
Technical Article

In-memory analytics – strategies for real-time CRM

Received (in revised form): 18th April 2011

Olaf Acker

is a partner in Booz & Company's Frankfurt and Dubai offices. He focuses on business technology strategy and transformation programs for global companies in the telecommunications, media and high technology industries, working in Europe, the Middle East, Asia and North America. He is a regular speaker on strategy and technology innovation on leading industry conferences.

Florian Gröne

is a principal with Booz & Company in Berlin. He works with communications, media and technology players on go-to-market strategy, customer-facing operating models and enabling technology architecture. He leads the firm's CRM Center of Excellence in Europe. His research has appeared in the *Journal of Management Studies*, the *Journal of Advertising Research*, the *Journal of International Business Studies*, among others.

Adrian Blockus

is an associate with Booz & Company based in Berlin. He specializes in IT, technology and strategy transformation programs for telecom and technology companies. He is an expert in business intelligence and data warehousing, advising large organizations in functional and business strategy, as well as BI and DWH technology architecture.

Carsten Bange

is the founder and managing director of the Business Application Research Center (BARC), the leading market analyst for business intelligence and data management in central Europe. He consults regularly on business intelligence and data management strategy, architecture and technology selection.

ABSTRACT For years, the process of devising customer data queries and creating business intelligence reports has been a lengthy one. That is because the information needed must be pulled from operational systems and then structured in separate analytical data warehouse systems that can accept the queries. Now, however, we are on the brink of true 'in-memory analytics', a technology that will allow operational data to be held in a single database that can handle all the day-to-day customer transactions and updates, as well as analytical requests – in virtually real time.

Journal of Database Marketing & Customer Strategy Management (2011) 18, 129–136.

doi:10.1057/dbm.2011.11

Keywords: CRM; in-memory; customer analytics; database

Correspondence:

Florian Gröne
Booz & Company GmbH,
Anna-Louisa-Karsch
Strasse 2, Berlin 10178,
Germany
E-mail: florian.groene
@booz.com

The advantages of in-memory analytics are many: Performance gains will allow business users to retrieve better queries and create more complex models, allowing them to experiment more fully with the data in creating sales and marketing campaigns, and to retrieve current customer information, even while on the road, through mobile

applications.¹ The resulting boost in customer insights will give those who move first to these systems a real competitive advantage. Companies whose operations depend on frequent data updates will be able to run more efficiently.² And by merging operational and analytical systems, with their attendant hardware and software,

companies can cut the total cost of ownership of their customer data efforts significantly.

To drive the shift to this new technology, CIOs must make sure the business understands its advantages in terms of better customer intelligence and lower overall cost. To do so, they must make a strong business case for the transformation – always a challenge with business intelligence (BI) systems – including ease of use, better analytical reports and better decision making. And they must devise a governance strategy to manage the technology’s rollout and monitor its use.

ANALYTICS ON THE FLY

Consider the plight of a marketing manager at a large telecom who wants to address a customer segment with a targeted offer and launch an event based campaign. The challenge he faces is that he needs to analyze millions of data records that are spread across multiple systems. He would have to call IT and file requests for excerpts of data, they would need to be quality checked, aggregated and enriched with attributes for the segmentation. Eventually, he might get back a set of records that is not suitable to answer the initial business question and would need to start another iteration of this cycle.

Now, however, a new technology called ‘in-memory analytics’ lets that marketer answer the same question from a tablet computer sitting on his lap while a meeting with the head of marketing is taking place. In real time, he can segment the customers based on his needs, refine his criteria and do simulations on the results he gets. New ideas for marketing campaigns and new product offers can be discussed efficiently on a fact base. As a result, the marketing manager spends less money on non-effective initiatives, increases the success rate of his campaigns and improves the telecom’s top and bottom line.

Until recently, large customer relationship management (CRM) systems depended on two separate database domains: The operational database maintained the day-to-day, high-volume transactional data, whereas the analytical database took the data needed to perform specific customer analyses and stored it separately.³ As a result, it was impossible to run real-time queries against the most up-to-date customer data. Thanks to major advances in the speed, cost, and sophistication of storage and memory technology and in the power of processors, however, the promise of real-time analytics – through which business users can access the full set of operational data when creating their reports – is finally becoming a reality, as shown in Figure 1.⁴

By giving business users access to truly live customer data, in-memory technology will transform how companies analyze and use that data. As such, it offers three significant benefits over traditional data warehouses:

- *Performance improvements:* Because users can now interact with and query data in memory, response time and calculation performance is dramatically improved. This increase in performance allows end-users to run more complex queries and gives them better modeling capabilities, adding up to greater business value (see ‘Inside the Technology’).
- *Customer value creation:* In-memory analytics provides the technology foundation that – applied effectively – can give business users instant self-service access to the information they need, providing a new level of customer insight that can be applied for core customer-facing processes such as service request handling or up-selling and cross-selling. Another benefit is the potential of improving traditional customer service offerings. A leading Australian ISP, for example, was able to

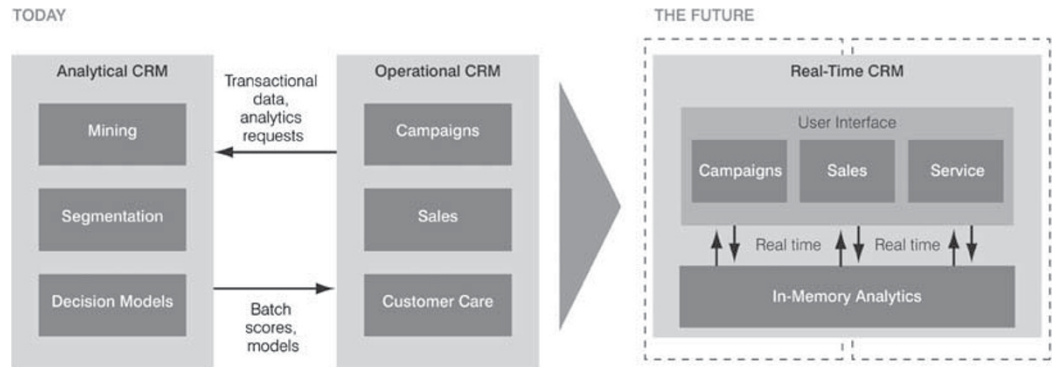


Figure 1: An integrated CRM architecture can speed up analytics requests.
 Source: Booz & Company analysis.

reduce 80 per cent of its technical customer service calls by an average of 1 min which in turn improved the customer service level.⁵

- **Lower costs:** Total cost of ownership and operations of storage infrastructure can significantly lower compared to traditional data warehouses, in part because more data can now be stored in one place, reducing data mart proliferation overhead. And although in-memory technology allows for the analysis of very large data sets, it is much simpler to set up and maintain. Rapid departmental deployments can free up IT resources previously devoted to responding to requests for reports.

INSIDE THE TECHNOLOGY

Until very recently, the effort to create, store and analyze critical transactional data related to all kinds of business activities was a cumbersome and expensive process. Operational data – the high-volume, transaction-heavy data generated through a variety of business processes, including sales, order management and customer care operations such as call centers – was maintained in huge data warehouses to ensure reliable performance and data integrity. And because the sources of the operational data typically varied significantly, maintaining it all in a single database with the homogeneous data model

that could serve as a ‘single point of truth’ proved very beneficial.

Meanwhile, the analytical data used to gather customer and performance insights, to segment customers, and to model and predict future behavior through customer usage and payment history, for example, was typically drawn periodically from the operational database and maintained separately. As valuable as that analytical data proved in boosting customer profitability and allowing more efficient up-selling and cross-selling, the architecture had some very real downsides.

Because it had to be duplicated periodically, the data in the analytical data warehouse was frequently at least a day or two – and sometimes as much as a month – out of date. A specially designed data mart had to be built for practically every new analysis request, which meant long deployment cycles, low project success rates and ever-growing data volumes at ever-higher cost. And the process introduced an additional layer of complex analytical software into the enterprise architecture, requiring additional training. Typical business users could only generate pre-defined standard analytical reports; anything more complex needed to be set up by a handful of power users.

Now, however, the once clear boundaries between operational and analytical databases

are beginning to blur – with true potential of convergence into integrated, real-time environments. With the emergence of multicore processors, increasing clock speeds and 64-bit technology, combined with the rapid decrease in the price of memory, data can be managed entirely in main memory.⁶ Although the idea of managing data in memory is not new, efforts to do so were hampered until recently by the fact that the old 32-bit architecture could address only 4 gigabytes of memory, and processors were not fast enough to give in-memory databases any real performance advantage.⁷ The transformation and convergence of database architectures will take time – but with new ways of column-based organizing, buffering and accessing the data, the performance improvements are significant.^{8,9}

The capacity capabilities of these systems are now equaling those of large disk-based databases. For example, a pilot implementation of a 40-terabyte in-memory database was recently completed, and theoretically, databases as large as 16 exabytes (16 384 petabytes) could be managed with in-memory technology,

based on today’s architecture. As shown in Figure 2, throughput is seven times higher, and the response time is virtually instantaneous.

