

1 Hello, everyone, and I'd like to welcome you to the IBM Midsize Business Analytics Summit for Spring 2010. My name is Ben Plummer. I am the Director of the IBM Cognos Midmarket Business Analytics Unit, and I want to thank you all for joining us.

2 So what is business analytics? What do we talk about when we really get down to the core of business analytics as we define it? It's more than just looking at information. It's more than just taking data and viewing it and reporting on it. It's really being able to turn that into insights into your business that are actionable and that can actually change and provide you value in terms of how you run your business. So it's far more than it's traditionally been in terms of simply looking at data. It's about collaborating with that information. It's about changing the way your business runs. It's about taking real measurable action, and as we go through this presentation today, it's important for you to understand that it's really changed a lot as we've brought the various components here at IBM and as the industry's matured, so I wanted to be sure we all understand what business analytics means when we talk about it, and I'm sure as we provide you the insights of the analysts and the customers, you'll get a better understanding of how this really has matured.

3 One of the things that's really changed our ability to interact and take this action is the advent and the creation of what we call in-memory analytics. And this has really changed the way companies are able to organize their information and their data and then begin to work with that data real-time to change the processes in their organizations, change the way they plan for the future of their organizations, and the way they actually let people collaborate with this data and work with it interactively to see the changes take place. And if you really drill down into this, and we're going to hear a lot more about this capability throughout the day, you'll find it provides you a lot of new insights

4 And it really allows you to do the things that support this kind of execution and it's about organizing your data, providing you the flexibility to actually twist and turn this data to gain the insights that you need, and then do rewrite-type activities to actually take that action. So, really, it supports that full cycle we discussed earlier and gives you that ability to really change the way your business is executed.

5 So with that, I could sit here and expound on it, but it's really a whole lot more exciting to hear it from the industry experts and the people who have to do this on a day-to-day basis and interact with companies like yours and advise companies like yours on how they can change their businesses with business analytics and in-memory technologies like the ones we've just discussed. And so with that, what I'd like to do is I'd like to turn this presentation over to Helena Schwenk, who's a Senior Analyst at Ovum, who really focuses in on information management, business analytics, in-memory technologies, and let her walk you through what she thinks the value is of in-memory analytics and providing you a much quicker insight into how your business works. So with that, Helena, I'd like to turn it over to you and allow you to walk these folks through your insights on this particular technology. Helena?

Thanks very much, Ben. I appreciate the introduction. Hello, everyone. It's a real pleasure to be here speaking today at this virtual mid-sized conference. This is really my opportunity to take you through some of the important issues concerning in-memory analytics and if we move on to the next slide,

6 we can look at the agenda and the format of this presentation. There are three parts to today's agenda. Firstly, we will look at why you should care about in-memory; why, perhaps, is it important to your company? Before moving on to talk about use cases for in-memory analytics. And then we'll wrap up the presentation by talking about the important factors, the kind of key takeaways for organizations considering selecting or deploying an in-memory solution.

7 So one of the questions I often get asked as an analyst is, why should I be interested in in-memory technology? What does it mean to me, and more importantly, my company? Well one of the simplest ways to judge whether in-memory analytic technology is suitable for your, your IT department, your business uses, is to look at it from a very high-level perspective, to try and discern the potential value and benefits it may bring. On the slide, you will see that I have listed some high level indicators, some key questions that provide a really good starting point. Of course, I'm not suggesting that this a replacement for a business case or requirements analysis, but really an easy and quick pointer to help you decide if in-memory analytics holds potential value or benefit for your company. One of the top-of-mind questions when considering in-memory is performance or data performance. So do your business users complain about slow reports and query response times? Does it take an age to get query results back before you can actually work with the data? Do queries sometimes time-out? Our research suggests that query performance, or rather, slow query performance, is one of the

most commonly cited complaints with existing BI deployments. So in this case, in-memory could provide a potential answer. Moving on, are business users frustrated by the level of interactivity of existing reports? Do they, for instance, want to drill down to a further level of granularity or detail? Do they want a different report view, or to summarize the data in another way, but are frustrated because this often means going back to IT and requesting this information? Well, equally, in-memory analytics has a part to play in facilitating self-service BI. On a similar note, will it be beneficial for users to be able to calculate or aggregate summarized data on demand, again giving endusers that freedom, that flexibility to calculate measures or values as they are needed or as they are requested rather than looking to IT to perform this task? So in this sense, in-memory not only allows faster analysis, and our users can perform queries in seconds that would have otherwise taken minutes, hours, perhaps even days with conventional tools, but it also facilitates self-service BI, without users having to wait in that IT queue. Again another question to ask is with the ability to conduct what-if analyses, to be able to change business scenarios, and look at the impact of these changes on your business plan, would that be a useful addition to your analysts, especially financial analysts? But more importantly, to be able to do this, to have this sort of functionality with little or no additional work from the IT department? Likewise, if we look at in-memory from the perspective of the IT department, I think it also holds potential benefits. Time-to-value, that is, the time it takes for the IT system to be judged as delivering value to the organization is an important factor or consideration for IT. In-memory analytics by virtue of its architectural underpinnings, offers potentially a most cost-effective and less time-consuming proposition for IT and can therefore increase this time-to-value because it can potentially eliminate the need to build OLAP cubes, or perhaps lessen the requirement for IT to get really heavily involved in support and maintenance tasks, such as performance tuning, the OLAP or planning environment.

