White Paper

Real-Time Business Intelligence

A new paradigm for rapid, cost-effective information aggregation from disparate data sources — designed to provide distributed users with real-time results for decision support and analysis.

May 2006

Prepared by:
Bob Ivkovic & Neil Ross
IT Architects, Principals

This page left intentionally blank for double-sided printing

1.0 AN INTRODUCTION TO IT ARCHITECTS

IT Architects (ITA) is a consulting firm dedicated to helping our clients derive business value by building the right information technology architectures in support of solution initiatives. ITA has a strong foundation in developing and deploying IT architectures across the various architecture disciplines, including business process, data, application and technical. ITA has led organizations in the deployment of Enterprise and Service Oriented Architectures to support system evolution planning and large-scale ERP implementations.

ITA is primarily focused on system delivery, architecture, and integration. ITA has been engaged to architect and develop integration solutions for ERP implementations such as SAP, JD Edwards, PeopleSoft, Siebel, and others. The firm has also successfully deployed system integration solutions using technologies from TIBCO, Vitria, Mercator (Ascential/IBM), Microsoft BizTalk and other EAI vendors; and specializes in the deployment of Web and other services based integration solutions built upon both the J2EE and .NET frameworks.

ITA has been engaged by clients to define architecture frameworks that are tailored according to an organization's IT maturity, technology mix, and business relationships. ITA has also assisted organizations with the transition and knowledge transfer required for ongoing success in deploying, scaling and managing IT solution initiatives.

ITA's website at www.itarchitects.ca provides further information pertaining to our services, frameworks and strategies. The corporate overview can be found at http://www.itarchitects.ca/ITA_Marketing_Presentation.pdf.

2.0 AN INTRODUCTION TO VISIPHOR

Visiphor is a software and professional services company specializing in the development and marketing of software products that enable integrated access to applications and databases. Visiphor has successfully deployed its integration technology as part of an application suite that automates law enforcement procedures and evidence handling across North America and abroad. This same technology has also been very successful at quickly crafting and deploying regional information sharing networks for justice. This real-time requirement to pull and aggregate data from multiple sources simultaneously has carried over to other industries including health care, financial services, government services, telecommunications and others.

Visiphor's core technology suite is referred to as the Briyante Integration Environment (BIE). Using industry standard Web services, Visiphor delivers a secure and economical approach to true, real-time application interoperability. Since data can be pulled simultaneously from multiple sources without storing it in an interim repository or warehouse, organizations have introduced BIE to complement their robust data warehousing, data analytics and operational reporting tools. This capability—often referred to as *federated query*—has allowed BIE to become an extremely attractive real-time business intelligence solution with significant ROI.

Numerous production deployments of BIE have demonstrated remarkable reductions in time, complexity, and risk associated with defining, implementing, and supporting integrated access to physically and technologically disparate data and systems. Visiphor is a Microsoft Gold Certified Partner and its software products adhere to widely adopted industry standards.

The Briyante Integration Environment may be classified as Enterprise Information Integration (EII) or Virtual Data Warehousing. ITA has recognized BIE as a unique and powerful tool that allows organizations on-demand access to disparate information—delivering it in an understandable format to the person making the request. This fully distributed, real-time access to a broad range of data can greatly enhance decision-making, offer a “snapshot” in time view of business activities, and validate similar information spread across all systems on the network. The BIE technology is also adept at eliminating redundant data entry into multiple systems.

3.0 IT AND BUSINESS CHALLENGES

Companies are faced with a number of challenges that impact their capability to access and use data distributed across both internal and external sources, including:

Increased Complexity – Architectural components are increasingly complex and sophisticated, often involving intricate inter-relationships and dependencies. Issues such as privacy, security, integrity, and regulatory demands such as Sarbanes Oxley are some of the key drivers. Duplication exists in almost all architectures, and as a result, the entire environment is becoming more of a challenge to holistically manage.

Increased Reliance on IT – Business units look to IT to provide leadership in understanding their needs and delivering viable IT enabled business solutions.

Accelerated Rate of Change – Like every other organization, IT is challenged to keep up with continuous change in technology, organizational structure, regulatory requirements, etc. IT requires more flexible architectures to respond more quickly to business needs, opportunities and threats.

Business Unit Silos – Autonomous units with unique business requirements and varied views of priorities often result in providing multiple "sources of truth".

Distributed Enterprises – Autonomous subsidiaries or branches located in different geographies with unique business requirements and their own enterprise applications and data stores.

