Enabling E-Markets Intelligence with the Black Pearl Knowledge Broker

Black Pearl, Inc

Mike Rogers

May 17, 2000

BPWP0001

Contents

Tables

Figures

Executive Summary

Nigel Rayner of the Gartner Group recently identified that the Internet is transitioning into its third phase of development: the Transaction Phase (1998-2003), where online activities move beyond simple catalog aggregation and lowest-price interaction. In this more complex market, customer loyalty is critical in maintaining and growing market share. Bricks and mortar companies can use Black Pearl's Knowledge Broker to extend their expertise into e-markets, dramatically increasing the scale, scope, and effectiveness of person-to-person relationship offerings by providing intelligent, personalized recommendations.

Opportunity

The real battle in the B2B marketplace isn't for capital, or financial resources. The latest estimate of the domestic B2B market by Forrester Research puts it at $130 billion in 1999, and $1.5 trillion by 2003. Nobel laureate Herbert Simon realized long ago that the real battle is for the scarcest of tomorrow's resources: attention.

> What information consumes is rather obvious: it consumes the attention of its recipients. Hence a wealth of information creates a poverty of attention, and a need to allocate that attention efficiently among the over-abundance of information sources that might consume it.
>
> Herbert Simon, *Scientific American*, Sept 1995

To succeed in the next decade, businesses need to occupy the attention high ground. In an era of easy, instant commodification and efficient markets delivering instant price knowledge to all actors along the value chain, the competitive advantage lies with companies that can ensure that their name is presented first or singly on a screen of data, and that they become synonymous in people's minds with high-quality, efficient service. Black Pearl believes the way to secure this advantage is to project existing business knowledge and expertise into electronic marketplace transactions using its enterprise-scalable, intelligent platform for delivering personalized recommendations at the point of transaction: the Knowledge Broker.

All marketplaces bring together buyers and sellers for the exchange of goods and services. Because of the limits of geography, time, or space, physical marketplaces have long relied on intermediaries to aggregate and logically categorize sellers' offerings and buyers' requirements, providing customers with the right product at the right place and at the right time. However, these intermediaries brought several drawbacks in terms of increased transportation costs and their own peculiar segmentation demands of inventory and circumstance.

Physical marketplaces are full of friction that eats into potential profits. Electronic marketplaces promised to reduce or eliminate the costs associated with these intermediaries through the annihilation of time and distance, the distribution of perfect pricing knowledge, and a multiplicity of on-demand electronic channels. The ideal of frictionless capitalism is becoming a reality.

The current B2B marketplace is organized in a network of overlapping and interoperating e-markets, where corporate entities and actors exchange goods, services, money, and information. But the ongoing e-market disintermediation has created a space where vendors find themselves competing for scarce attention resources while coping with the requirements of a new array of electronic intermediaries and an assault from low-end first-mover dot-com catalog commoditizers.

These new intermediaries and dot-coms create offerings aligned along new, dynamic horizontal and vertical channels. Unfortunately, this ability to repackage, bundle, and unbundle to create best value offerings runs up against limits to the precious attention resources of buyers. The availability and complexity of offerings has expanded exponentially, while decision-making and selection are limited by human agency and the software and interface tools available. Current search and evaluation tools simply overwhelm buyers and sellers, reducing the savings and potential of e-markets.

Black Pearl provides a leading-edge technology platform that allows traditional bricks and mortar companies to transfer their vast investment in partner and client knowledge and compelling value propositions online. Black Pearl's Knowledge Broker can dramatically increase the efficiency and effectiveness of transactions. By taking into account buyer and seller goals, preferences, behaviors and the e-market context, the Black Pearl Knowledge Broker produces concise and directed recommendations that increase the likelihood a transaction will occur and enhance long-term customer satisfaction. An array of technologies and approaches makes this possible, some based on proven business practices, others on emerging methods of knowledge discovery and business logic creation.

Initially, this white paper examines the state of the emerging e-markets, identifying the new challenges posed and ways to turn these from challenges to opportunities. Finally, this paper describes the performance of the Black Pearl Knowledge Broker platform and illustrates a typical vertical implementation within a telecommunications environment.

State of the E-Market

To project their value into emerging e-markets, bricks and mortar companies are finally taking the e-market plunge. They are hybridizing into clicks and mortar ebusinesses. What are some of the new environments and hybridized models they are encountering in the new e-markets?

United Technologies is the parent company of, amongst others, Carrier air conditioners, Pratt & Whitney engines, and Otis elevators. It established an online auction site to coordinate its supply purchases. With hundreds of millions of dollars flowing through this hub, it estimates average savings of 17%.

Sales Transactions

The fixed-price sales model was a necessary and successful compromise for the sluggish production and marketing conditions that existed during the Industrial Revolution. But within emerging e-markets, the flow of physical material can be shifted away from the flows of service and configuration within distribution channels. Instead of the older and unresponsive "stovepipe" method, sellers can offer and buyers can evaluate and search for specific offerings at many different points throughout the value Web.

Buyers and sellers no longer overwhelmingly meet with simple "point-and-shoot" one-to-one transactions but instead are increasingly reliant on acquiring and combining assets along value a value Web. Combining and evaluating these disparate elements and negotiating and paying for these at the instant of sale requires complex logistics and decision-making. Many buyers and sellers, however, are unprepared or unwilling to engage in the complex decision-making required to take advantage of these new pricing models.

Trading Communities

Inflexible, brittle-syntax Electronic Data Interchange (EDI) has already made several trading hubs possible in the early phases of e-markets. Newer, more loosely coupled XML-based trading communities are now proliferating for procurement, liquidation, and fulfillment activities. They take advantage of many different price and arbitration models, from sealed-bid auctions, bartering, and Dutch bidding.

As these communities move away from simple lowest-price offerings and the number and activities of members increase, issues of community such as trust, leadership, opinion-makers, politeness, quid pro quo, and reputation brokering enter the mix. Current technologies fail to address these complex variable contexts.

Securities Trading

Online trading models are increasingly transforming broker/client relationships from an older, sluggish hierarchy model into remediated many-to-many relationships. The quality and quantity of products has expanded enormously, with more potential returns but also more varied and complex risks and dependencies. Securities sellers need to be able to dynamically calculate the cost of securities at point of contact, factor in market externalities, add some value with respect to a specific client and their preferences, and resell at profit. Again, current technologies lack the sophistication and breadth to accomplish this.

Securities companies are painfully aware that the domain knowledge of their financial consultants generally resides within individual heads, and a race is on to codify this knowledge to make it transportable, resistant to personnel change, and more easily distributed throughout the organization.

Bi-directional Lag-time Reduced

As a rule of thumb, about 20% by quantity of an organization's purchases constitutes some 80% of the total purchase value. Obviously, trimming the quantity of purchase decisions required could return significant dividends. It is also common to hear business buyers complain that they can expend 80% of their time on 20% of the total purchase value.

> Strategists must now worry about learning to filter information, make sense of it, and act on it faster than others.
>
> Andersen Consulting, *Five Rules of the E-Economy*

This 80:20 rule becomes a significant drag on business productivity, particularly when coping with a flood of data with new e-markets. Increasing the potential velocity of business decision making, allowing organizations to grab mindshare and market share ahead of competitors, is now a primary execution goal for businesses moving into e-markets.

Greater Flexibility To Build 1:1 Relationships

The old promise of data mining becomes even more seductive in e-markets, given the correct application of technology. Instead of coping with data repositories and outdated, static deductions, immediate access to transactions -- coupled with models of cyclic demands and forecasts on a per-customer basis -- can lead to strengthening of buyer/seller relationships. Reducing the quantity of non-value-added time expended per-transaction makes it possible to offer valued-added-time and more relevant choices to buyers, increasing the likelihood of transaction success.

The 80:20 rule also crops up in the area of profit generation. Faced with a rapidly expanding market and a need to ramp-up quickly to support a huge influx of customers, companies operating in e-markets require some kind of decision support to identify the 20% of high-value customers and prioritize on their fulfillment and satisfaction to drive profitability.

The key to understanding the effective operating requirements within e-markets is that better information is not always more information but more relevant information.

Ernst & Young maintains that businesses that cannot speed up decision-making to keep pace with partners risk losing those business partners to more nimble competitors. Lag-time is the hidden cost of doing business in physical markets and lagged decision-making attains higher visibility in e-markets. Reducing costs and improving efficiency requires reducing lag-time.

> The often-caricatured vision of information overload poses a very real limit to the economy
>
> Ernst & Young,
> Limits to the New Economy,
> *Perspectives on Business Innovation*

E-Markets Segmentation

The purpose of segmentation is to enable companies to offer precisely targeted brands and baskets of offerings to behaviorally or conceptually similar groups of customers. In e-markets, behaviors change rapidly and so can segmentation models – and doing this requires flexibility and dynamic intelligence on the part of sellers. But e-markets create an excellent opportunity for sellers to synchronize with the cognitive demands and proclivities of buyers. In less fluid, static markets, sellers and intermediaries often forced buyers to cope with products and services organized along sometimes overly convoluted and irrational segments that bore little relation to how buyers actually think about their needs. E-markets and intelligent technologies allow new information brokering intermediaries to cognitively align product and service offerings along conceptually similar buyer-goal-driven paths.

This is the manifest aim of portals, to provide an encompassing purchasing and relationship experience personalized to individual buyer's requirements and that flows and adjusts itself to their needs in a logical and consistent fashion. Most portals fail this elementary fitness test, producing a flood of non-relevant and tangential information and choices for users.

Sellers think in terms of product and service offerings, whereas buyers think in terms of activities. Customers think about investing for specific purposes, with set goals. The categorization and segmentation of varied financial products should be irrelevant to them; they want to be presented with logical choices that best meet their needs. Sellers need some way to construct meta-offerings from across the continuum of financial product market space that best meet the very specific, personalized goals of individual customers.

Targeted Marketing

Marketing can benefit enormously from improved intelligence gathering and execution possibilities in e-markets, not to mention reduced costs to access new channels. Lead generation and rapid re-segmentation to suit market dynamics can be automated. With network-mediated cooperation, channel partners can combine on lead management and create virtual consortia that can satisfy complex buyer requirements at various touch points within the value Web. But this can be gridlocked in infoglut. Ideally, what is required is some way to reduced the lead quantity and improve lead quality using a contextual, situated scoring system to enable fast offerings and a greater likelihood of positive decision making.

In e-markets, valuations are increasingly being driven by the depth and extent of customer relationships. As a lowest-common denominator retailer of commodity stocks, E*Trade has been very successful to date. Yet its search for a merger partner in the bricks and mortar world occupies a central part of its strategy execution, with the most recent potential partner,

American Express, jilting it on the way to the altar. Previous potential merger partners included Goldman Sachs and JP Morgan.

