

Business Intelligence

Gartner Group

Datapro Summary

Business Intelligence (BI) is the product of analyzing quantitative business data, usually business transactions; but other sources of data can be used, for example, human resources data. It provides insights that will enable business managers to make tactical decisions, as well as to establish, modify, or tune the business strategies and processes in order to gain competitive advantage, improve business operations and profitability, and generally achieve whatever goals management has set. Users of an enterprise's BI data have traditionally been inside the enterprise, but with the explosion of the Web for conducting business; i.e., e-business, an enterprise's BI users can be external to the enterprise, as well. BI tools are software programs and systems that businesses use to analyze business data and provide reports and other visualizations to users. Some BI tools are also used to develop, deploy, and manage BI applications. This Perspective report explains BI technology, analyzes the various types of BI tools, and identifies the major vendors and products in each category.

—By Alan H. Tiedrich

Gartner Group

Technology Basics

BI is a Gartner concept and core topic of research, which addresses end-user access to, and analysis of, structured business data and information which is quantitative in nature (unstructured information contained in text documents could be considered part of BI, but this type of information can't be analyzed by what we have defined as BI tools). A term coined by Gartner in the late 1980s, BI is a user-centered process that includes accessing and exploring information, analyzing this information, and developing insights and understanding, which leads to improved and informed decision making. This involves an iterative process of accessing data (ideally stored in the data warehouse, data mart, or operational data store, but not necessarily) and analyzing it--thereby deriving insights, drawing conclusions, and communicating findings--to effect change positively within the enterprise. BI usage crosses the spectrum of users throughout the enterprise and includes rank-and-file workers, executives, analysts, and knowledge workers.

At one time, organizations depended on their Information Systems (IS) departments to provide them with both standard and customized reports. This goes back to the days of mainframes and minicomputers, when most end users did not have direct access to the computers. This began to change in the 1970s, when online host-based systems came into vogue. Even then, these systems were used primarily for entering business transactions, and reporting capabilities were primarily a limited number of predefined reports. IS typically was overburdened, and users had to wait for days or weeks to get their reports, if they needed reports other than the standard ones that were available. Eventually, Executive Information Systems (EISs), which were attuned to the decision support needs of executives and managers, were developed. With the advent of the PC, and

particularly networked PCs, basic BI tools gave users the technology to create their own basic routine and custom reports.

Types of BI Products

Today's BI tools categories include enterprise BI Suites (EBIS), query and reporting tools, advanced BI tools--primarily On-Line Analytical Processing (OLAP)/advanced analytic tools, and BI platforms for developing BI applications. The BI tools are used by end users to access, analyze, and report against data, which most frequently resides in data warehouses, data marts, or operational data stores. BI applications are developed by using BI development platforms, but these applications are not considered to be BI tools. An example of BI application is an executive information system (EIS).

Most modern BI tools to fall into two categories: enterprise BI suites and BI. Basic query and reporting tools have largely been absorbed into and superseded by the EBISs. Multidimensional OLAP engines, as well as relational OLAP engines, can be used as BI tools and also are the underlying infrastructure for BI platforms.

Data mining refers to a process rather than a technology, with the goal of discovering new correlations, trends, patterns, relationships, and categories. This is accomplished by sifting through large amounts of data stored in repositories, using pattern recognition technologies as well as statistical and mathematical techniques. Data mining iteratively applies different operations or transformations (e.g., feature selection, stratification, subsampling, clustering, visualizations, and regressions) to raw data, with two objectives:

- Finding representations that are particularly insightful to humans that result in a better understanding of the underlying business processes.
- Finding models that can forecast the outcome of situations or the value of given situations using historical or subjective data.

Unlike using BI tools, data mining is far less user-directed and instead relies upon specialized algorithms that correlate information and assist in discerning important (and otherwise unknown) trends, unguided by user bias and assumptions.

Query and Reporting Tools

Query and reporting tools--typically desktop tools--enable users to access databases, usually networked relational databases, but also multidimensional databases and local databases; to then do some basic analysis and produce reports, which could be displayed and/or printed. While the querying can be ad hoc, it also can be of a scheduled nature. There are reporting systems, typically server-based, that support scheduled querying and reporting. Some of these systems are geared to support "production reporting," which is typically an environment that produces reports from major enterprise systems for widespread distribution. These are successors to the "green-bar" reporting capabilities of mainframe and minicomputer systems that preceded the PCs. Desktop query and reporting tools have also been enhanced to provide some light OLAP.

EBISs

EBISs--a natural path for the disjointed query, reporting, and OLAP offerings--meet criteria including product scalability, usability, and manageability. They are integrated suites of the query, reporting, and OLAP tools. EBISs should have extensive scalability and extend not only to internal users but also to key customers, suppliers, and the general public. The products should also support a variety of users by providing extreme ease-of-use and requiring minimal training. EBIS products should aid administrators in the deployment and management of BI functionality, without adding significant new management resources. Because of strong Web affinity associated with EBISs, some vendors have described their EBISs as BI or Web "portals." These portal offerings typically provide a subset of EBIS functionality via a Web browser.

