The DPSC Electronic Marketspace:
The Impact of Intelligent Agents, the Internet and the Web on Electronic Commerce and Logistics.

by
Dr. Larry Kerschberg
Project Manager
and
Ms. Sonali Banerjee
Graduate Research Assistant

Submitted to American Management Systems, Inc.
Contract Deliverable

Knowledge Rover Laboratory for Logistics and Commerce
Center for Information Systems Integration and Evolution
Department of Information and Software Systems Engineering

George Mason University
MSN 4A4, CISE
4400 University Drive, Fairfax, Virginia

The DPSC Electronic Marketspace:
The Impact of Intelligent Agents, the Internet and the Web on Electronic Commerce and Logistics.

by

Dr. Larry Kerschberg

and

Ms. Sonali Banerjee

Center for Information Systems Integration and Evolution

George Mason University

Table of Contents

1. INTRODUCTION 3

1.1 PIONEERING APPLICATIONS 3
1.2 TRADITIONAL ELECTRONIC DATA INTERCHANGE SOFTWARE 5

2. ELECTRONIC COMMERCE ON THE INTERNET AND WEB 5

2.1 ELECTRONIC DATA INTERCHANGE TOOLS ON THE INTERNET 5
2.2 SECURITY ON THE INTERNET 6
2.3 SECURE COMMERCE ON THE INTERNET 7
2.4 INTERACTIVE TRANSACTION MANAGEMENT 7
2.5 ELECTRONIC CATALOGS AND TRANSPORTATION HANDLERS 8

3. SOFTWARE AGENTS FOR ELECTRONIC COMMERCE 9

3.1 AGENTS FOR ELECTRONIC COMMERCE 9
3.2 AGENT CONSTRUCTION KITS 9

4. KNOWLEDGE ROVERS 10

4.1 ENTERPRISE DATA AND INFORMATION REQUIREMENTS 11
4.2 ENTERPRISE ACTIVE SERVICES 13
4.3 KNOWLEDGE ROVER ARCHITECTURE 16

5. THE DPSC ELECTRONIC MARKETSPACE 20

5.1 A DESIGN OF THE ELECTRONIC MARKET SPACE 20
5.2 MARKETSPACE EXECUTIVE 20
5.3 ORDER REQUEST BROKER 21
5.4 FINANCIAL BROKER 22
5.5 QUALITY & PERFORMANCE EVALUATION AGENT 22
5.6 CUSTOMER SCENARIOS 24
5.7 INFORMATION FLOW 26

6. CONCLUSIONS 28

7. REFERENCES 29

Executive Summary

This document provides an overview of current research, development, and deployment of tools and systems for electronic commerce conducted over the Internet and the World Wide Web (WWW).

The paper begins by describing a number of innovative applications, including such industry leaders as General Electric and FedEx, who wish to use the Internet and the Web to link customers with suppliers through E-commerce. Credit card companies such as VISA and MASTERCARD wish to effect E-payments on the Web, and their efforts are discussed. Traditional forms of contracts and payments, such as Electronic Data Interchange, are now beginning to interoperate with the Web.

Next we present a discussion of the security and integrity issues involved in conducting Electronic Commerce over the Internet and WWW. It highlights the efforts of standards bodies, banks and other organizations to make E-commerce over the Internet a safe and trusted way of conducting business transactions.

We also discuss the role of intelligent agents in E-commerce, and review the DARPA-funded GMU research into Knowledge Rovers — a family of configurable software agents — and how they might be used to support the information and E-commerce needs of the Defense Personnel Support Center (DPSC).

Finally, we make a preliminary proposal for an agent-based approach to the DPSC Electronic Marketspace, and discuss the major agencies and their constituent agents.

1. Introduction

The rapid strides made in computer science and information technology has enabled business organizations to reach new customers and open new markets in an electronic marketspace. Electronic commerce, or E-commerce, is a more cost-effective way of doing business than building new stores or warehouses and hiring more people to run them. Electronic commerce today has several components:

1. interactive business and financial transactions in a virtual marketspace;
2. electronic cataloguing;
3. electronic order-tracking services;
4. automatic billing and payment services;
5. electronic transfer of funds;
6. vendor registration and electronic "brand naming";
7. automatic ordering, contracting and procurement by laying down plans and business rules which can then guide the contracting and procurement system;
8. advertising of products;
9. the data mining of consumer information for customer profiling;
10. customization of advertisements to have a more effective impact on prospective customers.

A fully-automated demand-driven inventory management and just-in-time delivery system can be achieved by integrating these components with guaranteed overnight (or on-time) delivery. The advent of the Internet, its popularity and accessibility, has radically increased the number of organizations doing business electronically and has brought about increased competition and made global commerce a reality.

1.1 Pioneering Applications

Some organizations that have either been pioneers in traditional electronic commerce or have provided new and innovative ideas for electronic commerce tools, services or standards are mentioned in this section.

The Sabre Group has developed the Sabre network for travel agents to create a complete flight (primarily on American Airlines), car and hotel reservation, including an online catalog of theater tickets, flowers, bon voyage gifts, etc. The Personal Computer (PC) counterpart, Easy Sabre, extends these booking facilities to customers. This network set the standard for similar services, such as United Connection, that allow PC users to reserve and purchase United Airlines tickets, book with hotels and car rental companies, etc.

The Allegiance Healthcare Corp. has developed an on-call EDI turnkey package called Analytical System for Automated Purchasing (ASAP) through which customers may purchase supplies using touch-tone telephones and receive electronic acknowledgments. Another package, ValueLink, an EDI-enabled stockless inventory management program, delivers hospital supplies on a just-in-time basis. Allegiance is now trying to develop a system that will help replenish supplies at customer locations (hospitals, nursing homes,

etc.) automatically by intelligently capturing information such as patient surgery schedules, hospital stock levels, etc.

Bell Atlantic reduced invoice processing costs through the use of an Evaluated Receipts and Settlement (ERS) program that eliminates supplier invoices from the vendor payment process [74].

Walmart uses data mining systems for information on buyers' habits at its different stores. "Smart" warehouses and logistic operations allow customer preferences to "pull" the goods in demand.

The Benetton Group of clothing stores and their partners, suppliers and outlets use EDI technology to obtain up-to date information of inventory levels, changes in fashions, and customer demands, that are met with robotic warehouses using just-in-time delivery system. Fruit of the Loom is another clothing company that uses the Internet to reduce inventory and provide personalized service by constantly monitoring information available from its distributors.

Virtual Vineyards [79] an online company on the Web sells wine and food products and provides information and expertise about that product. They are a Web-based virtual company using CyberCash Inc.'s secure financial transactions [21] and the Federal Express delivery system to conduct business transactions.

The Amazon.com Books Inc., is a virtual bookstore [2] with about 1.1 million titles where customers can post book reviews and authors can post self-conducted interviews. Readers can search by author, title and subject keywords or browse via subject links. When a reader makes a selection, titles of similar books are flashed. Amazon does not order most books until the customer enters shipping information and the sale is made, thereby reducing its inventory to a few best sellers although any number of books are available for browsing.

The General Electric (GE) Trading Process Network (TPN) [31], automates purchasing by GE. The GE Supply Net is a web-based ordering and order-tracking service which provides a front end to GE Supply's XPD 2000 order-entry system, part of the unit's just-in-time inventory and distribution system for overnight delivery.

The General Electric company is investing resources in electronic commerce [90]. The world's largest diversified manufacturer is accelerating its electronic commerce strategy with the goal of making the Web permeate its dealings with suppliers and customers. GE Capital Services, the company's $32.7 billion financial services unit, has created a consulting practice that will provide electronic-commerce services to GE Capital units and eventually to outside customers. GE will soon add seven business units to a Web-based automated procurement system.

The General Electric Trading Process Network on the WWW, http://www.tpn.geis.com/, automates purchasing by GE's electrical products distribution unit in Cleveland. Since joining the net in mid-1996, this business unit has reduced its purchasing cycle in half to seven days. Their web site allows potential suppliers to register. There is software available so that a potential supplier can create catalogs and interface to GE systems. In an expanding bid to exploit its considerable electronic commerce expertise, GE will also

release software to meld traditional EDI with the Internet. *This approach may be a model for the DPSC Electronic Marketspace.*

1.2 Traditional Electronic Data Interchange Software

Traditionally, Electronic Data Interchange (EDI) has been used to conduct online business-to-business commerce using Value Added Networks (VANs). EDI comprises the standards and services needed to perform electronic transactions such as purchase orders. However, EDI on VANs lacks the ability to do spontaneous online transactions. Also connectivity and proliferation is limited to other users of the EDI format and software, which runs on specified workstations, whereas the Web may be accessed from any browser-equipped PC.

EDI enables different computers to exchange transactions using standard ANSI (American National Standard Institute) format X12 series which specifies templates and standard fields for purchase orders, shipping documents, invoices, payments and hundreds of other data transactions. Typically, EDI software extracts flat files from mainframe financial or order processing systems and translates the data into EDI standard forms that are then transmitted as large batches of EDI forms over a dedicated phone line or a VAN. VANs ease the job of connecting with the many trading partners by collecting forms in an electronic mailbox, sorting, translating and forwarding them to recipients with the guarantee that they reach their destinations intact. EDI standards and technologies are dependable and since the communication is over private lines that have limited public exposure, they are relatively secure. Batch transfer of flat files efficiently move large volumes of data.

However, EDI software needs to synchronize protocols with all the partners' systems and the X12 format is rigid and complex, making new applications difficult to implement. Although the batch transfer process moves large chunks of data efficiently, it is a time-delayed process since any message generated by the transaction is stored in the mailbox until the next processing step, thereby making real time processing difficult if not impossible.

2. Electronic Commerce on the Internet and Web

Several organizations have used the Internet to publish and maintain electronic catalogs, to automate their transactions and to establish customer-business or business-to-business links. Most of these products are EDI-based and processed in a synchronized batch mode. The Internet has had a tremendous impact on electronic commerce because it enables continuous and easy contact with clients and suppliers. In this respect, it is a significant improvement over traditional EDI based on VANs.