E*Trade needs to acquire experience of its customers accumulated along many of their lifetime touch points, rather than at the single end-channel touch point it now occupies. Traditional financial brokers have decades of experience about making recommendations to customers, huge knowledge repositories that E*Trade needs to access. The clicks-and-mortar approach lets traditional companies combine their expertise with new technologies to provide enhanced service offerings. Combining new technologies and translating old expertise and knowledge bases into new economy will provide an amazing engine of success for the new clicks and mortar companies emerging within e-markets.

Buyer Profiling

The synergy between clicks-and-mortar re-engineering of knowledge repositories and ongoing online profiling within e-markets creates huge data models of customer activities. The application of technologies analogous to data mining allows sellers to predict with reasonable certainty within which purchasing activity cluster buyers fall and at which significant life event they are located. Also, situational behavior alters along a chronological sequence of purchases and according to where a buyer is located along significant life or business events such as marriage or inventory liquidation. The key is calculating the "lifetime value" of customers, or LTV.

> Jeremy Rifkin, writing in *The Industry Standard*, calls this The Age of Access, when "shifting from discrete market transactions that are limited in time and space to establishing relationships that extend in an open-ended way over time" becomes essential to successful operation in e-markets.

Typical firms churn 10-30% of their customers annually. The aim within e-markets, where the switching cost for buyers is massively reduced, must be to control this churn by pre-selecting better long-term, high-value customers while retaining borderline customers.

Higher Value Relationships

Emerging e-markets are characterized by efficient arbitrage that eliminates or minimizes price differentials. This commoditizes many products and services. To survive and prosper in e-markets, successful companies will have to add value with targeted information content to products/services, customizing and personalizing before resale. Bricks-and-mortar companies possess this added intelligence and are ramping up to transfer their expertise to the electronic realm.

The "M" Word: Mobile Commerce

In a recent article in *Business 2.0*, Patricia Seybold elaborates on how crucial and transformative the coming revolution in mobile commerce, or m-commerce, will be. Rapid marketplace adoption rates and an accelerating take-up of mobile Internet access devices means that wireless online usage will quickly dwarf wired Internet connections and business models within a few short years. This has dramatic repercussions for how companies will conduct business. Customers and clients will access accounts and profiles using several devices at multiple touchpoints. Businesses will have to design for multiple customer interaction scenarios across varied interaction media. And at the same time, they must ensure that customers retain easy access to satisfactory interaction endpoints while providing a consistent and seamless customer experience.

Marketing and product offerings will become more situationally focused and context-sensitive because clients will not have much patience for long-winded or high-volume recommendations or offerings considering the small size of their display screens. Marketing and selling into mobile channels also increases the demand for language and regional-specific content. This must be accomplished transparently, with culturally specific interaction scenarios taking local laws, regulations, and availability into account.

> The complexities involved in providing customized, targeted content and offers will increase exponentially. The job of organizing, managing, and targeting your content and transactions for different market segments has become much more challenging … In order to be ready for this multi-touch point customer in a multilingual, multicultural world, you're going to need to invest heavily in dynamic content management and configuration tools.
>
> Patricia Seybold, Dial M for Commerce, *Business 2.0*, April 2000

Black Pearl's Knowledge Broker is the only enterprise-level, multiple touchpoint encompassing platform that adequately solves some of these problems of adapting to m-commerce. Using an advanced architectural design that combines a holistic, ontology-based model of a customer that provides encompassing, full-spectrum viewpoints, the Knowledge Broker delivers true intelligence for e-markets.

Current Technological Challenges

Many of the technologies implemented in the marketplace during the previous decade required large IT investments and significant business and enterprise restructuring, yet failed to deliver results commensurate with their cost.

Application Servers

These monolithic solutions, based on CORBA or other developer-focused solution technologies, create a buffer between those within enterprises with business knowledge and those tasked with its implementation. This distancing leads to unwieldy, centralized hubs that are unable to quickly respond to changing market dynamics or new knowledge with updated or reformulated business rules. The Black Pearl Knowledge Broker, by contrast, can pull business data from a wide variety of heterogeneous sources (including RDBMS, IMS or VSAM, flat files, and so on) and represent the discovered inter-relationships using easy-to-understand business rules that can be rapidly amended by business experts.

Collaborative Filtering is the Personalization Story

This worked reasonably well for the B2C market that characterized early e-market commerce. Unfortunately, the assumptions and interaction model do not scale well up to many-to-many B2B trading networks. Storefronts and wide-spectrum portal approaches are based on the 1-to-many model and also rely on static, forced segmentation derived from physical markets. Collaborative filtering and this all-or-nothing presentation do not allow for adaptive realignment of offerings along buyers' cognitive activity clusters.

Heavyweight OLTP Centralization

The investment by enterprises in Online Transaction Processing (OLTP) produced expensive and fragile systems that exacerbated the stovepipe problem. They were slow to respond to changing conditions and difficult for people with business process and customer experience to quickly devise, submit, and receive feedback about off-center "what-if" scenario requests.

EDI's Inflexibility

Electronic Data Interchange (EDI) is not real-time and this "prevents trading partners from engaging in the kind of collaborative, interactive analysis … required to finalize a contract" (*The Global Advance of Electronic Commerce*, The Aspen Institute). Its "brittle syntax" demands a difficult and costly pairwise connection process between partners, with narrow conduits constraining business process interoperability and responsiveness. Nevertheless, by introducing the notion of loose, process-oriented, document-based coupling between organizations, EDI paved the way for the modern ideas of ontological-based information exchange between enterprises.

ERP's Life Cycle

Before the emergence of e-markets, Enterprise Resource Planning (ERP) delivered significant returns during the slower era of largely EDI-based transactions. But ERP efforts and projects now face obsolescence in today's rapid, maturing environment.

Nigel Rayner of the Gartner Group recently identified an evolutionary process where the large-scale ERP efforts of the 1990s are encountering diminishing returns in a changed marketplace. He describes the ongoing evolution of online markets using four different phases. These are illustrated in Table 1 and Figure 1.

Phase	Period	Description
Presence Phase	1996 – 1999	During this initial phase, it was generally sufficient for companies to merely establish Internet visibility to attract customers.
Interaction Phase	1997 – 2000	During this early phase, companies merely needed to supply basic search functionality on basic corporate Web sites to attract customers.
Transaction Phase	1998 – 2003	During this current phase, companies need to create "stickiness" and community to attract and retain customers, who now have a huge selection of indistinguishable, unremarkable, and generic sites and portals from which to choose. Complete and complex business transactions are now being done using the Internet, but higher-value transactions and stickiness derive from adding unique personalization features.
Transformation Phase	2000 – 2005	As we enter this phase, the market is transforming into a "collaborative commerce" model (or c-commerce). Traditional ERP is becoming marginalized. The integration requirements of Customer Relationship Management (CRM) and Supply Chain Management (SCM) with ERP will lead to companies offering "over-arching software architectures" that tie together these systems. When factoring the distributed delivery of mobile commerce (m-commerce) into the market, the ability to deliver holistic, high-quality, relevant recommendations derived from a 360°-view of the customer moves beyond useful. It becomes compelling.

TABLE 1: INTERNET ECOMMERCE FOUR-PHASE EVOLUTION

The Black Pearl Knowledge Broker is a significant step forward in enterprise-wide technology integration that enables companies to compete during the Transaction Phase of e-markets and provides a scalable stepping-stone into tomorrow's Transformation Phase environment.

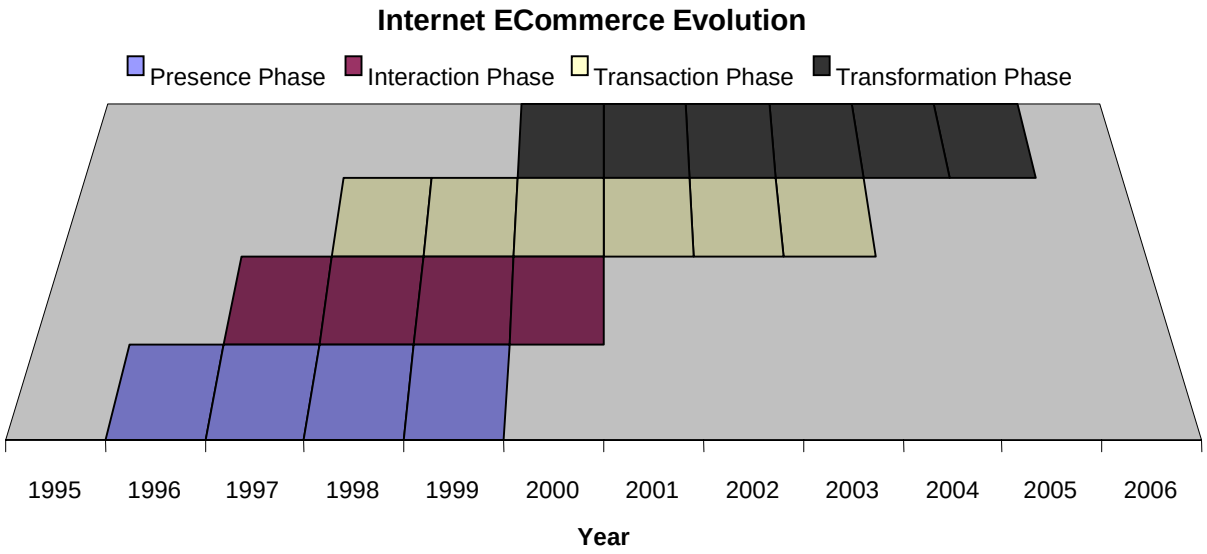

FIGURE 1: INTERNET E-COMMERCE EVOLUTION

E-Markets' Architectural Requirements

Any software solution that attempts to operate within the e-markets domain faces several significant architectural and technological hurdles and must meet certain requirements. Conventional enterprise-class software implementations (based on client-server and n-tier architectures) provide ineffective results in this domain. What are the functional criteria that successful enterprise-class e-market solutions must satisfy?

Persistence

An advanced enterprise-class system should be able to extract data and rules from RDBMS, VSAM, and IMS files, for example, without disturbing underlying data integrity or relationships while persisting its own system state. The Knowledge Broker accomplishes this through federating its datasources, with intelligent agents that traverse corporate data structures using JDBC technology. This data is cached within the Knowledge Broker and organized into an advanced, virtual data structure known as an ontology.

Reliability

An enterprise-class system should elegantly and cleanly handle congestion, time-outs, overloading, server fail-over, and data integrity. The Knowledge Broker accomplishes this by implementing its components as transactional – according to the atomic, consistent, integral, and durable (ACID) approach – through the use of the Java Transaction Service (JTS) technology.

Scalability

An enterprise-class system or portal should be able to handle huge numbers of components and concurrent requests, whose numbers can rapidly fluctuate across orders of magnitude. The only way to cope with this level of required instantiation and performance load balancing is through multiple servers capable of arbitrating resources between these components in a distributed fashion. To accomplish this, the Knowledge Broker's distributed server architecture uses state-of-the-art Java 2 Platform Enterprise Edition (J2EE) and Java Naming and Directory Interface (JNDI) enterprise technology.