OLAP/Advanced Analytic Tools

OLAP tools are server-based analysis tools that originally were based on multidimensional databases (MDDBs), but today can be based on relational data stores, most often overlaid by an indexing scheme that simulates an MDDB. MDDBs are databases constructed specifically to support analysis of quantitative data, along multiple dimensions. These databases hold this "multidimensional" data in a "pure" multidimensional form. OLAP technology enables users to organize the data in a hierarchical fashion in multiple hierarchies (for the same set of data). Most applications involve a time dimension, so that data can be analyzed over time to see trends. Other dimensions might involve geography, organizational unit, customer, product, and others. Data is aggregated, sometimes in the database and sometimes on the fly, in order to produce the desired views and reports. So, for example, data could be aggregated to monthly fiscal periods, and then accumulated to a higher level, usually quarterly and/or annual. MDDBs are optimized for multidimensional analysis and come with analytic functionality, providing good performance, but typically require a lot of time to load and expand the size of the source database many fold. Because MDDBs typically contain largely aggregated data, they come with a "reach-through" capability that enables access to the detailed data in the RDBMS that contains the cube's source data.

Several years back, some vendors began using RDBMSs to store multidimensional data in relational tables and new object types that support multidimensional analysis. The ROLAP database has the advantage of scalability and flexibility, but (typically) lacks the performance of MOLAP, although there are performance-enhancing techniques, like a star schema design. Although the MDDBs are still the most common ones for OLAP processing, more recently, this functionality is being built into or extended onto RDBMSs (not the same as storing the multidimensional data in relational tables). This has led to the terms MOLAP, ROLAP, and HOLAP (for hybrid products that can store native multidimensional data and also store data relationally). A "cube" is the conceptual model of measures that share dimensions and the hierarchies in these dimensions. MDDBs are accessed via APIs for multidimensional querying, while RDBMSs are accessed via SQL queries. SQL is limited as a multidimensional language, which ROLAP overcomes to some degree by the use of special engines that generate the SQL, which is more sophisticated than that which could be manually coded.

OLAP tools allow the user to explore the data in all the various dimensions inherent and built into the database (some allow ad hoc dimensions to be calculated on the fly). Users can select which measures to analyze and which dimensions to display vertically and horizontally on a cross-tab

chart. They can then "slice and dice" to focus on a particular combination of dimensions. By supporting a "drill down" capability, OLAP tools allow users to dig down a dimensional hierarchy and explore increasingly detailed levels of data. It also is possible to "drill across" through other dimensions. Another type of data manipulation in OLAP is "pivoting," or swapping rows and columns.

Desktop OLAP tools, now incorporated into EBISs, enable end users to view and manipulate multidimensional data, which might come from server-based ROLAP or MOLAP data stores. These tools have the ability to download cubes, so that they can function standalone when disconnected from networks. As part of EBISs, these desktop tools have been equipped with server-based processing capabilities, which provide a number of capabilities beyond their desktop capabilities, but do not rival the MDDB-driven MOLAP tools. These have far greater performance and analytical robustness.

BI Platforms

BI platforms offer complete sets of tools for the creation, deployment, support, and maintenance of BI applications. These are data-rich applications, with custom end-user interfaces, organized around specific business problems, with targeted analyses and models. BI platforms, although not as rapidly growing or widely used as EBISs, are also an important segment due to expected growth in BI applications. (BI platforms provide the environments and tools upon which BI application packages are built.) Due to RDBMS vendors building OLAP into their DBMSs, many platform vendors that offered multidimensional DBMSs for OLAP have been forced to migrate to a BI applications business, in order to survive. However, due to the RDBMSs providing BI functionality, these are actually encouraging the growth of the BI platforms market. Part of this is due to the greater ability to execute on the part of the DBMS vendors. If we consider the dimensions of value (ability to help an enterprise by creating unique applications that sets the enterprise apart from its competitors) vs. functional completeness (the higher the degree of functional completeness, the more well-suited a tool is for end users) for the various BI tools, we see that EBISs are highly functional but do not have the higher value of BI platforms, or of ready-made and custom BI applications. BI platforms are high value but typically not as functionally complete as EBISs.

