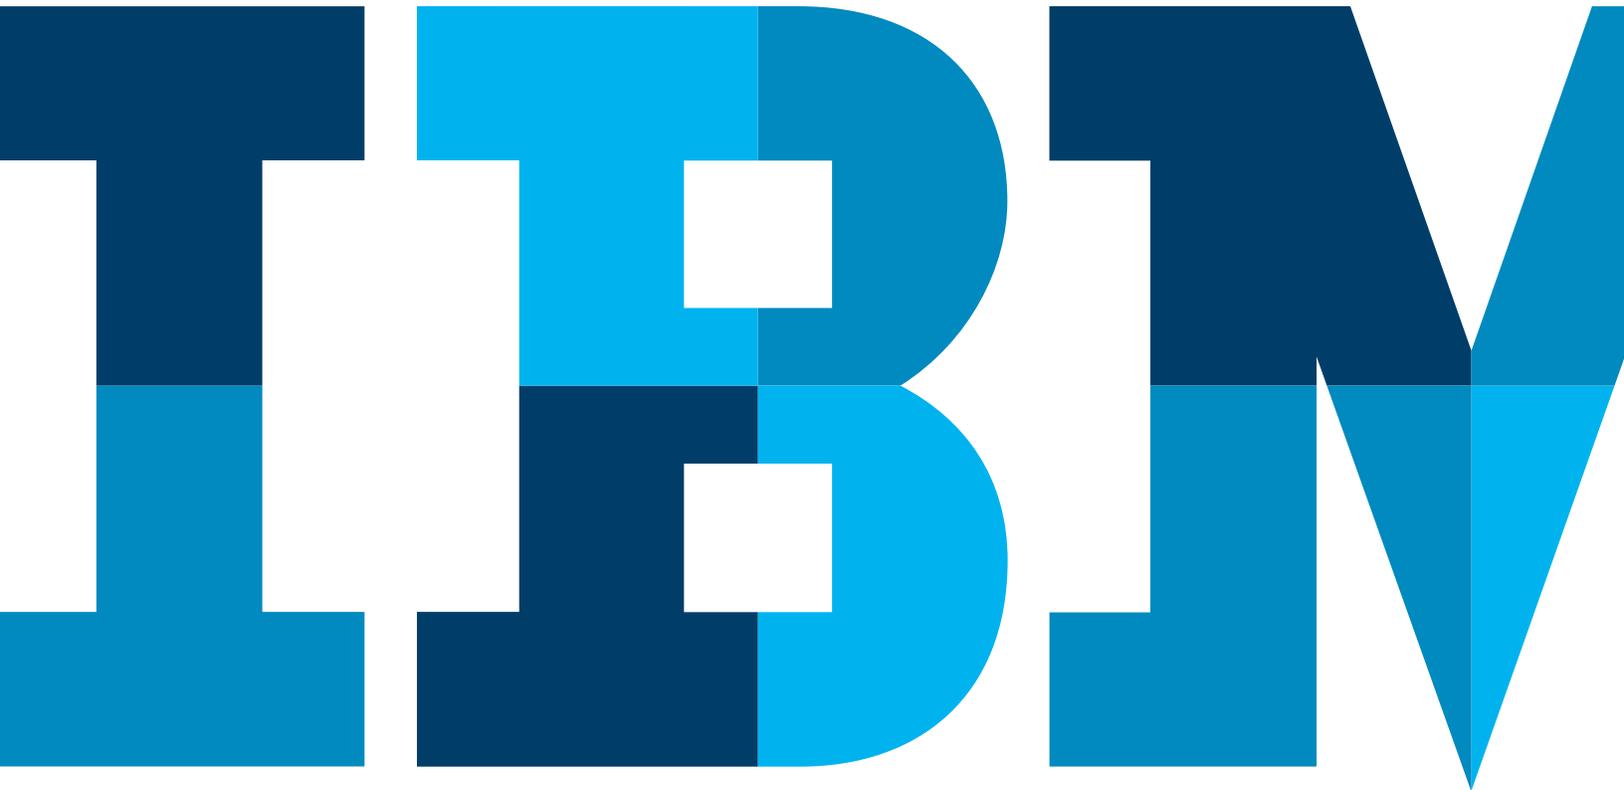


In-memory processing

Enhancing performance and analytics capabilities



Contents

- 2 Executive summary
- 2 Performance is a challenge in analytics
- 3 In-memory processing can provide better performance
- 4 Dynamic query: Flexible high speed query capability
- 4 Intelligent caching is the key
- 4 The advantages of dynamic query
- 5 Dynamic cubes: Technology for analyzing terabytes of data
- 5 Aggregates are key
- 5 The advantages of dynamic cubes
- 6 Cognos TM1: An “in-teractive” memory OLAP engine
- 7 Conclusion

Executive summary

The amount of data that users need to analyze for insights is growing by leaps and bounds. To make sense of that data, business analytics and business intelligence (BI) applications use a great deal of processing power, which can affect their performance. This is of little interest to business users, analysts and executives, all of whom expect to be able to access, analyze and process the data and intelligence they need as rapidly as possible.

To provide faster response times and more flexibility, your IT organization needs solutions that help reduce the bottlenecks that affect response times in analytical processing. These bottlenecks include disk input and output, network latency and

other activities related to accessing data from database tables and loading data into multi-dimensional cubes for querying. In addition, cube build times can be lengthy, which can affect the currency of the data being analyzed.

In-memory processing is a way to improve performance by greatly reducing the need for disk input and output. In some cases, it can even eliminate that need. Some forms of in-memory processing can even complete calculations more quickly, reducing the need to create aggregates or summaries in advance. And, finally, in-memory processing can power interactive visualizations of large and multidimensional datasets without requiring that you move the data into a specialized database.

But, one size does not fit all, and there are several options available for in-memory processing. The type of in-memory solution that fits a business need best is often determined by things such as business objectives, workloads, data sources and volumes of data.

This paper is intended for IT organizations that are researching options for faster query processing and want to understand how IBM® Cognos® software capabilities, underpinned by a dynamic query layer, use in-memory processing to accelerate performance.