Process Compatibility

An enterprise-class system should be able to easily adjust to dynamically changing business processes and operational process logic. Instead of the enterprise being forced to modify and constrain its processes to suit a sub-optimal system's hard-coded logic, optimal systems such as the Knowledge Broker completely separate the business logic from sometimes-idiosyncratic implementation logic. The Knowledge Broker allows those business specialists with greatest domain knowledge to directly create powerful knowledge bases and what-if scenarios using a simple-to-understand English-like interface. Its natural language interface enables Knowledge Broker users to directly and quickly tackle the hard problems of business without being distracted by deep-structure SQL syntax issues.

Discrete, Intelligent Functionality

An enterprise-class system should separate its logical functions into separate functional modules that can be individually accessed and managed. The Knowledge Broker's server functions can be distributed across multiple servers, and these servers can handle hundreds or thousands of client instances. Additionally, functionality within the Knowledge Broker is organized into logical modules, outlined in Table 2:

Logical Module	Description
Data Concept Editor	Manages database connectivity and data-concept mappings, such as *Stock Price.*
Business Concept Editor	Manages logical subcategories or combinations of data concepts that have business relevance, such as *Volatile Stock.*
Concept Relation Editor	Manages contextual relations between Business Concepts.
Rule Editor	Manages English-syntax rules such as *If Customer is High Risk then Recommend Buy Volatile Stock.*
Knowledge Discovery Center	Manages off-line data mining and what-if scenario construction using a simple, clean interface.

TABLE 2: KNOWLEDGE BROKER'S DISCRETE, INTELLIGENT FUNCTIONALITY

Adaptability

An enterprise-class system should support the many levels of hierarchy and networks of organization and responsibility that can exist within enterprises. The Knowledge Broker supports a hierarchy of rules representing, for example, legal rules (SEC- or EPA-mandates that must be obeyed), domain-specific (corporate policies and heuristics that should generally be obeyed unless countermanded by the issuing authority), and normal rules (individual hunches, corporate best practices, and so on that will be obeyed but can be broken if circumstances dictate).

Distributed Object Model

An enterprise-class system should present Application Programming Interfaces (APIs) with methods and events that allow object access to be distributed across multiple machines and networks. The Knowledge Broker uses JNDI and Enterprise Java Beans (EJB) services to make its components identifiable and accessible within and between enterprise networks. Its use of Java IIOP support also allows the Knowledge Broker to interoperate with industry-standard enterprise CORBA front- or back-office systems.

Web Support

An enterprise-class system should leverage Web technologies as B2B and inter-enterprise collaboration, value-chain management, and coopetition increasingly moves away from proprietary networks to TCP/IP and Web protocols. The Knowledge Broker's ontology (data, rules, performatives) storage format is XML, which makes for easy data interchange with new and legacy applications. The Knowledge Broker also features APIs allowing Java servlet, Java Server Page (JSP), and HTTP (the Web protocol) interfaces to its data inputs and outputs. This allows easy integration with common Web platforms and publishing systems, such as Vignette StoryServer[1].

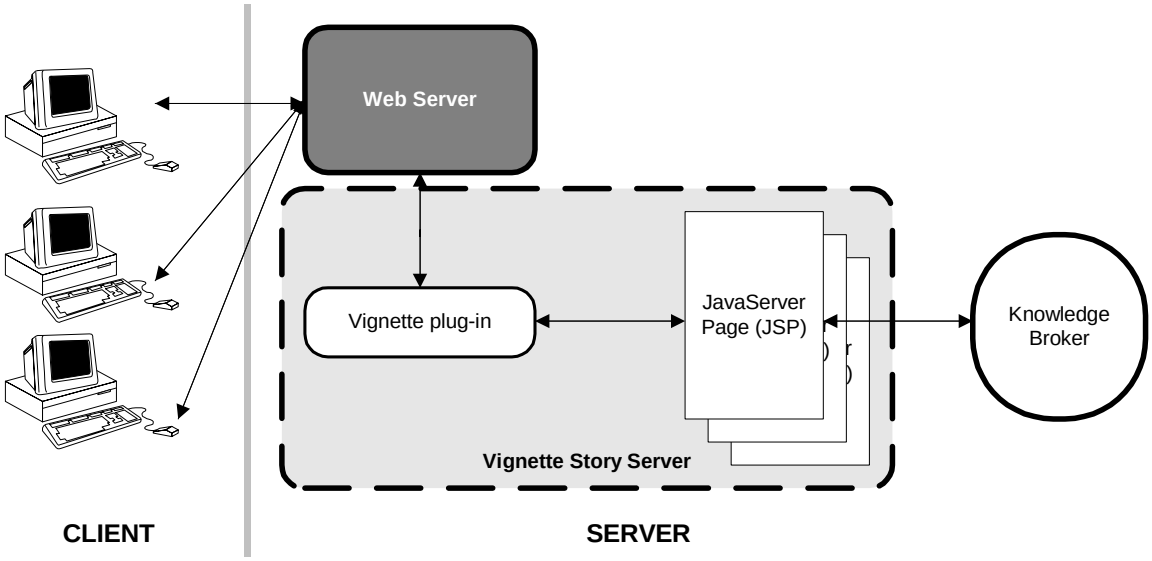

FIGURE 2: INTEGRATION BETWEEN KNOWLEDGE BROKER AND VIGNETTE STORYSERVER

[1] HTTP requests are sent from client browsers to an e-commerce Web server, which resolves requests for pages managed by Vignette through a plug-in. When an HTTP request is detected for a Java Server Page (JSP) that invokes Knowledge Broker functionality, the JSP page compiler compiles the JSP into a servlet, which passes the request to the Knowledge Broker. This returns a response to the servlet that finally returns an HTTP response to the client browsers.

Scaling for E-Markets

The only rational way to approach modern, transaction-phase e-market application development is using the n-tier, massively scalable, component-based, intelligent agent platform offered by the Knowledge Broker.

Client-Server Has Strict Limits

Conventional client-server architectures involve a table-database server and a client that uses application logic to interface directly with the server's tables, procedures, and presentation logic to channel the results to the user. The main drawbacks are that the client has no higher-level business object knowledge beyond the tabular data, maintaining application-specific data on the client is expensive, slow, and error-prone, and the implementation cannot scale well or restructure if the server becomes bottlenecked or the corporate organization changes.

Three-Tier Isn't Flexible

The three-tier, middleware approach of Enterprise Application Integration (EAI) and data warehousing strategies is an advance over client-server, but is still rooted in a pre-component, pre-Web paradigm. Placing the business logic into the application layer and retaining the presentation logic on the client increases maintainability, but the monolithic EAI systems themselves become a rigid bottleneck and limit the success and flexibility of Online Analytical Processing (OLAP). Scalability is costly and faces diminishing returns, and many EAI systems have limited multithreading models.

Knowledge Broker Can Do It

The Knowledge Broker's implementation of an n-tier, distributed and layered server platform stores the business logic in knowledge neighborhood clusters that can be configured to reflect corporate e-market demands. Knowledge Broker avoids bottlenecks by using intelligent, active agents to negotiate data access, load-balancing, and data federation. Enterprise interoperability (via EJB and JTS) also means that the Knowledge Broker's intelligent agents can access data and business logic in a highly mobile, autonomous, and distributed fashion. By separately partitioning its data collection, business logic, rules processing, and knowledge discovery among different, independent agents to facilitate multithreading and rapid instantiation, Knowledge Broker's performance scales elegantly from a single client and server right up to *n*-server and *m*-client implementations. Knowledge Broker's production of industry-standard XML and HTTP protocol output means that Knowledge Broker client development can leverage off Web standards, and presentation solutions can be created and modified cheaply, quickly, and efficiently.

Platform Solutions and Point Solutions

The Black Pearl Knowledge Broker platform encompasses several key technologies that provide a single-source e-market intelligence platform. This advanced technology platform provides a compelling application development environment that bricks and mortar businesses can use to extend the scale and scope of their expertise. They can create enhanced person-to-person transactions that identify and maintain high-value, long-lasting relationships.

The Knowledge Broker is not a point solution to a small-scale, departmental-level problem. Its potential solution domain is definitively enterprise-scale. Competing domain solutions occupy smaller areas of applicability, and do not scale well. The Knowledge Broker's problem domain encompasses and extends previous, isolated approaches while providing a coherent, unified approach to enterprise-scale problems relating to knowledge management and integration.