BI Tools Trends

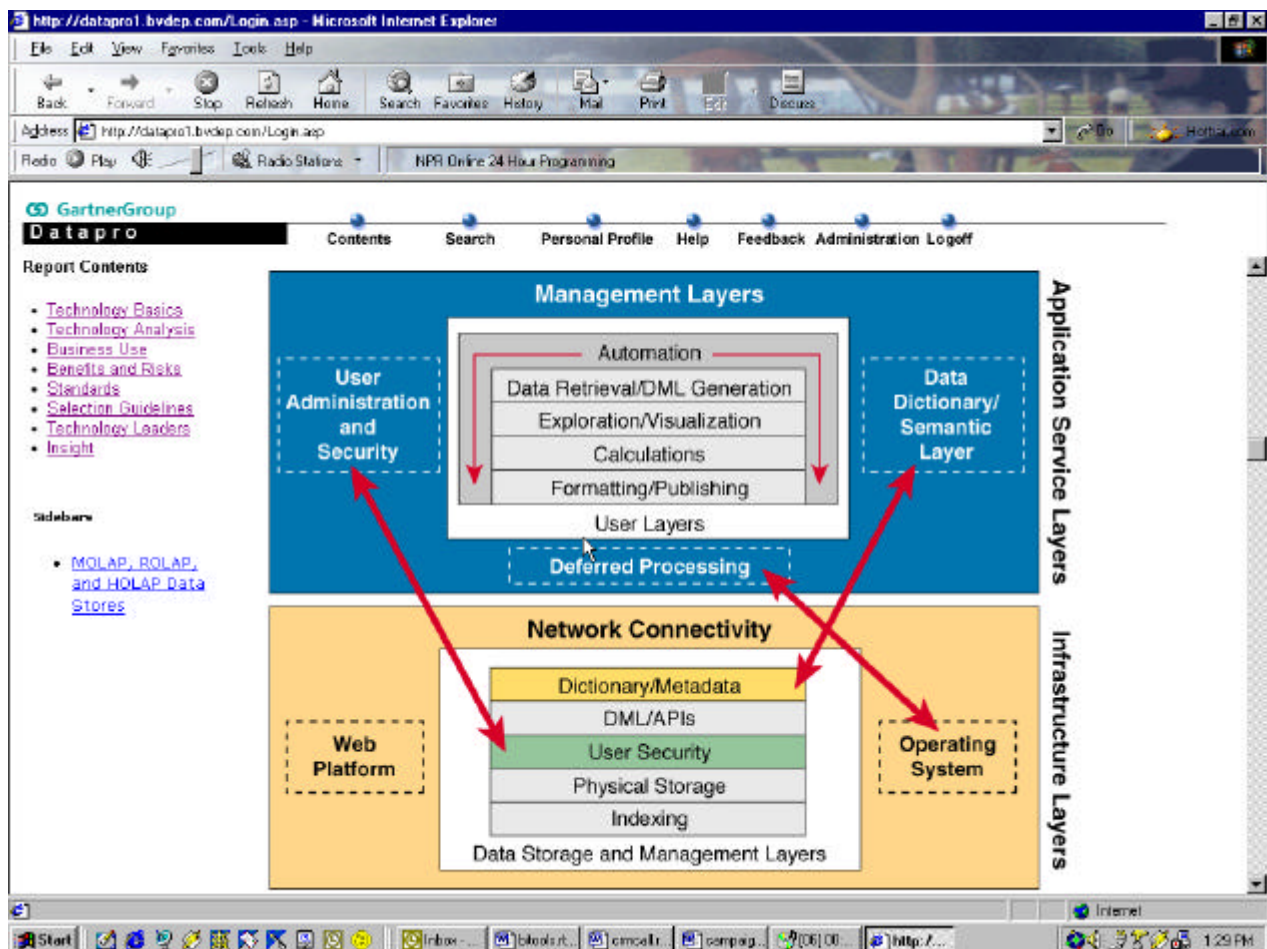
The BI market continues to evolve. The fastest growing category of BI tools is EBISs, reflecting heightened competition in the new economy, requiring better and faster decisions throughout an enterprise. The use of query, analysis, and reporting tools is declining as organizations replace or upgrade these with EBISs. Still, basic query and reporting tools remain the most ubiquitous BI tools, because they still satisfy most of the users' needs. These technologies are intended to support the masses of users requiring ad hoc database query, reporting, and basic OLAP analysis. The use of OLAP and other advanced BI tools, like data mining tools, is also growing. The vendors sometimes identify the EBIS products as BI portals, because the Web-enabled versions of these products provide a Web entry point into an enterprise's information. In fact, often, these BI portals also provide support for linkages to unstructured information too, although this typically requires some system integration to be done. Increasingly, EBIS products are more focused upon external constituents of an enterprise (e-business intelligence).

BI Architecture

An enterprise's BI architecture should be developed after the users' BI requirements have been determined and before BI tools have been selected. The BI architecture will have two major components, a BI information delivery architecture and a BI technology architecture. After determining BI information usage patterns, the information delivery architecture can be designed, based on these usage patterns and also on the type of deployment that is required. This can be any mix of desktop with network connectivity, desktop plus server-based, thin client Web-based, wireless devices, and mobile computing. The information delivery architecture will define the user interfaces, which more and more, are portals, which have personalization capabilities, so that they can be customized for each user or for user groups.

The BI technology architecture defines the infrastructure and components needed to support deployment, execution and management of the BI tools and applications, and the relationships of these components. A sound BI technology architecture will consist of two significant layers: infrastructure and application services. The infrastructure layers include data storage and management, and network connectivity. Application services include all of the BI services, such as query, analysis, and reporting and visualization engines, as well as security and metadata.

BI Technology Architecture



Source: Gartner.

