BI 2.0 – how real time Business Intelligence is irrevocably changing the way that we do business

In Search of Insight

CHARLES NICHOLLS

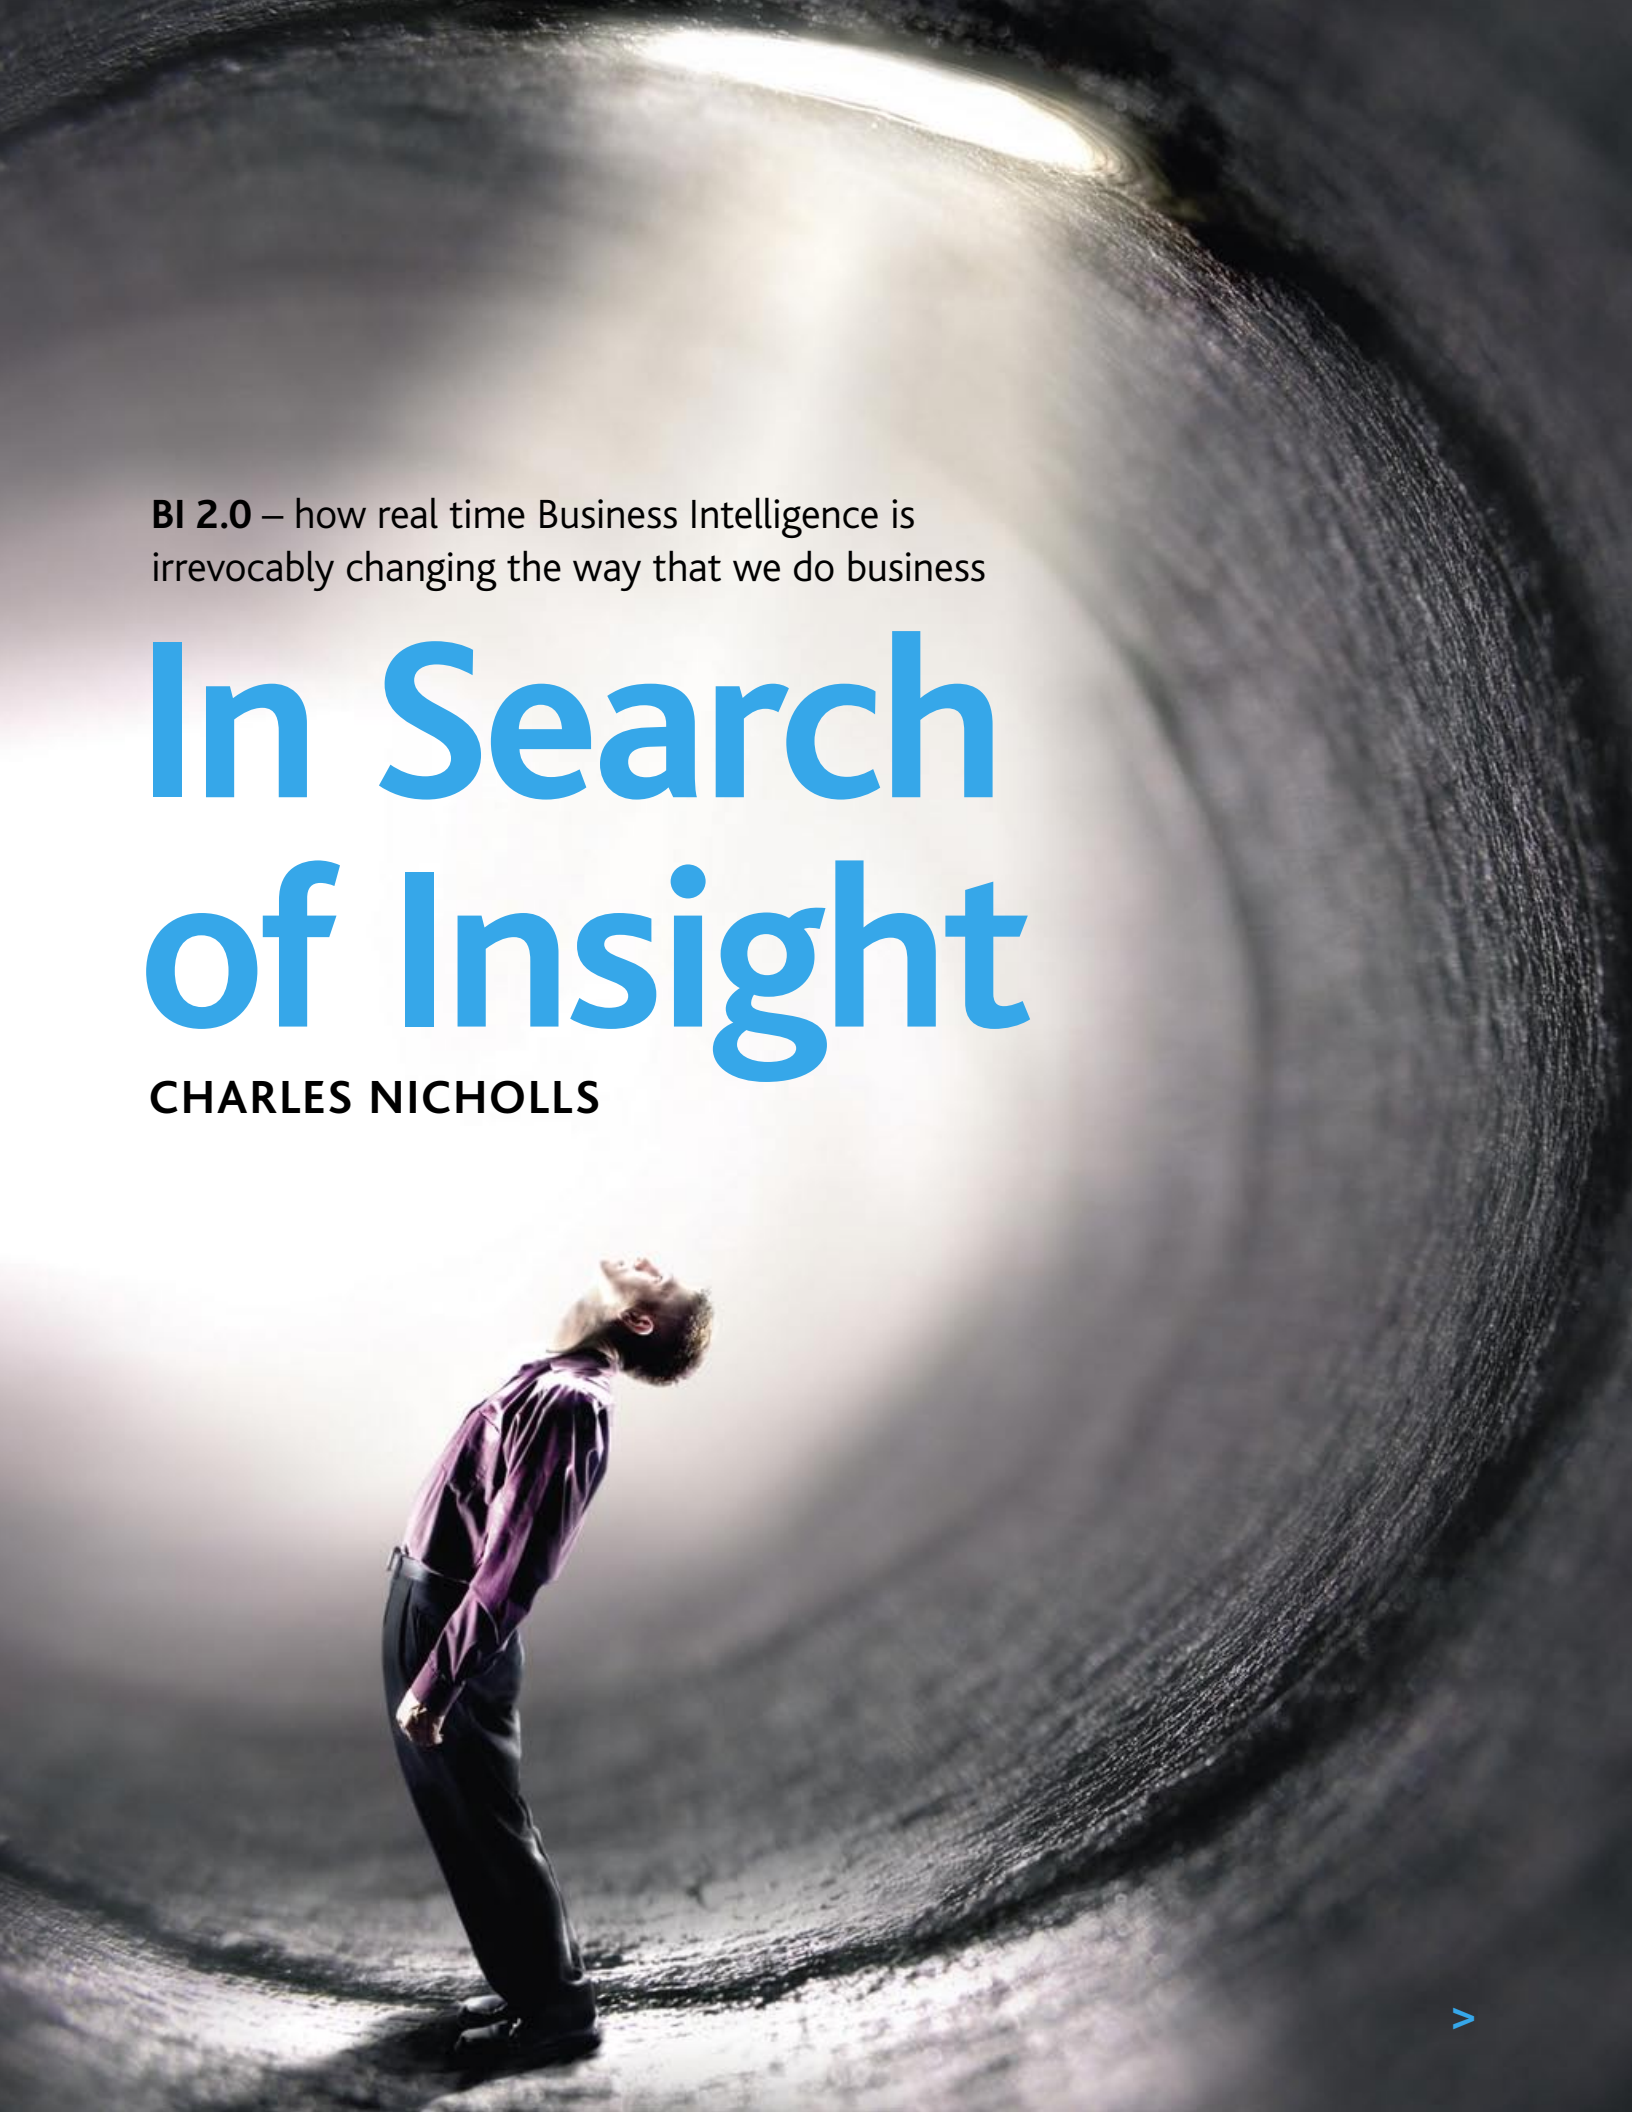

>

About this eBook

This eBook is written to share ideas about the challenges we face in using Business Intelligence to run our businesses in today's always on, real time world. It is free of charge, though the copyright remains with the author, Charles Nicholls, and with the publishers, SeeWhy Software.

You may post it in its entirety, or use sections, or reference the material on your website and in other documents without notifying the publishers, as long as you credit its content to the author and link to www.seewhy.com.

You may email it to friends and acquaintances that might be interested in the content and you may make copies without restriction.

Published By SeeWhy Software Limited

www.seewhy.com

Author's preface

This book originally started life as a series of notes I took to encapsulate an emerging idea. It was the end of 2002. The dot com bubble was truly burst, and I had just moved country again, this time from the US to the UK. My wife and I had just celebrated the birth of our second child, and as is often the case in these circumstances you challenge what you've been doing, and the assumptions that you live by.

It was in one such moment that I began to ponder on the Business Intelligence industry with all its unfulfilled promise, often long on vision and short on delivery. I recalled many meetings with senior executives around the world, each hoping BI would help them address new challenges as their businesses struggled to come to terms with a rapidly changing world.

The more you challenge the status quo, the faster you see the opportunities to make the world a better place. It was this process that started me on a journey that led inevitably to this book and to the formation of SeeWhy.

At a macro level, this book is not meant to be critical of people or the way that we do things. At its heart, it is optimistic about the future and the potential that exists by setting out a vision and road map for the future that we can aspire to.

In order to look forward to a brave new world of Business Intelligence (BI), we need to look critically at what works and what doesn't in the current generation of technologies. BI 2.0, the second generation of BI, will not replace these tools or way of doing things, but pushes BI concepts to new cross new frontiers. Chapter 3 looks at how we do things today, warts and all. By all means skip this chapter is you are familiar with the strengths and weaknesses of traditional BI.

It really makes the point that traditional BI does a good job at producing management information, but as soon as you try and extend the metaphor into operational or process centric BI you run into trouble.

I fundamentally believe that the exploitation of data will be a key driver for competitive advantage in the 21st Century. The means to capture data are well established, and are now a requirement of doing business. But data capture and processing, whether by enterprise applications or data warehouses are no longer means to achieve a competitive advantage. Reporting on the data and putting the information into the hands of the business users doesn't either.

Today, everybody does this.

The advantage undoubtedly lies with your ability to exploit what the data is telling you, and to do it before your competitors do. It's about acting on it in a time frame that can make a difference; it's about making day to day operations smarter. And in this area, current business intelligence systems are found wanting. It is also where the greatest opportunities lie.

Welcome to the world of intelligent business.

Introduction

We live in real time, minute by minute. News is no longer delayed by days but is streamed in real time. We bank online, and check our real time balances. We book flights with real time visibility of seat availability, and we select the seat we want, on line, in real time. All these transactions generate data.

Supporting our real time world is the microprocessor. They are everywhere, in almost every electrically powered device you can buy; from domestic appliances, cell phones, and cars, to the infrastructure we rely on for modern life. The next wave of miniaturization is already creating an 'internet of things' where devices and appliances are connected to the world over wireless networks, by RFID, each one constantly reporting their status. This too generates data - lots of data.

And to allow us to adapt our business models to today's real time world, software applications are now built using event driven technologies. Data moves around in real time over Service Oriented Architecture (SOA), using loosely coupled and highly interoperable services that promote standardized application integration.

And yet Business Intelligence (BI) today has not changed in concept since the invention of the relational database and the SQL query. Until now.

BI 2.0 is a term that encapsulates several important new concepts about the way that we use and exploit information in businesses, organizations and government. The term is also intrinsically linked with real time and event driven Business Intelligence, but is really about the application of these technologies to business processes.

