Intelligent Demand Signal Supply-Chain Information Management for the Pharmaceutical Industry

Executive Summary

Planning and execution processes in a supply-chain are subject to a high level of complexity because of the number of parties and issues involved, the number of relationships that exist among them, and the dynamic nature of the execution environment. The large volume of data flowing through a sizable computer-based supply-chain information management system that is based on rote data-processing principles tends to overwhelm the human users. As a result many opportunities for improving the efficiency of supply-chain processes and thereby reducing costs are overlooked by the human users, who are forced into a reactive mode.

The level of complexity is particularly acute in the pharmaceutical industry where the assessment, prevention and/or mitigation of risk within the supply-chain are a major and continuous concern. The risks, which are both financial and physical in nature, include increasing pressure from generic drugs as patents expire, fierce competition from off-shore and on-shore mail order companies, counterfeit products, contamination, recalls, theft, and unreliable demand forecasting. The situation is compounded by the general fragmentation of the pharmaceutical supply-chain, lack of in-transit and stock visibility, changing and non-uniform government regulations, and an inability to deal effectively with external influences such as personnel illness, weather, traffic congestion, unavailability of fuel, and natural disasters.

Concurrently the pharmaceutical industry is faced with the exponentially increasing volume of accessible data, commonly referred to as Big Data, which tends to cloud useful information and at the same time is forcing the industry to engage with its customers in the entirely new and untested Social Media arena. Similar data deluge symptoms are being experienced in other domains such as Internet searches where the number of website hits returned for a single query can easily exceed several million. To overcome the Big Data problem there is a need for computer software to be able to interpret data and automatically extract the information that is useful to the human user. However, the same data may have entirely different meanings depending on the context in which the data exists. Fortunately, advances in information technology over the past 10 years have made it possible to construct virtual models of real world context in a given application domain that can be processed and understood to a sufficient degree by a computer.

This paper describes the implementation principles and end-state capabilities of a computer-based intelligent supply-chain information management system environment that includes a virtual model of real world supply-chain context and multiple software agents that are able to interact with each other and the human users. Implemented in a service-oriented architecture (SOA) based infrastructure, the virtual context model provided by a multi-layer ontology and the collaborating software agents are able to continuously monitor the state of the supply-chain by interpreting the flow of data in the appropriate context. This allows the agents to rapidly re-plan in case of supply-chain interruptions, discover and act on opportunities for improvements, and identify patterns and trends based on the continuous analysis of historical data. As a result the human users are relieved from lower level data interpretation tasks and provided with actionable information for reactive and proactive planning and execution management functions. Order of magnitude improvements in efficiency and reduction in cost are achievable with such context-based information-centric software systems.

1. The Inherent Complexity of Supply-Chain Information Management

Planning and execution within a supply-chain can have all of the characteristics that are commonly associated with the family of *complex problems*. These characteristics include: many entities and issues that are related to each other; large volume of data that needs to be categorized and analyzed to extract useful information; the reliability of some of the data may be questionable; incomplete data in some areas requiring time critical decisions to be made with partial information; and, a dynamically changing and largely unpredictable execution environment.

Of particular concern to the pharmaceutical industry are the risks associated with product discontinuity, product shortages, product contamination, poor performance, patient safety, communication failures, theft, counterfeiting, non-conformance with government regulations, and competition. The fact that the assessment, prevention and/or mitigation of such risks within the supply chain is a major concern within the pharmaceutical industry is validated by the relatively large number of articles on this subject that regularly appear in journals such as SciRes, International Journal of Physical Distribution and Logistics Management, Business Process Management Journal, International Journal of Logistics Management, and so on. Of particular interest is an article by Breen (2008) that presents the results of a research workshop attended by 20 key stakeholders, including pharmaceutical manufacturers, distributers, pharmacy procurement specialists, and government regulators. The purpose of the workshop was for the attendees to identify pharmaceutical supply-chain risks and rate their criticality in terms of control, occurrence and impact on a scale of 1 to 10, with a score of 10 representing the highest risk (Table 1).

Table 1: Ratings of relative supply-chain risks by pharmaceutical industry stakeholders[1]

Risk Rating	Supply-Chain Risk Type
10	Fragmentation of supply chain; - no single source, multiple channels, lack of communication, and unilateral decisions.
9.0	Lack of visibility of stock.
8.5	Lack of forecasting capabilities such as unexpected increase in demand, inadequate capacity to meet unexpected demand, not being able to respond to demand, and general scarcity of timely information.
8.0	Inability to deal effectively with regulatory issues such as change of standards, drug recalls, and manufacturing licensing.
7.5	Inability to deal effectively with external influences such as personnel illness, weather, traffic congestion, unavailability of fuel, and natural disasters (both recovery and opportunities).
6.5	Too much data and inability to effectively deal with short term supply-chain planning.
6.0	Counterfeiting (would include e-pedigree and serialization requirements).

As shown in Table 1, the most serious concerns of the pharmaceutical industry are related to an apparent lack of meaningful information in support of effective planning and decision-making even though the volume of available data has greatly increased. These workshop results are consistent with the data deluge problems being experienced by most commercial and government supply-chains in an era of global connectivity. They point to the need for computer-based supply-chain tools with analytical capabilities that are able to automatically extract information from data and operate at the information management rather than data-processing level. This paper describes an enterprise-wide intelligent information management environment based on currently available computer hardware and software technology that is capable of providing the required level of support for the pharmaceutical industry.

[1] Breen L. (2008); 'A Preliminary Examination of Risk in the Pharmaceutical Supply Chain (PSC) in the National Health Service (UK)'; Journal of Service Science and Management, 1(2), (pp 193-198).

2. Desirable Capabilities of an Intelligent Supply-Chain Environment

From a general point of view there are at least two overriding requirements for an intelligent computer-based decision-making environment. The first requirement relates to the representation of information within the environment. The software must have some level of *understanding* of the information context that underlies the interactions of the human user with the environment. This is fundamental to any meaningful human-computer interaction that is akin to a partnership. The level to which this *understanding* can be elevated will largely determine the assistance capabilities and essentially the value of the software environment to the human user.

The second requirement is related to the need for collaboration. In a broad sense this includes not only the ability to interact with human stakeholders who play a role in the supply-chain, such as planning and management personnel, vendors, remote distribution centers, shippers, and customs officials, but also non-human sources of information and capabilities such as software agents with analytical capabilities. All of these interactions between human participants in the supply-chain processes, data sources, and software-based problem solving capabilities, must be able to be performed seamlessly without the user having to be concerned about access protocols, data formats, or system interoperability issues.

3. The Technical Approach

The desired capabilities of an intelligent supply-chain planning and execution environment call for a distributed system architecture that can be accessed from any physical location, is highly flexile, and totally transparent to the human user. In particular, the user must be shielded from the many protocols and data and content exchange transformations that are required to access capabilities and maintain seamless interoperability among those capabilities. Any member of the supply-chain team, once authenticated during the single sign-on point of entry, should be able to access those capabilities (e.g., intelligent decision-assistance tools and data sources) that are included in the authentication certificate. The focus of the human user should not be on systems, but on the capabilities or *services* that the computer-based environment can provide.

The notion of *services* is well established in our daily activities. Everywhere we see countless examples of tasks being performed by a combination of services, which are able to interoperate in a manner that results in the achievement of a desired objective. Typically, each of these services is not only *recomposable* but also sufficiently *decoupled* from the final objective to be useful for the performance of several somewhat similar tasks that may lead to quite different results. For example, a common knife can be used in the kitchen for preparing vegetables, or for peeling an orange, or for physical combat, or as a makeshift screwdriver. In each case the service provided by the knife is only one of the services that are required to complete the task. Clearly, the ability to design and implement a complex process through the application of many specialized services in a particular sequence has been responsible for most of mankind's achievements in the physical world.

3.1 Service-oriented architecture (SOA)

In the software domain these same concepts have gradually led to the adoption of Service-Oriented Architecture (SOA) principles. While SOA is by no means a new concept in the software industry it was not until Web services became available that the principles of this concept could be readily implemented. In the broadest sense SOA is a software framework for computational resources to provide services to customers, such as other services or users. A fundamental intent that is embodied in the SOA paradigm is *flexibility*. To be as flexible as possible a SOA environment is highly modular, platform independent, compliant with standards, and incorporates mechanisms for identifying, categorizing, provisioning, delivering, and monitoring services. The principal components of a conceptual SOA implementation scheme (Figure 1) include an Enterprise Service Bus (ESB), one or more portals to external clients with single sign-on facilities, and the enterprise services that facilitate the ability of the user community to perform its operational tasks. The ESB provides the communication bridge that facilitates the exchange of messages among services, although the services do not necessarily know anything about each other.

