From Business Intelligence to Enterprise IT Architecture

February 2011

Dr. Barry Devlin, 9sight Consulting
barry@9sight.com

Data warehousing / business intelligence has played a remarkable role in IT over the past few decades. It has remained a relatively closed ecosystem, both in terms of the vendors involved and the delivery teams within IT organizations. Signs are that that's all about to change as the boundaries between operational and informational computing become increasingly blurred and Enterprise 2.0 finally catches on in business. This article explores the business and technological drivers for this change and its far-reaching implications for architects, developers and vendors alike. And it describes the architecture that's needed to move the focus from business intelligence to the wider world of enterprise architecture.

From a series of articles published on B-Eye-Network.com from February 2010 to January 2011.

Contents

Part 1: Introduction

Part 2: Architectural Overview

Part 3: The Business Information Resource (BIR)

Part 4: The Second Axis of the BIR

Part 5: The Third Axis of the BIR

Part 6: Myths about decision making and process

Part 7: The Business Function Assembly (BFA)

Part 8: The Personal Action Domain (PAD)

Part 9: Implementation considerations for BI2

Part 10: Three Key Transition Steps

Part 1: Introduction

Consistency, integration and other dirty words

Reading some vendor materials and even analyst commentaries these days, you'd be forgiven for thinking that data consistency, cleansing, reconciliation and integration were swear words. Or maybe that all enterprise data has been made so clean at source in recent years that business intelligence can just get on with analysis of whatever comes to hand, and focusing on ever sexier ways to display the results. If only that were so…

It's worth recalling that among the key drivers for data warehousing, and later business intelligence, were information consistency and integration. Users required consistent reports and answers to decision support questions so that different departments could give agreed answers to the CEO. IT wanted an integrated set of base data to avoid time-consuming and costly reworking when the CEO got conflicting answers. This was a key emphasis I placed in the first ever article describing the data warehouse architecture, published in 1988 in the IBM Systems Journal[1].

Now, more than two decades later, far more weight is placed on speed and flexibility of decision making, and ever more sophisticated analytics. We see this in diverse technology trends from operational BI through data appliances to "post-relational" databases and stand-alone analytic and dashboard environments. Sometimes information consistency is assumed—"we have a data warehouse; that will take care of it." To the extent that the data in question passes through the warehouse, that's fine; but, a closer look often reveals that key parts of the informational environment stand stubbornly apart from the warehouse.

Financial, technical or data timeliness concerns often lead to "warehouse bypasses". In such cases, the requirement for clean, consistent data is conveniently forgotten as exciting new data marts and appliances are sourced directly from potentially inconsistent or unclean operational systems. And, of course, contradictory and incompatible results and reports emerge later to haunt the development and maintenance teams. In the history of data warehousing, this is at least the third iteration of the "independent data mart" disease. It emerges with particular virulence in times of economic stress, when vendors need quick sales and buyers are more willing to listen to stories of great bargains. And the cause remains the same: insufficient attention to cleanliness and consistency of the source data.

And then there are the resurgent direct query proponents, who seem to believe that "mashups"—to use the modern term for an old approach—of any data, anywhere, performed on the fly and all mixed together can solve all users' information needs. Enterprise information integration (EII), federation, virtual data warehousing and a variety of similar terms for the same technique have been around since the earliest days of BI. In the past, eliminating the need for a big store of reconciled data was a key driver. Today, it is the increasing business demand for ever-closer to real time data for decision making. Again, it also suits vendors of such tools to declare that little or no data modeling is required and deployment can therefore be rapid and low-cost. But, the oft-forgotten question is: does the data being combined in the real-time query actually belong together?

Add to this soft information[2]. Such information, from e-mails to voicemails, from Word documents to YouTube videos, from Tweets to Blogs, is growing dramatically in volume and is seen as increasingly important as a source of valuable knowledge for many aspect of business. It is also notoriously unclean, inconsistent and difficult to integrate. Typically, vendors coming from the search and enterprise content markets have only a limited understanding of the extent of data integration needed

---

[1] Devlin, B. A. and Murphy, P. T., *"An architecture for a business and information system,"* IBM Systems Journal, Volume 27, Number 1, (1988) http://bit.ly/EBIS1988

[2] I use the terms "hard" and "soft" information to distinguish between data that has been structured for traditional computer uses, such as data entry, computation, analysis and summarization, and data that is stored and used in a more free-form way, often called unstructured information. However, the phrase "unstructured information" is actually an oxymoron—information, by definition, has a structure that imbues meaning. Unstructured data is, in reality, nothing more than random noise!

and expected in BI. And their assumptions about the level of automation required in assignment of meaning, selection of relevant records and integration with hard information are often far less stringent than those encountered in BI.

What's happening in all cases is that a fundamental assumption that has underpinned data warehousing from the earliest days is coming under increasing stress. It was clear from the beginning, and it remains so today, that when decision making is taken to an enterprise level, a largely consistent, enterprise-wide base of information is required. The original data warehouse architecture assumed that this consistency could be created as data enters the warehouse. Modern business needs and long-standing practices largely invalidate this assumption. The practice of bypassing the warehouse when need be removes any opportunity to enforce consistency or enable integration. The business need for ever-closer to real-time decision making creates significant technical barriers to reconciliation of data if it does enter the warehouse. And the proliferation of soft data, especially from unreliable sources such as the Web, adds an even greater issue of defining what consistency means in the first place.

As this and other original postulates of data warehousing[3] are coming under increasing strain, the time comes to ask if the original architecture for data warehousing needs to be revisited. And it is to be hoped that the process will not be as painful as it was for Galileo when he suggested moving the earth from the center of the universe!

A changing landscape for IT

While ongoing issues of consistency and integration graphically illustrate some of the changing information landscape that data warehouse developers encounter, the rest of IT is facing equally daunting changes in business expectations and technology. Among the most important is service oriented architecture (SOA). Driven initially by a need for flexibility and adaptability in the operational environment, SOA is creating a new process-oriented, plug-and-play integration approach for operational applications. While there continues to be ongoing debate about the practicality of vendor approaches and the viability of fully implementing SOA in the short to medium term, there is little doubt that the days of enormous, custom-built, monolithic applications are numbered. They are simply too inflexible for modern business. Commercial off-the-shelf (COTS) applications such as SAP and Oracle Applications are already promoting SOA as a way to integrate and innovate within and around these packages.

Initially, SOA was aimed specifically at the operational environment, but its scope certainly extends to informational and collaborative applications in the medium to longer term. As business users become used to the concept that they can (or should be able to) link together existing services into a workflow they need to do their job, they will correctly begin to question the difference between these classes of function: Why can't we plug an analysis step into the workflow to understand the likely impact of delaying this shipment? How do we link into the e-mail system to notify a customer automatically of an order fulfillment problem? By now, these questions are commonplace and vendors increasingly offer point solutions in dashboards, business process management (BPM) flows and so on. But the underlying architectural questions are often avoided. What does all this mean for the traditional division between operational and informational data? Is process finally becoming a strong requirement in informational applications?

Web / Enterprise 2.0 approaches are also dissolving the old boundaries between operational, informational and collaborative function by reframing user interactions in a looser and more user-directed social environment. And as these old (and artificial) functional boundaries break down, so too does our traditional division between operational, informational and collaborative information.

[3] Devlin, B. "*Business Integrated Insight (BI²)—Reinventing enterprise information management,*" http://bit.ly/BI2_White_Paper

Today's business needs and technologies demand a new architecture

The above considerations, and more, lead to a daunting but obvious conclusion. Modern business needs and current technological trends demand a fully consistent, integrated set of information that spans the formerly separate worlds of operational, informational and collaborative activities. For the informational world, the original data warehouse provided such a store. But, unlike data warehousing, this information cannot be stored in a single database, but must be distributed throughout the entire IT infrastructure. Furthermore, Enterprise 2.0 makes it clear that the nature of this information is expanding from hard information to include a wide variety of more complex, soft information such as documents, audio and images. Consistency and integrity needs apply (to varying extents) to these information types as well. So, the key question being posed about information is: how can we create a new base of integrated and consistent information for the *entire* enterprise?

The simpler the answer, the better the solution. If you want to create a consistent, integrated information resource, you must stop creating duplicates of existing information that have to be managed to consistency, and you must eliminate, or substantially reduce, existing data duplication. The original data warehouse architecture did this. It proposed a logically single data store—the Business Data Warehouse—modeled at the enterprise level as the consistent and integrated source of all information for decision making. This simplicity was ultimately lost with the emergence of the layered architecture, due to a combination of database performance and enterprise modeling issues.

Nonetheless, the approach remains valid for the current much-expanded needs for integration. First, model all the information according to an enterprise-level model and then implement as far as possible in alignment to that model with minimal duplication. This is the approach proposed in Business Integrated Insight (BI^2), an architecture[4] that gathers all the information of the enterprise, hard and soft; operational, informational and collaborative into a single logical component. In addition, a second integration point, largely absent from the original data warehouse architecture, is absolutely mandatory today. A logically single, consistent and integrated set of processes that spans all aspects of business and IT needs is required to enable the flexibility and adaptability that modern business requires. This is also included as part of BI^2.

[4] Devlin, B. "*Business Integrated Insight (BI^2)—Reinventing enterprise information management,*" http://bit.ly/BI2_White_Paper

Part 2: Architectural Overview

In part 1, we've how current business trends demand greatly improved integrity in the information and processes that lie at the core of every enterprise today. Because of the ever widening scope and increasing speed of business, this need for integrity extends to issues of consistency across multiple data stores and types, timeliness of information, reliability of processes, support for social networking and, indeed, to every aspect of IT systems. The proposed solution was to look back to the original driving principles of data warehousing and understand how they apply in the current situation and where they need to be enhanced. Part 2 provides the initial view of a logical architecture, BI^2, based on these considerations.

Sense and respond and the (non-)evolution of business

In a previous series of articles entitled *"BI is dead, long live HEB"*, I introduced the powerful concept of *sense and respond*, first defined a 1992 Management Review article by Stephan H. Haeckel as an approach to managing an adaptive business in a continuously changing world[5]. Even three years ago, when I wrote these articles, it was clear that business was in need of (and technology was moving towards) an integration of all aspects of business activity. As seen in Figure 1, repeated from that series, I termed a business with such a level of integration as "highly evolved", because such a business would be able to plan for and react to faster and unpredictable change in its external environment. Such behavior requires advanced sensors of external events, a closed-loop feedback system for monitoring the outcomes of action taken, as well as more sophisticated organizational structures for decision-making and more highly networked social interactions than found in a traditional business.