At the heart of this architecture are events, specifically XML messages. Ultimately most modern processes themselves are actioned by events. Consequently, when you think about how to add intelligence into modern processes, the humble SQL query looks far from ideal.

The traditional data warehouse has enabled significant advances in our use of information, but its underlying architectural approach is now being questioned. It's architecture limits our ability to optimize every business process by embedding BI capabilities within. We need to look to event driven, continuous in-process analytics to replace batch driven reporting on processes after the fact.

In short, how can we build smarter business processes which give our organizations competitive advantage? How can we build the intelligent business?

This eBook sets out to answer this question, and to provide a roadmap setting out how we can get there. It's called BI 2.0.

Index

Chapter 1: Information Age comes of age

"Every two years the amount of information in the world quadruples, outstripping the original Moore's Law by a factor of two."

In 1965 Gordon Moore, the co-founder of Intel Corp., famously predicted that computer processor power would double every two years. This empirical observation became known as Moore's Law. He was uncannily accurate, and as a rule of thumb Moore's Law has been valid for the last 30 years, though it has been revised slightly, and recently it has been suggested that this growth in processing power may not be sustainable.

Whether it continues or not, we have all seen a true information technology revolution. The average 2006 laptop computer costing maybe $2000 is now as powerful as a 1980's mainframe computer which cost millions. Microprocessors are everywhere, in almost every electrically powered device you can buy; from domestic appliances, mobile phones, and cars, to the infrastructure we rely on for modern life. The next wave of miniaturization is already creating an 'internet of things' where devices and appliances are connected to the world over wireless networks, constantly reporting their status.

Sensors are appearing everywhere, tracking and monitoring everything from the products in our shops to the quality of the air we breathe. The next generation of sensors is even smaller, sometimes referred to as 'smart dust' – micro sensors that we won't even notice – which are being scattered across the planet to monitor such things as temperature and acidity of crops.

All this technology produces data. Lots and lots of data.

Drowning in data

At **SeeWhy** we've estimated that every two years the amount of information in the world quadruples, outstripping the original Moore's Law by a factor of two. The numbers are astonishing, sourced here from the **University Of Berkeley California**. Hosts on the internet double every six months, and the number of new web pages doubles every year. We produce 250GB of information each year for every person on the planet, but only .003% of this information is stored on paper. And according to the American Marketing Association, the majority of US based companies now have two to three times more customer data than they had a year ago.

Researchers conclude that one consequence of all this data is that *'the typical piece of information will never be looked at by a human being'*. And yet this raw data needs analyzing to put it into context, interpret its meaning and assess the impact on our businesses.

Of course the way we analyze data today is primarily manual, done by human beings. What you can surmise from this is that we are already drowning in data, and the problem is only going to get worse.

Despite all this data we lack insight.

Real Time World

We live in real time, minute by minute. News is no longer delayed by days but is streamed in real time. We are rarely out of touch, taking calls by mobile phone whenever and wherever, and email is real time and pocket sized, traveling with us wherever we go. We communicate by instant message and by texts. And from our mobile and wireless connected devices we can order products and services 24x7 or bid on auctions, in real time.

All of these real time transactions generate data – customer data, product data, sales data, even meta-data (data about data!)

Our expectations lie in the real time world. We are impatient. We expect everything to be 'always on', always up to date.

It would be unacceptable now for any of these services to be delivered in batch. You couldn't bid on auctions, or book flights. Your current account would show yesterday's balance. Instant messaging wouldn't be instant, but messages would be grouped together and delivered in a batch overnight, as would your email.

The point is clear: in an 'always on world' where we run our lives minute by minute, wirelessly connected to a wide variety of information sources, we cannot afford to run our businesses based on out of date information.

But we do.

The Age of Intelligent Business

Customers expect instant results, and don't want to wait for answers. We are already struggling to make sense of the data we have, and data volumes are growing significantly faster than processor capability. If we carry on doing the same things each year, we are going to have a problem: ever more data and ever less insight into our businesses.

