

Enterprise Knowledge Portals: What They Are and What They Do

Joseph M. Firestone, Ph.D.

Problems and Solutions in the Enterprise

IT applications are purposeful. They support and partially automate human participation in business processes. Portals are no different. So let's talk business processes and knowledge processes before we talk knowledge portals.

Every individual, team, or group within the enterprise encounters problems in the course of the workday. Every problem has alternative solutions. And every alternative solution is subject to criticism and to replacement if it performs less well than its competitors. The best problem solution is the competitive alternative that best survives criticism.

The Enterprise is a continuous, dynamic "swirl" of problem-solving knowledge-related Interactions from which knowledge is produced and integrated with the business processes of the enterprise. For a given problem generated by a business process, it is useful to abstract from the swirl and to conceptualize an iteration of a knowledge life cycle encompassing knowledge processes targeted on solving that problem. But before getting to such a knowledge life cycle model, we need to begin with the general concept of a business process of which both knowledge and knowledge management processes are sub-types

Business Processes

Figure One specifies the idea that any business process (including knowledge and knowledge management processes) may be viewed as a network of linked activities governed by validated rule sets, or knowledge. A linked sequence of activities performed by one or more agents sharing at least one objective is a Task. A linked, but not necessarily sequential set of tasks governed by validated rule sets, producing results of measurable value to the agent or agents performing the tasks, is a Task Pattern. A cluster of task patterns, not necessarily performed sequentially, often performed iteratively and incrementally, is a Task Cluster. Finally, a hierarchical network of interrelated, purposive, activities of intelligent agents that transforms inputs into valued outcomes, a cluster of task clusters, is a business process.




Figure One -- The Activity to Business Process Hierarchy

Any business process, task cluster, task pattern, or task must involve **decision cycles**, themselves composed of tasks and task patterns, through which agents execute their part in a business process or component. These decision cycles are focused on domain-centered tasks and task patterns in the enterprise, because such tasks and task patterns must be executed by agents (individuals, teams, and groups) in order to do work. It is through these domain-centered tasks that decision cycle tasks and task patterns affect the outcomes of task clusters and business processes.

The generic task patterns or **phases** of any decision/execution cycle are: Planning, Acting (including deciding), Monitoring, and Evaluating.

- **Planning** is a knowledge production and knowledge integration task pattern. It includes setting goals, objectives, and priorities, making forecasts as part of prospective analysis, performing cost/benefit assessments as part of prospective analysis, and revising or reengineering a business process. It involves capturing and using data, information, and knowledge to produce a plan, an instance of planning knowledge.
- **Acting** means performing the specific domain business process (cluster, pattern, or task) or any of its components. Acting involves **using the**

planning knowledge, along with other knowledge to make and implement decisions.

- **Monitoring** means retrospectively tracking and describing the business process (cluster, pattern, or task) and its outcome. Monitoring involves gathering data and information, modeling processes, and using previous knowledge to produce new descriptive, impact-related, and predictive knowledge about the results of acting. Monitoring is another knowledge production and knowledge integration task pattern.
- **Evaluating** means retrospectively assessing the performance of the business process as a value network [1]. Evaluating means using the results of monitoring, along with previous knowledge to assess the results of acting and to produce knowledge about the descriptive gaps between business outcomes and tactical objectives and about the normative (benefits and costs) impact of business outcomes. Evaluating is yet another decision cycle task pattern that produces and integrates knowledge in the business process.

There is a natural order to the four phases of any decision/execution cycle in a value network. Figure two illustrates the order of these phases, or task patterns of the decision cycle.




Figure Two -- Decision Cycle Phases and Their Interactions.

Three of these four phases require knowledge production and knowledge integration to solve problems that occur in each phase, and the fourth, the acting phase, uses the knowledge produced in the other three phases. So every decision cycle in every business process requires both knowledge processing (production and integration) and knowledge use. Knowledge use is not a separate task but rather is part of deciding and acting. But planning, monitoring, and evaluating are knowledge production task patterns of different types, each involving sequential patterns of knowledge production and knowledge integration.