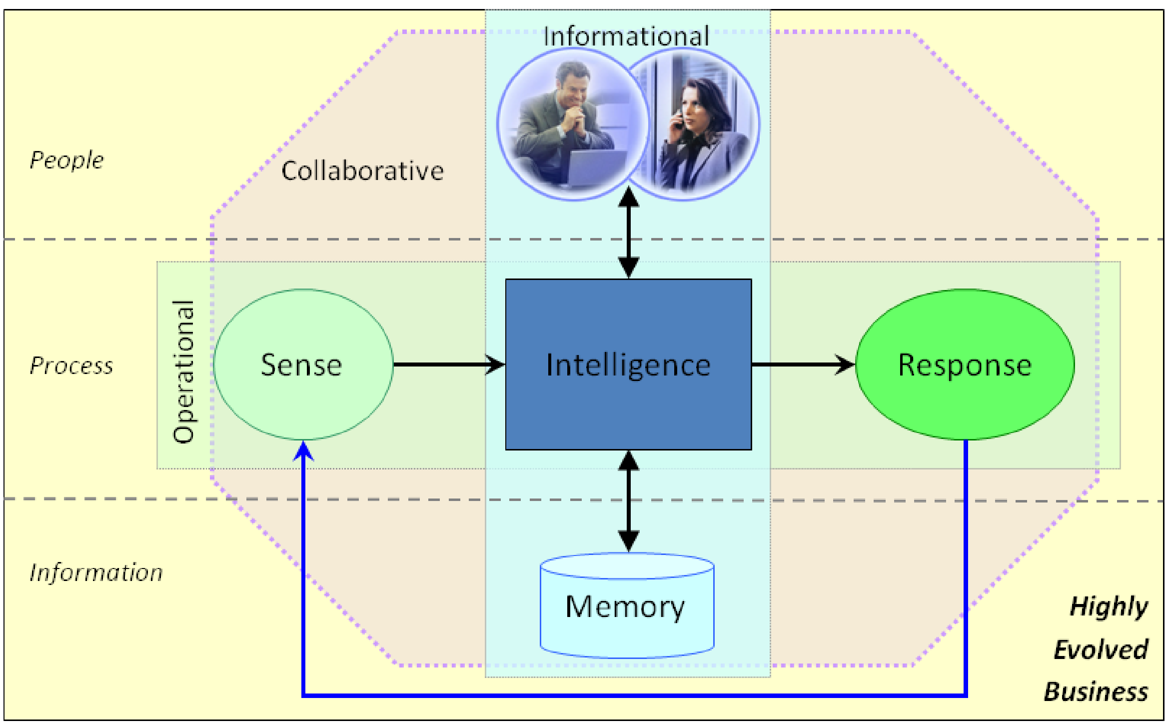

Figure 1: The sense and respond loop and the highly evolved business

[5] See *"Adaptive Enterprise: Creating and Leading Sense-and-Respond Organizations,"* Harvard Business School Press, 1999 and www.senseandrespond.com for further details.

In the intervening years, we've seen an increasing drive in business for speed of response and a demand for ever more timely and accurate sensing. However, there has been limited recognition of how this drive for speed demands extensive integration of information and deeply linked processes. Once a business has reached a certain size or level of complexity, advances in speed can no longer be achieved unless these integration issues are solved. This level of complexity has been reached in many businesses, and the combination of rapid action-taking combined with a lack of integration (of both information and processes) can lead to a rapid spiral into dysfunctional behavior and faulty decision-making. Such effects have been seen dramatically, for example, in the recent near-collapse of the financial industry.

In the interim, I've stopped using the term highly evolved business, because, in my opinion, business structures and approaches have not really reached the required level of sophistication yet. But, technology has certainly advanced, allowing us to begin to conceive of the architectural structures that are required to allow such business evolution to occur.

People, process and information

The sense and respond model described above provides the foundation for this new architecture by identifying the three levels at which integration must take place—people, process and information. As shown in figure 2, these are the three layers of the Business Integrated Insight (BI^2) architecture. This article introduces the structure and the layers; later articles will delve deeper into each layer.

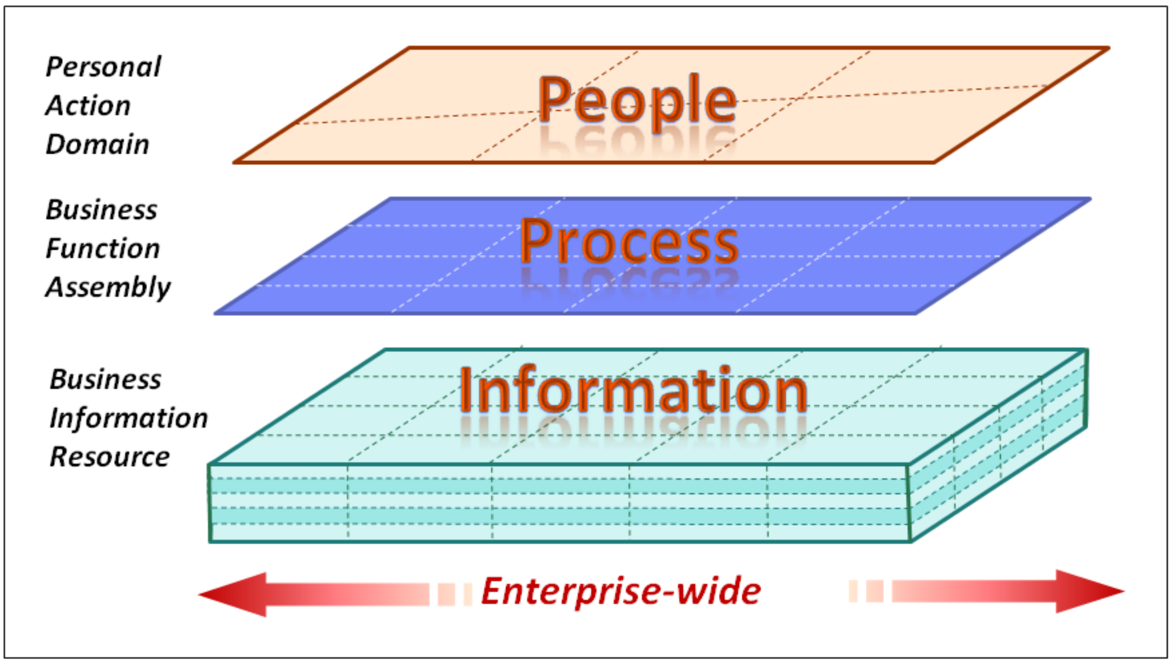

Figure 2: The three layers of the Business Integrated Insight (BI2) architecture

Information: the core of business

Let's begin with the information layer, as this is the most familiar layer to data warehousing and BI professionals. This layer, called the Business Information Resource (BIR), contains *all* the information needed by the enterprise to conduct its business. And when I say "all", I do mean *all*! This includes, of course, the information traditionally stored in a data warehouse and all the data marts, as well as the operational data used to run the business. It also covers soft information such as e-mails, documents and even images and other such content. It stretches from information stored in servers to that on users' PCs, as well as the cloud. And, in addition, it includes all information from external partners, regulatory bodies and even the internet used by the business.

The extent of this information is broader than most architectures would dare to include. Some might argue that it is foolhardy to include such a breadth of information. Whilst agreeing with the premise that the scope of information covered presents significant challenges in organization, definition, quality management, maintenance and other aspects, I contend that there is little choice but to address this scope. Information from each and every one of these areas contributes today to the fundamental and basic tasks of running and managing a business. If excluded from consideration in the architecture, any part of this information set can contribute to a level of disintegration in the information asset of the business. The extent of the contribution varies considerably depending on the information involved. The level and type of disruption relates to such traditional concepts as operational vs. informational data, data vs. content, internal vs. external or personal vs. centralized information. But we need a deeper and more formal analysis of how information is created, structured, used and managed.

Consider the impact, for example, that erroneous, personal spreadsheet information can have on the calculation of published business results. Or how messages on Twitter can contribute to knowledge of customer satisfaction with a product. The role and relevance of all information is modeled in the BIR, as will be described in part 3. For now, we simply recognize that the BIR structure is more complex than the layers above it—shown by its three dimensional depiction in contrast to the other two and the number of information classes represented by the cells in the layer. The bottom line is that any information that can add to or subtract from the overall integration of business information should be within scope of the architecture and the impact on such integration of any set of information must be well defined and understood.

Process: running a business with panache

The history of running a business has seen a steady increase in the level of automation and consolidation of the activities needed to achieve the goals of the enterprise. Early automation of the manufacturing process, for example, moved from artisan creation of individual products to an assembly line where each action was optimized and linked together in a defined and managed process. This approach integrated processes that were traditionally separate within the business, and today, just-in-time manufacturing moves process integration beyond the walls of the business.

Within IT, we have seen operational systems develop from individual, stand-alone applications to highly integrated systems such as SAP, creating integrated processes across vast swathes of the business. Service oriented architecture (SOA) has further driven both wider and more flexible integration by creating processes consisting of well-bounded and defined services linked by user-defined workflows. This development has two very significant implications.

First, activities previously beyond the normal operational processes of the business become available for integration within existing processes. One place where this clearly appears is in operational BI. Here, queries or analyses of specific short-term trends are embedded within operational process flows as services in the process. Users are thus encouraged to transition seamlessly between two very different ways of working as and when needed. Operational concepts of process and response time extend into the informational world. Similar overlaps appear between operational/informational and collaborative function. The result is that process orientation extends beyond operational applications to become pervasive across all types of computing activity.

Second, SOA enables users to define and change workflows. While more in theory than practice in current implementations, this functionality is becoming more pervasive and is a key direction for SOA to enable adaptive and flexible workflows unimpeded by IT bottlenecks. Closely related to this, but coming from Web 2.0, is the creation of mashups by business users. As a result, processes such as "application" design and ETL that were formerly the preserve of IT can potentially be performed by business users. The process layers must therefore cover both business and IT processes equally.

These considerations imply that the process layer in the BI^2 architecture, called the Business Function Assembly (BFA) has a considerably wider scope than might be traditionally assumed. At its heart, the BFA is a process definition, operation and management environment, as can be seen in SOA, business process management and similar tools. But with a scope extending into informational and collaborative activities, Web / Enterprise 2.0 function also plays a strong role.

People: acting together or falling apart

The top layer in the architecture, the Personal Action Domain (PAD) describes how users behave and interact with the two underlying layers of process and information. Of the three, this layer is the most loosely defined as yet. User motivation, intention and behavior are more the domain of psychology than IT. However, a model of how people act to manage, run and innovate in the business sphere is vital to the definition of workable processes that users actually *use* and the creation and maintenance of usable information.

In the future, the PAD layer must include the model and rules that describe behavior as well as some form of engine that uses this model to bridge from user intent to the processes required to deliver the intended outcomes. This model and engine must reflect the highly adaptive nature of human behavior in order to enable the highest possible level of innovation in the business. Today, however, this layer consists only of rather standalone components of user interface functionality. This layer, together with the BFA layer above, will be described in more detail in part 4 of this series.