Performance is a challenge in analytics

Business users today have more information at their fingertips than even before. However, even with this data, they continue to struggle to make the right decisions at the right time so they can head off the competition. Fast access to data, reporting and analysis can help them make timely and effective decisions. Therefore, business users want access to their data in “real time.” They are not interested in how they can get this rapid access; they simply want it to happen. They are seeking high-performance answers to their questions.

Most business users, however, are not getting the high performance they want. In 2011, Gartner analyst James Richardson reported that “Poor performance is prevalent in BI implementations. It is the third-most cited problem with disk-based BI platform deployments, with 21% of megavendors’ BI customers reporting this issue.”¹ Indeed, poor performance of BI software is a popular topic in blogs, and a 2010 Business Application Research Center survey of BI users identified poor performance and slow query response times as the second biggest problem with their BI deployments.² In other words, traditional database solutions are no longer sufficient to address BI application performance concerns.

In-memory processing can provide better performance

In-memory processing loads data and metadata into RAM instead of hard disks. In-memory is not a new concept; the concept of taking advantage of the speed of RAM has been around for some time. However, with the mainstream emergence of 64-bit architectures that enable larger addressable memory space and the rapid decline of memory prices, it is now more feasible for data to be analyzed in memory.

The Cognos platform includes a dynamic query layer that can use memory when an application requires it. With a 64-bit Java query engine, this dynamic query layer is the backbone of powerful in-memory capabilities that are flexible and that continue to expand. The query layer enables you to create applications that rely on optimized relational queries, optimizing the efficiency of communication with your data sources. It is designed so that you have the flexibility to choose how you want to manage data based on need. For example, you can choose to cache data from some queries in advance and opt to get the most recent data from other queries, all in the same report.

