



E-Book

Tips on evaluating, deploying and managing in-memory analytics tools

In-memory analytics technology runs queries against data stored in a computer's memory, potentially providing much faster analytical performance than traditional business intelligence tools that need to pull information from disk drives. And with the growth of 64-bit architectures and reductions in memory prices, in-memory analytics has become a more viable BI option. Interest in it is rising as a result – but before your organization takes the in-memory plunge, make sure you fully understand the technology and how to be successful with it. In this eBook, IT and BI professionals will get an overview of in-memory analytics along with comprehensive advice on implementing and using in-memory tools.

Sponsored By:





E-Book

Tips on evaluating, deploying and managing in-memory analytics tools

Table of Contents

[In-memory analytics offers potential BI speed boost for business users](#)

[Flexibility of in-memory business intelligence requires some vigilance](#)

[In-memory processing for analytics: Can you make the business case?](#)

[Challenges lurk on in-memory BI projects – proper preparation needed](#)

[Resources from Tableau Software](#)

In-memory analytics offers potential BI speed boost for business users

By Alan R. Earls, SearchBusinessAnalytics.com Contributor

In-memory analytics is getting a lot of attention from business intelligence (BI) users and vendors alike these days, and that isn't surprising. In-memory tools look at data sets stored entirely within a computer system's memory instead of requiring information to be pulled from disk drives, with the potential to dramatically reduce I/O cycles and thus greatly improve query response times for business users.

In addition, some forms of in-memory analytics software can simplify the data analysis process by doing away with the need for the indexes, aggregated tables and multidimensional cubes commonly required by traditional data warehousing and BI applications. That can ease the data management workload for IT teams and make it easier to develop and run queries, opening more analytic activities to end users without specialized skills and creating opportunities to implement self-service BI capabilities.

BI vendors offer a variety of in-memory analytics tools with vastly different architectures, ranging from spreadsheet-based products to high-end platforms that can handle large amounts of data. But the common thread between them, according to BI industry analysts, is speed. "In-memory is for people who really need superb performance and very low latency – but that is not most people," said Gartner Inc. analyst Merv Adrian.

Joshua Greenbaum, principal analyst at Enterprise Applications Consulting in Berkeley, Calif., said the expanding interest in in-memory analytics is the result of several converging developments in the IT industry. Some software vendors pioneered in-memory processing as much as 10 years ago, but the technology met with limited success until recently, Greenbaum said. One thing that is now helping to boost its adoption is the fact that memory has become "insanely cheap," he noted.

The shift toward 64-bit systems has also helped, since they can support vastly more than the 4 GB addressable memory limit of 32-bit machines. Also aiding the in-memory BI cause are high-performance but reasonably priced blade servers that work well for analytic

processing and can be easily combined to produce powerful multi-node systems, according to Greenbaum. Yet another factor, he said, is the development and growing use of columnar databases as alternatives to conventional row-based relational databases.

“Columnar technology has really matured, so putting that together with the hardware developments has created a perfect storm [for the in-memory approach],” Greenbaum said. He added that storing information in columns instead of rows enables substantial levels of data compression, making it more feasible to load large data sets into memory. It’s possible to run mainstream relational databases in memory as well, but Greenbaum said they don’t provide the increased compression efficiencies of columnar software.

Bring on the data for in-memory analytics

Because of the potential advantages of in-memory analytics, Greenbaum is seeing interest in the approach from large retailers, financial services firms and utilities as well as the U.S. military and other government agencies. For such organizations, “the use case is lots of data and the need to efficiently and quickly process that data,” he said.

Forrester Research Inc. analyst Boris Evelson thinks in-memory analytics is an important element in a broader trend toward making BI more central to how databases are designed, implemented and used. In addition to enabling business users to get information “at the speed of thought,” in-memory applications that minimize or even eliminate the need to aggregate data and build schemas can foster a more agile analytical process that is better able to adapt to changing business requirements than traditional BI is, Evelson said.

With such tools, “you bypass all those steps” involved in optimizing data for analysis and reporting, he added. “Your reports and dashboards are one and the same as the data models. Just by changing the reports or dashboards, you change the data model – it isn’t separate.”

However, Evelson cautioned that the functionality offered by in-memory analytics vendors ranges widely. Before making any buying decisions, organizations should be sure they know about the different categories of in-memory tools and what each is best suited for, he said.