2.1 Electronic Data Interchange Tools on the Internet

On the Internet, products such as Premenos's Templar can be used to send EDI transactions. Templar puts transactions into encrypted envelopes and electronically seals them so that only parties with the encryption key can view the contents of the envelopes that are sent securely over the Internet. Public key encryption ensures confidentiality, authentication, data integrity and non-repudiation of origin and return. However, all partners must install the same software and coordinate their upgrade cycles.

Actra Business Systems will soon release Business Gateway to convert conventional EDI software into Internet format such as the SSL encryption protocol.

The General Electric Information System (GEIS) has developed a system to meld traditional EDI with the Internet so that companies with EDI software can do business with companies on the web that do not have EDI. GEIS recently released GE Trade WEB, that lets companies that do not have conventional EDI software use Web browsers to communicate with EDI software.

2.2 Security on the Internet

There are confidentiality and security issues that must be addressed before the Internet can be used in large-scale automated business applications. Securing the Web and Internet has been a major roadblock in the advancement of electronic commerce.

As more companies conduct their business over the Internet, the required infrastructure and standards are being established. Commerce Net [17] is a consortium of 200 companies and organizations dedicated to promoting Internet commerce. It is an association to exchange logistics information and establish standards for electronic commerce on the Internet. They have proposed the Secure HTTP (S-HTTP) protocol to provide confidentiality, data integrity and authentication for business transactions.

The Internet has built firewalls, to protect individual networks from attacks by hackers, but the TCP/IP protocol suite used by all computers connected to the Internet is fundamentally lacking in security services at the lower layers of the protocol stack, that is, within TCP, IP and transmission protocols such as Ethernet. This allows such problems as eavesdropping, password “sniffing,” data modification, spoofing and repudiation to occur. To secure the Internet, the protocol architecture of IP includes an Authentication Header (AH) which provides authenticity and integrity using the MD5 [51] message-digest algorithm and Encapsulating Security Payload (ESP) [3], which provides confidentiality using the Data Encryption Standard (DES) algorithm [50].

A number of session-layer protocols have been proposed within the Internet Engineering Task Force (IETF) to support distribution of keys for use with almost any TCP/IP application. The most significant of these are the Simple Key Exchange Protocol (SKEP), Photuris and the Internet Security Association and Key Management Protocol (ISAKMP). The IETF has proposed Privacy-Enhanced Mail (PEM) and MIME Object Security Services (MOSS) to provide application-layer security [7].

The World Wide Web (WWW) [5] consortium currently has provided an HTTP extension designed to identify different payment types over the Internet using the Protocol Extension Protocol (PEP). In its Joint Electronic Payment Initiative (JEPI), it intends to create a framework for interoperability among various payment methods. JEPI would use the Universal Payment Preamble, a proposed negotiation standard for different payment methods such as Secure Electronic Transaction (SET) enabled software. The JEPI, by using the Universal Payment Preamble, would enable transactions using multiple protocols to interoperate over the Internet.

McAfee’s NetCrypto safeguards all types of TCP/IP transmissions and supports common encryption methodologies. The software manages encryption keys automatically, thus

eliminating the overhead of key management. However, secure hosts must be defined ahead of time. Also it does not provide sender authentication and encrypted transmissions can be destroyed, even if not read.

Netscape's Secure Socket Layer (SSL) which is embedded in Netscape Navigator or Microsoft's Private Communication Technology (PCT) may be used to build secure commerce systems. A combination of public key and symmetric cryptosystems provide confidentiality, data integrity and authentication security services just above the TCP layer.

2.3 Secure Commerce on the Internet

Visa and Mastercard software engineers have developed the Secure Electronic Transaction (SET) standard for securing credit card purchases over open networks like the Internet.

Smart card vendors such as Spyrus Inc. and Schlumberger Electronic Transactions Inc. will produce SET compatible cards beginning in Spring 1997.

First Virtual is a merchant bank established to facilitate secure credit card transactions on the Internet [39]. InfoHaus is First Virtual Holding's "virtual mall." First Virtual Holdings provides a payment system for Internet users based on email callbacks [10]. It's payment system uses the existing Internet protocols and instead of using cryptography, a high-level protocol is used. The protocol involves looking up the Personal Identification Number (PIN) in its database and finding the email address of the payer. An email message is then sent asking the payer to confirm the commitment to pay with a "yes", "no" or "fraud". Only on receipt of a "yes" reply is the financial transaction actually initiated. Since the Virtual PIN is useless off the Internet and requires email confirmation on the Internet, transactions are unaffected by simple attacks such as "sniffing." The valuable financial token, such as credit card numbers, account information, etc. , never appear on the Internet messages but are linked to the Virtual PIN after being retrieved from the database.

This is similar to the security measure suggested by AMS for the DPSC system wherein supplies are matched to contracts already in place.

However, instead of email verification, the confirmation is made against the database content.

Even though the Internet is an open medium and inherently is not secure, there are several tools available to provide secure transactions over the Internet.

2.4 Interactive Transaction Management

Today, there are several commercially available tools for Internet transactions, and there are service providers which specialize in Internet commerce. These interactive order management systems provide an Internet commerce "shell" or "store-front" which can be linked to a billing and payment system.

Open Market [62] offers commerce-oriented Web software including a commerce oriented web server to enable companies to conduct secure electronic commerce over the Internet. Their "distributed shared services architecture" separates the management of a

business transaction from that of its content. Companies can securely and centrally manage business transactions using content on multiple distributed Web servers. Open Market's architecture is scalable, flexible, offers robust security and supports open standards. Various companies including Time Warner, BBN, MCI, First Union, AMP, and Tribune use products such as OM-Transact, OM-Axcess, OM-SecureLink and ActiveCommerce DB.

OM-Transact is an application which enables companies to secure payment, complete order management, and provide online customer service over the Internet. It is a tool that may be embedded in applications built by Commerce Service Providers (CSPs). OM-Transact is built on an open systems approach. It supports all HTTP browsers, including those that support security protocols SSL or PCT.

OM-Transact software automates the actual mechanics of: 1) placing the order from the WEB browser, 2) transferring that order from a Web server to the vendors' order-entry and processing system, 3) tracking state, local and federal taxes, 4) calculating shipping charges, and 5) interfacing with accounting systems. OM-Transact on one or more servers, takes orders and other information through a firewall that connects to the Web.

A module of OM-Transact, Front Host, deals with the actual transaction, including taxes, postal issues, etc. The Front Host connects through a second firewall to an even more secure area, the Back Host that handles the actual financial transactions. OM-Axcess is an access management and reporting application designed to provide a simple and cost-effective way of building user authentication, authorization and session tracking.

By centrally managing the authentication and authorization of end users, OM-Axcess provides single sign-on access to proprietary information distributed across the Web. OM-Axcess can be used to provide information in a controlled environment, forge business partner relationships and collaborate in a distributed work environments. OM-SecureLink is a tool used to integrate catalogs and other Web contents with OM-Transact and OM-Axcess applications.

Internet store-front providers such as BBN and AT&T, host stores handling everything up to billing and charge on a per-transaction basis. Both these providers are themselves Open Market customers.

Connect Inc. [18] supplies Internet-based interactive commerce and order management software products. AMP Inc. launched AMP eMerce Internet Solutions, a service provider to businesses over the Web. Microsoft's Merchant Server is an NT product which provides its clients web commerce facilities for online retailing. BroadVision's One-To-One, Mercantec's SoftCart, Netscape's Merchant, IBM's Net.Commerce and Oracle's Merchant System are other Internet commerce servers.

2.5 Electronic Catalogs and Transportation Handlers

The infrastructure to support other aspects of the change in technology is currently provided by several commercial applications that we will briefly review.

Most organizations today have their web page on the World Wide Web (WWW) which provides up-to-date information about the company and its products. Catalog software developers like iCat Corp. [13] and the Vision Factory Inc. [80], have helped streamline

the process of publishing online catalogs, and feature additional capabilities for tracking and personalizing catalog sales.

Federal Express has a logistics, electronic commerce and catalog division that has developed Powership, an automated shipping system that allows organizations to place, track and bill their own FedEx orders in-house. FedEx-Ship software can be downloaded to desktops to schedule pickups, track and confirm deliveries and print bar-code labels. Inter-Netship is an Internet based service that allows customers to prepare shipping labels from the Web [26] without additional software or help from FedEx customer representatives.

United Parcel Service (UPS) and APL Ltd. are some transportation and freight companies with Web sites that allow shippers to track the status of packages and provide accurate shipment information.

3. Software Agents for Electronic Commerce

The advantage of using software agents for building Electronic Commerce frameworks and applications is that they are composable and configurable so that the application system for a particular enterprise can be built quickly. The agent system would provide the infrastructure needed for a seamless connection of enterprises and would serve as an "engine" or "shell" for transactions between databases involved in the different commerce applications.

3.1 Agents for Electronic Commerce

Some software agent systems that have been applied to the electronic commerce field are mentioned:

Kasbah, developed at MIT by Chavez & Maes [40] is a virtual marketplace on the Web, where users create autonomous agents to buy and sell goods on their behalf.

Broad Area Syndicated Information System (BASIS), developed by Genesereth at Stanford University [4] is an agent that is currently used in two application areas. CommerceNet supports the electronics industry in the Silicon Valley area with catalogues and information services, while the Electronic University is a virtual advisor to students on courses for their curriculum, projects, etc.

Mobile Unstructured Business ObjecT (Mubot) [61] developed by Crystaliz, Inc. consists of mobile collaborative business agents.

SMART STORE Virtual was developed by Andersen Consulting Research. An example is the BargainFinder Agent which gets album prices from nine virtual retailers.

Open Sesame [72] developed by Charles River Analytics, allows third-party vendors to add intelligence to applications.

3.2 Agent Construction Kits

Clearlake by Guideware Corp [16] can be used to navigate LANs, WANs and Internet to perform complex processes. The Clearlake Agent Manager communicates over the Internet via the Clearlake Agent Protocol to move agents between computers. Clearlake Personal Agents perform processes involving several individuals on the Internet, or to

interact with businesses and service providers. Enterprise Agents are designed to perform, coordinate and track complex, mission-critical business processes within a distributed organization. They also manage the allocation of people to roles within a process, access data in multiple databases and provide secure operation, authentication and digital signatures. They are extensible and provide process management and reporting.