Almost every business practice has adapted to shortening business cycles, except for the Business Intelligence world. Why do we run our business operations disconnected from the insights that could make us more money and reduce costs? Why should we use out-of-date information all the time? Clearly something has to change in the way that we (human beings) use technology to process and analyze data.

Businesses that can use this data to provide faster, better, cheaper, individually personalized services will be the inevitable winners of the next information age – the age of intelligent business.

Chapter 2: Business Intelligence today

"Users constantly complain of information overload. That is too much information that lacks real insight, and not enough time to make sense of it all. More reports faster, isn't the answer."

So how does traditional Business Intelligence fare in this age of intelligent business?

Almost every major business or organization the world over uses some form of Business Intelligence (BI) to run their business. BI enables us to plan and budget, control costs, figure out how to acquire new customers, and understand how to retain our existing ones; how to comply with regulators and ultimately to report results to shareholders. In short, BI is the eyes and ears checking the performance of our businesses and has grown into a $12bn a year business in its own right.

There is no doubt that BI has helped to publish information, held in static databases, to many middle and senior business managers effectively. It has enabled us to move from consolidating financials monthly, to daily.

We can now view data onscreen in published reports, nicely formatted, in place of the continuous computer paper printouts inches thick delivered on a trolley.

Insight not Information

But simply because you can now automate the distribution of data to users doesn't mean that in every case you should. Clearly you can generate an ROI that shows how much money you'd save compared with doing it manually, but users constantly complain of information overload.

Distributing reports out to the field doesn't change people's behavior. If the goal is to get them to manage their budgets more effectively, then training might be in order. Just sending a report, where the only indication that they're heading for an overspend is one number buried on page 26, is not effective at driving change.

The frustration by business users is most often stated as 'information arrives just too late to be really useful'. While at first blush this appears to be a timing issue, it's clear that information needs to be acted upon in order to be useful. Business people universally agree that they don't need more reports. What's lacking is real insight.

And this insight is needed at the front line in business operations where it can be used to improve performance. There is too much information that lacks real insight and not enough time to make sense of it all.

More reports faster, isn't the answer. But it's the answer that today's BI comes up with most frequently.

The time to insight

Without getting too technical, the real problem we're talking about is 'Latency'. Latency is the time taken from something happening or changing to the moment when we can do something about it. It is the root cause of the problem in the architecture used today for BI.

Latency, of course, is not a simple concept, but it is critical to BI. Analyst **Richard Hackathorn** produced a study of BI latency and concluded that there are three types: Data latency; Analysis latency; and Decision or Action latency.

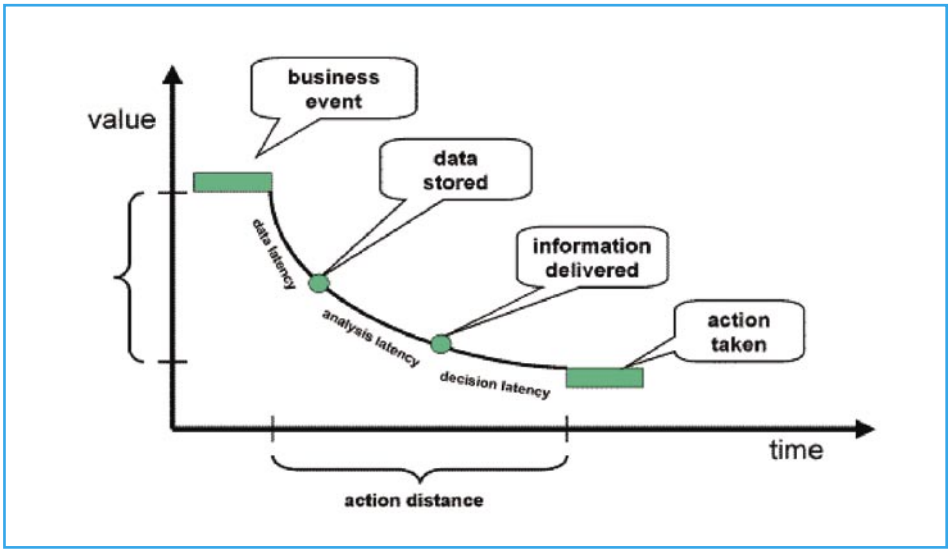

Reproduced with permission from Richard Hackathorn, Boulder Technology Inc.

Much of the recent work in BI has been to do with reducing the first category: data latency; feeding data faster into the data warehouse. The apocryphal example is the supermarket chain where the data is available for analysis in the data warehouse before the customer has left the car park.

Of course, if the customer happens to be driving away during the analyst's lunch break, while the data may be updated several times per day, the actual *analysis* is not going to be real time. In reality, no analyst is going to review the unique shopping habits of Mr. and Mrs. Smith and any impact on the business.

If the analyst were looking at this low level data, he might have noticed is that the Smiths have just defrauded the retailer by returning goods purchased on a cloned credit card, and the packet of cornflakes they bought was the last on the shelf so the shop is now losing sales until someone notices that the shelf is empty.

So the major problem with current BI is only partly the data latency; just as critical is the manual analysis process, and the time taken for real decisions to be made which can positively affect the business.

A retailer might survive these long delays. But in today's real time world, many businesses cannot afford to wait and are looking to BI to provide, not only the information, but the insight, the decision and in many cases the automated action.

So how does BI fair? To use an old scorecard favorite: must do better!

Chapter 3: Challenges for BI tools

"BI promised to 'democratize' access to information, to make it easy for ordinary business men and women to get the facts and figures they need to do their jobs more effectively."

The BI tools that we use today were designed to solve a problem from a previous generation of technology. Back in the early 1990's the relational database had become widely accepted, and as more data was stored in these databases, IT departments were spending time extracting it into reports for business users.

With the widespread adoption of Windows, BI tools took hold, promising to deskill the process of accessing data. Rather than teaching business users the complexity of SQL, here were tools that could equip them with the ability to create their own reports and analyses using drag and drop.

But all is not well with BI.

BI promised to 'democratize' access to information, to make it easy for ordinary business men and women to get the facts and figures they need to do their jobs more effectively. This was a big promise, but increasingly users of BI tools are frustrated with what they get, and how and when they get it.

Many organizations have adopted BI tools in quantity, buying thousands of user licenses to deploy up to 50,000 seats of some of the leading BI tools. There's a growing body of evidence that many of these licenses remain unused, and the majority fall short in meeting the expectations of their promise. Often the main benefit from deployment is merely a secure method of distributing standard reports.

Most BI users do not actually use the BI tool itself for anything other than viewing reports, cubes or dashboards produced by a central team. In many cases, users do not look at the standard reports produced, but use them as reference documents. The BI companies have attempted to provide more value by adding ever more features designed for users, but the majority of BI users do not use the features in the tools today interactively: they are essentially report consumers.

So where is BI headed? Based on current architectures, it's stalling.

What's needed is a new vision for BI. In thinking about this, it's first important to understand what are the challenges for today's BI technologies and what BI 2.0 could do to set it right.

"Reports are often not read, but are used as reference documents, referred to only to confirm a number."

Problem 1. Ease of use

BI tools are too hard to use. There. It's said and in the open. It's what every one knows but no one will admit to.

The difficulty of using these tools in the real world, as distinct from the slick sales demonstration, is caused not only by the tools themselves, but by the complexity of real world data. Data stored in data warehouses is almost always complex, leading to multiple ways to ask a given question. While BI tools generally try to simplify this complexity by providing end user metadata to use, multiple versions of a given calculation like 'revenue' make the system too complex for most business users. Typically companies end up with multiple definitions to suit different users.

Let's imagine a business person wants to ask a simple question such as "what is our divisional revenue for the current quarter?" Of course, the more important, forward looking question is "...and are we on track to hit target in all divisions?"

Sounds really simple, but in a real world production data warehouse a query involving 'revenue' needs to consider how revenue is calculated: there are probably multiple definitions that the user can use, such as 'gross sales', 'gross sales less discounts', or 'gross sales less discounts and returns' leading to confusion about what the numbers mean. This can lead to divisional numbers not adding up to group numbers because the comparison is not like for like.

It's clearly no surprise that users prefer to ask someone to calculate the numbers for them, or rely on standard 'reference' reports and dashboard which are distributed to them automatically. This explains why reports are often not read, but are used as reference documents, referred to only to confirm a number.

BI 2.0 builds on this by doing far more of the hard work up front, automating the analysis, highlighting the exceptions and taking or guiding the user directly to the area of the problem. BI 2.0 provides what the user really wanted: an alert about the problem, not a report to *find* the problem.

"This is akin to asking you to drive a car where your dashboard has a speedometer showing the average speed yesterday, and the fuel gauge gives the fuel remaining in the tank at the end of last week."

"What a difference a little time can make! By providing real-time information about what is selling, and in particular for this example, what is not selling, a $400m value has been unlocked."

Problem 2: Out of date & backward looking

Of course whether it's a dashboard or a report, the data is only as up-to-date as the last refresh, and only as current as the data in the warehouse at the time.

For example many retailers distribute reports to all of their stores overnight detailing items that were out of stock during the day. The idea is that the store manager picks up these reports in the morning and ensures the shelves are replenished before the store opens. But as we've seen above, users often don't look at standard reports, and store managers are no different. In fact there's lots of evidence that store managers are overloaded with reports, when they should be on the sales floor ensuring that the store is running smoothly.

Information overload isn't the only reason that those out of stock reports don't get looked at. In a classic case of shutting the stable door after the horse has bolted, these reports tell you what was broken yesterday. The store manager can see that just by walking the aisles and looking at the empty shelves.