BI Data Stores

An increasing number of organizations are using the SAP Business Information Warehouse (BW) as the primary BI data store--SAP BW, a relatively new product is benefiting from the ubiquitous use of SAP's R/3. The use of ROLAP data stores (relational databases) as the primary BI data stores is also growing rapidly, stemming from RDBMS's superior suitability for certain applications involving very large databases of detailed transactions and also due to vendors including OLAP functionality in their RDBMSs. Multidimensional data stores' use remains constant and the most prevalent, probably because such data stores and their tightly linked OLAP tools provide superior performance and functionality for certain types of applications, particularly where the use of aggregates is a major characteristic of the data. OLAP tools also have complex analytic calculations, making these tools optimal for certain applications.

BI Data Access

The growth of the Web to access BI applications is occurring, not surprisingly, at the expense of client/server computing. Client/server client-based processing is decreasing more drastically than client/server server-based, two-tier processing, reflecting that connectivity to corporate BI data is an important element of BI access, and nonconnected PCs will not be functional enough. E-mail distribution of BI reports is growing but will remain a minor factor, and mobile computing remains a miniscule BI platform because most stationary and mobile end users use the Web for BI access.

Metadata

Most BI tools that are marketed today use a metadata layer or metadata repository. Business metadata includes definitions of the data that is stored in the data sources, in business terms, so that users can use business terminology they are familiar with in creating queries and reports. It also contains business rules and calculations that have been defined for the business. There is also technical metadata, which enables the tools to actually access the physical data. RDBMSs also employ metadata. Metadata also is used by the data transformation tools that are used to build the data warehouses and data marts. When data warehouses and data marts are being constructed, automated metadata capture from the data sources is often possible, but sometimes it is forgone, leaving users to their own devices to capture the metadata into the transformation tool's repository. So, it can be a complicated situation, with several metadata repositories in existence in one organization. A lack of common metadata for tools--there are no standards for metadata--presents a significant IS challenge.

Technology Analysis

Business Use

The proper use of information is vital to the effectiveness and ultimate success of the enterprise. Reflecting a democratization of technology use and the emergence of information as a key business asset, the use of BI tools is growing most rapidly among administration and operations personnel. This trend indicates how pervasive technology has become in enterprises. No longer is such technology restricted to professionals or Information Technology (IT) experts; it has become

an everyday standard tool. Instead of just being used for transaction processing, IT's use as a customer service and revenue-enhancing tool is making it indispensable for performing "ordinary" previously clerical functions. This is elevating lower-level personnel into jobs that involve making decisions.

The majority of BI users need only basic information, and their jobs have a relatively small BI component (perhaps 10 percent). The most basic information is historical reporting, which is very generic; i.e., not specific to any user's needs, a common denominator type of information. It may consist largely of static published reports and could include OLAP views. Knowledge workers have a much higher BI content in their jobs, perhaps 30 percent, and business analysts, on the order of 80 percent. Such needs may be satisfied partly by parameter-driven reports, or at the highest level, by ad hoc queries and OLAP.

Business Forces Driving BI Tools Use

Perhaps it can be attributed to the explosion of personal computing and most recently to the Internet, but the pace of business has increased dramatically over the last several years. The term "Internet speed" is used to denote the pace at which businesses must act to survive and prosper. Executives, managers, and other high-level technology users have recognized that decisions must be made quickly and intelligently and must be based on accurate and timely data. Increasing revenue and improving customer service are also very important goals for enterprises, which are looking for ways to increase revenue not only from current lines of business but also by identifying new business opportunities and new business models. Although revenue growth is a fairly standard business goal, revenue growth and market share have become even more important since today's business strategy is to grow the business as large and as fast as possible, and to expand globally. The mergers of huge companies into mega-companies is accelerating this trend. How to serve customers better is an important goal, because businesses have recognized that acquiring new customers is typically more expensive than retaining customers. Making businesses more efficient and decreasing business costs also are important, reflecting the importance of optimizing the bottom line.

Rapid growth of BI technology use by executives demonstrates how aware they have become of the impact of technology on business and how important BI tools are to executive decision making. Executives and managers in our restructured companies cannot rely on staff (which they may not even have in our streamlined modern organizations) to assist them in gathering and analyzing information, but must accomplish these tasks personally. In addition, the rapid pace of today's commerce requires faster decision-making; executives cannot afford to wait for subordinates to "get back to them" with the required analysis. Lastly, decision making is an iterative process, which requires a "hands-on" approach. Fortunately, BI tools have been made friendlier so that personnel at all levels can use them with facility.

The most important driving force of business today, however, is the Internet, which is personified by e-business. Thus, we see business-to-business Web sites, business-to-consumer Web sites, marketplaces--all ways of doing business electronically. As integral components of these Web sites, analytic functionality (e-business intelligence) is vitally important to analyze the behavior of transactors on these Web sites, particularly their buying behavior. This BI information can be feedback into the systems powering the Web site in a closed-loop fashion, with the information being used automatically to modify the way the system handles situations or customers. A good

example of such a system is a customer relationship management (CRM) system; such systems are being endowed with intrinsic BI functionality. The second major use for BI in e-business is to support the extended enterprise--customers, suppliers, and other business partners. In this case, members of the extended enterprise have access to BI information, including not only that coming from their own business transactions, but also information based on transactions done by other entities (on a summary and anonymous basis). Thus, there is a rapidly growing need for extra-enterprise access to BI data, by the enterprise's customers, suppliers, and others.