CyberAgent SDK [20], FTP's autonomous, agile and adaptive CyberAgents carry out delegated tasks on a network. Traveling from desktop-to-desktop at the discretion of the sender, these intelligent agents execute remotely, gather data, make decisions, take corrective actions, and report. Moreover, they can modify themselves, collaborate as well as initiate or terminate other processes. Password security and access restrictions prevent agents from unauthorized use and travel. The CyberAgent SDK contains everything that software developers and system and network administrators need to construct, launch, receive and manage intelligent agents. Specifically, the SDK provides an integrated development environment, a CyberAgent Manager and sample agents. The development environment includes a Visual Editor, an agent scripting language featuring OLE and HTML support, debugging tools, sample source code, and a CyberAgent Wizard.

Object REXX and Agent Interaction Facility (AIF) from IBM constructs virtual "office buildings" which house agencies and agents that can pursue a variety of goals, such as information search and retrieval, change notification and product pricing.

Cognitive Adaptive Computer Help (COACH), also from IBM, models user behavior and uses an inference engine to provide both proactive and adaptive assistance in either the application-specific environment or general purpose system interfaces.

Tacoma [75] was developed at the Cornell University and the University of Tromso in Norway. Tacoma agents are written in Tcl. The single abstraction is the meet operation which an agent uses to invoke a second agent. Additional services are provided by other agents. For example, an agent meets with the rexec agent in order to migrate to a remote machine. Notable features of Tacoma are electronic cash, rear guard agents, broker agents that provide scheduling and directory services and the use of Horus for fault tolerance and group communication. Tacoma is being used in a large, distributed application called StormCast. A demonstration of Tacoma and StormCast is available on the WWW (only a standard WWW browser is needed).

The Interactive Development Environment for Agent Systems (IDEAS) [45], provides the facilities for implementing a federation of autonomous agents using the Tool Command Language Tcl [63] with the Tk Toolkit for the X Windows System on UNIX platforms. While each agent performs its reasoning using its own BinProlog [78] database, they co-operate and co-ordinate using messages similar to the KQML [33] format.

4. Knowledge Rovers

Knowledge rovers [41, 43] denote a collection of intelligent agents that can be configured for specific tasks and mission in support of enterprise transactions and information gathering. Rovers provide services such as, information access, query formulation, facilitation, brokerage, mediation, integration, database wrapping, etc., that enable the

modern enterprise to realize its information management needs, while serving its users in a timely, efficient and cost-effective manner.

Most enterprises have legacy systems and also access information on the Internet. Commerce on the Internet requires the ability to access diverse data residing in multiple, autonomous, heterogeneous databases and integrate or fuse that data into coherent information that can be used by the enterprise. The problem is complicated by the fact that the data may by multimedia with diverse temporal and spatial granularity and stored in diverse formats. The semantics of the data may conflict across multiple sources and the data may be of uncertain quality, and the reliability of the source may be questionable.

A federated approach to the Intelligent Integration of Information (I*3) using cooperating agents rather than a monolithic system is more amenable to the needs of information integration and electronic commerce. Thus, a system of cooperating information agents, each having the requisite knowledge to perform its tasks within the larger problem-solving framework provides a flexible architecture for Electronic Commerce.

4.1 Enterprise Data and Information Requirements

In this section we address enterprise data and information. These requirements are general and apply to most enterprises that have legacy systems and also access information on the Internet. We discuss the notions of data pull and data push requirements, review important concepts of information agents as they pertain to data and information architectures, and introduce the notion of active services that will be implemented through a family of configurable and cooperating agents.

4.1.1 Data Pull and Data Push Requirements

In our research in architectures for large-scale systems [35, 42, 49], we have identified two types of data requirements and their associated scenarios: Data Pull and Data Push.

Data Pull denotes a user-initiated information request, for example, in accessing data from archives, or in creating a logistics plan and then accessing multiple databases to instantiate that plan with actual data.

Data Push corresponds to the continual updating of databases used by the enterprise. These updates are under the control of the organizations that designed, built, control, and maintain the databases. The organizations have *autonomy* over and *responsibility* for the data, ensure its quality, and share the data within a federation of information systems [42, 44, 47, 73].

Thus, a system must support both data pull and data push, while providing reasonable service to all users. The approach proposed is a family of intelligent agents, called *knowledge rovers*, which provide services, such as, catalog access, product query formulation, facilitation, brokerage, negotiation, contracting, mediation, integration and source wrapping, to enable the modern enterprise to realize its information management needs, while serving its users in a timely, efficient and cost-effective manner.

4.1.2 Intelligent Agents for the Enterprise

Bird [8] has proposed an agent taxonomy based on two client/server classes. They are *Mobile Agents* (Clients) for Content, Communications and Messaging Services and Static Agents (Servers).

We add third type of agent, *Active View* Agents, which support active objects and views. Active view agents may act on behalf of users in *materializing* objects based on real-time events and conditions, as well as pre-defined rules provided either by users or by other agents.

Bird notes that distributed intelligent systems share many of the same characteristics of multidatabase systems [73], in particular, distribution, heterogeneity, and autonomy. Knowledge and data may be *distributed* among various experts, knowledge bases and databases, respectively. Problem-solving should be a cooperative endeavor [22, 38]. Multiple copies of knowledge and data, possibly in differing formats, must be maintained by the system.

There are several facets to the *heterogeneity* of information in systems [8]: syntactic, control and semantic.

1. *Syntactic heterogeneity* refers to the myriad of knowledge representation formats [36], data definition formats to represent both knowledge and data.
2. *Control heterogeneity* arises from the many reasoning mechanisms for intelligent systems including induction, deduction, analogy, case-based reasoning, etc.[9, 22, 24, 29, 30, 38, 76, 77].
3. *Semantic heterogeneity* [14, 25, 38, 83, 84, 93, 94] arises from disagreement on the meaning, interpretation and intended use of related knowledge and data.

A third characteristic of intelligent systems is that of *autonomy*. There are several aspects to autonomy; in the control structure of an agent, in the extent to which an agent shares information with other agents [29, 32], the manner in which an agent associates with other agents, and structural autonomy in the way an agent fits into an organization of agents for problem-solving [19, 27, 37, 44-48, 66, 76].

4.1.3 Enterprise Information Architecture

Bowman et al [11, 12] describe a three-layer architecture for scalable Internet resource discovery, proposed by the Internet Research Task Group. Table 1 denotes the three-layer architecture which provides access to heterogeneous repositories, including those on the WWW.

Figure 1 depicts GMU approach to the three-layer architecture for the application domain of logistics; the architecture is general enough to apply to most enterprises. The information architecture incorporates the three information layers consisting of:

1) *Information interface layer* where users access the system, formulate queries, collaborate in problem-solving activities, initiate pull scenarios and receive information from push scenarios. Users have access to their local databases and work through local views. We assume that collaboration mechanisms and tools exist at this layer;

Table 1: Internet Three Layer Information Architecture

Information Layer	Layer Service
Information Interface Layer	Users perceive the available information at this layer and may query and browse the data. This layer must support scalable organizing, browsing and search.
Information Management Layer	Responsible for the replication, distribution, and caching of information.
Information Gathering Layer	Responsible for the collecting and correlating the information from many incomplete, inconsistent, and heterogeneous repositories.

2) *Information management layer* where objects, mediated active views, and information in an Information Repository are integrated, managed, replicated, updated. This layer mediates between the information interface layer and the information gathering layer, allowing users to perceive an *integrated information space*, when in reality, data resides in multiple heterogeneous databases and information sources. A mediated view of data is provided at this layer and user views are materialized from the mediated view.

   The Real-Time Information Processing and Filtering process constantly monitors the system for events of importance to enterprise activities, and informs users, the mediated view, and the Information Repository should these events occur.

   The Information Repository contains meta-data and knowledge associated with enterprise resources. It is provided information by the Data/Knowledge Refinement, Fusion and Certification process, and constantly updates the repository. Here data from diverse, heterogeneous inter-networked information sources are mined, scrubbed, refined and evolved to produce high-quality information.

3) *Information gathering layer* where data from diverse, heterogeneous inter-networked information sources are accessed. Special rovers are used to perform the mediated access to local as well as Internet resources.

4.2 Enterprise Active Services

Table 2 presents a collection of active services for the enterprise. These services reflect our I*3 research, particularly in mediated data access through brokerage, facilitation and wrapping services. The GMU research group has developed a federated service architecture that provides Federation Interface Managers (FIMs) as active wrappers for information sources, an intelligent thesaurus [83, 84], temporal mediation services [6, 82], active views [1, 67-71] and inconsistency management services [53-60].

Data mining and knowledge discovery techniques are applicable to the study of data quality, user usage patterns, automatic classification, schema evolution, and system

evolution [52, 64, 65, 91-94]. We now present a synopsis of the various services. Readers may consult the relevant citations for additional details.

4.2.1 Thesaurus Services

The *intelligent thesaurus* [42, 83, 84] is an *active* data/knowledge dictionary capable of supporting multiple ontologies and allowing users to formulate and reformulate requests for objects. The intelligent thesaurus is similar to the thesaurus found in a library; it assists users in identifying similar, broader or narrower terms related to a particular search term, thereby increasing the likelihood of obtaining the desired information from the information repository. In addition, active rules may be associated with object types as well as their attributes and functions.

4.2.2 Federation Services

At GMU we have developed an I*3 federated distributed client/server architecture [34] in which the constituent systems maintain authority and autonomy over their data which at the same time sharing certain information with the federation. Client software and/or server software is provided to members so that they can interface existing information systems with the federation.

A Federation Interface Manager (FIM) is associated with each information system joining the federation. It is specialized into client and server FIMs. Each Client-Federation Interface Manager (FIM) consists of three subsystems, a Client Router, Client-to-Federal Translation Services, and Federal-to-Client Translation Services. The translation services map local data objects to the DPSC representation and vice-versa. The Client-FIM 1) accepts local transactions (including query requests) from local clients in the format used

Table 2: Enterprise Active Services for electronic commerce

Information Layer	Layer Service	Active Services
Information Interface Layer	Users perceive the available information at this layer and may query and browse the data. This layer must support scalable organizing, browsing and search.	Thesaurus Services Federation Services Catalog Services Transportation options
Information Management Layer	Responsible for the replication, distribution, and caching of information.	Mediation Services Active Views I*3 Services, Brokerage Order confirmation and status
Information Gathering Layer	Responsible for the collecting and correlating the information from many incomplete, inconsistent, and heterogeneous repositories.	Data Quality and Inconsistency Management Performance Assessment Correlating order to contract Correlating order fulfillment to billing and payment information

by the local information system, and 2) the router determines whether the transaction is for a local server or a remote server of the federation. If the destination is a valid server, the request is passed to the Client-to-Federal Translation Services so the request may be translated to standard federal transaction format prior to routing it to the appropriate server.

