Knowledge Management in the Service and Support Business

Kemal A. Delic
Hewlett-Packard
kemal_delic@hp.com

Umeshwar Dayal
Hewlett-Packard
umesh_dayal@hp.com

Abstract

As we are entering the New Millenium, we are witnessing a global shift toward a society of services. The evolution of the world's workforce structure indicates clearly the magnitude and importance of this mega-shift. The introduction of new, exciting technologies has led to a kind of New Economy that is based on the huge and rapid flows of data, information and knowledge. The Internet has had a most profound impact on business and society, enabling the quick spread of service industries and technologies. We observe that cost reduction and acceleration of business processes are the most obvious consequences of this singular phenomenon. In such an environment, the management of IT support services is becoming critical for business profitability.

1 Introduction

The management of support services includes the efficient deployment of people, processes and technologies so as to improve operational business parameters and financial indicators. Knowledge Management (KM) is a general umbrella term that encompasses several techniques and processes whose main common objective is to deal successfully with various inefficiencies in operational business processes and to create business value.

Intuitively, it is clear that knowledge plays a crucial role in human business activities, and that it has significant monetary value to an enterprise. However, measuring the impact of knowledge management has proved to be difficult, since there is no commonly accepted measurement methodology, and there are few reliable reports of the impact of knowledge management on business metrics [Del00].

We believe that knowledge management (KM) today [Dav00] plays a crucial role in the IT support and service environment. It directly impacts the productivity and training of support personnel, and this in turn justifies investments in knowledge management systems, programs and associated processes. To give an idea of the potential market size, industry estimates indicate that for each dollar spent on hardware and software products, seven to fourteen dollars are spent in associated training, support and services.

From the business point of view, KM impacts profit by reducing the costs of doing business and creates new streams of revenues through the introduction of new, innovative services. From a strategic point of view, it is seen as a potentially strong business growth generator (> 50%). Unfortunately the KM field is saturated with hype and buzzwords, so that real, documented KM success stories are rarely published. We will focus our attention on deployment of KM in IT support services as one of the most important practical domains, and we will describe our experience at Hewlett-Packard with the deployment of one KM system for IT support services.

IT support and services are typically organized as a multi-tiered operation, consisting of help desks or service desks, and operational service centers supporting the IT infrastructure and delivering various services. Customers obtain support and services by contacting the first line service desk. Problems that cannot be resolved in the first line are escalated to higher, but more expensive, levels of expertise. In addition to telephone or email-based help desks, enterprises are moving towards Web-based service portals, where customers can directly obtain access to IT support and services. In this paper, we will show in more detail how knowledge management can enhance the efficiency of traditional IT operations and enable the delivery of new types of services.

In the second section of this article, we introduce Knowledge Management. The third section deals specifically with KM in the Information Technology (IT) domain. In the fourth section, we outline the architecture of an IT KM system being deployed within Hewlett-

Proc. of the Third Int. Conf. on Practical Aspects of Knowledge Management (PAKM2000)
Basel, Switzerland, 30-31 Oct. 2000, (U. Reimer, ed.)
http://sunsite.informatik.rwth-aachen.de/Publications/CEUR-WS/Vol-34/

Packard Corporation. In the fifth section, we sketch the evolutionary architecture of IT support and service systems, concluding with a few insights about the future role of KM in the service industry.

2 Knowledge Management Defined

Knowledge is primarily embodied in human expertise and experience. It has first to be captured and expressed explicitly, then transformed and represented in various repositories, then disseminated and shared by knowledge workers who exploit it for making business decisions, and finally these actions may in turn lead to the creation of more knowledge. Knowledge Management can be defined as a set of generic processes aligned with the four principal transformation phases: gathering, organization, refining, and dissemination of knowledge (Figure 1). For the purpose of this article, we will consider Knowledge Management as a combination of several disciplines and techniques.

In Figure 1, we summarize a generic KM methodology and process that can be applied to the majority of known KM approaches in various domains such as consulting, customer service & support, training, human resources, finance, product creation, sales and marketing, and strategic management.

