

Data Integration: The Key to Effective Decisions



Author: MAS Strategies

Contributors: Cortney Claiborne, Darren Cunningham, Davythe Dicochea, Erin O'Malley, Philip On

Audience: Report developers, data warehouse managers, IT director, CIO

Contents

Executive Summary	ii
Introduction	1
Why Is Data Integration So Important?	2
The Value of Data Warehousing	3
Warning Signs: Does Your Organization Suffer from Poor Data Integration?	4
The Benefits of Data Integration: Combine “Islands of Information” to See the Big Picture	5
Approaches to Data Integration	8
Conclusion	12
Appendix	13
About MAS Strategies	14
About Business Objects	15

Executive Summary

To make sound decisions, an organization must establish a solid data foundation. This foundation must combine both current and historical data values from operational systems in order to identify trends and predict future outcomes. Data integration technology is the key to pulling this data together and delivering an information infrastructure that will meet strategic business intelligence (BI) initiatives. This information infrastructure consist of data warehouses, data marts, and operational data stores. The data integration solution can be developed either internally or acquired from a vendor.

While a homegrown data integration system frequently yields a quick and dirty solution that initially appears inexpensive, any upfront savings are often quickly lost as demands on resources and personnel begin to change. Vendor-supported packaged solutions, on the other hand, have withstood the test of time. They include capabilities such as metadata integration, ongoing updates and maintenance, access to a wider variety of data sources, and design and debugging options rarely offered by in-house solutions. Also important to note is the fact that most homegrown data integration solutions are almost never integrated with an organization's BI tools. Such integration is, however, available from commercial offerings either by adherence to industry standards and/or through integration with the BI tools in the data integration vendor's total product portfolio.

This paper will discuss the importance of data integration and help you identify the key challenges of integrating data. It will also provide you with an overview of datawarehousing, as well as summarize the benefits and approaches to integrating data.

Introduction

Imagine that you work with one of your organization's mission-critical operational systems. Your organization considers you the go-to person for any query or reporting request associated with this system.

What if your CEO were to ask you to modify one of your year-end reports to compare this year's numbers to those from the two previous years? Unfortunately, the operational system you are reporting off of only stores current-year detail records and prior-year summary balances. The summary balances from two years ago were purged from the system at the beginning of this year. When you try to explain this to the CEO, all you hear is, "So how long will it take to find the lost data?"

In another instance, the CFO would like you to change one of your reports to include data from another operational system—a system you are not familiar with. You point out that you will have difficulty satisfying his request, as you don't have access to this data and that, even if you did, it "defines its fields differently." You later discover that the CFO went on to ask an analyst associated with the other system to run a special report. The CFO then had someone on his staff reconcile this report with your report and manually combine the two in a spreadsheet.

In addition to the requests from the CEO and CFO, you receive another request to produce a report from commercial enterprise application software package that your company has recently implemented. Unlike the production system that you've worked with for years, this system contains seemingly strange files that appear to contain both system and user data. As you start to learn more about these files and try to decipher the meaning of apparently incomprehensible acronyms, you wonder just how you will access the necessary data to solve this request.

Are these scenarios familiar to you?

Can you reference historical values?

Can you currently combine data from several operational systems?

Is important data trapped inside proprietary applications?

Why Is Data Integration Important?

To be successful, your organization must run its operations effectively and efficiently—which requires the ability to analyze operational performance. If you can't see how you're performing, how do you know you are making the right business decisions? You need to be certain that you can take appropriate actions to build upon your successes, initiate any corrective measures, and effectively plan for the future. For an organization to thrive, or perhaps even survive, operations and analysis must work together and reinforce each other.

To draw valid conclusions, an organization needs to be able to analyze both current and historical data from multiple disparate sources.

With a bit of luck, the organization can consolidate the data from these disparate sources without resorting to "desperate measures."

Without the whole business picture, it's difficult to make sound and dependable business decisions. That's because good decision making requires a complete and accurate view of data. And integration of all your data sources is the start to getting the complete picture—and the key to not compromising your decision-making process.