Conclusion

So far, we have but skimmed the surface of the Business Integrated Insight architecture. However, the drivers are clear: enabling flexible and adaptive behavior within a highly integrated information and process environment. The goal is to enable business users to react rapidly and safely to the most subtle changes in the external environment and to clearly relate real world outcomes to specific decisions made and actions taken. Companies exhibiting such business integrated insight are increasingly gaining business advantage in a rapidly changing marketplace.

Part 3: The Business Information Resource (BIR)

Part 2 introduced the Business Integrated Insight (BI^2) architecture as an evolution of data warehousing / business intelligence. The scope of BI^2 is far broader than these previous approaches because of the growing need for a holistic and integrated view of all the information, processes and people's activities that comprise a modern business. Part 3 drills into the information layer, the Business Information Resource (BIR) to understand its drivers and to describe a key part of its structure.

Information topography and the information space

Figure 2 on page 5 shows the overall layered structure of BI^2 and the placement of the Business Information Resource in the overall architecture. As shown by its position as the foundational layer of the architecture, information is the basis for addressing the often competing needs of the business for integrity and flexibility, completeness and timeliness. And, as previously shown, these needs extend to *all* information accessed and used by a business—from operational transactions to data marts and from contractual documents to images on or from the Web. Which such a diverse set of information types, we need some form of classification to discuss their characteristics and to understand the opportunities and limitations of working with them.

Information topography provides this structure. Like physical topography, the information variety is all about understanding the "lay of the land". The objective is to create a map of the territory to enable navigation, decisions about use and so on. In physical topography, some map features are strictly measurable, such as map coordinates or height above sea level of a contour, while others are based on a human judgment: for example, is this land wet enough to be a marsh, is this a hill or a mountain. As we look at information topography, we also use a combination of measurable and more subjective criteria to classify the information. But while physical topography is a well-established discipline with well-known criteria, information topography is much more recent and less well defined.

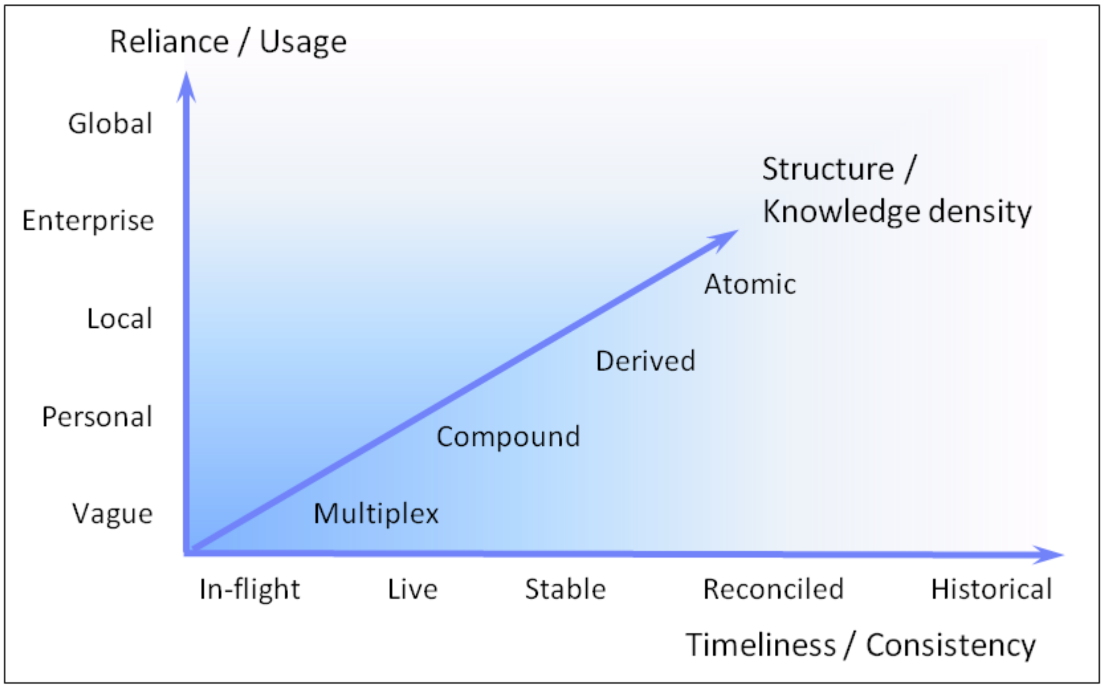

Figure 3: The three axes of the BIR information space

When discussing information types, a binary classification approach has generally been used in the past. Information is classified as operational or informational, hard or soft, real-time or historical, personal or public, etc. While this is a convenient way of talking, it cannot be easily used to decide how to treat a particular set of data; there are simply too many binary options. As a more usable approach, and in the absence of any established methodology, I have defined a three-dimensional information space, the three axes of which classify important characteristics of information. These three axes are shown in Figure 3 and are, of course, the source of the cubic grid representing the BIR in the overall architecture.

Before describing these axes in detail, there are two points of note. First, the axes represent to some extent combinations of related characteristics. I have done this to reduce the number of axes; I cannot draw a greater than three-dimensional space and only mathematicians can conceive of n-dimensional space (where n>3) anyway. As a result, one can envisage other axes describing the same space. Also, there are potentially additional axes describing characteristics that may be of interest. Consequently, this representation may evolve. Second, the axes are continua and the classifications along these axes generally merge from one to the next: they are subjective to a large extent. While this may be deemed unscientific, my experience suggests that this type of loose definition is the most that can be achieved in the real, business world and will provide a sufficient level of definition to architect the virtual and distributed information store that is the BIR.

The independent axes of the three-dimensional space are:

1. **Timeliness / Consistency:** Describes how information moves from creation for a specific purpose to broader, consistent usage and on to archival or deletion.

2. **Structure / Knowledge Density:** Describes the journey of information from soft to hard and the related concept of the amount of knowledge embedded in it.

3. **Reliance / Usage:** Describes the path of information from personal to enterprise usage and beyond and the implications for how far it can be relied upon.

The remainder of this article deals with the timeliness / consistency axis and the technological implications of this thinking. The other two axes will be considered in Part 4.

The TC axis: timeliness / consistency

Data warehousing has long described the concept of data moving from operational to informational or historical. This is the basic idea of timeliness. Operational data exists at a moment in time and is correct at that moment. It is not safe to assume that it is correct at some later time. Informational data, on the other hand, is deemed to be correct over a period of time—it has historical validity. It records valid states of the business over time. In order to do this, the consistency of data as it is combined from disparate sources has to be considered. Data has to be reconciled as it moves from the operational to informational class. This thinking leads directly to the TC axis and the classes depicted on it. And while the above is described specifically in terms of hard information, soft information can be seen in the same manner.

Although, in principle, timeliness and consistency are unconnected, in practice in IT, they are inversely related. Because logically related data is often created and manipulated in disconnected or dispersed processes, it can be very difficult and/or expensive to ensure that a newly created or changed data item is instantly consistent with related data items. A new invoice, for example, may be created with a local copy of a unit price at the precise moment when the master price file is being updated. Until reconciled, these data items are inconsistent, and if used together in an operational BI application, for example, will create errors or confusion. If we wish to use them together for more tactical or strategic decision-making, we must wait until they have been reconciled.

These considerations lead to the five broad classes on the TC axis. *In-flight* data consists of messages on the wire or the enterprise service bus, valid only at the instant it passes by. Traditionally, such data is recorded in a database, which we term *live*. (In high volume, high speed event processing scenarios, in-flight data may not be stored at atomic level.) Live data has a limited period of validity and may be inconsistent as mentioned above. *Stable* and *reconciled* data, which are stable over the medium term, are the next classes on the TC axis. In addition to being stable, reconciled data is also internally consistent in meaning and timing. *Historical* data extends the period of validity and consistency to, in principle, forever and wholly consistent.

Technology implications of the TC axis

In this section, we focus exclusively on hard information (atomic and derived classes on the knowledge density axis in Figure 3 on page 8) because of its familiarity to readers of BeyeNETWORK and its current high importance to decision-making. Similar considerations apply to soft information, although with less urgency.

The TC axis is an enhancement of the view that led to physical layering of the data warehouse architecture into operational, enterprise data warehouse and data marts. This physical separation of layers leads to a dubious conclusion: that the boundaries between these different types of information can be cleanly drawn. However, BI^2 defines the TC and other axes as *logical* and the classes on the axes are described as deliberately fluid. This reflects the real world. It can be very difficult to decide when a particular data item at a certain moment in its lifecycle is, for example, live or stable. Each class boundary poses a similar problem. More interestingly, a particular data item can logically change from one class to another (usually the adjacent one) without any physical movement or copying. The physical layering of data in the traditional warehouse architecture along these logical and porous boundaries is now seen to be somewhat arbitrary, and strict adherence to it can lead to unnecessary physical data copying and significant problems with timeliness, consistency and management of multiple data copies.

So the question arises: why were these physical layers instantiated in the first place? One answer is that technological limitations on database size, performance, optimization, etc. have mandated that multiple data stores were required. The logical and loose boundaries between information classes became, almost by default, the physical margins of the data layers. Given recent advances in database technology and hardware performance/cost ratios, the opportunity now exists to revisit those boundary decisions and database choices. Over much of the last decade and more, the dominant technology along the length of the TC axis has been traditional, row-oriented, general-purpose relational databases. For larger volumes of data and faster performance, most of these databases have taken advantage of some form of parallel processing.

The past few years has seen the emergence of column-based databases such as Vertica and ParAccel (not to forget the veteran Sybase IQ), which offer very significant performance improvements in query times. The vendors have focused on data mart offerings in the stable and historical information classes for a combination of technical and marketing reasons, but there is every reason to expect that columnar databases could mature sufficiently in areas of workload management and such that would allow them to make inroads in the reconciled data area. These developments could allow a substantial reduction in the complexity and data duplication in the EDW / data mart area.

But what of the long-standing physical boundary between in-flight / live (operational) information and the classes to the right? The recent emergence of hybrid databases—a combination of row- and column-based approaches—from Oracle, Vertica and others suggests that even this boundary may be breached. (Of course, fundamental issues around data cleanliness and completeness in operational data must also be addressed.) Such databases could potentially serve all classes of hard information equally for most small and medium-sized businesses. With further development of parallel processing and in-memory database support, they could possibly support all but the largest data sizes. Of course, at such an early stage of its development, much will need to be done to prove and harden this technology to enable businesses to invest with confidence in it; it is a serious change with wide-spread possible implications, both negative and positive. The risks of such a deep-seated change in technology will be weighed carefully against the enormous potential benefits of significant reductions in the number of copies of data, improved quality, reduced management costs and increased timeliness and availability of information.

Taking a more holistic view of information, together with recent advance in technology, we can see exciting near-term possibilities to simplify and streamline information for decision-making.