He also recommends asking vendors about specific product features and whether IT will need to be involved in the setup and maintenance process.

A frequent point of confusion regarding in-memory business intelligence is the fact that it isn't just a matter of eliminating disk drives, Evelson. "One of the most common questions I get is whether in-memory can be accomplished by simply running analytics on solid-state drives," he said. "That will certainly be faster than relying on disk, but you're retaining all the data modeling and all the I/O steps."

Even so, IT and BI managers shouldn't discount solid-state drives as a potential option, said Julie Lockner, an analyst at Enterprise Strategy Group in Milford, Mass. "You could get the answer you need in milliseconds from an in-memory system, but running a similar query through solid-state drives might get you the response in microseconds," she said. And for some organizations, that likely will be good enough, Lockner added.

All the benefits of in-memory analytics, without the usual limitations.



Analyze large amounts of data in minutes.

Tableau's evolution of in-memory technology optimizes across all the components of your PC so you can work with massive amounts of data, faster.

Get blazing fast query responses in seconds or less.

- Unlike solutions that force your data in-memory, it's your choice—switch between in-memory and live anytime.
- Traditional in-memory solutions force you to filter your data down to what can fit in RAM. Not Tableau.
- Once you load data into the Data Engine, Tableau starts instantly. No waiting to load data in-memory every time.
- It just works: there's no programming required to use the Data Engine.

Tableau is changing the way companies are analyzing and sharing their data.

Learn more at www.tableausoftware.com/inmemory



Flexibility of in-memory business intelligence requires some vigilance

By Alan R. Earls, SearchBusinessAnalytics.com Contributor

In-memory analytics holds out the potential of providing powerful data analysis capabilities to business users without requiring much in the way of IT support for building queries and creating reports. But there is a possible dark side to the in-memory business intelligence (BI) equation to keep in mind, according to analysts.

And it could be summed up in this way: With the added capabilities comes increased responsibility for end users.

The schema-less approaches supported by some in-memory analytics tools offer users lots of flexibility for querying and analyzing data, said Forrester Research Inc. analyst Boris Evelson. In traditional BI applications in which data is pulled from disk drives, the fixed data models, star schemas and aggregated tables built by IT and data warehousing teams control how information can be utilized as part of the analytics process. By contrast, Evelson compares in-memory analytics to the use of Excel spreadsheets.

"It's up to you how you use it," he said, referring to business users. "But if you don't know what you're doing, it can be a problem. You can make mistakes."

Products at the high end of the in-memory processing scale can further exacerbate that issue. With them, "the real challenge is fundamentally that this kind of firepower and the quantity of data it can handle require companies to take a different look at their business analytics processes," said Joshua Greenbaum, principal analyst at Enterprise Applications Consulting in Berkeley, Calif. He added that for many organizations, dealing with cultural issues can be the most difficult aspect of successfully using in-memory software.

Often, Greenbaum noted, companies don't have a huge base of experience for dealing with the potential ramifications of in-memory BI. He cited the example of a large water and sewer utility with which he recently worked as part of a consulting engagement. The organization currently doesn't use "smart" meters, but it plans to implement them in the

near future. "Once they have that technology in place, with more data and more up-to-date data, the question is what will they measure and what will they do from an analytics standpoint," he said. "It's not intuitive – you must think about it."

Putting in-memory power to good use

The same question applies to the adoption of in-memory analytics, he added: "What do we do now? That is the real crux of the problem. With in-memory, you have to figure out what this tremendous power means to running your business."

And the decision to implement in-memory analytics applications should be tied to specific business problems, said Mike Ferguson, managing director of Intelligent Business Strategies Ltd., a U.K.-based research and consulting firm that focuses on BI, data management and data integration technologies.

"I think the use of in-memory analytics is still in its early days – it's not what you would call mainstream yet," Ferguson said. But he sees in-memory technology as a natural next step for many BI applications, such as complex event processing (CEP) systems that capture real-time or near-real-time streams of data. In-memory BI tools that avoid round-trip I/O excursions to disks can enable users to "analyze data faster and be more automated and responsive, which could yield a significant return on investment," he said.

In a manufacturing company, for example, if a large customer cancels or changes an order, that can have implications for production schedules and inventory levels. "It may force you to reschedule work for other customers, so being able to respond quickly may avoid potential business disruptions," Ferguson said.

