

# The Open BI Appliance – New Thinking For New Opportunities

Merv Adrian, Principal, IT Market Strategy  
[www.itmarketstrategy.com](http://www.itmarketstrategy.com)

## ***Breaking the BI Logjam – a Continuing Challenge***

Year after year, CEOs and line of business executives assert they can't get information fast enough, that they don't trust it when they do, and that it doesn't give them enough of a guide to action. In a survey of over 2300 CIOs conducted by IBM in 2009, BI and Analytics ranked first among the elements selected for planning to enhance competitiveness.<sup>1</sup> This result was predictable: CIOs remain as interested in purchasing BI products as ever, even though most organizations already have some in production. And they are in place: a 2008 Forrester survey showed that 45% of firms with over 1000 employees reported having 3-5 BI products, and 16% said they had 6-9.

In fact, 2009 was the fourth successive year that BI topped Gartner's survey of CIO technology priorities<sup>2</sup>. The market passed the \$10B point in 2007 and has continued to grow in double digits even through a serious recessionary downturn in 2008-2009. Remarkably, primary research data year after year has shown 10-20% of organizations reported they were making "first-time" BI purchases in the year ahead. While mathematical logic suggests this is unlikely, the data clearly shows that internal awareness of existing solutions lags, and that BI is not used even where it is likely that products are in place. The simple truth is stark: *the promise of pervasive BI has still not been delivered.*

Nonetheless, for years, the BI market has continued to struggle to extend its mandate beyond specialists in IT. Despite its steady market growth, BI is still the province of power users, although the recent growth of analytic applications has shown one path to extended usage. Analytic applications provide ease of use driven by relevance within context, and they require no "BI setup," since they are already integrated with the data and processes they are used with. But analytic applications are tightly wedded to specific usage, and do nothing to address the more general problem of interactive, exploratory BI that can be used broadly without requiring complex setup, installation, integration and maintenance to keep all the components in step. Existing BI solutions, for all their growth, are still not in the hands of end users.

Is there a way out of this conundrum? One possible path appears to be emerging – the appliance market. Its ease of configuration and setup, and its focus on rapid delivery of value, shows promise. In the discussion that follows, we examine this notion and consider its applicability to extending BI usage.

---

<sup>1</sup> Source: IBM Global CIO Study 2009

<sup>2</sup> Source: 2009 Gartner Executive Programs CIO Survey, January 2009

## ***Market Readiness – Riding the Appliance Wave***

The IT market has shown its readiness for 'hardware appliances.' Storage appliances were a billion dollar market a decade ago, joining appliances for networking as a vehicle for rapid deployment, minimal configuration and low maintenance. Specialized appliances for security, encryption and a variety of web optimizations have become commonplace. In 2002, Netezza introduced the data warehouse appliance, and in the subsequent 7 years, grew into a nearly \$200M supplier, joined by large vendors such as IBM and Oracle, and smaller startups like Kickfire.

The appliance value proposition is simple: *an IT appliance is a purpose-built collection of components that are pre-integrated to minimize time and resources required for deployment and administration, and speed time to value.* Appliances come in many configurations: they may include server hardware, storage, and communications equipment along with the software that delivers the functionality. Hardware architecture has cooperated: the proliferation of standard racks, standard blades, standard interconnects, and standard storage form factors has made it simpler to configure systems before shipping. Vendors have learned to partner with one another around installation and support, and foundational software elements have increasingly joined in: commoditized stacks of operating system, web application server and even development tools are well established.

Why has the appliance become such an important solution? Complexity. The collection of components – the “stack” as it has come to be described – can be bewilderingly complex to assemble, integrate, and maintain. With multiple separate offerings - even when only one vendor is involved - the challenges for programmers, analysts and administrators, who strive to meet service level agreements with their constituents, while hiding complexities from them, can be daunting. The promise of an appliance is that time to value will be dramatically improved. Users will perceive results quickly, and more time can be spent on getting work done instead of on installation, on delivery instead of maintenance, on enhancement instead of mechanics.

The success of hardware appliances suggests an obvious question: *if we can quickly implement new storage, networks, websites, and data warehouses, why can't we do the same for BI?* The answer is simple: we can. The availability of commodity hardware also provides a platform for similar collections of software functionality that deliver the same benefits: pre-integration, minimal administration, and simplicity of usage. Let's look at the required elements for a BI solution in the form of a software appliance.

## ***Enter the BI Software Appliance***

The necessary components of a software BI appliance comprise a pre-integrated and optimized stack. This stack may be understood as consisting of three fundamental layers, which provide:

- **Connection** – this layer provides the operating environment for the layers above it, and reaches out to existing data sources including existing databases (including data warehouses and marts), applications, instrumented feeds, message buses, unstructured content in file systems, web pages, email repositories, news feeds, social networks – any information source that is relevant. It provides modeling and mapping tools for the import and governance of external sources. Connection management, encryption, and authorization are provided here. The latter is particularly important: in most organizations, combining multiple data sources can create a complex set of access authority challenges not accounted for in already-created internal data warehouses. Individualized access and authorization rules, including role-based concepts, are increasingly necessary. The software provider may also offer hardware, but partnering and certification with leading vendors of server hardware offers the best assurance of keeping the customer in step with the newest advances in storage, processor and network architecture.
- **Data Services** – this layer manages, transforms, and sometimes persists the data resources available. It provides data profiling, quality, modeling and governance to ensure that as more data is added, semantic models and standards evolve accordingly. As appropriate, scheduling of extraction, transformation and loading (ETL or ELT) are managed here for a data storage layer which builds and feeds the appropriate optimized data stores. This layer should leverage the underlying hardware to exploit such innovations as closely associated smart storage and in-memory technologies now available on commodity blade servers, delivering rapid, iterative processing capabilities to the interaction layer above. It also manages the metadata which describes available content in meaningful terms to facilitate exploration by business users.
- **Interaction** – this layer provides front end tools for reporting, query, analysis, modeling, data mining and scoring, dashboarding and collaboration. It may provide these within a business process context by interfacing with application environments, collaboration tools or office productivity products. The interaction layer provides delivery services ranging from portals and dashboards to broadcast and messaging services that deliver content to endpoints including smart phones and dedicated mobile devices as well as mobile computers. A powerful acceleration of time to value comes with business models that are role-appropriate, meaningful to the user's business process and industry, along with predefined performance and risk metrics, and preconfigured dashboards. These connect to metadata definitions and connections for data sources in the layers below. Connections to social and collaboration environments inside the organization, and those of partners and customers are provided here as well.

Even software appliances also need management facilities that handle such tasks as resource management, backup, and usage monitoring. These may be thought of as “adjacent to” the stack; users of the appliance will not see them unless they have specific administrative responsibilities.

Each of these layers can and should evolve continuously, and new features should become available to users with minimal reconfiguration or re-licensing in the appliance delivery model. For example, as SAP Business Objects has added new features such as tag clouds and hierarchies of self-defined “facets” to its tools, these have been made available to users other components of its stack like the BW Accelerator.

### ***How the Software Appliance Breaks the BI Logjam***

No solution can remove all complexity – the overwhelming number of possible data sources and connections will require some preliminary mapping even with the synergies described above. But the software appliance as described here creates enormous opportunities for rapid delivery of value and expansion of the user base. Its component layers attack the most challenging issues that create barriers to wide use, and shield users from the details of configuration, tool integration, version management, etc.

Beginning with the top of the stack, the interaction layer should deliver the most modern analytical constructs, perfected over several decades of BI development, combined with the rich understanding of business models and industries drawn from application expertise. These models, been hardened and refined by major vendors in the application software space, are now found in combination with BI tooling in preconfigured dashboards and KPIs. This greatly speeds initial approaches to business problems, but has hitherto been constrained by its boundaries: data that was not already defined within the model was difficult to find and combine for analytical purposes. The models themselves might or might not be available on-premise; with the emergence of the cloud as a platform for software delivery, this boundary too can fall. SAP’s delivery of models and templates through cloud-based on-demand connections is an example of this new thinking. Now advanced visualization, drillthrough, and charting, formerly confined to data that had been pre-loaded and pre-defined, can become more widely available within a richer, more appropriate context. In the today’s BI appliance, search innovations are added, permitting business users to explore their data-rich environments outside the predefined boundaries to locate other useful information and flexibly combine it with already defined elements.

The middle layer tackles the performance challenges of this scenario by combining capabilities for retrieving, filtering and processing these external data with powerful performance capabilities leveraging today’s multiprocessor architectures and larger memory spaces. By enabling more rapid, iterative analysis over a broader range of relevant data, the BI appliance supports more unfettered analysis, a broader range of “what-if” thinking on the part of business analysts who have repeatedly “hit the wall” in more constrained situations, bogged down by the time to find, retrieve, incorporate and manipulate other data. Data profiling, modeling and governance raise the level of trust in the data and facilitate continuous improvement in its accuracy. And social networking permits useful ideas to be reused – the models created, results obtained, and lessons learned can be shared via the connection to collaborative tools.

The connection layer provides the vehicle for rapid exploitation of emerging information sources as analysts experiment with them. Instrumented data streams, commercial data feeds, data from the firm’s web site usage logs – all can become a part of the wider net analysts cast to find new opportunities for topline growth, customer retention, and other business priorities.

## ***The Time for Software BI Appliances is Now***

All IT appliances have one common focus: to deliver business value rapidly, to the broadest audience, with the least setup and maintenance time. Thus, the notion of a “software appliance” is gaining currency. A BI appliance should offer a highly optimized, pre-integrated system for analyzing, modeling and delivering business answers and recommendations from data – and these are software tasks. Depending on the business objective, the right collection of pre-integrated software atop commodity hardware is all that is required to conform to our definition. Software appliances are thus an increasingly attractive way to acquire solutions to particular business problems.

Today, several vendors have assembled portfolios of software componentry that can be combined to deliver on the promise of the BI software appliance. SAP, for example, has enhanced its portfolio to drive high performance, open data analysis with its newest offering: BusinessObjects Explorer, Accelerated Version. Using SAP’s Data Services to provide ETL capabilities to access and retrieve information from any data source, this bundle lets the user combine the additional data with SAP NetWeaver Business Warehouse data via the BW Accelerator - in memory. The accelerator software provides fast indexing of large volumes of memory-resident data. The rich metadata in BusinessObjects Universes is combined with new metadata about the added information, and is made accessible to BusinessObjects Explorer technology for rapid, search-based analysis. Even individual user spreadsheets can be seamlessly combined with other data for richer exploratory analytics. SAP has added the additional data access into its bundle without creating new licensed products, taking a step towards the promise of an appliance-style deployment.

Such combinations show promise of delivering the vision of the BI appliance: as more elements already in the portfolio are bound in and integrated, ease of installation, deployment and use will continue to improve, and the BI user population will continue to grow. The logjam is on the verge of breaking up.

*This paper was funded by SAP, a client of IT Market Strategy*