

In-Memory Computing

Enabling Real-time Access to Big Data

March 2013

Nathaniel Rowe

March 2013

In-Memory Computing: Enabling Real-time Access to Big Data

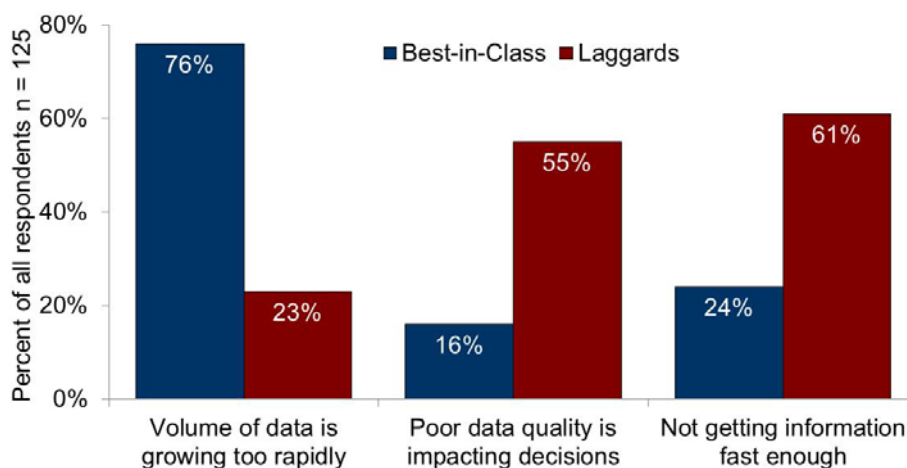
There are many tools available to help manage high volumes of data, quickly analyze information, and deliver prompt insight to decision makers. However, in Aberdeen's recent research study on *Data Management for BI* (January 2013) of 125 organizations, one technology stood out from the rest: in-memory computing. The latest enterprise-ready solutions using in-memory technology are shown to help Best-in-Class organizations analyze data faster, improve the quality of their business decisions, and use the insight to increase customer satisfaction and sales performance.

How the Best-in-Class use Enterprise Data Effectively

Modern enterprise data environments face a multitude of challenges. Data volumes are rising, more applications and data sources require integration, data quality must be addressed, unstructured data must be managed, analysis must be performed, and reports delivered. Within this maze of complex tasks, often times the "why" behind all this accumulated data gets lost. Ultimately, this information needs to serve a purpose; it has to drive business value. The end goal for any data environment should be delivering precisely the right information to the right people quickly enough for action to be taken.

In Aberdeen's January 2013 study on *Data Management for BI*, organizations indicated the three main obstacles preventing them from reaching this goal: 1) having too much data, 2) the information not arriving in time, and 3) poor data quality negatively impacting business decisions.

Figure 1: Why Aren't the Best-in-Class Worried About Speed?



Source: Aberdeen Group, January 2013

Research Brief

Aberdeen's Research Briefs provide a detailed exploration of a key finding from a primary research study, including key performance indicators, Best-in-Class insight, and vendor insight.

Maturity Class Definition

The January 2013 report on *Data Management for BI* defined the maturity classes with the following metrics:

- ✓ **Time to integrate data sources.** Best-in-Class organizations took 9 days, while Laggards took 137 days — 15 times longer.
- ✓ **Accessible business data.** The Best-in-Class improved the accessibility of their data by 35% last year. Laggards saw a 10% reduction in the same metric.
- ✓ **Data accuracy.** Best-in-Class companies reported 93% of their data to be reliable, while Laggards could only report 57%.
- ✓ **Delivery of information to decision-makers.** The Best-in-Class met the demand for actionable information 92% of the time, while Laggards met their delivery window only less than half the time (47%).

However, when examining how organizations at different maturity levels viewed the same obstacles, Aberdeen found the Best-in-Class (see definition in sidebar on Page 1) are far less concerned about data quality and speed. In fact, these top performers were up to 3.4-times less likely than Laggards to cite the issue as one of their top business pressures. This large gap means either:

- the Best-in-Class have substantially lower standards with regard to quality and information delivery, or
- they have successfully solved these problems.