This is a big problem in retail, accounting for a loss of 3 - 4% of total revenues on average across the sector. So for a $10bn retailer, that's a lost sales value each year of $400m. What's more frustrating about this is that **research** shows that often the stock is in-store, but just not on the shelf. In fact, in 70% of cases the problem is addressable by the store.

Now let's imagine a BI 2.0 world where information is never out of date, and as soon as an item goes out of stock in store staff can be notified; if the product is in the store, the shelf can be replenished, if not an emergency order can be created. The delivery mechanism of this information is unlikely to be a report, but more likely a constantly updated, real time electronic display of out of stock products in the stock room or pushed to a handheld terminal.

What a difference a little time can make! By providing real-time information about what is selling, and in particular for this example, what is not selling, a $400m value has been unlocked. Take a moment to consider what the potential is of an always on, always up-to-date approach in your organization.

"Query based architectures have served us well in an environment when the world was batch based. While there is a lot of batch file transfer still in existence, today's service oriented architectures fundamentally do not work this way."

Problem 3. Too slow

This difficulty of asking what seems like a simple question is compounded by the real-world performance of many data warehouse systems. Business users often complain about slow running queries. What seems a simple and obvious question from the user's perspective can often cause a significant quantity of data to be queried, resulting in unpredictable loads on the server, and slow performance for the user.

An example might be to ask the question 'what is our daily average sales so far this year, compared with last year?' This is an obvious question, but might result in a 360 million row query in a medium sized consumer facing business. This exposes one of the flaws of SQL: to compare like with like you need to extract all of the values in between the two data points. Of course if you want to ask the same question the next day, even though you've incremented one day, the database will recalculate 358m rows of data that were calculated the day before and then discarded, just to compare the latest value. Hardly optimal.

Of course a competent DBA can always build summary tables to speed queries, assuming that you're using an **aggregate aware BI tool**. But there is a trade off, as is so often the case in computing: summary calculations can cause a data explosion, which is explained well in an **article** by Richard Winter (check out the section on 'Combinatorial Explosion' for a succinct description of the problem). And these summaries can present a significant downside, particularly if you are aiming at a low latency database with a frequent refresh cycle because these summaries need to be rebuilt every time the data is refreshed. The complexity of maintaining large numbers of summaries in a rapidly changing database can introduce its own latency to the data.

Query based architectures have served us well in an environment when the world was batch based. We load data in batches. We run reports as a batch job. Analysis is batch based, where a query extracts a batch of data for analysis.

This batch orientation is a legacy due to a lack of real time processing technology and an era when there were no on-line users. The result was that transactions were grouped together and sent as a batch, perhaps by FTP between systems. We never wanted to send data in batch – real time was always optimal.

"Rather than reporting on the effectiveness of the process after the fact, BI should be used within the process as a way of routing the workflow automatically, based on what the customer is doing on the website."

Query performance continues to be an issue, especially where low latency is important. The closer you get to real time, the bigger the problem. While speed of response ultimately can be addressed by throwing hardware at the problem, it remains a factor in every BI user's day to day experience.

While there is a lot of batch file transfer still in existence, today's service oriented architectures fundamentally do not work this way. So BI 2.0 will have to address the issue of the query on a transaction by transaction basis, or even how to deliver BI without using queries. How we achieve this was one of the breakthrough moments on the road to creating **SeeWhy**.

Problem 4: Shifting sands

Whether you are using a BI tool against a static batch based data warehouse, or one updated incrementally several times a day, there are other issues that need to be considered specifically before considering using BI tools against real time data.

Let's imagine that you've upgraded your data warehouse to a real time basis. In the process you've spent millions doing so, and it's now apparent that, despite what you expected, it isn't real time, but rather, nearer real time that it was before.

Let's imagine that you're using your BI tool against this data warehouse, and you're looking to analyze sales by region. You run a query that pulls back some results at an aggregate level. Now you'd like to drill down to see a lower level of detail, but the underlying data has changed, resulting in an error: your regional sales no longer add up to your total.

The way some of the traditional BI products deal with this problem is that you 'prefetch' the data needed to build a complete cube of data. This effectively snapshots the database and gives you a cube in which you can slice and dice to your heart's content and all the numbers add up. Of course as soon as you want to drill down further than the edge of the cube, then you're likely to have different data which won't reconcile with the data you were just looking at.

By pre-fetching data you are running a bigger query on the database than is potentially necessary which takes us back to Problem 2: too slow. Additionally any decisions made based on the data in the cube are potentially incorrect because the data is out of date as soon as it has been extracted. Maybe this isn't a big deal for some applications, but it effectively limits automated actions being driven off the data. What if an error had been made in one of the operational systems? Without support for detecting and correcting transactions made in error, an action will stay uncorrected potentially causing significant damage to the business.

Here's a real example to illustrate the point. One Friday afternoon, a leading consumer packaged goods supplier to a large US supermarket received an unusually large call off order for pasta sauce. The order was sent by EDI directly into their systems, and an automated action was picked up which revised the forecast in the ERP system controlling production. The ERP system alerted managers that a significant extra quantity of pasta sauce was required, and the factory was shifted onto overtime working throughout the weekend to meet the perceived supply shortage.

Unfortunately for the company, the order was a 'fat finger' error, and was discovered only after a passing senior company manager stopped to ask why so many people were working late, and traced the spike in demand back to the supermarket who confirmed that they did not, after all, need a years supply of pasta sauce. In fact the order had been corrected only an hour after the original mistake, but unfortunately, with no intelligence in the loop, the system did not pick this up.

What was needed was a way of validating the order in the first place as a sanity check to ensure that the data was likely to be correct.

This particular order should have kicked out an exception to a manager to check that the order was correct. Even if this process had failed for whatever reason, if the BI system had been operating in real-time with the capability of triggering a roll back when it detected a erroneous transaction, then the automated action which caused the factory to spool up for weekend working would have been rescinded automatically as soon as the supermarket sent their cancellation over.

This is a practical example of building smarter business processes, where the process has some inherent intelligence. It illustrates the point that making decisions on dynamic data requires a fundamentally different approach, where each event can be automatically analyzed in real time.

Problem 5: Scalability

BI products have traditionally suffered from poor scalability, restricting their deployment and use to the more basic tasks with reduced data sets. Often the data in a data warehouse is limited to some level of summarization for practical purposes. For example, if you want to analyze where things go wrong in a customers purchase process, you need process-state data. Unfortunately process-state data is often not stored in the data warehouse, effectively restricting its applicability. While loading all the operational data in the first place would seem to be the ideal, it can increase the size of the database considerably.

This, together with a lack of real-time capability, effectively stops the data warehouse and the associated BI tools being used as part of a business process where an analytic process step can determine the workflow presented to a customer on a website, or to an operator in a call center, or a clerk in the shipping department.

But scalability isn't just about the amount of data. It is about the number of customers, or credit cards (millions), the number of product items (tens or hundreds of thousands), the number of staff and stores (thousands); in financial services it's the number of trades and traders, data feeds and counterparties; in telecommunications and utilities it's customers and call records, meter readings and bills, and so on. In fact it is the sheer volume of each of the dimensions of your business, and the interactions between them that creates a scalability issue for BI.

Put this into the context of today's real time world and it is clear that more data flowing faster needs more analysis more quickly. When the customer is on the website or speaking to your call center, you need to be able to react to what they are doing in real time. This is where the real competitive advantage now lies.

The current 'state of the art' is to base your analytics on past historical patterns – yesterday's out of date patterns. This is a fundamental flaw in the way we interact with customers. Where is the individual one to one, personalized and relevant conversation we heard so much about during the dot com bubble?

Current approaches rely on historical groupings of customers into static segments, often containing millions of customers in each grouping. There is nothing personal about it!

So scalability in BI 2.0 is about being able to handle the millions of dimensions of business and the relationships between them, at an individual transaction level in real time. Scalability on another scale altogether.

"The human brain is particularly bad at spotting the one bad event in a stream of thousands."

Problem 6: Manual

The other impact of not being driven by the day to day events happening in the business, is that the process inevitably relies upon people to notice that there's a problem. And that probably means that someone has to go and look.

There's lots of evidence that this is a fundamentally flawed process. People don't look unless they know there's a problem, and in some cases even then, they deliberately don't look!

'Manually looking' means that your business is limited to checking a tiny fraction of your business transactions. This is usually days or weeks after the event, and relies on an analyst having insight to understand what the aggregated data that he or she is looking at really means. An aggregate can hide a multitude of sins.

At this point you're probably bristling. "Our people are great," you're thinking. "They're diligent, and no computer system can ever be as good as one of our best people!"

You're right, but consider this. The human brain is particularly bad at spotting the one bad event in a stream of thousands.

Security guru **Bruce Schneier** describes how this is also a problem with airport security, where airport baggage screeners fail to spot guns, knives and weapons planted deliberately in baggage to test them. After only a few minutes of viewing bags passing in front of them, they lose the ability to recognize non-uniform patterns. We become blind to the exception, seeing only the uniform pattern.

Just so with data; there is simply too much to manually examine it all, and even if we attempted to do so, this 'exception blindness' would soon render the process completely ineffective.

So the result is that we have managers looking at reports, mainly as reference documents to confirm a problem they have discovered elsewhere. Their report probably looks only at aggregated data, missing the fine grain detail that is often so important.

