

SYBASE®

SYBASE IQ

IMPROVING THE ECONOMICS OF BUSINESS REPORTING

THE
ENTERPRISE.
UNWIRED.



TABLE OF CONTENTS

SYBASE IQ: THE ECONOMICS OF BUSINESS REPORTING 3

SPEED, SERVICE, AND RESPONSIVENESS 3

REPORTING IS ON THE RISE 4

NEW PRESSURES FOR INFORMATION TECHNOLOGY 4

HOW TO MEET THE CHALLENGE 4

WHICH APPROACH TO TAKE? 5

A CLOSER LOOK AT REPORTING SYSTEMS 5

ONE SIZE DOES NOT FIT ALL 6

THE ADVANTAGES OF SYBASE IQ 6

SYBASE IQ: A CLEAR WINNER 7

SPEED
SERVICE
RESPONSIVENESS

SYBASE IQ: IMPROVING THE ECONOMICS OF BUSINESS REPORTING

Each year, the amount of data your company handles grows exponentially. Your systems must work harder every day to satisfy increased demands for answers, insights, and reports. Data analysis and business reporting are consuming more and more of your systems and staff resources. It's time to evaluate better alternatives.

Clearly, there is a growing demand for business intelligence in all forms. People want and need better access to information to make smarter, faster decisions in response to changing circumstances or events. Reporting is one form of business intelligence that has become business-critical.

Once the domain of behind-the-scenes accountants and analysts, reporting is now seen as a way for companies to improve service, ensure quality, control costs, or prevent losses. Traditionally, reports are routine, scheduled documents that provide a snapshot of the operations. Increasingly, business reporting now encompasses frequent and complex “what if” analyses by knowledge workers using sophisticated analytical software tools.

SPEED, SERVICE, AND RESPONSIVENESS

Examples of the need for timely, accurate reporting can be found across many industries.

In healthcare, for example, payors need to match claims against eligibility data in order to prevent losses associated with intentional or unintentional fraud. It is not enough to see one or two snapshots, but payors must be able to easily interpret the changing patterns among hundreds and thousands of pieces of information about patients, providers, and suppliers.

In financial services, business reporting systems help providers perform sophisticated analyses to obtain a more complete picture of its customers' preferences, risk, and profitability. Armed with this information, financial services providers can better retain their customers through personalized customer service and customized product offerings. By doing so, they avoid the missed revenues that occur when customers are under-served—or worse—lost.

In deregulated telecommunications markets, customer service is everything. Telecommunications providers need to analyze massive amounts of data, including call logs and billing details, in order to develop the right array of calling plans and offerings that attract and retain loyal customers.

State governments must be able to provide accurate and on-time participation reports to ensure Medicaid funding. They must be able to understand and report who is participating in multiple programs such as food stamps, or Woman Infant and Children care programs, so that they can get the needed funds for distribution. If these reports are inaccurate or late, Federal funds are delayed or cancelled—a catastrophic result.

“SYBASE IQ IS TYPICALLY 128 TIMES FASTER THAN ROW-BASED LOCKING.”

—RICHARD WINTER, THE WINTER GROUP

REPORTING IS ON THE RISE

Reporting is not a new requirement for IT departments. Yet the proliferation of data and users, as well as new, more complex requirements has put reporting squarely on the “Top Five” agenda for many IT organizations.

- **More data.** Gartner predicts that in 2004, enterprises will be managing 30 times more data than they did in 1999. Data is growing organically—as business operations expand—but also because enterprises must keep more data online. In some cases regulatory requirements demand that more data be stored online, in other cases business users simply want longer and more convenient access to historical data for analytical purposes.
- **More end users.** Intelligence-driven decisions are becoming more prevalent, and decision support tools are becoming easier to use. These factors result in more users asking for more reports and ad hoc query support.
- **More complex queries.** The natural outcome of more users, with a greater need-to-know, and broader access to reporting tools is more sophisticated “what if” scenarios and more frequent reporting needs. Additionally, more stringent regulatory requirements are adding more in-depth analyses to reporting work loads.

NEW PRESSURES FOR INFORMATION TECHNOLOGY

These new demands are placing added pressures upon already overworked IT departments and systems resources.