As shown in Table 1, the situation is certainly the latter. The Best-in-Class not only set the bar in terms of data performance, but they lead the pack in regards to collecting information on a "Big Data" scale and gaining fast access to data analysis and insight.

Table 1: Best-in-Class Report Fast, Reliable Access to Big Data

Performance Metrics	Best-in-Class	Industry Average	Laggards
Annual data growth	100%	40%	34%
Demand for real-time analysis	46%	18%	14%
Ability to meet demand for data	92%	76%	48%

Source: Aberdeen Group, January 2013

The amount of data the Best-in-Class store is doubling every twelve months, a growth rate almost three-times larger than Laggards. As of the January 2013 study, they manage an average of 240 terabytes (TB) of data, which was 42% more overall volume than Laggards, who registered less than 135 TB. It comes as no surprise then, the issue of data volume is the most commonly cited business pressure by the Best-in-Class.

Furthermore, almost half of Best-in-Class (46%) required real-time access to their information. Yet despite their daunting task of processing such large amounts of data with blazing fast speed, the Best-in-Class report a staggering 92% success rate at meeting this insatiable demand for data. The Laggards, with their much smaller pool of data and slower demand, only meet their demand less than half the time.

The data ultimately delivered tends to be of high quality as well. The Best-in-Class, defined in part by their ability to keep their data clean, report an average of 93% of their information being accurate and reliable. This raises the inevitable question of: How? Given these constraints, what enables these top performers to quickly deliver quality data to the right people?

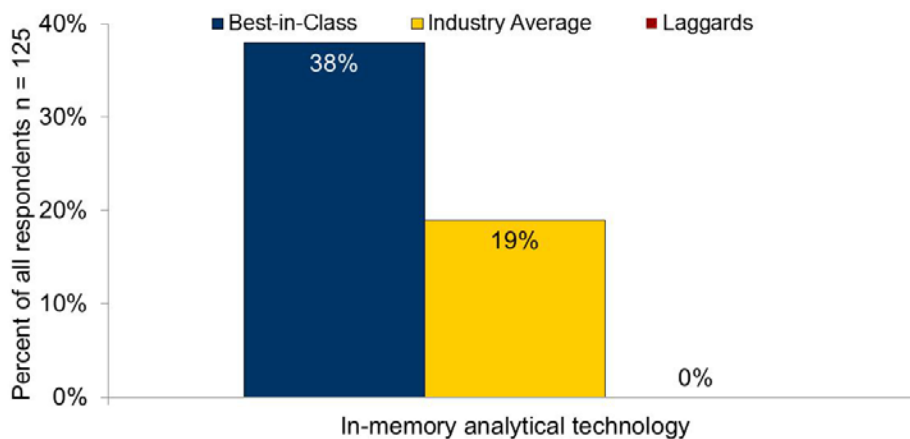
Definitions

√ **Big Data** refers to the problems of capturing, storing, managing, and analyzing massive amounts of various types of data. Most commonly this refers to terabytes or petabytes of data, stored in multiple formats, from different internal and external sources, with strict demands for speed and complexity of analysis.

In-Memory Computing To the Rescue

There are a number of technology enablers and business capabilities from the January 2013 study that the Best-in-Class are more likely to adopt than other organizations. Further information on these data management practices can be found in the [Data Management for BI](#) benchmark report. However, when it comes to the issue of delivering rapid access to large volumes of data, no technology solution correlated to Best-in-Class performance more than in-memory analytic tools (Figure 2).

Figure 2: In-Memory is a Best-in-Class Enabler



"Address the quality and consistency of your data at the source or as close to the source as you can. Pay attention to consistent business processes and to naming standards with properly identified global meaning across the enterprise."