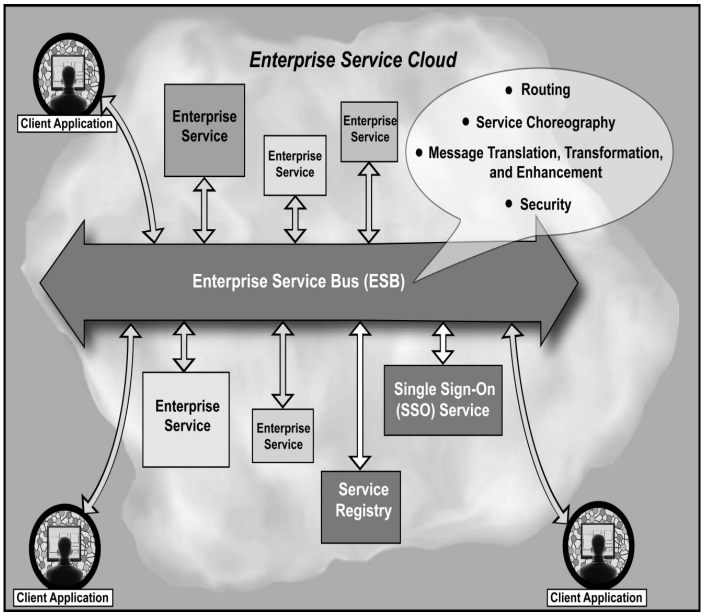

Figure 1: Principal SOA components

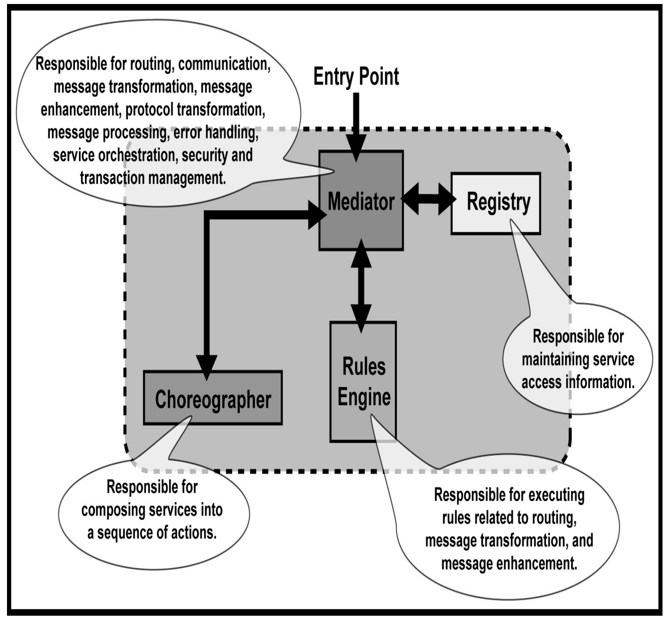

Figure 2: Principal ESB components

There are quite a number of commercial off-the-shelf ESB implementations that satisfy these specifications to varying degrees. A full ESB implementation would include four distinct components (Figure 2): Mediator; Service Registry; Choreographer; and, Rules Engine. The Mediator serves as the entry point for all messages and has by far the largest number of message management responsibilities. It is responsible for routing, communication, message transformation, message enhancement, protocol transformation, message processing, error handling, service orchestration, transaction management, and access control (security). The Service Registry provides the service mapping information (i.e., the location and binding of each service) to the Mediator. The Choreographer is responsible for the coordination of complex business processes that require the participation of multiple service providers. In some ESB implementations the Choreographer may also serve as an entry point to the ESB. In that case it assumes the additional responsibilities of message processing, transaction management, and access control (security). The Rules Engine provides the logic that is required for the routing, transformation and enhancement of messages.

3.2 Information-centric representation

The methods and procedures that we human beings utilize to make decisions and solve problems can be readily expressed in computer software as objects. For example, logistic planners develop distribution plans by reasoning about inventories, conveyances, routes, distribution centers, delivery windows, priority, weather, security, and so on. Each of these objects encapsulates knowledge about its own nature, its relationships with other objects, its behavior within a given environment, and the various constraints and requirements needed to effectively meet its individual performance objectives.

It is therefore apparent that a critical requirement for effective human-computer interaction in an intelligent supply-chain information management environment is the effective representation of the *context* within which the logistic planning and management activities are taking place. This can be accomplished utilizing an *ontology* (Figure 3). The term ontology is loosely used to describe an information structure that is rich in relationships and provides a virtual representation of some real world environment. As shown in Figure 4, the elements of an ontology include objects and their characteristics, different kinds of relationships among objects, often including the concept of inheritance and other characterizations (e.g., dynamic classification, multiple classification, incremental realization, and so on).

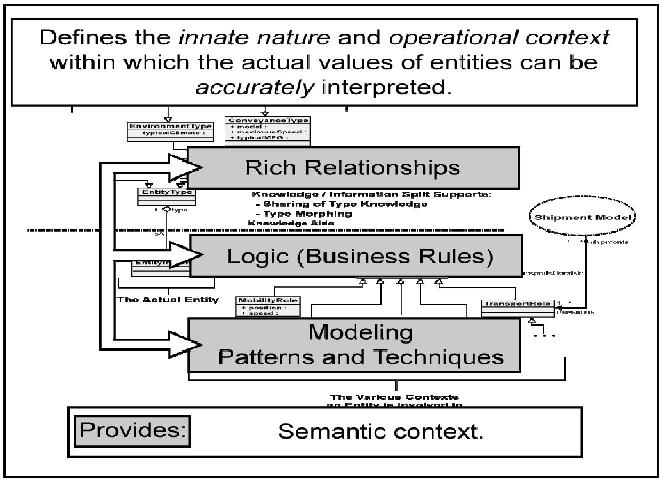

Figure 3: Ontology representation characteristics

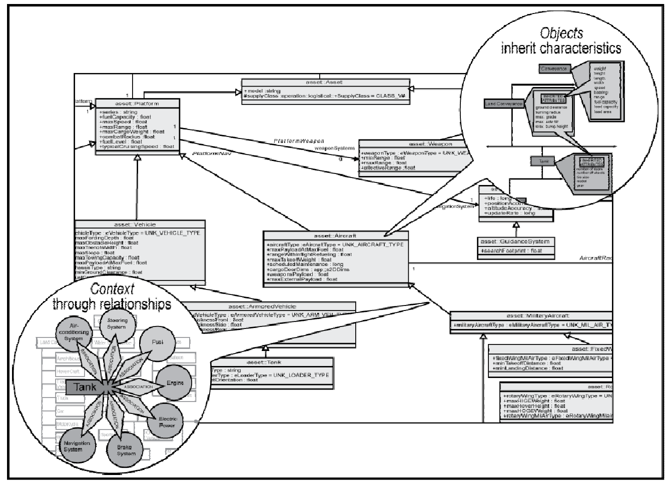

Figure 4: Ontology objects and concepts are machine processable

Since these elements of an ontology in combination with object-oriented computer languages and advanced modeling paradigms can be automatically interpreted by software, a computer-based information management environment can be endowed with at least a simplistic level of *understanding* of the real world *context* within which the required planning and execution decisions are being made. This level of understanding is sufficient to provide the necessary context for software agents to automatically interpret data, develop and evaluate plans, detect and explain the causes of conflicts, and generate alerts.

3.3 Software agents as intelligent tools

In software that incorporates an ontology-based representation of real world supply-chain planning and execution context, the intelligence of the information management environment is largely contributed by the inferencing tools that are available to the human user. Most of these tools will be in the form of services or self-initiating agents. While services are invoked to perform a discrete activity and return to their original inactive state after the activity has been completed, agents may be active on a continuous basis taking the initiative opportunistically whenever they determine that the situation warrants an action. Often these agent actions will invoke services. In this context opportunity is typically defined by the existence of sufficient information. For example, if an unexpected sizeable order for a particular drug exceeds local inventory stocks several agents will become immediately involved to check the inventory of other warehouses, automatically generate a shipment plan, or if no stocks can be found anywhere suggest an alternative product as a possible substitute.

There are many types of software agents, ranging from those that emulate symbolic reasoning by processing rules, to highly mathematical pattern matching neural networks, genetic algorithms, and particle swarm optimization techniques. While all of these have capabilities that are applicable to an intelligent supply-chain environment, the symbolic reasoning agents will normally play the most important role and bring the most immediate benefits when a virtual context model (i.e., ontology) has been constructed. Therefore, only symbolic reasoning agents that can interact directly with the ontology-based context model will be discussed in this paper.

For the following kinds of rule-based agents the reasoning process relies heavily on the rich representation of entities and concepts provided by the ontology.

Service Agents: These software agents, which are designed to be knowledgeable in a specific domain, and perform planning or assessment tasks in partnership with other agents (i.e., human agents or software agents), are often referred to as Service Agents. The manner in which they participate in the decision-making activities depends on the nature of the situation. Service Agents can be designed to react to changes in the problem state spontaneously through their ability to monitor information changes and respond opportunistically.

In an intelligent pharmaceutical supply-chain information management environment Service Agents would have knowledge and analysis capabilities in narrow domains such as inventory assessment, product quality validation, scheduling, weather data interpretation, drug pedigree and expiration date verification, and environmental conditions (e.g., temperature) monitoring. Typical analysis and inferencing characteristics of Service Agents include:

- Ability to generate alerts based on current state analysis.
- Ability to justify alerts, and analysis results with explanation facilities.
- Ability to broadcast requests for services to other agents.
- Ability to automatically generate queries and access data repositories.
- Ability to temporarily clone themselves to process multiple requests in parallel.
- Ability to undertake proactive explorations opportunistically.

Typical examples of Service Agents that would be useful in a pharmaceutical supply-chain are described in Appendix A.

Planning Agents: Planning is a reasoning activity that deals with the availability of resources and the actions that need to be taken to complete a given task. Consequently, Planning Agents are designed to reason about the problem state and produce a plan based on the current state of the supply-chain in conjunction with applicable constraints and objectives. The planning process involves matching constraints with available resources to produce a course of action that will satisfy the desired objectives.

Consequently, in a supply-chain environment Planning Agents deal with the broader issues of demand forecasting, maintaining inventory levels, assessing risk, measuring efficiency, and estimating order of magnitude costs. Typical analysis and inferencing characteristics of Planning Agents include:

- Ability to task Service Agents and request information from Mentor Agents.
- Ability to orchestrate evaluations involving several Service Agents.
- Ability to generate broad current state assessments on request or by alert.
- Ability to act on directions from human users and Coordination Agents.

Typical examples of Planning Agents for pharmaceutical supply-chain functions such as distribution planning, cost estimating, risk assessment, efficiency measurement, and opportunity recognition are described in Appendix B.

Mentor Agents: The purpose of a Mentor Agent is to temporarily provide a passive data element with active capabilities such as communication and limited self-determination. Mentor Agents are created either by human users or by Coordination Agents on a temporal basis to track a particular supply-chain object such as a requisition, container, pallet, or conveyance that is of special interest (Figure 5).