Part 4: The Second Axis of the BIR

The KD axis: knowledge density

The knowledge density axis is an enhanced view of the binary hard / soft information distinction first mentioned in part 1. Now, here's an intriguing thought: every piece of hard information you can think of was once soft. Address line in customer database: hard; customer sign-up form: soft. Order line in SAP: hard; customer call, e-mail or whatever placing the order: soft. ATM transaction: hard; customer deciding to use ATM: very soft. The fact is that all hard data originates in the real world, which is far more loosely structured than anything in the hard world of computer data. This thought may feel somewhat philosophical, but its relevance comes from the fact that all "data processing" (as IT used to be called) begins with a set of steps that gets us from soft information to hard. First analyze the structure and content of the soft information describing a real world event. Second, build a model of the key fields and their relationships that can be represented in a database (hard information) and third, create a process to ensure that the fields are filled every time the real world event occurs. In this way, soft information is transformed into metadata (field names, meanings and relationships) and hard data (values).

Note that I'm *not* proposing here that any or all soft information you come across should be transformed in hard information. Far from it; my thesis is the exact opposite. But, what I am describing are the fundamentals of the process by which the data used in computer-based applications is constructed. Such hardening of information is a prerequisite for all traditional data processing which requires very well-defined and understood facts, in largely numerical or categorical form, that can be created, updated, deleted, and manipulated—summarized, averaged, plotted—in precise, well-understood and repeatable ways.

From the above discussion, the knowledge density axis clearly represents a key IT process of structuring information into a form more amenable to computer processing. Elsewhere, I have called this the "structuredness" axis, but the spell-checker objects! The axis does represent the fact that the level of structure increases as we move along it. The label "knowledge density" reflects a perhaps more interesting observation that as we move from soft to hard information, more explicit knowledge is packed into smaller amounts of data. IT generally sees increasing knowledge density as a positive move—processing is faster, storage is smaller and algorithms are simpler. However, the down side of this process is that some of the more tacit knowledge residing in soft information can be easily lost. This happens when such knowledge has not been explicitly recognized in the modeling step, and thus is not captured in the hard data. The negotiations and documents exchanged leading to a large services contract, for example, contain far more information than is represented in the order entry line item. Such knowledge loss explains why it doesn't make sense to convert all soft information to hard data and throw the original content away.

Soft information exists in varying levels of complexity: hence its label "*multiplex*" on the KD axis. The complexity ranges from the simplest form of plain text, through formatted text, audio and image to video. A future version of this architecture may find it useful to explicitly separate the textual and audio categories from image and video. At the opposite end of the spectrum, hard information falls into two classes, *atomic* and *derived*. Atomic data contains a single piece of information (or fact) per data item and is extensively modeled. It is the most basic and simple form of data, and the most amenable to traditional (numerical) computer processing. Derived data, also usually modeled, typically consists of multiple data items that are derived or summarized from atomic data; the latter process may result in data loss.

As previously mentioned, hard data, both atomic and derived, stands separated from the metadata that describes it (and without which, it is meaningless!). But, where does metadata reside in this model? Traditional data warehouse architectures place metadata off to one side of the main "business" information. BI^2 explicitly defines metadata as a key component of the business information resource. And the *compound* class on the KD axis is where metadata resides. Strictly speaking, the compound class is defined as information that contains both hard and soft information elements together. Consider a "structured" text format such as XML. Raw textual and other soft information as well as tagged hard information reside in the same store. Tags express metadata and delineate structures within the soft information. Because of these properties, the compound class is ideal for business information that has intermediate or mixed levels of structure, as well as for metadata, which shows similar characteristics.

Technology implications of the KD axis

From a data warehousing point of view, soft information has long proven somewhat of a challenge. The initial approach is to treat it like any other data source—extract, transform and load any such information required by the business into the warehouse. Job done! But, not so fast. Relational databases (RDBs) provide only limited support for text and BLOB (binary large object) data. Even if we confine ourselves to text, where RDBs have increased search and analysis function over the years, the function and ease of use that users have come to expect from Google-like tools for text is largely lacking. Native XML databases and XML extensions of general purpose RDBs offer more extensive text functionality. But the big problem is that the growth in volumes of soft information already seen and predicted in the future militates strongly against any wholesale copying of soft information.

Content management systems (CMS) take a different approach to data warehouses when providing access to soft information. Rather than copying content from diverse sources, a CMS builds an inverted index containing pointers to all source occurrences of all significant words or phrases. In more advanced tools, various text analytic processes are also performed at this ingestion phase to add meaning via entity extraction, clustering, sentiment analysis, or classification. To some extent, each incoming piece of content is modeled on the fly. All later search and analysis activity occurs against this index with users routed to the required source only as the last step in the process. These indexes are the exact equivalent of metadata in a data warehouse environment, but their manner of use in CMSs shows how widely distributed soft information can be used without extensive copying into an informational environment.

In real data warehouses, metadata has, despite our best intentions, a peripheral role. Of course, it is used in all queries to identify tables, columns and relationships. However, such metadata resides in the heart of the RDB and is seldom even considered explicitly beyond the modeling stage. And most of us are acutely aware of the longstanding limitations on building and using business meaning, data sourcing and cleansing metadata that plague most data warehouses. As a consequence, analysis of hard information occurs directly against that information itself and relies only marginally on the metadata.

So, why does analytics occur so differently against soft and hard information? While there is a difference at the end-point—hard information analysis goes to the deepest level of detail in the data—the process of searching for relevant information and joining across different sources is very similar. And the CMS approach of extensive use of metadata provides significant advantages in terms of context, agility and ease of use. These considerations lead to the concept of a "Unified Information Store" for soft and hard information and are discussed in depth in my recent white paper[6] on this topic. The conclusion reached there is that while hard and soft information will continue to be best served by different storage options, with neither one being copied to the other, business users do require a consolidated and integrated view of both. This view is best achieved through the novel, direct use of metadata, both what data warehousing traditionally considers metadata and the advanced inverted indexes created in the content management environment.

The growing business need for meaningful convergence of hard and soft information has long presented data warehousing with particular challenges. Now, an understanding of the knowledge density axis of the BIR provides the basis for a modern approach to the problem. This approach steps back from copying soft information into the warehouse and instead provides integrated access through an enhanced and expanded set of metadata, which is an integral part of the Business Information Resource.

---

[6] Devlin, B., *"Beyond the Data Warehouse: A Unified Information Store for Data and Content"*, May 2010, http://bit.ly/UIS_White_Paper

Part 5: The Third Axis of the BIR

The RU axis: reliance / usage

As we've seen in parts 3 and 4, the axes of the information space were chosen to correspond to fundamental issues that arise whenever we want to use information beyond the context of a single application. The TC axis addresses the basic issues of timeliness and consistency of information that arise when we combine data from two or more sources. The data classes on this axis are the foundation for classical data warehousing. The KD axis arises when we look beyond traditional hard information (or "data" in common computer usage) and consider the increasing need to include softer information (or "content") within the scope of the architectures of we use to structure and manage business information. As we saw, this axis provides important new insights suggesting that soft information should *not* be copied into traditional data warehouse structures. Instead, metadata should be used to bridge between these different classes of information.

The RU axis addresses yet another fundamental information management issue. This issue emerged at the moment it became possible to copy corporate information into personally managed storage. The oft-discussed problems arising from the use of spreadsheets[7] in business illustrate the issue concisely. When a business user accesses data in the corporate BI system, there is a justifiable assumption about the reliability of that data. The accuracy, consistency and timeliness of the information is guaranteed by the information management processes and procedures embedded in the data warehouse and elsewhere. However, the moment this information is downloaded into a personal spreadsheet all reliability bets are off. The information can be changed deliberately—perhaps even with criminal intent—or accidentally. Erroneous calculations may be embedded in the cells of the spreadsheet. The level of reliance that can be placed in this data is considerably less than can be placed in its data warehouse source.

Reliability on its own is interesting. But, it's what you do with the information—usage—that makes this topic compelling. If our business person above uses his spreadsheet simply to check up on current sales and calculate the likelihood of reaching his month-end targets, the reduced reliability is relatively unimportant. If he brings the results to the monthly sales forecasting meeting, having the correct numbers and calculations is more important. And what if the senior vice president for sales likes spreadsheet so much that she decides it should be used by all her reportees for the next year? It becomes pretty obvious that widespread use of such low reliability information could potentially lead to some significant problems.

The RU axis provides a conceptual basis for examining this problem and deciding how one might best address it. As in the case of the other axes, we identify a number of broad classes of information, bearing in mind that the axis is a continuum and that the class boundaries are relatively fluid.

Let's begin in the more familiar territory around the centre of the axis and work outwards from there. The *local* scope of reliance / usage implies that a set of information is well understood and trusted by a defined set of users. This user set may be a department or business functional area, or may be a team temporarily assembled to undertake a particular piece of work. As we move to the *enterprise* class, we mean that information here can be relied upon throughout the enterprise in structure, meaning and content. Of the hard information stored in the BI environment, data marts typically fall in the local class, while the EDW clearly resides at the enterprise level. Note, however, that the RU classes should not be defined in terms of applications or databases; in fact, the axis represents a social / organizational construct describing how information can be used in the context of its reliability.

While the local and enterprise classes may be familiar, the *global* level is seldom explicitly recognized. Information at this level can be trusted and used in the widest arenas. At its most extreme, one could even envisage a universal class of, for example, physical and mathematical constants such as the mass of a proton or the value of pi. In a business sense, the global class refers to information that

[7] Eckerson, W. W. and Sherman, R. P., "*Strategies for Managing Spreadmarts: Migrating to a Managed BI Environment,*" TDWI Best Practice Report, First Quarter 2008, http://bit.ly/ZVx86

can be known and trusted between organizations, at a government or international level or within an industrial or geographical grouping. In today's interconnected and highly interdependent world, information at this level of reliance / usage is vital. However, as we have seen in the ongoing, international financial crisis, such trust is in short supply, based partially on the lack of reliable information.

Moving in the other direction along the RU axis, we come to the *personal* class, which consists of information defined and managed by one person, such as spreadsheets, documents and other files mainly residing on personal computers. As discussed earlier, the assumption that information which is relied upon personally can be used at local or enterprise levels causes significant data management and quality problems. In fact, similar problems arise whenever we try to use information created at one level of reliance / usage at a higher level on the axis.

Vague information, the lowest class on the RU axis, is information whose reliability and possible use is indeterminate or unknown. Information from the Internet often falls into this category. Compare Wikipedia, for example, with Encyclopedia Britannica. Both address the same need—reference information—but, Wikipedia information is vague: its provenance is unknown and its reliability open to question. Encyclopedia Britannica falls into the global class: its reliability is vouched for by reliable authorities. Wikipedia is free; Britannica costs money, and prior to the Internet used to cost a lot of money. This illustrates an important characteristic of the RU axis: moving information up this axis requires an investment in expertise, time and money.

