server Anywhere or Adaptive Server Enterprise. Only one node actually performs the load. When the load completes, the changes are replicated to the catalog on the other nodes. Only then will the other nodes become aware of and begin to read the new data. The delay between the time a write server completes an update action and when the query servers are aware of the new data can be substantial.

As regarded as Sybase IQ's capabilities are in supporting certain queries, there is a trade-off. That trade-off is in the area of data loading and "online transaction processing." The impact of this architecture trade-off relegates it to batch-loading abilities. Sybase IQ is dependent on Sybase Adaptive Server Enterprise (ASE) to handle general database processing and to replicate the data, implying both data latency to the query nodes and data redundancy.

² "Efficient Data Warehousing for a New Era," Winter, Winter Corp., 2000

Choosing a DBMS for Data Warehousing

As Sybase says, “ASIQ is not an OLTP system. It is not designed to do single-row inserts very fast. It is designed to load mass quantities of data via load table.”³ Unfortunately, labeling single-row inserts as an OLTP function is no longer appropriate. In today’s data warehouses, many need to be real-time with loading and inserting of a few rows frequently as opposed to infrequent loading of a batch. A key part of Teradata’s Active Data Warehouse initiative is that you can perform online updates and that you can stream data continuously into the system.

Once the database objects are set in place in Sybase IQ, changing them frequently requires the unloading and reloading of data. This is prohibitive to rapid application development environments as well as to making rapid changes to production. As with the queries, you need to get it right the first time or pay the price with Sybase IQ. This “big bang” approach has long ago been proven as the improper means to develop a data warehouse. Technology components should not restrict the ability to develop incrementally and iteratively.

Design setup also requires manually specifying data demographics such as cardinality for use by the optimizer. This is fine for most columns, where the cardinality and distribution of data doesn’t change much. But data demographics do change frequently for many columns and these dynamic columns are often the most queried in a data warehouse.

Parallelism is key to query performance success. Sybase IQ offers no intra-query parallelism, and the column vector parallel processing (the ability to scan multiple columns of the same table concurrently) that is offered is not guaranteed, but only implemented if the optimizer can determine that sufficient resources are available. In Teradata, all queries run in parallel. With regard to parallel loading, Sybase IQ only offers parallelism in the inserting and deletion of rows via the loader utility to the write server. In Sybase IQ, each query only runs within a single node so adding more nodes will not speed up a slow query.

A Review of the IQ Demonstration

Sybase recently loaded 48 TB of input data into a Sybase IQ database. They have rightfully compared this to 5 times the size of the U.S. Library of Congress. The machines Sybase IQ were on are very capable Sun Microsystems servers – the same set of Sun machines which partner with IBM’s DB2 DBMS and Oracle’s DBMS for TPC benchmarks. From a disclosure standpoint, five queries were run on this database as well.

As discussed above, data warehouses are more than the data put in the database. There are users, load cycles, complex queries, unpredictable queries, reports, data mining, etc. If reasonable, disclosed qualities as to be found in a production data warehouse environment could be put around a Sybase IQ database of 48 TB and it performs well, the data warehouse industry should quit ignoring Sybase IQ and take note!

Actually there is an opportunity for Sybase IQ to participate under disclosure – and at a lower level of scale and risk. That is with the TPC⁴ benchmarks. Specifically the 1 TB and 3 TB tests for TPC-H, which are opportunities to showcase abilities in data warehousing. These are the industry standard benchmarks for decision support. They test the simplest forms of data warehousing and decision support. The schema is very simple (8 tables with no data skew.) The very limited set of 22 queries is moderately complex at best and have been well understood and unchanged for years. Given the simplicity relative to real-world data warehousing, this should be an opportunity for Sybase to excel.

A key part of Teradata’s Active Data Warehouse initiative is that you can perform online updates and that you can stream data continuously into the system.

³ “Replicating Data into Adaptive Server® IQ with Replication Server: A Sybase® Adaptive Server® IQ Whitepaper,” Mumy, Sybase Business Intelligence Division, 2001.

⁴ Transaction Processing Performance Council, <http://www.tpc.org>

Choosing a DBMS for Data Warehousing

The whole reason for the existence of the TPC was the need for standardized ways to compare the performance of different products. The TPC benchmarks provide a way to (more or less) directly compare the performance of database products on a common workload. Instead of vague claims of being “10-1000x” faster than other databases, a TPC-H benchmark would allow customers to see Sybase IQ’s true performance.