Though your organization needs a complete view of operations, the data you need often resides in a variety of application systems that do not necessarily all use the same database management system. These application systems may only contain current data values. They may not store prior data values needed to provide historical context and to discover trends.

Data integration allows an organization to consolidate the current data contained in its many operational or production systems and combine it with historical values. And the creation

of a data warehouse (or, on a more limited scale, a single-subject data mart) facilitates access to this data. Collecting and consolidating the data needed to populate a data warehouse or data mart and periodically augmenting its content with new values while retaining the old is a practical application of data integration.

The Value of Data Warehousing

A data warehouse addresses a wide variety of analytical problems. While it can exist in a variety of forms, including data marts and operational data stores containing current rather than historical values (See Appendix), each implementation can create a data platform that can be used for analytical purposes. By consolidating, standardizing, and, in many cases, summarizing the data contained in multiple operational systems, an organization can analyze the combined data to achieve a “single version of the truth.”

By integrating data within a data warehouse, organizations can more effectively use this data for analytical purposes.

There are a multitude of benefits that result from integrating operational data within a data warehouse.

You can build a **data warehouse** to:

- Integrate current and historical data values
- Combine data from disparate sources
- Create a data foundation for analytical purposes
- Establish consistency throughout the organization
- Facilitate the adoption of corporate data standards without having to modify existing operational systems
- Provide historical breadth and enable trend analysis

And you can build an **operational data store (ODS)** to:

- Obtain a complete view of your customer
- Integrate current financial data for government reporting and compliance purposes
- Consolidate current information from multiple sources

WARNING SIGNS: Does Your Organization Suffer from Poor Data Integration?

The following situations are examples of poor data integration.

- Managers are arguing about the fact that analyses results differ—even though the data came from the same operational system.
- The CEO and CFO are uncomfortable signing off on the company's financial statements. The Sarbanes-Oxley Act requires that isolated financial data be integrated and that the CEO and CFO certify, subject to penalties that include imprisonment, the accuracy of their company's financial statements.
- Presentations that include an analysis prefaced by a statement such as, "...except for the data that we were unable to obtain from..." Or worse, a presentation that begins with, "Due to the discovery of data that was not included in last period's analysis, we are reversing our decision..."
- An analyst alerts management to a potential problem that was discovered while running a query against the data in an operational system. The analyst cannot, however, answer the follow-up question, "How long has this problem existed?"
- As a result of an out-of-stock condition for a critical part, an organization must expedite an order and purchase the item at a premium price. Once the order arrives, the organization discovers that another division had an excess quantity of the same part and was trying to sell it at a discount to balance its inventory.
- While analysts use a variety of business intelligence tools to generate reports from application systems, they re-enter relevant summary values into a spreadsheet for any analyses requiring data from more than one application.
- With a series of very convincing charts and graphs, an executive presents what appears to be a thorough analysis of the cause of a particular problem. However, while the format of the presentation qualifies as a work of art, the executive's credibility suffers greatly when someone says, "That's not what that data means, where in the world did you get that?"
- An analyst runs the same report each week against an application system. However, in order to see period-to-period comparisons, the analyst maintains a spreadsheet. Each week, he must manually add a new column and enter that week's report values.

The Benefits of Data Integration: Combine “Islands of Information” to See the Big Picture

Historically, operational systems, especially legacy ones, were created to solve a particular set of needs, each evolving into an independent island of information. These needs included order entry, shipping and receiving, payroll, manufacturing and inventory control, and customer support. In the case where a company deploys the same operational system in several locations, each location might still have its own database—a separate isle of information.

By integrating data across disparate operational systems, an organization can increase the effectiveness of its data access and analysis capabilities.

Essentially as an information silo, each of these operational systems tends to have its own data definitions. For example, the marketing department might define a customer as any organization that has ever purchased from the company, while the field service department only considers a customer to be an organization with a current maintenance agreement.

