

Sybase® IQ Accelerates Advanced Analytics Advanced Analytics as the New Competitive Weapon



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](http://www.sybase.com/detail?id=1054011)

At a time when companies in many industries offer similar products and use comparable technology, high-performance business processes are among the last remaining points of differentiation. Many of the previous bases for competition are no longer available. Unique geographical advantage doesn't matter in global competition, and protective regulation is largely gone. Proprietary technologies are rapidly copied, and breakthrough innovation in products or services seems increasingly difficult to achieve. Executing your business with maximum efficiency and effectiveness, and making the smartest business decisions possible is one of the few remaining bases for competitive advantage, and advanced analytics enables organizations to accomplish this.

Advanced analytics involves the extensive use of data mining, statistical analysis, forecasting, predictive modeling, and optimization techniques to drive strategic and operational decisions and actions. The analytics may be input for human decisions or may drive fully automated decisions. Organizations rely on advanced analytics for insight that goes far deeper than what traditional business intelligence systems typically enable. While traditional business intelligence systems deal with data access and reporting, and help answer questions such as what happened, what exactly is the problem, and what actions are needed, advanced analytics use the methods above to help you understand why events are happening, predict what will happen next, and identify the most optimal course of action.

The questions that advanced analytics can answer represent the higher-value and more proactive end of the business intelligence spectrum, and provide new ways to explore, understand, and validate the events and actions that will shape the success of your business into the future. Analytics can support almost any business process, and enable organizations to create distinctive capabilities in a wide variety of business domains. These can include obtaining better understanding of customers and competitors, fraud prevention, risk simulations and reduction, optimizing supply chain performance, maximizing the results of marketing investments, identifying and stopping revenue leaks, and scientific knowledge discovery.

BETTER ANALYTICS IMPROVES DECISIONS IN NUMEROUS INDUSTRIES

Companies across a variety of industries are discovering the power of advanced analytics for improving the effectiveness of business decisions that results in distinct competitive advantage.

In financial services, for example, advanced analytics applications support enhanced risk management. Industry dynamics such as the move to algorithmic trading, the credit crisis, decimalization, the increasing complexity of securities and the exponential increase in trading volumes has made it harder for financial services firms to make money. As a result, fund managers and traders are forced to move into riskier investment vehicles. Through better insight into historical data combined with current market activity, these firms can preserve profits while pursuing riskier strategies.

In deregulated telecommunications markets, operators now deliver a wide range of services and content to customers. Through the ability to quickly analyze massive volumes of customer data, these companies can simulate various packaging options to deliver the optimum set of services to maximize customer loyalty—and profits.

“WITH SYBASE IQ OUR USERS HAVE BEEN SURPRISED TO OBTAIN RESPONSES IN VIRTUALLY REAL TIME. THEY CALL THAT THE ‘GOOGLE’ EFFECT. THEY ARE ABLE TO OBTAIN QUICKER RESPONSES TO AD HOC QUERIES USING SYBASE IQ THAN THEY RECEIVED FROM PREFORMATTED REQUESTS IN THE PREVIOUS SYSTEM. AND THAT WITH A CLEARLY INCREASED VOLUME.”

PIERRE-ALEXANDRE PAUTRAT,
IT DEPARTMENT MANAGER

“WE MIGHT HAVE NEEDED A DISK WITH THREE TO FOUR TIMES MORE THAN THE CURRENT CAPACITY IF THE ENTERPRISE DATA WAREHOUSE HAD BEEN ESTABLISHED WITH ANOTHER TYPE OF DATABASE.”

YE GU HOI, LIFE INSURANCE
DEVELOPMENT TEAM MANAGER,
SAMSUNG SDS

In healthcare, advanced analytics promises the potential for combining data from physicians, individuals, labs, payers and beyond to transform the quality of healthcare. The move to electronic health records and the deployment of advanced clinical data systems drives an explosion in the amount of data that can be mined for intelligence. Providers can analyze a wide variety of variable including past medical history, medications, treatment plans, and environmental factors such as air quality or ozone level—and ensure the right levels of care are in place to meet predicted demand.

