

Sybase® IQ and SAP Netweaver® Business Intelligence

Accelerate operational intelligence with SAP® Business Information Warehouse
using Sybase's high performance analytics server

TABLE OF CONTENTS

- 1 Executive Summary
- 1 Business Intelligence Transforms Data to Decisions
- 2 Data Warehouses Are Fundamental to “Operational Intelligence”
 - SAP Business Information Warehouse (BW)
 - Challenges for Data Warehouses
 - Data Explosion
 - Growing Numbers of Users
 - Data Access: A Key Success Factor
 - Challenges for SAP BW Users
 - Slower Applications
 - Storage Becomes a Problem
- 4 Integrating Sybase IQ with SAP BW
 - Accessing Sybase IQ Data from SAP BW
 - Sybase IQ Architecture
 - Optimizing SAP BW With Sybase IQ
 - Efficient Storage of Third-Party Data
 - Affordable, Fast Online Archives
 - Complex Statistical Analytics
- 6 Conclusion

EXECUTIVE SUMMARY

Enterprises are fueled by information. “Operational intelligence” is measured by an organization’s ability to get information to the people who need it to make decisions or react appropriately to business events. Enterprises that rely on SAP NetWeaver® Business Intelligence (BI) to harness the value of operational data are facing a number of technical challenges created by exploding volumes of data, growing demand for data among business users, and unprecedented regulatory and legislative requirements. The underlying technologies were simply not designed to efficiently handle today’s data and user volumes.

Sybase® IQ is a leading high performance analytics server that is now certified to integrate with SAP Netweaver BI. Sybase IQ is especially designed for operational intelligence—capable of query and reporting performance up to 100 times faster and with greater economies than traditional relational databases.

The combination of Sybase IQ and SAP Netweaver BI helps customers integrate data from within and outside the enterprise and transform vast amounts of data into valuable information that can drive better decision-making and situation-specific actions.

Sybase IQ has clear advantages as an analytics server, especially when it comes to handling extremely large volumes of data. Patented storage and management capabilities result in extremely fast query response times. The unique structure and algorithms in Sybase IQ reduce the amount of storage required by up to 70%. Sybase IQ also requires little tuning and maintenance for reporting and queries. All of these attributes result in dramatically lower storage, maintenance, and operational costs and increase productivity. SAP Business Information Warehouse (BW) users can take advantage of the power of Sybase IQ in a variety of ways:

- Outside data can be uploaded directly to Sybase IQ, for greater storage and analytic efficiencies
- Legacy SAP data can be stored in Sybase IQ. This enables users to build online archives at little cost, complying with statutory availability requirements and supporting information lifecycle management (ILM) concepts
- New evaluations with unknown scope—such as multiple-year time series analyses—can now be performed on extremely large data pools
- Ad hoc queries are up to a 100 times faster
- Improvement in operational application performance

By integrating Sybase IQ with SAP BW, SAP users can cut costs, dramatically step up query performance, and take advantage of an enhanced feature set.

Businesses that already use Sybase IQ are now afforded the opportunity to make a strategic decision in favor of SAP, while protecting their investments.

BUSINESS INTELLIGENCE TRANSFORMS DATA TO DECISIONS

What products were bought in the last six months? Who purchased them? Which sales agents were involved? What type of cross-selling potential is there? How does the sales trend break down by product group over the last five years? What do daily sales look like in each of my sales regions?

Over the years, companies have employed a host of business intelligence (BI) solutions in their quest for answers to these types of questions. The goal of these BI initiatives is to integrate and consolidate data from various internal and external sources and then consolidate and condense it into information that can serve as a basis for making decisions.

What started as a tool to support sales, marketing and customer service departments has widely evolved into an enterprise-wide strategic platform. While BI systems are used in the operational management of divisions and business processes, they are also used to support strategic corporate decision-making. The dramatic change that has taken effect in the last few years is the growth in demand for operational intelligence across multiple systems and businesses—increasing the number of people who need access to increasing amounts of data.

