

Sybase® IQ 15 Accelerates Business Reporting



SOLUTION OVERVIEW

In a recent independent benchmark on a world's record one petabyte of raw data, Sybase IQ compressed data 85 percent resulting in a "green" data environment with 90 percent less CO₂ emissions over the data warehouse's life time.
www.sybase.com/detail?id=1054011

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 system and staff resources. It's time to evaluate better alternatives.

BETTER REPORTING IMPROVES INSIGHT

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 their customers' preferences, risk, and profitability. Armed with this information, financial services providers can develop richer, more customized service offerings to retain current customers and attract new ones. This capability enables financial services firms to protect commission-based revenue streams that have been eroding in the face of increasing industry automation.

In deregulated telecommunications markets, customer service is crucial. 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, as another example, 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 to distribute. If these reports are inaccurate or late, Federal funds are delayed or canceled—with potentially catastrophic results.

REPORTING IS ON THE RISE—GROWING IN BOTH VOLUME AND COMPLEXITY

Companies create reports for a wide range of purposes in order to better understand and act upon changing business events and conditions. This includes large numbers of standard reports which can range in the thousands per day, performance dashboards, or "key performance indicators" (KPIs), as well as ad hoc reports generated by a growing community of users seeking additional insight.

Query-intensive reports often compete with operational applications for computing resources, including customer relationship management (CRM), enterprise resources planning (ERP), or order entry systems. When reports access the same data used by day-to-day transactional systems, companies face a constant balancing act between report-generation performance and on-line transaction processing (OLTP) performance. This trade-off can result in significant efforts to tune and optimize reports, or to add computing hardware and storage capacity to handle the additional workloads without compromising the performance of either transactions or reports.

“FOR 18 MONTHS, WE HAVE RECORDED TWO-FIGURE GROWTH READY-TO-WEAR, WHILE THE REST OF THE MARKET IS RELATIVELY FLAT. SYBASE IQ CONTRIBUTES TO THIS GROWTH BY THE SUPPORT IT GIVES TO PRODUCT MANAGERS AND TO A COLLECTION THAT CONTINUES TO EVOLVE.”

ENSO GIANDOMENICO, ASSISTANT MANAGING DIRECTOR, PHILDAR

Even enterprise data warehouses are straining to meet increasing demand for business and corporate reporting. Additional reporting workloads are driven by the sheer growth in data, an increase in the number of people who want to analyze that data, and increasing complexity in the required reports and queries.

- **More data.** Organizations are managing much more data than they did even one or two years ago. 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 reporting purposes. Many of the larger reporting environments now incorporate tens or even hundreds of terabytes of raw data.
- **More end-users.** Intelligence-driven decisions are becoming more prevalent, largely driven by the growing popularity of dashboards or KPI reports. Decision support tools are also becoming easier to use. These factors result in simply more users asking for more reports and ad hoc query support.
- **More complex reports.** 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, whether scheduled or ad hoc. Tighter regulatory requirements, also, are adding more frequency and complexity to reporting workloads.

BUSINESS INTELLIGENCE INFRASTRUCTURES ARE UNDER PRESSURE

These new demands are placing added pressures upon already overworked IT departments and the systems established to support data warehousing or business intelligence reporting. These pressures can result in three undesirable outcomes related to business intelligence infrastructures: slower performance, missed reporting windows, and higher costs.

- **Slower performance.** The sheer volume of queries driven by this increased demand for reports and for faster, better decision-making 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 reports are run off of separate data warehouses, it is clearly unacceptable for end users to wait hours for results.
- **Missed reporting windows.** Even overnight batch reporting windows are not spared from the consequences of “information overload.” Often times, reports simply can’t finish in the time allotted given the traditional computing and storage resources available. This not only compromises Service Level Agreements (SLAs) between IT and business units, but can result in fines or penalties for non-compliance with regulatory agencies.
- **Higher costs.** It takes more system resources to handle more data and higher workloads. Traditional enterprise data warehouses or OLTP systems consume large amounts of CPU cycles to read every byte of every row of a large database and deliver the report result. In order to keep performance at target levels, more hardware must be added to the system. Reporting and decision-support also take more DBA time to tune queries, adding indexes and summary tables to ensure acceptable response times.

