

# Sybase® IQ: The Economics of Business Reporting

## TABLE OF CONTENTS

- 1 Sybase IQ: The Economics of Business Reporting
- 1 Speed, Flexibility, and Scalability
- 2 Reporting is on the Rise
- 2 New Pressures for Information Technology
- 3 How to Meet the Challenge
- 3 Which Approach to Take?
- 3 A Closer Look at Reporting Systems
- 4 One Size Does Not Fit All
- 4 The Advantages of Sybase IQ
- 5 Sybase IQ: A Clear Winner

## **SYBASE IQ: THE ECONOMICS OF BUSINESS REPORTING**

Each year, the amount of data your company handles grows exponentially, as does the number of users who need access to insights from this growing data. Your systems must work harder every day to satisfy increased demands for answers that support mission-critical business decisions. Data analysis and business reporting are consuming more and more of system and staff resources. It's time to evaluate better alternatives.

Clearly, there is a growing demand for business intelligence in all forms. Decision-makers 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 produced only by behind-the-scenes accountants and analysts, and received and acted upon only in the executive suites, today, reporting is recognized as a way for companies to improve service, ensure quality, control costs, and prevent losses by empowering decision-makers throughout the organization.

Traditionally, reports were routine, scheduled documents that provided a snapshot of operations within a department or across the enterprise. Increasingly, business reporting encompasses dynamically updated, visually rich dashboards and on-demand ad-hoc report generation for a wide variety of decision makers.

## **SPEED, FLEXIBILITY, AND SCALABILITY**

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. It is not enough to see one or two snapshots, payors must be able to easily interpret the changing patterns among hundreds or thousands of pieces of information about patients, providers, and suppliers.

In telecommunication markets, customer service is everything. Telecom providers need to arm hundreds of customer-service representatives with customer calling behavior reports culled from massive amounts of data from call logs and billing records to enable these customer care professionals to upsell and cross-sell calling plans and offerings.

In financial services, business reporting systems help providers perform sophisticated analyses to obtain a more complete picture of their 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 underserved—or worse—lost.

State governments must be able to provide accurate and on-time participation reports to ensure Medicaid funding is kept flowing. They must be able to understand and report who is participating in multiple programs such as food stamps and Woman, Infant, and Children care programs, in order to ensure appropriate funding. If these reports are inaccurate or late, Federal funds are delayed or cancelled—a potentially catastrophic result.

**“THE COLUMN BASED  
APPROACH IS UP TO 100 TIMES  
FASTER THAN TRADITIONAL  
ROWBASED DATABASES.  
SYBASE IQ IS THE MOST  
NOTABLE OF THESE.”**

**PHILIP HOWARD, BLOOR RESEARCH**

“SINCE ADDING SYBASE IQ, THE DEMAND FOR REPORTS HAS RISEN DRAMATICALLY AS USERS REALIZE JUST HOW MUCH MORE HELPFUL THE INFORMATION IS THAT WE CAN NOW PROVIDE. WE RECEIVE REQUESTS FOR 10,000 TO 20,000 MORE REPORTS PER MONTH, AND AD HOC QUERIES HAVE GROWN FROM 5,000 TO 10,000 PER MONTH.”

MARC GUILLARD, DBA,  
BNP PARIBAS SECURITIES SERVICES

## 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 at the top of the agenda for many IT organizations.

**More data.** Data is growing organically—as business operations expand—but also because enterprises must interact with more sources of data and 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 reporting 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 frequent access.** The natural outcome of more users, with a greater need to know, and broader access to reporting tools is higher frequency and volume of reports. Additionally, more stringent regulatory requirements are adding to the reporting workloads by forcing firms to ensure compliance by running frequent audit reports on their operational and historical data.

**More complexity.** Even as the volume of data grows, the number of users increases, and the amount of time available to create and access reports decreases, another significant challenge arises owing to the depth and sophistication of the reports that are produced. Decision-makers need an ever-finer view of the information that drives the business. Reports are growing not only in the amount of information they are drawn from, but the number of information relationships they display and the depth to which they explain those relationships.

## 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. A large volume of reporting queries consumes large amounts of CPU cycles, particularly when those queries hit every byte of every row of a large database before delivering the query result. In order to keep performance at target levels, more hardware must be added to the system.

Reporting and decision support also consume administrative cycles required to tune queries and add indexes and summary tables to ensure acceptable response times.

**Slower performance.** The sheer volume of concurrent ad hoc reports or intraday batch reporting processes can have a dramatic effect on system performance. When companies run reports off 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 minutes, or in many cases, hours, for query results.