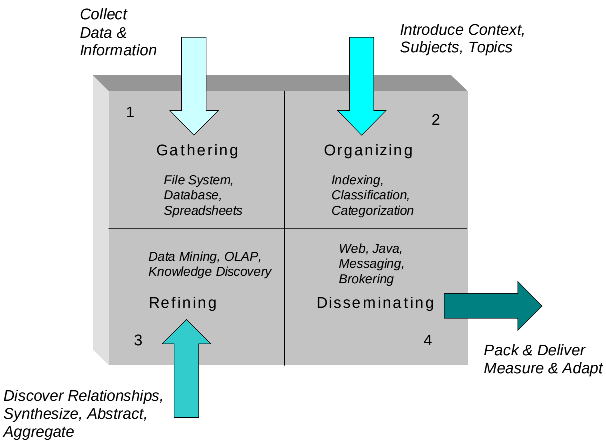

Figure 1. Generic Knowledge Management Process

The first stage is the gathering or capture of raw data or information collected from operational processes. This information may be in the form of structured data (e.g., relational databases, log files, event traces), or in the form of semi-structured or unstructured documents. The second stage organizes the information through indexing, categorization and classification into a domain-specific ontology consisting of contexts, subjects, and topics. The third stage includes a battery of automated refinery processes that include on-line analytical processing (OLAP), data mining and knowledge discovery, and information visualization techniques [Fay96, Min99]. The results of this stage are usually in the form of actionable knowledge that can be disseminated, shared, and delivered to knowledge workers for decision making. Today, much of this dissemination is Web-based.

We will try to classify typical KM scenarios as reported in industry and academia. The most frequent situation is known as <u>Knowledge Sharing</u> (Figure 2. - left) in which KM is seen as the technique to transfer expertise from the top performers or domain experts (20%) to the rest of the population (80%). In the IT support and services context, for example, this involves the transfer of problem solving knowledge from the more proficient service desk agents to other agents, thus improving the productivity of the whole population.

The next common scenario is known as <u>Knowledge Harvesting</u> (Figure 2. - right) where we collect knowledge codified in various repositories. A user harvests knowledge by querying the various repositories. In response, a list of solutions and answers is offered to the user, who will thus experience augmented problem-

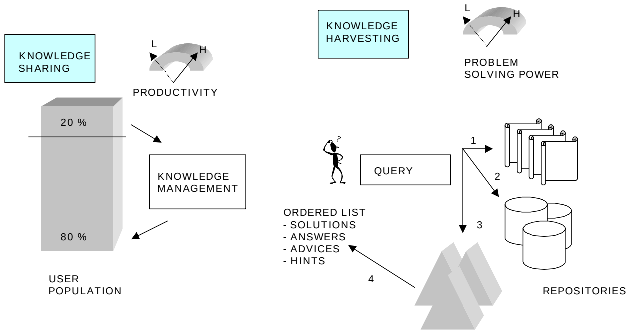

solving power.

Figure 2. Knowledge Sharing and Knowledge Harvesting

In the first two scenarios, the knowledge management system does not differentiate among users. However, the ultimate objective of IT tools is to be able to learn and adapt to individual users' habits, preferences and evolving needs. There are several products that are able to capture interactions with users via sophisticated adaptive algorithms in which the key asset is <u>Captured Knowledge</u> (Figure 3, left). Since this knowledge is captured in a suitable user model, it can be used to provide a user with solutions that are customized to his needs, and hence directly impacts the user's productivity.

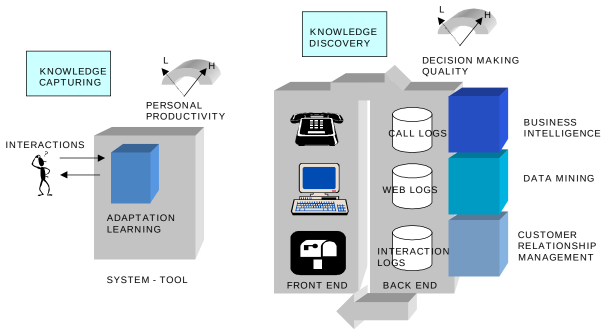

Figure 3. Knowledge Capturing and Knowledge Discovery

Finally, the most recent KM paradigm is known as Knowledge Discovery (Figure 3, right) wherein many different types of information about user interactions (e.g., transaction logs, case histories, web logs, call logs, traces of problem solving sessions) are amassed and analyzed with sophisticated large-scale algorithms. These create insights and recommend optimal business actions aimed at improving the quality of decision making.

