

IN-MEMORY ANALYTICS: Leveraging Emerging Technologies for Business Intelligence

FEATURING RESEARCH FROM GARTNER



IN-MEMORY ANALYTICS: LEVERAGING EMERGING TECHNOLOGIES FOR BUSINESS INTELLIGENCE

INSIDE THIS ISSUE

Introduction 1

Emerging Technologies Will
Drive Self-Service Business
Intelligence¹ 4

A recent Gartner research note highlights the advances in emerging technologies that will have a significant impact in making it easier to build and consume analytical applications. The first key finding highlights an important underlying architecture for enabling self-service analytics:

In-memory analytics will make it easier to build high-performance analytical applications against large data sets. Coupling in-memory analytics with interactive visualization will enable a broader class of users to explore data sets and discover insights.¹

We believe this research note reinforces the reality that some long held business assumptions are rapidly falling by the wayside. Chief among them is the belief that only the largest enterprises can have access to leading edge analytical capabilities.

While earlier generations of Business Intelligence (BI) and analytic tools were indeed sized, priced and configured to appeal to enterprise organizations, that's changing. Midsize enterprises that may have tried – and failed – to implement an enterprise solution, settled with using Excel-based spreadsheets for all their needs or, worse, opted themselves out of the market entirely, are no longer forced to do so. This is critically important for these organizations, as they often find themselves in the same competitive ring as the largest companies. They're forced, in a sense, to fight above their weight class and as a result require greater insight and agility to outflank their larger competitors. In-memory analytics will level the playing field for midsize organizations.

THE BENEFITS OF IN-MEMORY ANALYTICS

Traditional analytic tools run queries against a data warehouse with user queries being processed against the data stored on relatively slow hard drives. In-memory analytics leverages a significantly more efficient approach where all the data is loaded into memory. This results in dramatic improvements in query response and the end-user experience.





The in-memory approach is not new. The desire to take advantage of the speed of RAM has been with us for some time. Only recently, however, has the promise become a practical reality thanks to the mainstream adoption of 64-bit architectures that enable larger addressable memory space and the rapid decline in memory prices. Because of this rapidly changing infrastructure landscape, it is now realistic to analyze very large data sets entirely in-memory.

The benefits of in-memory analytics include:

- **Dramatic performance improvements.** Users are querying and interacting with data in memory which is literally millions of times faster than accessing data from disk.
- **Cost effective alternative to data warehouses.** This is especially beneficial for midsize companies that may lack the expertise and resources to build a data warehouse. The in-memory approach provides the ability to analyze very large data sets, but is much simpler to set up and administer. Consequently, IT is not burdened with time consuming performance tuning tasks typically required by data warehouses.
- **Discover new insights.** Business users now have self-service access to the right information coupled with rapid query execution to deliver new levels of insight required to optimize business performance. IT is no longer seen as a bottleneck.
- **Connect insight with action.** If the in-memory solution supports write-back capabilities, you have a powerful platform for building planning, budgeting and forecasting applications. You can then conduct “what-if” scenario modeling to understand the business impact of changes to key business drivers and respond rapidly by modifying plans directly.

IN-MEMORY ANALYTICS FOR MIDSIZE COMPANIES

IBM recently introduced an integrated reporting, analysis and planning solution purpose built for midsize companies. At the heart of IBM Cognos Express is an in-memory analytics server that delivers this power and flexibility, but without the cost and complexity of traditional data warehouse-based solutions. While IBM Cognos Express is new in market, the in-memory technology it uses has been derived from the proven IBM Cognos TM1 server, which was one of the first and most respected in-memory solutions (Applix TM1 was acquired by Cognos).

The unique capabilities of in-memory analytics provided by IBM Cognos Express include:

- **Read and write capabilities.** This is an essential requirement to connect your analytic insight with operational actions. With data entry and spreading capabilities, your plans, budgets and forecasts can now be built on top of an in-memory analytics server. You can now conduct “what-if” scenario modeling and make changes directly within the analysis tools.

“Slow query performance will stunt adoption faster than the buggiest code. Therefore we believe in-memory analytics will drive wider BI adoption, as it will be much easier for users to get the performance they demand for all analytical applications. There will be no need to wait for the IT bottleneck to break.”