When dimensional analysis, such as drilling up and down and navigating hierarchical data, is more appropriate, the dynamic query layer of the Cognos platform provides dynamic cubes, which extend the query engine with extensive use of aggregates to further accelerate performance. Dynamic cubes are meant to help you optimize the value of enterprise data warehouses, which often have exploding data volumes. Although some solutions require all data to be moved in memory in a “brute force” model, dynamic cubes enable you to scale up and accelerate applications gradually as your requirements grow.

When read-only analytics are not sufficient, the dynamic data layer is the engine that brings IBM Cognos TM1® capabilities into the Cognos platform environment. With its powerful rules-based in-memory cube technology, Cognos TM1 enables what-if analysis, and it is the engine that powers planning and consolidation applications, all with in-memory performance.

IBM developed its in-memory technology not only to meet different types of in-memory needs, but also to make in-memory solutions available without major hardware upgrades. The Cognos Analytics Server (Cognos TM1) might be what you need for what-if analysis, scenario modeling and other applications that require the ability to write back to the database. Or, you might choose the 64-bit Cognos platform query engine because you need extra speed for ad hoc analysis of information in your operational databases. Dynamic cubes can work best for your larger sets of data and provide interactive analysis capabilities that take advantage of predefined aggregates or summarized information. These capabilities use memory for performance and can help your IT organization better support fast analytics.

Dynamic query: Flexible high speed query capability

Dynamic query refers to the execution of BI content using the Java-based extensible query engine in the Cognos platform. The best practices learned over decades of BI software development were applied to its design. Dynamic query retains result sets and metadata captured from optimized queries to relational and OLAP data sources in a 64-bit, in-memory cache. It can reuse these result sets and metadata to minimize the wait times for future requests.

The dynamic query layer was developed to meet requirements for interactive reporting and ad hoc analysis. It employs sophisticated, multiphase query optimization techniques and can dynamically alternate between SQL and MDX processing, depending on what best suits the scenario.

Intelligent caching is the key

Responses and queries in any distributed server architecture can be affected by delays when the environment is under increased load. Limiting excess communications and further distributing certain types of data and processing can help reduce latency. Intelligent caching uses in-memory optimization to increase query performance and data cache reuse. As a result, response times are faster, business users are more satisfied and IT workload is reduced.

In addition, the intelligent caching available in dynamic query is security-aware. The Cognos Business Intelligence query service ties the result sets it stores in memory with the security profile of the user who submitted the original request. Confidentiality is preserved because cached data is not shared between users for which different security settings were defined in either IBM Cognos Framework Manager or the underlining data source.

The advantages of dynamic query

Dynamic query has a number of advantages. The advanced in-memory caching and aggregation can reduce data warehouse workload. Users are provided with a consistent experience, no matter what the data source. In addition, simple cache administration and query visualization tools help reduce total cost of ownership. IT organizations can also take advantage of improved query performance with the reduction in query planning and execution, along with lighter database server workloads.

Who can benefit from dynamic query?

Dynamic query is likely to be the right solution for your organization if:

- You have low or medium data volumes and an application source that consists of one or several operational or transactional systems.
- You have a need to access the latest information stored in your data warehouse or even in operational systems.
- You have some queries that are able to use information that has been cached and others that need to hit the latest data.
- You have complex data security requirements.

For additional information and a more technical understanding of the dynamic query mode, visit:

http://ibm.com/developerworks/data/library/cognos/infrastructure/cognos_specific/page529.html

Dynamic cubes: Technology for analyzing terabytes of data

Dynamic cubes are an in-memory cubing technology in the Cognos platform that uses the dynamic query engine. This technology is designed to provide high performance for large data volumes. It can optimize query performance on terabytes of relational data in a star or snowflake schema, which is the recognized industry standard. To achieve high performance, dynamic cubes use the power and scale of the relational database and a combination of caching, optimized aggregates and optimized SQL.