~ Manager, \$1 Billion - \$2.5 Billion North American Chemicals Company

Source: Aberdeen Group, January 2013

In-memory computing involves storing large blocks of data directly in the random access memory (RAM) of a server, and keeping it there for continued analysis. Without in-memory, when data needs to be analyzed a data look-up is performed on the relevant database or data warehouse, the requested dataset is transferred to the server and then cached temporarily in the RAM. This activity has a number of potential bottlenecks — the speed of the disk drive the data is stored on, the input / output (I / O) connection between the storage device and network, the network connection to the server, and even the software application running on the server itself. A dedicated in-memory solution removes all of these obstacles. With the price of RAM dropping and improvements in server technology, a single machine can now hold hundreds of gigabytes or even terabytes of RAM by itself, allowing for rapid data access even at high scale. Although the concept of in-memory analytics has been around for years, it is only recently that enterprise-ready solutions have started becoming mainstream.

Among the Best-in-Class, roughly two in five organizations (38%) have currently implemented this technology, showing they are early adopters and already benefiting from using these tools. In the January 2013 research study, none of the Laggards indicated they had adopted in-memory, making this technology a very strong indicator of top performance. Additionally, this technology attracts a large amount of interest; 37% of companies are in the

process of evaluating in-memory or planning to invest in a solution within the next two years.

For a more concrete analysis of the practical benefits of using in-memory, a previous research study on *In-memory Computing* (January 2012) compared companies that had fully implemented these solutions with companies that had not. The results are striking — users of in-memory process over three times the volume of data, at speeds over 100-times faster than their competitors. While companies without these tools average an hour and fifteen minutes to respond to a query or fulfill an analytic request, in-memory users have results in a matter of seconds (Table 2).

Table 2: More Data, More Speed, More Efficiency

Performance Metrics	Use in-memory (n = 33)	Don't use in-memory (n = 163)	In-memory Benefit
Median amount of data analyzed	14 terabytes	4 terabytes	3.5 times more data
Average response time for data analysis or query	42 seconds	75 minutes	107 times faster

"When it comes to data management and Big Data programs, just do it. It may be an expensive change process, but so is lost opportunities and lost business..."

~ CEO, Small IT Services Company, EMEA

Source: Aberdeen Group, December 2011

However, there are a few concerns to consider when evaluating in-memory options. First, servers are designed to flush memory-based data after a few CPU cycles in order to bring new data in, and when they lose power or are brought down for maintenance, all information stored there is wiped clean. Traditional storage — like disks — are a much safer option for long-term storage because they offer an array of essential enterprise-features like high availability (guaranteed levels of operational performance) and RAID — a redundant array of independent disks, meaning a relationship between drives that back each other up to replace data from a failed drive. Next, applications and management software running on servers often place limits on the amount of RAM available, as parts of the memory are reserved for other actions, such as storing the application. This means while a server might have several hundred gigabytes of RAM, only a portion of this can be used to store data for analysis. Finally, business intelligence (BI) and analytic software is not always optimized for in-memory processing. While these tools are essential for the easy visualization of information and presentation to end-users around the organization, they can serve as an unintended hindrance. The hardware might be capable of more, but the software doesn't allow it to live up to its potential.

There are two main categories of in-memory solutions currently available in the market, each addressing these problems in a difference way. The first involves a high-end, pre-built appliance. These involve a number of advanced, networked servers specifically designed to maximize their RAM capacity, then loaded with software tailor-made for high performance

analysis. However, these in-memory appliances are almost exclusively meant for large enterprises, and carry a price tag commensurate with the size of their intended customer.

The other approach involves BI or data management software written to unlock the in-memory potential of any commodity server. They attempt to address the common issues of storing data in RAM, like high availability and maximizing server capacity. While they might not provide the pure high-end results of a dedicated, tuned appliance, they can offer greater flexibility, a lower price point, and a feature set that is more than adequate for small to mid-sized companies. Aberdeen's report on [Big Data for Small Budgets](#) (December 2012) covers other options available to budget-conscious organizations looking to tackle large, complicated data tasks.