Although Sybase has TPC-H results, they will only release the results under NDA. As a result of this failure to publish, Sybase IQ’s benchmark results are not officially recognized by TPC. TPC benchmarks require full disclosure and performance comparisons to other products.

The tables in the Sybase demonstration are built using a telecommunications-type schematic. The 40-column fact data is divided into 16 tables, one each for a quarter of data making a total of 4 years of data. There are a total of 10 dimension tables, although none were used in the queries. Each query accessed only one of the quarterly fact tables and performed only either a sum or average (sum divided by count) on only one column of the table.

Queries in the IQ Demonstration

The applicable outputs of the test were query performance times.

The queries were as follows:⁵

```
select sum(PRE_BAL) from CALL_DETAIL_RECORD_1999Q3
```

```
select avg(AIR_DUR)
from CALL_DETAIL_RECORD_2000Q1
where APPLY_TIMESTAMP not in (
'2000-11-10', '2000-11-12', '2000-11-13'
, '2000-11-14', '2000-11-15', '2000-11-16'
, '2000-11-17', '2000-11-18', '2000-11-19'
, '2000-11-20', '2000-11-21', '2000-11-22'
, '2000-11-23', '2000-11-24', '2000-11-25'
, '2000-11-26', '2000-11-27', '2000-11-28'
, '2000-11-29', '2000-11-30', '2000-12-1'
, '2000-12-4', '2000-12-5', '2000-12-6')
```

```
select avg(TOTAL_CHG) from CALL_DETAIL_RECORD_1998Q1
```

```
select sum(AIR_DUR)
from CALL_DETAIL_RECORD_1999Q1
where PRIM_CEL not in ( '0', '268' )
```

```
select avg(DISC_CHG)
from CALL_DETAIL_RECORD_2001Q1
where SEC_CC not in ( '055', '039', '041', '049', '044', '033', '034', '351')
```

⁵ “Sun-Sybase DW Reference Architecture,” Raab, Info Sizing, 2002.

Choosing a DBMS for Data Warehousing

These ARE arguably representative queries from a data warehouse environment. However, in a very short time queries like these will become a very small, uninteresting part of the overall workload and will likely offer little or no business value. When end users are forced to run simple queries like this, then providing answers to real-world business questions will likely require dozens of these queries and considerable additional processing to combine results.

If these queries are part of an evaluation of a real data warehouse production environment running ad hoc queries, about 45 more queries would be needed to round out a representative environment with these representing 5 of the simplest queries.

The most glaring void with the queries is that there are no joins. Most data warehouse queries are joins. As a matter of fact, most data warehouse queries are, as a data warehouse Director recently said to me, “atrocious.” Business Intelligence needs are rarely simple.

Given the manner in which Sybase IQ stores data as columns, minimal I/O is required to retrieve all the values in a column and therefore aggregate functions such as these will perform very well.

Loads have minimal impact on the queries, which is good, but the readers will not see the loaded data until after the aforementioned versioning, potentially several hours after a real-time data warehouse would.

Although the queries access a very minimal percentage of the database making the size of data in the database irrelevant, actual disk used is 22 TB, signifying a healthy compression ratio of .46, or 54% space reduction, from the input data of 48 TB. How this shrinkage occurs is not disclosed. Compacting delimiter bytes and other “noise” input fields and compacting text into binary are ways to shrink input files without compression. Also for the queries referenced above, indexes would not need to be built per the Sybase recommendation and usual Sybase IQ implementation of one or more indexes per column. Far fewer would suffice. Likewise, a minimum number of actual columns would have needed to be loaded to support the 5 queries.

The queries represent what I’ve said above that Sybase IQ will perform well at – comparatively speaking – full scans of a single column of values. However, until fuller disclosure occurs, the test is useful to the data warehouse community as a partial verification of Sybase IQ’s theoretical database size limit and a confirmation of its good performance for simple queries in a highly controlled environment. By comparison, Teradata’s theoretical limit is well above a petabyte with current hardware technology. But much more important are the data sizes Teradata supports in mixed workload, active production environments, which will be addressed below.

The time required for creating, load preparation, loading and indexing is not disclosed. Daily maintenance loads, the ones that you live with every day, were not discussed in the accessible test writeups.

Based on this test, Sybase’s Vice President for Corporate Development Marty Beard states “Sybase no longer has to walk away from large, multiterabyte data warehousing deals.”⁶ You be the judge.

Given the manner in which Sybase IQ stores data as columns, minimal I/O is required to retrieve all the values in a column and therefore aggregate functions such as these will perform very well.