3 IT Knowledge Management

Different types of knowledge are encountered in IT domains: product knowledge, procedural knowledge, legal knowledge, behavioral knowledge, customer knowledge and topological knowledge Within the domain of IT support and services, knowledge management can be regarded as a process that impacts productivity and learning. It is typically achieved through the capture, articulation, and **reuse of relevant domain knowledge**. This strategy fits domains in which the problems encountered are simple, repetitive in nature, and for which standard solutions exist.

For the purpose of this article, we will define key IT terms: data, information, knowledge and wisdom. In the typical IT service domain, **data** is a collection of observed facts or events, such as "3 disk errors from server xyz have been recorded in the last 10 minutes". **Information** is derived from data by summarizing or aggregating data from several sources and over a period of time, e.g., "the failure rates of systems with a configuration similar to server xyz is 5% over a year; or, server xyz has been down 20% of the time in the past 3 months, and such failures affect 250 users in the production department". **Knowledge** is in the form of business rules or patterns derived from large collections of data and information, e.g. "3 disk errors within 15 minutes from systems similar to server xyz are predictive of server failure with 90% confidence". Deriving actionable business decisions and insights from knowledge gained over a huge population of IT systems and over very extended periods of time creates **wisdom**, e.g., "It is cost-effective to invest in a high availability solution with server fail-over capability for the production department, because a server outage there results in a loss of 500 person days of work." Business enterprises that are able to efficiently manage these kernel entities (data, information, knowledge, wisdom) are typically market leaders and consistent business winners.

Business enterprises are huge generators and consumers of data, information, and knowledge. Terabytes of data are

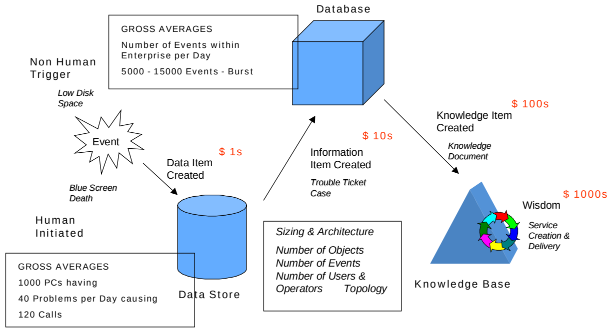

Figure 4. Event and Value Chain within Enterprise IT

created by machines and humans in their daily operational tasks. Generic event creation rates within the infrastructure of a typical enterprise are illustrated in Figure 4 in the context of PC support and service. One must take into account these figures when planning and sizing an enterprise architecture. From experience, we can attach a dollar value to each item in the chain and talk about its positioning in the value chain. So, if we spend $1 to capture event traces, then additional processing might bring us into the $100 range per knowledge item (e.g., document) if we are able to create one in a cost efficient manner. Some companies will still earn money by exploiting even a simple data capture process. Others will cash in on the wisdom acquired from careful knowledge management and transforming knowledge into the ability to deliver superior end-to-end services.

In the generic IT landscape, one typically encounters a multi-tiered operation consisting of help desks or service desks, and operational service centers supporting the IT infrastructure and delivering various services (Figure 5).

Activities in the help desk and service desk require knowledge sharing and reuse. One analyst usually covers 200-300 desktops, handles daily 20-30 phone calls, whose duration is typically 20 minutes and whose average cost is around $50. Service Desk analysts are focused on simple information inquiries, desktop problems, and application usage problems. Their work is usually supported by a problem-solving knowledge base. This knowledge base typically consists of a large collection of documents (case histories of problems encountered before and their resolution, product documentation, and the like), and other troubleshooting and diagnostic tools.

Operational centers within IT departments in big enterprises take care of more complex problems, for which unique solutions are created through the communication and cooperation of several human experts. (A typical enterprise has $3 billion in annual revenues, 9000 employees, 500 IT professionals, 4300 desktops, 470 servers spread over 40 sites.) Knowledge management here usually involves **indexing of human expertise** to enable **collaboration**. IT personnel in the operational center are usually focused on more complex domains such as network management, server management, and system management. They normally encounter complex but non-repetitive type of problems that are frequently specific to the particular set-up and the particular enterprise.