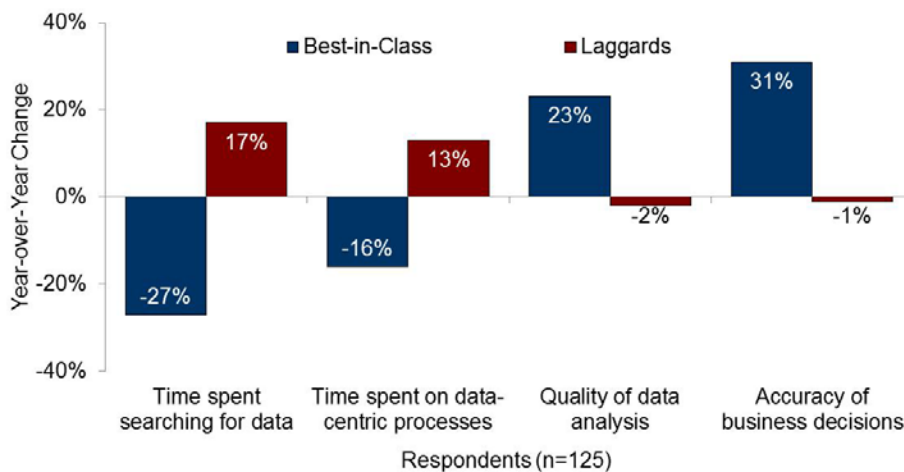
Business Benefits — Finding the Value in Your Data

While moving bits and bytes around faster than your competitor might be nice, it eventually comes down to the ultimate goal of data management: delivering precisely the right information to the right people quickly enough for action to be taken. The Best-in-Class companies in Aberdeen's [Data Management for BI](#) report (January 2013), thanks in part to their use of in-memory analytic technology, accomplish precisely that. As Figure 3 displays, when it comes to the essential operations of finding data, analyzing it, delivering the results, and finally acting on it, they showed incredible year-over-year performance improvements.

"Big Data is a huge opportunity to gather insight in what's going on in the enterprises ecosystem, especially customers. We need the technology to analyze and get the results we are asking for."

~ IT Manager, More than \$5 Billion, EMEA, Software Company

Figure 3: Faster Data Access Improves Visibility and Decisions




Source: Aberdeen Group, January 2013

The Best-in-Class report a 27% decrease in the average time spent looking for relevant data, and are much happier with the quality of data analysis performed. As a result, the accuracy of their business decisions improved by 31% in the last 12 months, and the time their employees spent on data-centric business processes became more streamlined. These processes

could include everything from customer interactions, to processing invoices, to collecting mortgage applications, to managing inventory. A 16% reduction in the time each employee spends on these tasks can quickly add up to a lot of man hours saved, hours that can be retasked to other critical jobs.

The business benefits don't end there. When powerful analytics link to a specific business operation or goal, all the performance benefits of getting fast, reliable, actionable intelligence becomes much more pronounced. One of the most common use cases for analytics is understanding customers — their behavior, preferences, demands, and preferred channels. The Best-in-Class not only outperform other organizations in pure data management metrics, but by applying these tools to their sales and marketing division, report incredible improvements in customer interaction.

Figure 4: Better Data Analysis Improves Customer Interactions



Source: Aberdeen Group, January 2013

In the last year, these top performers reported double digit increases in customer retention and satisfaction. By better understanding the needs of their customer base, and reacting quickly to these demands, these organizations distinguish themselves from their competitors and drive a 13% improvement in sales performance.