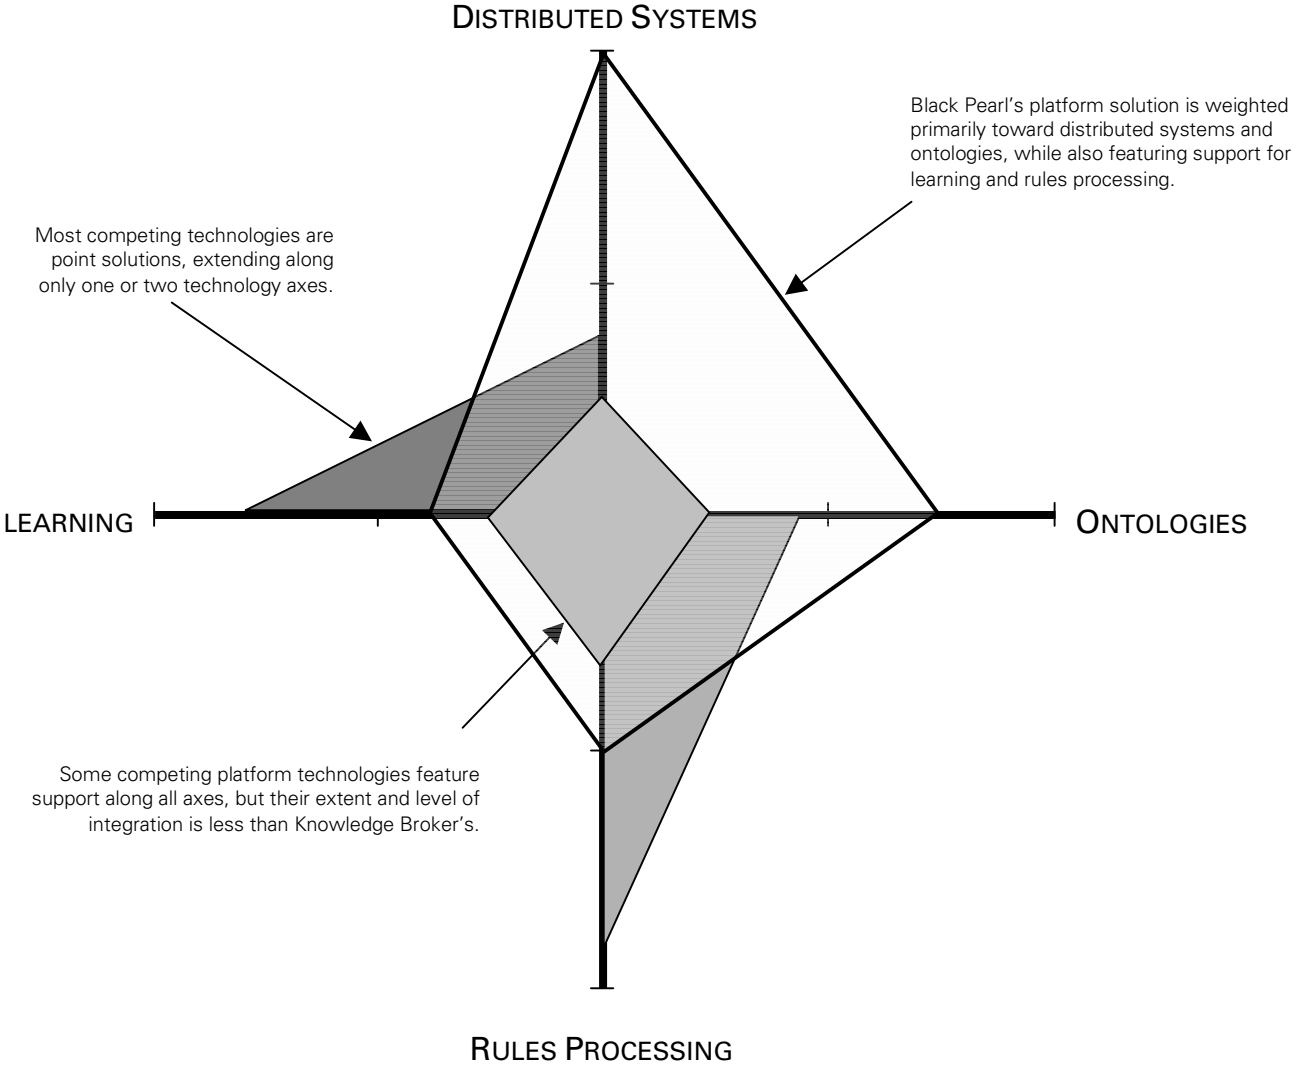

FIGURE 3: THE KNOWLEDGE BROKER IS A PLATFORM SOLUTION

What Knowledge Broker Is Not

- Knowledge Broker is not merely a real-time trading platform.
- Knowledge Broker is not merely a pure recommendations engine.
- Knowledge Broker is not merely an academic ontology or machine-learning project.
- Knowledge Broker is not merely a distributed systems solution.
- Knowledge Broker is not merely a pure rules engine.
- Knowledge Broker is not merely a catalog or content aggregator.

An Intelligent Agent Platform Solution

The Knowledge Broker is a technology for implementing an intelligent agent e-markets platform. It is built on a synergy of underlying, enabling technologies that unite in concert to move beyond current enterprise application solutions.

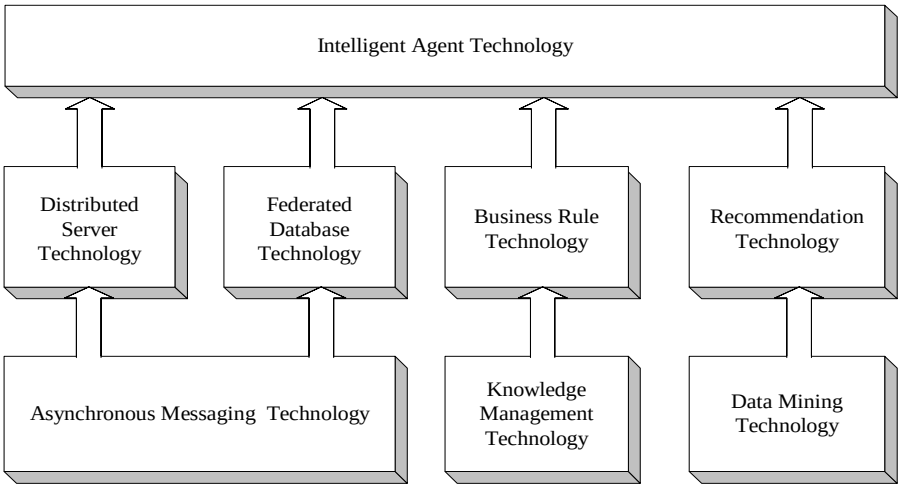

FIGURE 4: KNOWLEDGE BROKER'S INTELLIGENT AGENT ARCHITECTURE

Converging the Goals of Buyers and Sellers in Real-time E-Markets

Intelligent Commerce Requires Ontologies

One major advantage of the Black Pearl Knowledge Broker is its ability to capture and encapsulate business data, rules, and best practices within a flexible, standards-based ontology that provides for easy and consistent browsing and examination. It creates this powerful logical structure without modifying the underlying data or rules, avoiding expensive, IT-heavy, and error-prone data warehousing issues.

> The main purpose of an ontology is to enable communication between computer systems in a way that is independent of the individual system technologies, information architectures and application domain.
>
> What is an ontology?
> Ontology.Org

Ontologies are essential to enable real-time e-markets. Ontologies are unified data models that incorporate and reconcile disparate data sources and performative goals into a consistent, logical structure using standardized knowledge representation technologies that facilitate straightforward knowledge translation and goal transfer.

Facilitate Differentiation, Avoid Centralization

Ontologies are a qualitative advance beyond the simple catalog aggregation or product hierarchy offerings from most B2B or B2C vendors. These rigid and centralized offerings are divisive, often difficult to traverse, and impose artificial bundling and segmentation structures. Frequently, this rigid hub-like structure derives from these systems' reliance on EDI (or simple EDI-influenced XML semantics).

> XML syntax is designed for representing an encoded serialization, and thus has a very limited range of expression for modeling complex object semantics … Interoperable computing solutions imply the existence of a sharable ontology, or common set of object semantics … shared ontologies are not the domain of XML.
>
> Robin Cover,
> *XML and Semantic Transparency*,
> OASIS

Ontologies make it possible for sellers to represent and relationally integrate different activities and offerings to better represent the conceptual activities that buyers are engaging in at particular transaction touchpoints. The friction and cognitive load experienced by the buyer is reduced, and the likelihood of a successful or more profitable transaction is increased, coupled with greater long-term buyer satisfaction.

Broker the Goals of Buyers and Sellers in Real-Time

In a challenging transaction environment characterized by dynamic on-the-fly segment reconfigurations, technologies that intercede at the intersection between buyers and sellers must provide a rapid response trust infrastructure that is impartial and empowering without being intrusive, constraining, or skewed. Because of its scalability and advanced buyer/seller requirements reconciliation abilities, Black Pearl's Knowledge Broker meets these primary prerequisites for operation in demanding e-market environments. The need for this even-handed, massively scalable trusted architecture was recognized from the earliest days of online e-commerce.

> Electronic brokers will be required to permit even reasonably efficient levels and patterns of exchanges. Their ability to handle complex, albeit mechanical, transactions, to process millions of bits of information per second, and to act in a demonstrably even-handed fashion will be critical as this information market develops.
>
> Resnick, Zeckhauser, Avery, *Roles for Electronic Brokers*, MIT Center for Coordination Science

Assess Opportunities for Sellers and Ensure Transaction Success

While improving the quality and relevance of the choices made available for buyers, Black Pearl's Knowledge Broker emphatically maintains and enhances the best interests of the potential seller in the transaction. With the ability to intelligently encode best practices, winning strategies, market conditions, and in-house expertise into a compelling basket of recommendations, deploying the Knowledge Broker improves "first-call effectiveness", better satisfying customer needs on first contact (in terms of sales, support, queries, and so on). Organizations that deploy the Black Pearl Knowledge Broker transform into metamediaries, capable of operating intelligently at various touchpoints aggregated across horizontal and vertical domains.

> Metamediaries make metamarkets possible by assembling the products and services in a metamarket into a seamless, integrated bundle. They serve as a single point of contact between buyers and sellers in a metamarket, and greatly improve the efficiency and effectiveness of exchanges. At their heart, metamediaries are providers of trusted advice and information that customers need to make better decisions for a cluster of activities.
>
> Mohanbir Sawhney, Making New Markets, *Business 2.0*

By instantaneously incorporating variables such as future costings or revenue on a per-buyer basis, projected profitability, and so on, the Knowledge Broker lets clicks and mortar companies project their value into every transaction and break through into wider e-markets.

Knowledge Broker Architecture

The Knowledge Broker is a distributed reasoning server that organizes, filters and delivers information to users who are looking to gather meaningful information so that they can match customer interests with seller expertise at the moment of customer contact, and with full consideration of buyer and seller goals, preferences, behaviors and market context.

The Black Pearl Knowledge Broker works by allowing online Web users, knowledge workers, and executives to analyze and make decisions about relevant products, services and available assets.

The Knowledge Broker architecture permits rapid organization of multiple "what if" scenarios and allows the user to quickly change from viewing one potential course of action to another. Using intelligent agents distributed across an asynchronous, federated server architecture, the Black Pearl Knowledge Broker reads data as needed from existing enterprise databases and the Internet, then converts the data into intelligent recommendations through the application of business rules. The rules are constructed using ordinary terms already used by the enterprise, and they can be changed on-the-fly by a business user keeping maintenance costs low and responsiveness high. Users can switch between planning, geographical and text views of the information. The information used to determine decision-making comes from a variety of structured and unstructured data sources including the World Wide Web, relational databases, and VSAM or IMS files. All of the data, previously retrieved and tabulated manually, is now available and organized from a single Knowledge Broker interface.

Asynchronous Framework, Distributed Multi-Threaded Agents

The Black Pearl Knowledge Broker uses asynchronous communicating agents that facilitate a loosely coupling, flexible, dynamic transactional environment, and inter-application and inter-enterprise communication. This lightweight, loosely coupled trading model moves away from closed-system, CommerceOne-style hubs that require large IT infrastructure, RDBMS aggregation, and OLAP systems.

The Knowledge Broker automates decision making by representing knowledge and the explanations that support that knowledge through these multiple inter-communicating agents. Distinct, and often autonomous, Knowledge Broker agents are intelligent and perform within a system in response to messages from other applications or nodes.

The distributed multi-point communication model allows the trading and transaction functions to be dispersed within organizations, creating a federated, transactional space that can adjust to closely mirror to market structures and buyer segmentations.

Another key achievement is atomicity. Participating companies maintain control over distribution and organization of their internal data and translate or publish data only when dealing with other members of the value chain at various touchpoints. This better suits modern, rapid business processes and practices and allows companies to project their unique value into the chain instead of being commoditized or homogenized.