– Source: Gartner RAS Core
Research Note G00152770
*Emerging Technologies Will Drive
Self-Service Business Intelligence,*
Kurt Schlegel, 8 February 2008



- **Centrally managed data, business hierarchies, rules and calculations.** These are all powered by the in-memory analytics server, which facilitates fast and precise data loads from a variety of sources to gain timely insight into key corporate data. Dimensions can also be manipulated as needed to change the hierarchy, delete elements, and alter element aliases.
- **Empower business users to analyze any combination of data.** This intuitive set of tools allows users at all levels of the organization to point at any source of data to model and build custom cubes and dimensions on-the-fly.
- **High impact visualizations.** A powerful Web-based analysis tool allows quick, thorough analysis of complex data by swapping, stacking and switching dimensions in any combination. With a range of high impact visualizations to support your findings, you can easily share business insights throughout your organization.
- **Extend and transform Excel.** You get to keep the familiar Excel front end, but also augment it with a powerful in-memory analytics engine. It's a perfect combination for multidimensional analysis and strategic planning tasks, enabling a new level of insight and action.
- **Designed for modern 64 bit architectures** to support very large data sets such as analyzing profitability all the way down to the SKU level.
- **Easy to install, easy to use, and easy to buy.** IBM Cognos Express now brings the power and capabilities of in-memory analysis within reach of all companies.

In-memory analytics makes this possible, delivering a level of performance on mainstream technology platforms that was once only available to large enterprises with big budgets and IT staff. It delivers the kind of flexibility today's front line workers, managers, directors and C-level executives need to manage their workload and drive the future of their business. That it accomplishes all this while maintaining a solid growth path is unheard of in this space. No one else offers a solution that starts where IBM Cognos Express starts – inexpensive and simple to implement, use and maintain – and ends up as large and capable as a fast-growing organization needs it to be.

Source: IBM

¹*Gartner RAS Core Research Note G00152770 Emerging Technologies Will Drive Self-Service Business Intelligence, Kurt Schlegel, 8 February 2008 (Complete Gartner research note included in this document, see page 4)*



EMERGING TECHNOLOGIES WILL DRIVE SELF-SERVICE BUSINESS INTELLIGENCE

This document describes how emerging technologies will make it easier to build and consume analytical applications. However, these innovations could undermine the authority of central IT in maintaining business intelligence (BI) standards.

KEY FINDINGS

- In-memory analytics will make it easier to build high-performance analytical applications against large data sets. Coupling in-memory analytics with interactive visualization will enable a broader class of users to explore data sets and discover insights.
- Integrating search into analytical applications will make it easier to index large amounts of structured data, and to deliver reports and information to users via a mechanism (search) that is much more intuitive than traditional ad hoc query and reporting.
- Software as a service (SaaS) will play a role, particularly in midsize enterprises, in delivering analytical applications, with very little required by IT.
- Service-oriented architecture (SOA) — as well as advances in visual development capabilities such as mash ups — will make it easier to build sophisticated analytical applications without deep experience in development and programming.

RECOMMENDATIONS

- IT shouldn't try to fight these emerging technologies by prohibiting them in an effort to enforce standards. This policy didn't work with spreadmarts and it won't work with these emerging technologies that have the potential to dwarf the spreadmart problem.
- IT should incorporate these emerging technologies into the standard BI architecture when possible to prevent business units from adopting them to create "rogue" analytic applications.
- However, these emerging technologies will inevitably introduce some new analytic applications built independently from a central BI architecture. In these situations, the central BI team needs to clearly communicate which performance measures should be used to run the business.
- Build a governance strategy that incorporates a potential explosion in the number of analytic applications. Such a strategy should include an inventory of analytic applications that includes clearly defined owners and use cases.
- Central BI teams that are overwhelmed with BI project requests should exploit these emerging technologies as a component of a self-service BI strategy to reduce costs and speed up delivery.