DRIVING FACTORS

Several factors – involving improved analytical speed and performance and better analytical results – are driving the push to in-memory technology at the enterprise level.¹⁰

BUSINESS DEMAND

The delays that typically arise out of the periodic extraction, transformation and loading of data from the operational to the analytical systems may be generally acceptable in doing trend analysis and forecasting. But traditional data warehouses simply cannot keep pace with today’s business requirements for fast and accurate analytical data, especially in situations where mobility is becoming the norm. In every industry, customers now expect instant responses to their requests and questions; in this environment, in-memory technology allows companies to create an

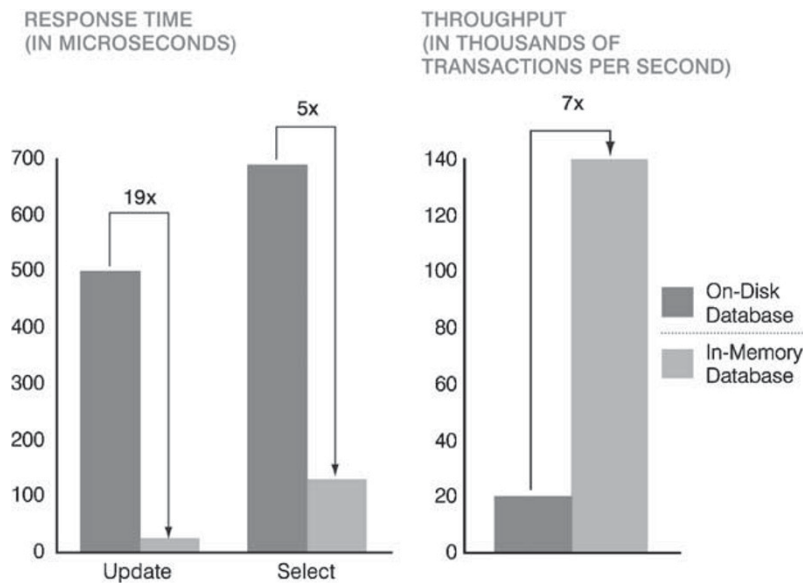


Figure 2: In-memory technology offers vastly superior response time and throughput. Source: Booz & Company analysis.



Figure 3: Selected benefits of real-time analytics across different industries.
Source: Booz & Company analysis.

entirely new level of customer experience, and it gives users instant access to the data they need to provide online self-service, real-time customer segmentation and dynamic pricing.

Time-sensitive industries like airlines and transport logistics will now have access to real-time information in running their operations, and the resulting increase in efficiency and agility can be leveraged to improve customer experience across different touch points, and optimize business cost structures (see Figure 3).¹¹

PERFORMANCE OF ANALYTICS

Most analytical applications have moved beyond the spreadsheets and tables offered by traditional reporting tools and now use interactive data visualization as the end-user interface, which allows many more people in the organization to make use of these systems. However, the new interfaces, which offer users interactive dashboards and the ability to perform much more intuitive tasks, demand very fast response times, as

users now expect instantaneous results. As in-memory analytics allows data to be accessed directly from memory, query results come back much more quickly than they would from a traditional disk-based data warehouse. The time it takes to update the database is also significantly reduced, and the system can handle many more queries, as Figure 4 shows.

GROWING DATA VOLUMES

The sheer amount of transaction data being digitally captured and stored is increasing exponentially, as are unstructured forms of data such as e-mail, video and graphics. According to one estimate, 0.8 zettabyte of data was produced in 2009 – if a gigabyte were the size of a sesame seed, a zettabyte would equal the diameter of the sun – and that is expected to rise to 35 zettabytes by 2020.¹² At the same time, tighter regulations involving the tracking of financial transactions and customer data put the onus on organizations to maintain this data and keep it available for years. Much of this data still resides on legacy systems,