Mergers and Acquisitions – Today's fast-paced business environment results in numerous mergers and acquisitions. Executives want to see consolidated information and systems delivered as soon as possible due to the mission criticality of various applications and data.

IT Cost Control – Maintaining control of IT costs is a constant concern while still facilitating changes quickly and efficiently.

Enterprise Projects – Many companies are striving to position themselves for the eventual replacement of legacy and other systems. These systems impact the enterprise as a whole, and as such, must be deployed as enterprise-wide. The project teams responsible for the deployment of new systems and applications can only do this successfully by aligning to enterprise requirements and facilitating a certain amount of interoperability between best-of-breed applications (best supported by an enterprise architecture initiative).

Managing the EAI Potential - Putting EAI in place can also lead to concerns that the complexity will be difficult to manage over time. How will organization manage its existing and future EAI integration requirements?

The Enterprise Architecture Potential – Recognizing that enterprise architecture can help to identify synergies and opportunities that can translate into true business value, how can long planning and deployment timelines be avoided in order to support business agility?

All of these challenges are making it increasingly more important to access geographically and technologically disparate data sources that use numerous data representations—to do so simultaneously, on-demand and in real-time. There is a requirement to have the flexibility to deliver information requests quickly without disturbing existing, critical IT infrastructure.

4.0 THE BRIYANTE VALUE PROPOSITION

ITA has formed a strategic alliance with Visiphor to address a recurring request from clients who tell us that they are looking for a rapid, cost-effective solution to aggregate information from disparate data sources and provide distributed users with real-time information for decision-support and analysis. The Briyante Integration Environment (BIE) is the answer because it facilitates a new approach to an old problem.

The BIE consists of a Server and Design Studio, a standards-based data sharing toolkit and development platform (patent pending). It unifies access via the internet to information assets through a single interface, and is characterized by:

- a visual development interface that abstracts much of the complexity involved in connecting to various data sources, querying the source data, and delivering it to a user or another system in an understandable format
- one-click generation of .NET Web services that publish and consume data
- a unique dynamic SQL engine that provides on-the-fly adjustments to submitted queries in order to suit the format(s) of the source system data

4.1 Basic Advantages of the Visiphor BIE Solution

Visiphor's BIE offers an effective solution for business users across the organization because BIE bundles the benefits of unified data access in ways that no other solution can:

- Data access is fully distributed, on-demand and in real-time
- Data migration and conversion are unnecessary
- No modifications are required to existing systems
- Existing business rules can be applied to data presentation

- The underlying complexity involved in gathering and massaging data is inconsequential
- Offers complete flexibility in modifying data access targets and selection criteria on a routine basis, thereby testing, scaling and expanding the system as needed.
- Data ownership is maintained at the source; can be economically deployed at the department level; and can be migrated across the enterprise
- Superior scalability in a standards-based, secure architecture; and an iterative, “divide and conquer” approach to application deployment
- Capability to access data as old systems are retired and new systems introduced

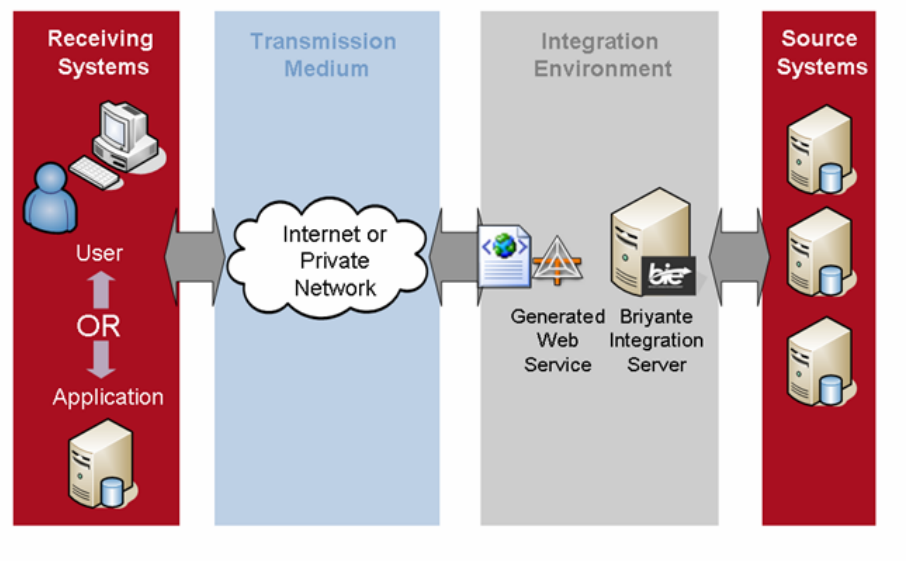

5.0 REAL-TIME BUSINESS INTELLIGENCE SCENARIOS

The following describe some typical business scenarios where real-time information is needed to support business processes by enhancing the quality of decision-making and analysis. ITA and Visiphor's Briyante Integration Environment (BIE) can respond to these situations with the rapid deployment of cost-effective solutions.

