KEMAL A. DELIC
UMESHWAR DAYAL

Decision makers throughout your enterprise need an IT architecture that serves their needs — rather than the other way around. Here is a view of emerging enterprise analytic systems that uses BI and analytic requirements as the point of origin

A New Analytic PERSPECTIVE

In the first 50 years of the history of computing, we've seen the deep and comprehensive infusion of computing systems into various business domains. Today, computing systems represent an indispensable infrastructure, with which we run, manage, and coordinate business operations. We envision that the next decade and beyond will bring a new era of ubiquitous computing systems. We believe that this progress will lead to the better use of interoperating enterprise analytic systems: in other words, a vision of CRM, ERP, enterprise application integration, and enterprise knowledge management applications all interwoven through analytics. We predict the morphing of contemporary enterprises into more intelligent enterprises, which will operate in various business markets that mimic natural ecosystems, evolving and adapting based on their quest for survival.

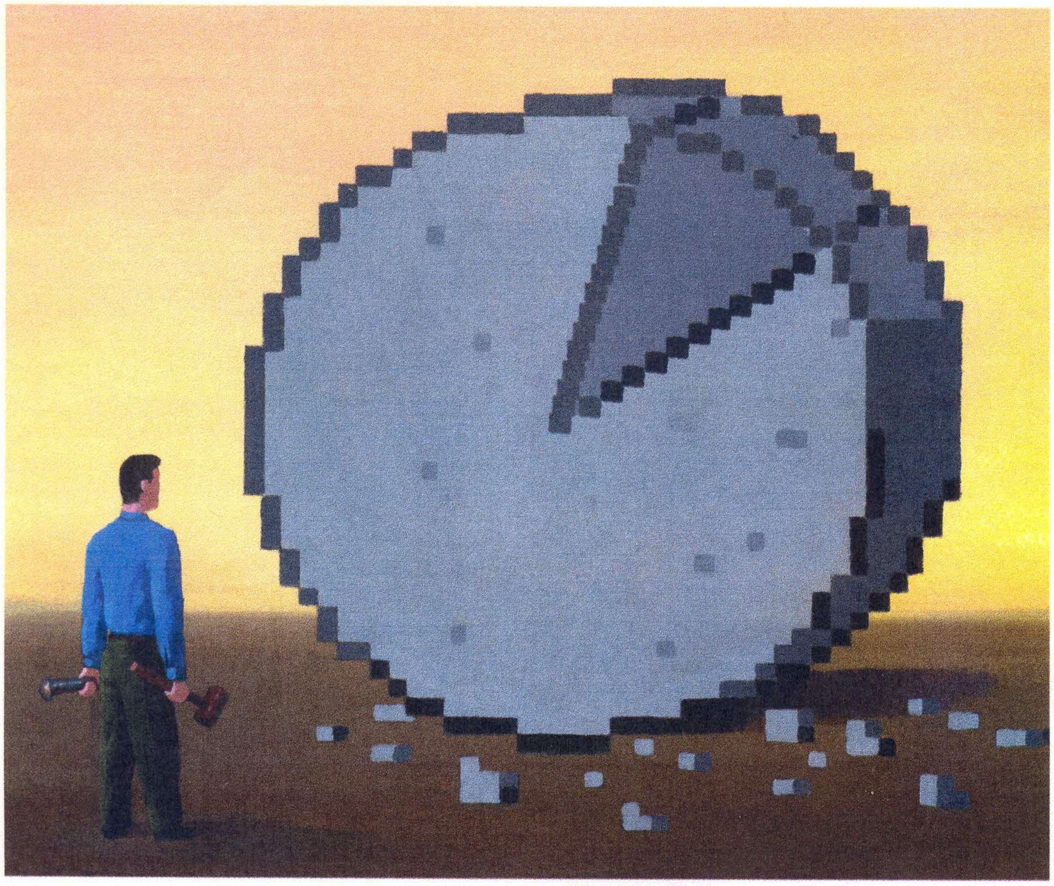

FEATURE

IMPERATIVE 1

1

Enable decision makers across the organization to act with the highest level of confidence and knowledge.

For a complete list of imperatives for the intelligent enterprise, visit IntelligentEnterprise.com.

Jonathan Evans/Photodisc Green

Today's enterprises function nonstop under an endless and variable stream of events that end with either the closure of a business transaction or a reaction to originating events. Data and information repositories, logs, and traces capture transactions and events, thereby providing the foundation for higher-level applications — namely enterprise analytics. Large enterprises operate today with the reality of huge data volumes and vivid business dynamics. They transfer terabits of data daily while accumulating terabytes of data and information and processing them with teraflops of computing power. To function smoothly, enterprises must build their architectures carefully to govern distinct enterprise layers for event management, transaction processing, and the rendering of enterprise analytics.