⁶ “Forty-eight Terabytes and Counting,” Forsyth, Sybase Magazine, 2Q02.

Choosing a DBMS for Data Warehousing

Engagement Expectations

Sybase has won awards for its technical support. Additionally, Sybase customer retention and renewal are very high. However, it has the lowest mean satisfaction rating of all BI/DW DBMS vendors⁷ (NCR has the highest).

When engaging Sybase IQ, Sybase admits that you should know the nature of the queries beforehand: “You should know in advance how data in the columns will generally be queried.”⁸ I disagree that you should know the queries before building the data warehouse. A data warehouse is designed to be interactive and allow for unrestricted access to the data. A data warehouse should support known and on-the-fly queries. After all, do you know the questions that your users are likely to ask 2 years from now? Five years from now?

A data warehouse is designed to be interactive and allow for unrestricted access to the data. A data warehouse should support known and on-the-fly queries.

My experience is that business users do not know what queries they will need answers to in 6 weeks! In fact, it is often the case that within days or weeks of a new DSS system being delivered (one that has been designed around a finite set of queries/data and is usually application specific), the end users will realize that they need something different and something more.

Expect any “OLTP”-oriented required work to be done in Sybase ASE, not Sybase IQ, for better performance reasons. You will usually have an ASE database with Sybase IQ to handle the workload that’s not meant for Sybase IQ. Keep in mind the OLTP-nature of many data warehouses today. Data warehousing is not just for batch-loaded month-end reporting anymore.

In case you are considering upsizing within Sybase’s product line, keep in mind Sybase IQ was not engineered together with ASE. They are completely different products with limited code and SQL syntax portability. Making code ports within these products, 2 products in Sybase’s own product line, is comparable to porting between any 2 other DBMS. This would be a good time to re-evaluate your data warehouse DBMS.

Also Sybase IQ relies on a different application for error handling. That application has major error handling challenges that have been known to cause problems.

Finally, Sybase IQ is a specialized DBMS with limited ISV support and a different architecture. And row-oriented DBMS offer much more portability of skills between DBMS. Expect to use specialists from Sybase with Sybase IQ.

Compromised Architecture Results

Limitations in DBMS capabilities will impact your ability to adhere to the architecture of your choice. This will drive up your support costs as you begin to work within the limitations of the DBMS by down-scoping the data and workload given to any single DBMS instance and creating data marts and sharing data across them and the data warehouse much sooner than you planned for.

Alternatively, a costly DBMS-replacement project could be undertaken that would include a DBMS evaluation, parallel production runs, staff training and code migration – all expensive. Get this important decision right the first time.

⁷ Database Solutions III, survey.com, 1999.

⁸ Sybase IQ’s Administration and Performance Guide Version 12.4.3 page 167, Sybase Corporation, 2001

Choosing a DBMS for Data Warehousing

Sybase IQ's true place in data warehousing

Unlike Teradata, Sybase is not focused on data warehousing. And unlike Teradata, Oracle, Microsoft or IBM, the major players in data warehousing, Sybase does not have major dependencies within its product set on its DBMS. Sybase's main focus is its portal, mobile and application development solutions.

Based on the "Consolidated Statement of Operations by Segment for Q2 2002 for Sybase, Inc.," the Business Intelligence Division had an operating loss of \$3,141,000 United States dollars, based on licensing fee revenues of \$261,000 United States dollars.

Sybase (other than BI) continues to show revenue and profit increases year over year while acquiring new companies, maintaining a healthy cash reserve, and continuing to carry no debt.

The viability of Sybase IQ in data warehousing, however, is as a specialist data mart useful when the performance of known, simple, non-concurrent queries is the overriding selection factor. But this is like saying it's a good thing that 2 of the 12 eggs in the carton aren't broke.

Sybase IQ is a reporting tool that performs well on quick searches. The basic architecture is read-only with every column indexed and a high degree of data compression and replication. The Sybase replication technology is strong and they have done an especially good job of getting data to the read servers. They however, are not closing the loop and providing strong replication from the read servers back to the OLTP systems that feed them.

The viability of Sybase IQ in data warehousing, however, is as a specialist data mart useful when the performance of known, simple, non-concurrent queries is the overriding selection factor.