Advanced Prediction, Planning, Reasoning, and Learning Technologies

The Black Pearl Knowledge Broker integrates several key intelligent technologies within several clearly defined functional areas. These are accessed with an easy-to-use yet powerful and customizable GUI.

In almost any real-world application, an intelligent agent cannot be expected to have all the information needed to solve a problem. Therefore, the Knowledge Broker implements mobile agents using a variety of cutting-edge technologies to classify, predict and learn. Its constituent intelligent agents learn which external agents and datasources to trust, to cooperate with and to avoid based on prior experience and user preferences.

The Knowledge Broker's ontology consists of conceptual components that can be manipulated by users through an adaptable, easy-to-use Graphical User Interface (GUI). The components are shown in Table 3:

Components	Description
Datasource Connections	URLs, strings, or pathnames that describe disparate data sources.
Datasources	Disparate Web, relational database, and flat files containing data pertinent to business processes. These are organized component-wise within a tree-like structure.
Data Concepts	Data representation terms that map directly to one or more of the datasource components.
Business Concepts	More complex data representations, defined by users, that incorporate business knowledge relevant to particular industries or companies.
Concept Relations	Define the connections and relations between Business Concepts, and their interrelation.
Predictive Models	Encapsulate "what-if" scenario testing and allow for easy, rapid knowledge discovery and data exploration. They can be embedded or made to interact at will with other components of the ontology.
Rules	Allow users to define their business processes and experience using a simple, English-like syntax. Knowledge Broker users can relate Data Concepts, Business Concepts and Concept Relations, and Predictive Models with each other. Thus, the Knowledge Broker ontology conforms to existing business practices, instead of forcing enterprises to restructure their business processes to fit within rigid data schema.

TABLE 3: KNOWLEDGE BROKER ONTOLOGY COMPONENTS

The relation of these components within the powerful Knowledge Broker ontology is illustrated in Figure 5: The Knowledge Broker Ontology:

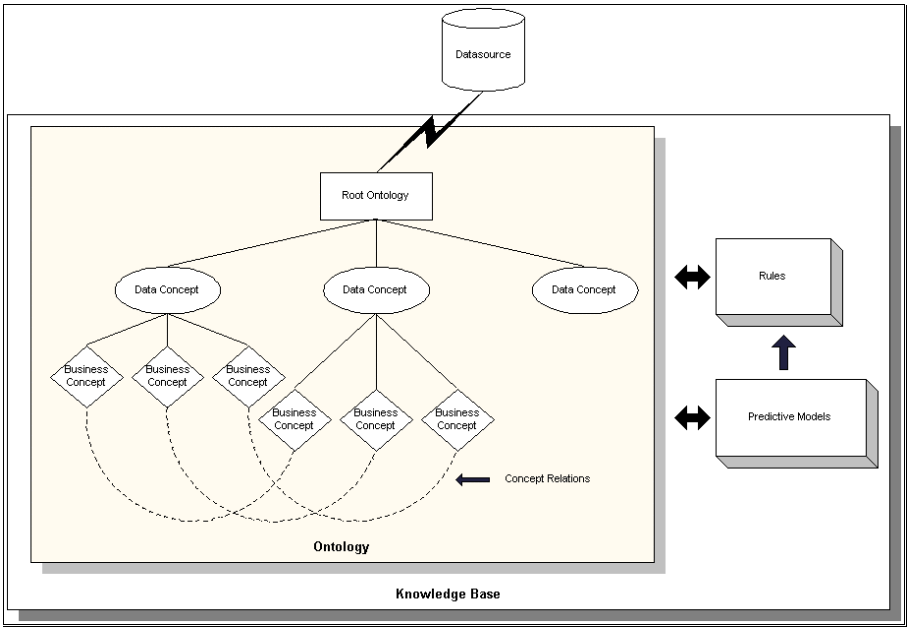

FIGURE 5: THE KNOWLEDGE BROKER ONTOLOGY

Federation and Industry-Standard Technologies

Distributing transactions across multiple constituencies in the value chain alleviates the expensive stovepiping of materials and data common with dumb logistics and EDI. Multi-point transactions empower local decision-making and relationship building, creating a "stickier" and more knowledgeable experience when dealing with buyers.

Black Pearl's Knowledge Broker uses server-side inferencing and supports a transaction-processing environment that can be federated across multiple servers, reducing load and easing bandwidth constraints. This also allows companies to establish transaction servers close (in Internet-space ping times) to trading "action" or negotiation hubs, as required.

The Black Pearl Knowledge Broker supports a wide array of industry-standard technologies to provide for easy integration and inter-operation with technology solutions and platforms from across the B2B landscape.

Flexible Data Model

Black Pearl's Knowledge Broker uses the XML content description metalanguage as a highly extensible, situation context-specific tagging format for data interchange and transformations. Adoption of this *lingua franca* throughout the B2B landscape is converting closed trading partner hubs into open, interoperable markets.

Distributed Java, Distributed Transactions

Black Pearl's Knowledge Broker supports enterprise interoperability and messaging using the industry standards Enterprise JavaBeans (EJB), Java Transaction Service (JTS), Java Message Service (JMS), and Java Naming and Directory Interface (JNDI). These are described in Table 4:

Java Technologies	Description
EJB	Support for the EJB component architecture specification ensures enterprise-wide session containerization and interoperability within development and operating environments.
JMS	Support for JMS enterprise messaging technology facilitates enterprise- and network-wide point-to-point and publish/subscribe messaging services.
JTS	Support for JTS object transaction technology delivers CORBA transaction services and Internet Inter-ORB Protocol (IIOP) interfaces. This provides a solid foundation for application context propagation, performance monitoring, measurement, and implementing Quality–of–Service (QoS) metrics.
JNDI	Support for JNDI provides access to a unified interface to multiple, heterogeneous naming and directory services throughout the enterprise. This provides the infrastructure for the Knowledge Broker's high scalability, load-balancing, and fail-over redundancy with minimal configuration overhead. This is accomplished through a naming service lookup that provides clients with logical view of server transaction space that is not tied to physical servers.

TABLE 4: DISTRIBUTED JAVA TECHNOLOGIES

KQML/KIF

Support for the Knowledge Query and Manipulation Language (KQML) facilitates data and performative goal interchange between network agents, allowing negotiation and synchronization of information about search, retrieval, subscribing, and offerings. Support for the Knowledge Interchange Format (KIF) ensures communication and negotiation between different e-markets agents using a common content language optimized for autonomous, asynchronous operation.

Black Pearl's Concerted Array of Fundamental Technologies

The Knowledge Broker's winning e-market capabilities and ontological approach are enabled by six key technologies that implement logical functions. These are performed by separate services (through agents) that can be individually managed, operated, and instantiated. These functions are:

1. Event Notification
2. Classification
3. Rules Processing
4. Distributed Transactions
5. Autonomous Intelligence
6. Learning

Event Notification

Event notification allows the Knowledge Broker to recognize that an event has occurred and dispatches an agent to take some action. This allows the Knowledge Broker to send information or complete a transaction via intelligent agents in an asynchronous fashion.

Classification

Classification is used to classify groups of customer profiles or portfolios into communities based on similar properties. After communities are formed, the Knowledge Broker can predict what products/services may be most appropriate to an individual given the collective interests of the community.

Rules Processing

Rules processing produces actions and discovers new knowledge. A key implementation goal for the Knowledge Broker is that the recommendations for next logical products or actions should not be a "black box" to an online buyer or seller. Therefore, simple English-like rules produce recommended products or actions during a transaction. A rule explanation capability is also useful, so that Knowledge Broker users can monitor decisions within the market context to see what rules produce what key effects.

Distributed Transactions

Distributed transactions cope well with information and business vocabularies spread across multiple constituencies within a sales chain. This maps well onto sales transactions, where buyers and sellers are negotiating and committing to terms within a network of inter-connected product and customer data sources.

Autonomous Intelligence

Autonomous intelligence allows mobile software programs to carry with them an itinerary of transaction data (buy, sell, trade, and so on) and is critical to a complex, electronic market. When an agent carries an optimized state to negotiate a best buy or best sell, the buyer or seller in the transaction can negotiate the best terms exploiting multiple channels. Knowledge Broker's intelligent software agents implement autonomous Java-based Remote Method Invocation (RMI) and other industry-standard methods of marshalling business process data and performative requests "over the wire" for remote processing and fulfillment.

Learning

The Knowledge Broker learns through combining its technologies into a best-of-breed intelligent platform. Each of these technologies plays a distinct role in the Knowledge Broker's ability to provide recommendations at e-market transaction touchpoints. When profiling buyers during active e-commerce negotiations and activities, the Knowledge Broker guides clients through a series of moves to achieve the best outcomes despite market volatility. Most transactional applications fall short of this by only pushing, rather than interpreting, the latest information or preferences. The Knowledge Broker's intelligent and flexible architecture enables real-time customer profiling against the current situation.

Some of the learning technologies additionally deployed to enhance the real-time learning and recommendation ability of the Knowledge Broker are described in Table 5:

Technology	Description
Search	The Knowledge Broker facilitates searching using natural language, an optimal interface for human-computer interaction when dealing with complex problems in the business logic domains. The ability to express and query business rules in simple English empowers those with the most comprehensive business knowledge to directly engage with the system's recommendations and reasoning and get concise, meaningful results quickly.
Reasoning	The Knowledge Broker incorporates rules technology that allows agents to reason about their environment and plan actions that lead to a goal (say, recommending buying or selling a product). The rules processing environment is extremely fast, efficient, and has proven scalability up to enterprise-level requirements of thousands of rules-per-second. With support for incorporating legacy rulesets and business rules knowledge discovered within heterogeneous enterprise-wide sources by the data federation agents, the Knowledge Broker can deliver real-time, appropriate, and intelligent responses to complex and/or rapidly changing market conditions.
Prediction	The purpose of the Knowledge Broker is to intervene at a touchpoint while the contact exists and make real-time recommendations. Its Knowledge Discovery Center implements simple data mining and predictive model building in a supporting, offline role where real-time demands are not critical. Using a straightforward and intuitive interface, business users can test out "what-if?" scenarios and discover new customer classification schemes. For real-time utilization, users can also embed successful predictive models within the Knowledge Broker's Reasoning component, or use proven knowledge technologies to deduce new classifications and business rules for possible incorporation.
Optimization	The Knowledge Broker can integrate with custom-built or industry-standard optimizing modules, such as Black-Scholes or Markowitz. This makes for easy integration with enterprise-specific optimization and operational strategies and technologies.
Planning	The Knowledge Broker can be used to support the concept of specific and personalized basket and spread scenarios for individual customers and can provide a dynamic recommendation strategy to maintain or enhance these personal baskets in response to changing market conditions.