BUSINESS INTELLIGENCE INFRASTRUCTURES ARE UNDER PRESSURE

The value of advanced analytics is substantial and can be measured in terms of competitive advantage, significant cost savings, and greater revenues. Yet this value can be undermined by the challenges that advanced analytics presents. The nature of advanced analytics places additional pressures on business intelligence infrastructures due to three factors: more data, complex queries, and ad hoc 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 advanced analytics requires more data to be available online. The more data that is accessible for analysis, the more accurate predicted outcomes are likely to be. For example, performing analysis of trends sufficient to make meaningful predictions requires years’ worth of data. Historical data that would have once been archived now must be available online to support these advanced analytics applications, and in fact many data warehouses are now incorporating tens or even hundreds of terabytes of raw data.
- **Complex queries.** Advanced analytics requires complex queries that compare and contrast different data sets in order to spot trends, correlate cause and effect, or recognize patterns. “How has this new product affected sales of earlier products?” or “which customers are more likely to respond to this new campaign?” represent complex queries. Advanced analytics also considers a wide range of predicates—selection criteria such as those based on sex, age, income, location, or other factors. The growing amount of data mentioned above also makes these complex queries more challenging because of the sheer volume of data that needs to be analyzed to identify root causes of events, data relationships, and the likely outcomes of potential business decisions.
- **Ad hoc queries.** Advanced analytics requires an ability to explore vast amounts of business data with an unconstrained “what if?” approach. Due to this fact, ad hoc (vs. planned or scheduled) queries comprise a significant amount of the query workload, can be very large, and can potentially be run at any time by any user. Since the timing and volume of these ad hoc queries are unpredictable, the “right” capacity needed for them is elusive, raising the twin risks that resources will fall short of demand or will be under-utilized much of the time—with disastrous consequences to the business and other users.

All of these demands of advanced analytics consume significant computing resources and are placing added pressures upon already overworked IT departments and the systems established to support data warehousing or business intelligence. They can result in significant efforts to tune and optimize the data warehouse, or to add significant hardware and storage capacity to handle the additional workloads.

These pressures can result in two undesirable outcomes related to business intelligence infrastructures: slower performance and higher costs.

- **Slower performance.** The sheer volume and complexity of queries required for advanced analytics can have a dramatic effect on system 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 queries are run off of separate data warehouses, it is clearly unacceptable for end users to wait hours for results.
- **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 query result. In order to keep performance at target levels, more hardware must be added to the system. Advanced analytics also take more DBA time to tune queries, adding indexes and summary tables to ensure acceptable response times and meet Service Level Agreements (SLAs).

Hence, the growing reliance on advanced analytics as a new competitive weapon, and the pressures it places on traditional business intelligence infrastructures, is driving the need for a faster, more efficient way to manage and store business data.

ANALYTICS SERVERS PROVIDE RELIEF TO OLTP AND EDW ENVIRONMENTS

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

Analytics servers are often 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 analytical needs of the business. Without close watch, however, these servers can become inefficient and bloated.

SYBASE IQ—A SMARTER APPROACH TO ADVANCED ANALYTICS

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

Unlike traditional databases that were designed to support business transactions, Sybase IQ was architected for advanced analytics. 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 analytical system than in the original operational system.

Traditional 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 makes advanced analytics practical—with extremely high performance for analyzing massive amounts of data, without requiring special tuning or expensive hardware and storage resources.

THE ADVANTAGES OF SYBASE IQ

For businesses looking for more accurate and cost-effective analytics solutions that deliver winning business results by attracting new customers, growing revenues, and predicting outcomes and risks in a volatile and competitive world climate, 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 analytics—to increase revenue, customers and margins.

Sybase IQ is the only analytics server that delivers dramatically faster, more accurate analytics—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 analytics tasks, for the rapidly expanding BI 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.