A Knowledge Life Cycle Model

So decision execution cycles performed by agents in executing tasks and task patterns in business processes are in large part sequentially ordered knowledge production and knowledge integration processes. Figure three provides an overview of a Knowledge Life Cycle model begun in collaboration with Mark McElroy, Edward Swannstrom, Douglas Weidner, and Steve Cavaleri [2], during meetings sponsored by the Knowledge Management Consortium International (KMCI), and further developed recently by Mark McElroy and myself [3]. Knowledge Production and Knowledge Integration are core knowledge processes in the model. Knowledge Production is stimulated by problems produced by business processes. It produces Validated Knowledge Claims (VKCs), Unvalidated Knowledge Claims (UKCs), and Invalidated Knowledge Claims (IKCs), and information about the status of these. Organizational Knowledge (OK) is composed of all of the foregoing results of knowledge production. It is what is integrated into the enterprise by the Knowledge Integration process.




Figure Three -- The Knowledge Life Cycle Model (Overview)

The knowledge integration process, in turn, produces the Distributed Organizational Knowledge Base (DOKB) and the DOKB, in turn, has a major impact on structures incorporating organizational knowledge such as normative business processes, plans, organizational culture, organizational strategy, policies, procedures, as well as information systems. Coupled with external sources these structures then feed back first to behavioral business processes through the acting phase of decision cycles which then generate new problems in the planning, monitoring, and evaluating phases of decision cycles to be solved. Then, second, the feedback generates knowledge production to solve these new problems - which is why it's called the Knowledge Life Cycle (KLC) model.

Drilling down into knowledge production (figure four), the KLC view is that information acquisition, and individual and group learning, in the service of problem-solving, impact on knowledge claim formulation, which, in turn, produces Codified Knowledge Claims (CKCs). These, in their turn, are tested in the knowledge validation sub-process, a sub-process of criticism including, but not limited to, empirical testing, which then produces organizational knowledge.

Knowledge validation is the key sub-process that distinguishes knowledge production from information production. It is the sub-process of criticism of competing knowledge claims, and of comparative testing and assessment of them, that transforms knowledge claims from mere information into tested information, some of which passes organizational tests and therefore becomes,

from the organizational point of view, knowledge. Individual and group learning may involve knowledge production from the perspective of the individual or group, but from the perspective of the enterprise, what the individuals and groups learn is information, not knowledge until it has been validated. Similarly information acquired may be knowledge from the perspective of the external parties it is acquired from, but not knowledge to the enterprise acquiring it, until it has been validated as such.




Figure Four -- The Components of Knowledge Production

Figure Four also illustrates that knowledge validation has a feedback effect on individual and group learning. This occurs because individuals and groups participating in knowledge claim validation are affected by their participation in this process. They both produce organizational knowledge in the form of validated knowledge claims and experience change in their own justified beliefs as an outcome of that participation.

Drilling down into knowledge integration (figure five), organizational knowledge is integrated across the enterprise by the broadcasting, searching/retrieving, teaching, and sharing sub-processes. These generally work in parallel rather than sequentially. Not all are necessary to a specific instance of the KLC. All may be based in personal non-electronic or electronic interactions. Here is a glossary of the major terms used in the KLC Model.




Figure Five -- The Components of Knowledge Integration

Sidebar One: KLC Glossary

Codified Knowledge Claims - Information that has been codified, but which has not yet been subjected to organizational validation.

Distributed Organizational Knowledge Base - an abstract construct representing the outcome of knowledge integration. The DOKB is found everywhere in the enterprise, not merely in electronic repositories. It is distributed over all of the agents and all of the repositories in the enterprise.

Experiential Feedback Loops - Processes by which information concerning the outcomes of organizational learning activities are fed back into the Knowledge Production phase of an organization's knowledge life cycle as a useful reference for future action.

Individual and Group Learning - A process involving human interaction, knowledge claim formulation, and validation by which new individual and/or group knowledge is created.

Information About Invalidated Knowledge Claims - Information that attests to the existence of invalidated knowledge claims and the circumstances under which such knowledge was invalidated.

Information About Unvalidated Knowledge Claims - Information that attest to the existence of unvalidated knowledge claims, and the circumstances under which such knowledge was tested and neither validated nor invalidated.

Information About Validated Knowledge Claims - Information that attests to the existence of validated knowledge claims and the circumstances under which such knowledge was validated.

Information Acquisition - A process by which an organization either deliberately or serendipitously acquires knowledge claims or information produced by others external to the organization.

Invalidated Knowledge - A collection of codified invalidated knowledge claims.

Invalidated Knowledge Claims - Codified knowledge claims that have not satisfied an organization's validation criteria. Falsehoods.

Knowledge Claim - A codified expression of potential knowledge which may be held as validated knowledge at an individual and/or group level, but which has not yet been subjected to a validation process at an organizational level. Information. Knowledge claims are components of hierarchical networks of rules, that if validated would become the basis for organizational or agent behavior.