What’s more, each system might have its own set of value lists. For example, an organization might store the codes for the North, East, South, and West sales regions as “1, 2, 3, 4” in one system and then as “N, E, S, W” in another. Units of measure might differ among systems with weight represented as

pounds in one system and kilograms in another or with monetary figures represented as dollars in one and Euros in another. In a multinational organization, a system in the United States might represent the date June 12, 2004 as 06/12/2004 while the same system in the United Kingdom might store this date as 12/06/2004.

Analyze Current Values and Trends

In an operational environment, organizations deploy query and reporting tools along with production reports to determine current status. They generally summarize the data and only maintain historical data values for a limited time, if at all. Though an operational system provides the most current values, these values may not be appropriate for tracking and analyzing how something has changed over time.

Consider, for example, analyzing your own household budget. While it is important to know the balance in your checking account prior to writing a check, you need to compare a series of ending month balances over a period of time to determine your savings trends. Viewing your checkbook balance at any point in time is the equivalent of reporting against an operational system; and comparing the ending balances over several months is the equivalent of accessing a data warehouse.

A single observation does not a trend make! Just as it takes two points to determine a straight line, it takes a series of values collected over a period of time to determine a trend.

Even when operational systems store historical values, the level of detail and amount of stored historical data is usually quite limited. For example, you might summarize prior year’s sales for a

customer in a single number, with the underlying detail needed to determine “sales by customer by month” no longer available (or more likely, archived on backup offline storage). One of the advantages of a data warehouse is that even if you store summarized data and not more detailed records, you can still save multiple summaries (e.g., daily sales by customer, daily sales by product, daily sales by store).

In a production environment, values are constantly changing, as most transactions typically update one or more data values. There is nothing more frustrating than performing an analysis against an operational system only to find that ten minutes later, as a result of new transactions having hit the system, you now get different results. You can avoid this problem by capturing a data snapshot and storing it in a data warehouse. While the values stored may not be up-to-the-second, they are usually collected at well-defined, cut-off cycles (e.g., monthly, weekly, daily, hourly). And this ensures the validity of period-to-period comparisons.

Data doesn't really flow. Rather, it tends to reproduce itself spontaneously with each new extract request. And unless the proper disciplines such as those involved in data warehouse integration are in place, mutations are likely to occur with each new generation.

Treat Data As a Corporate Asset

While many organizations speak of their data as a corporate asset, its quantity, unlike other assets, is not necessarily limited. Data is the one asset that can grow and reproduce almost without limit, while often mutating in the process. As it happens, despite the widespread use of data-flow diagrams, data doesn't really flow from one system to another. Rather, a copy (perhaps somewhat transformed and/or summarized) is sent to the second system while the first system still retains the data. This can lead to inconsistent data in each system and as a result, inconsistent decision making.

Integrating the data from the various operational systems within a data warehouse serves to create a “single version of truth” so that you can treat data as the enormous asset it is. To do this effectively, the lineage of the data, including its origin and/or derivation, must be readily available and not lost.

Get a 360-Degree View of Your Business

Consider multiple divisions of an organization. Each division has its own purchasing system and wants to maximize the discount it receives from its vendors. Although each of the separate divisions most likely has the purchasing details it needs to negotiate a discount with each vendor, the organization as a whole could likely negotiate better discounts if it were able to aggregate the total amount it purchased from each vendor across all of its divisions. For example, take a three-division company. If each division places a \$5 million purchase order each year with the same vendor, it could instead base its discount negotiations on the fact that the company as a whole spends \$15 million a year with the vendor.

Integrating data from multiple systems creates an environment where the whole is worth far more than the sum of the individual parts.

Because it is able to determine the total amount it purchases from the vendor, it would likely receive a higher percentage discount than each division would have received by negotiating independently.

A similar situation exists when an organization needs to integrate customer-centric data from its various systems to obtain a complete or “360-degree view” of the customer. For instance, it is necessary to combine data residing in call center, accounting, and sales systems.

Or take a situation where the analysis involves combining data contained in two or more applications. To determine profits, a company might need to integrate data from applications that track revenue with data from applications that track expenses to view the total picture.