The concept of Mentor Agents brings several potential benefits. First, it increases the granularity of the active participants in the problem solving process. As agents with collaboration capabilities, agentified data elements can pursue their own objectives and perform a significant amount of local problem solving without repeatedly impacting the communication and coordination facilities utilized by the higher level components of the distributed supply-chain system. Typically, a Mentor Agent is equipped with communication capabilities, process management capabilities, information about its own nature, and objectives. Second, the ability of Mentor Agents to task Service Agents greatly increases the potential for concurrent activities. Multiple Mentor Agents can request the same or different services simultaneously.

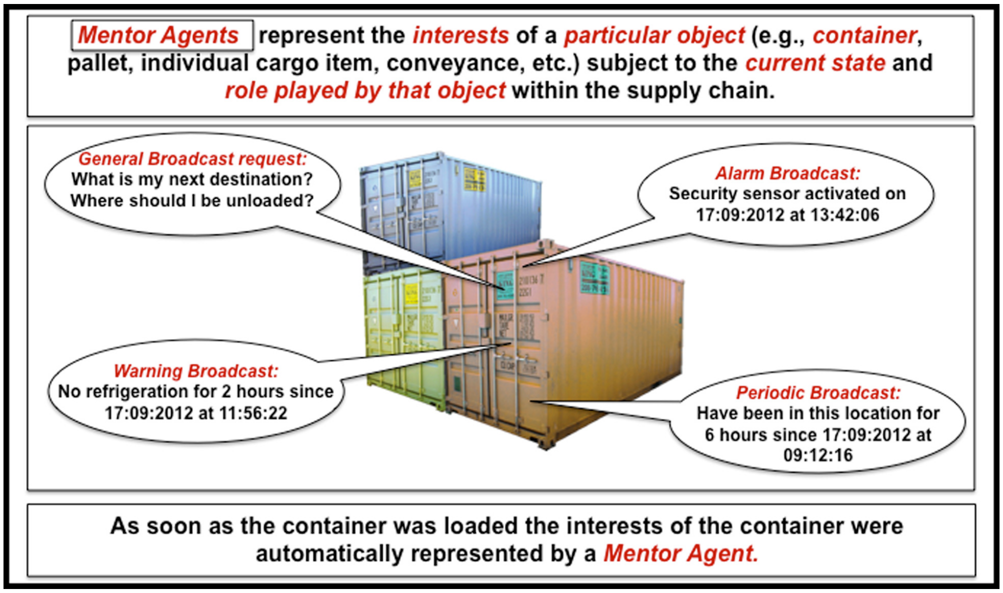

Figure 5: Mentor Agent representing a particular *container* in a shipment

Third, groups of Mentor Agents can negotiate among themselves in the case of matters that do not directly affect other higher level components or as a means of developing alternatives for consideration by higher level components. Fourth, by virtue of their communication facilities Mentor Agents are able to maintain their relationships to other aspects of the current state of the supply-chain. In this respect they are the product of *decentralization* rather than *decomposition*. In other words, the concept of Mentor Agents overcomes one of the most serious deficiencies of the rationalistic approach to problem solving; - namely, the dilution and loss of relationships that occurs when a complex problem is decomposed into sub-problems. In fact, the relationships are greatly strengthened because they become active communication channels that can be dynamically created and terminated in response to the changing state of the problem situation.

In summary, the capabilities of a Mentor Agent that is temporarily created in an intelligent pharmaceutical supply-chain environment would normally include one or more of the following:

- Some understanding of its needs as derived from the context model (i.e., ontology).
- Ability to orient itself geographically and geometrically (i.e., location).
- Ability to communicate and request services from Service Agents.
- Ability to communicate and negotiate with other Mentor Agents.
- Ability to pursue interests proactively leading to alternative recommendations.

Coordination Agents: This group of agents is responsible for facilitating collaboration among human users and software agents. Consequently Coordination Agents require the most intelligence because they need to be able to assess the impact of decisions in individual domains on the particular course of action under consideration (e.g., distribution plan), as well as the overall problem space (e.g., demand signal analysis and planning).

Particularly in a supply-chain planning and management environment the most important and demanding role of Coordination Agents is to facilitate collaboration by activating agents and alerting human users of the need for interaction. Another important function of Coordination Agents is the recognition of conflicts.

The emphasis here is on the detection and identification of the causes of a conflict by the agent rather than its resolution. The resolution of a conflict usually involves higher level decisions that have the potential for impacting other areas of the supply-chain. Therefore, apart from very mundane conflicts that could be resolved automatically, the human user should at least be provided with an opportunity to resolve conflicts with wider consequences.

Typical examples of Coordination Agents for pharmaceutical supply-chain functions such as collaboration, conflict detection and analysis, threat assessment, recall management, and the coordination of distribution plans are described in Appendix C.

Control Agents: Corporate guidance and governance is the province of Control Agents. They are concerned with the measurement of performance, the prevention of security breaches (e.g., theft), the monitoring of priorities such as patient safety, regulatory compliance, and the identification of supply-chain trends.

The role of Control Agents to identify trends warrants further discussion. The detection of supply-chain trends is almost exclusively considered to be a human role in existing supply-chain planning and management networks. As a result, due to the large number of transactions that are involved in sizable supply-chains and the dynamically changing nature of the execution phase of operations, many opportunities for proactive planning are overlooked. Particularly under major unforeseen demand conditions that might occur immediately following a natural disaster, or when unforeseen events seriously disrupt distribution plans, the human decision-maker is forced into a reactive mode.

Control Agents with access to pattern matching tools such as neural networks can provide powerful trend detection capabilities. Since such tools are able to operate unobtrusively in background on a continuous basis they are able to address the following kinds of questions that are of interest at the executive level of supply-chain management:

- What quantity of any particular product or drug type has been delivered to a specific geographic region or location over a given time period?
- What were the principal choke points where raw material deliveries, production sequences, or distribution shipments have been delayed during a given time period?
- What has been the average time that recalls or certain kinds of pilot programs, promotions, and distribution shipments have taken over a given time period?
- What have been the relative geographic densities of sales of specific products or certain kinds of drugs over a given time period?
- What (if any) are the relationships between customer sentiment gleaned from social media sites and product sales?

Typical examples of Control Agents that would be appropriate for an intelligent pharmaceutical supply-chain are described in Appendix D.

3.4 The system environment

Conceptually, as shown in Figure 6, the logistical context provided by a multi-layered ontology allows the various groups of agents to monitor and act on the data that flows on a continues basis through the pharmaceutical supply-chain.

The primary functions of the Planning Agents are focused on the organization of task sequences into schemes devised to meet certain objectives. This will often involve the generation of alternatives that are then evaluated by mostly Service Agents in respect to cost, resource constraints, risk, and opportunities for improving efficiency and reducing costs. During execution the focus shifts from planning to re-planning when the original planning assumptions have changed within the dynamics of the supply-chain environment.

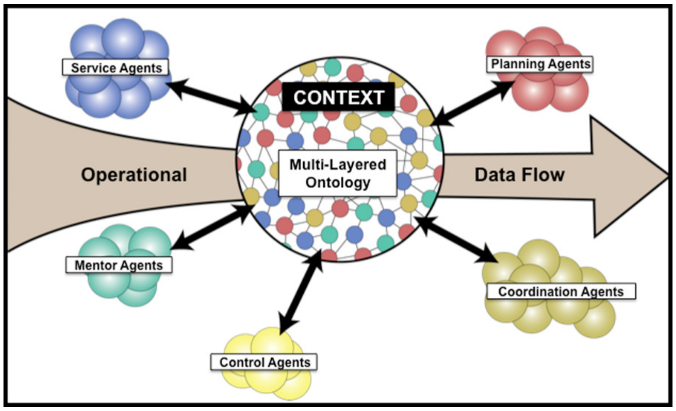

Figure 6: Context-based intelligent tools

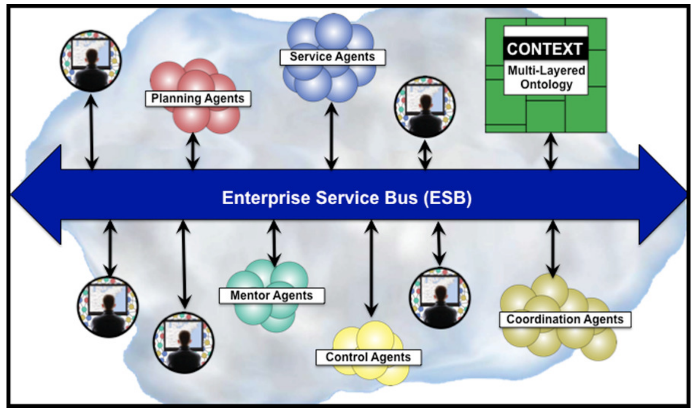

Figure 7: SOA-based system architecture

The Coordination Agents are responsible for facilitating collaboration, exploring the availability of resources such as raw materials for production and inventory stock levels of products. Any major undertaking such as the framing and execution of a recall program, pilot study or distribution plan will be coordinated by the Collaboration Agent. If there are several concurrent major undertakings, which is highly likely, the Collaboration Agent will be automatically cloned so that each project can receive the attention that it deserves. In this way the demand for large quantities of particular drugs that might follow a natural disaster or the outbreak of an epidemic cant be coordinated without disrupting ongoing production, promotion, marketing, and distribution operations.

For example, if the Opportunity Agent identifies that two distribution plans have overlapping features such as geographical product origins and customer destinations then the Collaboration Agent will immediately explore the possibility of combining transportation modes and arrangements. This exploration may involve several Service Agents such as the Inventory, Transport, Scheduling, and Physical Conditions Agents to determine whether the modification of the original plans is feasible and cost effective.