All IT processes are supported by computerized systems, so that huge volumes of data are created and stored daily. The process of transforming data into information and knowledge is not completely automated, and recent investigations in text mining and business intelligence represent efforts in that direction. Knowledge is **extracted and captured** in a model able to emulate certain human behaviors. Such knowledge could be embedded in products (giving differentiating or innovative features) or it could drive various processes.

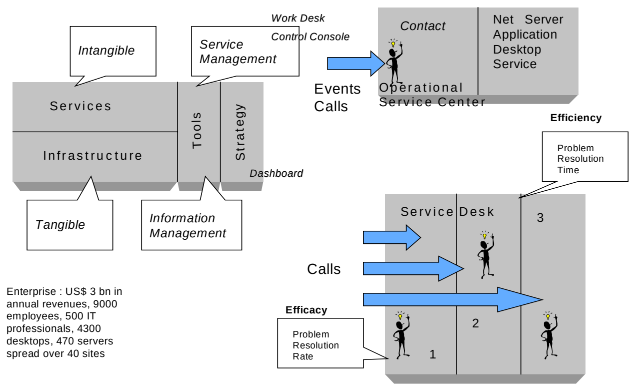

Figure 5. Generic IT Landscape: Service Desks And Operational Centres

Strategic managers and senior managers are typically supported with knowledge automatically extracted from different repositories deployed within a decision support system [9]. For these managers, the primary purpose of discovered knowledge is to create business insights (for example, why are users from one line of business more efficient than others). We believe that certain core technologies such as data mining and machine learning, Bayesian reasoning, neural networks and genetic algorithms, information retrieval and natural language processing, information visualization, and intelligent agents, some of which have being actively investigated during the last 50 years, will play an important role in the future development of decision support services. Accumulated experience and insight from these large fields of research are beginning to appear in innovative applications and tools. The notion of "knowledge" provides the glue for combining and exploiting techniques from these different research disciplines.

Enterprise information systems should be designed with the user model explicitly included into all aspects of the design. The primary point of focus for user interaction is a user view or portal, which provides the working context for the user, guides the user through his tasks, serves up relevant data, information and knowledge to aid him in decision making, and (ideally) evolves and adapts to the user's needs. The machine emulates intelligent activities such as answering questions, giving advice, solving problems, and playing what-if scenarios.

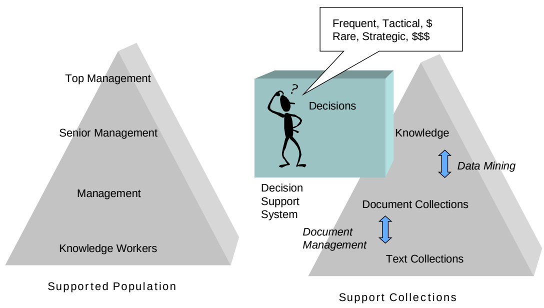

Figure 6. Cooperative Systems: Users and Collections

In Figure 6, we sketch segmented user populations and repositories used to support decision making. All users are considered to be knowledge workers, the principal difference among them being the frequency of decision making and the associated value-at-risk. Typically, we provide decision support systems for two main areas: first, for users who have to make frequent and not too risky or costly decisions (service-desk agents, for example); and second, for users who have to make infrequent but risky or costly decisions (a CIO, for example).

4 An IT Support & Service Architecture

At Hewlett-Packard, we have implemented an IT architecture for support and services that is based on the principles of knowledge management outlined above [Del98]. This architecture is depicted in Figure 7. In their daily work, our help desk analysts use HP WiseWare[1], a knowledge based system that contains various types of knowledge documents covering well over 70 standard products, as well as customer specific problems, for a total of about 150,000 solutions. Approximately 70 to 80 percent of the analysts use the WiseWare knowledge base regularly. To illustrate the problem-solving power of WiseWare, let us mention that it is equivalent to a library containing 2000 books of 100 pages, which could be packed onto 20 bookshelves, each containing 100 books. From the word count perspective, WiseWare contains circa 47 million words, which is more than any currently available on-line encyclopedia. A search engine enables pinpointing of the relevant content with 2,3 word long query (long-term average length of the query), which is equivalent to finding in the above library the appropriate subchapter (10-15 pages) or even a single page in a few seconds. On average, 8-15 percent of the users provide feedback (annotation) about content usage, and this feedback drives content management tasks.