5.1 Oil & Gas - Energy Trading

Many energy companies have diversified their operations to include production, transportation and marketing activities. For example, larger oil and gas company activities include oil and exploration, production, mid-stream processing, transportation, storage, and marketing, as well as power generation and marketing.

Energy traders involved in the volatile power and gas markets need immediate and current information to make split-second decisions. The problem is the information they need may be stored in several different information technology systems within the operations, marketing, and finance departments. Operational systems for production, transportation, storage, delivery and external market information are dispersed and are provided by different vendors, making it difficult to get an instant and accurate picture of the current situation. Today, a trader may need to view several different screens of data from unrelated systems or get hard copy reports to assemble key information. This takes time and valuable opportunities can be missed.

Energy traders need real-time information to make decisions to maximize profits. For example, they need to know their current aggregate position on the buy, sell, and hedge contracts they have in place, the value of these deals against current market pricing, and the current credit position of

certain counter-parties against pre-approved limits in order to decide whether to extend, liquidate, or hedge their positions to maximize returns or limit their exposure.

5.2 Oil & Gas - Well Data Management

A well means different things to different people within an organization depending on whether they are engineers in the field, corporate well data administrators, or accountants who managing AFEs for a company's wells. Oil & gas companies have failed miserably integrating their field well, corporate well and AFE systems using EAI technology. Not only is it difficult to get wells and well data synchronized across these systems, but once the interfaces are developed they slowly lose their effectiveness as they encounter new well data instances and associated anomalies. As a result, users of these systems do not have the time nor the patience to figure out why an interface fails or cannot handle a specific well data message. Departments and individuals often go back to manually entering data into each system because it is easier than figuring out the root cause of the failure.

Well data users need a way to pull related data from each system and present it for both specific reporting purposes and an ad-hoc basis in order to answer unexpected questions from various departments. A well means different things to an organization. This is why it is important to be able to aggregate well data in different systems to provide the correct aggregation of well data in a coherent but yet functionally specific format that offers the correct meaning for the user.

The Well Life Cycle

Defined ownership of well data across different business areas during the well life cycle is critical, but in reality does not exist. Many organizations have been ineffective in sharing production data with operations, production accounting, marketing and other business areas. Thus, an opportunity exists to reduce cycle times from well prospecting to production

by establishing proper well data governance and sharing. Furthermore, this will optimize reserves management by providing access to consolidated well data for proper forecasting.

5.3 Inventory Management

The product or equipment inventories of organizations can be managed in several different geographic locations and warehouses. Although each warehouse has the means to manage its inventory, the organization has been challenged to manage inventory across all of its warehouses at an enterprise level. To do so requires a consolidated view into often diverse inventory systems and data stores in order to control inventory levels and provide alert mechanisms for reorders or and product transfers on a per warehouse basis. Unless organizations have the ability to access disparate inventory records across the enterprise, they will be unable to streamline inventory management operations.

5.4 Health Care - Unified Patient Record

Health care facilities hold pieces of the overall patient health record in multiple sources. They are faced with the onerous task of pulling together records for a given patient across multiple applications. Furthermore, they are challenged to gather this information in a timely fashion for access in busy clinical settings. This same problem exists at the state or provincial level where similar data has to be pulled from all systems in the state or province to effect a single patient health record. There are probably hundreds of information sets that would be applicable – from an executive dashboard that a CEO would view, to a health safety dashboard, to other dashboards for operations, staff scheduling, transcription workflow, etc.

This is even more critical when you consider all the different health care providers, pharmacies and medical entities that need access to complete and timely patient records in order to provide adequate health services to their patients and customers, not to mention their need to abide by government health regulations and policies set forth by the medical industry.

5.5 Health Care - Patient Research

There are initiatives underway to provide real-time visualization of research data, and the amalgamation of research data with data from regional health systems. The primary issue with health care records comes down to inconsistent use of terminology, even within a single health care facility. The terms “episode, encounter, visit, service event, service item,” etc. are all differently defined and used interchangeably. Unfortunately, many systems do not understand that differences in label semantics, and so confusion and subjective interpretation can be the results.