What is significant is that all of these actions can be undertaken automatically and concurrently for hundreds of planning, execution and trend analysis tasks on a continuous near real-time basis. When events that have the potential for disrupting the supply-chain occur the human users have the necessary tools and actionable information available to take immediate and effective action. At the same time the Control Agents are systematically analyzing past events and actions with a view to identifying patterns and trends within the supply-chain.

The system implementation framework is based on SOA principles (Figure 7), with interaction among the various loosely coupled applications and services managed transparently to the human users by an ESB. While many of the agents operate concurrently in an opportunistic mode, the workflow of logistical operations is essentially sequential in character. In a SOA-based system environment the orchestration of such sequences is normally performed by a Business Process Management (BPM) facility.

> ***Business Process Management (BPM):*** BPM is a method for actively defining, executing, monitoring, analyzing, and subsequently refining manual or automated business processes. In other words, a business process is essentially a sequence of related, structured activities (i.e., a workflow) that is intended to achieve an objective or larger task. Such workflows can include interactions between human users, software applications or services, or a combination of both. In a SOA-based information management environment this orchestration is most commonly performed by the Choreographer component of the ESB (Figure 2). Based on SOA principles, a sound BPM design will decompose a complex business process into smaller, more manageable elements that comply with common standards and reuse existing solutions.

While BPM and SOA concepts are closely connected, they are certainly not synonymous. Rather, they are complementary. Described more precisely, a SOA-based system environment provides the enabling infrastructure for BPM by separating the functional execution of the business process from its technical implementation. Conversely, BPM offers even the most well architected inventory of SOA functionality (i.e., services) specific objectives. The business process models identified as part of the BPM approach will effectively align the software capabilities produced to the actual needs of the users. Too often enterprises suffer from a distinct mismatch between available software functionality and actual user needs.

In addition to those components discussed above, an effective logistics decision-support environment includes a number of other principal components, including:

- A web-based application portal that provides the human user with an integrated, highly interactive canvas (i.e., view) across what may otherwise be a disparate collection of services, information sources (e.g., GIS, databases, etc.), intelligent agents, and external systems. Further, benefiting from the strong presence of BPM principles and functionality complementing the overarching SOA-based enterprise, this rich user-interface is purposefully organized around the very business processes that are relevant to the specific role of the human user. In other words, orienting the various flavors of the user-interface around relevant business processes provides specific users with a graphical, highly-interactive (essentially customized) user-interface that is designed and engineered in terms of the very workflows, terminology, and practices that comprise that user's tasks, objectives, and practices (i.e., business processes). The result is a convenient, highly efficient dashboard that fosters an effective partnership between the human users and the software capabilities designed to assist them.
- An ontology service that builds, maintains, and exposes its evolving context to agents and other services that are context-dependent. Such informational services can support synchronization of interested clients (i.e., software agents or human users) with changes occurring within the context they manage via asynchronous service requests that can live for extended periods of time. The result is a means by which clients can subscribe to, and consequentially be notified of, particular events and conditions of interest as they occur.
- An inference service that may comprise a number of agent communities. An agent community is a collection of related agents in a given domain such as the Planning Agents, Mentor Agents, Service Agents, Coordination Agents, and Control Agents described in Section 3.3. Each agent utilizes applicable ontology services and other types of services to examine and analyze the current state of a particular transaction sequence or larger supply-chain context.

3.5 The user environment

From the human user's point of view the intelligent pharmaceutical supply-chain environment described in this paper is highly interactive and proactive. Not only are the users able to conduct searches for data where the search keys are known (i.e., directed searches) but they are also able to conduct semantic searches when the queries can be only vaguely formulated. In those cases agents with data mapping capabilities will search through one or more databases and return to the user approximately matching query results with computed certainty factors.

At the same time the user is automatically alerted to both opportunities for taking advantage of events that could lead to greater efficiency or lower costs and events that either are already or could potentially disrupt the supply-chain. Since agents are continuously monitoring both the flow of transactions and external events that stimulate the operational environment, as well as the planning and execution sequences that are triggered by this flow, many of the opportunities for effective intervention that are likely to be overlooked in current data-centric management systems will be brought to the attention of the human user through agent warnings and alerts.

In this respect an intelligent supply chain information management environment is both reactive and proactive. For example, if any particular recall process, pilot program, or distribution plan is running behind schedule then this will be noted and recorded in a *warning* report by an agent. If a product shortage has arisen or is likely to arise in the near future due to a particular event then this will be noted by an agent and the user will receive an *alert*. However, agents are also continuously analyzing past events, transactions and actions to identify patterns and trends, so that these can be related to current or expected near term conditions within the supply-chain. This type of analysis may involve multiple Control, Coordination, Planning, and Service Agents, with the objective of identifying potential supply-chain events and disruptions proactively. For example, repeated shipment thefts, disruptions, or delays in a particular region may suggest the need for considering an alternative inter-modal distribution plan.

Data access: As suggested by the severity ratings assigned in Table 1, much of the management time in the pharmaceutical industry is spent on determining the current state of a largely fragmented supply-chain that lacks visibility of resources (e.g., inventory stocks), reliable demand forecasting capabilities, operational efficiency, ability to deal effectively with external influences, and communication across functional domains.

In the intelligent supply-chain environment described in this paper a capability, such as visibility, comes as a by-product of the ontology-based context model that represents the concepts, principles and entities that the human operator reasons about as objects with characteristics and relationships. This allows the human user to lodge queries about any particular plan, product, or transaction and pursue such queries to reasonable depth, with the objective of receiving answers to the following kinds of questions: Where is this shipment right now? What quantities of this product are currently in inventory and where? When will the following milestone be reached in this pilot program?

Figure 8: Displayed symbols are objects

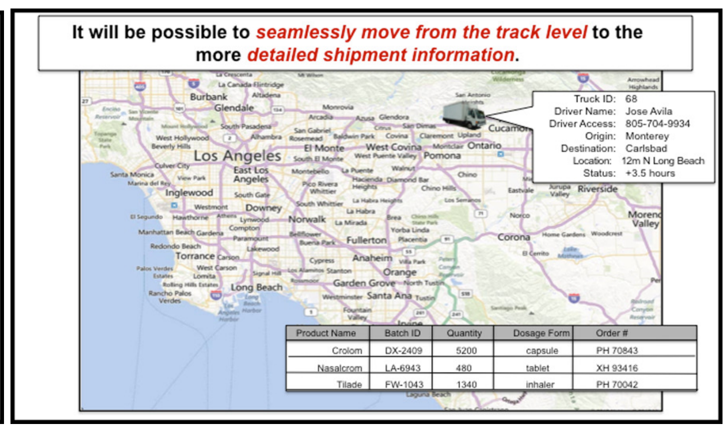

Figure 9: Information on request

The need for visibility also extends to resources that are in-transit. As shown in Figure 8, to obtain information about any of the symbols displayed on the map the user simply clicks on the particular symbol (e.g., conveyance, distribution center icon, city, or route) displayed on the dashboard user-interface. A second mouse click allows the user to drill down to more detailed information. For example, in Figure 9 the user is able to seamlessly move from the summary information relating to the current location, destination, priority, and expected delivery window of a shipment, to the details of the individual contents of the truck.

Not only are users able to search on multiple keys such as product name, active ingredient, dosage form, strength, location, and so on (Figure 10), but they can also conduct semantic searches. As shown in Figure 11, the user may describe the required information in fairly vague terms when the exact identification of the item is not known. For example, the user may know only the disease and restrict the search to Continental

United States (CONUS) distribution centers and a particular dosage form, as well as specifying an expiration date and the price ceiling.

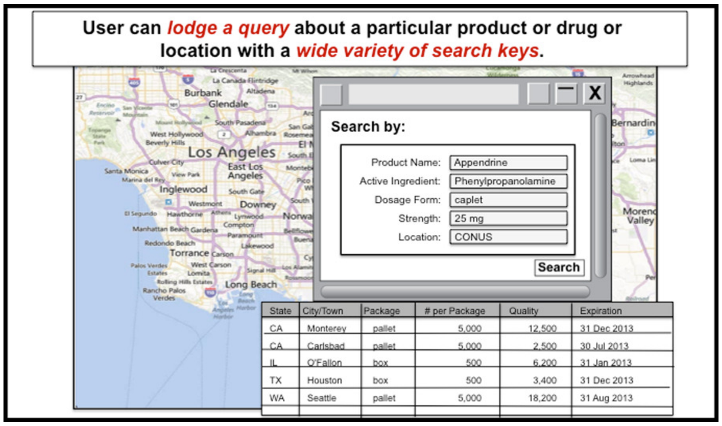

Figure 10: Ability to search on multiple keys

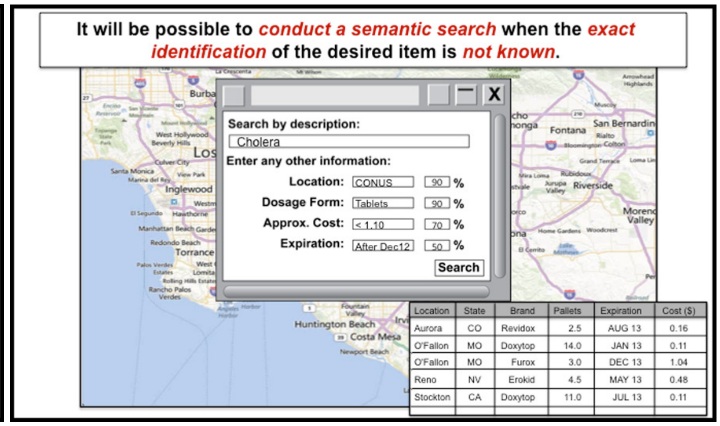

Figure 11: Search with partial information

Similarly, either by clicking on a displayed graphic symbol or by employing direct or semantic search capabilities the user is able to obtain a summary of the inventory of all of the distribution centers in a particular geographical region or drill down to the current inventory of a particular distribution center. The same data is of course also available to agents. However, in this case the required queries are generated automatically for use in the evaluation of alternative plans, the assessment of risk, the determination of costs, and any other supply-chain information management task that the particular agent may be engaged in.

