



Five Technology Trends for Improved Business Intelligence Performance

The demand for business intelligence data only continues to increase, putting BI vendors and practitioners under the gun to develop new technologies and improve existing processes in order to boost the performance of BI systems and maximize their return on investment.

The good news is that the surge in BI development, driven partly by the various mergers and acquisitions that have taken place in the industry over the past few years, is providing BI users with better tools than they previously had access to. And several major BI trends are helping savvy companies deliver the analytical data their business users need without clogging up system performance.

In this E-Book, data warehousing and BI expert Rick Sherman explains how you can take advantage of these trends to ensure that you get the right information to the right people at the right times, so they can make the best possible decisions.

Readers of the E-Book will learn how:

- In-database analytics can speed up data retrieval and analysis while also improving the security of corporate information.
- Analyzing data in memory may help reduce BI complexity and eliminate many of the performance bottlenecks faced by users.
- BI appliances offer a lower total cost of ownership yet still provide high-performance access to BI data.
- Data virtualization can simplify the process of integrating information from disparate sources and help to improve data quality.
- Operational BI enables companies to more closely entwine data analysis with ongoing business operations.

Sponsored By: **SYBASE®**

Five Technology Trends for Improved Business Intelligence Performance

By Rick Sherman, SearchDataManagement.com contributor

The demand for relevant, accurate business data never diminishes. Business intelligence (BI) vendors and practitioners are constantly under the gun to develop new technologies and improve old processes to help companies get the information they need for data analysis purposes. The good news is that the surge in BI development, which has been especially pronounced after all the mergers and acquisitions by vendors in recent years, is providing business users with better analysis tools than ever before.

Several significant BI trends are helping savvy businesses address the ever-increasing information demands pushing up the size of data volumes and clogging the throughput of BI systems. These trends include in-database and in-memory analytics, BI appliances, data virtualization and operational BI. Let's discuss them one by one.

In-database analytics

You've got a large amount of data, you need it to be completely accurate and you need to process it really quickly. A situation like that used to mean compromising among quantity, quality and speed. With in-database analytics, however, there's no compromise because the data is processed within the database, avoiding the data movement that slows response time.

You move your compute-intensive BI functions -- reporting, analysis and data mining -- to an enterprise data warehouse (EDW) built on an analytics database platform. Such platforms provide parallel processing, partitioning, scalability and optimization features geared toward analytic functionality. Their cornerstone is a shift from general-purpose relational databases to columnar databases designed specifically to handle analytic workloads.

This approach lets you consolidate your analytical data marts into the EDW. Not only are data retrieval and analysis much faster, but corporate information is more secure because it doesn't leave the EDW. Also, in-database analytics lets you share analytic code and models across the enterprise for ad hoc analysis.

Enterprises with intensive processing needs for functions such as fraud detection, credit scoring and risk management can use in-database analytics to get the high-throughput operational analytics they require. And all organizations will find that the in-database approach can help them make better predictions about future business risks and opportunities, identify trends and anomalies, and make informed decisions more efficiently and affordably.

In-memory analytics

In-memory analytics is an old idea with a new twist. It always has been possible to run programs in memory, but 32-bit architectures limited how much data could be processed, and the cost of memory was too high to justify it for most analytical uses. You also couldn't safely keep data in memory for long.

In-memory analytics helps eliminate many of the traditional BI complexities and performance bottlenecks, such as the time it takes to access disk drives and carry out I/O activities.

Now, that has all changed. Modern 64-bit architectures have a larger addressable memory space, and memory costs have declined to the point where doing analysis in memory is more practical. You can even keep data in memory indefinitely while backing it up externally in case your server goes down.

In-memory analytics helps eliminate many of the traditional BI complexities and performance bottlenecks, such as the time it takes to access disk drives and carry out I/O activities. Memory access speeds can be as much as a million times faster than those of hard drives. By processing in memory instead of on disk, companies can give business users the query performance they crave on tasks like dashboard analytics and performance management.

Performance isn't the only advantage. In-memory analytics can also save significant development time by eliminating the need to store pre-calculated data in OLAP cubes or aggregate database tables. Data elements are simply loaded into memory, with multiple tables joined by matching names. Enterprise reporting and analytics, as well as application maintenance, are also simplified.

BI appliances

Rapid growth in data volumes and the number of users and/or queries can cause companies to experience performance or cost issues with their data warehousing implementations. But BI appliances can offer a lower total cost of ownership (TCO) along with high-performance data access. As the name suggests, BI appliances combine hardware, a database and BI tools on a platform that's tuned for analysis workloads. The devices come in pre-configured packages offering ease of installation, simplified operations and smoother maintenance.

The architecture of BI appliances leverages low-cost commodity hardware and software to reduce costs, although they generally include proprietary components designed to improve analytical performance. The appliances also make it easier for business users to access and analyze information, paving the way for companies to use their data to become more responsive to market opportunities and to guide the development of new products and services.

While not a replacement for large-scale BI implementations, BI appliances may be a great antidote to the numerous data analysis "shadow systems" that business units and workgroups have built and proliferated in large companies. In small and medium-sized companies, where IT resources often are limited and budgets constrained, the appliances can provide a more cost-effective solution than a traditional data warehouse.