Minimize communication errors by ensuring that every department speaks the same business language and uses the same data definitions.

Discover and Reconcile Differing Data Definitions and Business Rules

Every department understands its own data. It’s the data from other departments that always seems to be wrong and in need of reconciling to fit the individual department’s needs.

A great benefit of any data integration effort is the discovery that different parts of the same organization do not necessarily speak a common language or use the same business processes. When ultimately standardized, the long-term benefits of having a common set of business rules and common set of definitions and terms can greatly improve efficiency and effectiveness.

For example, when determining departmental productivity using “cost per employee” as a metric, do two part-time employees, each working a four-hour day, count as one employee or two? The answer is likely to differ by department and unless an organization-wide definition is established, departmental comparisons are not meaningful.

Once an organization recognizes differing definitions and standardizes on enterprise definitions, a data warehouse can facilitate their implementation. It may be impractical to modify every operational system to reflect the enterprise standard. However, it is possible to transform the data extracted from each operational system to conform to the enterprise standard definitions and value lists as the data is loaded into the warehouse. When an analyst uses a data warehouse, he is comparing “apples to apples;” when an analyst directly accesses two or more operational systems for analysis purposes, he is often comparing “apples to oranges.”

Maintain the Response and Performance of Operational Systems

The days of having to “submit queries and run reports against the production databases only between noon and 1 pm or after 6 pm” are hopefully long past. Yet, running queries or reports against the database used by an online application can still negatively impact the performance and user response time of that application. Performance counts. If an analysis request negatively impacts the response of an operational system, the analysis request will be deferred, perhaps permanently! With a data warehouse, you offload the query to an environment where the database can be optimized for this purpose.

Integrated data provides a framework that helps organizations:

- Deliver a complete view of a customer
- Offload the processing burden on operational systems
- Standardize business processes and data definitions
- Combine current and past values from disparate sources in order to see the big picture

Approaches to Data Integration

Once you recognize the benefits of and need for data integration, you have to determine how best to move forward. The two basic approaches to data integration are:

1. Develop and build your own in-house solution.
2. Acquire a commercial offering.

You should carefully consider the pros and cons of each.

In-House Development

Organizations that develop their own data integration solutions frequently do so in a somewhat piecemeal fashion, without any overall data integration strategy. They generally assign an analysis request that requires access to data from multiple sources to the IT department. A programmer then writes the code necessary to access and integrate all of the data.

Developing the ETL process by hand is difficult; maintaining it is time-consuming and expensive.

If the programmer is fortunate, the source systems are well-documented, the content of the data fields conform to the documentation, and each of the individual systems use the same value lists and code sets to represent the individual values of common data elements. If this is not the case, the programmer's task quickly expands to include data value transformations. This frequently causes the schedule to slip, especially if the data mappings are not simple one-to-one transformations.

Satisfying the initial consolidation requirement is only the beginning of the overall integration effort. As any experienced programmer knows, the initial coding effort is followed by ongoing support and maintenance especially if a new analysis request requires additional data fields or the file structure of the source systems changes. One of the givens in any applications environment is the ongoing need to respond to change; another is that "quick and dirty" one-time coding efforts frequently evolve into scheduled production jobs.

Moreover, a series of uncoordinated, individual integration tasks, even if each one were successfully accomplished, ultimately result in an assortment of uncoordinated (and usually undocumented) solutions that, collectively, become quickly unmanageable. The problem is further compounded if a different programmer is responsible for each individual data integration solution—as most programmers have their own individual programming idiosyncrasies and may have even used different programming languages

Programmer turnover is another factor to consider. While programming the initial extract program may involve creativity, future maintenance of these programs is often a thankless task. In general, programmers prefer new challenges and the original authors of the extract program may no longer be available to maintain them. And even if they are, they may not go out of their way to mention their initial involvement in the creation of the extract programs.

Purchasing Commercial Data Integration Solutions

While an initial data integration request can lead to an initial decision to develop the code in-house, software vendors have recognized the market requirement for packaged data integration solutions that can access a wide variety of data sources and perform a variety of transformations and aggregations. These data sources can encompass a variety of relational, mainframe legacy data structures, XML data structures, and message queuing systems.