Impact of external factors: Both the formulation and execution of shipment plans is impacted by external factors such as weather conditions, customs requirements at border crossings or points of debarkation in foreign countries, location of criminal or enemy activities, availability of indigenous transportation, terrain, traffic conditions, and so on. In this respect an intelligent toolset is able to accept several on-line data feeds and combine the imported data with sufficient context to allow agents to automatically reason about the implications of the external factors. Candidate data feeds include:

- Weather forecasts on a regional and local level. The Weather Agent is able to translate raw weather data into a weather report that provides actionable information to a human user and is machine processable for inferencing purposes by other software agents.
- Indigenous transportation systems and services (e.g., major roads, railways, ferries, commercial airline routes, third party logistic services providers (3PL)) in regions and local areas that are available as shipment alternatives.
- Supplies, conveyances, fuel, and related transportation resources available at transportation hubs and distribution centers.
- Location of criminal and/or political activities that pose significant security threats.
- Infrastructure objects such as distribution centers, warehouses, railway stations, ferry stations, airports, ocean ports, fuel stations, and so on.

Pattern recognition: As the scale of the adaptive toolset progressively encompasses a more significant portion of the supply-chain enterprise the intelligent agents will have access to an increasingly larger set of historical data. This will allow the implementation of agents with sophisticated analysis and case-based reasoning capabilities. Such analytical tools, operating in a collaborative manner, will be able to analyze the efficiency, as well as the occurrence and impact of events in past undertakings such as recalls, pilot programs, thefts and the effectiveness of security countermeasures, foreign shipments and custom procedures, product demand patterns and their influence on production plans and inventory levels, impact of changes in government regulations on demand, efficiency and profitability, identification of counterfeiting patterns, and influence of social media strategies and investments on sales. In particular, answers to the following kinds of questions can be generated automatically:

- What quantities of any particular product have been delivered over a given time period, what shortages are likely to arise, and when?
- What were the principal causes of delays in shipments of products to foreign countries?
- What were the principal choke points where shipments have been delayed during a given time period? Where are choke points likely to occur in the future based on current market forecasts?
- Where have shipments been intercepted by criminal or political action (in foreign countries) over a given time period and what are the risk factors that should be applied to future shipments?
- What has been the average time that certain kinds of shipments have taken over a given time period and how do these times relate to future distribution plans?
- What have been the relative densities of air, surface and rail movements over a given time period and how do these densities relate to supply-chain performance?

3.6 Agent collaboration and decision-assistance

Historically, computer-based data-processing systems have been designed to be activated and controlled by human users. In this respect they may be characterized as *passive* decision-assistance environments that with few exceptions respond only when tasked by a human user. For example, the user enters the requirements for certain goods to be shipped between two geographical locations and a movement plan is either interactively formulated or automatically generated if more sophisticated tools are available. In other words, the user directs the system to assist in some predefined manner and the system generates the appropriate response or result to the best of its capabilities. If the users do not request the system to undertake any tasks then the system will be essentially idle.

A context-based (i.e., information-centric) software system with inferencing capabilities provided by agents is in contrast an *active* decision-assistance environment in which data cleansing, monitoring, analysis, planning and re-planning, pattern identification, and exploratory processing will occur on an on-going basis. In fact, under certain circumstances the system environment may be intensely active while the human users are largely inactive. The activities of the system environment are activated at least as much by the data that flows through the system on a continuous basis (Figure 6) as by the interactions of the human users with the system environment. This is largely made possible by the virtual model (i.e., multi-layered ontology) of the real world supply-chain context that allows the agents to autonomously and concurrently interpret and analyze the data flow in the appropriate context.

As an example of a typical sequence of supply-chain management events we will assume the following scenario as the direct consequence of a natural disaster. A major earthquake has occurred in North Africa, with severe devastation in Morocco, and the onset of a typhoid epidemic has created an urgent demand for particular drugs. Initial laboratory tests have identified a multidrug-resistant (MDR) strain of typhoid. One of the company's products, which we will refer to as X-12 for convenience, is known to be effective against this MDR strain. Consequently an urgent order (i.e., *Requisition*) for 20 pallets of X-12 has been received from the Red Cross, for immediate shipment to Casablanca. As soon as the *Requisition* enters the information management system the various software agents are automatically triggered to respond.

As shown in Figure 12, the Priority Agent sends an *alert* to the Collaboration Agent suggesting that coordination will be necessary due to the high priority of the request. The Collaboration Agent promptly broadcasts its coordinating role in the Morocco Project to all agents and starts monitoring the requisition (Figure 13). It immediately requests the Disaster Agent to ascertain the company's current world-wide inventory of X-12. The Disaster Agent concurrently invokes the Inventory Agent and the Location Agent to determine the quantity of X-12 supplies in stock and whether there are any X-12 shipments in-transit. In addition, the Location Agent queries the Closure Agent to see if any X-12 supplies have been delivered to Africa within the past six months. While the Inventory Agent is responsible for tracking supplies that are currently stored in distribution centers and warehouses, the role of the Location Agent is to keep track of in-transit and recently delivered supplies (see definition of Service Agents in Appendix A).

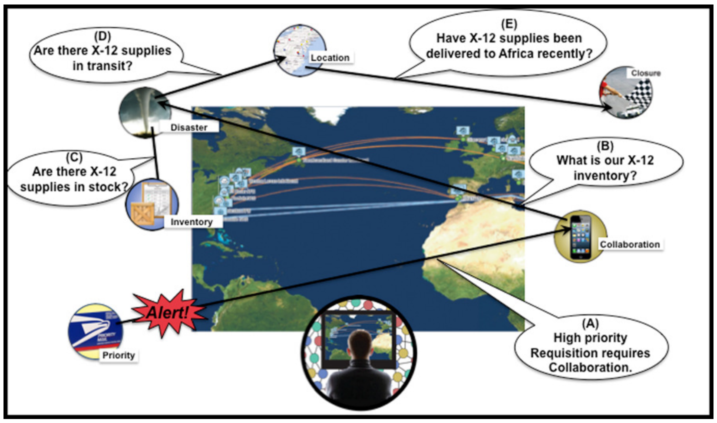

Figure 12: Are the requested X-12 supplies available in inventory?

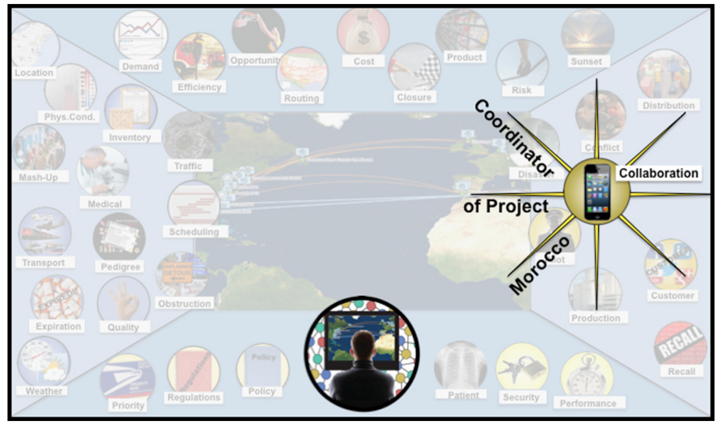

Figure 13: Collaboration Agent alerts all agents to its coordinating role in the Morocco project

In Figure 14, the Collaboration Agent determines on the basis of the reports received from the Inventory Agent and the Location Agent that there are only 10 pallets of X-12 in stock and another 5 pallets in the field that could be potentially available. Specifically, the Inventory Agent has reported that there are a total of 6 pallets of X-12 in a Seattle warehouse and 4 pallets in a Charleston distribution center. Concurrently the Location Agent has determined that there are 3 pallets of X-12 on-board a cargo ship headed for Dakar (Senegal, Africa) and 2 pallets of X-12 were delivered to a warehouse in Lagos (Nigeria) three weeks ago. This detailed information is reported by the Collaboration Agent to the user, who directs the Disaster Agent to explore the feasibility of diverting the in-transit shipment of X-12 (3 pallets) to Casablanca and orders the Collaboration Agent to arrange for the shipment of 1 pallet of X-12 from the Lagos warehouse to Casablanca.

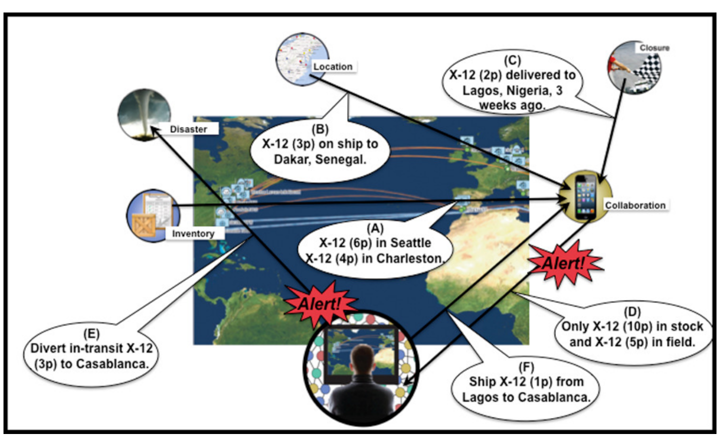

Figure 14: Maximum of 15 pallets of X-12 available

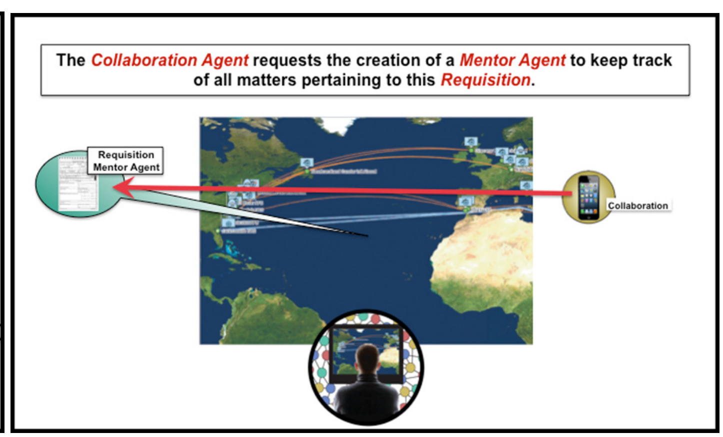

Figure 15: Mentor Agent is assigned to *Requisition*

In the meantime, the Collaboration Agent requests the creation of a Mentor Agent for this *Requisition* (Figure 15). The Mentor Agent keeps track of all matters pertaining to the *Requisition* such as: order details; origin and destination of supplies; shipment arrangements; and, current status.