The Proven Teradata Solution

Parallelism, Optimizer, Scalability

All database functions in Teradata are done in parallel all of the time. These functions include: Table Scan, Index Scan, Indexed Read, Join, Sort, Row Redistribution, Parse, Catalogue Read, Index Creation, Insert, Delete, Update, Load, Backup and Restore. All units of parallelism participate in each database action. There is no conditional parallelism within Teradata.

Of special note are the following three database functions:

- **Table Scan** – Not only do all table scans run in parallel under Teradata, but Teradata also employs a technique called synchronous scan. If query 1 starts a scan of table A and query 2 needs to scan table A, query 2 will "piggy back" on query 1's scan. When query 1 finishes, the scan continues at the beginning of the table for query 2 until it reaches the point where it originally picked up the scan. During this process other queries that need to scan table A will pick up the scan in progress and drop off when they have scanned the entire table.
- **Insert** – Whenever a query requires a transient or temporary table to be built (in preparation for a join etc.) that table is populated via inserts. If an RDBMS does not do inserts in parallel all of the time, there can be serious query performance implications.
- **Parse Step** – If a query needs to scan table A and Table B and Table C prior to a join step that joined all three tables, Teradata ensures they are done at the same time. Some RDBMSs would first scan A, then scan B, then scan C. Although running these scans which have no real dependency on one another serially does not seem rational, at small data volumes it might not adversely impact performance

All database functions in Teradata are done in parallel all of the time.

Choosing a DBMS for Data Warehousing

much. But what if Table A contains 500,000,000 rows, Table B contains 350,000,000 rows, and Table C contains 100,000,000 rows? Teradata's intelligent optimizer will run steps that can be run together in parallel. In the above example, Teradata would start all three scans in parallel. When scans of tables B and C finished it would begin the join step as the scan for table A finished.

Teradata's optimizer is grounded in the knowledge that every query will be executing on a massively parallel processing system. It knows the system and understands that the amount of parallelism is the same for all operations. It has row count information and knows how many of a table's rows each parallel unit will be managing so I/O can be costed more effectively, and considers the impact of the type and number of CPUs on the node. The optimizer calculates ratios, which compare the number of parallel units on a node against the available CPU power. It uses this information to build an efficient query plan.

You don't need the fastest node with Teradata because you will be achieving high performance through adding nodes.

The NCR BYNET, which scales linearly to 512 nodes, and has fault tolerant characteristics and was designed specifically for this parallel processing environment.

You will need a scalable and highly available node-to-node interconnect. The NCR BYNET, which scales linearly to 512 nodes, and has fault tolerant characteristics and was designed specifically for this parallel processing environment. The BYNET is a linearly scalable, high performance, fault tolerant, self-configuring, multi-stage network. It supports guaranteed delivery, point-to-point and broadcast connections at the hardware level. It also provides direct access to high performance messaging services with rich additional functionality for parallel and distributed processing applications.

You will need highly available hardware components. You will depend on your data warehouse for reliable, highly available, mission-critical information. Your hardware must be configured as a high availability cabinet system to meet these needs. The cabinet features many points of redundancy, either via dual chassis or dual components in an individual chassis. Examples include dual BYNET switch chassis, dual BYNET cards, dual UPS, dual AC power feed capability, hot-pluggable fans and disks and redundant power supplies in all chassis.

Hot-pluggable components allow you to replace components without affecting your applications. If a component fails, built-in redundancy allows the application to continue running. Depending on the criticality of your solution, by adding internal uninterruptible power supplies and additional levels of redundancy, you can enhance the reliability features.

To grow a Teradata system from 100 GB to 500 GB or 2-4 TB or more is simply a function of adding nodes and disk storage.

To grow a Teradata system from 100 GB to 500 GB or 2-4 TB or more is simply a function of adding nodes and disk storage. The database only needs to "know" of the additional hardware resources to utilize them. Once new hardware has been added to the system, the NCR Support Technician assigns new AMPs and then issues a "PDE reconfig" command to automatically distribute the data.

Choosing a DBMS for Data Warehousing

Real-Time Capabilities

The Active Data Warehouse concept at Teradata is a focus on the timeliness of data. Real-time, active, data warehouses, take this to the maximum by eliminating data latency. Teradata databases are designed for maximum concurrent applications including historical trend queries and real-time loading.

Teradata manages contending requirements for resources through dynamic resource prioritization that is customizable by the customer. When the system is lightly used, those uses, high priority or not, receive heavy allocations. When new jobs appear with high priority, the allocations change according to the specified priority scheme. Incoming data feeds and computations that affect customer interactions – while the customer is still in the store, on the phone or on the website – can be performed quickly, ensuring the maximum effectiveness with the interaction.