The Server-FIM consists of Federal-to-Server Translation Services, Server-to-Federation Translation Services, and the Server Router. When a transaction arrives at a destination, the Federal-to-Server Translation Services translates the federal transaction into the local transaction format of the server. After translation the transaction is sent to the server router. The router receives all transactions, both locally and remotely generated, and logs them. Transactions are then queued for processing. Once the server has processed the transaction, the server router sends the response, if necessary, to the Server-to-Federal Translation Services, for server to DPSC translation, and then routes the response to the appropriate Client-FIM, where the response is translated into the local format.

4.2.3 Mediation Services

Mediation refers to a broad class of services associated with I*3 [86-89]. At GMU we have focused on the mediation of temporal data of differing granularity. This is of particular importance in the context of multidimensional databases and data warehousing applications, where historical data is integrated and analyzed for patterns and interesting properties.

A *temporal mediator* [6, 81, 82] consists of three components: 1) a repository of *windowing functions* and *conversion functions*, 2) a time unit thesaurus, and 3) a product query interpreter. There are two types of windowing functions: the first associates time points to sets of object instances, and the other associates object instances to sets of time points. A conversion function transforms information in terms of one time unit into that in terms of some other time unit. The time unit thesaurus stores the knowledge about time units (e.g., names of time units and relationships among them). The time-unit thesaurus stores concepts such as the seasons, fiscal year definitions, and calendars, and effects translation of these time units into others.

Users pose queries using the windowing functions and desired time units using a temporal relational algebra. To answer such a user query, the query interpreter first employs the windowing functions together with the time unit thesaurus to access the temporal data from the underlying databases and then uses the time unit thesaurus to select suitable conversion functions which convert the responses to the desired time units. Thus, a temporal mediator provides a simple, yet powerful, interface that supports multiple temporal representations in a federated environment. Temporal mediators may also be used to compare historical databases such as those needed for auditing and data warehousing purposes.

4.2.4 Active Views

Active views are motivated by the need to mediate between users and the plethora of enterprise data and information being generated. Active views can be used to define complex objects, events and conditions of interest.

Thus, active views mitigate the need for users to constantly issue queries to verify the existence of certain events or condition in the enterprise, or to be bombarded constantly with irrelevant information. The active view automates this task by compiling user object specifications into rules monitored by active databases [15, 85], or polling queries issued against traditional enterprise database systems. These rules can also be given rovers that comb the WWW seeking relevant information. Active views also provide active caching [1, 69] of materialized views so as to support subscription services, notification services, and to update user profiles. Further, materialized views can be updated automatically. Active views are discussed in more detail later in this paper.

4.2.5 Data Quality and Inconsistency Management

In any environment of multiple information resources one would expect that sources would overlap in providing similar but inconsistent data. Inconsistencies are detected during the process of integration, and *harmonization agents* are engaged to resolve them. The concept of the harmonization agent is incorporated in the Multiplex [54, 57] proof-of-concept system which considers the reliability and quality of the conflicting information sources, and resolves conflicts in a way that increases the overall value of the information [60].

4.3 Knowledge Rover Architecture

The concept of Knowledge Rovers serves as a metaphor for the family of cooperating intelligent agents that support the information architecture. The notion of a rover is that it can be configured automatically with appropriate knowledge bases (ontologies), task-specific information, negotiation and communication protocols for its mission into cyberspace. Figure 2 superimposes the various agents in the Knowledge Rover Family onto Figure 1, and roles are described below:

4.3.1 Knowledge Rover Types

Executive Agent — is a *coordinator* for a group of agents. It is informed of significant events. A significant event can lead to the activation of new agents. For example, if the enterprise is notified of disaster-relief request, then the executive agent would coordinate with other agents in implementing the relief scenario.

User Agents — acts on behalf of a user, and is responsible for assisting users: 1) in browsing catalogs and information holdings such as the information repository, 2) in the intelligent formulation of queries, and 3) in the planning of tasks within a mission-specific scenario such as provisioning logistic support for a disaster relief effort.

Real-time Agents — are mission-specific, defined and configured to process incoming data, and update the appropriate database or notify the appropriate users. The real-time agents are autonomous, communicate with each other using a pre-defined protocol. Real-time agents are responsible for monitoring the external environment, interacting with other systems, or acting on inputs from users. When an event is detected by a real-time agent, it is signaled to the relevant agents.

Facilitation Agents — provide intelligent dictionary and object location services. For example a facilitation agent [32] might accept a request from the Executive Agent to find all *external* providers of ‘antibiotics,’ and it might respond with the pharmaceutical

producers and suppliers for the region in question. Other agents such as *knowledge rovers* (defined below) could then arrange for the items to be requisitioned, retrieved, and paid for. A knowledge rover could also post a request for bids, accept responses, make contracts, and provision the requested items.

Mediation Agents — are configured to assist in the integration of information from multiple data and information sources, having diverse data formats, different meanings, differing time units, and providing differing levels of information quality. Mediators [86, 87] are configured to accept queries from the Executive, translate the queries into the query language of the appropriate database system, accept the retrieved result, integrate it with results from other sources, and return the information to the Executive for presentation to the User Interface Agent.

Active View Agents — are created to monitor real-time events from the environment or from databases, and to use these events to initiate actions that will result in the update and synchronization of objects in the Data Warehouse and also in local views maintained at user workstations. These agents are configured to perform very specialized monitoring tasks and have a collection of rules and actions that can be executed in response to events and conditions that occur in the environment or in the databases.

Information Curators — are responsible for the quality of information the Information Repository. They assist in evolving the data and knowledge bases associated with enterprise information resources. They work with knowledge rovers to incorporate newly discovered resources into the information repositories.

Knowledge Rovers — are instructed carry out specific tasks on behalf of the executive, such as to identify which vendors have a specific item on hand. This would involve obtaining information from several vendors. The knowledge rover dispatches *field agents* to specific sites to get the relevant information. If the knowledge rover gets similar information from more than one source, it may ask a *mediator* such as Multiplex [57] to

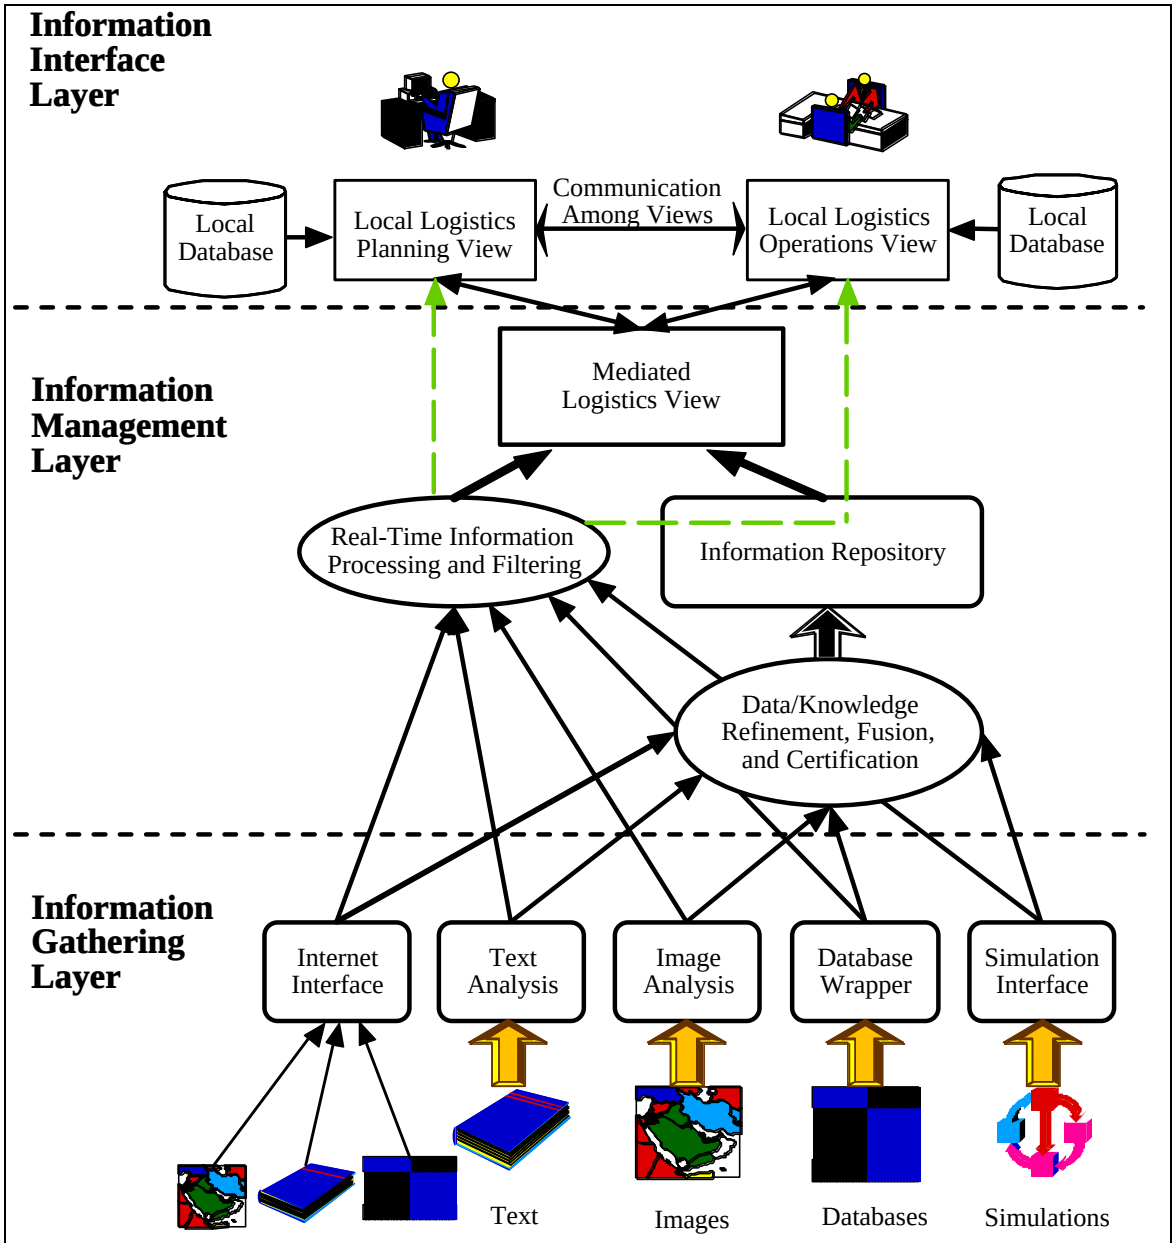

Figure 1. Enterprise Three-Layer Information Architecture

resolve the inconsistency. The knowledge rover reports back to the Executive Agent. The rovers are also responsible for Internet resource discovery. These new information sources and their data are analyzed to determine the adequacy, quality and reliability of retrieved information and whether it should be incorporated into the information repository.