The user is now aware that there are insufficient supplies of X-12 available to fulfill the Red Cross order. Even if the company's policy on maintaining minimum reserve stock levels were to be violated and all of the X-12 supplies recently shipped to Lagos in response to local orders were to be reallocated, there would still be only 15 pallets of X-12 available. Consequently, as shown in Figure 16, the user now directs the Medical Agent to advise whether the company has other products that could substitute for X-12. The Medical Agent recommends another product, Y-24, as an acceptable substitute for X-12. The user then directs the Collaboration Agent to use Y-24 as an alternative to X-12. This immediately triggers a sequence of interactions to determine the availability of Y-24 supplies. The Disaster Agent queries the Inventory Agent for Y-24 stock levels and the Location Agent for in-transit and recently African deliveries of Y-24.

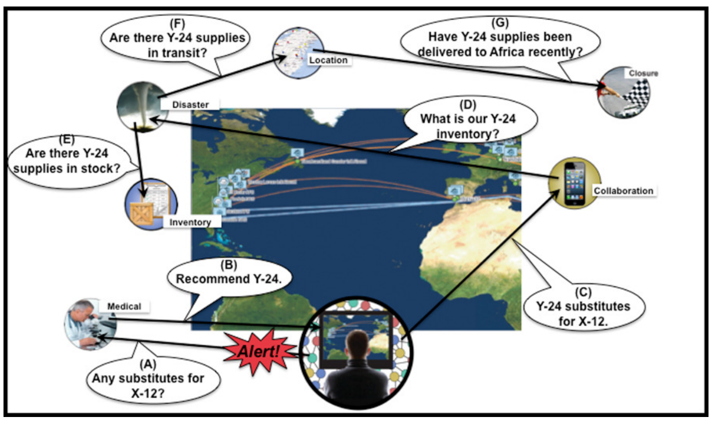

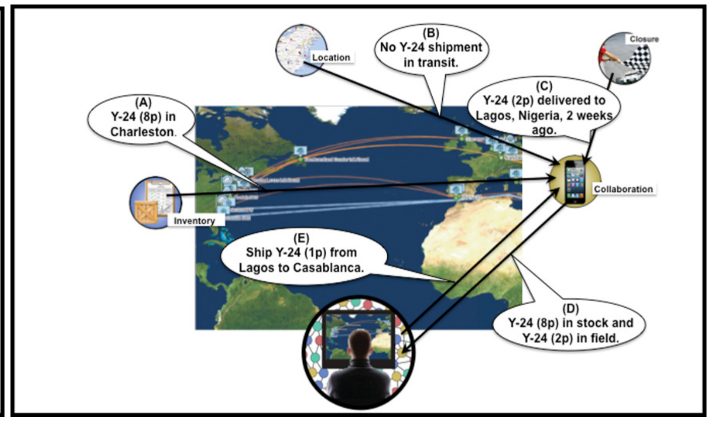

Figure 16: Product Y-24 can substitute for X-12   Figure 17: Maximum of 10 pallets of Y-24 available

Based on the reports provided by the Inventory, Location and Closure Agents the Collaboration Agent is able to report to the user (Figure 17) that there are 8 pallets of Y-24 in the Charleston distribution center and that 2 pallets of Y-24 were included in the recent shipment of X-12 to the Lagos warehouse.

As shown in Table 2, in addition to order details, customer information, shipment arrangements, and current status, the Mentor Agent now holds the following information about the potential source of X-12 and Y-24 supplies for the *Requisition*:

Table 2: Source of supplies information held by the *Requisition* Mentor Agent

Current Location	Total Inventory	Mandated Reserves	Immediately Available
Seattle	X-12 (6 pallets)	X-12 (1 pallet)	X-12 (5 pallets)
Charleston	X-12 (4 p); Y-24 (8 p)	X-12 (1 p); Y-24 (1 p)	X-12 (3 p); Y-24 (7 p)
In-Transit	X-12 (3 p)	---	X-12 (3 p)
Lagos	X-12 (2 p); Y-24 (2 p)	X-12 (1 p); Y-24 (1 p)	X-12 (1 p); y-24 (1 p)

[Mandated reserve level is 1 p (p = pallet)]

The Collaboration Agent is now ready to make the necessary shipping arrangements (Figure 18). It directs the Transport Agent to organize the shipping of 5 pallets (X-12) from Seattle and 10 pallets (i.e., 3 X-12 pallets and 7 Y-24 pallets) from Charleston to Casablanca. In addition, it directs the Routing Agent to prepare surface transportation plans for shipping 2 pallets (i.e., 1 X-12 pallet and 1 Y-24 pallet) from Lagos by road to

Casablanca. The Routing Agent asks the Risk Agent to assess the risk of road transport in Africa between Lagos and Casablanca, which in turn also involves the Obstruction Agent in respect to potential road closures. Even before receiving a response from the Risk Agent both the Transport Agent and the Routing Agent request the Cost Agent to provide estimates of the cost of shipping 10 pallets by airfreight from CONUS to Casablanca and 2 pallets by 3PL from Lagos to Casablanca, respectively.

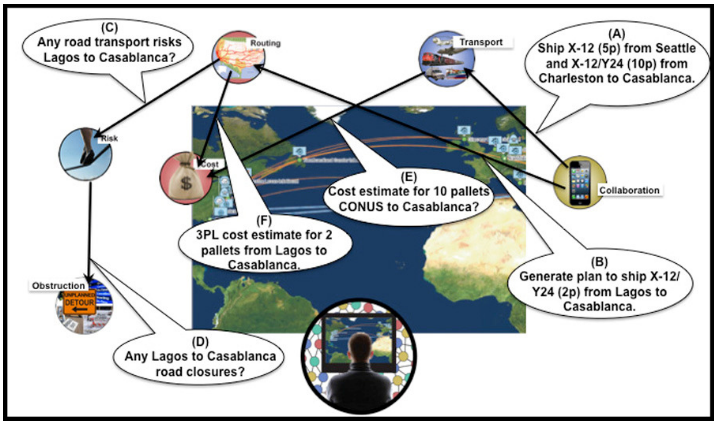

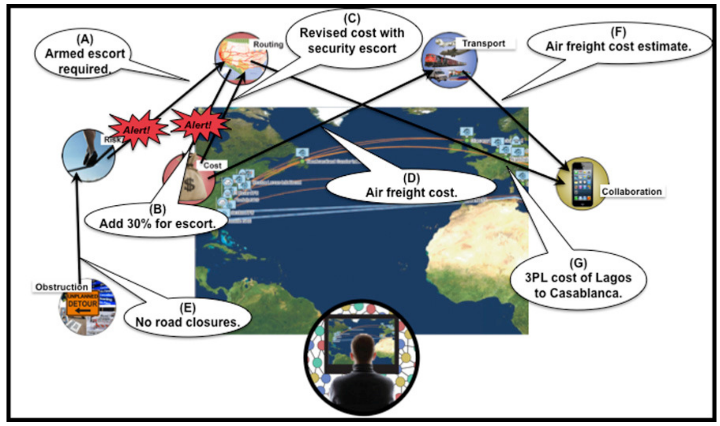

Figure 18: Planning transportation to Casablanca

Figure 19: Escort requirements increase 3PL costs

In Figure 19, the Risk Agent informs the Routing Agent that road transport in Africa will require an armed escort. At the same time the Obstruction Agent verifies that there are currently no known road closures on the route between Lagos and Casablanca. The estimates for airfreight and escorted 3PL transport are passed onto the Collaboration Agent by the Transport and Routing Agents, respectively.