However, he pointed to another potential problem: In many cases, it's hard to determine in advance the full value that an in-memory analytics deployment will provide to an organization. "You couldn't necessarily say how much retail banking risk management is improved by in-memory, though it is likely to make a contribution," Ferguson said, adding that he sees the technology "as part of the natural, evolutionary improvement in database and analytic practice rather than as an instant game changer."

The big picture is that in-memory deployments generally are doable, Evelson said. But even though in-memory tools should be more user-friendly than conventional BI software is, he thinks it's still a good idea for IT to keep its hand firmly on the tiller when it comes to issues such as security and data cleansing.

In-memory processing for analytics: Can you make the business case?

By Alan R. Earls, SearchBusinessAnalytics.com Contributor

Once you've decided that in-memory analytics could be a valuable technology for your organization, the next steps are building a business case and then evaluating and selecting software.

While in-memory processing has great potential for business intelligence (BI) uses, Boris Evelson, an analyst at Forrester Research Inc. in Cambridge, Mass., warns that organizations should take a hard look at whether it's really necessary for their operations. "If you're at a traditional, fairly slow-moving business, where things don't change very fast, you may not need in-memory," he said. "However, if your business is fast-paced and changes frequently – particularly in front-office areas such as marketing – in-memory can be very helpful."

Evelson recommends thinking of the business case for an in-memory BI project partly in terms of increased self-sufficiency for business users. Everything that can be done with in-memory tools can also be done with traditional data warehousing and BI systems, he said – but in most cases, the latter approach requires more formalized processes and greater involvement on the part of IT, BI and data warehousing professionals.

"What is the value to your organization," Evelson asked, "of not having to engage a professional programmer [to work on an analytics project]? And what are the things he or she could do if they weren't tied up with that traditional analytics technology?"

Going from those kinds of considerations to selecting in-memory tools can be complicated. In a recent Forrester report, Evelson pointed out the wide variations in functionality and features offered by in-memory analytics vendors and said it's important to understand the distinctions between the different types of technologies as well as their pros and cons.

Some of the specific issues he cited for evaluation include the extent to which different products are able to handle data sets too large for a given memory space; whether an in-

memory application's database can be accessed by BI tools from other vendors; and whether in-memory software uses the multidimensional cubes common in traditional online analytical processing or instead contains a fully loaded index "that does not require a fixed OLAP model."

A proper mix for in-memory processing

To ensure that an in-memory processing investment delivers useful analytical results, the system should provide business users with rapid access to data across a variety of sources, recommends Rob Fiorillo, senior vice president for BI at Cincinnati-based consulting firm Itelligence Inc. "It should be very open and not focused on just one data source," Fiorillo said.

He added that in-memory business intelligence software should be tightly integrated with an organization's conventional BI tools to provide a combination of analytics capabilities. And when comparison-shopping between in-memory products, Fiorillo suggests looking closely at their data compression ratios; companies "should expect to get anywhere from a 20:1 to 50:1 reduction" in space requirements via compression, he said.

Julie Lockner, an analyst at Enterprise Strategy Group in Milford, Mass., said that most of the organizations she has had contact with are trying to push the BI envelope and become faster at analyzing data and more responsive to business developments. But she added that in-memory analytics projects need to be thought through, beyond simply saying the technology makes sense for an organization. "Your business requirements should come with cost-justifications before you even consider in-memory," Lockner said.

Moving on to starting an implementation, Lockner advised that an in-memory analytics project should be structured much as you would approach any development program, with agility as a key goal. She recommended the use of agile development methodologies that provide end users with features and functionality on an incremental basis, instead of the traditional waterfall development approach. The project team needs to be skilled in programming as well as analytics, architecture and perhaps even statistical analysis in order to tackle the deployment with maximum effectiveness, she said.

However, Lockner noted that people who have all the required skills don't come cheap: Typically, they can command salaries of \$150,000 and up, or hundreds of dollars per hour for external consultants. Like other types of BI initiatives, she said, in-memory analytics deployments can get expensive – another good reason to build a solid business case before getting started.

Challenges lurk on in-memory BI projects – proper preparation needed

By Alan R. Earls, SearchBusinessAnalytics.com Contributor

Like other types of technology projects, deployments of in-memory analytics software present some potential challenges. But IT analysts and consultants say that for well-prepared organizations, those challenges shouldn't be too daunting.