- **Higher costs.** It takes more system resources to handle more data and higher workloads. Analytical queries consume large amounts of CPU cycles to read every byte of every row of a large database and deliver the query result. In order to keep performance at target levels, more hardware must be added to the system. Reporting and decision-support also takes more DBA time to tune queries, adding indexes and summary tables to ensure acceptable response times.
- **Slower performance.** The sheer volume of ad hoc queries or intra-day batch reporting processes can have a dramatic effect on systems performance. When companies run reports off of the same operational systems that drive the business, slow response times have a direct impact on revenues, productivity, and customer satisfaction. Even if reporting systems are separate from operational systems, it is unacceptable for end users to wait hours for query results.
- **Missed reporting windows.** Even overnight batch reporting windows are not spared from the consequences of “information overload.” Reports simply cannot finish in the time allotted given the resources available. This not only compromises Service Level Agreements between IT and business units, but can result in fines or penalties for non-compliance with a regulatory agency.

HOW TO MEET THE CHALLENGE

In many instances, queries and reports are simply performed by operational systems in a mixed-use environment, processing two distinct types of workloads—either processing transactions, or processing queries.

In cases where the query support workload is affecting transactional performance, many IT organizations move operational data to a separate system to serve the analyst or decision-making community of users. These are also called Reporting Servers, or Operational Data Stores. Typically this is a one-way copy of data, which is refreshed at designated intervals depending upon how current the data must be to serve the analytical needs of the business.

“THIS DATA WAREHOUSE IS BY FAR ONE OF THE LARGEST I’VE EVER SEEN AND IT’S GOING TO GROW CONSIDERABLY EVERY YEAR. I AM CONFIDENT THE HARDWARE AND TECHNOLOGY [SYBASE IQ] WE HAVE SELECTED WILL BE ABLE TO KEEP UP WITH OUR DEMANDS AS WE GROW IT TO THE 30 TERABYTE LEVEL.”

—TIMOTHY S. GEARY, DATA RESOURCE MANAGER, INFORMATION TECHNOLOGY, NATIONAL PRODUCT GROUP, NIELSON MEDIA RESEARCH

Enterprise Data Warehouses are often created when there is a need to aggregate and standardize data from multiple operational systems to serve the decision-making community, and to support powerhouse analytics and queries.

WHICH APPROACH TO TAKE?

Mixed-use operational systems are simpler to manage, as there is only one copy of the data to maintain. However, decision-support queries are very different in nature from OLTP processing, and can consume many CPU cycles in order to complete. Users want to know why their queries or reports take so long, and DBAs can spend a great deal of time tuning systems for faster queries.

Enterprise Data Warehouses are an excellent way to create a shared, centralized view of information that is contained in multiple operational systems throughout the organization, and deal with complex issues such as data aggregation, standardization, and bi-directional updates. Data Warehouses are strategic investments, and high ROI is associated with very high numbers of queries and users.

Reporting systems are often a low-risk way to preserve the performance of operational systems by separating two distinct workloads and optimizing each system for its particular task.

A CLOSER LOOK AT REPORTING SYSTEMS

Reporting systems involve a certain amount of redundancy, as operational data is essentially copied and placed on a separate server and storage repository. Without close watch, it is easy for reporting systems to become inefficient and bloated, driving up unnecessary costs.

To evaluate the cost effectiveness of a reporting system solution, keep in mind the following:

- **How much data will you store?** The need to store more data, for longer periods of time, will have a significant impact of the amount of storage hardware. According to Gartner, storage is the fastest-growing capital cost within the data center, with storage increasing more than 30% per year.
- **How much data will you refresh, and at what intervals?** More data creates other complications, for example it takes longer to load and refresh onto reporting systems. Customer service suffers when batch update windows are not long enough, or intra-day updates take systems offline or interfere with decision support response times.
- **How many queries will you support?** Growing query workloads demand more attention from DBAs who need to tune and optimize systems to perform efficiently for all users.

ONE SIZE DOES NOT FIT ALL

If you are thinking that you will simply copy your operational database onto your reporting system and use the same underlying data management system, think again.

First of all, traditional RDBMSs are optimized for OLTP, not reporting or analytics. OLTP-designed RDBMSs require complex, space-consuming indexing and summary tables to perform well. These indexes and summary tables actually explode data sizes, often requiring three to ten times more data storage on the reporting system than the original operational system.

Traditional RDBMS are also more complex to implement for decision support environments. They require more time to load and refresh a reporting system, due to labor-intensive steps of creating backups, tables, and indexes. It also takes a lot of tuning to maintain query performance with a traditional RDBMS—diagnosing, testing, and tuning queries over and over again.