TABLE 5: KNOWLEDGE BROKER'S ADVANCED LEARNING TECHNOLOGIES

The Knowledge Broker GUI

The Knowledge Broker features a fully featured GUI designed for ease-of-use and rapid information retrieval, categorization, and business rule entry by individual users. The Knowledge Broker GUI dynamically adjusts depending on the primary user task being performed. Although Knowledge Broker users can utilize custom-designed forms, applications, and so on, the GUI will be the primary access interface for most people. The GUI's six functional modules are the:

1. Launcher
2. Data Concept Editor
3. Business Concept Editor
4. Concept Relation Editor
5. Rule Editor
6. Knowledge Discovery Center

The GUI is designed to allow users to easily manipulate the key components of the Knowledge Broker ontology (see Table 3: Knowledge Broker Ontology Components on page 24)

The Launcher

The Launcher is the container GUI. It enables users to access the other functional modules, set preferences, and load and save changes to the ontology.

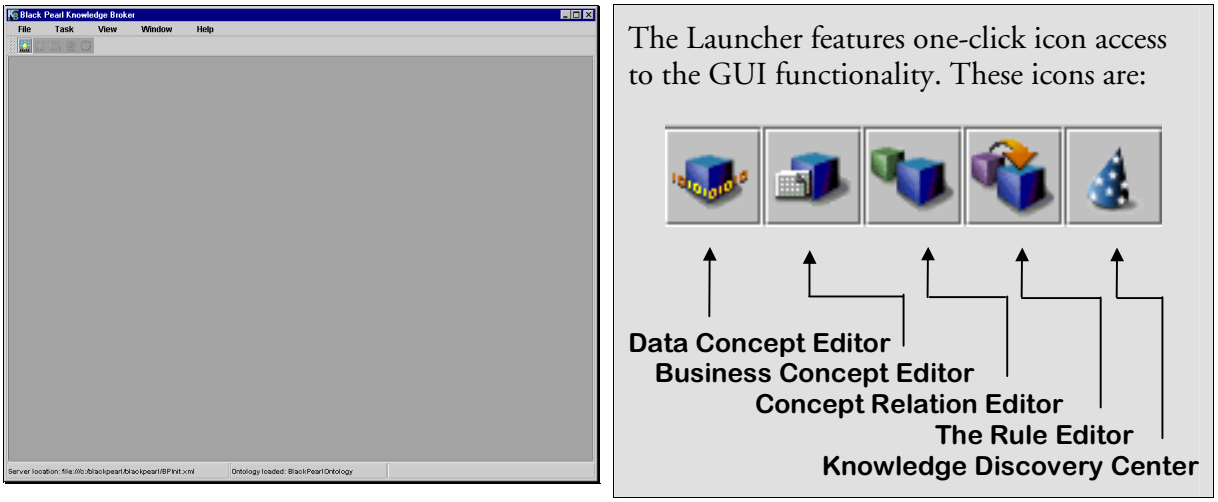

FIGURE 6: THE LAUNCHER AND THE KNOWLEDGE BROKER MODULE ICONS

The Data Concept Editor

The Data Concept Editor enables users to attach to datasources, identify data concept items and associate these with business concepts using simple or complex (star, cascade) mapping.

For example, the request below attaches a datasource connection to a particular Oracle schema, and analyzes all tables and columns:

```
jdbc:oracle:thin:@bpd1:1521:orcl/PARKER/*/*
```

A Knowledge Broker data discovery agent then extracts an Equity table, with column entries such as Beta, Industry, Sector, CUSIP, Code, and so on, and these become datasource components within the ontology. Business users can choose to extract and map only those components that interest them onto data concepts, or alternatively use automatic data concept generation. They can also perform complex calculations on the datasource components and embed these results within the ontology.

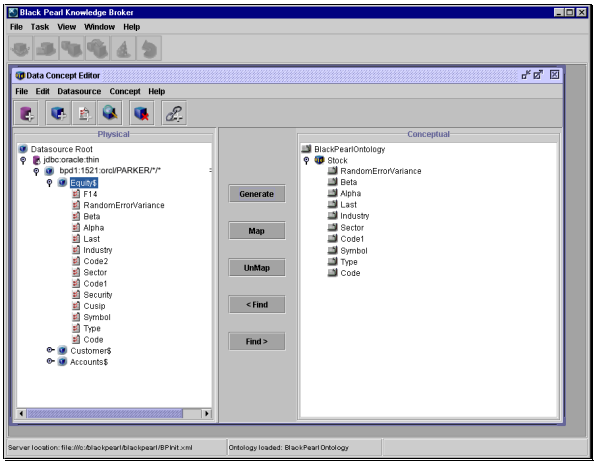

FIGURE 7: THE DATA CONCEPT EDITOR

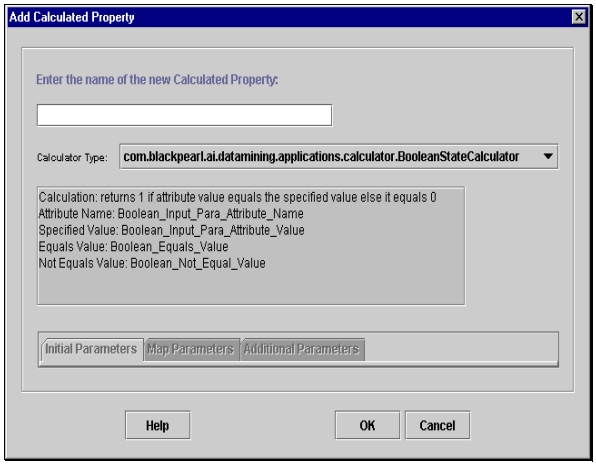

FIGURE 8: ADDING CALCULATED PROPERTIES

The Business Concept Editor

The Business Concept Editor allows business users to create business concepts by applying filters to one or more data concepts or other business concepts. This powerful facility allows users to create, say, a Volatility Stock category within the Stock data concept by specifying only those objects with specific *Beta* (or other) characteristics. By applying a different filter on, for example, *Liquidity Preferences*, users can also categorize customers as High, Moderate, or Low Risk. The data is being reshaped within the ontology to more closely reflect business experience and processes.

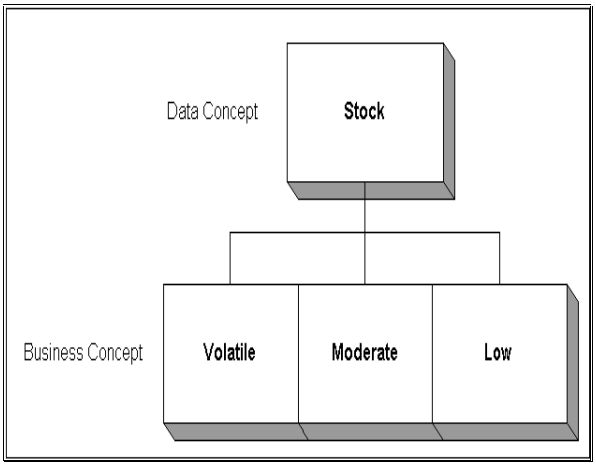

FIGURE 9: VOLATILE STOCK CATEGORIZATIONS

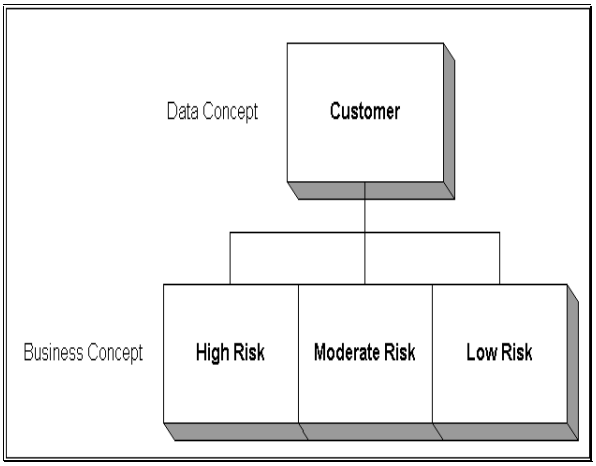

FIGURE 10: CUSTOMER RISK PREFERENCES

This reshaping is easily accomplished in the ontology with point-and-click simplicity, using drop-down selections and user-entered values to select only those data instances within the ontology that satisfy desired criteria.

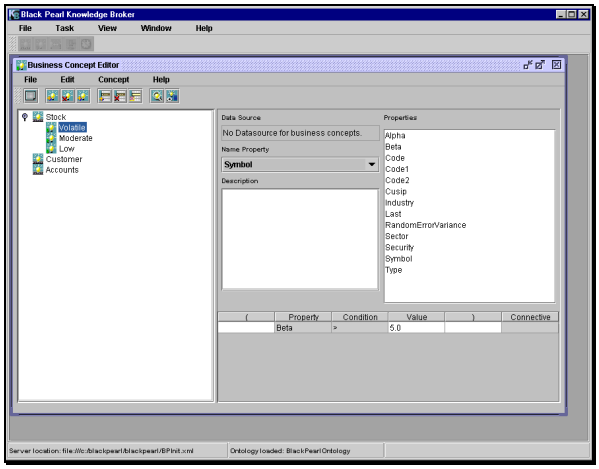

FIGURE 11: THE BUSINESS CONCEPT EDITOR

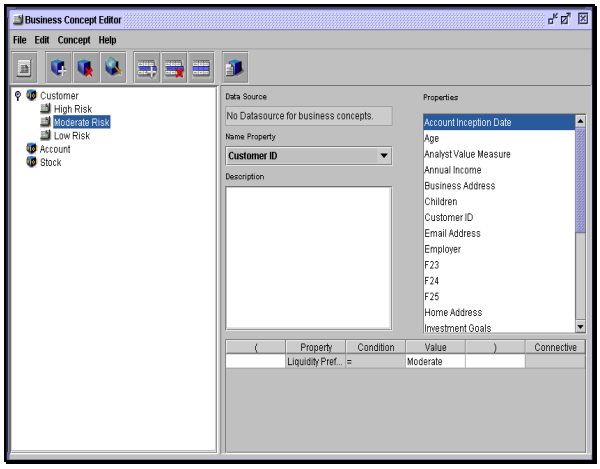

FIGURE 12: CREATING CUSTOMER CATEGORIES

The Concept Relation Editor

The Concept Relation Editor enables you to declare that a single business concept has multiple meanings, dependent on the context. For example, if you have Volatile Stock business concept, you could define a concept relation that enables Volatile Stock to be on List A if the context is a High Risk Customer, and on List B if the context is a Low Risk Customer.

Defining a concept relation enables you to avoid writing a rule to determine whether the Customer prefers Low Risk or High Risk. By retaining this contextual knowledge, the ontology minimizes the number of rules needed to model business problems and processes and avoids the combinatorial rule explosion common in many non-ontological business systems.