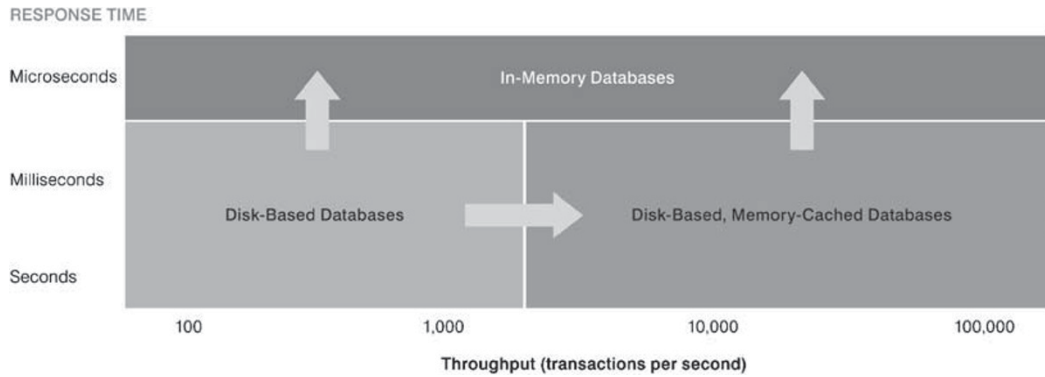


Figure 4: Performance comparison of different database types.
 Source: Booz & Company analysis.

which are costly to operate and maintain. In-memory analytics allows such data to be accessed rapidly on an *ad hoc* basis, without having to build additional complex data marts and load data into them. Instead, these systems allow users to connect to legacy data stores, populate an *ad hoc* database, conduct the analysis and then discard the in-memory database once the analysis is complete.¹³

SPEED OF DEPLOYMENT

Given the rapid growth of data volumes and the proliferation of applications dependent on databases, companies are struggling to manage the many BI efforts being developed throughout their organizations. In many cases, for instance, users simply want access to their specific transactional systems for reporting and analysis, without the need to deploy a full data mart. In-memory analytics removes the need to build complex performance layers such as multidimensional cubes within the data warehouse; instead, users can run their analytical applications directly against an in-memory performance layer.

GREATER INTELLIGENCE

In addition to the real gains in performance and speed offered by in-memory analytics, these new systems can significantly improve

the quality of the business and customer intelligence they generate. And they can transform how that intelligence is delivered, and to whom. The benefits include the following:

- *Improved decision making:* The ease of use of in-memory technology allows anyone in the organization, from business analysts to managers, to build their own queries and dashboards with very little technical expertise. Control over critical data shifts away from those who manage it to the stakeholders who own and use it, allowing them to make better business decisions.
- *Richer insight:* The significantly greater processing speed and calculation performance of in-memory technology lets end-users develop richer, more complex models, enabling better customer segmentation and more powerful campaign planning in the CRM space, for example. The result is significantly greater business value for the system as a whole.
- *Increased efficiency:* Converting to in-memory technology as a platform for analysis allows a whole technological layer to be removed from the enterprise architecture, reducing complexity and the infrastructure the traditional systems required. Furthermore, the source data has to be created or populated only

once and then is immediately available for any kind of analysis. Consequently, organizations can operate at a higher level of performance, deliver more reports per hour and free up capacity on the source systems for other operational purposes.

- *Self-service BI*: In-memory analytics allows any user to easily carve out subsets of the enterprise BI environment for convenient departmental usage. Work groups can operate autonomously without affecting the central data warehouse workload. In addition, in-memory technology enables a much greater degree of *ad hoc* analysis within the organization and allows users to source data rapidly, build analytical applications and conduct specific investigations. Once the analysis is no longer required, it can be disposed of easily. Quick response times and strong visual interfaces also enable mobile BI applications, which can be used by salespeople to gain a complete view of customers, based on real-time customer sales data, while on the road.

STEPS FOR THE CIO

The virtues of in-memory analytics are many, but as with any new technology, it is the responsibility of the CIO to make these virtues clear both to top management and to business users. Most important, in-memory analytics must be seen as part of a broader BI strategy that takes into account its overall business value and the underlying technology architecture, while remaining aware of the challenges inherent in every major new technology.¹⁴

The top priority is to educate the business as to the value and advantages of in-memory analytics, as well as to the costs and risks involved. In many organizations, business users of analytical tools have grown accustomed to relying entirely on the IT department to perform its ‘data magic in the basement’, a process that can take days.

Moreover, users frequently avoid using traditional BI tools because of their inherent complexity and difficulty. With in-memory analytics, as we have seen, response time can be virtually instantaneous, and users have the ability to design their own queries. It is critical to make clear to the business that a significant portion of the value of in-memory technology lies in its ability to open up these bottlenecks and offer users greater access to fresh data and increased query flexibility.