Business specialists need complex, multi-dimensional analyses, but businesses are also in search of key performance indicators (KPIs) which can be used by both departmental users and management. KPIs can be displayed on dashboards as graphical symbols such as traffic lights or thermometers. In addition, users want

real-time access to this data so that they can monitor processes with the smallest possible lag and intervene rapidly whenever KPIs deviate from their target values. To link strategic and operational perspectives, users must be able to “drill down” highly consolidated figures into detailed bits of information so that they can perform in-depth analyses.

It is hardly surprising that BI is becoming increasingly important. In 2004, demand for BI figured in the Gartner survey of the top ten CIO priorities for the first time. Two thirds of the population polled in a study carried out by InfoWorld and market researcher IDG in 2005 ranked BI as “highly important” to “critical.” According to this study, businesses invest more than \$8 billion US dollars in BI software every year.

DATA WAREHOUSES ARE FUNDAMENTAL TO “OPERATIONAL INTELLIGENCE”

Data warehouses are intended to create an integrated, uniform data pool to support business intelligence and analysis across the enterprise.

Enterprise-wide data warehouses compile and prepare operating data from diverse systems, platforms and data sources, including business applications (enterprise resource planning, customer relationship management, supply chain management, and other applications), databases, and other sources. These are often isolated data silos that have grown over the years. Ultimately, companies want an integrated, uniform view of all available data as a basis for extensive analyses.

SAP BUSINESS INFORMATION WAREHOUSE (BW)

In companies that use SAP systems, this role of data integration is increasingly assumed by the SAP® Business Information Warehouse (BW), a component of SAP NetWeaver BI. According to a survey of IT managers conducted by Munich Polytechnic in 2005, SAP BW is currently the predominant driver for NetWeaver, SAP’s new technology platform.

SAP BW is used to control, monitor, and maintain all processes employed to collect and store data. It is designed to provide support for processing large amounts of data, thus facilitating enterprise-wide data warehousing, reporting, and analyses. SAP offers its own ETL tools for the convergence, standardization, and enhancement of data from various internal and external sources. Data is extracted from operational applications or directly from relational databases, spreadsheets, or files, processed using pre-defined conversion rules, and readied for the data warehouse. Metadata repositories contain information on data properties such as sources, history, or technical characteristics.

Various approaches can be used in evaluating this basic data. SAP BW supports multi-dimensional models for complex OLAP analytics (InfoCubes), a two-dimensional operational data store (ODS) with small data volumes consisting of relational spreadsheets for the operational management of daily operations, as well as data warehousing for querying large volumes of historical data. Users access these models via SAP’s own BI tools (BEx Analyzer or Web Analyzer).

CHALLENGES FOR DATA WAREHOUSES

The extent to which data warehouses are considered successful hinges on data quality and, more importantly, the speed at which desired results are delivered. Performance is put to the utmost test by two phenomena: the explosive growth of stored data and the rapid increase in user demand.

Data Explosion

Data explosion is driven by the rising volume of operational data and the swelling databases required for long-term strategic analytics. Numerous companies integrate third-party data in their data warehouses, such as market and stock exchange information, some of which they source from data feed providers. In addition, they store increasing amounts of historical data to support sophisticated analyses.

Ever stricter legal requirements force businesses to store more and more historical data. Businesses are compelled to maintain data archives for longer periods of time and, more importantly, grant rapid access to them whenever the need arises. A major thrust in this direction came from the Sarbanes-Oxley Act, which was passed in the United States in July 2002. Besides being applicable to US-based corporations and companies with international operations and ties to the US, the Act introduced new standards affecting corporate governance in businesses around the world. A similar ordinance is set to be adopted via EU directives in Europe as well.

Additional regulations mandating intensified archiving are included in the International Financial Reporting Standards (IFRS) adopted by the EU in 2002 and the Basel II directives. Basel II requires borrowers to guarantee the availability of all information relevant to their ratings, and banks to keep comparative databases and capacity for ad hoc analyses available. On top of that, businesses are faced with a plethora of national record-keeping rules imposed by regulatory authorities.

This is why companies have increasingly turned to information lifecycle management (ILM) concepts, which aim to provide the right data, anytime, anywhere. ILM systems take the entire data lifecycle into consideration.