Today, most approaches to enterprise analytics start from the inside out: from the perspective of data and information sources themselves. We suggest starting from the outside in — from the actual users' needs and their typical decision-making behavior. In this article, we'll outline three broad categories of users and discuss how an enterprise analytic architecture might serve them. We'll argue that understanding how they make decisions — and in this way, clarifying their data and information needs — will help articulate a better conceptual enterprise analytic architecture.

MODELING COMPLEXITY: APPROACHES

Business enterprises are extremely complex systems, difficult to manage and hard to model. Nonetheless, we model business enterprises to better understand how to manage them, where and how to evolve them, and how to improve business and financial performance. Our experience tells us that at least two-thirds of enterprise complexity is attributable to organizational complexity, while one-third is due to IT complexity.

Taking an abstract, holistic view, we can stratify enterprises into three layers:

1. Business
2. Systems and tools
3. Data and information sets.

In this simplified view, IT systems in layers 2 and 3 function as mediators and serve as the business conduit ("fabrics") for monitoring business operations, grasping the holistic enterprise situation, and executing the necessary actions that optimize a limited set of top business parameters, such as revenue, profit, growth, and market position. Corporate managers and senior executives are typically in the business cockpit and therefore take these actions.

Another useful way to consider an enterprise is to split it into two distinct parts: static or slowly changing (such as the IT infrastructure); and dynamic (such as the typically enormous number of operational parameters), which includes data and information flows within and outside the enterprise. Both parts are critical to supporting corporate decision-making. The static part acts as an overall context, based on a current set of operational parameters. However, decision-making is dynamic by nature: It has urgency and involves different types of people and styles.

DECISION-MAKING: A MATTER OF STYLE

Perfect decisions require perfect decision makers — and timely support from perfect data and information sets. We know reality is slightly different. Real-world decision makers live in an imperfect world. To give structure to this kind of world and understand groupings of potential consumers of analytic artifacts, we can think about three generic groups of decision makers: executives, managers/analysts, and employees. We will briefly discuss each group's decision-making characteristics.

A major challenge for executive decision makers is dealing with huge volumes of weakly structured data sets and other imperfect information sets in a timely manner. Top executives should be able to operate across terabytes of raw enterprise data that is aggregated, abstracted, and transformed into higher-density artifacts (such as maps, simulations, and insights) that can be presented to them via dashboards and other summaries.

Managers and analysts will be exposed to analytic artifacts created from gigabytes of data before reaching decisions. Finally, lower-level employees also encounter a daily inflow of megabytes of multimedia data and information. Anecdotal evidence indicates that these three generic groups of decision makers make (respectively) intuitive, rational, and routine decisions (see Figure 1).

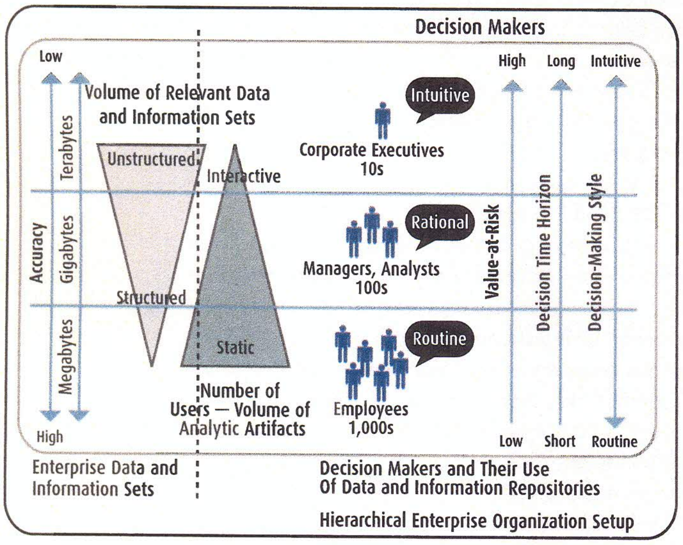

FIGURE 1 *Decision makers and their use of data and information repositories.*

ENTERPRISE ANALYTIC SYSTEMS: SUPPORTING DECISIONS

Enterprise analytic systems cover different domains, address varying value-at-risk entities, and require different analytic artifacts to improve and accelerate decisions. In practical terms, supporting decision-making may mean exposing decision makers to expectations, predictions, and forecasts so that they can rationalize their options. To outline possible consequences and envision scenarios, developers must consider users' expectations. For example, cause-and-effect models will produce predictions. Randomness is a key factor in forecasting, which makes the creation of large systems for modeling and forecasting outcomes and behaviors difficult and complex. We can consider each generic group of users more closely:

Executives are senior employees responsible for strategic decisions. Perhaps without knowing it, they digest a huge amount of data and information to make a strategic decision about how to improve corporate client profitability, for example. Executives also make intuitive decisions, and do so with input from the enterprise data and information domain, and with external information about market conditions as well as their views of partners and suppliers. The time horizon for their strategic decision-making is usually weeks or months.