The right data integration solution ensures better decision making by delivering credible, timely, and accurate information to business users.

With the possible exception of data integration tools available from some database vendors, most commercial data integration software populates a wide variety of target databases. Some database vendors have limited the scope of their data integration software to only populate their own databases.

Some commercial products also have domain knowledge and an understanding of the business rules needed to automatically generate the code to effectively access data contained in packaged enterprise applications software solutions. For example, they can generate SAP proprietary ABAP code to access data contained in SAP's R/3 and mySAP applications or have the ability to work with flexfields contained in Oracle Applications.

A commercial data integration solution that can work directly with third-party packaged software applications minimizes, or even avoids, many of the problems associated with continually modifying and retesting homegrown integration programs. This retesting of an in-house developed solution is required whenever there are changes to the packaged application software. Commercial data integration solutions usually do this as part of their normal maintenance.

By including a library of prepackaged, but extensible, data transformation functions, leading commercial data integration products are capable of performing data transformations and aggregations. This minimizes the need for custom coding and code maintenance. A robust offering will have interactive debugging facilities that allow the data integration staff to monitor the data flowing through each transformation, establish conditional breakpoints, and view and profile live data flows.

Data quality is an important component of any data integration implementation, and packaged data integration solutions are frequently integrated with data quality and data profiling capabilities provided by the data integration vendor or through a third-party partner. In some cases, the data integration vendor will OEM and support best-of-breed third-party data cleansing software, making the fact that it was developed by another vendor relatively transparent to the deploying organization.

Most packaged data integration solutions are also designed to leverage and integrate metadata. This is accomplished by conforming to standards such as the Object Management Group Common Data Warehouse Metamodel (OMG–CWM) thus allowing the data integration software’s metadata repository to exchange metadata with other CWM-compliant metadata repositories used by third-party design and business intelligence tools. When a single vendor supplies a range of integration, design, or business intelligence tools, it can easily share metadata across all of its products and greatly facilitate the overall ease of integrating the individual components.

Impact analysis, or the ability to determine how a change to a source system data field can affect a business intelligence report or analysis, is only possible through metadata integration and the resultant ability to track data lineage. Data lineage is especially important when a target field is derived from multiple source system fields. A good commercial data integration solution facilitates change data management by providing strong impact analysis capabilities including “what-if” developer scenarios.

A commercial data integration solution needs to meet both current and future performance requirements. This may be accomplished through a variety of mechanisms such as parallel processing technology and workload balancing. Change data capture techniques can track and extract only those changes that have occurred to relevant fields in the source data files since the prior extract.

Ideally, a commercial data integration solution will have several editions, with varying price points dependent on functionality (e.g., varying by number of servers, number of supported databases, number of supported users) so to address the economic constraints of departmental- and enterprise-scale deployments. If multiple editions are available, they should be upwardly compatible with each other.

In almost all situations, a packaged data integration solution is significantly less expensive in the intermediate—and long-term.

Data Integration Build Versus Buy—The Bottom Line

As a general rule, unless the data integration task is truly a “one-time” effort, organizations should strongly consider a packaged data integration solution. The short-term initial costs associated with an in-house programming effort are likely to be less than the acquisition cost of a packaged product. But on-going maintenance costs and the indirect costs associated with an inability to respond quickly to change will just as quickly consume the initial cost savings.

In addition, the productivity resulting from the ability of most commercial data integration packages to integrate and share metadata with other data warehouse tools is something that most in-house solutions simply do not consider or provide. Of particular importance

is the ability to share metadata with modeling and design tools and the business intelligence tools that will access the data warehouse. Commercial data integration packages are also likely to be integrated with or include data quality and data profiling technology—functionality frequently overlooked by homegrown, in-house development efforts.

When considering data integration tools offered by a database vendor, it's important to recognize that one of the major strengths of a database vendor's own data integration product can also be one of its major weaknesses. That is, a vendor often optimizes its solution for populating its own database. In fact, some offerings, with the possible exception of also generating flat files, can only populate a vendor's own database.