HP RuleWare is an extension of WiseWare that provides procedural knowledge about what a customer is entitled to and procedures for handling customers. Currently, HP RuleWare contains circa 6000 rules, which capture knowledge about customers.

These two knowledge repositories represent the middle layer of our architecture. All user interaction traces are captured in the back-end system, Search & Access Mine, which is a huge repository of interaction traces containing several million sessions that are collected from several global servers, pre-processed, transformed and stored in the appropriate format. Such transformations typically create tables, charts and graphs, and drive the computation of various business indicators.

The design of HP WiseWare combines Web technology with an indexing engine to enable good coverage of Wintel-related problems for help desk analysts. A well controlled and ISO 9001-certified knowledge management process (Knowledge Refinery) enables a constant feed of relevant source material for HP WiseWare. As the user population has shifted from help desk analysts towards service desk and channel partners, the nature and complexity of the service calls and support has changed as well. For instance, there is an increased emphasis on self-support. The productivity of end users can be greatly increased by providing assistance for simple, repetitive problems. Exploiting real-time information from the user's environment can enable the expansion of another line of support services (e.g., predictive support).

[1] HP WiseWare, HP RuleWare, HP Search & Access Mine are internal HP products

Realizing that the integration of the whole support landscape (self-healing, self-support, technical support and administrative support) is an obvious future necessity, we designed the architecture to cover all relevant problem areas with competent user support aimed at the integration and correct management of data, information, and knowledge.

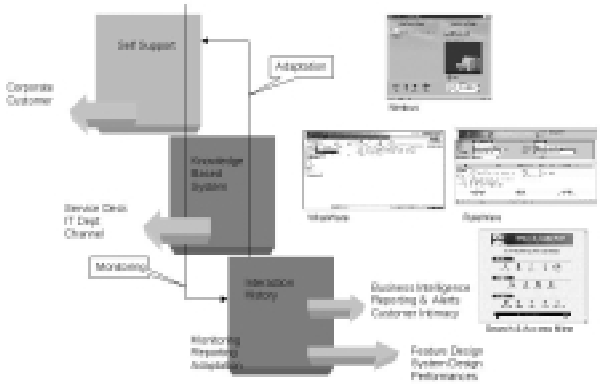

Figure 7. Layered Knowledge Management Architecture

The upper layer of the architecture contains a front-end system (which in our architecture is a Web-based portal called Nimbus) that delivers support and simple advice enabling end-users to self-solve their problems, while gathering information from the user's environment and capturing usage patterns. All user interactions with the knowledge-based system are captured. These include all queries launched, documents retrieved and accessed with user ratings of the documents.

The back-end layer of the architecture contains traces (i.e., logs) of all user interactions with the knowledge-based system. This enables the profiling of individual users, drives the adaptation (personalization) of the system's responses to meet each user's needs, and the reporting and alerting functions that are provided to the business managers. Profiling knowledge is stored in database tables and associated procedures.

To summarize, this layered architecture contains problem solving knowledge in various document repositories; knowledge about user behavior is captured in user models; and interaction models provide the basis for strategic decision making. Judicious management and exploitation of knowledge enables an evolution from problem solving assistance toward a decision support system [Tur98] for users and managers. Consequently, IT support processes will be governed by business metrics and measurements, instead of rough indicators and intuition.

5 Enterprise Service Architecture

IT support services are just one of the touch points that typical enterprises have with their customers and partners. In general, within an enterprise we find several different repositories containing data, information, and knowledge about different aspects of the business. For instance, document repositories contain project, product, process, and workflow information. Structured databases contain financial data, accounting, sales and marketing figures, operational and business tables. Messaging systems contain communication and collaboration traces. Interaction histories contain log files and web-access interaction databases. Various data & knowledge feeds come in and out of the enterprise. Ideally, we would like to create an integrated architecture that will augment the support services we described earlier with many other types of decision support services for various corporate users and for various applications such as customer relationship management, marketing campaign design and optimization, consulting, and so on.