Key Takeaways

The methods of data management are rapidly evolving to meet the changing demands of the modern business world. For organizations struggling with controlling their growing volume of data and delivering intelligence to their decision makers in a timely fashion, Aberdeen recommends the following:

- **Build the business case for an in-memory solution.** The unique ability to analyze data directly in the RAM of a server allows for rapid processing of even large-scale data. These solutions correlate strongly with the Best-in-Class' ability to reliably deliver real-time data access. Furthermore, users of in-memory solutions

"As a data mining specialist, my key takeaway is that a Data Scientist must really understand the business prior to building any Big Data process. To be effective in building analytical solutions that meet the needs of the organization, you really need to have an intuitive grasp of how the operational processes function. The same is true for Big Data processes in the sales / marketing world. The first step must be getting a full assessment of the business itself — whether that's an operational process or understanding the market scope for a given product."

~ Data Mining Specialist, Mid-sized Transportation / Logistics Company, North America

analyze 3.4-times the amount of data over 100-times faster than companies without these tools.

- **Choose the in-memory solution that is right for you.** Make sure that potential stumbling blocks like high availability, server utilization, and properly designed applications are all addressed. Consider your budget constraints, and the amount of processing power you need to support your business units. Then target the solution that best fits those constraints.
- **Link your analytic platform to a business process or goal.** The ultimate goal of data management and analysis is to drive value by supplying the right information in a timely fashion. Identify the most pressing areas that your organization needs to improve upon, like sales, marketing, or supply-chain efficiency. Then collect the raw data on these operations and feed it into a powerful analytic engine (like in-memory) to give your decisions makers the insight they need to make the right changes.

For more information on this or other research topics, please visit www.aberdeen.com.

Related Research

[Big Data Trends in 2013](#); February 2013

[Data Management for BI: Getting Accurate Decisions from Big Data](#); January 2013

[The Big Data Imperative: Why Information Governance Must be Addressed](#); December 2012

[Big Data for Small Budgets](#); December 2012

[Go Big or Go Home? Maximizing the Value of Analytics and Big Data](#); September 2012

[The State of Big Data: Video Benchmark](#); July 2012

[Agile or Fragile? Your Analytics, Your Choice](#); July 2012

[Beyond Agile Analytics: Is Agile Data Integration Next?](#); June 2012

[Managing the TCO of BI: The Path to ROI is Paved with Adoption](#); May 2012

[Enabling Access to Big Data with Data Integration](#); April 2012

[High Performance Organizations Empower Employees with Real-Time Mobile Analytics](#); April 2012

[Mobile BI 2012: Accelerating Business on the Move](#); March 2012

[The Little Elephant in the Big Data World: Hadoop 1.0 Goes Live](#); March 2012

[Operational Intelligence - Part 1: Driving Performance with Tactical Visibility](#); February 2012

[In-memory Computing: Lifting the Burden of Big Data](#); January 2012

Author: Nathaniel Rowe, Research Analyst, Enterprise Data Management
(nathaniel.rowe@aberdeen.com)

For more than two decades, Aberdeen's research has been helping corporations worldwide become Best-in-Class. Having benchmarked the performance of more than 644,000 companies, Aberdeen is uniquely positioned to provide organizations with the facts that matter — the facts that enable companies to get ahead and drive results. That's why our research is relied on by more than 2.5 million readers in over 40 countries, 90% of the Fortune 1,000, and 93% of the Technology 500.

As a Harte-Hanks Company, Aberdeen's research provides insight and analysis to the Harte-Hanks community of local, regional, national and international marketing executives. Combined, we help our customers leverage the power of insight to deliver innovative multichannel marketing programs that drive business-changing results. For additional information, visit Aberdeen <http://www.aberdeen.com> or call (617) 854-5200, or to learn more about Harte-Hanks, call (800) 456-9748 or go to <http://www.harte-hanks.com>.

This document is the result of primary research performed by Aberdeen Group. Aberdeen Group's methodologies provide for objective fact-based research and represent the best analysis available at the time of publication. Unless otherwise noted, the entire contents of this publication are copyrighted by Aberdeen Group, Inc. and may not be reproduced, distributed, archived, or transmitted in any form or by any means without prior written consent by Aberdeen Group, Inc. (2013a)