Build Versus Buy Decision Criteria

Data Integration Considerations	Build	Buy
Initial start-up cost	Lower	Higher
Continuing operational cost	Higher	Lower
Ongoing support and maintenance	In-house responsibility	Vendor responsibility
One-time "quick and dirty" task	Consider	May be overkill unless "one-time" task becomes ongoing request
IT staff requirements	Higher	Lower
Data sources / data targets	Single / single	Multiple/multiple, Multiple/single, Single/multiple
Data sources include third-party enterprise application software	Changes to data source will require IT effort	Changes to data source likely handled by data integration vendor
Complex transformations	Limited; IT must create custom code	Comprehensive; customized user code can usually be added to vendor-supplied library
Impact analysis and data lineage	Limited	Usually included
Data quality functionality	Frequently overlooked or relatively limited	Usually included, sometimes via OEM partnerships
Metadata sharing with third-party tools	Frequently overlooked	Likely to conform to industry standards
Integration with vendor's other products	Not applicable	By design; also provides single source of support

Conclusion

Reliable data is the basis for sound decision making. And data integration is also the key to information accountability—do users of business intelligence tools feel that are basing their decisions on data that is trustworthy? The best business intelligence tools are of little value if the data they analyze is not complete and accurate.

Operational and analytical systems complement each other. Organizations must effectively deploy both in order to succeed. For analytic purposes such as trend analysis and forecasting, it's necessary to collect time-stamped data values from multiple sources in a data warehouse or data mart. For operational purposes, it's frequently necessary to report off of the data resident in an operational system. Organizations can use an operational data store to consolidate current data values from multiple operational systems.

Data integration technology is used to bring this data together. In fact, data integration and data quality solutions are the keys to achieving information accountability. While some organizations choose to develop their own in-house data integration solutions, those that use packaged software solutions can benefit from the vendor's expertise and experience in working with multiple, and sometimes esoteric, data sources. This also frees up their staffs for more productive tasks. Additionally, commercial data integration products usually provide metadata interoperability with other tools and track data lineage. Regardless of how obtained, data integration enables data warehouses, data marts, and operational data stores—which all provide organizations with the means to make reliable business decisions. Successful data integration is a key factor for an organization's ultimate business intelligence success.

Appendix

Many consider Bill Inmon the father of data warehousing. In his book *Building the Data Warehouse*¹, he defined a data warehouse as, “a subject-oriented, integrated, nonvolatile, time-variant collection of data in support of management’s decisions.”

To appreciate the data content differences between a data warehouse and an operational application (or production system) environment, the following table may be helpful.

Data Characteristic	Production Application	Data Warehouse
Data use	Operational	Analytical
Level of detail	Detailed	Detailed and summary
Data currency	Real-time, latest value	Multiple historical generations
Longevity	Relatively brief	“Forever”
Stability	Dynamic	Static
Scope of definition	Application-wide	Enterprise-wide
Orientation	Application	Subject
Data operations	Capture/Update	Read
Data per transaction	Limited	Large
Database optimized	For update	For access

Source: Updated from “Data Warehouse—Concepts and Implementation Strategies” presentation, M. Schiff.

While these characteristics are not meant as absolutes for each environment, they represent general statements as to what is typical of each environment. For example, although data warehouse content is obviously updated with new values each time a new snapshot is added, the general use of the data in the warehouse is for read-only analysis purposes. A data warehouse typically adds new, time-stamped values of existing data elements; a production system usually modifies existing values. For example, a production application for payroll might contain the salary of each employee; a data warehouse might contain the salary history for each employee. When an employee receives a salary change, the new value would replace the old value in the payroll system while an additional record, containing the new salary and effective date, would be added to the data warehouse content where it would reside along with the previous salary and quite likely all past salary amounts (or at least a reasonable history) for each employee as well.