As a result, average data warehouse volumes climb 125 percent a year. Terabytes are no longer a rare unit of measurement.

Growing Numbers of Users

The second challenge is presented by the rapidly growing number of users who need access to data for analytical purposes. Since data warehouses drive both long-term strategic decision-making as well as numerous time-critical operating decisions, they are employed by an expanding number of people with different needs. According to SAP, over 80 percent of potential SAP BW users are information users, who put data warehouses to work mainly in the field of operational reporting.

New BI user groups are being tapped outside of companies as well. Based on the study by InfoWorld and IDG cited earlier, 60 percent of the businesses surveyed want to make their BI data accessible to customers as well, with a third and one fourth of them seeking to give access to suppliers and regulatory authorities.

DATA ACCESS: A KEY SUCCESS FACTOR

An extremely powerful infrastructure is needed to handle the increase in data volume and user figures. It is the prerequisite for efficient analytics, making it the key to data warehouse acceptance within the companies that deploy them. Fittingly, the OLAP Survey 2004 has technical challenges taking center stage. This year, for the first time, tool performance was placed above corporate policy and data quality on the list of the most severe problems faced by businesses.

Today's focus lies on data access, as evidenced by a survey of US and Canadian IT and business executives conducted by market researcher TechRepublic in summer 2005¹. The clear majority of those polled calls for:

- A reduction in query response time—from hours to minutes (90 percent agreement or significant agreement).
- Simultaneous access to both real-time and historical data (nearly 90 percent agreement).
- Faster access to information (over 90 percent agreement).
- Placing importance (average to significant) on accessing large amounts of data (over 90 percent agreement).
- Placing importance (average to significant) on data loading (88 percent agreement).

CHALLENGES FOR SAP BW USERS

The challenges faced by data warehouses in general are mirrored in SAP BW. Legal archiving requirements are growing rapidly. Many businesses store historical data both from the SAP system as well as from a host of third parties. For instance, pharmaceutical companies load information including drug lists and batch numbers to their BW databases. On top of that, this data is often arranged chronologically. External data frequently accounts for a higher share than SAP data.

This has two serious consequences for SAP users: analytics performance deteriorates, and storage becomes a problem from a cost and political perspective.

Slower Applications

Queries are forced to analyze far more data than is actually necessary. For instance, querying 20 days of sales can often mean sifting through database spreadsheets containing multiple years' worth of information. Traditional relational database technology is not optimized for analytics. During each task, the entire data set must be read, including the fields that are in fact completely irrelevant to the query. This highly inefficient use of computing resources prolongs query response times and overall application performance.

Storage Becomes a Problem

Average net data volumes have more than doubled in recent years. Moreover, conventional relational databases need added storage space for indices, aggregates, and other technical parameters. Usually, anywhere between 2.5 and 6 terabytes of storage space is occupied to store a net 1 terabyte of data in SAP BW.

As a result, storage costs rise extremely quickly. Data volume growth outstrips improvements in the price-performance ratio of magnetic storage media. Furthermore, every investment in data storage translates into additional costs for data management and maintenance over the hardware's entire lifecycle. Expenses skyrocket as the database structure must be monitored and optimized to accommodate oversized (and often superfluous) data. Not surprisingly, 75 percent of those polled in the 2005 TechRepublic survey cited previously rated storage cost reduction "important" or "very important."

Statutory data archiving regulations create additional problems for many businesses. Since companies are required to rapidly retrieve archived data on request, they can no longer make do with conventional archiving concepts. Traditional relational databases call for fast, expensive hard drives with DB engines for access control. This is why a number of companies use cheaper storage media driven by proprietary archiving system controllers to manage historical data. While more affordable, these solutions have limited access features, accompanied by far inferior performance.

SAP BW offers two options: the Archiving Development Kit (ADK) which stores data using archiving techniques, but forces data to be reloaded to BW for evaluation; and Nearline storage technology, enabling direct SQL access to historical data, eliminating the need to reload data to BW. However, the performance of these two systems cannot be compared to that of online systems.