Practicalities of the RU axis

In the case of the other two axes, I've discussed technical implications; here I prefer practicalities. While it's true that the RU axis springs largely from technological advances—personal computers and, more recently, the World Wide Web—the movement and management of information along this axis depends as much on improving organizational procedures and personal behaviors as it does on any particular use of technology. Let's continue to focus on the most obvious example of reliance / usage issues—the spreadsheet problem.

Both the joy and sorrow of spreadsheets spring from an early design decision to allow the user to manipulate all the spreadsheet data and create any desired calculations almost entirely free from any computational, tracking or auditability restrictions. As an exploratory tool, this makes perfect sense. And given its origin in the world of personal computing in the late 1970s, with its embedded opposition to the centralized management approach of mainframes, such restrictions would probably have been anathema to early developers and users. Unfortunately, from the viewpoint of reliance / usage, this opposition to tracking and auditability appears to continue to this day. With the widespread and presumably enduring use of such internally unregulated spreadsheets, the options open to IT departments for managing the promotion of personal spreadsheets to group or enterprise class usage are mostly limited to the implementation of best practices and procedures in the organization.

Ideally, what is needed here is a technology that blends the ease-of-use of spreadsheets with an underlying tracking and auditability infrastructure in an environment that supports innovation by individual users and manages the process by which such innovation is first shared and then promoted to enterprise use under managed and controlled conditions. I've discussed these considerations extensively elsewhere[8] in the context of a relatively new and widely praised analytic tool called Lyza[9].

Conclusion

The provenance of information used in decision-making has always been broader than the classical operational systems often considered as its sole source. With the increasing dependence on the Internet and other external sources, particularly for soft information, determining and tracking the reliability of information is key to managing its valid usage and improving its quality. The procedures

[8] Devlin, B., "*Playmarts: Agility with Control, Reconnecting Business Analysts to the Data Warehouse,*" December 2008, http://bit.ly/lN5UV and "Collaborative Analytics, Sharing and Harvesting Analytic Insights across the Business," June 2009, http://bit.ly/otVWP

[9]See www.lyzasoft.com for further product information.

for doing this—promoting information up the RU axis—are largely organizational and social in nature. However, I also believe that the inclusion of well thought out social networking function, close attention to tracking and auditability needs as well as support for promotion procedures in technology is possible. Such an approach is vital to allow IT to cost-effectively manage and control the much-needed promotion of information and its associated processes up the RU axis, so that personal innovation can be safely made available in the wider business as repeatable and reusable procedures in a timely manner.

Part 6: Myths about decision making and process

Part 2 of this document introduced the Business Function Assembly (BFA). As we saw there, a fundamental precept of the BFA is that everything, literally *everything* that is done within a business is part of a process. This thought is probably business-as-usual for developers and users of operational systems. In fact, to suggest otherwise to these folks might well be considered heresy. Traditional processes and business models focused on the strictly defined and regimented sets of procedures and steps required to automate manual processes. Whether talking about an automobile assembly line or the tasks involved in assessing and paying out on an insurance claim, the old approach of work-study (as described by F. W. Taylor in the early 20th century) and, subsequently, workflow software emphasized the definition of individual tasks and their linking together into optimized procedures or processes. Such processes were analyzed in depth and defined in advance, implemented in workflows and presented to users to follow slavishly and without question. Operational or production systems, of course, still benefit enormously from such approaches, gaining efficiency, productivity and reliability. However, for our this discussion, the most important aspect of such workflows is the inclusion within them of decision points—steps in the process where the workflow bifurcates (or trifurcates or more) on the basis of a decision. This decision, however simple or automated, is nonetheless a decision. Hold that thought!

Myth #1: Decision-making is a "no process zone"

In the informational world, the idea that decision-makers work within a process is anathema to many. Where, they ask, is there space in a process for creative thinking? How can users be innovative if confined to the fixed steps of a workflow? Such thinking, I believe, may demonstrate a somewhat outdated understanding of what processes and workflows imply.

Decision-making has historically been seen as a standalone activity. Use cases typically revolve around managers analyzing sales performance, business analysts seeking to understand historical trends or predict future directions and even executives tracking key performance indicators (KPIs). The focus is on the analysis and hopefully creative thinking that follows. Very seldom is the question asked: what happens next? The simple answer is that an existing process either continues as it was or is changed. The change may be simple and within the existing process structure—such as reassigning salespeople to different accounts—or it may involve a fundamental change in the process itself—implementing a new leads analysis step within the sales process. The point is that there was an existing process which came to a decision point, an analysis was performed, a decision made and the process resumed. In times of yore, the time between the decision point and process resumption could often be quite extended. In today's business environment, speedy decision-making is a high priority. Operational BI pushes this to the limit, attempting to reduce the time gap to minutes in many cases; and with automated processes, where decision-making is moved from people to software algorithms, sub-second response time as seen in automated stock trading, for example.

This business need for ever timelier decision-making in ever more situations raises a fundamental architectural question for application development: is there a valid boundary to be drawn between the operational and informational environments? While there certainly continues to be a substantial set of decisions with relaxed time frames (and we'll come back to them in a moment), my experience suggests that many decisions that were previously considered to be tactical—requiring resolution within a day or two—have gradually moved into that gray area between operational and informational applications—between half a minute and a couple of hours. From an implementation viewpoint, such "near-operational" decisions pose a greater challenge to data warehouse developers than to the providers of operational systems. In the latter case, with decision points occurring naturally in the workflow, implementations based on a service oriented architecture (SOA) or similar approaches can simply call whatever service is required to instantiate the decision point, assuming for the moment that the relevant information is available. The process flows naturally through the decision point from the before to the after situation. The data warehouse developer, on the other hand, is faced with the prospect of taking a previously asynchronous procedure—populating the warehouse and downstream data marts—and somehow bridging the gap between the operational processes before and after the decision point.

Of course, the operational system's assumption that the relevant information is available at the decision point is invalid in cases where current, cross-application or reconciled data is required to make the decision. Hence, the emergence of operational BI and an increasingly blurred boundary between operational and informational systems. From an architectural point of view, the increasingly prevalent blurring of this boundary was a key driver for the concept of BI^2.

Now, back to longer-running decisions. Is there any process involved in them? Or, perhaps more accurately, is there any process around them? Realistically, almost all business decisions of any significance demand the involvement of a number of people. Even the traditional business analyst, often seen sitting alone for hours in front of a screen full of spreadsheets, graphs, tables and pie charts, usually collaborates extensively with other people in the business both before and after this PC-centric phase. The prior work involves meetings with business people to understand their needs and investigations with peers and IT to source the required data. Subsequent work involves reviews with peers and managers, as well as presentations at meetings, all of which often drives substantial rework on the analysis, both revisiting business requirements and acquiring further data. This is work that all relies on a process, however light-weight or *ad hoc* in nature.

Similar considerations also apply to decision-making at a more senior level. Such work is highly collaborative, involving the creation of specialized teams, research into market trends, gathering existing documentation from a wide variety of sources, meetings, rework of initial and interim analyses, and so on. As anyone who has been involved in such decision-making knows, this is a process (and sometimes, an exceedingly tedious process at that!) Not only is it a process, but it is a process that is repeatable, at least in part. The next time a similar decision is required, a similar procedure—perhaps with subtle differences—will be followed, including some or all of the same people. In this, we can see a workflow with steps and connectors, and even decision points, but of a more collaborative and variable nature than traditional operational workflows.

With a little further thought, we can see a similar process applies to the case of our business analyst above and, in fact, to all decision-making. Decision-making, rather than being a "no process zone", is actually more correctly an *adaptive workflow*. So, let's restate the myth as:

Modern Premise #1: Everything a business needs to do today, whether operational or informational in nature, has a process or workflow of varying flexibility underlying it.

Myth #2: Decision makers are always seeking information, and the more the better

The vast majority of data warehousing and BI tools focus almost exclusively on data and information. Vendors sell them based on how much data they can handle, the analytic functions they provide and the types of decisions they can support. IT buyers tend to evaluate these tools in similar terms. You have to step beyond the pure bi market to performance or process management to find tools that position BI function within a process.

However, my experience of interviewing business users during the requirements gathering phases of data warehouse projects is that it often takes considerable effort to get most users to define their information needs in any detail. Again and again, these decision-makers describe what they're trying to do, what they hope to achieve and the steps they envisage in getting there. Surely this is process thinking, pure and simple?

Similarly, the drive to analyze and store ever increasing volumes of information comes as much from vendors and IT as it does from the business users themselves. For sure, there is a small cadre of statisticians and business analysts who understand the value of "big data" and know how to use it to best effect. But, the vast majority of managers and other day-to-day users of business intelligence tools work with relatively small quantities of data and only when there is no other option do they consider increasing the volumes of data analyzed.

Which leads us to:

Modern Premise #2: Business users are focused on results and the actions they need to take to get them; process trumps information—always.

Myth #3: Business and IT processes are fundamentally separate in nature

Business processes started out as written check-lists and simple manuals that users referred to until they knew the procedures they needed. And they rarely took out the book except to check on a rarely-used procedure. As the procedures were documented, the opportunity arose to automate them using tools to help define, manage and edit them as needed. Then IT took over and created an application that supported the process. This is all very well as long as any possible process change has been embedded in the supporting application. Of course, that's seldom the case. The business people, and especially the marketing and sales folks, often come up with bright ideas to include in the process, or perhaps just get bored with the way it is! When this occurs, IT processes are fired up. We have requirements gathering, change management, coding, testing, etc. And when the code is delivered, the business people run the new process. This is fine, except that sometimes the business folks change their minds faster than IT can deliver. Hence, workflow tools and business process managers appeared, leading on to SOA where users can change the workflow as needed and swap out one service and swap in another to react flexibly and quickly to market changes.

If you look closely at what has occurred here, a set of activities—(re)defining the workflow and replacing functional components of the "application" have just been moved from IT's responsibility to that of the business. Former IT process just became business process. Of course, the part where IT designs and writes the services remains firmly in the IT process area. Except when we look at what's happening in the social networking, collaborative world of Web 2.0 and beyond. Here the users are taking more and more control of what they want to see, right down to the definition and delivery of the services themselves.

Of course, IT will always have to run the processes for managing the infrastructure, deploying applications and so on. Really? I notice a lot of substantial hardware and software sitting beneath users' desks, being managed by the users themselves.

The point here is not whether this is good or bad. The reality is that users have been becoming more sophisticated in their ability to use computers since PCs were introduced in the 1980s. And to some extent, computers have become friendlier. Many activities that once demanded trained IT personnel are now easily handled by business users. The time has come to re-examine the business-IT boundary. By the way, the same argument applies to data, but here we're focusing on process. And here comes:

Modern Premise #3: Business and IT processes are all part of the same process environment; they all exist to serve business ends and need to be fully interlinked and interchangeable.