There are also times when an organization needs to collect data from several operational systems for additional operational purposes such as determining current part quantities across all of its inventory control systems. This data warehouse variant is commonly referred to as an operational data store. While it differs from the classic data warehouse as it stores relatively current values and minimal history, the process of bringing this data together is another classic example of data integration.

Bill Inmon and Claudia Imhoff highlighted this difference in their book, *Building the Operational Data Store*¹ when they defined an operational data store as a “subject oriented, integrated, current valued data store, containing only corporate detailed data.”

¹W.H. Inmon and Claudia Imhoff, 1996, John Wiley & Sons, Inc.

About MAS Strategies

MAS Strategies specializes in helping vendors market and position their business intelligence and data warehousing products in today's highly competitive market. Typical engagements include SWOT analyses, market research, due diligence support, technology white papers, public presentations, and helping organizations evaluate tactical and strategic product and marketing decisions. MAS Strategies also assists user organizations in data warehouse procurement evaluations, needs analysis, and project implementations.

With over 30 years of industry experience as a developer, consultant, vendor, industry analyst, and end-user, MAS Strategies' principal and founder, is an expert in developing, marketing, and implementing solutions that transform operational data into useful decision-enabling information. In addition to his responsibilities at MAS Strategies, he is vice president of the Data Warehousing and Business Intelligence service at an industry analyst firm where he provides tactical market intelligence and analysis.

His previous positions included Executive Director - Data Warehousing and Advanced Decision Support for Oracle Corporation's Public Sector Group and Director of Software AG's Data Management program where he was one of the industry's earliest proponents of the data mart concept. In 1984, while at Digital Equipment Corporation, he formulated the architecture for one of the first successful data warehouse implementations. In previous positions as IT Director and Systems and Programming Manager he acquired practical, first-hand, knowledge of the technical, business, and political realities that must be addressed for any successful systems implementation or product launch.

MAS Strategies' founder earned his Bachelor and Master of Science degrees from MIT's Sloan School of Management where he specialized in operations research as an undergraduate, and in information systems as a graduate.

For further information about MAS Strategies, visit its web site at:
www.mas-strategies.com.

About Business Objects

Business Objects is the world's leading business intelligence software company. Business intelligence enables organizations to track, understand, and manage enterprise performance. The company's solutions leverage the information that is stored in an array of corporate databases, enterprise resource planning (ERP), and customer relationship management (CRM) systems.

Popular uses of BI include enterprise reporting, management dashboards and scorecards, customer intelligence applications, financial reporting, and both customer and partner extranets. These solutions enable companies to gain visibility into their business, acquire and retain profitable customers, reduce costs, optimize the supply chain, increase productivity, and improve financial performance.

In December 2003, Business Objects completed the acquisition of Crystal Decisions, the leader in enterprise reporting. The combined product line includes software for reporting, query and analysis, performance management, analytic applications, and data integration. In addition, Business Objects offers consulting and education services to help customers effectively deploy their business intelligence projects.

Business Objects has more than 24,000 customers in over 80 countries. The company's stock is traded under the ticker symbols NASDAQ: BOBJ and Euronext Paris (ISIN: FR0004026250 - BOB). It is included in the SBF 120 and IT CAC 50 French stock market indexes. Business Objects can be reached at +1 800 877 2340 and www.businessobjects.com.

**Americas**

Business Objects Americas
3030 Orchard Parkway
San Jose, California 95134
USA
Tel: +1 408 953 6000
+1 800 877 2340

Asia-Pacific

Business Objects Asia Pacific Pte Ltd
350 Orchard Road
#20-04/06 Shaw House
238868
Singapore
Tel: +65 6887 4228

Europe, Middle East, Africa

Business Objects SA
157-159 rue Anatole France
92309 Levallois-Perret Cedex
France
Tel: +33 1 41 25 21 21

Japan

Business Objects Japan K.K.
Head Office
Yebisu Garden Place Tower 28F
4-20-3 Ebisu, Shibuya-ku
Tokyo 150-6028
Tel: +81 3 5447 3900

For a complete listing of our sales offices, please visit our web site.

► www.businessobjects.com