EIS

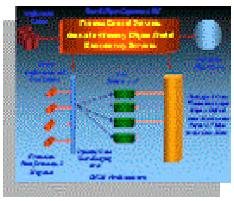

DKMS Briefs

Joseph M. Firestone

DKMS Brief No. Eleven: My Road to Knowledge Management through Data Warehousing

I'm going to begin this brief with a story about how I entered Knowledge management (KM). I'll then comment on this and some alternative roads to KM.

My odyssey from Data Warehousing to KM began in 1991, when, after an 8 year absence from analytical work, I joined a small database marketing firm located in Washington, DC and discovered the world of database marketing, which later was one source of Customer Relationship Management (CRM). I was hired as the firm's Director of Statistical Services, and soon found myself immersed in data mining, fuzzy logic, survival analysis, neural networks, and customer retention and lifetime value models. That work led me to two discoveries that put me on the road to KM. First, I discovered client/server enterprise computing and soon after Object Technology. And second, I discovered data warehousing, a field that coupled with database marketing, and data mining placed me back in touch with my career-long interest in Executive Information Systems and Decision Support Systems (DSS).

As the '90s passed and I became more deeply immersed in data mining and data warehousing, and as I observed the progress in enterprise computing, I began to realize that data modeling and data-centric approaches to Data Warehousing and DSS were inadequate for representing the complex systems models, causal models, and process models necessary for sophisticated decision support. Also, as the data warehousing IT environment became increasingly more complex with diverse databases, application servers, and business process engines, I concluded that the old data-centric approaches would not suffice to model the change processes characteristic of more complex data warehousing systems. That is, I believed that a Dynamic Integration Problem (DIP) was developing in DSS systems, that this problem was due to the data-centricity of data warehousing practices and software capabilities, and that an object layer was sorely needed to better integrate DSS systems.

Spurred by this conclusion I set out to persuade others that an Object-Oriented (OO) approach to data warehousing was the right way to go. In August 1997, I wrote a White Paper called “Object-Oriented Data Warehousing.” You can still find it posted at www.dkms.com. I don't think either that paper, or my views on O-O Data Warehousing were very successful in influencing that field. But they certainly made me develop a coherent distributed object/component view of computing.

When I contemplated that system, and realized that it would contain formal models that were validated by human investigators in the process of data mining, I concluded, further, that this distributed computing system would support the production of “knowledge” - i.e. empirically validated models. Therefore it could be considered a *Distributed Knowledge Management System* (DKMS). Later that same August, I wrote a White Paper called "Distributed Knowledge Management Systems: the Next Wave in DSS" (see www.dkms.com), a revision of the O-O Data Warehousing paper, to record my new insight.

I don't know if that paper was any more successful than the first one. But it did one very important thing for me. It, along with some then very current material from a special issue of CIO magazine, persuaded me that I was working in a field broader than Data Warehousing, a field called Knowledge Management that had been developing very well during most of the 1990s, and whose development I had missed. So, I began to read everything I could about KM, and also began a transition, by now complete, from someone practicing data warehousing to someone practicing KM.

And, really, this transition suited my previous interests better than Data Warehousing had. My DSS and data mining, and data warehousing background were all still relevant, but, in addition, my previous work in General Systems Theory, Social and Political Theory, Philosophy of Science, Epistemology, Value Theory, Theory of Measurement, Operations Research, Sociology, Social Psychology and Conflict Theory were all enormously relevant as well. Who ever said, "you can't go home again"? Here I was, at home in Knowledge Management, at last.

Considering my story in retrospect, some comments on it may be instructive. First, what led me to KM was the same thing that led me to Data Warehousing, namely an interest in developing solutions to the problem of providing effective decision support. But this implies an assumption, which you may or may not share, that “effective decision support” requires much more than timely and flexible access to comprehensive, properly modeled, and high quality historical data supplemented by excellent “slicing and dicing” and reporting capabilities. Specifically, I had become dissatisfied with data-centric approaches to Data Warehousing because I found them lacking in providing support for formal modeling and formal models. And I conjectured that (a) better support for modeling than that provided by Data Warehousing, even supplemented by Data

Mining capabilities, was needed for decision support and (b) KM as a field was friendly to such modeling.