Knowledge Claim Formulation - A process involving human interaction by which new organizational knowledge claims are formulated.

Knowledge Integration - The process by which an organization introduces new knowledge claims to its operating environment and retires old ones. Knowledge Integration includes all knowledge transmission, teaching, knowledge sharing, and other social activity that communicates either an understanding of previously produced organizational knowledge to knowledge workers, or the knowledge that certain sets of knowledge claims have been tested, and that they and information about their validity strength is available in the organizational knowledge base, or some degree of understanding between these alternatives. Knowledge integration processes, therefore, may also include the transmission and integration of information.

Knowledge Production - A process by which new organizational knowledge is created. Synonymous with "organizational learning."

Knowledge Validation Process - A process by which knowledge claims are subjected to organizational criteria to determine their value and veracity.

Organizational Knowledge - A complex network of knowledge and knowledge sets held by an organization, consisting of declarative and procedural rules (validated knowledge claims).

Organizational Learning - A process involving human interaction, knowledge claim formulation, and validation by which new organizational knowledge is created.

(Business) Structures Incorporating Organizational Knowledge - Outcomes of organizational system interaction. The organization behaves through these structures including business processes, strategic plans, authority structures, information systems, policies and procedures, etc. Knowledge structures exist within these business structures and are the particular configurations of knowledge found in them.

Unvalidated Knowledge Claims - Codified knowledge claims that have not satisfied an organization's validation criteria, but

which were not invalidated either. Knowledge claims requiring further study.

Validated Knowledge Claims - Codified knowledge claims that have satisfied an organization's validation criteria. Truth from the viewpoint of the organization.

Knowledge production and knowledge integration, their sub-processes, task clusters, etc., like other value networks, are partly composed of decision cycles through which agents execute their roles in these value networks. This means that planning, acting, monitoring and evaluating also apply to knowledge processes and to activity in the KLC. These processes are executed by agents performing decision cycles. It is the role of Knowledge Management to support these agents.

The KLC, Knowledge Processes and Knowledge Management

What is the relationship between managing the KLC or its knowledge processes and Knowledge Management itself?

To answer this question we need to decide whether managing knowledge refers to managing knowledge processes, managing the outcomes of these processes, or managing both? It has recently been stated [4, P. 87] that *"It's not knowledge management, stupid, it's knowledge PROCESS management."* But this is surely too simple. While KM is a process that manages the knowledge processes of the KLC, since those processes produce knowledge outcomes including the knowledge base, it is also true that KM indirectly manages knowledge outcomes. Or, to put the situation another way, knowledge management is most directly knowledge process management, and only indirectly knowledge base management. The knowledge processes in question are given in the KLC. So knowledge management is both process and outcome management.

The Nature of Knowledge Management

There are many available definitions of knowledge management [5], but few specifications that bring the definitions a step closer to analysis and measurement. I define KM as human activity that is part of the Knowledge Management Process (KMP) of an agent or collective. This reduces KM to the definition of KMP. And the KMP, in turn, is an ongoing, persistent, purposeful interaction among human-based agents through which the participating agents aim at managing (handling, directing, governing, controlling, coordinating, planning, organizing) other agents, components, and activities participating in the basic knowledge processes (knowledge production and knowledge integration) into a planned, directed, unified whole, producing, maintaining, enhancing, acquiring, and transmitting the enterprise's knowledge base. This definition is

another way of stating the idea that KM is management of the KLC and its outcomes, but the idea of KM still needs further specification.

Note first that the KMP is a business process. I break down the KMP [6] into three task clusters: interpersonal behavior, knowledge processing behavior, and decision-making behavior.

Interpersonal behavior may be further categorized into:

- Figurehead or ceremonial KM activity (focuses on performing formal KM acts such as signing contracts, attending public functions on behalf of the enterprise's KM process, and representing the KM process to dignitaries visiting the enterprise);
- Leadership (includes hiring, training, motivating, monitoring, and evaluating staff. It also includes persuading non-KM agents within the enterprise of the validity of KM process activities); and
- Building external relationships -- another political activity designed to build status and to cultivate external sources of support for KM.

KM Knowledge processing behavior includes:

- KM knowledge production (different in that it is here that the rules for knowledge production that are used at the level of knowledge processes are specified);
- KM Knowledge Integration (affected by KM knowledge production, and also affects knowledge production activities by stimulating new ones).