Field Agents — are *specialized* rovers that have expertise in a certain domain, for example, pharmaceuticals, and knowledge about domain-specific information holdings at one or more sites. For example, a field agent could be tasked to monitor all aspects of a single item, say an 'antibiotic' produced by several manufactures and distributed by several vendors. They negotiate with the local systems through their *wrapper*, or Federation Interface Manager, retrieve appropriate data, and forward it to the appropriate requesting agent.

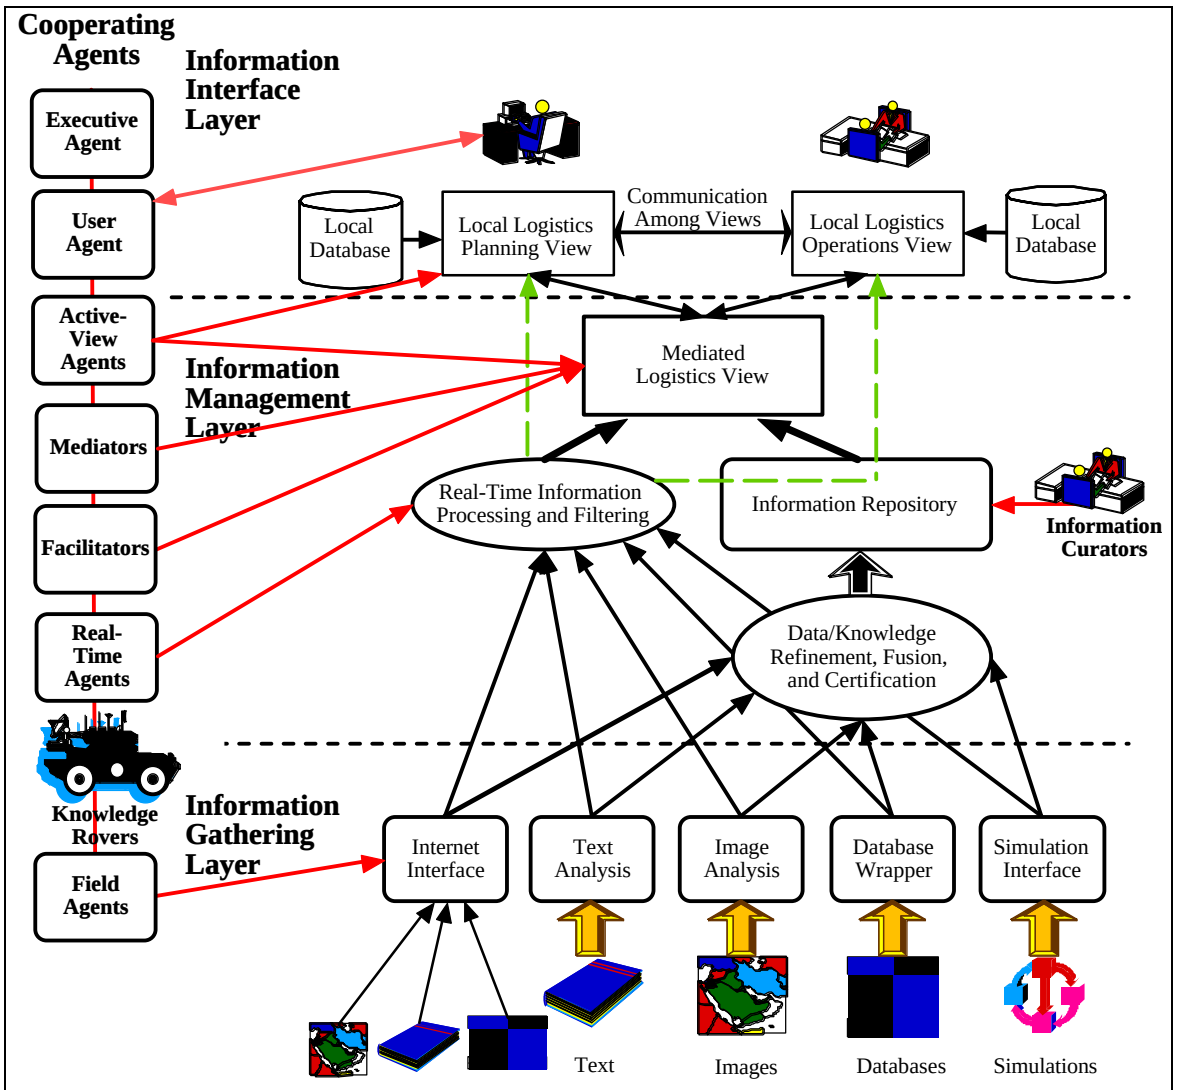

Figure 2. Cooperating Agents for the Information Architecture

5. The DPSC Electronic Marketspace

We have adapted the Knowledge Rover architecture and tailored it to the current vision of the DPSC Electronic Marketspace concept. The design consists of a collection of cooperating multilayered agents. This electronic system is designed primarily to eliminate the need for a large DoD warehouses of food, clothing and pharmaceuticals for the different supply sites.

Since this has *strategic implications*, it is very important to closely monitor inventories and transactions, track the supply of critical items, and take necessary corrective action in a timely fashion. This is assisted by the deployment of special rovers, called sentinels that monitor the system, inform logisticians and decision-makers and trigger alternative contingency plans to take care of a possible crisis. Figure 3 is an overview of the entire system.

5.1 A Design of the Electronic Market Space

Each agent subsystem, called an *agency*, may be considered as a multi-layered agent or a federation of agents with a specific goal and functional role in the marketspace. Therefore each can be considered as a collection of cooperating autonomous agents with particular expertise. A detailed version of the architecture with the various agents and sentinels is shown in Figure 4.

The following discussion outlines the constituent agents and the role they play.

5.2 Marketspace Executive

This agency handles planning and monitoring of the marketspace. It ensures that potential crisis situations are detected early, and avoided by taking suitable action. This also acts as an interface for the organization's planners and strategists who are aware of the long term goals, expectations and strategies.

5.2.1 Strategic Planner

The strategic planner is the interface to the marketspace executive agency. It allows users to provide domain knowledge, define long term plans and objectives, define rules and constraints, identify critical items and situations that need to be monitored via sentinels, define the cause-effect relationships, establish rules and courses of action for the sentinels, and obtain reports on all aspects of the system at any time. The desired features of the interface would be to provide heuristic support for all information searching scenarios and providing suggestions for building better queries.

5.2.2 Sentinel to Track Critical/Strategic Items

This agency continuously monitors the state of the system, tracking the critical and strategic items via sentinels and active views.

5.2.3 Readiness Assessor Sentinel

This agent ensure that strategic plans can be fulfilled by the resources. It is a deamon that ensure that crisis situations are avoided and contingency plans are in place in critical situations.

5.2.4 Demand Forecasting Manager

This agent breaks down plans into resources and material that need to be provided. It also builds the contract terms that will be necessary .

5.2.5 Sentinel Tracking Contract Terms

This agency is a deamon that monitors the system to alert against any deviation from the contracts made so that necessary alerts are flagged and triggers are fired. This allows the system to detect any deviation from plans already in place so that necessary action can be taken in a timely fashion to avoid crisis situations.

5.3 Order Request Broker

This agency handles the day-to-day functions of the marketspace. It ensures that transactions run smoothly, and takes action to correct any inconsistency. This also acts as an interface for the organization's operators by providing optimized customer order - broker mapping, suggestions for building better queries, improving overall order fulfillment and delivery turn-around time, and by allowing the creation of personalized user profiles. This agency may be considered to have the following agents:

5.3.1 Negotiator

This agent would not only provide optimized broker-supplier mapping, but would also negotiate with a registered supplier and form a contract in accordance with the specifications and terms specified by the customer.