WHAT YOU NEED TO KNOW

Much of the innovation in the BI platform market will come from emerging technologies that make it easier to build and consume analytical applications. Lack of both end-user and developer skills is frequently cited as a major barrier when deploying BI applications. Anecdotal evidence suggests no more than 20% of users in most organizations use reporting, ad hoc query and online analytical processing (OLAP) tools on a regular basis. Emerging technologies, such as interactive visualization, in-memory analytics, BI integrated search, SaaS and SOA will help overcome this skills gap in both the construction and consumption of analytical applications. These technologies will help reach the 80% of users that aren't using analytical applications today, helping to grow the BI platform market and reducing the labor required by central teams to deliver BI projects. However, they will also introduce a potentially troublesome consequence by marginalizing the role of IT in delivering many analytical applications. Many of these technologies have the potential to dwarf the spreadmart problem, making it easy for rogue business units and even individual users to create their own analytical applications that scale bigger and look better than anything IT is building today. Organizations will need to reconcile the benefits of these technologies against the potential to undermine ongoing efforts to standardize BI architecture. The reality is that central IT has very little power to prevent independent business units (and users) from adopting these technologies. As with spreadsheets, the answer is not prohibition. Instead, organizations should look to incorporate these technologies into the standard BI architecture, and promote self-service BI as a means of delivering analytical requirements more quickly and with less centralized effort.

STRATEGIC PLANNING ASSUMPTION(S)

By 2012, emerging technologies such as interactive visualization, in-memory analytics, search, SaaS and SOA/Web 2.0 will marginalize IT's role in building BI applications.

ANALYSIS

With the increasing parity of core BI functionality such as OLAP, reporting and ad hoc query, many people are asking where the innovation in the BI area comes from. One answer to that question is ease of use. Despite the success of BI in many organizations, most users find BI difficult to consume. Moreover, most IT organizations are overwhelmed with BI requests to meet business requirements. In both cases, there is a need to make analytical applications easier to build and consume, to overcome this skills gap. More investment in training will help, but investing in certain emerging technologies such as interactive visualization, in-memory analytics, search, SaaS, and SOA and Web 2.0 is absolutely key.

Interactive Visualization will be quickly accepted over the next two years as a common front-end to analytical applications, driven by the ubiquity of rich Internet applications. Interactive visualization, with its attractive display, should be more widely adopted by those users who aren't accustomed to the grid style of analysis and reporting offered by relational databases and spreadsheets. By definition, interacting and exploring data using interactive visualization requires only the most intuitive of tasks, such as clicking on a pie wedge, or circling the dots on a



scatter plot. Since interactive visualization is actually fun to use — more fun than a spreadsheet, that is — it will drive much wider adoption. For the most part, interactive visualization can be added to a BI platform architecture without a significant investment.

In-memory analytics doesn't so much make it easier to consume analytical applications rather than make it easier to build them; essentially it removes the need to build a complicated disk-based performance layer, such as relational-based aggregates and multidimensional cubes. Business units won't need IT to build these complicated disk-based performance layers if they can use analytical applications running on top of an in-memory performance layer. It is difficult for IT to build disk-based performance layers for all user requests, and most users are unwilling to use BI without it. Slow query performance will stunt adoption faster than the buggiest code. Therefore we believe in-memory analytics will drive wider BI adoption, as it will be much easier for users to get the performance they demand for all analytical applications. There will be no need to wait for the IT bottleneck to break. While the initial investment will be costly from a hardware perspective, it will inevitably reduce overall costs in the long run, as fewer labor resources will be required to build summary and aggregate tables. For example, many SAP NetWeaver BI customers have spent hundreds of thousands on the BI Accelerator, which uses in-memory analytics. These customers plan to significantly reduce their emphasis on building aggregates to improve performance.

BI-integrated search implies using two use cases. The first, and simpler of the two, is to apply a search-based mechanism (instead of a hierarchical folder navigation structure) to the library of thousands of reports that exist in most BI content repositories in large organizations. This simple application of search technology will make it easier for users to find the information they seek if the reports already exist. This move could cut costs derived from calls to a BI help desk by end users looking for information. The second use case is to apply search to structured data sets to retrieve and explore data, even when the reports don't exist. Keywords from a search string can be mapped to an index of structured data sources — results can be returned and filtered very quickly (given the fast index). This would be a very easy and fast method of providing powerful ad hoc query analysis capabilities for data retrieval and exploration, without requiring much IT investment. However, it is unlikely that search would be sufficient beyond ad hoc query. It doesn't provide the broader range of functionality such as comprehensive report formatting, consistent metadata, scheduling and broader analysis capabilities. Moreover, the practise of applying search to create ad hoc reports is just emerging and has only been demonstrated successfully in a few dozen organizations. Nevertheless, these first and second use cases will make it much easier to deliver a limited set of BI functionality with significantly less effort than a comprehensive BI platform would require. Furthermore, the population at large has already voiced its affinity for search as an easy and fun-to-use technology (similar to visualization) that should drive much wider adoption of analytical applications.