Decision making behavior includes:

- Changing knowledge process rules (involves making the decision to change such rules and causing both the new rules and the mandate to use them to be implemented);
- Crisis Handling (e.g., meeting CEO requests for new competitive intelligence in an area of high strategic interest for an enterprise, and directing rapid development of a KM support infrastructure in response to requests from high level executives);
- Allocating Resources (KM support infrastructures, training, professional conferences, salaries for KM staff, funds for new KM programs, etc.);
- Negotiating agreements (with representatives of business processes over levels of effort for KM, the shape of KM programs, the ROI expected of KM activities, etc.).

The nine categories within the task clusters are the task patterns illustrated in figure one. Further specification of KM involves breaking down the task patterns further. In brief, the nature of knowledge management is that it is a complex process composed of the above task clusters and task patterns, executed by agents through planning, acting, monitoring, and evaluating activities.

Enter the Enterprise Knowledge Portal

About eighteen months ago (March 20, 1999) I introduced the *concept* "Enterprise Knowledge Portal" (in "Enterprise Information Portals and Enterprise Knowledge Portals," [7]) to the IT world by defining it and distinguishing it from the Enterprise Information Portal (EIP). The *term* had been used earlier (March 5, 1999) by Hummingbird Communications, in a press release [8] announcing their impending acquisition of the PC DOCS group, and in an article in DataPro Industry News by Karen Shegda and Allan Tiedrich [9] entitled "Knowledge Management = Access + Collaboration + Retrieval + Analysis. But they declined to define or characterize the EKP except to say that it would integrate structured and unstructured content, and provide a single point of access to all relevant enterprise information. On March 25, IDC published its "Sourcebook for Knowledge Superconductivity" [10]. In this report Gerry Murray, then IDC's Director of Knowledge Management Technologies (currently, IDC's Vice President of Business Strategy and Development) distinguishes four types of corporate portals, including EKPs and offers some definitions (which I have discussed previously in ("Defining Enterprise Information Portals," [11]).

The introduction of the EKP concept was followed by remarkably little activity on EKPs. Google searches on "Enterprise Information Portal" (1015 hits), "Enterprise Information Portals" (943 hits), "Enterprise Knowledge Portal" (56 hits) and "Enterprise Knowledge Portals," (66 hits) tell the story of EKP vs. EIP during 1999 and the first two months of 2000. Other than urls reflecting writings by Gerry Murray or myself, and third party links to what we've done, there are few evidences of EKP activity on the web, and only six vendors that bill themselves as having EKP products. Hummingbird promised an EKP in March 1999, but by this March had released an EIP instead. During its first year, the EKP space gave every evidence of being swallowed up by its EIP parent.

In spite of the initial failure of the EKP to grab mind share in the portal space, the future is bright for it because a number of the claims of primary benefits of EIPs, such as ROI, competitive advantage, increased effectiveness, and accelerated innovation, assume that information delivered by EIPs is correct information. But the risk associated with EIPs is that if that's not the case, these four benefits are lost. An overriding justification for implementing an EKP, rather than an EIP, is to *secure* these four benefits and to minimize decision making risk by increasing the *quality/validity* of information supplied by the portal. To see this clearly, we need to examine the idea of the EKP.

What is the EKP?

To understand the concept fully, it is best to begin with definitions of business and enterprise information portals. According to Wayne Eckerson [12] a Business Portal is an application that “provides business users one-stop shopping for any information object they need inside or outside the corporation.” It also provides shared services such as “security, metadata repository, personalization, search, publish/subscribe,” etc., as well as a common look and feel to the portal gateway.

According to Merrill Lynch's Shilakes and Tylman [13]: "Enterprise Information Portals are applications that enable companies to unlock internally and externally stored information, and *provide users a single gateway to personalized information* needed to make informed business decisions. " ". . . an amalgamation of software applications that consolidate, manage, analyze and distribute information across and outside of an enterprise (including Business Intelligence, Content Management, Data Warehouse & Mart and Data Management applications)."