The CFO of a computer chip manufacturer described this problem perfectly: at the end of each year they find that there is a shortfall in actual revenues compared with the contracted volumes with their customers. This shortfall can be measured in hundreds of millions of dollars. When a customer wants to buy chips, a contract is established where the price is a function of volume, and a call off system is set up where each call off order can be sent electronically via EDI. When

"The CFO of a computer chip manufacturer described this problem perfectly: at the end of each year they find that there is a shortfall in actual revenues compared with the contracted volumes with their customers. This shortfall can be measured in hundreds of millions of dollars."

the orders come in they are posted automatically into the order and production systems, and as a sequence of orders flow in from the customer, all looks well.

What actually happens is that some customer orders fall just slightly short of the rate they should be ordering. The order is just smaller than it should be making it tough to spot, (this is known as the Just Noticeable Difference and calculated at approximately 7%) but over many call off orders across multiple contracts the short-fall adds up to hundreds of millions in revenue short-fall. The devil is in the detail.

The only way to fix this type of problem is to screen every call-off order automatically, and immediately flag the exceptions up to a manager who can address the problem. This problem cannot easily be spotted at an aggregate level, but requires a process to check every EDI order automatically. This is really an event driven description: when the call off order comes in, it needs to be analyzed individually in the context of past orders, expected future orders and the contract values. What's needed is intelligence, embedded within, or running alongside, the call off order process to automatically check and validate every order.

"Rather than reporting on the effectiveness of the process after the fact, BI should be used within the process as a way of routing the workflow automatically, based on what the customer is doing on the website."

Problem 7. Disconnected from processes

This last example illustrates perfectly that BI should play a role in many of the day to day processes that are part of the modern business. Whether it's checking data entered by customers on a website, or orders coming in over EDI, validating a booking on a travel site, or cross selling a customer in a call center, there is a compelling case for making these processes more intelligent. Rather than reporting on the effectiveness of the process after the fact, BI should be used *within* the process as a way of routing the workflow automatically, based on what the customer is doing on the website. This is a closed loop scenario where an action can be completely automated enabling the process to be driven by intelligence to an optimized outcome.

Of course, BI today is typically disconnected from processes: based on a data warehouse which is out of date, reports are the wrong delivery mechanism for an in-process intelligence capability. A BI tool can be presented as a Web Service, but this does not solve the problem that the underlying data is out of date and that the database still needs to be queried.

The data warehouse has no place in day to day processes, or in an event driven or service oriented architecture (SOA). Why mix architectures when logically the events themselves should be analyzed either in parallel to the flow, or as a process step? For these processes, analysis of events should be 'event driven' i.e. the analysis is of the event in context of history, not a query on a database.

Chapter 4: Real time BI and the dash for dashboards

"The issues here relate to speed: how often will the data service be used? What is the load? Is the data structured in the right way to get rapid access, given the nature of the processes involved?"

From all that we've said, you might conclude that the requirements for BI have changed. Indeed, the traditional data warehouse architecture doesn't suit the way that we now seek to use intelligence as part of our real time, day to day operations. This doesn't mean that the data warehouse is going to go away, merely that it's not necessarily suited to all of the BI tasks that are required by contemporary business.

One area in which this is particularly true is SOA. Today's applications are built using Service Oriented Architecture (SOA), using loosely coupled and highly interoperable services that promote standardized application integration and reuse.

A key component of SOA is the underlying asynchronous message infrastructure; an Enterprise Service Bus (ESB) or Enterprise Application Integration (EAI) system. ESB's are typically based on Web Services standards and provide foundational services for SOA via an event driven messaging engine.

At the heart of this architecture are events, specifically XML messages. Ultimately the processes themselves, perhaps managed in a Business Process Management tool, are actioned by events.

The traditional view is to present your data warehouse as a service which can then be used by the other applications and processes connected to the ESB. The challenge here, as highlighted by Gartner in Chapter 5, is that data warehouses are data centric, storing data ready for analysis and reporting, not aligned to the process at hand. The issues here relate to speed: how often will the data service be used? What is the load? Is the data structured in the right way to get rapid access, given the nature of the process or processes involved? Will there be any ad hoc queries running on the data warehouse at the same time?

Moreover, since the data warehouse often does not hold process- state data, let alone real time process state information, it usually doesn't have the information you need to make a decision.

The intelligence element here is that we need to compare these events with historical data in real time. Is the data entered likely to be correct given this customer's past behavior? Is this a reasonable EDI order or is it likely to be an error? Is this booking highly likely to be fraudulent? Given previous transactions this customer has made, is it a normal deposit or an unusually large one that indicates a cross sell opportunity?

These questions are important to answer, and typically have three characteristics in common:

- The analysis needs to take place in real time either as a process step or in parallel with the process;
- and therefore the analysis needs to be automated;
- and the analysis itself involves comparison of the event with historical data at a low level of granularity.

These characteristics typically mean that traditional approaches, such as query based BI tools, Business process management tools or rules engines are incapable of addressing the problem effectively on their own.

Effectively we're talking about a significant extension to the way that we've always thought about BI.

One technology that has been touted as a significant extension to BI is the dashboard. Far removed from the 'Executive Information Systems' of old, still they present an interesting challenge to the understanding of real time BI.

The dash for dashboards

Is this good? While the dial may be in the green, it can easily mask variable performance at a lower level which needs acting upon urgently.

Maybe it's a modern marvel to have your information displayed automatically on a screen, and it's certainly easier than writing SQL queries or getting someone to do it for you. There's no doubt that business users' appetite for dashboards is significant, driven not least by the alternatives which are much less accessible.

But, as we've said, most of the data warehouses that feed the dashboards are out of date to some degree. For the business users in operations, out of date data may be *interesting*, but not very *useful* for running the business today, however well it is presented.

The dashboard is also presenting an aggregated view of the business, an executive summary if you like. If it's too granular, then it loses its value – after all who wants to sit in front of a real time dashboard and watch a needle twitch? So real time dashboards tend to use aggregates and averages.

But, just as in a report, an average can hide a multitude of sins. The devil is in the individual transaction level detail, and in its context. It's all very well looking at a dashboard to see what your call volumes are right now but without knowing what it normally is, or what it should be, that doesn't help much.

It is akin to asking you to drive a car where your dashboard has a speedometer showing the average speed yesterday, and the fuel gauge gives the fuel remaining in the tank at the end of last week. Oh, and you can only use the rear view mirror because there is no windscreen!

Clearly you wouldn't drive your car like this – well at least not for long! But this is how we run businesses today, based on a rearward view of the business. This reliance on out of date data by the BI industry ensures that the vision of BI 'in every business process' just isn't possible.

So while you might need a dashboard, what you do with it is also highly relevant. It makes no sense to have someone sit and look at a real time dashboard all day. Oh look, the needle's gone red… no wait, its OK, its back in the green again. What happens when they go to lunch? How do you identify the low level problems which tactically cost the business money?

Of course a dashboard is often seen as the way forward to present data to end users, and there is no doubt that dashboards go part of the way in addressing users' frustrations. In reality a dashboard is often little more than an on-line version of a report, probably with significantly less information, but in a more accessible manner.

Dashboards deliver the "what do I need to look at" which are the symptoms, but completely fail to address the cause. Remember that a dashboard is a bit misleading: we're using a real time metaphor of gauges and dials which we associate with real time instrumentation, where the underlying data is often static and out of date.

In many cases the dashboard may prompt us to look at the reference report to try and understand the detail of what's causing the dial to be 'in the red.' But by definition, this is a manual process: if an analyst has to go and look, whether at a dashboard or a report, there's a pretty good chance that it won't get noticed at all, let alone in time to react.

Of course the real point here is that the decisions that need to be taken in seconds have to be automated and automatically flagged to the user. Many dashboards include an 'alert' system which will trigger a notification when a particular value exceeds a threshold. Of course, it's usually up to you to set what the threshold is; in a complex business with many dimensions, this could result in millions of rules and thresholds, and a maintenance nightmare.

Chapter 5: End of the data warehouse?

When Gartner Group proclaimed the Death of the Database at the end of 2005, it generated a controversy surrounding the database's future role. You can listen to the podcast here. Blogs were written on both sides of the debate: explaining how wrong they were; and how right they were. Message boards filled with chatter debating the merits. Gartner was both right and wrong at the same time.

What they were really getting at, was that the role of the database is changing, and in particular where we persist data, and how it's utilized. From a business point of view, this utilization is about building adaptability to change into IT. Gartner contends that the business logic should not necessarily be stored in a database, but should be closer to, and relevant to, the business where it is to be used. So in short, we should put the relevant data where it needs to be used, and view storage as support for business process tasks, not related to the database that it happens to be stored in.

Specifically relating to BI, Gartner state that:

> *"Business Intelligence is something that's going to become part of every application. I want to be able to do analysis on my data relative to the current events that are taking place. Therefore the functions of Business Intelligence are going to be embedded in applications. So that functionality has to persist and therefore the ability to access the data to perform the functionality has to be accessed as well".*

What Gartner is getting at is that a central repository of information, such as a data warehouse, will not serve real time applications well. Therefore the application logic and the persistence of the data required to give context to the business events, needs to be within, or very close to, the application itself.

So Business Intelligence needs to be centered on business processes and less detached from day to day operations. Gartner concludes:

> *"Business Intelligence, and business decisions are not just going to be based on historical data, but sources of data that are immediate, mid term and long term all combined into a process that... gets built into the applications. Today [this] is called real time event processing... which [over time will] become the norm"*
>
> *"Business Intelligence is going to become pervasive, not as a stand alone application, but as something that exists throughout the application infrastructure".*

If Gartner is right, this clearly has ramifications for the BI industry: BI tools are detached from processes, but clearly there is a need to make processes smarter – using BI techniques. So what does all this mean for the data warehouse specifically? Does this signify the end of the data warehouse?