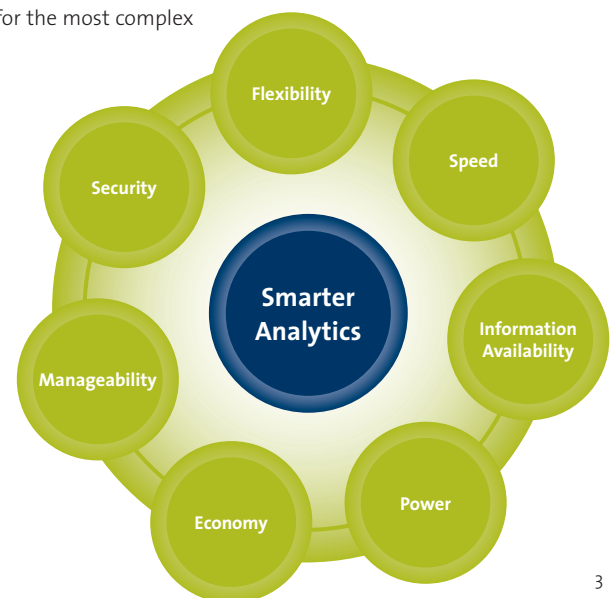
In-database analytics capability takes that performance to a new level while ensuring greater security of the data. Sybase IQ delivers in-database analytics three ways. First, enhancements to an extensive library of built-in numerical and analytical functions, including date conversion, greatly expand the readily available capability resident in Sybase IQ. Access to these analytics functions is standards-based, and the in-database logic is extensible by anyone with a good working knowledge of SQL. Additionally, built-in ANSI SQL OLAP extensions allow aggregation analysis on large data sets, yielding quick results for computations such as correlation and covariance. Moreover, new libraries of pluggable analytical algorithms from statistical and data mining software partners who have certified their products with Sybase IQ's in-database analytics capabilities, bring the entire range of statistical and predictive analytics capabilities into Sybase IQ.

“SYBASE IQ PROVIDES AN EFFECTIVE DECISION-MAKING SUPPORT SYSTEM FOR MANAGEMENT, AND HAS GREATLY IMPROVED OUR UNDERSTANDING OF PATTERNS. IT HAS ALSO IMPROVED OUR DECISION-MAKING CAPABILITY, ALLOWING US TO BE MORE CONFIDENT IN ADJUSTING TO FUTURE MARKET COMPETITION.”

GAO JIANBO, PROJECT MANAGER,
CHINA TELECOM

“WITHOUT SYBASE IQ, GAINING RAPID QUERYING SPEED AND AD HOC NAVIGABILITY WITH SUCH HUGE VOLUMES OF DATA WOULD BE HIGH IMPOSSIBLE—BUT THE CAPABILITY OF SYBASE IQ IS TRULY OUTSTANDING. THE RUN THE TYPE OF AD HOC QUERY THAT OUR USERS DELIGHT IN ON EVEN THE MOST POWERFUL DATABASE ENGINE AT THESE LEVELS COULD TAKE UP TO 18 HOURS. ON SYBASE IQ, THE SAME QUERY TAKES UNDER A MINUTE.”

JAN SPENCER, DIRECTOR OF SYSTEMS
DEVELOPMENT, AVIS EUROPE



Moreover, Sybase IQ includes infrastructure for analytics applications, making deployment of new analytics capability quicker and easier.

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: 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 a faster, efficient, cost-effective advanced analytics system 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">• In-database analytics capability• Column-based architecture for low I/O and fast retrieval• Query parallelism per node• Query speed 10X – 100X faster than traditional row-based DBs• Patented indexing technology
Economy	<ul style="list-style-type: none">• Supports all standard connectivity protocols• Supports all standard hardware platforms/disk technologies• Patented Data Compression (reduces storage needs up to 90%)• Configurable tablespaces
Manageability	<ul style="list-style-type: none">• Range partitioning• Index advisor utility guides index selection for top performance• XML & GUI-based query plan generator for diagnostics• Versatile GUI-based admin console for server management
Power	<ul style="list-style-type: none">• One-touch multiplex grid configuration• High performance from any schema• SQL logic deployable via 3GL, stored procedures, or cursors• Supports hundreds or thousands of users & petabytes of data• In-database analytics capability• Foundations for analytical applications
Flexibility	<ul style="list-style-type: none">• Independent compute power and storage capacity scalability• On-the-fly schema changes without query blockage or downtime• Handles structured and unstructured data• Works with industry-leading BI ISV products
Information Availability	<ul style="list-style-type: none">• Multi-node loading• Enhanced large table loads• Direct loads from clients• No impact to query performance while loading• Supports rules-based information lifecycle management
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• FIPS (on Unix), Kerberos, and IPv6 support