Conclusion

I've deliberately spent some time on fundamental, and indeed possibly revolutionary, process thinking here, because my experience is that data-oriented people like many readers of BeyeNetwork don't think too much about this aspect of IT. I believe that this needs to be rectified. If data warehousing professionals don't step up to the needs of process-orientation, the process-oriented approach of operational systems will encroach into BI without any consideration for the unique needs of decision-making. That would do a great disservice to the business.

Part 7: The Business Function Assembly (BFA)

Beyond the myths we knocked in part 6, there is another important change in the way business function is delivered that is widely accepted—the move from monolithic applications to plug-and-play assemblies of piece parts.

Traditionally, applications were written with a single, well-integrated purpose in mind. The thinking originates in the batch world. An application is loaded into memory, it reads input, writes output and (hopefully) completes with a zero return code. Without user interaction during runtime, applications were designed to handle as many functional steps as possible within the context of the information they read in the logic contained therein. With the move to interactive computing, the same application model was largely carried over to the new environment. Over time, monolithic applications moved to a more structured approach where functions and sub-functions were designed and programmed with well-defined boundaries and information-passing interfaces. But they still remained largely monolithic. By the late 1990s, it was becoming increasingly clear that these black-box applications, however well-structured internally, were too inflexible for rapidly changing business needs and too costly to maintain and upgrade. What to do?

Enter service orientation and, by 2001, the formalization of a service oriented architecture (SOA). I've described SOA in some detail in my 2007 article. Of course, by now there are many different vendor-dependent flavors of SOA. Beyond that, Web 2.0 has introduced further variations on the concept. Regardless, the principle is the same—software components of sufficient granularity to be recognized and used by business users are instantiated as services that have independent existence, published interfaces and can be searched and embedded, removed and replaced in larger components with minimal effort. While I have been engaged in many debates about the viability and practicality of SOA and similar architectures, in my view software developers have little choice today but to adopt one of these approaches if they are to satisfy modern business demands for flexibility, reuse and user empowerment.

This modern drive towards complete service orientation of business function enables the key concept of assembling and disassembling process flows in today's business. It is for this reason that the process layer is called the Business Function Assembly (BFA).

The process space

In line with the information layer (BIR) of the architecture we've already seen (and the people layer, yet to come), I describe the business function in terms of a process space. As shown in Figure 4, this space is two-dimensional, described by the Business Effect (BE) and Active Scope (AS) axes. Each axis is divided into four broad categories and, as in the case of the BIR, each axis represents a continuum of characteristics rather than discrete categories. Thus, although the categories are described as if cleanly bounded, in reality, they merge and blend one into the next.

Business function basics

Each piece of business function is created and exists for one, and only one, reason—to enable the business to achieve some goal or objective. The goal may be to understand and keep ahead of market trends, to keep a supermarket shelf fully stocked or to build and deliver Airbus A380 airplanes. An objective may be as small as recording a customer's change of address or as large as taking over its principal competitor. Goals and objectives come from the people layer of the BI^2 architecture, the Personal Action Domain (PAD); only humans can decide what a business' goals should be and set objectives for achieving them.

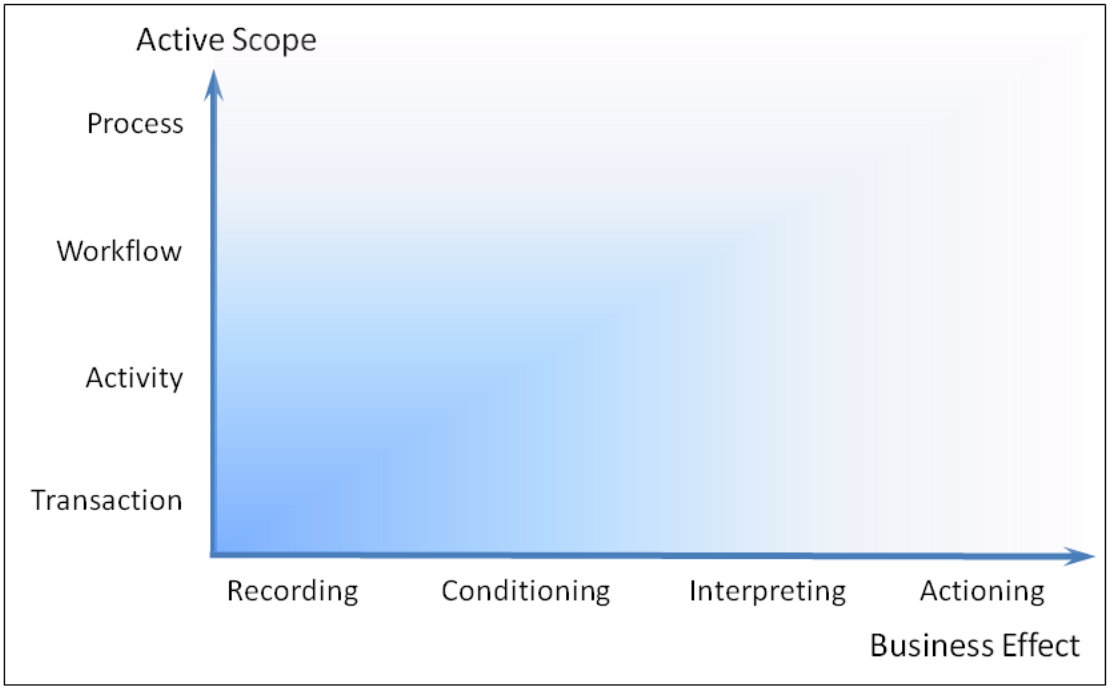

Figure 4: The two axes of the BFA process space

The role of the process layer is therefore to orchestrate the activities needed to achieve the goals and objectives of the business. Such activities can be performed by people, machines or software and, of course, automation increasingly shifts the burden from people to machines and software. Recognizing that an activity can, at different times, be performed by any one of these implies that the architecture must include all three within its scope. For example, the task of closing a sale may be performed by a salesperson, a vending machine or website. If the process layer is to accurately and completely represent the business, each of these possibilities must be represented. This representation occurs entirely in the BIR. Whatever its source—person, machine or software—where the fact of the sale comes together is in the information asset. The BFA is, in effect, an intermediary between the human intentions of goals and objectives represented in the PAD and the information that records and measures them in the BIR.

The BE axis: business effect

The BE axis therefore describes how a particular function which supports a goal or objective of the business interacts with the BIR. In traditional programming, this interaction was often described by the acronym CRUD—create, read, update, delete—conceived largely for transactional databases. However, this categorization is both too technology oriented and too directed towards traditional hard information storage. Modern business function must take account of soft information, often with very different consistency and commitment approaches; event data, some of which may never be permanently stored and far more beyond.

Furthermore, as discussed in Part 4, the BIR also contains metadata, another key aspect of the information asset of the business. The BFA must therefore also act on metadata, which defines and describes the structure and meaning of the information layer. This also follows directly from reality #3 listed earlier.

As a result of these considerations, we come to the following four categories on the BE axis. *Recording* function produces new business entities or instances; for example, create a new purchase order (hard information) or create an entirely new type of order (metadata). Such function operates on the leftmost side of the BIR. Recording function must create a permanent record of information for longer term analysis and auditability; current thinking about event data (in-flight information) that suggests some such data is never landed on disk may need to be reconsidered. *Conditioning* function modifies existing entities or instances; for example, update an order or archive an old order (hard information) or extract begin and end dates from a contract document (soft to hard information conversion). The effect of conditioning extends over the entire BIR, as it is the mechanism by which information is transformed and moved along the various axes of the BIR. *Interpreting* function uses existing information to monitor and understand what is occurring in and outside the business, allowing hypotheses about causes and effects to be formed. This function does not modify existing information, but can create new information based solely on existing information. *Actioning* function applies insight to interpretative hypotheses and existing information to take action. Such actions link back directly to invoke recording, conditioning or interpreting functions, thus closing the sense and respond loop, both across the full breadth of the axis and between individual categories.

The AS axis: active scope

As previously described, the BFA conforms to SOA principles and this is reflected directly in the AS axis. In a typical SOA, functionality is structured as well-defined, callable and largely immutable services that perform meaningful *activities* at a business level. Services may be built upon lower-level services recursively, and linked together into *workflows.* They can be inserted, replaced or removed with minimal technical expertise from a workflow in a plug-and-play manner. The AS axis reflects these characteristic categories, but extends them in both directions.

Transactions represent the atomic category on the axis and are the fundamental level at which function and information interact. Traditionally, transactions on hard information have strong ACID (atomicity, consistency, isolation and durability) characteristics. In BI^2, this generally remains the case for hard information. However, because of the breadth of information types in the BIR—atomic, derived, compound and multiplex—a much broader mix of transaction types must be included with differing ACID characteristics. For example, in the NoSQL paradigm (implemented in Google's BigTable and Amazon's Dynamo, for example), often only weak consistency such as eventual consistency or single data item transactions may be offered. At the other end of the axis, *processes* describe the highest level of workflows, spanning entire business areas.

Moving along the axis from the transaction end to the process end, definitional control moves from IT to the business. Processes and, increasingly, workflows are fully business defined. Transactions are clearly the province of IT and services mostly so. All functional categories, and their actions and interfaces, are described by metadata in the BIR.

Conclusion

As the middle layer, the BFA plays a key role in the BI^2 architecture, bridging users' needs to and from the information resource of the business. As the ultimate record of the business, the BIR is the storehouse of work-in-progress (operational information, both hard and soft), the record of information and function structure (metadata) as well as the long-term memory (for analysis, audit and archival). However, as previously mentioned, users typically think more in terms of activities and actions than of data and information. The BFA is thus the impedance matching layer from business users' thinking to the information heart of the business.

Part 8: The Personal Action Domain (PAD)

As we reach the top layer of the BI² architecture, we've moved far from familiar territory for most data warehouse implementation teams. One might say we have moved beyond the borders of IT. In fact, in the first article I wrote on BI², I said as much: *"…[the PAD] actually doesn't represent anything that can be physically instantiated in a computer environment!"* In the intervening year or so, I've changed my mind—it's a prerogative no longer reserved to women. But, I should explain why.

In my initial thinking, I wanted a place in the architecture to address what really goes on inside business users' heads and, indeed, hearts. The Personal Action Domain (PAD) thus represents the intent and behaviors of all users of the system, from executives to front-line staff to IT, within the company or external to it. To that extent, the PAD is beyond the physical computing resources of the company; it is actually the people who run and interact with the business. Understanding the scope of needs, goals and behaviors of users leads us into the domains of psychology, behavioral science, sociology, motivation and a host of areas far removed from IT, as we shall see in this article. So, fully separating these considerations from those of IT makes a lot of sense. However, it has become increasingly clear as I've drilled further into the details of the architecture that IT functionality is required to bridge from this internal and personal level to the BFA. So, rather than adding another layer to the architecture to represent this, I've revised the scope of the PAD to include this function. Ah! The joys of developing and publishing an architecture in parallel!