The answer is no and yes, at least as we know it today. The data warehouse doesn't go away, however its role will change significantly over time.

Today, in many organizations, the data warehouse is perceived as the sole central repository for historical data, and the place you go to analyze data.

But what is clear is that the data warehouse will continue to be the repository for historical data – the system of record if you like. If you need to analyze historical data which has been gathered from multiple systems, the data warehouse is the place for the data.

What's less clear is the role that the data warehouse will play going forward in analyzing data as part of a business process. It's clearly a near impossible task to make the data warehouse 'always on,' available 24x7, real time information on demand, available to analysts and systems alike. Take a moment to consider all the compromises involved here: The more that you think this through, the more you end up heading down the road with Gartner.

"The information needed for day to day operational decisions needs to be held in a way that is centered on the business process in hand, perhaps as part of the applications themselves."

The information needed for day to day operational decisions needs to be held in a way that is centered on the business process in hand, perhaps as part of the applications themselves. Given all of the challenges related to real time data warehousing, it seems obvious that this analytic functionality relies on data 'stored' in memory, and persisted into near real time storage.

This may cause alarm from the data warehouse purists: this persistence is another store of historical data, duplicating what is already stored in the data warehouse. Or is it? This is not really true: these persistence stores will never become the system of record, but are there to enable in-memory analytic systems to recover in the event of some failure – after all they need to be designed for real time and high availability to cope with the demands of real time, process oriented business.

This data is also important in providing historical contextual information for the analysis of current business events. This is the *'immediate, mid term and long term'* data referred to by Gartner. In order to understand the significance of real time events, we need historical data to put it into context.

The other point about these persistence stores is that they usually contain data which is rarely held in the data warehouse. This data ranges from the lowest level of granular detail, to process state data which is usually

not deemed relevant to the typical data warehouse requirement. After all, the lower the level of aggregation in the data warehouse, the larger the database. Often, one level of de-aggregation can cause a data explosion in the data warehouse, so in most organizations, much thought has been given as to how much data is needed for the analysis in hand.

This gets to the root of the issues. Data warehouses and their related architectures were never designed to provide in-process intelligence and be available, with real time data, 24x7. Assuming you have the data, you can probably report on the effectiveness of a process, but providing analytics *within* a process is another matter altogether.

These two approaches are really complementary despite Gartner's inflammatory statements about the death of the data warehouse. Having a clear view about which approach to use when is the start of a meaningful strategy.

Chapter 6: Six mega-trends set to impact the data warehouse

"If you run a data warehouse and don't think that ID theft is going to impact you, you need to think again."

If you manage, or are responsible for, a data warehouse or Business Intelligence implementation, then there are six key mega-trends that are set to impact you over the next few years. Each of these will place stress on your ability to enhance your systems as increasing regulation, in particular, begins to bite. This is not an exhaustive list – you can probably come up with a few of your own.

It will be no surprise that these trends will pull you in multiple directions. The data warehouse has long been considered a compromise solution to a constantly moving target. This compromise, between minimizing impact on operational systems, load times, user query times and of course cost, is set to get more and more difficult to manage as the data warehouse gets asked to do many things that it was never designed for in the first place. While every situation is different, some of this will no doubt resonate.

Data security

You cannot have failed to notice that identity (ID) theft is a fast growing problem. Since the ChoicePoint initial disclosure, between February 2005 and June 2006 an additional 190 disclosures of personal data were made, on a potential total 88 million consumers.

If you run a data warehouse and don't think that ID theft is going to impact you, you need to think again. In the retail sector, MasterCard and Visa have mandated a set of data security standards (the **Payment Card Industry Data Security Standard**, or PCI for short) which impacts all businesses processing and storing credit card information. Every retailer and many financial services companies are affected. If your data warehouse has credit card information in it, it must be encrypted in order to comply.

This is just the tip of the iceberg. It's not just credit card data, or the retail sector that's going to be affected. Several national governments across the world, including the US, are in the process of drafting legislation that will mandate robust processes around the access to, and the encryption of, personal data related to individuals, whether staff, customers, or businesses.

Whether we like it or not there's a raft of government regulation coming down the pipe that will impact the way we use and store data in the data warehouse in the next few years.

"There's always an analyst somewhere that will argue that five years of data is essential for their analysis."

Data longevity

Increased government regulation is also having an impact on the amount of time that we need to store data, and the types of data to be stored. Whilst some of these regulations are broad, such as Sarbanes Oxley, others are specific to particular industries. For example, in financial services, Regulation NMS in the US and in Europe MiFID is set to impose a significant burden on the retention of data to be able to demonstrate that these new regulations have been complied with.

Within our businesses, users want to keep data for ever longer periods, despite its questionable worth. How valuable is pre 9/11 data for comparison purposes? The world is running on ever shorter cycles, yet businesses are reluctant to archive data. There's always an analyst somewhere that will argue that five years of data is essential for their analysis.

Even if your users are not interested in older data, your audit and compliance teams most certainly are. Accountability by the CEO and CFO for the quality of information published by companies has caused the audit function to be pushed to center stage.

"The advent of new sources of data, and specifically RFID tagging, means that data volumes are set to grow dramatically."

Data volume

As we touched on earlier, if there's one truism in data warehousing, it's that data volumes only grow. It's a one way street. There are three primary drivers: compliance; more sources of data; and de-aggregation (a need to store lower levels of granularity). Each of these is causing volumes to grow. For example, analyst firm **The Winter Corporation** has been tracking the average size of Microsoft SQL Server data warehouses. In 2001 the average size was 293 GB but in only 5 years this has grown to 3.0 terabytes on average, growing at over 200% per year. Whatever the numbers, data volumes are growing, pushed ever upwards by these three drivers.

Of course in some markets this growth is likely to be exponential. Take retail and consumer packaged goods where some of the larger data warehouses are found today. The advent of new sources of data, and specifically **RFID** tagging, means that data volumes are set to grow dramatically. Item level RFID tagging has the potential to cause a data warehouse size to expand ten fold over time, if all of this additional data needs to be in the warehouse. This of course begs the question about whether this data should be put into the data warehouse, but it is clear that some level of analysis will need to be performed on RFID data.

More users

Analyst firm IDC estimates that only 15% of staff in enterprises have BI tools, but it is clear that deployments to broader communities will continue. BI is no longer restricted to internal users but information in the form of reports is often shared with customers and partners. In many industries there are growing demands to share more than reports across entire supply chains, with data exchange and sharing becoming common.

Pioneers are also experimenting with sharing more than just static reports – increasingly customers are demanding more interactive access to information and even personalized dashboards. This trend looks set to continue putting pressure on the back end infrastructure to deliver high availability and high quality of service to customers and partners.

The dashboard replaces the report

OK, so this is a bit more of a reach, but there is an underlying trend here which needs capturing. Dashboards are the delivery mechanism of choice for "reports". Business people far prefer to have a personalized, relevant dashboard than go and look through standard reports. There are already BI tools that can take a report and turn it into a dashboard, so before long it will be common practice for reports to be replaced by ever more complex and interactive dashboards.

Of course the BI industry has some way to go in delivering against this. The personal, flexible and interactive dashboard has not yet been delivered, but the trend is clearly set.

BI becomes real time, mission critical

Business has become increasingly global. Maybe this has impacted your organization in only a small way so far, but macro trends such as outsourcing call center staff to India, and manufacturing to China, will affect most businesses over time. But whether we like it or not, globalization is here to stay. The impact on the data warehouse is significant.

As a consequence of having users accessing your systems on the other side of the world, 24x7 operation becomes essential. Whereas perhaps you currently rely on an overnight window to update your data, in a globalized company your time window disappears, forcing you to compromise service levels for some markets, or move to incremental updating of the data warehouse. Most data warehouses were never designed for this type of operation, and therefore there is considerable uncertainly as to how to transform your data warehouse into a robust 24x7 system.

"We are moving inexorably towards mission critical decision processing where BI is part of the process of doing business, not just reporting on yesterday's business."

Of course, it's not just globalization that is making BI mission critical. Business is increasingly being done electronically, whether financial trades on a stock market, bank payments, orders placed over EDI or online. The net result is that BI needs to adapt to these changes and the way that we do business today. EDI orders need to be validated, bank payments checked and trades confirmed. BI has a role to play in all of these processes, and as it begins to do so, we move inexorably towards mission critical decision processing where BI is part of the process of doing business, not just reporting on yesterday's business.

Chapter 7: Vision for the new BI 2.0

la·ten·cy (lāt'n-sē) *n.*

"The period of apparent inactivity between the time the stimulus is presented and the moment a response occurs."

The conclusion of Chapter 1 said that "Almost every business practice has adapted to shortening business cycles, except for the Business Intelligence world. Why do we run our business operations disconnected from the insights that could make us more money and reduce costs? in an 'always on world' where we run our lives minute by minute, wirelessly connected to a wide variety of information sources, why should we run our businesses based on out of date information? Clearly something has to change in the way that we (human beings) use technology to process and analyze data."

So what does BI 2.0 change and why?

In Chapter 2 we discussed how latency has such a dramatic effect on BI processes. New BI 2.0 architectures transform the way that we access data, analyze data and drive decisions.