I think the bottom line here is that both endusers and IT benefit. The former get this self-service freedom, that kind of improved performance that translates to reduced user wait times, and the latter spend less time on BI development and support tasks like query and performance tuning. So if any of these questions rings true for you or perhaps has peaked your interest, then I would recommend that you investigate in-memory analytics further as part of this virtual conference.

8 I'd like to take a step back now, and look at what makes in-memory different from, say, traditional architectural approaches to BI and planning. So I'll start off with a definition. The term in-memory describes the notion of loading and processing data in RAM, that is, random access memory, rather than on disk, and this provides a big performance boost, as one of the biggest bottlenecks in typical BI deployment is slow disk or even slow database access. Another bottleneck is

disk input/output. And one of the reasons that in-memory is faster is because it has a lower I/O burden compared to on-disk methods. As way of an example, if we briefly compare two query approaches, where the traditional approach, users are required to build the query, then they might send it to the database for processing. They might have to wait for results before they can actually work with the data, and this could be an iterative and a repetitive process. With an in-memory approach, all the data, once it's first loaded into memory, users can access and interact with it straight away, resulting in performance speeds that are on order of magnitude faster than accessing that data from disk. But if you're sitting there thinking, hasn't this in-memory been around for a while? Then you're not mistaken. In-memory technology has in fact been around for more than a couple of decades and the technology, in many ways, still provides the same benefits in terms of faster access compared to on-disk as it did back then. However, the difference today is that in-memory technology is becoming more scalable, more faster, and cost-effective than its earlier counterparts. And one of the factors influencing this change is the declining cost of RAM. But equally, the popularity of in-memory can also be linked to incremental hardware advances, such as improvements in 64-bit computing, faster processor speeds, multi-core, all of which, when combined together, help lower barriers to adoption and improve the price performance ratio of the technology. The bottom line here is it's the combination of the declining cost of RAM with these technological innovations, is helping in-memory move out of its niche status, into an increasingly mainstream and viable technology platform for mid-sized organizations.

9 I'd like now to talk about some of the use cases for in-memory analytics, and how it can accelerate this time to insight. You can see from the slide that we've touched on some of these points before, such as improved self-service and what-if analysis, but I really want to hone in on one or two areas to help provide a clearer understanding of the potential value or use cases of in-memory. And to do so, I think it's important to talk about different deployment scenarios. First off, if we take a company who is perhaps new to BI, they're looking to dip their toe into a BI or planning project for the first time, then I believe an in-memory solution could provide a great starting point. Apart from the many benefits it could bring to business users, in terms of self-service, such as faster querying time, being able to slice and dice large datasets and on-the-fly analysis, it potentially delivers a quicker deployment time compared to taking the more traditional data-mart or OLAP cube implementation route. And this is because the task of defining, of building, and also pre-calculating an OLAP cube which let's admit it, is quite a time-consuming task, is greatly simplified with an in-memory approach. With in-memory, there's no need to define upfront the type of queries that need answering through, for instance, predefined indexes, through predefined calculations, query drill paths, or through aggregate tables and summarized data. In fact, for those new to BI, an in-memory solution could, in certain circumstances, but not all, replace the need for a data mart or OLAP cube. Likewise, without the requirement