Dynamic cubes add an in-memory relational OLAP component and a cube designer to the dynamic query mode server to provide a multidimensional view of a relational data warehouse with accelerated performance. Querying data with OLAP functionality enables drilling down or drilling up for improved data navigation. Dynamic cubes are built into the query service of the dynamic query mode to load data directly from relational data warehouses. Each cube definition represents a dimensional view of a fact table in a star or snowflake schema, and you can then use virtual cubes to join multiple cubes or facts. This provides access to a wider area of analysis, which is seamless to your business users. Data control is achieved by caching only the data that is required and by moving appropriate calculations and filtering operations to the database.

Aggregates are key

Aggregates make a significant difference in performance. They can reduce the need to compute summaries dynamically based on millions or billions of fact table data. Dynamic cubes can create in-memory aggregates and they are aggregate aware, which means they can route queries to in-database aggregates or summary. As a result, they achieve performance gains that improve the processes of running reports or navigating data.

In addition, the Cognos Aggregate Advisor analyzes performance logs and makes recommendations for performance improvements with optimized aggregates. In-memory aggregate recommendations can be immediately submitted for implementation without reauthoring or tweaking the model. As more recommendations or additional data assets are added to memory, performance can continue improving over time. This high-performance environment helps users quickly analyze and report on volumes of data for faster decisions and minimizes the need for IT to manually analyze which aggregates to create.

You can use in-memory aggregates in place of in-database aggregates, or you can use both, depending on the maturity of the database design and other considerations such as user load, data volumes or hardware availability. Using both can increase the performance of queries by spreading the load between the application and database tiers. How the aggregations are used and managed is transparent to business users, which do not change even if you change the aggregate strategy.

The advantages of dynamic cubes

Dynamic cubes can drive high-performance analysis of large data volumes without moving the data, while helping to reduce IT costs. When using aggregates, this high-performance environment enables business users to analyze and report on volumes of data up to 80 times faster than dimensional analysis without the benefit of aggregates. Users benefit greatly from in-memory data assets and aggregate awareness. IT benefits from the ability to address targeted performance objectives without reauthoring.

After a dynamic cube is created and published as a package, it can be accessed and reported on as if it were any other type of traditional OLAP-style database. In addition, dynamic cubes can be integrated into your existing Cognos infrastructure. Major capital outlays such as large hardware appliances are not required to take advantage of this in-memory solution. Instead, you can gradually grow your environment as needs dictate. In addition, no separate installation or license is required. Existing entitlements are sufficient.

Who can benefit from dynamic cubes?

Dynamic cubes are likely to be the right solution for your IT organization if:

- Your data source is a relational data warehouse that is structured in a star or snowflake schema.
- The business users in your organization need to analyze terabytes of data.
- You are seeking a solution with low latency and fast performance.

For more information about Cognos Dynamic Cubes, see the IBM Dynamic Cubes User Guide 10.2.0 at:
http://pic.dhe.ibm.com/infocenter/cbi/v10r2m0/nav/5_6

Cognos TM1: An “in-teractive” memory OLAP engine

Cognos TM1 is a high-performance OLAP engine that enables on-demand, in-memory analytics for data that changes frequently or is input by users. It is often used for real-time scenario modeling, what-if analysis and enterprise-wide budgeting and planning. Cognos TM1 can be thought of as an “in-teractive” memory solution because of its write-back

capabilities, instant recalculation based on data changes, intuitive gestures for data exploration and rich visualization. Data written to the server is immediately available to all users without any need to prebuild or preaggregate cubes. Complex calculations and consolidations can be done on the fly and are cached for optimized query performance.

Cognos TM1 is a database in and of itself; it does not rely on an underlying relational database. Data is modeled in cubes that are loaded entirely in RAM when the server starts up. Rule-based or consolidated values are calculated on demand and remain in memory for cached reuse until the underlying data is changed, at which point the engine recalculates the results. The in-memory nature of Cognos TM1 is what makes this recalculation routine very fast compared to disk-based alternatives.