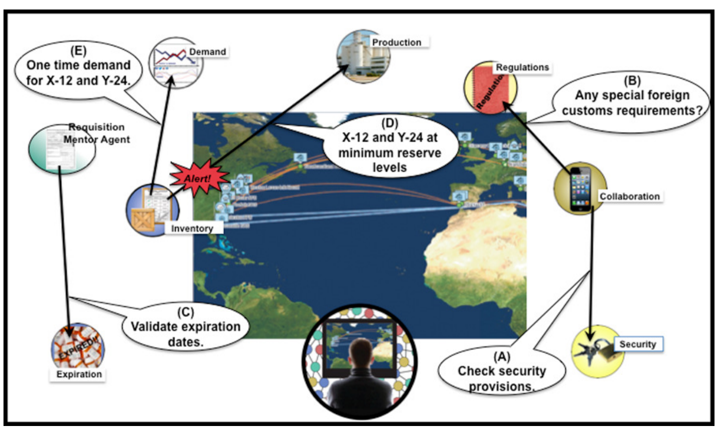

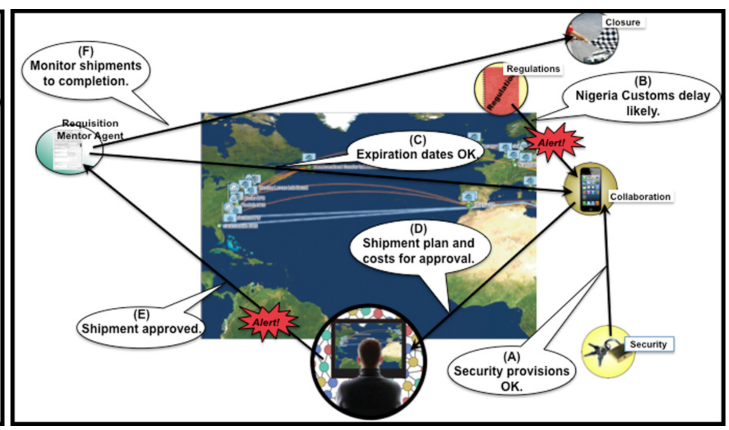

Figure 20: Pre-approval checks and inventory alert

Figure 21: Final decision in minutes instead of days

The Collaboration Agent now conducts the final checks prior to sending the complete shipment plan to the user for approval. As shown in Figure 20, while the Collaboration Agent requests the Security Agent to review the plan's security provisions and the Regulations Agent to report back on Morocco and Nigeria Customs requirements, the Mentor Agent asks the Expiration Agent to confirm that the expiration dates of the X-12 and Y-24 batches are in order. In the meantime the Inventory Agent sends an alert to the Production Agent indicating that X-12 and Y-24 inventory levels are now at a minimum. However, to avoid future overstocking the Inventory Agent also advises the Demand Agent that the sudden increase in demand for the X-12 and Y-24 drugs is an anomaly caused by a natural disaster and should therefore not trigger a revision of the future demand forecasts for these two products.

In Figure 21 several confirming reports are received by the Collaboration Agent. The Security Agent responds that the security arrangements, including the armed escort provisions for the surface shipment of the two X-12 and Y-24 pallets from the Lagos warehouse to Casablanca, are adequate. Also, the Mentor Agent having received a positive response from the Expiration Agent is able to report that the expiration dates are in order. However, the Regulations Agent sends an alert indicating that there may be Custom delays in Lagos since the Nigerian authorities sometimes insist on confirming not only the authenticity of medical drugs coming into their country but also drugs moving out of Nigeria. The Collaboration Agent now sends the shipment plan together with cost estimates to the user for action. The user approves the plan by sending an alert to the Mentor Agent, which as a final action directs the Closure Agent to monitor the execution of the shipment plan to completion.

What is particularly noteworthy about this scenario is that most of the tasks that are currently being performed manually by human planners and subject matter experts in existing supply-chain operations can be automated. The benefits of the automation are at least threefold. First, decisions that can take days and weeks when they have to be performed manually can be made in minutes. Second, the underlying information (e.g., stock levels, cost estimates, regulation checks, and risk assessments) that has to be collected before an informed decision can be made is likely to be more reliable due to uniformity and accuracy. Third, since all of the tasks are performed within the system environment they are automatically archived. The continuous analysis of this valuable data source provides a basis for trend detection and the automated generation and update of demand, inventory, production, and risk projections.

The counterargument that might be made by skeptics who believe that these scenarios are more wishful thinking than based on the realities of the current state of information technology is that the ability of the analytical components (i.e., the agents) to perform effectively is predicated on the availability of relevant data. For example, the Routing Agent in the above scenario cannot prepare a useful plan for shipping products from Lagos in Nigeria to Casablanca in Morocco without access to data relating to transportation alternatives, 3PL shippers, and road conditions. In a fragmented legacy system environment this information is held in multiple unconnected databases and human minds. On the other hand, an integrated system environment that is used by all stakeholders encourages the sharing of data because it is quickly recognized that such data sharing increases efficiency, reduces operational costs, and therefore becomes a major profitability factor. Furthermore, external information relating to weather, traffic, road conditions, infrastructure, and news-worthy events has over the past decade become increasingly available in the public domain on Internet websites that can be easily accessed by a computer-based information system. To exploit the global connectivity and computer-based capabilities of the information age it is necessary for the pharmaceutical industry to rapidly transition from its largely fragmented data-processing systems to a fully integrated, intelligent information management environment.

3.7 Execution scenario examples

The following two execution scenarios will serve as examples of how an intelligent supply-chain information management environment can not only detect and automatically react to unforeseen events, but also use the collected information to automatically modify demand planning proactively.

In the first execution example (Figure 22) we will assume that while packing the Y-24 batches in Lagos in preparation for their shipment to Casablanca it is discovered that the expiration dates on two of the batches are so far in the future that they cannot possibly be correct. As soon as this alarming information is entered into the system the Collaboration Agent sends an alert to the Security Agent to check whether these batches could be counterfeit? This immediately generates a sequence of interactions involving the Quality and Pedigree Agents. The Pedigree Agent verifies that the audit trail of batches XA-968 and XA-985 of product Y-24 are in fact authentic and communicates this information to the Collaboration Agent. On receiving this assurance the Collaboration Agent directs the Expiration Agent to initiate correction of the expiration date and after being notified by the Expiration Agent that the corrected expiration date is adequate for the Morocco Project informs the user that the expiration data has been corrected.

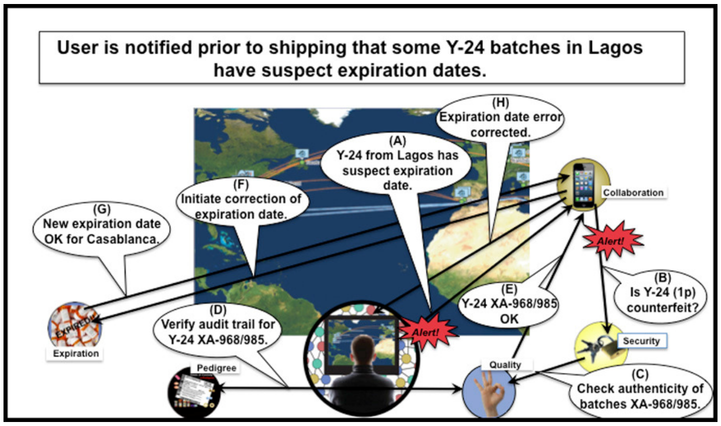

Figure 22: Suspect expiration dates detected

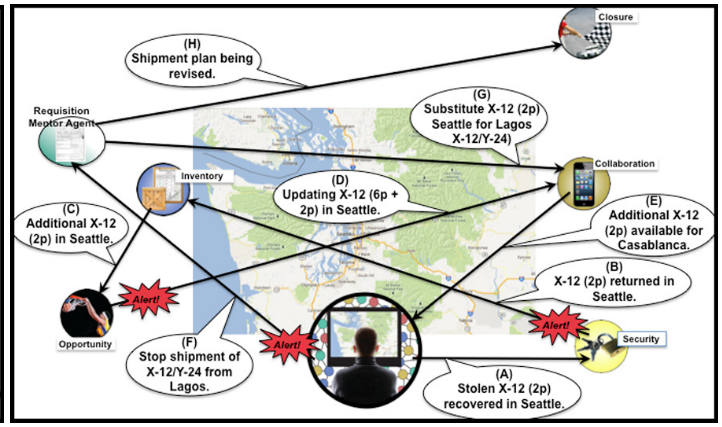

Figure 23: Police recovered stolen X-12 in Seattle

In the second execution example depicted in Figure 23, the company is informed by the Seattle police that two stolen X-12 pallets have been recovered in Seattle. The user enters this information into the system and immediately the Security Agent sends an alert to the Inventory Agent to update the current stock level of X-12 in Seattle (i.e., from 6 pallets to 8 pallets). This action is noted by the Opportunity Agent which now alerts the Collaboration Agent. On being informed of this event the user decides to take advantage of the sudden increase in X-12 stock and directs the Mentor Agent to cancel the shipment of the two pallets of X-12/Y-24 from Lagos to Casablanca. Instead the Collaboration Agent will now direct the Transport Agent (not shown in Figure 23) to increase the airfreight shipment from Seattle to Casablanca by 2 pallets of X-12. The Mentor Agent advises the Closure Agent that the shipment plans for the Morocco Project are being revised.

To complete this scenario; - on request of the Transport Agent the Cost Agent will revise the cost estimate for shipping 12 pallets instead of 10 pallets from CONUS to Casablanca by airfreight and transmit the revised cost estimate to the Collaboration Agent. The latter will send the modified shipping plan to the user for approval. The user will approve the new plan and send an alert to the Mentor Agent, which informs the Closure Agent of the updated shipping plan. The user is able to acknowledge with satisfaction that rapid and decisive action has allowed the company to save the costly and risky overland shipment of 2 pallets of drugs from Lagos to Casablanca.

5. Conclusions

The high complexity of planning and management tasks in the global pharmaceutical supply-chain is due to the multitude of issues involved, the relationships among those issues, the frequency of changes during execution that threaten to disrupt the supply-chain, the wide range of disparate risk factors, and the diversity of the players involved[2]. Management of this compound complexity requires the assistance of the kind of intelligent software system environment described in this paper.

To summarize, there are essentially two principal requirements for such an environment. The first requirement is a rich contextual representation of supply-chain information. This can be provided by a virtual model of the real world context within which the information management tasks such as the preparation of demand forecasts, the effective management of inventory levels, the design and operation of effective distribution networks, the

[2] The players or stakeholders in the pharmaceutical supply-chain typically have very different objectives. For example, the manufacturer is interested in high efficiency at minimum cost, brand recognition, risk mitigation, patient safety, accurate demand forecasting, and customer sentiment; the distributor is concerned about narrow profit margins, inventory levels, and security; the shipper is concerned about conveyance and driver reliability and route conditions; while the customers expect to receive their orders on time and in an undamaged state.

formulation of multi-modal shipment plan, reacting to unforeseen events, preparing proactively for potential future events, and so on, can be expeditiously performed. The importance of this virtual model of real world context must not be underestimated. As a core requirement it provides the basis of most of the assistance capabilities of the intelligent information management environment described in this paper. Without access to the context provided by the multi-layered ontology the different groups of software agents defined in Section 3.3 and the Appendices could not function as intelligent tools in the manner described in Sections 3.6 and 3.7.