Hence, the growth in data levels, end-user reporting volumes, and report complexity, and the pressures this places on traditional business intelligence infrastructures, are driving the need for a faster, more efficient way to manage and store business data.

REPORT SERVERS PROVIDE RELIEF TO OLTP AND EDW SYSTEMS

In cases where report query workloads are affecting the performance of OLTP systems or Enterprise Data Warehouses, many IT organizations offload critical data to a separate report servers to support the analyst or decision-making community of users.

Reporting servers often provide a low-risk way to preserve the performance of operational systems or EDWs by separating distinct workloads and optimizing each system for its particular task. Relevant data is essentially copied and placed on a separate server and storage repository, and refreshed at designated intervals depending upon how current the data must be to serve the reporting needs of the business. As the amount of data and reporting requirements increase, however, these reporting servers can become inefficient and bloated.

SYBASE IQ – TURBO-CHARGING REPORT SERVERS

Sybase IQ is a highly optimized analytics engine used for business intelligence, advanced analytics, predictive modeling, stringent regulatory compliance, and rapid reporting.

Unlike transactional databases that were designed to support business transactions, Sybase IQ was architected for reporting and analysis. Transactional databases require complex, space-consuming indexing and summary tables to perform query-intensive workloads well. These indexes and summary tables actually explode data sizes, often requiring 5 or 10 times more data in the reporting system than in the original operational system.

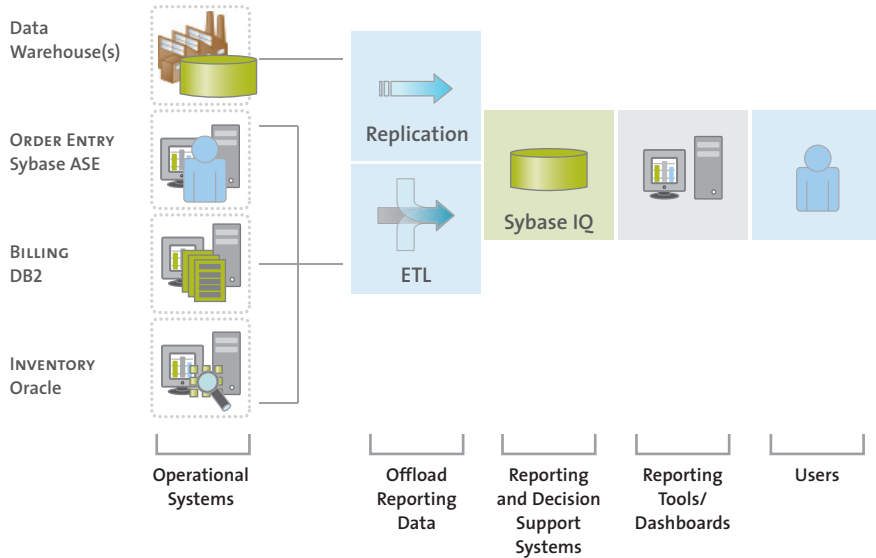


Figure 1. Sybase IQ is a powerful analytics server designed and optimized for fast query processing and more efficient utilization of systems resources for enterprise reporting systems.