EIPs:

- Use “push” and “pull” technologies to transmit information through a standardized web interface;
- Provide “interactivity” -- the ability to “ ‘question’ and share information on” user desktops;
- Exhibit a trend toward verticalization in applications including packaged applications with targeted content toward industries or corporate functions;
- Integrate disparate applications and data/content stores into a single system;
- Access both external and internal sources;
- Support bi-directional information exchange from sources;
- Use data and information acquired for further processing

An EKP is an enhanced Enterprise Information Portal (EIP). It is an EIP that:

- is goal directed toward knowledge production, knowledge integration, and knowledge management, and also
 - Focuses upon, provides, produces and manages information about the validity of the information it supplies,
 - Provides information about your business and meta-information about the degree to which you can rely on that information,
 - Distinguishes knowledge from mere information,
 - Provides a facility for producing knowledge from information

- orients one toward producing and integrating knowledge rather than information

The next sections specify these in more detail

AN EKP is an EIP

An EKP shares the characteristics of other EIPs. It is a particularly comprehensive version of an EIP however, incorporating a personalized browser-based interface, structured data management, unstructured content management, and collaborative, as well as knowledge production, knowledge integration, and knowledge management functionality. It also requires an integrative architecture incorporating Knowledge Claim Objects (KCOs) encapsulating knowledge claim data, metadata describing the validity characteristics of these knowledge claim objects, and methods producing behavior of the objects. EKP architecture will be described in greater detail below.

Knowledge Processing Means Knowledge Production and Knowledge Integration

This should be clear from the previous knowledge processing/knowledge management framework. See also, "The Metaprise, The AKMS, and The Enterprise Knowledge Portal," [14], and Mark McElroy's "The Second Generation of KM," [3], for other statements on aspects of the framework. The point is that to be an EKP, an EIP application must implement use cases that support knowledge production and knowledge integration. That means the use cases must support information acquisition, individual and group learning, knowledge claim formulation, knowledge claim validation, broadcasting, searching/retrieving, sharing, and teaching. In each of these areas it must also support planning, monitoring, and evaluating phases of decision cycles. This again underlines how comprehensive EKPs are compared to EIPs.

The Knowledge Management Process

One distinguishing characteristic of the EKP is that it must support the Knowledge Management Process as defined earlier. This means that the EKP must implement use cases that support the nine task patterns, tasks, and activities of the KMP and also the phases of decision cycles encompassed by the KMP. There is no similar requirement for EIPs.

Provides, Produces, and Manages Information about the Validity of Information It Supplies

KNOWLEDGE AND INNOVATION: JOURNAL OF THE KMCI

Knowledge is validated information. EKPs distinguish knowledge from mere information by providing information about the results of tests of the validity of any piece of information. That means that EKPs *must track and store such meta-information while EIPs in general need not*.

This requirement is one that distinguishes EKPs from other EIPs. *It is a requirement that greatly expands the diversity and volume of metadata found in EKPs as compared with EIPs, since the validity meta-information must record the full history of the discussions and interactions that transform information into knowledge.*

Provides Business Information Along with Meta-information About the Degree To Which You Can Rely On It

Validity information about a knowledge claim is meta-information about that claim. This validity base includes meta-information comparing the knowledge claim against competing knowledge claims. This meta-information tells you the degree to which you can rely on the target knowledge claim compared to its competitors. It tells you the relative strength of the knowledge claim compared to its competitors. Thus, *EKPs record the history of the competitive struggle among ideas (knowledge claims) put forward to solve problems within the enterprise. EIPs need record no such history.*

Distinguishes Knowledge From Information

By providing validity information (meta-information) about knowledge claims, EKPs provide information on the relative strength of knowledge claims. The stronger the claim, the closer it approaches organizational knowledge and the stronger the support it provides for decisions. The weaker the claim, the more closely it approaches false organizational information and the weaker the support it provides for decisions.

Provides a Facility for Producing Knowledge From Information

By providing services for knowledge claim formulation and validation and tracking and storing the results of validation activities in knowledge claim objects, EKPs provide a facility that supports producing knowledge from knowledge claims, or, put another way, supports producing better or worse validated knowledge claims from unvalidated ones). Since knowledge claims are information and knowledge is validated knowledge claims, it follows that EKPs provide a facility for producing knowledge from information.

Orients One toward Producing and Integrating Knowledge Rather than Only Information

Because the EKP supports the full set of knowledge life cycle activities including, and most critically, individual and group learning and knowledge validation, it orients one toward knowledge production. Because new knowledge results from use of the EKP, enterprise information integration processes will be oriented toward integrating knowledge and validity information as well as business information.

EKP Architecture and Components

The EKP provides a:

- knowledge worker-centric, knowledge work flow-oriented, single point of access to enterprise data and content stores, and applications supporting knowledge production, knowledge integration, and knowledge management;
- personalized desktop browser-based portal that with the assistance of an integrative, logically centralized, but physically distributed *Artificial Knowledge Manager (AKM)*, composed of distributed *Artificial Knowledge Servers (AKSs)* and *intelligent mobile agents*, is connected to all enterprise mission critical application sources and data and content stores;
- secure, seamless, single-logon capability for all network, application, and service resources.