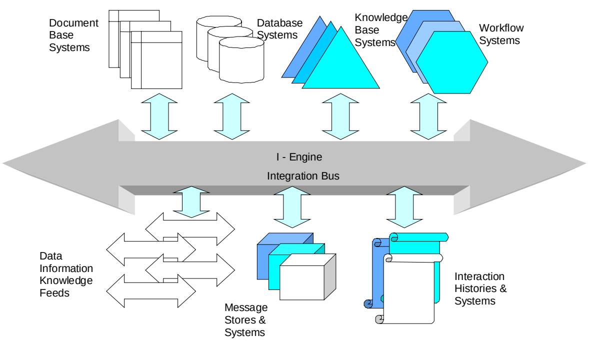

Figure 8. Integrated Enterprise Service Architecture

This requires an architecture that supports the integration of the various repositories, the creation of business intelligence from the information captured in these repositories, and the creation of services that exploit this business intelligence. We believe that basing such an enterprise integration architecture on a publish-subscribe paradigm offers important benefits. In this paradigm, the different components communicate via messages. Components can subscribe to particular types of messages (events) published by other components. The publish-subscribe middleware provides message brokering, i.e., transparency between publishers and subscribers, asynchronous delivery of messages, scalability and high availability.

We see this message-brokering paradigm evolve into a knowledge brokering paradigm, where components can publish and subscribe to knowledge (IT knowledge, business knowledge, legal knowledge, financial knowledge), not just to messages or events. Ultimately, this will lead to a dynamic marketplace for electronic services [Qim99] within an enterprise, and even across multiple enterprises, in which service providers advertise their capabilities and the types of knowledge they can provide, and consumers can dynamically find service providers that can satisfy their requirements.

6 Conclusions

Knowledge is a crucial ingredient for enterprise IT service & support. Knowledge may appear explicitly in documents, it may be embedded in algorithms, tools and processes that assist users in automatic problem solving. Knowledge Management is still more an art than an established scientific theory or commonly accepted methodology. It combines several techniques and technologies aiming to capture, articulate and disseminate knowledge [IDC99, Gar99]. In different areas, it takes different forms, but its absence will always be very palpable. Companies that are able to transform data and information (costs) into knowledge assets (values) can claim that they have successfully applied Knowledge Management.

References

[Dav00] Thomas H. Davenport, Laurence Prusak, *Working Knowledge*, Harvard Business School Press, May 2000.

[Del00] Kemal A. Delic and Birgit Hoellmer, Knowledge-Based Support in Help-Desk Environments, *IEEE IT Professional*, Vol. 2, No. 1, pp. 44-48, January/February 2000.

[Fay96] Fayyad, U., Piatetsky-Shapiro, G, Smyth, P., and Uthurasamy, R., (eds.), *Advances in Knowledge Discovery and Data Mining*, AAAI Press/ MIT Press, 1996.

[Min99] Hao, Ming; Dayal, Umesh; Hsu, Meichun; Baker, Jim; D'Eletto, Robert, A Java-based Visual Mining Infrastructure and Applications, HP Labs Technical Report HPL-1999-49 available at http://www.hpl.hp.com/techreports/1999/HPL-1999-49.html

[Del98] Kemal A. Delic and Dominique Lahaix, Knowledge Harvesting, Articulation, and Delivery, *HP Journal*, Vol. 49 No. 2. pp. 74-81. May 1998.

[Qim99] Chen, Qiming; Hsu, Meichun; Dayal, Umeshwar; Griss, Martin, Multi-Agent Cooperation, Dynamic Workflow and XML for E-Commerce Automation, HP Labs Technical Report HPL-1999-136 available at http://www.hpl.hp.com/techreports/1999/HPL-1999-136.html

[IDC99] IDC, Bulletin "Knowledge Management Factbook", September 1999.

[Gar99] Gartner Group, Conference Presentation, Knowledge Management - Everything and Nothing, 1999.

[Tur98] Efraim Turban and Jay E. Aronson. *Decision Support Systems & Intelligent Systems* (5th ed., Prentice Hall, 1998; ISBN 0-13-740937-0)