5.3.2 Automatic Invoicing

This agent is the supply receiver and payment requester. It consults a database that contains records of contracts made and supplies obtained against these contracts. It

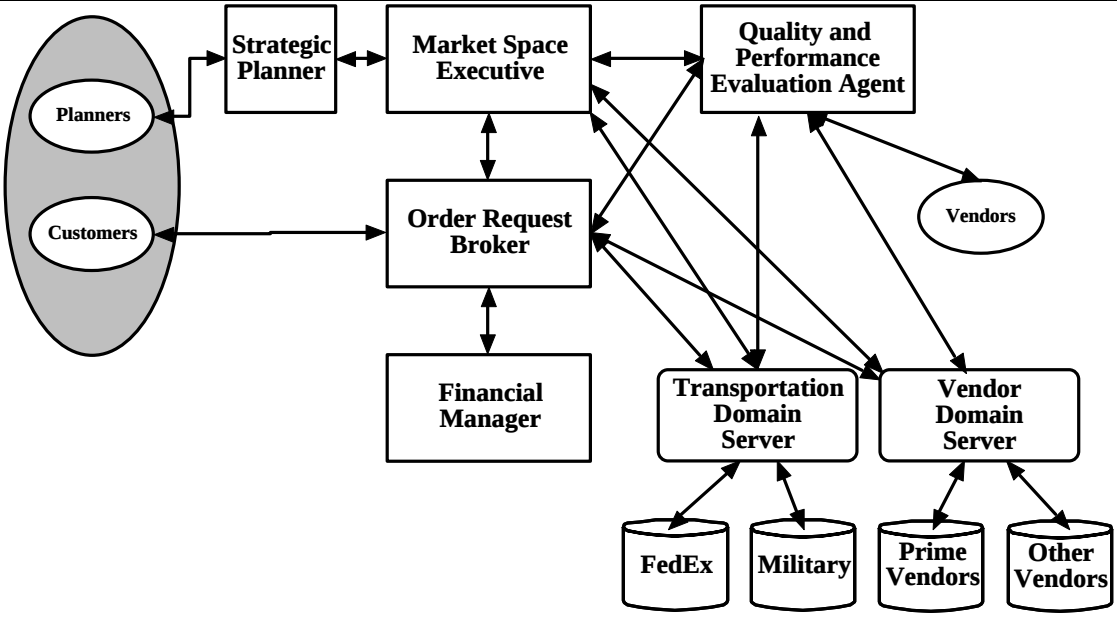

Figure 3: Vision for the DPSC Electronic Market Space

acknowledges the receipt of merchandise based on UPC code scans and a contract identification code which indicates the partial or complete fulfillment of a particular contract. On the receipt of merchandise it sends a request for payment to the appropriate agent in the financial broker.

5.3.3 Order Tracking Manager

This agent takes over the contract after negotiations have been completed by the negotiator and tracks progress made, and maintains contact with the supplier to ensure that the contract will be fulfilled according to schedule. In case of any problem indicated by the supplier, this agent sends the pertinent information to the sentinel tracking contract terms in the marketspace executive as well as the quality and performance tracker in the quality and performance evaluation agent.

5.4 Financial Broker

An area vital to the electronic commerce architecture is a secure payment scheme. There are a number of institutions and organizations involved in research to develop not only a secure form of payment but also to make the Internet a secure place for such a business transaction. Some of the major institutions include CyberCash [21], DigiCash [23] and First Virtual Bank [28]. Besides the security issues that need to be addressed, a viable interface has to be developed to make payments for purchases with a pay-on-delivery paradigm in mind. This would involve automated invoicing, electronic transfer of funds and performing of audits. AMS also has expertise in this area through its EDI activities.

The financial broker is the agency that handles the financial transactions of the organization. It consists of the following agents:

5.4.1 Electronic Invoice Payment and Funds Transfer

This agent removes the need for invoices to be prepared and approved. When the automatic invoicing agent acknowledges the receipt of merchandise corresponding to a particular contract, funds are automatically transferred to the supplier account and record of the financial transaction are sent to the auditor. Depending on the contract terms, the payment may be made incrementally as items are delivered or deferred until the order is completed.

5.4.2 Auditors

This agent maintains records of all the financial transactions and performs audits.

5.5 Quality & Performance Evaluation Agent

This agency not only provides the "yellow pages" of the organization, but also tracks performance of the vendors, builds a level of confidence in different suppliers, verifies supplier claims and obtains referrals. The agency consists of the following agents:

5.5.1 Vendor Registration Facilitator

This agent enables new vendors to make their services available to the client. The agent serves as a broker and establishes the logistics between the vendor and the customer.

5.5.2 Supplier Claim Verification Agent

The services that this agent should provide include the verification of supplier authenticity by obtaining and checking references. Small contracts in non-critical areas may be recommended (to the negotiator in the order request broker) by this agent in order to build trust in a newly registered vendor.

5.5.3 Quality & Performance Tracker

The service that this agent provides includes extraction of information and updating the knowledge stored in the supplier-meta-database. As the system is put into use, a variety of information such as previous product specification, user preferences, suppliers used, and services used can be logged at this tracker site. Techniques such as data mining can be used to extract relationships between the queries, suppliers, clients and information obtained from the order request broker and this information can be used to provide an evaluation of vendor performance. This evaluation would be taken into consideration by the negotiator in the order request broker when selecting a vendor for a new contract.

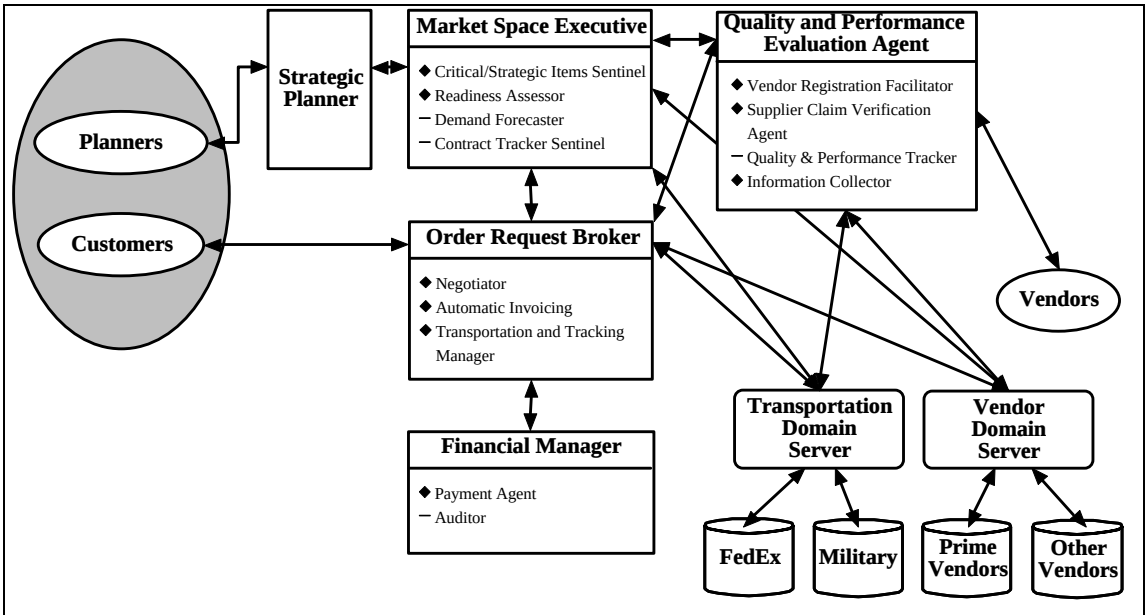

Figure 4: Detailed Architecture for the DPSC Electronic Marketspace

5.6 Customer Scenarios

This section presents scenarios showing how the agent community would cooperate in providing information to logistic users.

Customer	Example Scenario	Related Breakthrough concepts	Highlights
92 Airborne	Deploying to Bosnia where floods have caused an outbreak of cholera. Assumptions: Supplies are prepackaged and delivered prior to deployment; Information is available.	Automated Readiness Models available to the Marketspace Executive; Electronic Cataloging; Data Warehousing; Information Sharing; Electronic commerce and communication.	Soldier on the ground in Bosnia uses laptop computer to hook up to a satellite and access DPSC's electronic marketspace; Searches electronic catalogs for price, delivery, stock availability, etc; Orders directly from vendor; With insight into logistics requirements from operations plans and readiness models, DPSC anticipates requisitions models, and facilitates orders rather than reacting from "Ground Zero"; In-place business arrangements; facilitate delivery of goods without scrambling. **Examples**: A soldier can type in "cholera" or "Bosnia/ wintertime" and obtain a list of suggested items to order from the table of equipment (TOE) from the readiness model; Soldier requests either generic - "aspirin" or specific brand such as BAYER or ALLEVE.
Florida National Guard Unit	Deploying troops for hurricane emergency relief. Assumptions: 1. Supplies are not prepackaged. 2. Transporting supplies is more difficult. 3. National Guard requirements are not usually covered under typical business arrangements. Improvisation is key to success.	Electronic cataloging; Data warehousing; Automated readiness models available to the Marketspace; Information sharing; Electronic commerce.	Guardsman in Florida uses laptop computer to hook up to satellite and access DPSC's electronic marketspace; Searches electronic catalogs for price, delivery, stock availability, etc; Orders directly from vendor; DPSC uses historical supply models and requirements; Prime Vendor information is available in data warehouses; DPSC explores piggybacking on existing business arrangements.

Army Base	Reserves being activated	Information Sharing; Automated Readiness Models available to the Marketspace; Data Warehousing; Real Time Processing; Shared Production.	**External:** Customer contacts one source to get support for a given mission and DPSC's ability to fulfill needs. **Internal:** Type in "cold, desert" to determine requirements, recommendations, incidentals for all needs, now and in the future; System reports stock on hand; Stock immediately available is determined and shared production helps evaluate options for remaining requirements; Asset balances updated real-time as stock is allocated; Web site is culmination of all information needed.
Air Force Base (AFB)	Out-of-stock condition: Prime vendor in same town. Goal: Get stock from prime vendor to AFB.	Information Sharing; Electronic Commerce using established Business Arrangements; Electronic Catalogs; Data Warehousing;	**External Perspective:** Customer wants one contact (either prime vendor or DPSC) for guaranteed receipt of goods within hours. **Internal Perspective:** Prime vendor relationship established via "Best Value" Procurement; Real-time processing of requirements; Funds verified; Violations flagged and turned around immediately; Vendor and Order Request Broker are notified automatically via "intelligent gateway"; Requisition visible to contractor so he/she can ship goods; Requisition automatically re-edited to reflect alternate supplier if Prime Vendor cannot deliver; Information sharing within the system provides flexibility and automated intelligence to best determine where stock should come from; Delivery order validated against contract terms, processed automatically against contract, sent electronically to contractor, updated automatically and simultaneously to reflect new totals - signed electronically; Information on past performance automatically recorded; Best shipping strategy determined by automatically evaluating excess stock on hand and customer priority; Contract terms (e.g. delivery schedule) are strictly monitored and enforced.

5.7 Information Flow

Information required by the different agents in order for them to perform correctly, comes from different sources. The information exchange takes place by sending messages from one agent or object to another. The flow of information is critical to the agent performance. In this section we considered the different types of messages and the routes that they would follow.