New legal requirements force businesses to adhere to sophisticated ILM concepts—the more critical the data is to the company and its processes, the more accessible it must be. Archived data can often no longer be stored on hard disk-based long-term storage solutions or low-cost magnetic tapes, giving rise to the need for expensive online storage solutions, which cause costs to soar.

INTEGRATING SYBASE IQ WITH SAP BW

The certification of Sybase IQ with SAP Netweaver BI gives SAP users new capabilities for taming the data explosion. Sybase IQ, as an infrastructure-based solution to data management within the SAP Business Warehouse, is designed specifically for the query-intensive analytics that are fundamental to operational intelligence.

Using data models from the SAP® Business Information Warehouse (BW) data warehouse and analytics module, users can perform multidimensional analyses that include Sybase IQ data stores. Sybase IQ is accessed directly via the JDBC-enabled SAP BW Universal Data Connect 3.5 interface. This eliminates the need to physically convert and transmit data to the BW database—although this option is available to customers on request.

From version 3.5 onwards, SAP BW supports JDBC, the standard Java interface. It stages all data stored in Sybase IQ in SAP BW, without having to load it in the BW database or convert it. Furthermore, BW data can be piped into Sybase IQ and processed there to take some of the workload off SAP BW.

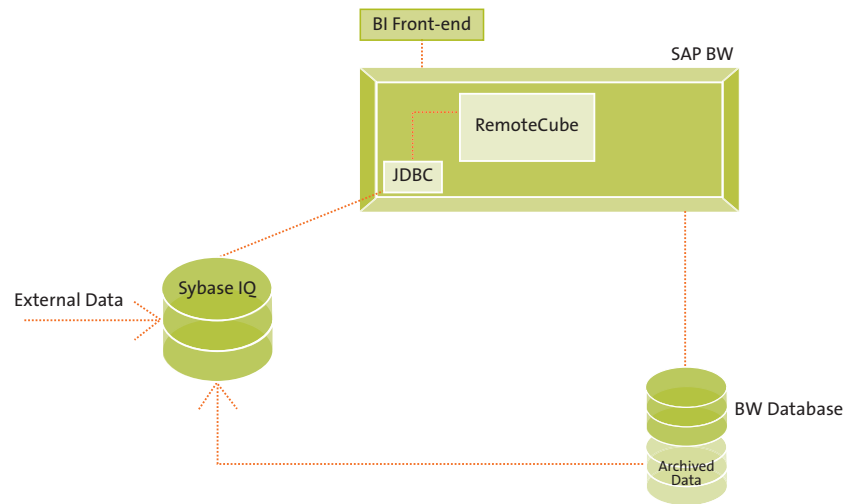
ACCESSING SYBASE IQ DATA FROM SAP BW

Data residing in Sybase IQ is accessed via the multi-dimensional models included in SAP BW (RemoteCubes). They enable data visualization from various perspectives including time, place, product, etc. Cubes are usually assigned to business management topic clusters. Multi-dimensional analyses can be used to find answers to complex business-related questions.

Moreover, MultiCubes enable simultaneous access to data in SAP BW and Sybase IQ.

Through their own user interfaces (workbenches), administrators define the SAP elements that make up the cubes. InfoSources (information units that form logical entities from a company perspective) and their data sources can now reside in Sybase IQ as well.

Users can invoke cubes using SAP's BI front-end tools or certified third-party products.



Architecture of an Integrated SAP BW and Sybase IQ Solution

SYBASE IQ ARCHITECTURE

Sybase IQ is designed to handle analyses of huge data volumes. It is a fully relational, SQL-savvy server which uses a patented internal storage and management technique.

Data is organized in columns instead of rows. This can be especially important in environments creating data with thousands of fields inside a row as is often the case in factory environments. Every field can be used as an index, eliminating the need to define conventional indices. Instead of reading the entire row, queries only read the value in the appropriate column. This reduces the number of I/O operations by more than 90 percent—negating the need for expensive, high performance disk storage. Furthermore, data is reduced by up to 70 percent of its original size for storage, significantly reducing data management operational costs.