The second requirement is collaboration among the human users, as well as interaction between the human users and the intelligent software tools (e.g., agents) and, as discussed in Sections 3.6 and 3.7, between the intelligent tools themselves. Effective collaboration between any two parties assumes at least some commonality of purpose. Between human parties this commonality is based not only on the understanding that each party has of its own objectives, but also on some level of understanding of the objectives and needs of the other party. In addition, there is a distinctly opportunistic aspect to collaboration. While the general requirement for collaboration and even the protocol that must be adhered to during the process of collaboration may be prescribed, the events that will initiate collaboration are largely unpredictable.

Similar principles of collaboration apply to the interactions between the human users and the software agents, and among the software agents themselves. The human users will expect the agents that they interact with to have some semblance of common understanding of the context[3] of the interaction. This applies equally whether the user is requesting an explanation of an agent-generated result or queries the agent for specific information. Similarly, agents need some understanding of context to determine under what circumstances they should send an alert to human users or other agents. Clearly, the prerequisite for this semblance of understanding is the existence of a virtual model of real world context at the software level.

The current state of technology in software development provides the means for implementing a distributed, collaborative, intelligent, information management environment. Service-oriented architecture (SOA) concepts provide the framework and the guiding principles for developing distributed, service-based systems. The field of ontology representation is sufficiently mature to support the expressive modeling of domain knowledge as the enabling foundation for intelligent software tools or agents. As shown in the scenarios described in Sections 3.6 and 3.7 such agents can continuously monitor the supply-chain, participate in decision-making processes within specific domains, gather and present relevant information to the human user, and opportunistically communicate with human users and other agents.

[3] The *context* within which the interaction occurs is provided by the ontology representation of the real world operational background that pertains to the use-case of the particular supply-chain transaction. A higher level of intelligence would be required for the agent to have some level of understanding of the *purpose* of the interaction. While some indicators of purpose may be implicitly embedded in the context model provided by the ontology, inferencing is required to extend understanding from context to purpose. For relatively simple situations such inferencing can be performed by software agents. In more complex situations the participation of the human user will be required.

Appendix A: Typical Service Agents

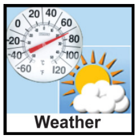

1. The ***Weather Agent*** has the ability to interpret and translate unstructured data from public (Internet) weather sites into a weather report that has meaning to both the human user and the computer (i.e., is machine processable).

2. The ***Traffic Agent*** has the ability to interpret and translate unstructured data from public (Internet) vehicular traffic report sites into traffic congestion warnings and alerts that have meaning to the human user and the computer (i.e., are machine processable).

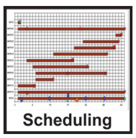

3. The ***Scheduling Agent*** is capable of integrating inter-modal movements, taking into consideration the delivery dates of shipments at distribution centers and final destinations, as well as the availability of surface and air transportation, and delivery windows.

4. The ***Quality Agent*** is capable of validating the authenticity (pedigree) of a drug based on serial/batch identification, as well as confirming in conjunction with the *Physical Conditions Agent* that any necessary environmental conditions (e.g., temperature) have been maintained.

5. The ***Inventory Agent*** is responsible for monitoring the inventory of distribution centers and therefore has the ability to access data sources and formulate queries on an on-going basis, as well as in response to requests for inventory information from other agents and human users.

6. The ***Transport Agent*** is capable of selecting the most suitable mode(s) of transportation from origin to destination, including the services of third party logistic services providers (3PL).

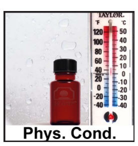

7. The ***Physical Conditions Agent*** is capable of monitoring any necessary environmental conditions such as temperature that are required to maintain the integrity of a specialty drug.

8. The ***Location Agent*** has the ability to determine the location of supplies that are currently in-transit or have been recently delivered in significant quantity to a customer. These supplies could be potentially available to the manufacturer in an emergency by redirection or through negotiation with the customer.

9. The ***Mash-Up Agent*** is capable of generating a web application that combines data and/or existing Internet functionality (e.g., Google Earth) from multiple sources into an action report, such as an on-the-spot view of a local event (e.g., disaster area survey or loading dock activities).

10. The ***Expiration Agent*** monitors the aging date of drugs and will issue warnings and alerts based on user specifications.

11. The ***Obstruction Agent*** is capable of monitoring road and airport closures due to flooding, weather and other conditions, and issue warnings and alerts if relevant to a particular distribution shipment.

12. The ***Medical Agent*** is responsible for identifying substitutes for products in case of shortages.

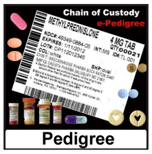

13. The ***Pedigree Agent*** is responsible for maintaining an audit trail of a drug from manufacture through distribution to the customer.

Appendix B: Typical Planning Agents

1. The ***Routing Agent*** has the ability to plan and re-plan multi-modal routing alternatives under time critical conditions, taking into consideration route conditions, efficiency, cost, and risk.

2. The ***Cost Agent*** has the ability to rapidly estimate the cost of alternative courses of action such as distribution plans, individual shipments, and product stocking strategies.

3. The ***Risk Agent*** has the ability to assess the risks associated with a given shipment, distribution plan, recall process, or pilot program based on past experience and specified risk factors.

4. The ***Efficiency Agent*** is responsible for monitoring the compliance of distribution plans, recall processes, pilot programs, and other performance-sensitive undertakings with established schedules.

5. The ***Opportunity Agent*** will monitor the current state of the supply-chain with the intent of identifying potential process improvements and/or distribution efficiencies that may result in resource savings (e.g., time, cost, material).

6. The ***Closure Agent*** is responsible for determining when the execution of a plan consisting of a sequence of measurable actions, tasks, or shipments can be considered to have been completed.

7. The ***Sunset Agent*** will analyze historical product sales, monitor customer sentiment on social media site(s), and assess market competitiveness factors to issue product phase-out planning recommendations.

8. The ***Demand Agent*** is capable of preparing demand forecasts for drug products or raw materials based on historical data, current requirements, and identified trends.

9. The ***Product Agent*** has the ability to recommend inventory reserves and stocking levels at distribution centers based on forecast estimates produced by the *Demand Agent*.

Appendix C: Typical Coordination Agents


Conflict

1. The ***Conflict Agent*** is capable of detecting conflict conditions that may arise among agents and within the supply-chain in general and the distribution network in particular, and identifying the likely causes.


Collaboration

2. The ***Collaboration Agent*** is responsible for facilitating collaboration by activating agents and alerting the human users to the need for interaction. In its coordinating role it also keeps an historical record of the individual actions that pertain to its assigned tasks.


Disaster

3. The ***Disaster Agent*** is responsible for coordinating the mitigation and exploitation of natural (e.g., hurricane) and man-made (e.g., terrorist) disaster situations by adapting predefined contingency plans to actual conditions.

Recall

4. The ***Recall Agent*** is responsible for coordinating the preparation of recall plans based on predefined templates, adapted to specific incident conditions, and monitoring the execution of the adapted recall plan.


Distribution

5. The ***Distribution Agent*** is capable of matching transportation needs, based on product lists, special environmental requirements, shipment schedules, and associated resource constraints, to produce an economical distribution plan.

Customer

6. The ***Customer Agent*** is capable of integrating customer sentiment information extracted by specialized *Service Agents* from social media sites and injecting these considerations into on-going planning and execution processes.


Pilot

7. The ***Pilot Agent*** is responsible for monitoring the execution of pilot program plans in respect to schedule compliance, actual and potential disruptions, and changes in conditions during execution that were not considered during the planning stage.

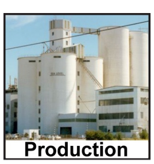
Production

8. The ***Production Agent*** is capable of estimating production needs based on current inventory levels, demand forecasts, and established strategic planning factors that can be quantified.

Appendix D: Typical Control Agents

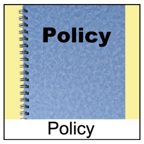

1. The ***Policy Agent*** has the ability to abstract the principal features of a plan to a conceptual level to identify potential or actual non-compliance with corporate policies.

2. The ***Performance Agent*** has the ability to apply metrics for evaluating not only the efficiency of an individual plan but also assess its impact on the overall operational efficiency of the supply-chain.

3. The ***Priority Agent*** is responsible for monitoring the assigned priority of orders and/or shipments and drawing high priority transactions to the attention of the *Collaboration Agent*, as well as alerting other agents and/or human users if high priority transactions are subject to delay.

4. The ***Security Agent*** exercises the dual role of monitoring internal security inspection records (issues warnings and alerts on poor inspection results and omitted inspections) and assessing the theft potential of a shipment based on shipment value, route and transportation mode risk, product type, and relevant historical theft records.

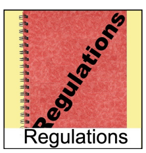

5. The ***Regulations Agent*** is capable of monitoring a global Regulations Repository with the objective of alerting other agents and/or human users to changes in government regulations, providing semantic search capabilities to human users, and assessing conflicts with regulations that may be identified by the *Conflict Agent*.

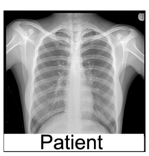

6. The ***Patient Agent*** is responsible for patient safety in respect to Recall Program notifications, sample tracking and updating, monitoring of Internet sites such as the USP-ISMP Medication Errors Reporting Program (MERC) website, and maintaining records of medication errors.