Managers and analysts are employees with specialized skills responsible for tactical decision-making. To make decisions about how to optimize cash flow, for example, they digest a modest amount of data and information. Characteristically rational decision makers, managers and analysts focus on expertise-related subjects involving relevant, but limited data and information sets. Their decision-making timeframe is typically days or weeks.

Employees making routine decisions about how to optimize employee communication costs, for example, digest a limited amount of data and information related to their domain. Employee decisions typically have a timeframe of hours or days.

Different users of enterprise analytics use different artifacts. Senior executives and analysts typically require interactive content (sometimes called "active content"), while larger communities of employees require passive (or static) analytic content. Each generic group we've outlined also exhibits different usage patterns. Analysts typically need a powerful client/server setup; managers need browser access to analytics from a portable Web-connected machine; and lower-level

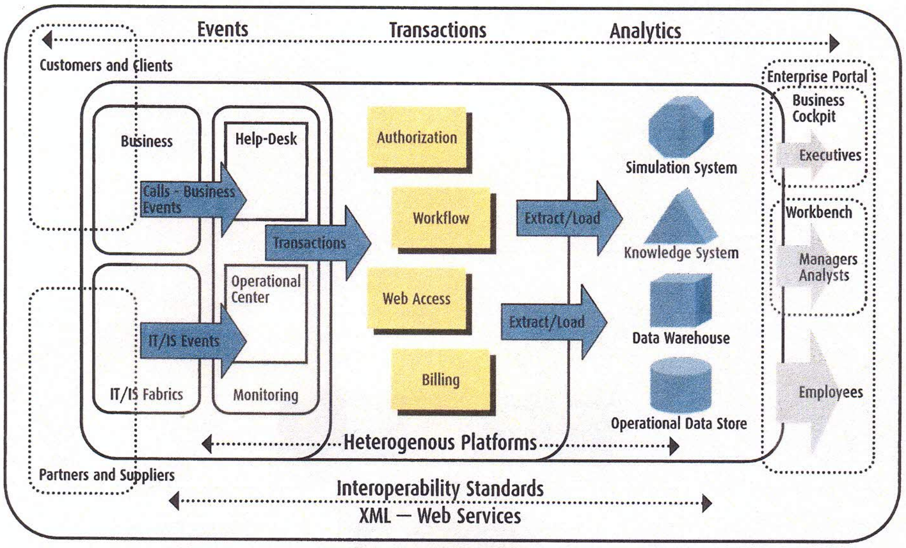

FIGURE 2 *Stratified enterprise architecture.*

employees might need analytics right on their desktops or accessible via temporarily connected PDAs, mobile phones, or other portable devices. These are important considerations for the design and architecture of analytic systems.

We can now outline what we propose as a stratified enterprise architecture arranged into different computing layers that support different needs and decision-making styles.

THE ARCHITECTURE

The typical artifacts for an enterprise analytic architecture are estimates, forecasts, scenarios, views, and rules. They are the usual result of human observations, reasoning, and intelligent analysis. We will call computing systems specifically designed to render analytic artifacts "enterprise analytic systems." Today, these are typically data warehouses, knowledge bases, and simulations. Their objective is to make sense out of large data and information sets and pack final results into the appropriate form-factor for the decision maker.

An enterprise analytic artifact might be delivered in the form of a graph, view, document, program, interactive graphic, simulation, or model. Some of these artifacts are passive, while others are interactive. Some

FINANCIAL SERVICES: THE LEADING EDGE

THE FINANCIAL SERVICES industry is shaping up to be the first to profit from new technology and architectural developments in enterprise analytic systems, serving both clients and customers. Let's consider an example of a large North American bank that's trying to grasp the incredible volume of events and transactions, and estimate the number of potential users and consumers of analytic artifacts. This bank processes more than 5 billion transactions per year (about 150 transactions per second), which are initiated through 5,000 point-of-sale devices and automated teller machines. In its data center, more than 30 high-end servers and 1,000 multiprocessor servers support not only customers and clients but also bank personnel, who use 25,000 desktops and laptops. The bank has more than 8 million retail customers using Web banking services.

Giving advice to a large number of banking customers about how they could better manage their money in current accounts might be an example of a noninteractive, high-volume analytic artifact. Supplying business clients with different investment scenarios would be an example of an interactive analytic artifact with low-volume but high practical potential.

It will take an intricate computing structure to create, transfer, and render enterprise analytics. Algorithms will need to capture precious expertise representing the most dense and valuable intellectual property. Our guess is that data center capacities and enterprise grids will be the fabrics necessary to cope with complexity, volume, and dynamics of enterprise analytic systems. In financial services, the number of events might be in the tens of billions; the number of transactions might reach 5 billion per year; and the number of static analytic artifacts might reach into the tens of millions per year. We could see intricate interactive analytics reaching hundreds of thousands of instances per year for business customers, representing the next generation of financial services that would gradually replace simple and increasingly outmoded reporting systems.