Query, Reporting Tools, & EBISs

With the variety and number of BI users, enterprise-wide access to corporate BI data is a high priority. Today, in spite of the availability of advanced analytic tools and applications, basic query, analysis, and reporting tools are still the most sought after BI tools. With the democratization of information, the number of BI users is increasing dramatically at executive and operations personnel levels. A great deal of the need is for these users, who will be satisfied by basic query, analysis, and reporting. EBISs are the successors to the basic query and reporting tools and are supplanting or extending them, to provide support for varying levels of users, with a variety of query, reporting, and OLAP capabilities, with minimal training.

OLAP Tools

OLAP is used primarily for analyzing aggregated/derived multidimensional data. It supports sales analysis and forecasting, financial budgeting and forecasting, risk analysis, trend analysis and a variety of other applications. OLAP and other advanced analytic capabilities carry a high priority. These BI tools, while used by a smaller user population than EBISs or basic query and reporting tools, have the possibility of returning higher value in that they can help organizations spot trends that alert them to problems or opportunities.

Data Mining

Data mining is used for many different types of applications, like fraud detection, churn management; forecasting customer behavior for targeted marketing or cross-selling, i.e., the likelihood of a future event (e.g., response to an offer, churn, failure); affinity modeling, i.e., what other products people bought and who bought a particular product; the detection of outliers; link analysis to define the relationships between events and other occurrences, e.g., for looking at all kinds of fraudulent behavior (e.g., front-running, money laundering, insurance fraud, and e-auctions)--in industries as diverse as insurance and telecommunications, basically anywhere. The two most popular types of data mining are clustering and decision trees; other common ones include neural networks and regression. The latest opportunity for data mining is in mining click-stream data, initially focusing on operational issues (e.g., peak hours and the top 10 pages hit), but future applications will include profiling together with segmentation and more advanced trend analysis, perhaps followed by more advanced personalization and CRM integration (e.g., churn and complaints), and then CRM aspects will be tied back to operational aspects, inventory management, and dynamic pricing.

BI Platforms

The older high-end BI platforms dominated the market when end-to-end environments based upon 4GLs were the norm. These were targeted at IS organizations that internally developed user applications. Today, many IS (and end-user) organizations prefer packaged applications over internally developed solutions. Hence, BI platforms will be used increasingly by third-party application developers (e.g., VARs as a mechanism to deliver "best practice" BI applications) directly to end users.

Benefits and Risks

The benefits of using BI tools are explored in detail in the Business Uses section, so these will not be repeated here. Suffice to say that the appropriate use of BI tools will mean the difference between the life and death of many enterprises; between stagnation and growth; between lackluster results and outstanding financial performance; between excellent, personalized customer service and impersonal, shoddy service; and between optimizing the relationships with parties outside of the enterprise and losing the possible benefit of working with suppliers and others as business partners. BI is just that important!

As to risks, the risks are not so much in the technology, as they are in properly assessing the enterprise's true BI needs and then in selecting the most appropriate vendors and products. The major technological risk is that the technology is changing so rapidly. Vendors that created their products before the Web became *avante garde* now find themselves with legacy software that has to be rationalized with new Web-enabled versions of their products. Newer products built for the Web are from vendors that are newer, and there is a trade-off between getting the latest technology and getting the most stable vendors. E-business has great implications for BI tools, and vendors are still building e-business intelligence capabilities into their products and integrating BI data and results back into the e-business applications. The latest technological shift to affect BI is the advent of wireless applications. BI tools vendors have adding capabilities to allow BI applications to be accessed via wireless technology. Naturally, these new technologies bear some risk until they are well proven. Lastly, many of the BI tools vendors are small companies, or medium sized, apart from a few large companies and the RDBMS vendors. Therefore, one of the most important risks is the vendor's ability to execute, so this is something to consider.

Some of the biggest risks related to the use of BI tools are based on data. One of the biggest risks of using BI tools is related to data quality; i.e., that the data being used has not been properly cleansed. Because business domains within an enterprise often choose their own BI tools, an enterprise may end up with multiple BI tools in use, as well as multiple data marts with data that may not be defined in a common way or with metadata that is not compatible. This can lead to different domains drawing different conclusions about the same data. If data warehouse architecture for the enterprise is carefully planned, much of this can be alleviated.