In addition, some BI appliances use open source software, which allows them to be aggressively priced with an even lower TCO. Nearly any tool that helps simplify BI is a good thing and may move us closer to the elusive concept of pervasive BI that has been the holy grail of our industry; BI appliances are a good example.

Data virtualization

It's not unusual for a company to have several database management systems as well as enterprise applications from different vendors. Pulling together data from all those systems can be a challenge, especially when some of

the information that the business needs is scattered in unstructured and semi-structured data stores. Data quality also is put at risk because data may not be consistent from system to system. The problems are exacerbated by the fact that data volumes are growing bigger and bigger all the time.

Data virtualization is an approach to data management that decouples information from applications and stores it at the middleware layer. You can then integrate data from multiple disparate sources for consumption by nearly any front-end business tool, including portals, reports, applications and search engines. The location and source of the data no longer matter; the idea is that if all applications access data on this single virtualized layer, they'll all be using the same version with no data consistency issues.

Because nothing is simple in our industry, there are different names for data virtualization: virtual data federation, high-performance query or enterprise information integration, information as a service, and data as a service.

No matter what you call it, with data virtualization, it's critical to make sure that the data you put on the virtual layer is the most "truthful" version. This can mean that the middleware layer is where your data quality analysis needs to take place. The foundation of this middleware layer is the metadata that describes the data from both business and technical perspectives to ensure consistency and enforce security.

Data virtualization is a proven data integration technique that can help you leverage your information assets while increasing the quality of your data.

Operational BI

Traditional BI does a great job of providing a backwards-looking view of the business. Data from disparate sources is brought together in a data warehouse for financial reporting and analysis. While this works for strategic planning and high-level decision making, it isn't always ideal for the rest of the enterprise. To make business intelligence capabilities more broadly accessible and applicable, many companies are moving toward operational BI, in which analysis is more closely entwined with business operations and processes.

The traditional approach involves gathering piles of data and analyzing it later. But if you wait to analyze an end-to-end process, it can be too late to benefit the business. With operational BI, analysis can take place in tandem with business processing, so problems can be spotted sooner. That enables the creation of a performance and feedback loop in which decision makers can analyze what's happening in the business, act upon their findings and then see the results.

All of the other trends mentioned above may be able to provide the missing technical ingredients for an enterprise wishing to delve into operational BI. The technology choice depends on a company's specific business needs and IT preferences. Regardless of the approach taken, the common thread is that an enterprise has to break free from the traditional bounds of both ERP reporting and data warehousing to offer the data currency demanded by operational BI.

With operational BI, analysis can take place in tandem with business processing, so problems can be spotted sooner.

A word of caution: Not every process in a company needs real-time data. In fact, most don't. With that in mind, you need to determine which users really require up-to-the-minute data, can justify the expense of providing it and can actually handle getting data delivered to them more than once a day. The others may get along perfectly well with historical data, or a combination of that and real-time data.

Conclusion

The inability to get the right information to the right person at the right time has plagued companies for years. As a result, they miscalculate inventory levels, fail to nurture repeat customers and miss the signs of preventable business problems. But by adopting some of the newer BI technologies and approaches, enterprises can arm themselves with tools that may be able to help them avoid those pitfalls and maximize the ROI of their corporate data.

About the author: *Rick Sherman is the founder of Athena IT Solutions, a Boston-based firm that provides data warehouse and business intelligence consulting, training and vendor services. In addition to having more than 20 years of experience in the business, Sherman is a published author of more than 50 articles, an industry speaker, a DM Review World Class Solution Awards judge and a data management expert at SearchDataManagement.com. Sherman can be found blogging at The Data Doghouse and can be reached at rsherman@athena-solutions.com.*



SMARTER ANALYTICS IS THE POWER THAT DRIVES THE NEW BUSINESS REALITY

Unlimited headroom for data and users. Incremental scalability to grow and adapt. Faster, more accurate answers to complex questions. Deeper insight into risk and opportunity. Agility to adjust processes and programs. Investment for today and tomorrow. That's Sybase IQ. High-performance analytics proven in over 3,100 unique installations at 1,700+ companies.

www.sybase.com/bi

SYBASE®

Resources from Sybase

[White paper: Sybase IQ Technology Overview – Bloor Research](#)

[Techcast: Sybase IQ 15: More Flexibility to Adapt and Grow](#)

[Webcast: Sybase IQ 15: Smarter Analytics](#)

About Sybase:

Sybase is an industry leader in delivering enterprise and mobile software to manage, analyze and mobilize information. We are recognized globally as a performance leader, proven in the most data-intensive industries and across all major systems, networks and devices. Our information management, analytics and enterprise mobility solutions have powered the world's most mission-critical systems in financial services, telecommunications, manufacturing and government. With Sybase, enterprises can manage the high volumes and variety of data, analyze this vital information and mobilize it so people can conduct business, regardless of where they are located or what devices they are using. And since all Sybase solutions are built using open standards, organizations can leverage the IT investments they depend on today, along with the ones they'll need tomorrow.