Queries and other requests can be scheduled or can be conditionally executed based on the current workload.

Teradata also supports a host of other performance enhancement features that enable it to process queries quickly and keep the information factory flowing. Examples include join indexes which pre-join data across tables and automatic maintenance of summary tables.

Built on fault-tolerant hardware, Teradata is designed to withstand component failures. Its fallback capability keeps replicated data on physically separate disks, keeping data available and recovery automatic.

Continual feeding without table-level locks with Teradata utilities can be done with multiple feeders at any point in time. And again, the impact of the data load on the resources is customizable. The process ensures no input data is missed regardless of the allocation.

TCO Leadership

According to Gartner research,⁹ “Teradata is the most efficient DBMS in handling complex and diverse workloads compared to DB2 and Oracle.” Sybase IQ was not mentioned in the report. Furthermore it states “because of the sophistication of Teradata’s query optimizer, less DBA effort is required for the data model design with concern for performance requirements diminished” and “many of the Teradata utilities create a self-tuning environment (e.g., automatic defragmentation) that requires significantly less database maintenance.”

Meta Group reports “Informix and Teradata DWs seem to require the lowest relative project support staffing levels.”¹⁰

In Teradata, data placement, free space management, data partitioning, data reorganization, index reorganization, workspace management, query tuning and workload management are automatic. When you fail to look beyond hardware, software and maintenance when considering TCO, and fail to look at administration, management and risk issues and the far less need for tuning, you miss out on significant components of TCO.

Incoming data feeds and computations that affect customer interactions – while the customer is still in the store, on the phone or on the website – can be performed quickly, ensuring the maximum effectiveness with the interaction.

In Teradata, data placement, free space management, data partitioning, data reorganization, index reorganization, workspace management, query tuning and workload management are automatic.

⁹ “Data Warehouse Administration: TCO ‘Rules of Thumb’”; K. Strange, Gartner, 11/1/2001

¹⁰ “Data Warehouse DBMS Support Costs Either an Arm or a Leg”; Doug Laney, Meta Group, 6/19/2000

Choosing a DBMS for Data Warehousing

By comparison, a recent white paper¹¹ defined and talked about the need for low TCO. In regards to Sybase IQ, the paper stated the message of Sybase IQ's ability to respond to queries in seconds and minutes instead of hours with OLTP-based systems, physical storage requirements reduction, the lack of necessary tuning and linear scalability. All definitely are indicators of TCO. And OLTP systems truly are notorious for poor query performance. Selected queries on Sybase IQ can beat those OLTP run times, but this is not TCO.

Continued Relevancy Today

As a testament to Teradata's continued relevancy in data warehousing, Teradata also is the DBMS platform for 6 of the 13 winners in the 2002 Data Warehousing Institute (TDWI) Best Practices program. Nominations were judged by a panel of experts, including practitioners, consultants, industry analysts, journalists, Journal of Data Warehousing editors, Institute Fellows, and TDWI management.

The past several years have seen tremendous investment in Teradata's data warehouse offerings. Teradata also continues to invest in areas of relevancy to data warehouses today and the future, guided by a vision from some of the industry's most well-respected individuals.

Reference Account Review

To complete the research for this paper, I spoke with the Director of a Global Enterprise Data Warehouse for a multinational conglomerate with arguably the most complex data warehouse in the world, to get an up-to-date confirmation of Teradata's capabilities.

Here is some background on the implementation and some excerpts of that discussion:

- The data warehouse has 6000 tables on 15 TB of disk with 700 concurrent users running sales and marketing reports, all management committee reporting, feeds to demand planning, sales and operations planning, demand management, etc. The whole functions of the supply chain and marketing are enabled with the data warehouse. Every business event is in the data warehouse.
- Most queries are very complex multi-table joins. 100% of the queries are joins. 100% of the queries access more than one column.
- They currently load 4 times per day. The global warehouse is loaded when the region is off-hours. They will soon have potentially many updaters at the same time as they move to real-time. Some applications today update the data warehouse directly.
- They chose Teradata based on a test that was performed on it and other leading data warehouse DBMS. Teradata was the only one that could finish the test and Teradata did it without changing one SQL. According to the Director, Teradata "scales, scales and scales." And it doesn't require disk fragmentation, index reorganization, etc. which, if these things were necessary to do, would cause a question about the ability to get to real-time.
- If doing a DBMS evaluation for data warehousing today, he would use standard business queries that users would be using day in and day out. Single-table queries are not reality and would not even be part of the test.
- This company doesn't do data compression in the data warehouse. It is not important to them because, like many companies, they are not disk constrained but are CPU constrained.
- They specifically pursued a DBMS that could handle all of their data in a central data warehouse repository. They see no need to distribute the warehouse since Teradata is available. Since they specifically pursued a central data warehouse strategy for cost containment and data sharing reasons, the consequences of a DBMS that could not handle the centralized strategy would be unfathomable.