Standards

OLE DB for OLAP

Microsoft OLE DB for OLAP is a set of objects and interfaces that extends the ability of OLE DB to provide access to multidimensional data stores and enables users to perform sophisticated data analysis through fast, consistent, interactive access to a variety of possible views of the underlying information. OLE DB for OLAP enables developing and accessing multidimensional data providers (MDPs), which present data in multidimensional views, and tabular data providers (TDPs), which present data in tabular views. It is possible for a single data source object to support both tabular and multidimensional presentations. OLE DB for OLAP Component Object Model (COM) interfaces to create client applications. OLE DB for OLAP allows independent software vendors and corporate application developers to depend on a single interface for accessing multidimensional data, regardless of the vendor or source. With this technology, OLAP applications can uniformly access both relational and nonrelational data stored in diverse information sources, regardless of location or type.

OLE DB for Data Mining

This specification addresses data mining-specific provisions of OLE DB support, which was enhanced with the OLE DB for Data Mining specification to provide integration of third-party data mining algorithms.

Selection Guidelines

The selection of BI tools should be done only after an analysis of the business and enterprise needs have been performed and a BI technology architecture has been designed. It is likely that a single BI tool will not support the needs of all users. Therefore, several tools may need to be selected to perform various BI functions. EBISs include a breadth of BI functions required by end users, so selection of an EBIS may satisfy some or all of the BI needs. If BI applications need to be developed, then a BI platform may need to be selected. It is possible that both an EBIS and a BI platform may be needed. Other needs might be satisfied by selecting an OLAP server or a data mining tool. Again, the types of tools needed will be determined only after an analysis of the users' BI needs. As a general statement, Gartner has found that clients are looking for several things when selecting BI tools, so these should be considered in addition to other specific criteria defined below for the two major BI tool categories (EBIS and BI Platforms): Internet accessibility, end-user ease of use, and long-term company viability.

EBISs

Gartner criteria for evaluating EBISs include scalability, usability, and manageability:

Scalability

An EBIS needs to extend beyond the enterprise, offering basic BI functionality. Its architecture must be designed to distribute processing flexibly without undue complexity or placing an unreasonable burden upon IS departments. Attributes should include distributed architecture; load balancing and failover support; and extranet, intranet, and Internet support.

Usability

An EBIS must support a wide range of users (i.e., with different needs, styles, and levels of sophistication). It must help keep user training to a minimum with straightforward and consistent interfaces. The main attributes are support for multiple user styles, common user objects, and interchangeable and interoperable components.

Manageability

An EBIS must support strong management/administration facilities that aid IS personnel in installation and ongoing support. Since resources will continue to be scarce, and EBIS implementations are likely to be large, these facilities must introduce a new level of automation and abstraction to reduce the cost of deployment. The facilities must include common infrastructure, security, scheduled processing, and rich administration.

BI Platforms

Gartner has defined the characteristics of a BI Platform to include a modular, distributed architecture, supporting relevant standards like XML, OLE DB for OLAP, and providing total Web deployment. It must be open and extensible, so that third parties are encouraged to and can easily add functionality. The vendor must provide a strong third-party support program. This will assist the vendor in building a business of sufficient size, which is essential to the success of BI Platform vendors. Of course, the product must have comprehensive BI functionality.

Modern Platform Architecture

The BI platform must provide support for modularity (e.g., components) and an accepted distributed computing model. Platforms should support de facto and de jure BI and other standards--e.g., LDAP, Extensible Markup Language (XML), OLE DB for online analytical processing (OLAP), Common Object Request Broker Architecture (CORBA), Common Object Model/Distributed Component Object Model (COM/DCOM)--and total Web deployment.

Third-Party Extensibility

The BI platform must provide the ability to add significant new functionality without fundamentally altering the core BI platform with a useful and well-documented set of application programming interfaces (APIs) and support the dominant applications development paradigms (e.g., Java).

Vendor Support Programs

The BI platform vendor must support the creation, promotion and advocacy of third-party extensions and product offerings, when possible, as part of an overarching third-party support program.

Critical Market Expansiveness

The BI platform must be widely used to offer a "safe choice" for application developers--e.g., value-added resellers (VARs), independent software vendors (ISVs)--as an established and stable vendor/product.

BI Features

The BI platform must provide BI-specific functionality such as database access capabilities (e.g., SQL), multidimensional (OLAP) data manipulation, modeling functions (e.g., what-if analysis), statistical analysis (e.g., ANOVA), and graphical presentation of results (e.g., charting).

OLAP Servers

Of MOLAP- and ROLAP-based tools, the MOLAP-based tools offer the most analytic functionality and the best query performance, provided aggregated data is suitable for the organization's requirements, as opposed to large volumes of detailed records. They offer good scalability of numbers of users and somewhat less, but still good, data scalability. ROLAP tools typically provide the best data scalability, and they are a better solution than MOLAP when large volumes of detailed records must be accessed frequently. Scalability of numbers of users is only fair at best, and query performance and analytic functionality are fair. HOLAP tools provide the best of both worlds in a single solution, but these do not always offer the full-blown capabilities of both ROLAP and MOLAP. They also place on IS the requirement of administering both a relational database environment and a multidimensional database environment. A HOLAP environment, while offering flexibility, also requires decisions to be made about which data elements should be placed in the relational database and which should be placed in the multidimensional construct. Help in the form of better administrative, monitoring, and tuning tools are likely to come from RDBMS vendors who have added or are adding MOLAP support. Desktop tools can provide excellent scalability and query performance, but this comes at the expense of analytical functionality (which is lacking) and data scalability.