to build or maintain a cube, IT needs to spend less time and effort on set-up that the support burden is also reduced, and so is potentially the type of skill sets required to support the environment. IT, therefore, perhaps has more time to focus on the other urgent development and maintenance tasks fighting or vying for their attention. And there is also less need for IT to kind of get involved with some of the performance tuning tasks such as designing and building aggregate tables or intensive database tuning tasks that are sometimes required through a more traditional architecture. Furthermore, another area where IT requirements change is the data footprint. Many vendors use sophisticated compression techniques, resulting in a requirement that is greatly reduced compared with storing the same data on disk. Another scenario I think we should look at, is that in-memory could provide a possible complement to an existing enterprise data warehouse-based BI system. If users often complain about query and report performance on their existing BI system, or perhaps that it takes too long for IT to respond to their information requests, then again, an in-memory solution could help here. It could act as a highly performant targeted, downstream BI engine, designed to provide a subset of data for a particular use case or a particular group of users. And one of these potential use cases could be if a company wants a quick solution to a specific analytic problem, maybe for financial reporting or expense planning for example. And in the case of a planning solution, the combination of an in-memory platform that provides that really fast underlying architecture with the ability to perform what-if analysis, provides a pathway for organizations to simulate different business scenarios, and understand the impact this has on their financial plan. It can, for example, allow users to test alternative business scenarios, such as an increase or a decrease in wholesale prices, and to be able to generate a new set of financial reports or statements based on these impacts that show the impact on revenue, cash, and expense and so on and so forth, thereby allowing companies to become more responsive to changing markets conditions. Furthermore, if this functionality is combined with write-back, and that is the ability to update the underlying data in real-time, these changes can be seen immediately, without any delay, which again, can increase the relevance and accuracy of a financial plan.

10 Moving on to the last part of my presentation, which summarizes the most important factors I would encourage you to think about when considering in-memory. Firstly, the move towards in-memory should always be considered in the context of the business requirements. We've touched on the merits of faster performance and on-the-fly analysis using in-memory, but sometimes these things are often hard to quantify by themselves. The bottom line is that any improvements in speed, in agility, flexibility, must be set in the context of how it delivers value to the business, and this must not be lost. I would therefore encourage both IT and the business units to work together to understand where the performance bottlenecks are, and how in-memory can help by, for instance, allowing the business to answer difficult, important, or challenging business

questions that perhaps weren't possible before. And as part of this process, it's also important to understand which particular user groups or areas of the business in-memory would most benefit. It may be, for instance, that in-memory is not for everyone and it's not for every user. Secondly, for those companies who are perhaps newcomers to BI or planning, then it is important for an in-memory solution to be able to grow with your requirements, whether this is in terms of analytic functionality, the ability to cater for reporting requirements, analysis, and planning, but also in terms of scalability, in terms of catering for growing volumes of data and changing user requirements, and how these will impact the suitability of in-memory. Likewise, for those who already have a BI deployment in place, then a key consideration is that you don't create yet another data silo that comes with all the associated efforts and costs of supporting a separate technology platform and set of tools, but also provides yet another version of the truth and increases the risk of data discrepancies. As always, there are many variables and factors that will decide on a company's choice of BI and planning architecture, but we would encourage using a metadata, an abstraction layer, to help improve consistency across various business terms and definitions across various different BI components. I would also encourage you to quiz vendors on the level of support they provide here. And the last point in this section is how write-back functionality, when combined with in-memory, could be an extremely useful feature for planning, budgeting, and forecasting solution, as it provides that level of interactivity and immediacy that allows users to see the impact of changes on the plan, without having to wait around or be dependent on IT. Moving on, one topic that is a perennial bug bear for BI is data integration and quality, and it's often an issue that tends to be swept under the rug, which is why I've put it on my slide here. All BI and planning solutions need a solid, a robust, and trustworthy data foundation to stand any chance of successful adoption, and an in-memory solution is no different here. And that is why a data integration and data quality strategy that understands how data is moved, how it's integrated, accessed, and used, will be a great asset to your in-memory project. And finally, I would encourage you to audit the vendors on your short list, and also to test some of their claims. For instance, in terms of understanding what level of 64-bit support they provide. Do they, for instance, support your chosen operating system? Likewise, it's important to ask vendors for reference customers, companies who reflect the types of application, user groups, and data scalability requirements that closely resemble your needs. And finally, in terms of future-proofing the investment, we would encourage you to consider how the in-memory product, or vendor of choice, can evolve, how it might be able to roll up into a broader enterprise-class BI platform. How much integration work, for instance, is needed in order to leverage other parts of that vendor's product stack or perhaps how much effort is needed to integrate with third-party tools. We believe planning these sorts of requirements at an earlier stage in the project lifecycle can reduce some of the pain, some of the heartache, and rework needed later on as business requirements grow, how they change, and how they evolve. That pretty much concludes my part of the presentation. I would like to thank you all very much for listening, and I'm now going to hand back to Ben.

Thank you, Helena, that was an absolutely fantastic outline for what I believe many mid-sized companies need to look to gain the value of business analytics and really leverage the power of in-memory technologies. And many of the concepts and the ideas that you've brought forward during your presentation have really been the foundation and cornerstone of the development of the product that I'd like to introduce you to now.