The Concept Relation Editor enables you to declare that a single business concept has multiple meanings, dependent on the context. For example, if you have a Popular Selections business concept (whose parent concept is the Music data concept), you could define a concept relation that enables a particular Popular Selection to appear on one list (for example, List A) if the context is an Older Customer. However, if the context is a Teenage Customer, then the Popular Selection concept will be adapted accordingly and appear on a different list (for example, List B).

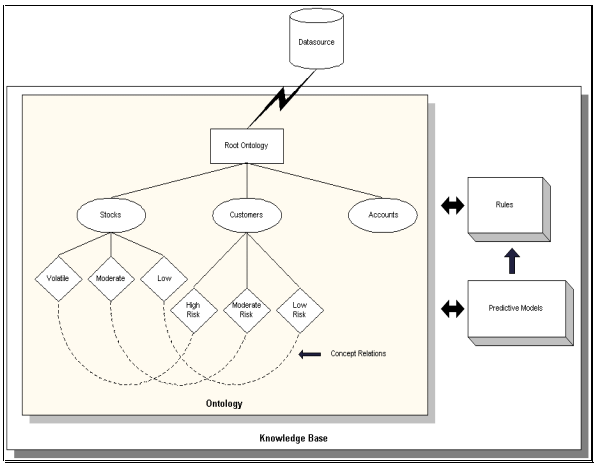

FIGURE 13: CONCEPT RELATIONS IN THE ONTOLOGY

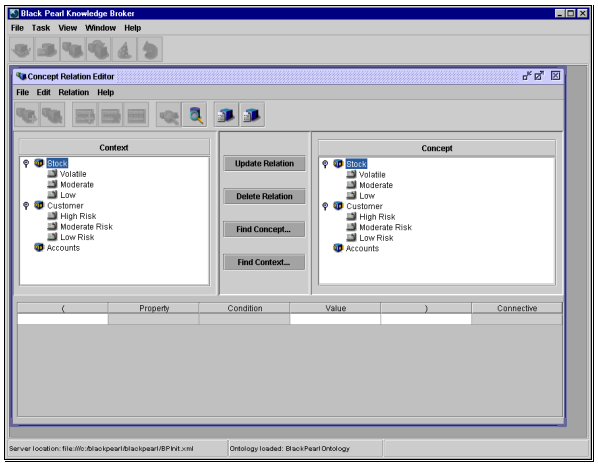

FIGURE 14: THE CONCEPT RELATION EDITOR

The Rule Editor

The Rule Editor enables users to define rules that describe two things: how your business typically responds to individual transactions and what situations will require different responses to seemingly similar transactions.

For example, you may want to define a process where you recommend volatile stocks to customers who tolerate high risk, moderate stocks to customers who prefer medium risk, and low volatility stocks to customers with a low risk tolerance. The can be easily performed within the Knowledge Broker GUI using straightforward point-and-clicking within drop-down selections.

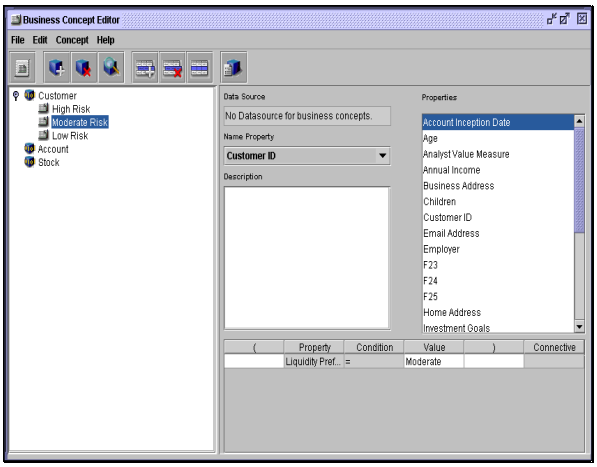

FIGURE 15: THE RULE EDITOR

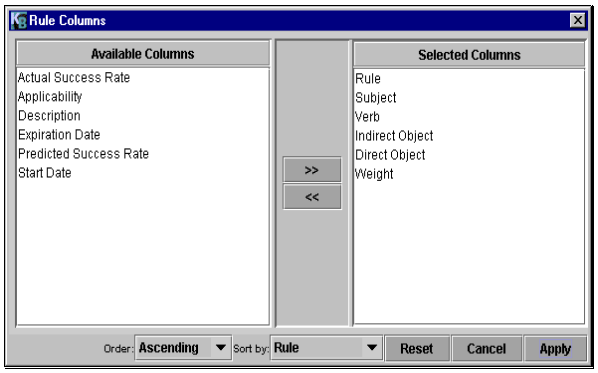

FIGURE 16: SPECIFYING APPLICABLE RULE PROPERTIES

The Knowledge Discovery Center

The Knowledge Discovery Center helps business users to create models of the real world by incorporating relevant data and business Concepts and analyzing historical data to determine patterns, relationships, and trends that can be used to predict future values and generate real-time recommendations. The models use datasets that can include customer demographics and histories, marketing and demographic information, corporate transactions and database information—for example, credit bureau records.

For a trivial example, a business user might want to determine the likely tax liability of a customer so that personalized tax advice can be offered at the point of interaction. The most likely input variable to consider would be Annual Income, to see how that affects Tax Bracket. The business user can enter the Knowledge Discovery Center and avail of its many easy-to-use analytical tools that facilitate such explorations. The business users can use the GUI to identify specific attributes they want to study further, and analyze these initially for clues as to behavior or characteristics.

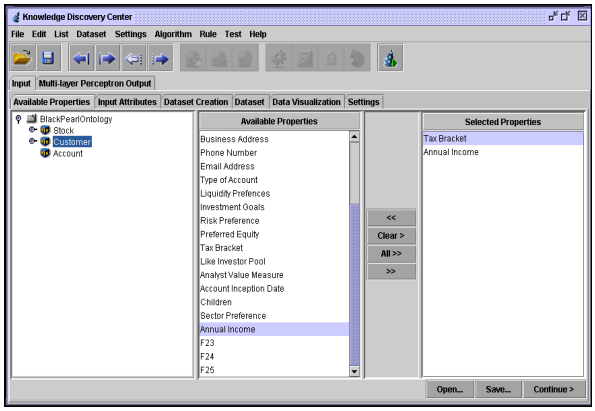

FIGURE 17: THE KNOWLEDGE DISCOVERY CENTER

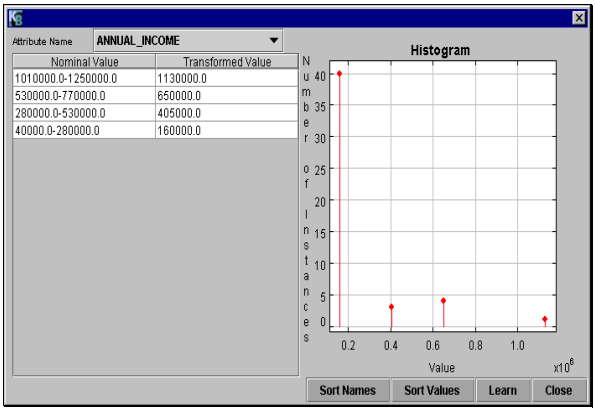

FIGURE 18: ANALYZING CUSTOMER'S ANNUAL INCOME

Now the business user easily selects which attributes they think are primary, and which are dependent.

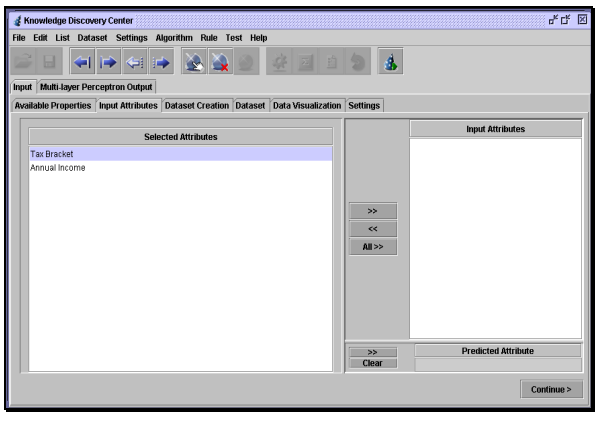

FIGURE 19: SELECTING THE MODEL'S ATTRIBUTES

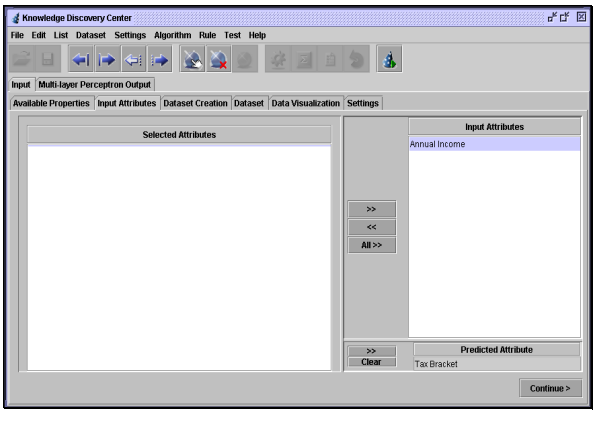

FIGURE 20: IDENTIFYING THE PRIMARY ATTRIBUTE

Business users can generate a mathematical model that can be easily embedded within the ontology as a calculated property or a rule. This allows complex scenarios and difficult calculations that would run too slowly in "real time" to be moved "offline" and tested and refined carefully. Following this testing period, the optimized models can be re-incorporated within the real-time Knowledge Broker for instantaneous results.

Alternatively, an easy-to-use and transparent interface guides business users through the creation of a decision tree that can discover business rules applicable to the problem that can be incorporated by the user within the ontology. The suitability of either model depends on the business problem being analyzed and the preferences of the user.

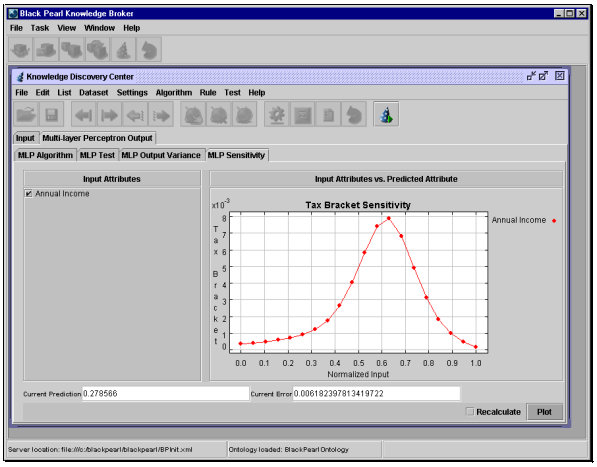

FIGURE 21: MLP'S MATHEMATICAL-BASED OUTPUT

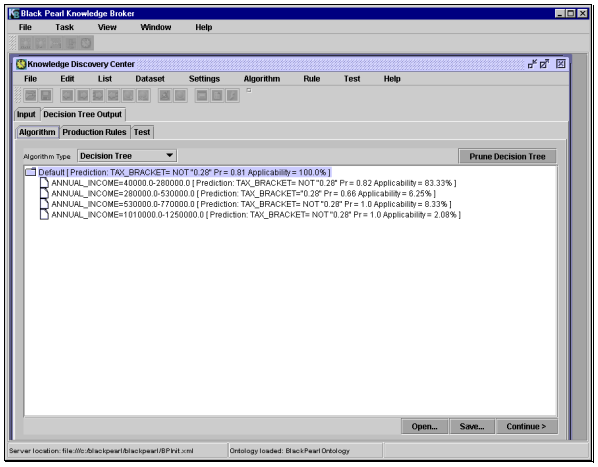

FIGURE 22: DECISION TREE'S RULES-BASED OUTPUT

The Knowledge Broker in Action

There are a huge number of vertical markets where the Knowledge Broker's typical output of a small set of easily digested, dynamic, high quality information and decision recommendations can increase the efficiency of execution, the likelihood of transaction success, upselling, profitability and client satisfaction.