Steve Kent, director of business intelligence at Burlington, Mass.-based Collaborative Consulting, said one of the big requirements for successful in-memory BI initiatives is avoiding data quality and integrity issues. "The key is to have good information," he said. "You must implement [data quality] business rules and have clean data if you are going to get good information out of the process."

Managing the relationship between in-memory analytics tools and data warehouses is another area that requires proper attention, according to Kent. And there are basic data management issues to take into account. For example, some in-memory applications eliminate the need for data aggregation and the building of multidimensional cubes to support data analysis. But, Kent said, "in the case of extremely large data sets – for example, in the pharmaceutical field – you might still want to aggregate data."

Another key to success is doing your homework on business needs prior to an in-memory analytics deployment, said Julie Lockner, an analyst at Enterprise Strategy Group in Milford, Mass. In her view, organizations should work to make sure they have a solid set of business requirements that, barring unforeseen circumstances, won't change too much during the course of a project.

"Everyone is talking about agility and BI, and agile data warehouses and agile analytics, but being agile is costly," Lockner said. If you know in advance that your business needs are likely to change frequently, she added, you should try to create a development environment that will make it easier to manage those changes. Overall, Lockner advised, it's important for an organization as a whole to realize the level of investment that will be required to make in-memory processing work for BI and analytics.

Once you've got the technical and business-requirements details mastered, the challenge of training is next, right? Actually, said Forrester Research Inc. analyst Boris Evelson, one of the premises of in-memory business intelligence is that in many cases it shouldn't involve as much end-user training as conventional BI technologies do.

Know what you're doing on in-memory BI

"If your users know their data, and if they know something like Excel, in-memory should require as little training as possible," Evelson said. To help avoid any potential problems, he added, in-memory analytics should be targeted primarily for use by more experienced business users – particularly if you're implementing one of the higher-end forms of in-memory software. "It isn't something to give to someone who doesn't know what he or she is doing," he said.

On the other hand, there may be more training and education issues for IT professionals themselves, said Mike Ferguson, managing director at Intelligent Business Strategies Ltd., a U.K.-based research and consulting firm. IT staffers involved in managing in-memory analytics applications need to assess the impact that the new tools will have on server and network performance, Ferguson noted. That includes developing an understanding of whether any IT infrastructure components need to be strengthened to support in-memory operations and whether some elements are no longer required.

For example, Ferguson said that integrating in-memory online analytical processing (OLAP) databases with data warehouse appliances could make it possible to eliminate some physical data marts. "That simplifies the architecture in a data warehouse environment, because it removes the need to extract data from a data warehouse and move it into data marts and on into cube data stores," he said. As a result, there are fewer steps involved in preparing data for analysis and fewer data stores to manage, which should lower the cost of supporting analytics activities within an organization.

Over the next few years, Ferguson expects to see the development of massively parallel processing appliances that can support a combination of solid-state drives (SSDs) and in-memory data storage for BI and analytics uses. In such systems, he said, SSDs would

replace conventional disk drives as the primary storage medium, with the most heavily accessed data going into memory. But that “will not happen overnight,” he added. “We need the prices of SSDs to fall.”

Ultimately, no matter how an in-memory analytics system is structured from a technical standpoint, the success of a project depends on choosing the right use case for the in-memory tools, Evelson said. BI applications and business processes that require a lot of rigor and adherence to standardized analytical routines – ones in which you really don’t want business users to step out of well-defined boundaries – are not good candidates for the in-memory approach, he cautioned.

About the Author: *Alan R. Earls is a Boston-area freelance writer focused on business and technology.*

Resources from Tableau Software



[Whitepaper: Analytics At The Speed Of Thought](#)

[Whitepaper: In-Memory or Live Data – Which is better?](#)

Why choose? Tableau allows you to switch between in-memory and a live data connection anytime.

[Download a free trial.](#)

About Tableau Software

Tableau Software, a privately held company in Seattle WA, builds software that delivers fast analytics and visualization to everyday businesspeople. Our mission is simple: help people see and understand data. Tableau's award-winning products integrate data exploration and visualization to make analytics fast, easy and fun. They include Tableau Desktop, Tableau Server and Tableau Public.

We understand the needs of businesspeople, non-technical and technical alike, when it comes to retrieving and analyzing large volumes of data. As a result, Tableau has already attracted over 65,000 licensed users in companies from one-person businesses to the world's largest organizations.