The distributed AKM balances processing loads across the enterprise and provides for dynamic integration of the portal system in the face of change. The EKP system is managed, in part, by the AKM's ubiquitous intelligent agents serving all application servers, data and content stores, and clients in the enterprise. The EKP system, provides a new work environment for enterprise knowledge workers, one that is aligned with, and supports and partially automates, their individual and collaborative work flow in creating, distributing and using data, information, and knowledge, and in making and implementing decisions and actions.

More specifically, the EKP in operation provides: a wide range of functionality (including structured data management, unstructured content management, collaborative processing, information processing, information management, knowledge processing, and knowledge management); a wide range of data and content stores as sources of previously developed information and knowledge; and an integrative object/component-based portal architecture. Here is an introduction to knowledge portal architecture and components.

Figure six presents an overview. Note the complexity of the EKP system with respect to diversity of data and information stores, and application servers, the presence of structured data and unstructured content sources, the publication and agent capabilities, the web server and portal capabilities, the text and data mining capabilities, the collaborative capabilities, and the dynamic integrative

capabilities provided by the AKM with its agents and logically centralized but physically distributed AKS servers and intelligent agents.

Note also, that this view of the EKP does not emphasize its front-end aspects or format. In the EKP the action is in the middleware and in how it functions to support knowledge production, knowledge integration and knowledge management.




Figure Six -- The Architecture of the Enterprise Knowledge Portal

Legend for Figure Six

ETL = Extraction, Transformation, and Loading
DDS = Dynamic Data Staging Area
DW = Data Warehouse
ODS = Operational Data Store
ERP = Enterprise Resource Planning
Query = Query and Reporting Server
OODBMS = Object-Oriented Database Management System
OLTP = Online Transaction Processing
MOLAP = Multi-dimensional Online Analytical Processing
ROLAP = Relational Online Analytical Processing

Though figure six illustrates the diversity of components that may enter an EKP, it does not specify its necessary generic components. These are:

- Browser and e-mail clients
- The Avatar -- a client-based intelligent agent
- The portal application server(s),
- The access management system
- Knowledge Claim Objects
- The enterprise Artificial Knowledge Server(s) (AKSs),
- Complex adaptive system (cas) intelligent agent platform
- The formal knowledge production application server(s) and its associated clients supporting analytical and statistical modeling, KDD and Data Mining, Text Mining, Simulation, impact analysis and forecasting,
- The collaborative processing application server, and
- A persistent storage component.

Among these components, the avatar, knowledge claim objects, the AKS, the cas agent platform, and the collaborative processing application server are all unusual and deserve further analysis in specifying EKP architecture.

The Avatar

The purpose of the Avatar, a client-based cas agent is to create a highly automated, enterprise knowledge-enabled, self-learning/adaptive portal interface that is personally and dynamically tailored to each user. The types of automation involved here include:

- assimilation of the user's local environment, personal preferences and cognitive patterns;
- learning from, interacting with, and utilizing enterprise knowledge communicated by the Artificial Knowledge Servers;
- providing the user autonomous negotiating capabilities with the rest of the EKP system;
- supporting performing personal and collaborative workflows, by
 - (1) providing workflow memory, accessible from the portal interface adapting in accordance with anticipated needs of the user, and
 - (2) by providing access to user cognitive maps; and
- producing knowledge claims and submitting these to the EKP system.

The knowledge claims of the Avatar represent local knowledge, in contrast to the rules that have been validated by the EKP system's network of Artificial Knowledge Servers.

Knowledge Claim Objects

An important class of objects in the EKP system is the knowledge claim object (KCO). A KCO is distinguished from an ordinary business object by the presence of validity metadata encapsulated in the object. Such metadata compares the KCO to alternative, competing KCO's, and may be expressed in many different forms.

The "metadata" may be qualitative or quantitative, or it may be in the form of textual content. In relatively infrequent but important special cases, the metadata may involve quantitative ratings of a knowledge claim compared to its competitors. When the KCO is accessed by a user, data, metadata, and methods are all available, so the user can evaluate the KCO as a basis for decision against competing KCOs. This capability is not available in EIPs, which express knowledge claims as data or business objects only.

The Artificial Knowledge Server

The distributed AKS provides Process Control Services, an Object Model of the EKP system, and connectivity to all enterprise information, data stores, and applications.