“AS WE STORE MORE DATA IN SYBASE IQ, WE USE THE MAINFRAME LESS AND LESS. THIS HAS LED TO THE STATE SAVING \$100,000 PER MONTH IN MAINFRAME DISK STORAGE COSTS.”

—BRENT SANDERSON, IT MANAGER, STATE OF UTAH

THE ADVANTAGES OF SYBASE IQ

Unlike other data management platforms, Sybase IQ has been designed from the ground up for reporting and decision support, resulting in a number of proven advantages:

- **Faster.** Sybase has a unique column-based structure, as opposed a row-based structure, along with patented indexing technology that allows Sybase IQ to deliver ad hoc query performance up to 100 times faster than traditional RDBMSs.
- **Less Expensive.** Sybase IQ also has sophisticated compression algorithms that reduce storage needs anywhere from 30 to 60 percent. Independently audited tests have confirmed that to store 155 Terabytes of raw input data, Sybase IQ required 55 Terabytes of physical storage. OLTP databases are row-based, with additional indexes (data about the data) that explode the data storage requirement. In this example, one Petabyte (or 1,000 Terabytes) of physical storage would be required by an OLTP solution. OLTP databases are for transactions, NOT large searches through Terabytes of data.
- **Easier.** Sybase IQ is easier to maintain than traditional databases and does not require time- and resource-intensive tuning to obtain excellent performance. Like all Sybase technology, Sybase IQ is built upon open standards, ensuring integration and interoperability with existing systems and familiar reporting tools.
- **More Scalable.** Sybase IQ offers near-linear user and data scalability to support 100s and 1,000s of users and Terabytes of data.

“THE RESULTS ARE ASTONISHING; SYBASE IQ RESPONDS TO 57% OF AD HOC QUERIES IN LESS THAN ONE SECOND... 1.5 TERABYTES OF RAW DATA REQUIRES 500 GM OF STORAGE IN SYBASE IQ, COMPARED TO 6 TERABYTES WITH ANOTHER RELATIONAL DATABASE.”

—MICHAEL VAN LAETHEM, IS PROJECT MANAGER, FORTIS BANK

SYBASE IQ: A CLEAR WINNER

Sybase IQ infuses organizations with fast, flexible access to information. With Sybase IQ, you can analyze business performance, detect fraud, gauge marketing results, manage customer relationships, and ensure financial controls—in ways never before possible. The economics of Sybase IQ put faster, efficient, cost-effective reporting and decision support systems well within reach.

FIG 1: COMPARING SYBASE IQ TO TRADITIONAL DATA MANAGEMENT SOLUTIONS

Speed	10 to 100 times faster than OLTP systems for analytics.
24/7	Faster data loading. No impact to query performance while loading.
Cost	Traditional RDBMSs require more storage than Sybase IQ, adding 3 to 10 times more storage costs.
Ease of Deployment	Designed for analytics; no need to re-architect OLTP systems. Tune once and you're done.
Supports familiar reporting tools	Standards-based and works with any ODBC- or JDBC-compliant tools.
Ongoing Ease-of-use	Standard SQL, any schema, low maintenance.

“SYBASE IQ REDUCED LOADING AND INDEXING FROM 30 MINUTES TO TWO AND A HALF TO THREE MINUTES. QUERY SPEEDS WERE 20 TO 50 TIMES FASTER THAN ORACLE. TIME TO ADD A COLUMN WAS REDUCED FROM FOUR HOURS WITH ORACLE TO 15 MINUTES WITH SYBASE IQ.”

—JEFF BUTLER, ASSISTANT DIRECTOR, BUREAU OF TRANSPORTATION

FOR MORE INFORMATION, CONTACT US TODAY AT INFOBI@SYBASE.COM.
AND VISIT US ON THE WEB AT WWW.SYBASE.COM/PRODUCTS/BI/SYBASEIQ.



SYBASE®

Sybase Incorporated
Worldwide Headquarters
One Sybase Drive
Dublin CA, 94568 USA
T 1.800.8.SYBASE
www.sybase.com

Copyright © 2004 Sybase, Inc. All rights reserved. Unpublished rights reserved under U.S. copyright laws. Sybase and the Sybase logo are registered trademarks of Sybase, Inc. All other trademarks are property of their respective owners.® indicates registration in the United States. Printed in the U.S.A. L02555 08/04