RESOURCES

Delic, Kemal A. and Umeshwar Dayal. "The Rise of The Intelligent Enterprise," *ACM Ubiquity*, vol. 3, no. 45.

Delic, Kemal A. "Enterprise Models, Strategic Transformations and Possible Solutions," *ACM Ubiquity*, vol. 3, no. 20.

Grigori, D., F. Casati, U. Dayal, and M.C. Shan. "Improving Business Process Quality through Exception Understanding, Prediction, and Prevention," *Proceedings of VLDB 2001*, Rome, 2001.

Han, J. and M. Kamber. *Data Mining: Concepts and Techniques*, Morgan Kaufmann, San Francisco, 2000.

Inmon, W.H. *Building the Data Warehouse*, Second Edition, John Wiley, 1996.

Chaudhuri, S. and U. Dayal. "An Overview of Data Warehousing and OLAP Technology," *ACM SIGMOD Record*, March 1997.

Quinn, James Brian. *Intelligent Enterprise*, New York, The Free Press, 1992.

KEMAL A. DELIC

[kemal_delic@hp.com] is a lab scientist with Hewlett-Packard's operations research and development and a senior enterprise architect with relevant experience in knowledge management, Bayesian nets modeling, and real-time intelligent systems.

UMESHWAR DAYAL

[umeshwar_dayal@hp.com] is principal laboratory scientist with HP Laboratories, where he leads a team performing research in data mining, knowledge management, and business process management.

may be delivered from the single analytic server, while others may require grid-level computing power. Typically, the computing power and resources engaged are inversely proportional to the volume of rendered artifacts. Some business intelligence analytics software vendors are already able to deliver tens of thousands of analytic artifacts to corporate employees' desktops.

The purpose of the enterprise analytic systems ranges widely, from providing holistic views of the enterprise state and subsystems to the tasks of optimization and forecasting. The earliest and most typical deployment was in identification and reduction of various inefficiencies leading to optimization of overall enterprise performance. Highly abstracted, enterprise architecture can be dissected into event, transaction, and analytic layers, each having different purposes, objectives, and design constraints (see Figure 2, page 30).

Looking at Figure 2, interactions with customers, clients, partners, and suppliers create a constant flow of business and IT events that should be accurately and timely forwarded to appropriate event-management applications. Within IT, the architecture must address IT events that will enable the monitoring and management of IT fabrics: network, routers, switches, servers, operating systems, applications, desktops, notebooks, and personal mobile platforms.

Transactional applications aggregate, transform, and capture a fraction of this endless event stream as transactions. Business processes are also a critical part of enterprise operations. They create the host of interrelated transaction records typically guided by the established policy or automatic execution of the business processes. Business processes represent the key business chores executed daily in each enterprise. Transaction persistence and recovery are necessary features of this middle layer; these features play a key role in billing and auditing processes, for example.

Analytic applications operate on amassed data and information sets; these implement sophisticated (parallel) algorithms on specialized data repositories, such as operational data stores, data marts, data warehouses, and knowledge bases with the purpose of delivering accurate analytic artifacts in the appropriate form-factor. Enterprise portals are rendering engines for different user populations. Enterprise analytic systems should support better decision-making across the range from daily operational decisions to strategic and disruptive changes.

Enterprise event, transaction, and analytic subsystems serve as the business communication, coordination, and cooperation conduit with markets and suppliers. They're created on heterogeneous platforms for which interoperability standards are now sorely absent. Web services are addressing the standards and interoperability needs in an innovative way, although the proliferation of informal standards and proprietary solutions continues. Analytic system functioning is based on advances in various computing domains, especially marked by developments in data mining and business intelligence.

TECHNOLOGY DEVELOPMENTS

The headwaters of most analytic technologies today are in the research of information theory and artificial intelligence, which has helped to create data mining, data warehousing, online analytic processing, and business intelligence. These technologies have combined with advances in hardware and improved software creation methods, leading to the creation of the impressive superstructures. The results have been huge, dynamic programs typically encountered in large-scale simulations and real-world modeling for weather prognosis, traffic simulations, and financial analysis.

As enterprise analytic systems move out of research, industry-specific requirements will begin to dictate specialization as well as new methods and approaches. (See the sidebar, "Financial Services: The Leading Edge," page 30) Advances will likely stimulate the creation of technologies that better fit the needs of enterprise analytics. We believe we'll see a distinctive analytic architecture develop, as we saw when the industry responded to the need for management of enterprise events and transactions. The new technology will support the set of analytic tools and methods to analyze the past, assess the present, and predict the future according to key enterprise business parameters. ie