The basics of the human mind

So, before we move back onto the solid ground of IT, let's look briefly at some interesting facts about how we humans work, or to be more precise, how we perceive the world and how we react to those perceptions.

There is a largely unspoken assumption among business consultants and designers of IT systems that business, and the people who are that business, operate in a fully rational and largely predictable manner. Business processes are structured around logical if-then-else workflows. Information gathered consists entirely of that which is measurable and quantifiable. But, if you stand back and honestly watch the behaviors and reactions of your peers in IT or business, not to mention yourself, how much of what you see or do is truly rational, logical and predictable? Would you say 90%? Or would it be 50% or less? I'm not aware of any studies that validate my stance—and please send me references if you know of them—but my personal experience as a manager and a consultant, and the experiences of those I've asked, favor the latter range.

Most psychological theories of human development posit that rational thought only emerges around the age of seven and matures slowly over the succeeding years. Prior to that, learning is akin to a raw data download from the environment, mostly the home situation, with impressions unfiltered by any reasoning or conclusions about how the world works accepted in total from those of our parents. Those of us who are parents recognize just how quickly and comprehensively toddlers ingest all that they see and "learn" from it. We also recognize in ourselves how difficult it can be to reset some of these earliest learnings, even when our rational minds have proven them wrong. Simply watch where your mind goes in moments of crisis, the old beliefs that are instantly activated and the primitive reactions that arise unbidden and are later rationalized away! If you are interested in exploring further, I can suggest Daniel Siegel's excellent 2010 book, *"Mindsight"*.

So, how does this relate to the world of business and our IT architecture? In business, the implication is that at least some, or maybe many, decisions are based less on rational interpretation of solid information and more on a nebulous set of feelings, reactions and thoughts about the actual situation, the people involved and the perhaps unrelated context of what has happened in the previous hour. A tailback on the highway on the way to a meeting or a customer who unconsciously reminds you of your favorite aunt may have as much to do with closing a sale as the basic facts presented at the time. The level of innovative collaboration occurring in a brainstorming session is well-known to correlate with the personalities of the participants and their social skills. In the context of decision-making, smaller, shorter-term and more clear-cut decisions tend to be more fact-based, while larger, strategic decisions are more strongly influenced by soft factors, especially where the available information is less clear or conflicting. And, while increasing information availability or quality can certainly tilt the balance towards the rational in many cases, psychological studies leave little doubt about the continued importance of intuition or "gut-feel" for the most successful decision makers.

Our emerging enterprise IT architecture thus needs to reflect the reality that decision making is as much art as science. At a conceptual level, this is reflected in the axes of the people space as we'll see in the nest section. The underlying technology of this layer takes account of this particularly through Web 2.0 approaches and collaborative tools. However, this is still very much an emerging area.

The people space

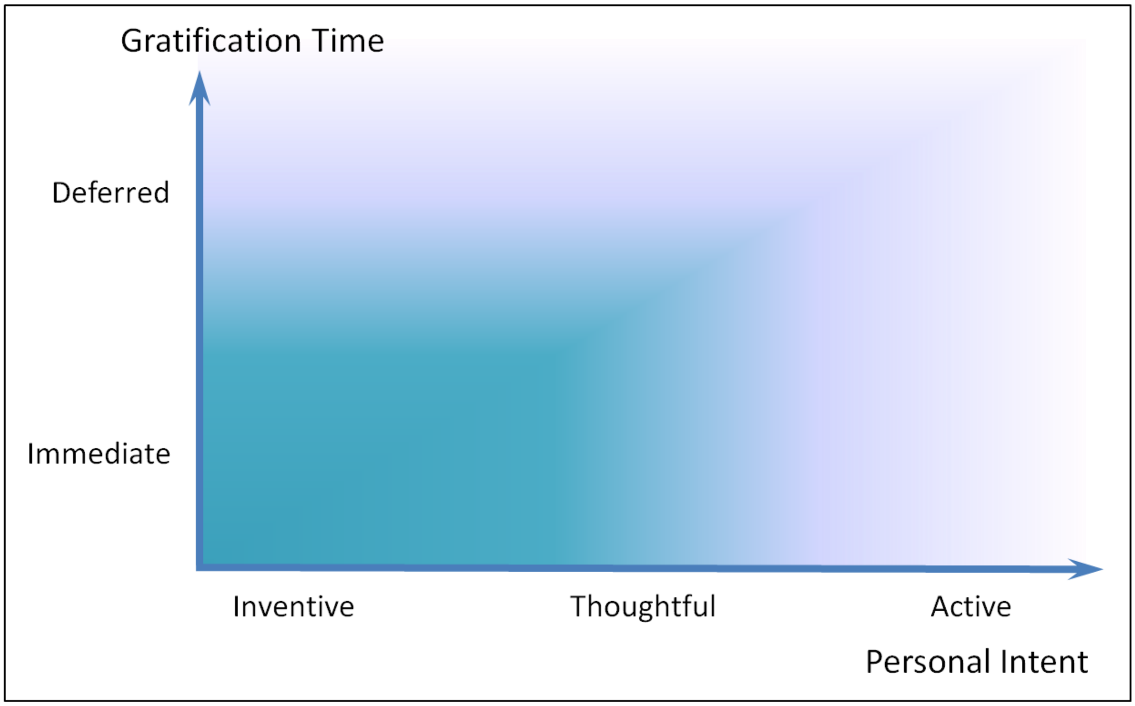

Figure 5: The two axes of the PAD people space

In line with the information (BIR) and process (BFA) layers of the architecture we've already seen, I describe the function needed in the PAD in terms of a people space. As shown in Figure 5, this space is two-dimensional, described by the Personal Intent (PI) and Gratification Time (GT) axes. Each axis is divided into a number of categories and, as in the case of the BIR and BFA, each axis represents a continuum of characteristics rather than discrete categories. Thus, although the categories are described as if cleanly bounded, in reality, they merge and blend one into the next.

The PI axis: personal intent

Modeling what a particular user intends in a specific moment is a vital step to understanding what type of business function is required in the BFA and what information is needed in the BIR to satisfy that intent. *Active* intent initiates action, causing something to happen. It thus corresponds to the type of behavior typical of running the business. Psychologically, it is termed "left-brain"—logical and fact-based, operating at a detailed level in a rapid and directive manner. It is thus very amenable to being part of a well-defined and rigid process or workflow. While most directly linked to the actioning function in the BFA, it also links to the recording and conditioning functions.

At the opposite end of the spectrum, is *inventive* intent, which is innovative and creative. Such behavior is "right-brain" oriented—intuitive, symbolic and creative in nature. In this intention mode, users create or recreate the business, its processes and its information content. This type of thinking is typically much slower and more passive than that involved in active intent and resists being locked into a predefined or rigid process. It operates largely internally in people, and doesn't link directly to a category of the process layer.

Between these two modes is *thoughtful* intent, which gathers information of any and all types for use primarily by the interpretive function of the BFA, although it also links to the recording and conditioning functions. A combination of left- and right-brain functions, this mode of thinking operates at medium speeds and particularly exercises social skills and is often highly collaborative in nature. It corresponds most closely to business intelligence in traditional thinking. Thoughtful intent requires an adaptive process to enable it—to provide a structure for collaboration, with the flexibility to change track when necessary.

The GT axis: gratification time

The gratification time axis simply reflects the normal reality of the world that some things you can have now, while others you have to wait for. In the modern business, there is a growing desire and need for *immediate* gratification, which places added emphasis on access to in-flight and live data. However, such data may not be internally consistent, as previously described, implying that *deferred* gratification may have to be endured.

Conclusion

The Personal Action Domain is, in many ways, the most important layer of the architecture: it is where all business really happens—in the human world of thoughts, emotions and actions. It is also the layer that has received least attention from IT consultants and architects so far, and even what I present here is very much a work in progress. Of all three layers, I suspect it is the one that will undergo most change over the coming months. So, stay tuned!

Part 9: Implementation considerations for BI^2

Waiting for the Evolution

Way back in part 1, I described the compelling rationale for a new approach to information usage in business that extended far beyond the traditional bounds of Business Intelligence and right to the edges of the entire IT architecture. Information, I argued, and the processes around it are fundamentally broken, still based on a set of basic postulates about business and IT that were once appropriate but that are no longer true. There are, however, an equally compelling set of arguments as to why we cannot hope to move from the old approach to a new one. These arguments can be summed up effectively as: *it's too hard!*

While it is undoubtedly true that moving from the current BI architecture to BI^2 is hard, it certainly doesn't follow that it is too hard. We need to ask: too hard compared to what? Too hard for whom?

The alternative to moving to a new architecture is mainly a case of staying with what we've got and continuing to add to it in an *ad hoc* fashion, as we've now been doing for nearly twenty years. The alternative to BI^2 is to stand aside and watch competing and overlapping approaches develop in our current three-way-siloed operational / informational / collaborative environment. The alternative is a continuation of the current user-unfriendly, business-destructive conglomeration of systems that have evolved over the past two decades and more. Which leads clearly to the answer to the second question: it may be too hard for IT to change it, but isn't it really so much harder for business users to live with it?

I contend, however, that find a way we must, no matter hard it may be, to move from our current situation for three simple reasons:

1. The current approach supports today's business needs only poorly at best and is incapable of being easily or rationally extended to meet even currently foreseen new requirements
2. As business demands more integration and flexibility, the current three silos will each try to address these needs separately and independently, leading to three incompatible and highly overlapping systems. This is already happening and is both highly expensive and very confusing
3. The cost and effort of maintaining the current environment continues to increase, even in the face of reductions in IT spend, leading to spiraling reductions in the ability of IT to respond to new business needs

We can no longer wait for the revolution; we need to evolve.

Evolve or else

Some modern theorists propose that evolution proceeds through a series of punctuated equilibria. In the equilibrium states, evolution is gradual and incremental, successful adaptations build upon one another in slow, steady, but non-disruptive progression. At the punctuation points, radical changes occur relatively rapidly. Old patterns die out quickly and new ones rise to dominance. Mainstream thinking continues to insist that the changes occur by random chance; a few radical scientists perceive an emerging pattern in evolution towards increasing complexity and self-awareness. Applying these ideas to the world of BI and IT technology evolution is instructive, as technology drives the direction of most business development today.

In BI, we have witnessed a period of approximately twenty years from the invention of data warehousing in the late 1980s where the dominant technologies have been relational databases, ETL propagation and stand-alone, analysis-focused BI tools. In each of these areas, change over most of this period has been steady and incremental—an equilibrium state. But now, we can see evidence of radical and rapid change; we seem to be at a punctuation point in all three layers of the BI^2 architecture.