**Missed reporting windows.** Even batch reporting windows are not spared from the consequences of “information overload.” Reports simply cannot finish in the time allotted given the resources available. Additionally, some environments are required to produce reports from near-real time data and face fines or penalties from regulatory agencies for noncompliance.

## HOW TO MEET THE CHALLENGE

In many instances, reports are simply generated by operational systems in a mixed-use environment, processing two distinct types of workloads: transaction and decision-support queries.

In cases where the reporting workload is affecting transactional performance, many IT organizations find themselves in an impossible position where both operational and reporting requirements cannot be satisfied. If both sets of requirements cannot be met, it is most frequently the query load that is not fully met, leading to compromises in the service level agreements (SLAs) between IT and business units. In such a situation, it may be a good idea to move operational data to a separate system to serve the decision-makers. Such a deployment is called an Operational Reporting Server and is often used just by one department or business unit. Typically, this is one-way copy of data, which is refreshed at designated intervals depending upon how current the data must be to serve the reporting needs of the business. Enterprise Report Servers, on the other hand, are often created when there is a need to aggregate and standardize data from multiple operational systems to serve the enterprise-wide decision-making community, and to support hundreds or even thousands of users.

## 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 Report Servers 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. Enterprise Report Servers are strategic investments, and high ROI is associated with very high volumes of queries and users.

Operational Reporting Servers 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—transaction or decision-support queries.

## 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 careful management, 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 on 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 intraday 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.

**“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**

**“BY COMPRESSING THE RAW DATA IN A FAR MORE INTELLIGENT FASHION THAN OTHER MARKET ALTERNATIVES, SYBASE IQ HAS SAVED US 40 PERCENT WHEN IT COMES TO DATA VOLUME, INCLUDING INDEXES.”**

ALMOG ZAHAVI, DIRECTOR OF SOFTWARE DEVELOPMENT AND CONTROL SYSTEMS, CELLCOM

### **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. Traditional relational database management systems (RDBMSs) are optimized for OLTP, not reporting. OLTP-designed RDBMSs require complex, space-consuming indexing and summary tables to perform well on reporting workloads. 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 RDBMSs are also more complex to implement for reporting environments. They demand more time to load and refresh a reporting system due to labor-intensive steps of creating pre-aggregates and indexes. They also require constant tuning to make them perform well for reporting queries.

A customized data warehouse appliance can be a decent alternative to a traditional RDBMS, but these tend to be cost prohibitive for smaller reporting environments and many of them lack the flexibility required to support hundreds of concurrent reports and users.

### **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 reporting query performance up to 100 times faster than traditional RDBMSs. For a dedicated (operational or enterprise) reporting environment, Sybase IQ offers a flexible architecture and enhanced data loading features to ensure that, in even the most demanding environments, answers are available to decision-makers when needed—within seconds or minutes.

**Less Expensive.** Sybase IQ also has sophisticated compression algorithms that reduce storage needs anywhere from 30 to 85 percent. Independently audited tests have confirmed that to store 1 Petabyte of raw input data, Sybase IQ required 260 Terabytes of physical storage. OLTP databases are row-based, and require additional indexes (data about the data) that explode the data storage requirement. In this example, more than 1.5 Petabyte (or 1500 Terabytes) of physical storage would be required by an OLTP database. OLTP databases are for transactions, NOT large searches through terabytes and petabytes 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 industry leading reporting and dashboarding tools. Also, Sybase IQ does not mandate specialized schemas (dimensional modeling) to perform well.

**More Scalable.** Sybase IQ offers near-linear user and data scalability to support hundreds or thousands of users and terabytes or even petabytes of data.

## **SYBASE IQ: A CLEARWINNER**

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

**“THE CURRENT SYSTEM WAS NOT EFFICIENT UNTIL SYBASE IQ CAME ALONG. SINCE WE’VE (USSECURITIES AND EXCHANGE COMMISSION) BEEN LIVE WITH THE NEW SYSTEM, THERE HAS BEEN NO UNPLANNED DOWNTIME. NONE. IT HAS BEEN 100% SATISFACTORY.”**

**SAM FOSTER, PRESIDENT,  
FOSTERSOFT, USSEC CONTRACTOR**



SYBASE, INC.  
WORLDWIDE HEADQUARTERS  
ONE SYBASE DRIVE  
DUBLIN, CA 94568-7902  
U.S.A.  
1 800 8 SYBASE

[www.sybase.com](http://www.sybase.com)

Copyright © 2009 Sybase, Inc. All rights reserved. Unpublished rights reserved under U.S. copyright laws. Sybase and the Sybase logo are trademarks of Sybase, Inc. or its subsidiaries. All other trademarks are the property of their respective owners. ® indicates registration in the United States. Specifications are subject to change without notice.  
04/09 L03195

**SYBASE®**