SaaS, outside of Web site analytics, has yet to be widely embraced in the BI and performance management space, but we think this is likely to change in the mid market. Mid-market




companies lack the sunk costs that large enterprises have already invested in a BI infrastructure. They have similar requirements to integrate, report and analyze data from numerous systems, but they don't have the staff or infrastructure to pull it off. Plus, mid-market companies are often in volatile business cycles where revenues could grow quickly or come crashing down. All these traits make SaaS a suitable match for mid-market companies. The wider adoption of SaaS for various analytical applications will drive greater user adoption, particularly in the mid market. Customers seeking SaaS for BI and data warehouse (DW) functionality should negotiate contracts with as short a time frame as possible, paying monthly or quarterly. These short-term contracts will reduce much of the investment risk, and emerging BI SaaS vendors have the incentive to lower these barriers to encourage adoption. It isn't proven that SaaS is cheaper in the long run; in fact, it could be as expensive, if not more so, than building BI and DW internally. However, it does provide a much cheaper entry path to start delivering on a BI strategy, and offers an easier method of scaling back the investment if a tough economic climate diminishes the budget. SaaS will also have a dramatic impact on the BI strategies of large organizations, but it is less likely to be a wholesale replacement of a BI and DW infrastructure. Instead it will be the targeted delivery of analytic applications or even simple reports for a very specific vertical or horizontal subject area. We expect there to be an explosion of data aggregators that collect data from numerous competitors within an industry, normalize the information and provide it back to these same competitors in the form of external benchmarks.

SOA, coupled with a move toward model-driven architecture, based on a visual "drag-and-drop" development style, will make it easier to build BI applications. SOA-driven development will not be without IT. IT will be a major player in building SOA-based analytical applications. However, SOA will make it easier for BI to become embedded within a wider array of business applications and will therefore encourage wider user adoption. On top of that, a proliferation of "drag-and-drop" style BI application development — requiring few programming skills using tools such as LogiXML, QlikTech or Business Objects' Xcelsius — will drive a resurgence in departmental analytical application development. This will, in turn, encourage adoption and usage, but also has the potential to engender more rogue deployments that buck standards set by a central BI team in IT. There isn't much that central IT teams can do to stop this trend, so we recommend promoting visual-based development of BI to exploit as many of the positive impacts (such as less demand on central IT) as possible. At the same time we recommend minimizing the potential for undermining centrally-defined BI standards through collaboration between enterprise BI architecture teams and de-centralized teams building BI projects.

Overall, these five emerging technologies provide enormous innovation, making it easier for users to build and consume BI applications driving overall adoption with less of a strain on centralized IT resources. It should be noted that ease of use from emerging technologies is not the only innovation in the BI space. In particular, significant benefits will come from architectural innovations such as tying BI more closely to corporate strategy and embedding analytics in the business process.

Source: Gartner RAS Core Research Note G00152770, Kurt Schlegel, 8 February 2008



ABOUT IBM COGNOS BI AND PERFORMANCE MANAGEMENT

IBM Cognos business intelligence (BI) and performance management solutions deliver world-leading enterprise planning, consolidation and BI software, support and services to help companies plan, understand and manage financial and operational performance. IBM Cognos solutions bring together technology, analytical applications, best practices and a broad network of partners to give customers an open, adaptive and complete performance solution. Over 23,000 customers in more than 135 countries around the world choose IBM Cognos solutions. For further information or to reach a representative: www.ibm.com/cognos

Request a call

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

Copyright © 2009 IBM.

In-Memory Analytics: Leveraging Emerging Technologies for Business Intelligence is published by IBM. Editorial supplied by IBM is independent of Gartner analysis. All Gartner research is © 2009 by Gartner, Inc. and/or its Affiliates. All rights reserved. All Gartner materials are used with Gartner's permission and in no way does the use or publication of Gartner research indicate Gartner's endorsement of IBM's products and/or strategies. Reproduction and distribution of this publication in any form without prior written permission is forbidden. The information contained herein has been obtained from sources believed to be reliable. Gartner disclaims all warranties as to the accuracy, completeness or adequacy of such information. Gartner shall have no liability for errors, omissions or inadequacies in the information contained herein or for interpretations thereof. The reader assumes sole responsibility for the selection of these materials to achieve its intended results. The opinions expressed herein are subject to change without notice.