- Telecommunications services (provisioning, consumer retailing, and so on).
- Financial services (including insurance brokerage, investing, real estate, and so on).
- Logistics (supply chain management, scheduling, purchase orders, and so on).
- Healthcare (provisioning, pharmaceuticals management, managed health services, and so on).

The Knowledge Broker can be integrated within a vertical B2B portal site, or developed in-house to assist brokers or agents in a decision-support and real-time assistance role. Knowledge Broker users can utilize the GUI directly, or interact with the Knowledge Broker using APIs, custom Web/client entry forms, and other defined and documented entry points.

A typical interaction with a system powered by the Knowledge Broker interaction might develop as follows.

The Financial Plan Generator (FPG) uses the Knowledge Broker to identify those investment products and services that best fit client needs. This is accomplished by matching investment products and services with corresponding client profiles. The Knowledge Broker provides the technological infrastructure for two integrated modules within the FPG:

- The Profiler, a data collection module.
- The Selector module, a rules-based product recommendation engine.

Where incomplete client profile data exists, the FPG uses the Knowledge Broker's ability to embed and extrapolate information and decisions from previously calculated and tested predictive models to establish high probability product and services recommendations.

In Table 6 on page 38, a financial consultant captures client response to the recommended financial plan and enters the information through the FPG Profiler, updating client profile records that are stored in the Knowledge Broker's ontology. The client's responses are integrated within the ontology by the Knowledge Broker and used to enable the FPG to learn over time to make more intelligent and successful recommendations for both full and partial client profiles.

Client Action	FGP System Action
A client initiates a call, requesting financial services. The financial consultant establishes a client connection to the FPG by, for example, logging on to a Web page.	The FPG Web page asks the financial consultant to establish their identity for security reasons, then requests client identification details.
The financial consultant enters suitable client identification information.	The Knowledge Broker-enabled Profiler attempts to return client profile data input screen to the FPG. Existing client profile information (including past recommendations, if any) are retrieved from the ontology and presented. If this is a new profile, the FPG alerts the financial consultant that this appears to be a "new" client.
If this is a *new* client, the financial consultant collects and enters the required client profile data. If this is an *existing* client, the financial consultant updates the client profile as required by the Profiler.	The Profiler displays a client interaction and activity screen. The financial consultant may edit and/or save the client profile information
The financial consultant confirms the profile data with the client and submits this identification and profile to the FPG for client instance retrieval.	The Profiler completes the write to the client profile record in the Knowledge Broker's ontology. The FPG presents a Generate Financial Plan/Exit action box (Exit terminates the session, with only client profile data updated).
The financial consultant clicks the Generate Financial Plan button and sees an extensive, annotated, and hot-linked financial plan.	The FPG invokes the Selector functionality provided by the Knowledge Broker. Accessing previously provided Asset Allocation user preferences; the Selector uses custom modeling to determine this particular client's risk tolerance and recommended asset allocation. Then, accessing pre-determined products and services, client profile information and demographics, stored client preferences, market dynamics data, the Selector returns a basket of recommended products and allocations. The FPG produces an extensively annotated, hotlinked financial plan that includes risk tolerance, asset allocation, and recommended products and services. This is also recorded in the client profile in the ontology.
The financial consultant can click on a concise description and explanation of every item (and its applicability) in the on-screen display.	The FPG uses personalization and publishing links to display required financial information, breakdowns, and metrics to the financial consultant.
The financial consultant discusses any or all of the proposed items with the client. The financial consultant tags the client's explicit choices (acceptance or rejection), optionally completing explanation fields that encapsulate client comments.	The FPG passes the client responses to the Knowledge Broker for incorporation in the specific client preferences instance in the ontology. Through extension, this client's preferences improve the general accuracy of the Knowledge Broker through their incorporation within the more general rules and predictive models contained in the ontology.
The financial consultant ends the session with the client, closing their client session with the FPG. The client's choices may have already been executed, or are now pending.	The FPG closes this client session and waits for another financial consultant to initiate a client session.

TABLE 6: A TYPICAL INTERACTION WITH THE KNOWLEDGE BROKER

The Black Pearl Alliance Program

The Black Pearl Alliance Program seeks to deliver a robust set of application and systems integration partnerships to provide the industry's most comprehensive e-intelligence solution. The partnerships that are being established range from platform applications to content management and up to CRM solutions. These ISV partnerships benefit from leveraging the Black Pearl platform's distributed data access and rules-based interaction management to provide a consistent and real-time 360° view of the customer experience. Figure 23 provides a view of the types of application partnerships being pursued by Black Pearl (several companies are highlighted to provide additional information):

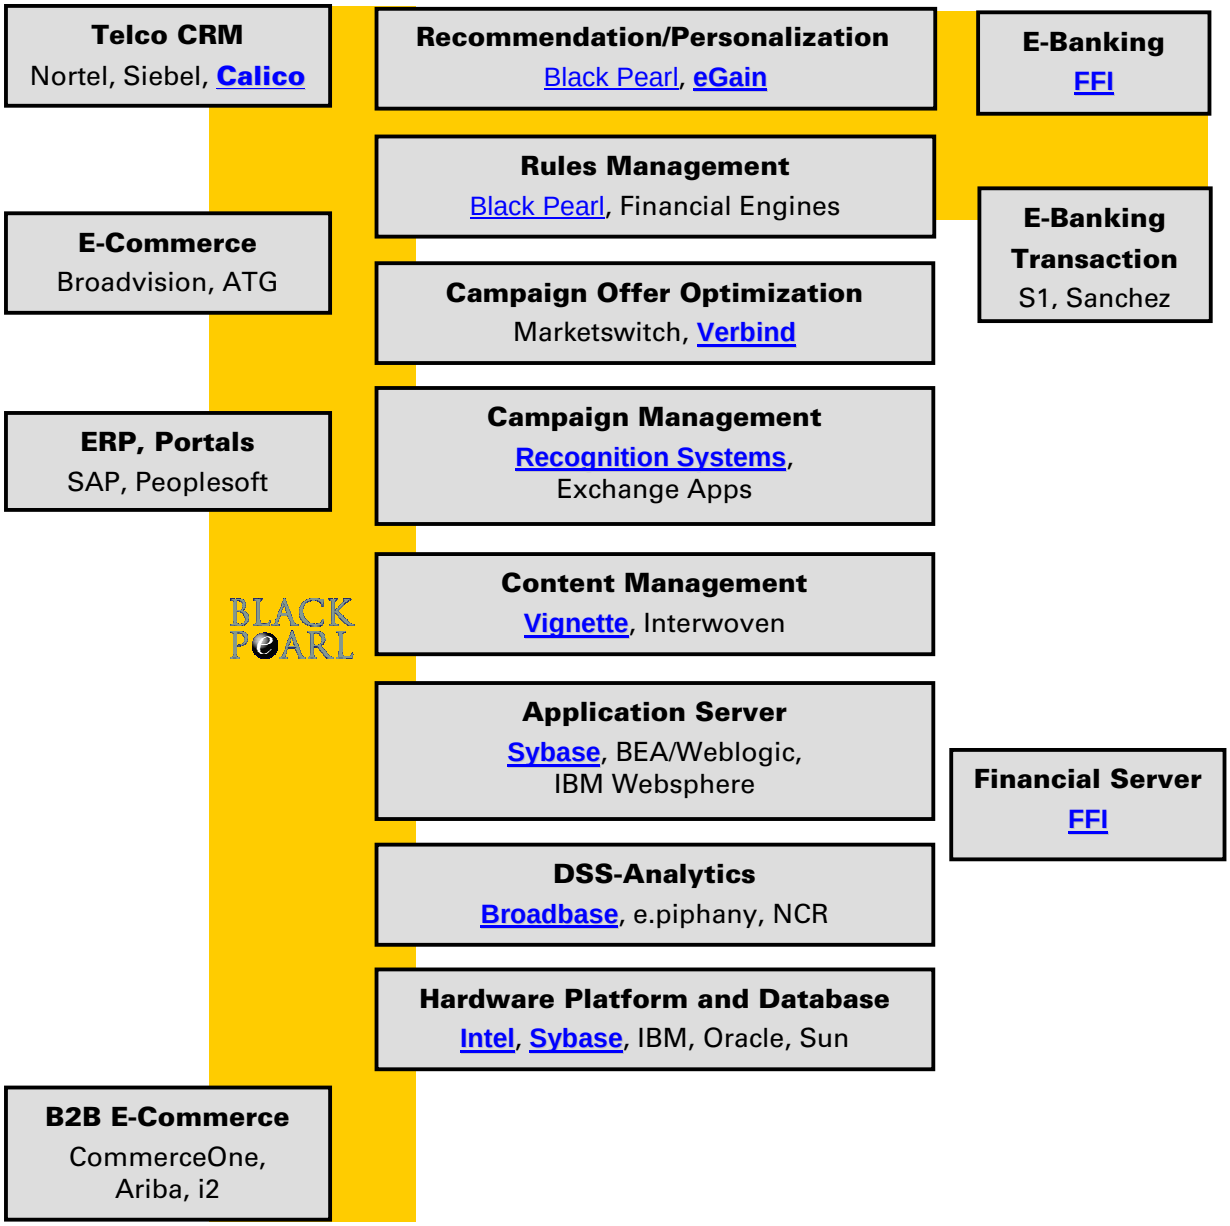

FIGURE 23: APPLICATION PARTNERSHIPS & BLACK PEARL ALLIANCES

Contact Information

If you would like to explore the possibility of working with Black Pearl to deliver the richest set of e-markets solutions for financial services, telecommunications and business exchange market segments, you can contact our Alliances Organization by e-mailing alliances@blackpearl.com. If you want to evaluate the Black Pearl Knowledge Broker for implementation, please contact sales@blackpearl.com.