The goal of BI 2.0 is to reduce all three of these latencies to milliseconds, effectively zero in practical terms, and therefore to maximize value. Today this value is often an opportunity missed, or an expense, or a risk that has to grow large enough to be spotted manually.

Data latency can be reduced dramatically by using in-memory processing in place of storing data on a disk. Analysis latency can be reduced by automating the interpretation of the data, so that people do not need to look at every item, only the problems. In order to do this, of course, you need to be able to understand which business events are problems or could become future problems. And finally, decision latency can be eliminated in many operational (and tactical) decisions by automating the actions.

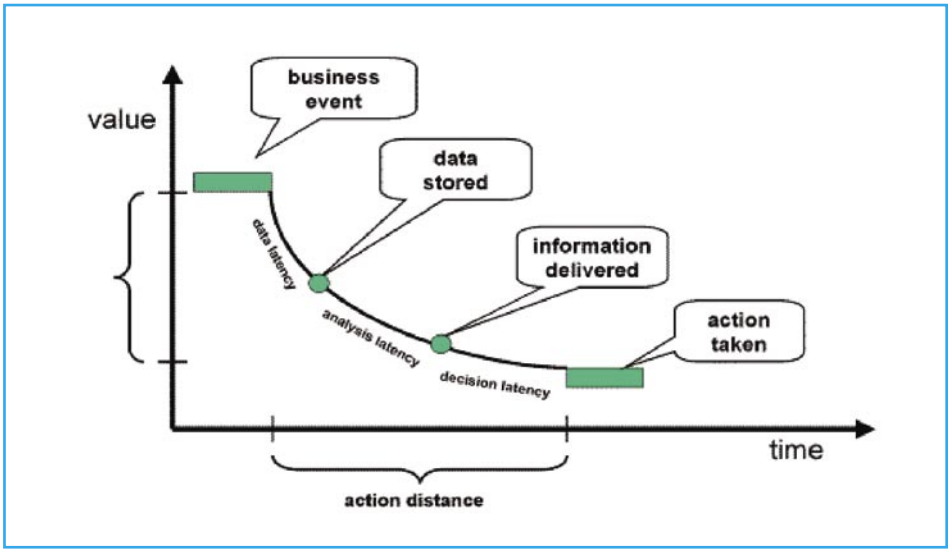

Reproduced with permission from Richard Hackathorn, Boulder Technology Inc.

"So in a BI 2.0 world, as stuff happens in your every day business, the BI system can automatically store the data in memory, analyze and interpret the event's significance, and in many cases, initiate a corrective action without any human involvement."

So in a BI 2.0 world, as stuff happens in your every day business, the BI system can automatically store the data in memory, analyze and interpret the event's significance, and in many cases, initiate a corrective action without any human involvement. And BI 2.0's ability to do each of these steps in real time is integral to any automated corrective action, or to the generation of an exception notification. In short, your business can become more intelligent, as processes become smarter.

This is often referred to as a 'closed loop' and while the idea is not new, existing BI architectures based on data warehouses don't fare well.

Indeed, many people haven't made the connection between real time and closed loop – without real time, the value of an automated action diminishes rapidly. For example in the CRM world, it's well known that a real time, highly targeted response to a specific customer event will achieve a 50% increase in response rate compared with an offer made days later.

So in order to deliver BI 2.0, very different architectures are required from the traditional data warehouse centric 'extract-transform-load-query-analyze' approach of today. BI 2.0 architectures are event based – i.e. these systems analyze streams of event data, not static data stored in data warehouses.

Many processes, of course, cannot be modeled easily, and cannot be automated. For example, in security applications, human intervention is essential. If a member of staff needs training, disciplining or dismissing as a result of their actions, this must be done in person.

What is needed under these circumstances is to detect the problem as rapidly as possible, and present the investigator with a concise alert, containing all of the facts needed to support the decision. It's critical that these alerts make the human decision process efficient – everything needed to make the decision should be contained in the alert itself. Today, much of day-to-day operations exception handling is characterized by reference to multiple operational data sources and reference reports to supplement and cross check the information.

While this may be the goal, in many businesses today the critical event remains elusive due to the volumes of data involved.

Event Driven BI

BI 2.0 is fundamentally different to traditional Business Intelligence. Data is not stored in a database; queries are not run against the database; data is not extracted for analysis. BI 2.0 uses Event Stream Processing.

It's worth just stopping to digest that.

"While this may be the goal, in many businesses today the critical event remains elusive due to the volumes of data involved."

We have become so accustomed to all BI being based on queries running against a data warehouse or some other underlying database, that this new BI 2.0 architecture needs considering carefully. After all, this is a radical departure. Event Stream Processing, as the name implies, processes streams of events in memory, typically in line or parallel with business processes. Its core capabilities are the ability to analyze massive quantities of data in real time, event by event.

Typically, this means looking for scenarios of events (patterns and combinations of events in succession) which are significant for the business problem in hand. The outputs of these systems are usually real time metrics; real time alerts; and the initiation of real time actions in custom or third party applications. The net effect, is that analysis processes are automated and do not rely on human interpretation and action. This real time, automated and continuous analysis provides dramatic returns and competitive advantage for those that have adopted it.

So how does it work?

"This real time, automated and continuous analysis provides dramatic returns and competitive advantage for those that have adopted it."

If you consider where data is in your organization at any point in time, it resides in applications, data warehouses (having been extracted from applications), and is flowing through middleware. Since extracting data from the source applications slows down their performance, BI 2.0 gets data "in-flight" directly from middleware.

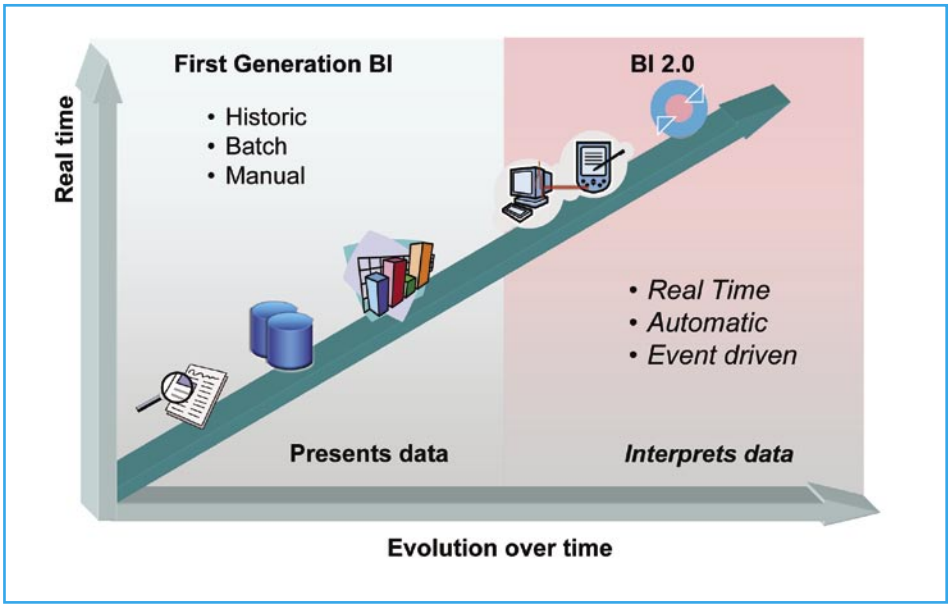

 |

By tapping into a flow of events, typically from industry standard middleware, BI 2.0 can analyze these streams of data to perform millions of calculations in real time that would be unthinkable in a traditional data warehouse environment. Middleware such as IBM MQ, MS MQ, Tibco, JMS, any EAI or ESB platform, SAP's Netweaver or indeed Cisco's AON blade in a network environment, can easily create a stream of events for analysis, similar to a third party listening in on a conference call.

"Knowing there is a problem sooner buys the organization more reaction time. You can act while the customer is still on the phone, or on the website."

This analysis is performed in memory – not written out to disk – and compares every individual event to what is normal for that unique account, customer, SKU, IP address etc. to identify those that represent problems or opportunities. It also enables a trend within a stream to be monitored, and in the most advanced systems, to project this trend forward automatically to understand the impact on the business in the near future.

Scalability

The ability to analyze large quantities of data in real time sounds very expensive, and indeed it would be in the traditional data warehouse world if you chose to follow this route. But there are major differences in the way that data is analyzed in Event Stream Processing technologies which fundamentally change the economics: each event is analyzed individually, not as a batch, and this, together with advanced in-memory computation techniques enables very high throughput rates of events to be analyzed on commodity hardware.

For example, if you were to calculate the sum of transactions by unique credit card over the last 30 days, in the old world, this would involve extracting data from a data warehouse from the previous 30 days for each credit card involved. Let's assume that that there are 10m credit cards involved, and that each card makes one transaction per day. In the old world your query would pull back 30 x 10,000,000 or a total of 300m rows of data. This sequential processing would then involve 10m sums of 30 data points. If the same calculation was needed every time a new transaction is made, the entire calculation will have to be repeated.

In the BI 2.0 world, this doesn't happen. Using 'delta math', products like **SeeWhy** calculate the impact of the changes on the sum. So while 10m sums are being held in memory, when a transaction happens on one card, the individual sum for that card is calculated by adding the new value onto the existing total, and dropping off any values from the other end of the sum as appropriate. Moreover, advanced caching techniques mean that the values in between the end points don't need to be stored, dramatically increasing efficiency, and reducing memory usage. So in fact you don't need to store all 10m data points to do the calculation.