5.6 Historic Customer Billing

Many companies recognize that it is extremely difficult and expensive to migrate historic customer data from a legacy system to a new one. This scenario is most common in customer billing systems where the rate structures, price plans and product catalogues do not resemble the data structures supported by the successor application. At the same time, there is a need to periodically access historic customer billing information to respond to billing inquiries and complaints. Many corporations are opting to leave their legacy applications running in “read-only” mode where information is still accessible but no new transactions are processed in the legacy application. This requires multiple sign-ons and knowledge of how to navigate both systems to obtain the needed information.

Customer service representatives are often faced with a situation where they need access to legacy data in order to respond to customer billing inquires and complaints. It can be quite ungainly to do so in conjunction with information derived from multiple systems.

5.7 Customer Relationship Management - General

Many companies deploy best-of-breed solutions for capturing orders using a Customer Relationship Management (CRM) solution, or alternatively the order lifecycle is managed through a workflow application; or provisioning is supported using a separate provisioning application and then implemented through deployment and engineering applications.

When a customer wants to know the status of an order, several inquires to different organizational units and queries to their supporting applications are required to provide this information. This is not only time consuming, but is expensive and an inefficient means of responding to a predictable recurring informational need of a customer.

5.8 Customer Complaint Handling

Many companies deploy several customer complaint management or trouble ticketing applications: one for IT, another for customer complaints and others for internal use. In many situations, several trouble tickets are opened to respond to a single situation as several participants are involved in the resolution of the reported problem. Each participant tracks their individual involvement in their own trouble ticketing application. Several queries and verbal confirmations are required to determine the overall status of a trouble ticket when dealing with a customer.

5.9 Order Management - Data Validation

Significant damage is done to a customer relationship when a customer orders a product only to find out some time later that the product cannot be supplied because the customer's service address is not within the service area boundaries. There is no easy way to tie customer addresses to service area locations. Thus, companies are forced to complete an order transaction only to contact the customer later and tell them that they cannot fulfill their order or contract.

5.10 Customer Management - Data Duplication

Another typical problem is that customer data exists in several applications. When a customer requests updated information such as change in the name and phone number of a site contact, there is a need to determine the status of this information in each application to ensure that the updated information has been applied consistently. Customer service reps usually

scramble to find the correct data and correct it—while IT tries to figure out which source system has the correct data and how to reconcile it across the enterprise.

5.11 Composite Views of Customer - Data Aggregation

A customer calls into a contact centre requesting information about all of the products and services that they have purchased. The customer has several questions about the diverse group of products and services. The Customer Service Representative (CSR) needs several minutes to sign-on to different applications to conduct name searches to find the customer's accounts and then make detailed queries to obtain baseline information.

During the development of the baseline, the customer becomes impatient and starts asking questions to which the CSR cannot respond, which in turn creates an unsatisfactory customer experience and a stressful call for the CSR. This wastes time and adversely impacts the customer relationship. The CSR has no way to aggregate the data and bring it into a single view to answer the customer's questions or to make every day decisions.

5.12 Customer Migration - Data Consistency

An existing customer informs the organization that they are relocating to another location that is also serviced by the organization. The Customer Service Rep is pleased that they have retained a valued customer. The CSR needs to develop a composite understanding of the billing history, purchased products, credit status, services used, open orders, etc., which are all managed in separate applications.

The CSR must conduct multiple log-ons and name searches, and then consolidate customer information to develop a full understanding of the customer profile. Again, the result is an unsatisfactory customer experience due to the protracted period to perform what is perceived to be a simple request.

5.13 Law Enforcement

Shortly after a bank robbery in his patrol area, a police officer performs a routine traffic stop. When the officer queries his local department's database about the driver, no hits occur because there are no outstanding warrants or a history of interaction with the department. If only the officer had additional access to a regional information sharing system so he could view information from the motor vehicle department, probation database and other police departments' record management systems. This would allow the officer to determine if the driver is using any known aliases, which could be cross-matched against the vehicle plate. A picture from the motor vehicle department would confirm the driver's identity along with his record of robbery and violent offences in other neighboring municipalities. (Note: Visiphor has developed such regional information sharing systems in King County, Washington; Charlotte, North Carolina; and British Columbia, to name a few).

5.14 Scheduling Consolidation

Regardless of the type of organization, every enterprise has scheduling dependencies and requirements for schedule consolidation across divisional or project boundaries. Large organizations should have the ability to consolidate multiple project plans into a single program plan.

