Chapter 3

Enterprise Knowledge Clouds: Next Generation Knowledge Management Systems?

Jeff Riley and Kemal Delic

Contents

3.1 Introduction

The field of knowledge management (KM) has been through several cycles of hype and disappointment and has created important disputes along the lines of "knowledge" being the philosophical discourse and science and "management" being the empirical and experiential teaching. In reality, it has created a booming business for technologists, consultants, and a wide variety of technology vendors. Seen today, after a few decades, it still seems that the term "knowledge management" remains undefined, fuzzy, and disputed. Indeed, the very definition of "knowledge" and the distinction between data, information, and knowledge is poorly defined and still not well understood by many. Still, we are all aware of the rudimentary elements of "knowledge reuse" in a wide variety of business operations. Some will even hint that the contemporary Internet is a type of "knowledge bazaar" where individuals and corporations can shop for all manner of "knowledge consumables." In this chapter, we discuss the possibility that the treatment of knowledge management systems (KMS), which represent an intricate part of many business enterprises, has yet another chance of reappearing in totally new technological, market, and social circumstances.

In Section 3.2, we sketch the context in which we see the emergence of massive, globally dependable infrastructure(s) used by several hundreds of millions of users across the globe. We position business aims and interest in this subject and narrow those into enterprise needs for knowledge to operate. We then outline a generic knowledge management architecture within contemporary business enterprises, which typically appears in the form of the enterprise stack application, hosted in data centers.

After observing the current deficiencies and projecting future developments, we depict a high-level architecture for the Enterprise Knowledge Cloud (EKC) as a collaborative, cooperating, competing mega-structure providing computing, networking, and storage services to various "knowledge producers and consumers"—such as devices, people, and applications. Some architectural and design landscapes are provided for illustration. We conclude with a no-nonsense list of things we expect to observe happening as a sign of the mega-shift from the industrial to post-industrial world of the twenty-first century. We believe that EKCs are potential breakthrough applications marking an enterprise technology transition all the way from mainframe computers to networked PCs, to grids and emerging computing clouds.

3.1.1 Emerging Cloud Computing Infrastructures

The U.S. National Institute of Standards and Technology (NIST) draft definition of cloud computing is as follows:

> Cloud computing is a model for enabling convenient, on-demand network access to a shared pool of configurable computing resources

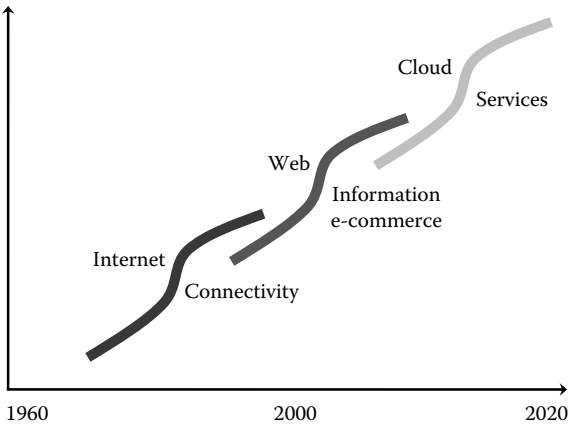

Figure 3.1 The Internet evolving into the cloud. (Adapted from Delic, K.A. and Walker, M.A., *ACM Ubiquity*, 9, 31, 2009.)

> (e.g., networks, servers, storage, applications, and services) that can be rapidly provisioned and released with minimal management effort or service provider interaction. (Mell and Grance 2009)

Based on this definition, it can immediately be seen that cloud computing encompasses infrastructure, which to some extent at least already exists today in the form of the World Wide Web ("the Web") providing a wide variety of information technology (IT) services that can be purchased on-demand (computing cycles, storage, and network) in a highly simplified procedure. Over time, just as "the Internet" has evolved into "the Web," the Web will evolve into "the cloud" (Figure 3.1).