Source of Information	**Information Description**	**Receiver of Information**
Customer Strategic planner	Order info (requisitions, priorities, etc.)	Marketspace Executive Negotiator Supplier Claim Verification Agent
Customer	Order ratification	Negotiator
Supplier Claim Verification Agent Established business arrangements from vendor database	Supply Sources	Readiness Accessor Strategic planner User
Readiness Assessor	Recommended alternative sources of supply prioritized to meet user order criteria	Strategic planner
Vendor database	Historic contract and vendor performance info	Strategic planner Demand forecaster
Order Request Broker	Contract and vendor performance info	Strategic Planner Quality and Performance tracker Vendor database
Vendor database	Historic vendor performance measures and uncertainties Vendor's expected contract performance parameters	Readiness Assessor
Quality & Performance Tracker	Dynamic, real-time vendor performance estimates and uncertainties/confidence levels Tradeoff relationships between vendor cost/price and performance	Readiness Assessor
Demand Forecaster	Demand data over the course of the planning scenario Demand priorities established	Strategic planner

Source of Information	Information Description	Receiver of Information
Strategic Planner Negotiator Readiness Assessor	Customer ordering criteria Real-time requisitions Subset of established business relationships that meet legal and regulatory conditions	Negotiator
Negotiator	Running contracts	Order Tracking Manager Quality & Performance Tracker Automatic Invoicing
Automatic Invoicing	Partial fulfillment of contract	Payment Agent Quality & Performance Tracker Order tracking Manager
Demand Forecaster	Estimated risk of completing the transaction	Strategic Planner
Readiness Assessor	Best supply support alternative for demand Prioritized list of alternative supply support Measure of fit with customer ordering criteria	Strategic Planner
Order tracking Manager	Deviation from schedule Delays in delivery Incomplete fulfillment of contract	Quality & Performance Tracker
Payment Agent	Record of funds transferred	Auditor
Vendor	Electronic Catalog Payment methods	Vendor Registration Facilitator
Supplier claim Verification agent	New supplier information	Readiness Assessor Vendor database
Readiness Assessor Strategic Planner	Information about unconventional requirements	Information Collector

6. Conclusions

This paper has presented the current state of information technology in the area of Electronic Commerce, and its impact on the vision of the DPSC Electronic Marketspace. The traditional tools of electronic commerce have been discussed and the impact of the Internet and World Wide Web on Electronic Commerce has been considered. We have briefly examined some of the commercial software tools and techniques that are available.

Although intelligent software agent research and development is underway, there are no commercially available agents to fully support the DPSC vision. However, we feel that agent technology is maturing rapidly, and this technology can be especially useful for the construction of sentinels and for automating many of the Web-based transactions associate with the customer-supplier relationships.

We have presented the George Mason University approach to information integration and intelligent software agents through an information architecture consisting of the information interface, management and gathering layers. Intelligent *active services* are presented for each layer and the role of knowledge rovers is discussed. Knowledge rovers represent a *family* of cooperating intelligent agents that may be configured to support enterprise tasks, scenarios, and decision-makers. These rovers play specific roles within an enterprise information architecture, supporting users, maintaining active views, mediating between users and heterogeneous data sources, refining data into knowledge, and roaming the Global Information Infrastructure seeking, locating, negotiating for and retrieving data and knowledge specific to their mission.

The architecture proposed by GMU for the DPSC Electronic Marketspace, is based on a federation of multilayered agencies, consisting of autonomous cooperating agents capable of doing business on the World Wide Web.

Rapid advances in Internet and Web security paradigms and mechanisms, E-commerce software services, and their use by large industries such and General Electric and FedEx, indicate that secure E-commerce on the Web, assisted by intelligent software agents such as Knowledge Rovers, is not only feasible, but we will see numerous applications developed in the near future.

7. References

[1] R. Alonso, D. Barbara, and H. Garcia-Molina, "Data Caching Issues in an Information Retrieval System," *ACM Transactions on Database Systems*, vol. 15, 1990.

[2] Amazon, "http://www.Amazon.com," .

[3] R. Atkinson, "IP Encapsulating Security Payload," *Internet RFC 1827, Proposed Standards*, 1995.

[4] BASIS, "http://logic.stanford.edu/cit/cit.html," .

[5] T. Berners-Lee, R. Cailliau, A. Loutonen, H. F. Nielsen, and A. Secret, "The World-Wide Web," *Communications of the ACM*, vol. 37, pp. 76—82, 1994.

[6] C. Bettini, X. S. Wang, E. Bertino, and S. Jajodia, "Semantic Assumptions and Query Evaluation in Temporal Databases," ACM SIGMOD International Conference on Management of Data, San Jose, CA, 1995.

[7] Bhimani, "Securing the Commercial Internet," *Communications of the ACM*, pp. 29-35, 1996.

[8] S. D. Bird, "Toward a taxonomy of multi-agent systems," *International Journal of Man-Machine Studies*, vol. 39, pp. 689-704, 1993.

[9] A. Bond and L. Gasser, "Readings in Distributed Artificial Intelligence," . San Mateo, CA: Morgan Kaufmann Publishers, Inc., 1988.

[10] N. S. Borenstein, "Perils and pitfalls of practical cybercommerce," *Communications of the ACM*, vol. 39, pp. 36-44, 1996.

[11] C. M. Bowman, P. B. Danzig, D. R. Hardy, U. Manber, M. F. Schwartz, and D. P. Wessels, "Harvest: A Scalable, Customizable Discovery and Access System," University of Colorado, Boulder, Technical Report CU-CS-732-94, March 1995.

[12] C. M. Bowman, P. B. Danzig, U. Manber, and M. F. Schwartz, "Scalable Internet Resource Discovery: Research Problems and Approaches," *Communications of the ACM*, vol. 37, pp. 98—107, 1994.

[13] Catalog, "http://www.icat.com/," .

[14] S. Ceri and J. Widom, "Managing Semantic Heterogeneity with Production Rules and Persistent Queues," IBM Research Laboratory, San Jose, CA, Technical Report RJ9064 (80754), October 1992.

[15] S. Chakravarthy and J. Widom, "Foreword to Special Issue on Active Database Systems," *Journal of Intelligent Information Systems*, vol. 7, pp. 64, 1996.

[16] Clearlake, "http://www.guideware.com/," .

[17] CommerceNET, "http://www.commerce.net/," .

[18] Connect, "http://www.connect.inc.com," .

[19] M. Cutkosky, R. Englemore, R. Fikes, M. Genesereth, T. Gruber, W. S. Mark, J. Tenenbaum, and J. Weber, "PACT: An Experiment in Integrating Concurrent Engineering Systems," In *IEEE Computer*, vol. 26, 1993, pp. 28-37.

[20] CyberAgent, "http://i-shop.iworld.com/CyberAgent_SDK.html," .

[21] Cybercash, “http://tear.cybercash.com/,” .

[22] R. Davis and R. Smith, “Negotiation as a Metaphor for Distributed Problem Solving,” *Artificial Intelligence*, vol. 20, pp. 63-109, 1983.

[23] Digicash, “http://cybercash.com/,” .

[24] R. Engelmore and T. Morgan, “Blackboard Systems,” In *The Insight Series in Artificial Intelligence*, T. Morgan, Ed.: Addison-Wesley Publishing Company, 1988, pp. 602.

[25] A. Farquhar, R. Fikes, W. Pratt, and J. Rice, “Collaborative Ontology Construction for Information Integration,” Knowledge Systems Lab, Computer Science Department, Stanford University, Palo Alto, CA, KSL-95-63, August 1995.

[26] FedEx, “http://www.fedex.com,” .

[27] T. Finin, “CIKM '94 Workshop on Intelligent Information Agents,” , T. Finin, Ed. Baltimore, MD, 1994.

[28] First_Virtual_Bank, “http://www.fv.com/,” .

[29] R. Fritzen, T. Finin, D. McKay, and R. McEntire, “KQML — A Language and Protocol for Knowledge and Information Exchange,” International Distributed Artificial Intelligence Workshop, Seattle, WA, 1994.

[30] L. Gasser and M. N. Huhns, “Distributed Artificial Intelligence,” In *Research Notes in Artificial Intelligence*, vol. II. London and San Mateo: Pitman; Morgan Kaufmann Publishers, Inc., 1989, pp. 519.

[31] General_Electric_Information_Systems, “http://www.tpn.geis.com,” .

[32] M. Genesereth and S. P. Ketchpel, “Software Agents,” *Communications of the ACM*, vol. 37, pp. 48-53, 1994.

[33] M. Genesereth and Others, “Specification of the KQML agent communication language,” 1993.

[34] H. Gomaa and G. K. Farrukh, “An Approach for Configuring Distributed Applications from Reusable Architectures,” IEEE International Conference on Engineering of Complex Computer Systems, Montreal, Canada, 1996.

[35] H. Gomaa, D. Menascé, and L. Kerschberg, “A Software Architectural Design Method for Large-Scale Distributed Information Systems,” *Journal of Distributed Systems Engineering*, vol. September, 1996.

[36] T. Gruber, “A Translation Approach to Portable Ontology Specifications,” *Knowledge Acquisition*, vol. 5, pp. 199-220, 1993.

[37] R. V. Guha and D. B. Lenat, “Enabling Agents to Work Together,” *Communications of the ACM*, vol. 37, pp. 126-142, 1994.

[38] C. Hewitt, “Open Information Systems Semantics for Distributed Artificial Intelligence,” *Artificial Intelligence*, vol. 47, pp. 79-106, 1991.

[39] InfoHaus, “http://www.mo.net/walkrich/infohaus.htm,” .

[40] Kasbah, “http://pattie.www.media.mit.edu/people/pattie/pattie-proj-list.html,” .

[41] L. Kerschberg, “Knowledge Rovers: Cooperative Intelligent Agent Support for Enterprise Information Architectures,” In *Cooperative Information Agents*, vol. 1202, *Lecture Notes in Computer Science*, P. Kandzia and M. Klusch, Eds. Berlin and Heidelberg: Springer-Verlag, 1997.

[42] L. Kerschberg, H. Gomaa, D. A. Menascé, and J. P. Yoon, “Data and Information Architectures for Large-Scale Distributed Data Intensive Information Systems,” Proc. of the Eighth IEEE International Conference on Scientific and Statistical Database Management, Stockholm, Sweden, 1996.