What-if analysis processing and scenario modeling capabilities are key

The ability to examine data and find trends that can help predict the future is paramount to business success. With what-if analysis, users can quickly model scenarios based on factors that might change. The quick recalculation of an entire budget based on a changing driver might call for hundreds or even millions of calculated and consolidated values to be updated interactively, which can be too big a burden for disk-based applications. Because TM1 has all the data in memory, it can quickly recalculate new scenarios based on user input. The write-back capability means users are working interactively with the data and seeing their changes as they happen without having to wait for a rebuild of the data set.

The advantages of an in-memory OLAP engine and who can benefit

Cognos TM1 enables users to create as many what-if scenarios and plan versions as they like, all in near real-time. It also includes tools to create personal models and scenarios right on the desktop, with the ability to share them for greater collaboration. New calculations and metrics can be added on the fly, and users can see the results incorporated into their analysis and reporting immediately.

Who can benefit from Cognos TM1?

Cognos TM1 is likely to be the right solution for your organization if:

- You need to interact with data in memory for applications such as planning, budgeting and forecasting.
- You have servers with high memory capacity, medium volumes of data and highly volatile data that requires a great deal of recalculation.
- You need to generate what-if scenarios and predictive models for analysis.

For additional information about Cognos TM1 software, visit the Cognos TM1 information center:

<http://publib.boulder.ibm.com/infocenter/ctm1/v10r1m0/index.jsp?topic=%2Fcom.ibm.swg.ba.cognos.ctm1.doc%2Fwelcome.html>

Conclusion

IBM's in-memory processing technology offers a number of benefits, including:

- **Performance improvements.** Users can query and interact with data in memory, which has proven to be significantly faster than accessing data from a disk.
- **A cost-effective approach to data management.** IBM in-memory technology relieves IT of the burden of performance-tuning tasks typically required by data warehouses. In addition, it is possible to deploy these solutions without a major hardware upgrade.
- **A powerful analytics platform.** In-memory processing provides the power to conduct sophisticated, complex relational and OLAP analysis without sacrificing speed or performance. You can slice, dice and navigate data more effectively. And, if you use Cognos TM1 for in-memory processing, you are able to write back, which is a very special ability.

For more information about IBM's in-memory processing technologies, visit:

ibm.com/software/analytics/cognos/platform/
ibm.com/software/analytics/cognos/products/tm1/



About IBM Business Analytics

IBM Business Analytics software delivers actionable insights decision-makers need to achieve better business performance. IBM offers a comprehensive, unified portfolio of business intelligence, predictive and advanced analytics, financial performance and strategy management, governance, risk and compliance and analytic applications.

With IBM software, companies can spot trends, patterns and anomalies, compare “what if” scenarios, predict potential threats and opportunities, identify and manage key business risks and plan, budget and forecast resources. With these deep analytic capabilities our customers around the world can better understand, anticipate and shape business outcomes.

For more information

For further please visit ibm.com/business-analytics

Request a call

To request a call or to ask a question, go to ibm.com/business-analytics/contactus An IBM representative will respond to your inquiry within two business days.

© Copyright IBM Corporation 2012

IBM Corporation
Software Group (or appropriate division)
Route 100
Somers, NY 10589

Produced in the United States of America
September 2012

IBM, the IBM logo, and ibm.com are trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the web at “Copyright and trademark information” at www.ibm.com/legal/copytrade.shtml.

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

This document is current as of the initial date of publication and may be changed by IBM at any time. Not all offerings are available in every country in which IBM operates.

THE INFORMATION IN THIS DOCUMENT IS PROVIDED “AS IS” WITHOUT ANY WARRANTY, EXPRESS OR IMPLIED, INCLUDING WITHOUT ANY WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND ANY WARRANTY OR CONDITION OF NONINFRINGEMENT. IBM products are warranted according to the terms and conditions of the agreements under which they are provided.

1 “Need for Speed Powers In-Memory Business Intelligence,” Gartner Group,” 9 June 2011.

2 Jeff Kelly, “Survey: Poor data quality most common business intelligence problem,” TechTarget, 20 Sep. 2010.



Please Recycle