As part of the education process, CIOs should identify and point out particularly valuable business opportunities that in-memory technology offers. These might include the ability to go beyond traditional BI reports to create powerful applications such as what-if analyses, interactive filtering and pattern discovery, all in an easy, intuitive fashion. These capabilities should be actively promoted in order to foster a high-performance decision-making culture. But to accomplish this, many of the organizational processes with which both business users and the IT department are familiar will need to be rethought – an effort that, like any major change, must be planned and executed carefully.

Calculating the business case for any BI effort, traditional or otherwise, is difficult. The new technology will require a significant up-front investment in new storage hardware and software, and in the training needed for both the business and IT staffs to make the best use of it. Moreover, companies will need to maintain their old warehouses as they implement and test the new system, an additional up-front cost that must be taken into account. It will also be necessary to conduct a thorough cleansing of all customer data to avoid contaminating the new system with bad information.

Once such systems are installed, most companies will struggle to calculate the tangible cost-of-ownership benefits, such as overall infrastructure savings or lower

administrative labor costs. In order to build a better business case, CIOs must gain an in-depth understanding of the different types of BI applications and user segments involved, as well as the extent of the administrative maintenance effort required. In-memory analytics can then be better integrated into the overall BI tool strategy and positioned to either replace or complement current BI solutions.

The proper role of in-memory analytics in a company's overall BI architecture is critical. Do not try to convert the entire architecture to in-memory technology all at once. Instead, develop a thorough investment road map that includes both a plan for incorporating in-memory technology in the standard architecture when possible – in order to prevent business units from adopting it as part of a 'rogue' IT effort – and a strategy for switching over to the in-memory technology on a department-by-department basis.

Building a governance strategy that can effectively manage the potential explosion in the number of analytical applications is essential. Such a strategy should include an inventory of analytical applications that clearly defines owners and use cases and that can serve as the basis for a wider rollout of in-memory analytics throughout the organization. An established 'BI competence center' with the authority to drive standardization and exercise governance will be invaluable.

CONCLUSION

By combining operational and analytical databases into a single instantly available warehouse, in-memory analytics will give business users access to a whole new realm of crucial customer information, transform how they use that information, and thus

give them a real competitive advantage in the race to gain better customer insights more quickly. As with any BI effort, however, the new technology's virtues must be sold to business users, and its use must be monitored and managed carefully to ensure that all users are getting the most out of it.

REFERENCES

- 1 Gartner Research. (2010) The consumerization of BI: Emerging technologies to watch in 2011, <http://www.gartner.com>.
- 2 Plattner, H. and Zeier, A. (2011) *In-Memory Data Management – An Inflection Point for Enterprise Applications*. Berlin, Germany: Springer.
- 3 Peterson, M., Koch, V., Gröne, F. and Vo, K. (2010) Online customers, digital marketing – The CMO-CIO connection. *Journal of Direct, Data and Digital Marketing Practice* 11(3): 219–221.
- 4 Plattner, H. and Zeier, A. (2011) *In-Memory Data Management – An Inflection Point for Enterprise Applications*. Berlin, Germany: Springer.
- 5 ENP Newswire. (2011) Internode deploys Oracle coherence for faster customer service resolution. Press Release, 7 February 2011.
- 6 Plattner, H. and Zeier, A. (2011) *In-Memory Data Management – An Inflection Point for Enterprise Applications*. Berlin, Germany: Springer.
- 7 Boncz, P.A., Kersten, M.L. and Manegold, S. (2008) Breaking the memory wall in MonetDB. *Communications of the ACM* 51(12): 77–85.
- 8 Clementsen, D.S. and He, Z. (2010) Vertical partitioning for flash and HDD database systems. *Journal of Systems and Software* 83(11): 2237–2250.
- 9 Deng, Y. and Zhou, J. (2011) Architectures and optimization methods of flash memory based storage systems. *Journal of Systems Architecture* 57(2): 214–227.
- 10 BARC Business Application Research Center. (2010) The BI survey 9, <http://www.bi-survey.com>.
- 11 Peterson, M., Gröne, F., Kammer, K. and Kirscheneder, J. (2010) Multi-channel customer management – Delighting consumers, driving efficiency. *Journal of Direct, Data and Digital Marketing Practice* 12(1): 10–15.
- 12 Jacobs, A. (2009) The pathologies of big data. *Communications of the ACM* 52(8): 36–44.
- 13 Cummings, J. (2009) Thinking outside the storage box. *Network World* 26(4): 27–30.
- 14 Gartner Research. (2010) Trends that business intelligence leaders can not afford to ignore, <http://www.gartner.com>.