- Process Control Services include:
 - In-memory proactive object state management and synchronization across distributed objects
 - Component management and Workflow Management
 - Transactional multithreading
 - business rule management and processing,

- KCO management and processing and
- metadata management
- The In-memory Active Object Model/Persistent Object Store is characterized by:
 - Event-driven behavior;
 - An EKP-wide model with shared representation;
 - Declarative and procedural business rules;
 - Caching along with partial instantiation of objects;
 - A Persistent Object Store for the AKS; and
 - Reflexive Objects and KCOs
- Connectivity Services required by the EKP are:
 - Language APIs: C, C++, Java, HTML, XML, CORBA, DCOM
 - Databases: Relational, ODBC, OODBMS, hierarchical, network, flat file, XML, etc.
 - Wrapper connectivity for application software: custom, CORBA, or COM-based
 - Applications connectivity including all the applications in the enterprise whether these are mainframe, server, or desktop - based

CAS Agents

The second type of component comprising the EKP's AKM is the intelligent agent (IA). EKP IAs are lightweight, intelligent, efficient, specialized Business Process Engines (BPEs) that provide some memory and a small amount of processing power at almost no cost. They also provide bi-directional communication, an inference engine, an ability to model semantic networks, and an ability to learn through reinforcement of semantic network connections and creation of new nodes. It is this last ability to learn that makes them intelligent.

Intelligent agents alone cannot yet create the virtual enterprise. For complex processing and an enterprise wide view, the AKS is also indispensable. But IAs provide distributed load balancing to processing in the AKM. They are necessary partners in providing the processing power needed for implementing the EKP. When we add agents to the AKS to create the AKM, we provide software wiring for the enterprise that connects its central brain components (the AKSs) to its sensors (the agents). The result is a flexible and scalable AKM that can integrate the various components of the EKP into a virtual enterprise.

Collaborative Processing Application Server

Eight categories of collaboration are: prioritization, planning, project management, distributing expertise, training, problem solving, knowledge production, and workflow. Each of these areas represents non-trivial functions that are currently realized in complex applications, that to some extent represent distinct functional sub-spaces, and that could each be wrapped into a collaborative or knowledge portal either separately or in combination with one of the other categories.

Products such as Instinctive's eRoom or Intraspect, that focus on collaboration in general or collaboration in support of project management, are very different from products that intend to support strategic planning implementations such as Engenia's Unity. And these are very different from products that provide for group collaboration on analytical modeling and/or data mining, or that provide for a team approach to prioritized decision making, such as Expert Choice. And these, in turn, are very different from products such as Ernst & Young's ERNIE, and Orbital Software's Organik Knowledgeware that allow knowledge workers to access the expertise of "Gurus" in specialized fields.

In short, the term collaborative application server covers a variety of applications. In the context of the EKP, collaborative processing requirements should encompass all categories specified above, because all are related to knowledge outcomes or knowledge processes. EKPs then, are particularly comprehensive in their collaborative processing functionality.

The Essence of The EKP

Every individual, team, or group within the enterprise encounters problems in the course of the workday. Every problem has alternative solutions. And every alternative solution is subject to criticism and to replacement if it performs less well than its competitors. The best problem solution is the competitive alternative that best survives criticism. And it is that alternative that is the enterprise's best knowledge at any given moment.

The set of problem-solving interactions in an enterprise constitutes a continuous, dynamic "swirl" from which knowledge is produced and integrated with the business processes of the enterprise. The essence of the EKP is its support for facilitating this knowledge "swirl" and its management. The EKP supports every phase of the KLC and every KM activity. It supports business processes with new knowledge production and integration. It supports collaboration focused on knowledge production by providing the history of knowledge-validating activities in the enterprise. It also supports knowledge production through automated arbitration between local knowledge claims and regional and global knowledge claims. But most of all by supporting every phase of the KLC, and by subjecting knowledge claims to competition more efficiently than ever before, it subjects

them to evolutionary forces, and thus supports the acceleration of innovation, the growth of knowledge, in the EKP-supported enterprise.

e-Business Knowledge Portals

Portal technology is currently in full migration to the field of e-business. EIP technology is equally applicable to trading communities, trans-enterprise and externally-facing enterprise applications of all types. Eventually when true EKP products and solutions are implemented, they too will be used in e-business. Especially in trading communities, and in communities of practice in medicine, pharmacology, architecture, engineering, science, more generally, and in other areas where the distinction between true and false information is central.