[43] L. Kerschberg, S. Jajodia, H. Gomaa, and A. Motro, “Knowledge Rovers: A Family of Configurable Software Agents,” In *Proposal Funded by DARPA's Advanced Logistics Program*. Fairfax, VA: Center for Information Systems Integration and Evolution, George Mason University, URL: http://www.isse.gmu.edu/KRG/, 1996.

[44] M. Klusch, “Using a Cooperative Agent System FCSI for a Context-Based Recognition of Interdatabase Dependencies,” Workshop on Intelligent Information Agents, CKIM Conference, Gaithersburg, MD, 1994.

[45] M. Klusch, “Utilitarian Coalition Formation Between Information Agents,” In *Cooperative Knowledge Processing*, S. Kirn and G. O'Hare, Eds. London: Springer-Verlag, 1996.

[46] M. Klusch and O. Shehory, “Coalition Formation Among Rational Information Agents,” Seventh European Workshop on Modelling Autonomous Agents in a Multi-Agent World (MAAMAW-96), Eindhoven, Netherlands, 1996.

[47] M. Klusch and O. Shehory, “A Polynomial Kernel-Oriented Coalition Algorithm for Rational Information Agents,” Second International Conference on Multi-Agent Systems, Kyoto, Japan, 1996.

[48] P. Maes, “Designing Autonomous Agents: Theory and Practice from Biology to Engineering and Back,” In *Special Issues of Robotics and Autonomous Systems*. Cambridge, MA, London, England: The MIT Press, 1990, pp. 194.

[49] D. A. Menascé, H. Gomaa, and L. Kerschberg, “A Performance-Oriented Design Methodology for Large-Scale Distributed Data Intensive Information Systems,” First IEEE International Conference on Engineering of Complex Computer Systems, Florida, 1995.

[50] P. Metzger, P. Karn, and W. Simpson, “The ESP DES-CBC Transform,” *Internet RFC 1829, Proposed Standards*, 1995.

[51] P. Metzger and W. Simpson, “IP Authentication using Keyed MD5,” *Internet RFC 1828, Proposed Standards*, 1995.

[52] R. S. Michalski, L. Kerschberg, K. Kaufman, and J. Ribeiro, “Mining for Knowledge in Databases: The INLEN Architecture, Initial Implementation and First Results,” *Journal of Intelligent Information Systems*, vol. 1, pp. 85-113, 1992.

[53] A. Motro, “Accommodating Imprecision in Database Systems: Issues and Solutions,” In *Multidatabase Systems: An Advanced Solution to Global Information Sharing*, A. R. Hurson, M. W. Bright, and S. Pakzad, Eds.: IEEE Computer Society Press, 1993, pp. 381-386.

[54] A. Motro, “A Formal Framework for Integrating Inconsistent Answers from Multiple Information Sources,” Department of Information and Software Systems Engineering, George Mason University, Fairfax, VA, Technical Report ISSE-TR-93-106, October 1993.

[55] A. Motro, “Intensional Answers to Database Queries,” *IEEE Transactions on Knowledge and Data Engineering*, vol. 6, pp. 444-454, 1994.

[56] A. Motro, “Management of Uncertainty in Database Systems,” In *Modern Database Systems: The Object Model, Interoperability and Beyond*, W. Kim, Ed.: Addison-Wesley Publishing Company/ACM Press, 1994.

[57] A. Motro, “Multiplex: A Formal Model for Multidatabases and Its Implementation,” ISSE Department, George Mason University, Fairfax, VA, Technical Report ISSE-TR-95-10, 1995.

[58] A. Motro, “Responding with Knowledge,” In *Advances in Databases and Artificial Intelligence, Vol. 1: The Landscape of Intelligence in Database and Information Systems*, vol. 1, L. Delcambre and F. Petry, Eds.: JAI Press, 1995.

[59] A. Motro, D. Marks, and S. Jajodia, “Aggregation in Relational Databases: Controlled Disclosure of Sensitive Information,” European Symposium on Research in Computer Security, 1994.

[60] A. Motro and P. Smets, “Uncertainty Management in Information Systems: from Needs to Solutions,” . Norwall, MA: Kluwer Academic Publishers, 1996, pp. 480.

[61] Mubot, “http://www.crystaliz.com/logicware/mubot.html,” .

[62] OpenMarket, “http://payment.openmarket.com/,” .

[63] J. K. Ousterhout, *Tcl and Tk Toolkit*: Addison-Wesley, 1994.

[64] G. Piatetsky-Shapiro and W. J. Frawley, “Knowledge Discovery in Databases,” . Menlo Park, CA: AAAI Press/MIT Press, 1991.

[65] J. Ribeiro, K. Kaufman, and L. Kerschberg, “Knowledge Discovery in Multiple Databases,” ISMM International Conference on Intelligent Information Management Systems, Washington D.C., 1995.

[66] D. Riecken, “M: An Architecture of Integrated Agents,” *Communications of the ACM*, vol. 37, pp. 106-116, 1994.

[67] A. Segev and J. Park, “Updating Distributed Materialized Views,” *IEEE Transactions on Knowledge and Data Engineering*, vol. 1, pp. 173-184, 1989.

[68] L. Seligman, “A Mediated Approach to Consistency Management Among Distributed, Heterogeneous Information Systems,” In *Information and Software Systems Engineering*. Fairfax: George Mason University, 1994.

[69] L. Seligman and L. Kerschberg, “An Active Database Approach to Consistency Management in Heterogeneous Data- and Knowledge-based Systems,” *International Journal of Cooperative and Intelligent Systems*, vol. 2, 1993.

[70] L. Seligman and L. Kerschberg, “Knowledge-base/Database Consistency in a Federated Multidatabase Environment,” IEEE RIDE — Interoperability in Multidatabase Systems, Vienna, Austria, 1993.

[71] L. Seligman and L. Kerschberg, “Federated Knowledge and Database Systems: A New Architecture for Integrating of AI and Database Systems,” In *Advances in Databases and Artificial Intelligence, Vol. 1: The Landscape of Intelligence in Database and Information Systems*, vol. 1, L. Delcambre and F. Petry, Eds.: JAI Press, 1995.

[72] O. Sesame, “http://www.csm-usa.com/sesame.htm,” .

[73] A. Sheth and J. Larson, “Federated Database Systems for Managing Distributed, Heterogeneous, and Autonomous Databases,” *ACM Computing Surveys*, vol. 22, pp. 183-236, 1990.

[74] Sivori, “Evaluated Receipts and Settlement at Bell Atlantic,” *Communications of the ACM*, pp. 24-28, 1996.

[75] TACOMA, “http://www.cs.uit.no/DOS/Tacoma/index.html,” .

[76] M. Tambe, W. L. Johnson, R. M. Jones, F. Koss, J. E. Laird, P. S. Rosenbloom, and K. Schwamb, “Intelligent Agents for Interactive Simulation Experiments,” In *AI Magazine*, vol. 16, 1995, pp. 15-39.

[77] M. Tambe and P. S. Rosenbloom, “Event Tracking in a Dynamic Multi-Agent Environment,” USC — Information Sciences Institute, Technical Report ISI-RR-393, September 11, 1994.

[78] P. Tarau, “BinProlog 3.0 User Guide,” *University de Moncton, Canada*, 1994.

[79] Virtual_Vineyards, “http://www.virtualvin.com,” .

[80] Vision, “http://www.thevisionfactory.com/,” .

[81] X. S. Wang, S. Jajodia, and V. S. Subrahmanian, “Temporal Modules: An Approach Toward Federated Temporal Databases,” ACM SIGMOD International Conference on Management of Data, Washington, D.C., 1993.

[82] X. S. Wang, S. Jajodia, and V. S. Subrahmanian, “Temporal Modules: An Approach toward Federated Temporal Databases,” *Information Sciences*, vol. 82, pp. 103-128, 1995.

[83] D. Weishar, “A Knowledge-Based Architecture for Query Formulation and Processing in Federated Heterogeneous Databases,” In *Information and Software Systems Engineering*. Fairfax, VA: George Mason University, 1993, pp. 230.

[84] D. Weishar and L. Kerschberg, “Data/Knowledge Packets as a Means of Supporting Semantic Heterogeneity in Multidatabase Systems,” In *ACM SIGMOD Record*, 1991.

[85] J. Widom and S. Ceri, “Active Database Systems: Triggers and Rules for Advanced Database Processing,” : Morgan Kaufmann Publishers, Inc., 1995.

[86] G. Wiederhold, “The Roles of Artificial Intelligence in Information Systems,” *Journal of Intelligent Information Systems*, vol. 1, pp. 35-56, 1992.

[87] G. Wiederhold, “Foreword to Special Issue on the Intelligent Integration of Information,” *Journal of Intelligent Information Systems*, vol. 6, 2/3, pp. 93-97, 1996.

[88] G. Wiederhold, S. Jajodia, and W. Litwin, “Dealing with Granularity of Time in Temporal Databases,” In *Lecture Notes in Computer Science*, vol. 498, R. Anderson and others, Eds.: Springer-Verlag, 1991, pp. 124-140.

[89] G. Wiederhold, S. Jajodia, and W. Litwin, “Integrating Temporal Data in a Heterogeneous Environment,” In *Temporal Databases: Theory, Design, and Implementation*, A. U. Tansel, S. Jajodia, and others, Eds.: Benjamin/Cummings, 1993, pp. 563-579.

[90] C. Wilder and M. K. McGee, “GE The Net Pays Off,” In *Information Week*, 1997, pp. 14-16.

[91] J. P. Yoon and L. Kerschberg, “A Framework for Constraint Management in Object-Oriented Databases,” International Conference on Information and Knowledge Management, Baltimore, MD, 1992.

[92] J. P. Yoon and L. Kerschberg, “A Framework for Knowledge Discovery and Evolution in Databases,” *IEEE Transactions on Knowledge and Data Engineering*, 1993.

[93] J. P. Yoon and L. Kerschberg, “Semantic Query Reformulation in Object-Oriented Systems,” In *International Conference on Deductive and Object-Oriented Databases*, vol. 760. Phoenix, AZ: Springer-Verlag, 1993.

[94] J. P. Yoon and L. Kerschberg, “Semantic Update Optimization in Active Databases,” Proceedings IFIP WG2.6 Working Conference on Database Semantics (DS-6), Atlanta, 1995.