Technology Leaders

[Business Objects, Inc. BusinessObjects 2000](#)

Business Objects for a long time, and even now, has been one of the two major independent tool competitors in the EBIS market. BusinessObjects 2000 is a comprehensive EBIS, providing client/server and Web client functionality that are very similar to one another in two companion products--BusinessObjects and WebIntelligence. Business Objects' strategy is targeted towards e-Business Intelligence, or the marriage of the Internet with BI. Although BusinessObjects has been Business Objects' flagship product and the foundation of the company, WebIntelligence is where Business Objects charts the future of the company. BusinessObjects 2000 includes the key products BusinessObjects InfoView 5.1 (the BI portal to both BusinessObjects and WebIntelligence), WebIntelligence 2.6 (thin client Web-based with a distributed architecture), BusinessObjects 5.1 (full-client desktop and Web-enabled), BusinessObjects Broadcast Agent 5.1 (optional), and BusinessObjects Set Analyzer 1.2 (optional). WebIntelligence Wireless Edition supports wireless access via WAP-enabled devices to mobile users. Business Objects has set up a subsidiary (Ithena) to focus on BI applications.

[Information Builders, Inc. WebFOCUS Business Intelligence Suite](#)

Information Builders has been a leading provider of query, analysis, and report writing technology for many years. The scope of the WebFOCUS product line places it in both the EBIS and BI Platform categories, giving it dual status. Information Builders' WebFOCUS Business

Intelligence Suite is a scalable, integrated EBIS designed for enterprise-level query, analysis, and reporting over heterogeneous platform environments to clients in any location. The WebFOCUS BI Suite includes WebFOCUS Reporting Server for reporting and analysis; a WebFOCUS Maintain Server for building client/server and Web applications; the WebFOCUS Report Broker to manage reporting and distribution; ERP modules for J.D. Edwards, SAP, PeopleSoft, etc.; and WebFOCUS InfoCube, a library of business intelligence templates for comprehensive analysis of sales, marketing, profitability, and more. WebFOCUS provides multidimensional analysis of data, with change of sorting criterion, multiple reporting, graphing, drill-down, user configurable reporting formats; on-demand paging, with hyperlinks to pages of other reports; and the capability to manage output effectively and be distributed via the intranet, e-mail, networked printers, or to wireless devices. The application set leverages data resources from repositories whether they are relational databases or ERP packages such as SAP or PeopleSoft, or whether they are stored in mainframe systems such as IMS or CICS. The system provides a development toolset that is fully integrated with Microsoft Office 2000; BackOffice 2000; and Microsoft OLAP Services for dynamically generating information using cube data; Excel pivot tables; and many Office documents in a variety of formats, including XML, HTML, CGI, Excel, PDF, and Excel 2000 for display in Web browsers or for use with desktop applications. The WebFOCUS system is scalable from a wireless palm-top to desktop, to workstation, to mainframe and can run on NT, Linux, Unix, OpenVMS, OS/390, MVS, and CMS.

[Microsoft Corp. SQL Server 2000 Analysis Services](#)

Microsoft SQL Server 2000 Analysis Services is a robust BI Platform that runs on Microsoft Windows NT platforms only. It can access other relational databases through OLE DB support and allows front-end tools to access it through OLE DB for OLAP. While Analysis Services is part of the SQL Server 2000 product, Analysis Services is a middle-tier OLAP server that functions independently of the SQL Server RDBMS. In addition to providing analytic functionality, Analysis Services also has data mining capabilities, albeit simple ones, coming with two data mining models: a decision tree model and a clustering model. Microsoft Analysis Services was designed with a number of scalability- and performance-enhancing features, including cube data compression (Unlike competitors' cubes, which explode the source data into much larger OLAP data sets, Analysis Services compresses the data into smaller OLAP cubes.), partial data aggregation with on-the-fly aggregation, and partitioning, to name a few of the most significant. In addition, it has extensive support for parallel processing and load distribution, through cube partitioning, parallel querying against partitioned cubes, as well as wizards to help determine how to trade off storage considerations for performance. Analysis Services supports MOLAP, ROLAP, and HOLAP; write-back, which is essential for budgeting and forecasting applications; and even has a realtime cube update option. Its extensive functionality can be further enhanced and automated by custom development. Although Microsoft just entered this market at the end of 1998, it has established a strong presence and position in this market.