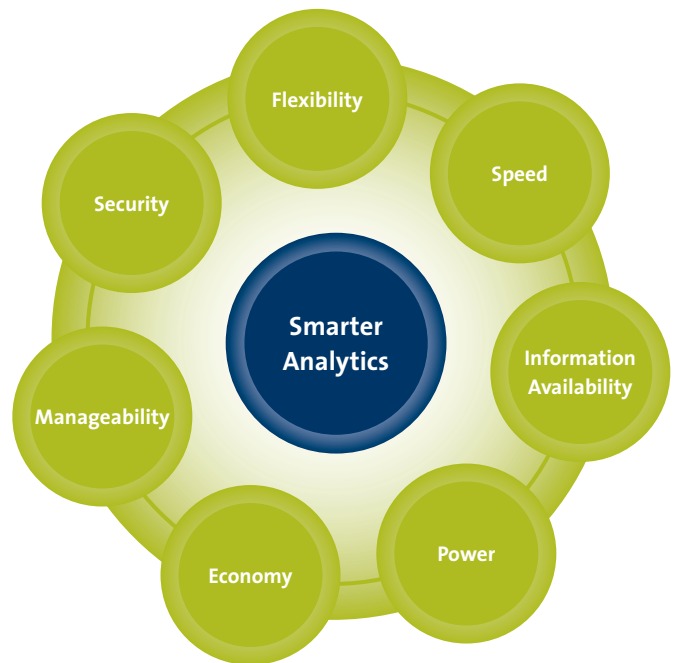
Transactional databases are also more complex to implement for decision support environments. They require more time to load and refresh, due to labor-intensive steps of creating backups, tables, and indexes. It also takes a lot of tuning to maintain query performance with a traditional database—diagnosing,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

Sybase IQ turbo charges report servers—with extremely high performance for feeding data to reporting tools and dashboards, without requiring special tuning or expensive hardware and storage resources.

THE ADVANTAGES OF SYBASE IQ 15

For businesses looking for more accurate and cost-effective reporting solutions that deliver winning business results by creating a complete picture of customer preferences, meeting regulatory requirements, and providing up-to-the-moment visibility into operational results, Sybase IQ is the world’s leading column-based analytics server, delivering the smartest approach to enable enterprises to turn raw data into actionable information through comprehensive reporting.

Sybase IQ is the only reporting engine that delivers dramatically faster, more accurate reports—to all your users, from all your information, on your terms. Sybase IQ’s column-based core architecture and innovative features provide flexibility and scalability in multiple dimensions, enabling unsurpassed query performance for the most complex reporting tasks, for the rapidly expanding user base, on the largest datasets—all within an environment that is both affordable, eco-friendly and manageable at the implementation stage and throughout the information lifecycle.



“SYBASE IQ SURPASSED THE SPEED OF OUR INTERNAL CUSTOM APPLICATION BY TWO OR THREE TIMES, AND IT WORKED SO WELL WE ESTABLISHED SYBASE IQ AS A BUSINESS INTELLIGENCE ARCHITECTURE STANDARD THAT THE REST OF THE COMPANY CAN NOW CONSIDER.”

MARC GUILLARD, BNP PARIBAS SECURITIES SERVICES DBA

All of this is evidenced by award-winning results: more than 1,500 customers with more than 3,000 projects worldwide, top industry analyst support, media and pundits, blogs, and benchmark results.

SYBASE IQ 15: 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.

For more information, contact us today at infobi@sybase.com or visit www.sybase.com/iq.

Speed	<ul style="list-style-type: none"> • Column-based architecture for low I/O and fast retrieval • Architected for queries, not transactions • Patented indexing technology • Query speed 10X - 100X faster than traditional row-based DBs
Economy	<ul style="list-style-type: none"> • Compresses data as its loaded (reduces storage needs up to 90%) • Reduced hardware/energy costs due to compression capability • Runs on standard systems including Linux, Windows, & Unix
Manageability	<ul style="list-style-type: none"> • Ease of deployment; no need to re-architect OLTP systems • Tune once, and you're done
Power	<ul style="list-style-type: none"> • One-touch multiplex grid configuration • High performance from any schema • Support hundreds or thousands of users and petabytes of data • Tested on trillions of rows of data
Flexibility	<ul style="list-style-type: none"> • Independent compute power and storage capacity scalability • Standards-based; works with ODBC/JDBC-compliant tools • Supports popular BI tools (Microstrategy; Cognos; BO) • Easily co-exists within existing BI ecosystems
Information Availability	<ul style="list-style-type: none"> • Faster data loading • Multi-node loading • Direct client loads • No impact to query performance while loading
Security	<ul style="list-style-type: none"> • Password authentication/encryption for data-in-flight • Database and column encryption for data-at-rest • Common Criteria certified at EAL-3 • User settable passwords • FIPS (on Unix), Kerberos, and IPv6 support

Figure 2. Comparing Sybase IQ to Traditional Data Management Solutions