Second, I was attracted by Object Technology because it seemed to provide a better capability for working with analytical models than more standard approaches to Data Warehousing. Still more specifically, I realized that conceptual and knowledge networks could be represented as objects much more conveniently than they could be represented in other modes. And this realization too, pointed toward KM, since one aspect of it is a concern with the principles and technology used to manage information and knowledge networks.

Third, I also associated an O-O Data Warehousing approach with KM because I associated it with Knowledge Production. That was the primary consideration that led me to identify myself with KM as a field, to begin my transition to it, and to rename my approach Distributed Knowledge Management Systems. In retrospect however, I now see that this identification was a mistake, one due to my own initial lack of understanding of KM.

In the beginning, I made no distinction between Knowledge Production, and more broadly, Knowledge Processing, and KM. Readers of my first IKM column on "KM and IT" will recognize that KM is management of Knowledge Processing and its outcomes and that it does not, in general, refer to Knowledge Processing itself. Since my first formulations of Distributed Knowledge Management (DKM) architecture were about an envisioned software product that would support Knowledge Processing, I incorrectly named it a KM System when I should have called it a Distributed Knowledge Processing System (DKPS). Therefore I incorrectly concluded that my DKMS formulation placed me in the KM field. I could, instead, have contented myself with a characterization that was closer to Data Warehousing. Since then, I've modified DKM architecture to also support activities that manage Knowledge Processing and its outcomes, and further developed my architectural approach in the context of work on Enterprise Knowledge Portals (EKPs). So now I am squarely within KM. My initial mistake turned out to be a fortunate one given the richness of the field and of my activities within it. These range from IT applications, to theory about the scope and boundaries of the field of KM, to KM Metrics, to Methodology.

My road to KM through Data Warehousing is only one of many available. Two alternatives are suggested by: (1) a desire to "get the right information to the right people at the right time"; (and 2) a desire to get beyond "arid IT-based" concerns and to take the human-side of decision support into account. I'll comment on each of these.

First, "getting the right information to the right people at the right time" illustrates another conceptual mistake people often make about KM. Unless "the right information" in this expression refers to knowledge, the expression is clearly about information Processing, or perhaps, if interpreted very liberally, about

Information Management (IM). Of course, this comment implies that there is a meaningful distinction between “knowledge” and “information”. But KM must assume that such a distinction is meaningful. If it were not, there would be no KM.

Historically the above slogan, however mistaken it may be, has attracted many people to KM and has been a guiding precept of IT-based approaches to KM. My collaborator Mark McElroy and I have argued that this idea is the basis of the first generation of KM, which primarily views KM as an IT field. In spite of this history, this idea, insofar as people believe it, represents a fundamental weakness in the field, since one day, people will realize that “getting the right information to the right people at the right time” is about IM and not KM.

Second, the desire to get beyond “arid IT-based” concerns and to take the human-side of decision support into account, is about a view of KM that sees knowledge as subjective and personal in character, largely “tacit” or “implicit”, and as distinct from codified expressions, which are really not knowledge, but only information. Knowledge is frequently viewed as “justified true belief” in this approach, a definition that has been the dominant one in philosophy since Plato, but which has been under vigorous attack since at least the 1930s. People who take this road to KM, view it as primarily an applied social science discipline, whose role is to “enable” better knowledge creation and sharing by facilitating the “conversion” of tacit and implicit knowledge to codified expressions.

The problem with this road to KM is that (a) in viewing knowledge as “justified true belief” it makes it dependent on the “knower” and therefore basically subjective. And (b) in restricting knowledge to beliefs in the mind, it neglects the role of management in providing a framework of rules and technology for testing and evaluating codified expressions or knowledge claims and thereby creating a basis for producing objective knowledge. In a number of other places, I’ve specified two types of knowledge found in organizations: surviving beliefs and surviving knowledge claims. In restricting attention to facilitating expressing surviving beliefs alone, this road to KM misses one of its major objectives: to enhance Knowledge Production and, in this way, indirectly improve the quality of surviving knowledge claims used in future decisions.

[Portions of this brief were previously published in the Preface of Joseph M. Firestone, *Enterprise Information Portals and Knowledge Management*, Burlington, MA: KMCI Press/Butterworth-Heinemann, 2003.]