[MicroStrategy Inc. MicroStrategy 7](#)

MicroStrategy has proven a tough competitor for data warehousing DW/BI scenarios, especially where leveraging traditional RDBMS strengths for scalability, flexibility, and administrative controls are important. MicroStrategy Inc.'s MicroStrategy 7 is the culmination of a rearchitecting of the product into a componentized multitier environment. It has a number of significant improvements in functionality, ease of use, deployability, performance, and application

development. In addition to its long-time classification as an EBIS, Gartner recently has classed MicroStrategy in the BI Platforms category, giving it dual status. Using a three-tier architecture, it is a suite of integrated components, which distributes processing among multiple servers. MicroStrategy Web is MicroStrategy's client interface that provides OLAP capabilities via a Web browser, while MicroStrategy Agent provides these capabilities in a Windows environment. MicroStrategy 7 has an intrinsic relational database foundation, which leverages the functions and capabilities inherent in RDBMSs. In MicroStrategy 7, mathematical, statistical, and financial functions either can be performed in the RDBMS or in the MicroStrategy 7 analytical engine. MicroStrategy 7 cannot access OLAP servers; e.g., Hyperion Essbase. MicroStrategy 7 is targeting the wireless BI applications market. MicroStrategy doesn't offer analytic applications, but relies on partnerships for this (MicroStrategy has an applications group, which right now focuses solely on eCRM applications. There is a product available now, eCRM6, which is built with the MSTR 6 platform. MicroStrategy is currently developing eCRM7 on top of the 7 platform.).

Oracle Corp. Express

Oracle Corp. offers Express' integration of relational and multidimensional data and support for both Unix and Windows NT platforms, combined with an Oracle Web strategy, data warehousing strategy, and its reputation as a leading RDBMS vendor. Oracle Corp.'s Express products constitute a BI Platform and support analytical business processes, such as modeling, forecasting, statistical analysis, sales and distribution analysis, and financial analysis. Express Server, an OLAP server, is a key component of Oracle's business intelligence (BI) tools, which also include Oracle Express Objects, an object-oriented graphical tool for developing custom BI applications; Oracle Discoverer for performing ad hoc queries; and Oracle Reports for reporting. Express Server operates on both client/server and Web platforms. Express Server is a HOLAP tool, able to support both RDBMSs and/or its own MDDB. Express client tools include Express Objects and Express Analyzer, an analytical tool that allows end users to run and extend Express Object applications. Express Server is also the underlying technology enabling Oracle's packaged BI applications: Oracle Financial Analyzer, Oracle Sales Analyzer, and Oracle Demand Planning. Oracle has ambitious plans for integrating Express technology with the Oracle RDBMS in the forthcoming Oracle 9i and also for improving its application development support with BI Beans and Java OLAP API--Java object-oriented API.

The SAS Institute SAS System and Solutions

The SAS System is a comprehensive data collection and storage system, with extensive data analysis and application development capabilities. SAS Software Solutions, including over 50 software tools, support all the activities of data warehousing and business intelligence processes. The SAS BI/OLAP solution, which constitutes a BI Platform, is based on The SAS System and the SAS Software Solutions. The foundation of SAS Software Solutions, The SAS System, is an integrated suite of information delivery tools, which enables companies to transform enterprise-wide data into business intelligence. Base SAS, the foundation of the SAS System, provides an application development environment for data access, management, analysis, and presentation. The SAS BI/OLAP solution creates a data warehouse of information. Access to this data, which can be stored as MOLAP or ROLAP, is controlled by the SAS Data Warehouse metadata. SAS Software Solutions offer a wide range of query, analysis, and reporting capabilities for the data warehouse, including EISs, data mining, neural networks, data visualization and discovery,

statistical analysis, econometric time series forecasting, operational research, quality control, geographic information systems, and market research. The SAS data mining solution enables users to apply exploratory statistical and visualization techniques, select and transform the most significant predictive variables, model the variables to predict outcomes, and confirm a model's accuracy. The SAS Solution for e-Intelligence is a comprehensive suite of data integration, analysis, and reporting applications that integrates data from all channels--including e-commerce systems. SAS e-Intelligence Solutions extract data from Web-based applications and organize it into the data warehouse, so the data can be analyzed along with traditional corporate data using SAS analytic tools.

Insight

For many years, IT was used primarily to process business transactions and to provide reports on historical information. While this information was helpful, it was too aged to help with making decisions that required more timeliness and better analysis tools. BI tools have come into their own in the last few years, and the rapid pace and complexity of business today, together with the availability of IT to support BI processes, has both enabled and forced enterprises into leveraging the business data that is captured through business transactions processing. The growth and availability of relational databases and data warehousing began to support a shift into converting data into business intelligence. The concurrent growth of multidimensional databases and OLAP technology also furthered the ability to create business intelligence information. BI tools have evolved from rudimentary query and reporting tools into enterprise BI suites and BI platforms. Most recently, the growth of the Web for e-business and the business trends of customer relationship management and supply chain planning have extended the enterprise to include suppliers, customers, and other business partners. All of these parties need and are being provided with access--via the Web--to BI being assembled by the enterprise. Fortunately, great advances in BI technology now offer a selection of alternative BI tools that can handle many of the needs and are being constantly enhanced with additional capabilities. The biggest challenges for enterprises are to define their BI goals, their user requirements, a suitable BI technology architecture, and then to select the best tools or applications and implement them.