As a testament to Teradata's continued relevancy in data warehousing, Teradata also is the DBMS platform for 6 of the 13 winners in the 2002 Data Warehousing Institute (TDWI) Best Practices program.

¹¹ "Total Cost of Ownership – TCO," Hendershot, AMG, Inc., available at <http://www.sybase.com/tcwhitepaper1>, 2002.

Choosing a DBMS for Data Warehousing

- When adding to the data warehouse, they interview users to understand the correct business relationships. But they do not try to get to the specific user queries nor do they nearly know at design time all queries that will eventually be run.
- Furthermore, if a vendor didn't participate in TPC benchmarks, then he would consider that they didn't have what it takes to perform the benchmark and he wouldn't consider the vendor for data warehousing.
- Lastly, it is not acceptable in this environment to force off users and restart the system when you need to make a small change such as adding a view.

Industry Quotes

A side-by-side comparison using the criteria list in Criteria for a Data Warehouse DBMS Selection section validates the overall superiority of Teradata over Sybase IQ. The industry agrees.

In closing, please note the following quotes from industry.

- “In my opinion, the database vendors that play in the VLDW space today appear to stack up (at lowest TCO to a highest TCO as follows): NCR Teradata, Oracle 9i, IBM DB2 UDB, Sybase IQ, and Microsoft SQLServer.”
– *Daniel Linstedt is Chief Technology Officer of Core Integration Partners, Inc.*
- “Everyone knows that NCR's Teradata subsidiary and IBM own the VLDB market.”
- “The evidence of its capacity to start small and scale to VLDB propositions, Sybase insists, puts IQ Multiplex ahead of its competition, with the possible exception of NCR.”
– *Chris Forsyth, Executive Editor, Sybase Magazine, 2Q02*

Based on a multipart questionnaire to identify the world's leading databases with participation in over 30 countries worldwide, Winter Corp.'s Database Scalability Program reports:

- “Teradata customers ranked first among all commercial, non-proprietary sites. It manages more data and processes more concurrent queries than any other DBMS.”
- “In two categories, Peak Workload, All Environments and Unix Only, the first, second and the third most powerful, non-proprietary DBMS are Teradata sites.”
- “Within the list of Top Ten respondents for all six categories, there were more Teradata users than users of any other DBMS.”
- “The data show conclusively that the Teradata Database is the product of choice when an extremely large and heavily used data warehouse is needed.”

The Gartner Group Magic Data Warehouse DBMS Quadrant Update dated April 29, 2002 reported:

- “...Teradata ahead of the other database management system products, Oracle and Database 2 Universal Database, in the 'ability to execute' positioning.”
- “Teradata continues to lead the market in technical capability.”
- Sybase can be found in the Niche Players quadrant with this commentary: “Sybase is here primarily because of its Sybase IQ product and its appropriateness for certain types of applications, such as database marketing where segmenting large customer lists is based on a few elements of criteria.”

About the Author

William McKnight is founder and president of McKnight Associates, Inc., a consulting firm specializing in data warehousing solutions. William is an internationally recognized expert in data warehousing and CRM with more than 15 years of experience architecting and managing information and technology services for G2000 organizations.

William is a frequent and highly rated speaker at major worldwide conferences and private events, providing instruction on customer intimacy, return-on-investment, architecture, business integration, and other business intelligence strategic and architecture issues. He is a well-published author and a columnist in *Data Management Review* for the column "The CRM-Ready Data Warehouse."

A regularly featured expert on data warehouse/business intelligence and CRM at major conferences, William is widely quoted on data warehouse issues in publications such as *Computerworld* and *InformationWeek* and has been featured on several prominent expert panels. An expert witness, skills evaluation author and a judge for best practices competitions, William is the former vice president of a recognized best practices information management program with Teradata.

William was a Version 1 developer of the DB2 DBMS product at IBM Santa Teresa Labs.



5960 West Parker Road
Suite 278, #133
Plano, TX 75093
(214) 514-1444