The end result of these different techniques for doing these calculations means that Event Stream Processing engines are very efficient, and can typically process data an order of magnitude faster than query based approaches. This speed means that calculations can be done automatically as part of a process, whereas before this could not be considered.

If you can calculate metrics at an individual level, really fast, and in real time, then why would you introduce any latency? Concepts such as 'right time' have only been introduced because of the difficulty traditional architectures have in doing real time calculations.

Once this restriction is removed, then real time is the obvious route: after all, knowing there is a problem sooner buys the organization more reaction time. You can act while the customer is still on the phone, or on the website. Or perhaps a critical re-order cut off point is met, meaning that new supplies will be delivered to the store the next day. Clearly real time is the obvious goal of most organizations; it's just traditionally been very hard and very expensive.

BI 2.0 changes all that for good. There is no reason to rely on out of date information any more.

BAM	BI 2.0
Requires a defined process	Designed for any business metric
Low volume of events	High volume stream processing
Simple calculations	Delta processing
Aggregated, summary data	Lowest level of granularity
Manually set thresholds	Automatic thresholds
Presentation of data	Automated continuous analyses
Designed for process analysts	Designed for business
Displays simple process metrics	Displays business problems
Fixed time windows	Sliding time windows

Chapter 8: Intelligent processes

As you have probably already worked out from the thrust so far, there is a lot that today's BI has to do to keep up with where business is headed. We've witnessed waves of investment in office and process automation. It should come as no surprise that businesses seek ways to optimize that automation. Automating a dumb process, means that you can now take dumb decisions all that much faster!

The majority of processes aren't modeled, but are rather less explicitly defined. BI 2.0 needs to be able to work in both well defined process and in the less well defined areas.

Businesses want, and need, to make processes smarter. This is much easier than you might think – the means to do so are all embodied in the heterogeneous messaging technologies underpinning many of today's transactional systems and the current generation of SOA and event driven architectures. As the world has embraced business process management and event driven architectures, these in turn form the underpinnings for the next generation of process oriented Business Intelligence capabilities.

It's also worth noting that many processes themselves cannot be modeled and explicitly defined in a Business Process Management tool. In fact the majority of processes aren't modeled, but are rather less explicitly defined. Business users very often cannot describe their business processes accurately or specifically enough to be modeled. Yet despite this, these operational processes still need intelligence. So BI 2.0 needs to be able to work in both well defined process and in the less well defined areas.

The need for more intelligent processes drives some of the key requirements for new BI 2.0, summarized by the following characteristics:

Event driven

Automated processes are essentially driven by events, and therefore it's implicit that in order to create smarter processes you need to be able to analyze and interpret events. This means analyzing data, event by event, either in parallel with the business process, or as an implicit process step.

Real time

Real time is essential in an event driven world, but this principle is restated here due to any confusion surrounding 'right time'.

Real time is a critical element to analyzing events: without this, it is very hard to build in BI capabilities as a process step, and severely limits your ability to automate actions.

Batch processes by comparison are informational – i.e. they report on the effectiveness of the process, but cannot be part of the process itself unless time is not critical. Any application that involves trading, dynamic pricing,

demand sensing, security, risk, fraud, replenishment, or any form of interaction with a customer (for example via a web site or call center) are all time critical processes and require real time processing.

Sure, you can always do them slower, but optimizing these processes is where the competitive advantage lies, and to do this you need real time.

Automates analysis

This is less intuitive, but in order to automate day to day operational decision making, you need to be able to do a lot more than simply by presenting data. Putting real time data up onto a dashboard merely means that you have to pay some one to watch it!

Watching the needle twitch on a real time dashboard isn't all that useful. The challenge is how you turn real time data into something useful – something actionable.

In short you need to be able to automatically interpret data, dynamically, in real time. What this means in practice is the ability to compare each individual event with what would normally be expected based on past or predicted future performance. So BI 2.0 products need to understand what 'normal' looks like at both individual and aggregate levels, and be able to compare individual events to this automatically. Without this it is hard to automate decision making.

"Watching the needle twitch on a real time dashboard isn't all that useful. The challenge is how you turn real time data into something useful – something actionable."

Forward looking

Understanding the impact of any given event on an organization needs to be forward looking. So for example, "will my shipment arrive on time?" or "Is the system going to break today?" are both forward looking interpretations. This adds immediate value to operations teams that have a rolling forward looking perspective of what their performance is likely to be at the end of the day, week month etc.

Process oriented

To be embedded within a process, in order to make the process itself inherently smarter, requires BI 2.0 products to be process oriented. This is not process based (where the process has to be explicitly modeled in a BPM tool), but oriented around optimizing the outcome of a particular process where the process itself may or may not be explicitly defined. This process orientation is a prerequisite of any closed loop business intelligence where actions can be automatically driven from the results of analysis, or relevant operations staff alerted if the decision cannot be automated. So both closed loop and process orientation are key components of BI 2.0

Scaleable

Scalability is a cornerstone of BI 2.0 because of the event driven architectures on which it is based. Event streams can be unpredictable and can occur in very high volumes. For example, a retailer may want to build a demand sensing application to track the sales of every top selling item for every store. The retailer may have 30,000 unique items being sold in 1,000 stores, creating 30m store/item combinations that need tracking, and may be selling 10m items per day. Dealing with this scale is run of the mill for BI 2.0 – in fact this scalability itself enables new classes of applications which would never have been possible in a BI 1.0 world.

Conclusions

BI 2.0 represents both a bold new vision, and a fundamental shift in the way that we use information in business. It extends our definition of BI beyond the traditional data warehouse and query tool to include dynamic in-process and automated decision making.

In the past we have relied on out of date information, attempting to fix problems after the fact. BI 2.0 changes that for good. Now we can build business intelligence capabilities right into the processes themselves – in short: build smarter processes. These processes don't have to be explicitly defined, but rather are the fabric of the way that we do business.

The impact of this is sending ripples across the business world – transforming our ability to exploit ever increasing quantities of information in ever more sophisticated ways, automatically, and in real time.

Competitive advantage can no longer be gained just by implementing packaged applications and tools to automate processes. Simply storing data in data warehouses is no guarantee of any return – in fact there is lots of evidence that shows that data warehouses are underutilized, disconnected from process, and while they deliver lots of data out to the business, they often fail to give the business any real insight.

Leading companies are focused on automating processes, on building intelligence into processes – smarter processes if you like. Intelligence embedded within processes is the new frontier, where the competitive battles are increasingly being won or lost.

Find out more about BI 2.0

You can find out more about BI 2.0, its characteristics, applications and successes by visiting www.seewhy.com.

About the author

Charles Nicholls is Founder and CEO of SeeWhy Software and a Business Intelligence visionary. He incorporated the company in 2003 with a vision to create a new generation of Business Intelligence to revolutionize the way we analyze and use data.

Recognizing that traditional Business Intelligence is out of date and disconnected from day to day operations, Nicholls founded SeeWhy to 'eradicate dumb processes from our businesses.' Always on and never out of date, this next generation of BI enables users to build smart processes by building in intelligence into the standard operating procedure. This capability is now being referred to as "BI 2.0" and represents a fundamental shift in the way companies and government can now do business.

Under Nicholls' leadership, SeeWhy has grown from a ground-breaking idea into the technology leader in real-time event driven Business Intelligence. SeeWhy is a
Red Herring 100 Europe company, was recently picked as a leading 'cool vendor'
in the Business Intelligence space by analysts Gartner Group, and rated a 'Global Innovator' by analysts GuideWire Group.

Nicholls, a 20-year veteran of the software industry in the US and Europe, is internationally recognized as one of the pre-eminent thinkers in the Business Intelligence space. In 2003 The Massachusetts Institute of Technology selected Nicholls as a global leader in 'Intrapreneurship' following his highly innovative creation of Business Objects Analytics.

Nicholls is also the author of the controversial eBook "In search of Insight" which has established a new agenda for the Business Intelligence industry. In addition Nicholls is also inventor of several patents in the Business Intelligence space.

In his role as an advisor to leading global companies, Nicholls has worked with many of the world's top companies on their BI systems, including Amazon.com, BellSouth, Capital One, Citibank, Chase, Diageo, Ebay, GE, JPMorgan, Lands End, Marks and Spencer, MasterCard and Sainsbury's.

Before creating SeeWhy, Nicholls was an executive officer of real-time fraud detection solutions provider HNC Software. During his tenure HNC was successfully sold to FairIsaac Corporation (NYSE FIC) for $800m.

Prior to HNC, Nicholls spent 6 years at Business Objects (NASDAQ BOBJ), where as a leading BI visionary he founded and led the Business Objects Analytics division and 'Skunk works' engineering team.

Nicholls also founded Ithena, an award winning Silicon Valley CRM analytics startup, which Nicholls merged with Business Objects in 2000. During his tenure he drove revenues from concept to $28m in three years.

Published By SeeWhy Software Limited

About SeeWhy Software

SeeWhy is the first real time Business Intelligence platform for the event driven enterprise. SeeWhy continuously analyzes and interprets streams of individual business events, to alert you immediately to opportunities and risks and enable everyday decisions to be automated. Scalable for the most demanding applications, SeeWhy puts 'instant insight' into your business processes, giving you real time visibility into current and future performance. SeeWhy has been selected as one of the top private European software companies for 2006 as part of the Red Herring 100 Europe, named Global Innovator by Guidewire Group and highlighted as a cool company by Gartner Group. SeeWhy Software was incorporated in 2003 and is headquartered in Windsor, UK.

www.seewhy.com