11 The IBM Cognos Express Business Analytics product, and many of the things you heard Helena discuss in terms of benefits to IT and benefits to the user communities, we believe we've incorporated into the IBM Cognos Express product, which really gives you the foundational capabilities of delivering access to information through enhanced query and reporting, much like you heard Helena discuss, giving you the insights you need in terms of being able to move through that data, either drilling down or reorienting yourself to provide you the appropriate insight and leverage that in-memory technology, and then, and something that we believe is unique to IBM Cognos Express in terms of our implementation, is allowing you to actually take action with that data and execute on the planning and forecasting by leveraging write-back directly into that in-memory environment. And you heard Helena talk about all of these capabilities being important to mid-sized companies and important to deployments of business analytics solutions, and when we set out to deliver the IBM Cognos Express product, those things were all core and part of what we built into this technology. So I'd like to take a little bit deeper look at some of the things that we also tried to implement in IBM Cognos Express and perhaps tie those back to some of the things you heard Helena talk about as well.

12 You heard Helena talk about the benefits of in-memory technologies to IT and to endusers, and one of the things that we absolutely wanted to be sure we delivered with the IBM Cognos Express product was lightening the burden, in terms of what IT has to be involved with, and providing endusers with a completely easy-to-use environment where everything they needed to do was kind of contained in one package, so from an IT perspective, one of the things that's very difficult and has held back many business analytics deployments was the installation. We've really worked very hard to simplify the installation, allow the initial creation of cubes or data stores for analysis to be done by the users so IT doesn't have to get as involved as early or as often, and this really simplifies the use and the installation of the product. One other area that Helena didn't discuss, but I know is important to many small companies, is how they purchase products. And we've really tried to focus on making the IBM Cognos Express product very easy to

buy, by allowing our partners, who have the domain expertise and who work with companies like yours or like mid-sized companies, to provide the product directly or to allow the product to be purchased and tried online, so people can get a feel for it before they actually bring it in. This simplifies the way the product's brought in, it makes it easier for you to purchase it when you do make your decision, and this easy-to-install, easy-to-use, easy-to-buy concept really allows for a streamlined interaction with the IBM Cognos Express product and IBM itself.

13 Foundational to IBM's Cognos Express product is a multi-dimensional environment. We have, at the core of this, integrated one of the probably best-known technologies around in-memory, both query and reporting, as well as write-back, and that was the TM1 engine. That's been brought in and is core to the product and really allows us to do all the things that you heard Helena talk about and I pointed out earlier, which is, the ability to organize your data, consolidate it, quickly begin to gain value from it, share it across the organization consistently, so there is one single version of the truth, and then allow you to do the analysis you need and the write-back that you're going to need to change the way your business operates. So, this multi-dimensional data store really gives us that foundational component that you heard Helena expound on in her presentation, and really gives IBM Cognos Express a unique way to get up and running in your organization, organize your information, and then allow you to actually change the way your business runs.

14 So this in-memory data store is really the foundational piece that allows the IBM Cognos Express product to deliver very rapid queries, dashboards, reports that allow you to gain that insight into your business that you're going to need and then this core technology allows you to twist, turn, reorganize, reorient the data to allow you to do the types of analysis that let you get to root cause problems in your business or areas where your business is excelling and you want to understand why, and then this in-memory analytics capability allows you to actually take action because of its support for write-back, which is one of the areas that you heard Helena talk about, and one of the areas that we found our customers believe differentiates IBM Cognos Express from other products out there. It allows you to actually execute on that same set of data to establish plans and forecasts and budgets and really change the way your business operates. And this cycle, supported by this in-memory analytics capability, really makes IBM Cognos Express unique and really sets it apart in the industry. This, coupled with the ability to do any type of activity you want,

15 from simple querying and reporting to full planning, all in an integrated, centrally



managed, and administered environment, really makes IBM Cognos Express a unique platform for mid-sized businesses that allows them to get up and running with their data very quickly through our easy-to-install, allows users to serve themselves, regardless of whether it's reporting, planning, visualization, any type of data activity they'd want to do, really gives them the full ease of use to move freely across all of these capabilities. And then our pricing, packaging, ability to work with partners to get solutions into the market truly give you the ability to get an easy-to-buy, an easy-to-implement kind of solution in. So IBM Cognos Express leverages many of the things you heard Helena talk about today. Our in-memory analytics capability really allows us to have a unique foundation that IBM Cognos Express was built on and we feel like this is really the only product out there today that allows you to do everything from reporting analysis through planning, in a mid-sized company, to allow you to change the way your business operates, and leverage the value and power of in-memory analytics.

16 With that, I'd like to thank everybody for joining us. I'd like to thank Helena once again for providing what I thought was a truly fantastic definition of the value of in-memory analytics, and the concepts that really drive the value of this into mid-sized companies, So thank you, everyone, once again.

END