Conclusion

The Enterprise Knowledge Portal is an application on the verge of development. The technology it requires is in existence now. The cost of its development is low as software applications go, since its implementation is largely a matter of systems integration. On the other hand, the benefits associated with the EKP are great. They are nothing less than realization of the promise of the EIP to achieve increased ROI, competitive advantage, increased effectiveness, and acceleration of innovation.

As I indicated previously, EIPs are risky because (neglecting data quality applications which involve relatively superficial quality issues) they fail to evaluate the information they produce and deliver for quality and validity. Nothing, including EKPs, can ensure certainty about information, models, or knowledge claims. But EKP applications incorporate a systematic approach to testing and validation that produces quality assured information. In the category of portal technology they, not EIPs, are the best we can do. They, not EIPs, are the future of portal technology.

References

[1] Verna Allee, "Reconfiguring the Value Network," available at: http://www.vernaallee.com/reconfiguring_val_net.html.

[2] Edward Swanstrom, Joseph M. Firestone, Mark W. McElroy, Douglas T. Weidner, and Steve Cavaleri, "The Age of The Metaprise," Knowledge Management Consortium International, Gaithersburg, MD, 1999, available at <http://www.km.org/metaprise/MetapriseGrp.htm>.

[3] In e-mail and telephone communications.

- [4] Mark McElroy, "The Second Generation of KM," Knowledge Management (October, 1999), Pp. 86-88, also available at <http://www.macroinnovation.com/papers.htm>
- [5] See Yogesh Malhotra's compilation at <http://www.brint.com>
- [6] Joseph M. Firestone, "Enterprise Knowledge Management Modeling and Distributed Knowledge Management Systems," available at http://www.dkms.com/White_Papers.htm.
- [7] Joseph M. Firestone, "DKMS Brief No. Eight: Enterprise Information Portals and Enterprise Knowledge Portals," Executive Information Systems, Inc., 1999, available at http://www.dkms.com/White_Papers.htm.
- [8] Hummingbird Communications, Inc. Press Release on pending acquisition, March 5, 1999, available at: <http://www.hummingbird.com/press/1999/pcdocs.html>.
- [9] Karen Shegda and Allan Tiedrich, "Knowledge Management = Access + Collaboration + Retrieval + Analysis," DataPro Industry News, 1999, available at <http://gartner5.gartnerweb.com/public/static/datapro/industry/indnews28.html>.
- [10] "Enterprise Knowledge Portals to Become the Shared Desktop of the Future," IDC Press Release March 25, 1999 <http://www.idc.com/Data/Software/content/SW032699PR.htm>.
- [11] Joseph M. Firestone, "Defining Enterprise Information Portals," Executive Information Systems, Inc., 1999, available at http://www.dkms.com/White_Papers.htm.
- [12] Wayne Eckerson, "Business Portals: Drivers, Definitions, and Rules," The Data Warehousing Institute, Gaithersburg, MD, April 1999
- [13] Christopher C. Shilakes and Julie Tylman, "Enterprise Information Portals," Merrill Lynch, Inc., New York, NY, November 16, 1998
- [14] Joseph M. Firestone, "The Metaprise, The AKMS, and The Enterprise Knowledge Portal," available at http://www.dkms.com/White_Papers.htm.

Biography

Joseph M. Firestone, Ph.D. is Vice-President and Chief Scientist of Executive Information Systems (EIS), Inc. Joe has varied experience in consulting, management, information technology, decision support, and social systems

KNOWLEDGE AND INNOVATION: JOURNAL OF THE KMCI

analysis. Currently, he focuses on product, methodology, architecture, and solutions development in Enterprise Information and knowledge Portals, where he performs Knowledge and knowledge management audits, training, and facilitative systems planning, requirements capture, analysis, and design. Joe was the first to define and specify the Enterprise Knowledge Portal Concept. He is widely published in the areas of Decision Support (especially Enterprise Information and Knowledge Portals, Data Warehouses/Data Marts, and Data Mining), and Knowledge Management, and has recently completed a full-length industry report entitled "Approaching Enterprise Information Portals."

Joe is a founding member of the Knowledge Management Consortium International (KMCI), a member of its: Executive Committee, Board of Directors, Metaprise Project, and Governing Council of the KMCI Institute. He is also the Editor of the new journal "Knowledge and Innovation: Journal of the KMCI," and Chairperson of the KMCI's Artificial Knowledge Management Systems Special Interest Group. Joe is a frequent speaker at national conferences on KM and Portals. He is also developer of the web site www.dkms.com, one of the most widely visited web sites in the Portal and KM fields. DKMS.com has now reached a visitation rate of 95,000 visits annually.