5.15 KPI Measurement

Organizations always have a need to measure their performance based on a set of Key Performance Indicators (KPIs). Many have implemented so-called KPI "dashboards" to capture and display KPI data at both the operational and executive levels. These dashboards can often blend Web pages, spreadsheets, Sharepoint sites, proprietary interfaces and other formats to present information. Unfortunately, these multiple presentation layers are non-unified and inconsistent, so they rarely deliver the required information in a form needed to make calculated and timely decisions. Furthermore, organizations do not always target the right KPIs to measure actual performance metrics because they do not have the means to

retrieve them. For example, an organization's cash-flow management function may be premised on the tracking of purchase order requisitions and submissions. An organization may not have the ability to tap into multiple systems to see which POs have been submitted and completed, and determine how this is going to affect cash flow. However, this may be possible by implementing a dashboard that allows the user to view data brought in through a federated query against multiple systems. A Web query can be used to receive data as an XML message into a spreadsheet program such as MS Excel so further analysis and data manipulation can be performed without building a specialized application.

5.16 Business Continuity and Systems Evolution Planning

Organizations in all industries are evolving their application portfolios as they are continuously replacing legacy systems with ERPs and package solutions. Thus, these organizations are in constant flux by moving and converting data into transitional or semi-permanent data repositories. The business is disrupted while project teams build interim solutions to get appropriate data to the end user. The fact that organizations are expected to evolve their applications and systems to meet changing business requirements is a problem, as internal business users, customers, partners and suppliers rely on both existing and requested information access. Organizations are expected to evolve their technology to meet the demands of the changing business environment while providing minimal impact to the business.

6.0 THE SOLUTION

One option is to build a data warehouse where this information is aggregated by a batch or near-real-time process. This approach is time consuming, costly to build, costly to maintain, inflexible, and significantly increases the complexity of the environment.

Another option might be to wait for the applications or data sources to be integrated before making the information available. This is not something most businesses would favor or accept, since this usually involves time-consuming hard coding connections and transformation routines between applications and systems.

A better option is to make use of the advanced integration and data sharing solution from Visiphor, which can solve these and other problems for organizations in various industries. The Briyante Integration Environment (BIE) provides an easily-to-use flexible tool that securely bridges the “islands of information” from dispersed legacy and proprietary data stores, and assembles it dynamically (using pre-defined business rules) into a single Web interface that provides the mission-critical data end users need. The source data remains in its native application database, so business unit managers retain the ownership and integrity of their data.

For example, in the case of energy trading, data from specific gas pipeline receipt and delivery points, storage facilities, contracts in place, market supply/demand, and pricing conditions can be assembled into a powerful real-time overview that enables traders to capitalize on situations worth thousands or millions before the competition can react. Timely reports can be delivered in spreadsheet, PDF, HTML, or XML formats.

6.1 The Benefits and Outcomes

The Briyante solution offers a real-time information delivery solution that has a rapid ROI even at a department level, and is easy and quick to deploy with a minimum of technical and end-user training due to its drag and drop technology for creating data views. It can access data from any

application with OLEDB and ODBC connectors, and uses highly secure SSL encryption to deliver read-only information via the Internet to any authorized user's desktop, laptop, or PDA/cell phone. The BIE uses industry standards, such as platform-independent Service Oriented Architecture built on the Microsoft .NET framework, so it is scalable and flexible enough to be migrated from a departmental, to a divisional, and ultimately to an enterprise-wide deployment as business needs dictate. The value proposition from a business standpoint is Visiphor's product ability to provide solutions that are flexible, responsive to change, and have the ability to abstract complexity.

7.0 CONCLUSION

ITA can cost-effectively design and implement Visiphor's Briyante technology to solve the problems illustrated in the business scenarios presented in this document, in addition to many others. Our company has the knowledge and expertise to develop the effective solutions for these business problems. ITA and Visiphor, as a team, have a unique product and deployment capability to empower your organization's people to access and use relevant and current information to make decisions and meet customer demands.

Visiphor's BIE technology is designed to access real-time information quickly and easily from disparate data sources without the latency or transactional constraints imposed by EAI, nor the analytical and financial constraints imposed by data warehouses. BIE provides the data needed to impose custom analytics and improve business intelligence.

Visiphor's leading-edge technology easily performs federated queries information aggregation at run-time through a single interface. New systems can be rapidly and effectively brought online as they are introduced into the integrated environment. This gives organizations the agility and flexibility to access whatever data they need, whenever and wherever they need it. Now you can connect what matters.