The architecture of Sybase IQ results in the following advantages:

- Response times are accelerated substantially.
- Data can be loaded in real time.
- Ad hoc queries require no special tuning.
- A dramatically reduced storage footprint results in major reductions in administration costs—where conventional relational databases occupy 2.5 to 6 TB of storage space to accommodate a net 1 TB of raw data, Sybase IQ reduces this storage space to between 0.25 and 0.9 TB. In other words, to house the same amount of raw data, conventional data warehouses have to be 6 to 10 times bigger than Sybase IQ.
- Costs can be cut even further by using cheaper storage media. Sybase IQ is the only database product capable of running on low-cost ATA hard drives, while matching expensive enterprise hard disks for performance.
- Moreover, Sybase IQ is extremely scalable. The database can be used by a handful of users just as easily as by several hundred or thousand, and it is capable of handling infinite volumes of data, from a few gigabytes to several hundred terabytes.

Sybase IQ is a tried-and-tested technology that has proven itself in numerous installations over a number of years, assuring SAP users of a reliable and safe data management solution.

OPTIMIZING SAP BW WITH SYBASE IQ

Integrating Sybase IQ with SAP BW enables SAP users to save costs, while considerably speeding up queries. At the same time, they receive enhanced feature sets and can realize new information management initiatives.

There are several ways of taking advantage of Sybase IQ, all of which can be used simultaneously:

Efficient Storage of Third-party Data

External data no longer needs to be transferred to SAP BW, as it can be loaded directly into Sybase IQ where it remains available for any type of analysis via cubes.

Sybase IQ enables data to be managed and evaluated much more efficiently. Response times drop dramatically when complex analyses are run. High data compression cuts storage costs. Maintenance work is thus reduced as well. Data volumes, which are often quite large, can be loaded much more quickly.

In addition, the SAP BW database is relieved of some of its workload. Expensive storage space is freed up, and the BW database achieves performance improvements since less data must be read for analyses.

Affordable, Fast Online Archives

Businesses can easily transfer their historical data from SAP BW to Sybase IQ. This allows them to build online archives which can be rapidly accessed when the need arises.

Companies can thus cut both storage and maintenance costs. Data is compressed and then stored on low-cost media. In turn, this frees up storage space in SAP BW, which can be used for other tasks.

Businesses using Sybase IQ fulfill strict regulatory requirements regarding the availability of archived data. Moreover, they have a much easier time of introducing ILM concepts while achieving availability and cost goals at the same time.

Complex Statistical Analytics

Together, data within SAP combined with external data constitute a huge statistical base. Sybase IQ enables the efficient analysis of terabytes' worth of data, markedly widening the time spans that can be evaluated. Time series analyses of multiple years of data, which were completely unthinkable using conventional database technology, can now be carried out at high speed.

Businesses can now track sales over a large number of years, for example, through detailed, complex time series charts. Transaction data and external market information can be analyzed based on the user's choice of criteria, with a view to detecting trends and extrapolating forecasts. Retailers and wholesalers can identify hot sellers and the customer groups that purchase them, as well as links between sales successes and events. Financial institutions can perform much more efficient multi-dimensional simulations to forecast trends and plan future investment strategies. Telcos can evaluate orders, contracts, invoices, usage figures, customer care projects, etc., in order to obtain knowledge to drive product development. The list of possibilities is endless.

CONCLUSION

The power of Sybase IQ as an analytics server used in conjunction with SAP BW gives information-driven enterprise entirely new options for taming the data explosion:

- Dramatically improve the efficiency and performance of existing business intelligence infrastructures.
- Eliminate the need to make additional hardware or software investments to handle massive data volumes and growing analytical workloads.
- Perform multi-dimensional analysis of operational data with direct access to data stored in Sybase IQ.
- Lower the cost of managing and storing data on Sybase IQ.
- Perform analyses that were previously unthinkable or unaffordable with traditional relational database technology.

Sybase IQ and SAP BW give companies a more efficient, more powerful solution for creating operational intelligence out of the torrent of data and events that flow into and out of an organization each day.

For more information visit www.sybase.com/sap.

¹ TechRepublic: Information Management Trends in 2005. ©2005 TechRepublic, Inc. www.techrepublic.com