Information: Relational databases make an evolutionary leap

Relational databases have long been based on a physical architecture where the rows of a database table are stored sequentially on hard disk and a variety of indexing techniques provide speedy access to individual records. It should be noted that Codd never specified this physical architecture when he developed relational theory; in fact, one of his primary goals was that the logical structure and physical storage should be completely independent of one another. However, this physical implementation worked well, particularly for operational / transactional databases with substantial levels of create, update and single-record access activity. Much of the development of relational databases since then has been incremental improvement in the indexing, query optimization and so on.

In the past five years, the relational database landscape has rapidly and dramatically shifted. It is worth noting that few of the individual changes driving this shift are radical or even new in themselves, but the combination has moved the equilibrium state. On the surface, the most obvious shift is a rethinking of the physical architecture that rearranges the sequence in which data is stored from row-based to columnar. However, a deeper technology shift at the hardware level is simultaneously at work. This is the emergence of commodity multi-core processors capable of parallel processing and cheap solid-state memory. Taken together, these shifts have combined to enable very substantial improvements in performance for BI-type usage. Even at the level of relational theory, new thinking is being applied in some databases. While the columnar storage technique has been the darling of the past few years, it is likely to be a shadow of the change in relational databases that will occur as in-memory, massively parallel processing databases, perhaps based on advanced relational theory, become mainstream.

Add to this the database demands arising from the explosive growth (estimated as greater than 60% CAGR by IDC's Digital Universe studies) of soft information in a highly distributed environment with relaxed consistency constraints. While some consultants and vendors blithely talk of storing such information in BLOBs in the traditional relational environment, soft information has fundamentally different characteristics than the hard information for which relational theory developed. A more realistic approach would be to build links or relationships, which may be quite different from those described by conventional modeling, between the traditional hard information stores and the mass of soft information. Such an approach is, indeed, mandatory to radically reduce duplication of information. IDC estimates that 75% of soft information is duplicated.

The quantum leap in database technology described above can support efforts to significantly reduce the number of copies of information stored in the enterprise, where such duplication arises primarily from performance considerations.

Process: ETL – Evolving To Linkage

The change in ETL technology has been considerably less dramatic than that seen in databases. While parallel processing has the potential to speed up ETL considerably, and certainly does in some cases, the impact has been relatively insignificant in data warehousing. One reason for this may be the fact that many warehouse developers still insist on writing bespoke scripts and programs to do ETL. More importantly may be the fact that the real challenge in ETL is about maintenance of the ETL processes as business demands and information structures change.

However, the fundamental problem for ETL as a species is that its habitat is under threat. The main, but still developing, threat comes from the evolution of the database environment which reduces the need and desire to copy information. On an ongoing basis, consolidation and re-architecting of operational data sources also reduces the opportunity for ETL to thrive.

But the biggest threat to ETL comes from federation or virtualization (or mash-up as it's sometimes called today). From a business viewpoint, federation offers one significant advantage over ETL—instant gratification. As sources of information proliferate, and most of them are under the direct control of business users, the tendency for these users is to join them on the fly. At a technology level, increasing network speeds and bandwidth support the desire to access remote information in real-time. IT is well aware of the many potential problems that this process may cause, but is often unable to resist. The situation is reminiscent of the IT war against spreadsheets; resistance is futile!

As the volume and variety of soft information grows, and if we accept that it should not be copied into relational databases, the role of federation grows and, proportionally, ETL declines. While both technologies have their strengths and weaknesses, federation's star is in the ascendant. Note also that this and similar trends are taking place in the operational world. Service oriented architecture (SOA) and master data management (MDM) both involve serious attempts to link operational activities together in real-time.

People: Evolving self-awareness and social networking

I characterize the traditional BI tools provided to end-users as being about "the data, the whole data and nothing but the data"! They stem from a world of analysis of information dominated by scientists and accountants, who can discover a wealth of meaning in the trends and details to be found in ever larger volumes of data. This approach continues to bear fruit and will not be totally abandoned anytime soon. But, another force is at work. Social evolution is at hand.

Today's decision makers are increasingly moving from data- and information-centric to knowledge- and wisdom-centric on the data-information-knowledge-wisdom spectrum. The underlying trend is that decision making is becoming increasingly personal and directly influenced by the social context of the time. Knowledge is often labeled as tacit; wisdom usually defies categorization. Social networking tools from e-mail (don't sneer, it was an important step!) to Facebook and Twitter, Amazon book reviews to personal blogs, have all driven decision making in the direction of personal opinions and peer influence. You may regret the end of the age of reason, but the tide is coming in. Decisions, and especially larger and longer-term ones, are made increasingly in a collaborative environment where analyzing data is only the first step.

This trend has been slower in emergence that those in the information and process spheres, but don't be fooled. The current and next generations of decision makers are going to spend a lot less time sitting in front of Cognos, Business Objects or even Tableau graphic data representations. They are sharing information and opinions in large volumes already, and much of that information and opinions emerges chaotically from the social network. So, any architecture that tries to impose a "single version of the truth", as traditional BI does, is doomed to failure.

Conclusion

The technology-driven evolution is thus pretty clear. We are at a punctuation point. What emerges in BI and, indeed, in the broader field of IT support for business will be significantly different than what dominated before. The dinosaurs are facing extinction; the mammals' time has come.

In the next and final part of this article, I'll be looking at some of the immediate steps you can take to prevent your BI systems going the way of the dinosaurs!

Part 10: Three Key Transition Steps

This article has shown that Business Integrated Insight is an approach that covers all aspects of traditional Enterprise IT Architecture and, indeed, extends its scope into areas such as collaborative computing and social networking that are often seen as ancillary to the mainstream of enterprise IT. Its implementation affects all aspects of IT support for the business and how business and IT interact. Undertaking change on this scale requires careful planning and in-depth change management processes. While considerable technological work is required, implementation will also demand substantial organizational change—between business and IT as well as within both IT and all business functions. The considerations for such a level of change are far beyond the scope of a short article, but the following is a set of immediate actions that you can undertake to begin to prepare for the process.

(1) Deduplicate across your information resources

As a BI expert in your organization (I assume that most readers either are or aspire to be BI leaders in their companies!), you are almost certainly aware of extensive duplication of data that exists within your data warehouse and related environments. As we've seen in previous articles, reducing duplication of information is a key driver and goal of BI^2. But are you aware just *how much* of your data is duplicated? And why?

Duplicate data exists between your data warehouse and data marts, as well as among the data marts themselves. If you have more than one data warehouse—despite my and other DW experts' unceasing advice to the contrary—duplication exists among these warehouses. Since these are resources under your own control, where better place to begin?

Reducing (or, better, eliminating) duplication between data marts is an ideal starting point. However, first you need to understand the underlying reasons for the duplication. Among the most common causes are performance and politics; and in order to tackle the duplicate data, you need to understand and address the underlying cause. Duplication for performance reasons is easier to eliminate because it is likely that the decision was taken when the performance of the hardware / software available for your data marts was much poorer than today. A well-considered upgrade may enable significant consolidation of data marts with attendant reductions in ongoing running and maintenance costs. Where duplication occurred for political reasons, more care is required. However, organization or personnel changes may well have occurred since the initial implementation, and attitudes may have softened. But, in either case, we are looking at reduced cost of ownership as justification.

Similar considerations apply to deduplicating back from data marts to the DW or even across DWs. Although costs here are likely to be higher and thus harder to justify, the ongoing running and maintenance costs that can be reduced may be higher than you think—they tend to creep up over time; so when was the last time you checked?

Another key area of duplication, and one that is growing rapidly today, relates to soft information. As this category of information grows in importance, more and more is being duplicated into the informational environment. This is a default behavior in many organizations, rather than one that has been specifically decided. It may also be a small amount of duplicated data today, but it will almost certainly grow rapidly. This is an area where a new approach using virtualization can be of significant value and the tools are improving rapidly. This leads nicely to the second immediate action.

(2) Examine your BI architecture

Most organizations' BI architecture is relatively old, at least in its guiding principles, which often date back ten years and more. While these principles may have been appropriate back then, there are two considerations that should be taken into account today: (1) how far from the original principles has the current implementation strayed and (2) when last were these principles and their implementation reviewed.

In the case of the former, based on my experience, many organizations' current implementation bears little resemblance to its stated architecture. Additional components have been added. Data now flows in diverse paths with different departments responsible. Tools implemented for particular purposes have further extended their roles. Systems that were planned to retire when the architecture was first designed live on, often with renewed growth and vigor. Depending on the level of consciousness of the IT organization as these changes occurred, their impact may be limited or significant.

If the architectural principles and their implementation have not been reviewed in the last two or three years, this would be a good time to do so. Much of what we have discussed in this series has been inspired by the significant technology changes of the past few years. These changes have the potential to reduce IT costs in some cases or to enable new business behavior in others.

And if your organization no longer has a BI architect or your BI architect hasn't looked beyond a relational database and ETL tool in years, engaging an external expert would probably be a good idea.

(3) Review BI and IT architecture fit

As you move forward with BI^2, looking beyond the bounds of the BI department is mandatory. We have seen that while the starting point for this architectural work was data warehousing, its end point includes the operational and collaborative environments as well. Understanding the intersection of BI with both of these environments will be vital.

Consequently, this would also be a good time to renew and expand your relationships with the IT departments responsible for the other aspects of the IT environment. This will enable closer collaboration in key areas of the architecture. Areas of high importance needing early attention include any SOA-like activities ongoing in the operational environment and social networking initiatives in the collaborative environment. Both of these areas have high impact on current BI systems in any case, but will also provide some of the initial implementation points for Business Integrated Insight.

Conclusion

This brings us to the end of our journey from BI to a new Enterprise IT Architecture. We started among the fields of information in today's BI systems and discovered that much has changed in business and technology since I first defined the data warehouse architecture in the mid-1980s. And, perhaps surprisingly, many of the changes we saw spread far beyond the bounds of traditional business intelligence needs and approaches. So, we journeyed into the lands of business process and even dived into the somewhat murky waters of human behavior.

What we discovered, in a word, was *change*. The way business works has changed—decision making has become deeply integrated with action-taking. The way information is used has changed—from being on the edge of the business, it has become central to all aspects of it. The technology has changed—from batch to real-time, from disconnected to integrated. And the way people work has changed—from operating in a command and control paradigm to taking charge of their own information and interactions in a collaborative environment. So, as spiritual thinker and writer, Neale Donald Walsch, declares in a broader context: *"when everything changes, change everything"*[10]. Business Integrated Insight offers to do just that.

[10]Walsch, N.D.,, *"When Everything Changes, Change Everything",* EmNin Books, Ashland, OR, (2009)