We predict future growth in which we will see a huge number of common devices interconnected and totally new applications emerging. It will most likely emerge as a hugely re-scaled version of today's Internet. This growth will likely be stimulated via innovative applications starting to proliferate: a well-known social network has provided a platform for 4,000 applications written by 80,000 developers in just 6 months; the Amazon Elastic Compute Cloud (Amazon EC2) has 330,000 registered application developers.

We observe that the cloud infrastructure is global, highly dependable, and supports innovative business models and new types of social phenomena such as blogs, MySpace, Facebook, YouTube, and Twitter, not to mention the myriad multiplayer, role-playing games and virtual reality sites.

3.1.2 Collective Intelligence

Collective intelligence is a phenomenon that emerges from the interaction—either collaboration or competition—of many individuals. By some estimates, today there

are 80 million people worldwide writing weblogs ("blogs"). The blogs are typically topic-oriented and some attract important readership. Authors range from large company CEOs to housewives and young children. When taken together, the cloud computing infrastructure that hosts *blogospheres* looks like a big social agglomeration providing a kind of "collective intelligence." But it is not just blogs that form the collective intelligence—the phenomenon of collective intelligence is nurtured and enhanced by the social and participatory culture of the Internet, so all content developed and shared on the Internet becomes part of the collective intelligence. The Internet then, and the content available there, appears as an omnipresent, omniscient, cloud-like giant infrastructure—as a new form of "knowledge management." Today this represents a massive collaboration of mostly people only, but very soon in the future we may envisage intelligent virtual objects and devices collaborating with people—this is already beginning to happen to some extent with Internet-attached devices starting to proliferate. Thus, rescaling from the actual ~1.2 billion users to tens or hundreds of billions of real-world objects having a data representation in the virtual world is probably realistic. A real danger, and a real problem to be solved by knowledge management practitioners, is how to sort the wheat from the chaff—or the knowledge from the data and information—in an environment where the sheer amount of data and information could be overwhelming.

3.1.3 Intelligent Enterprise

Business enterprises today use the existing Internet infrastructure to execute various business operations and provide a wide variety of services. As we see the shift of all nonphysical operations versus the Internet, we observe a new type of enterprise emerging: we call it the *Intelligent Enterprise* (Delic and Dayal 2002).

The Intelligent Enterprise is able to interact with its environment and change its behavior, structure, and strategy—behaving actually as an intelligent entity. It is able to adapt to rapid changing market circumstances, gradually change its business model, and survive into the next market cycle. The Intelligent Enterprise as we see it is characterized by its ability to learn from and adapt to changes in its environment and reinvent itself, sometimes with surprising results. In order to keep up with the rapidly changing demands of doing business, most enterprises implement increasingly complex IT solutions. Although implemented to make the enterprise more efficient, coupled with the organizational complexity of such large enterprise business, the technical complexity introduced by the many and varied IT solutions helps create pockets of inefficiencies within the organization. We see future Intelligent Enterprises deriving efficiencies through the automation of their core business processes and the exploitation of knowledge inherent in their organization. Their ability to respond quickly to changes will improve significantly as the knowledge base and "intelligence density" within the enterprise grows and problem-solving capabilities improve dramatically. Intelligent Enterprises will form dynamic partnerships with other enterprises to create dynamic business ecosystems,

which will be self-managed, self-configured, and self-optimized. In short, future enterprises will become smarter—more intelligent—and by doing so will evolve automatically into organizations more suited to their changing environment.

We postulate that the emergence of collective intelligence in the cloud computing infrastructure will influence markets and established businesses, allowing—even encouraging—Intelligent Enterprises to emerge and reshape the contemporary approach to enterprise knowledge management (EKM). Next, we describe the current state of EKM.

3.2 Enterprise Knowledge Management: Architecture and Technologies

Constantly evolving markets exercise pressure on business enterprises to continually evolve and improve. One of the most widely used business paradigms is about EKM—as a means to capture and express tacit human knowledge into an explicit form (externalized knowledge or content) that could be later (hopefully) reused. Various schools of thought were proposed, several assistive technologies were developed, and an important number of successful EKM stories were reported. From our experience, the best domain for EKM is in the enterprise IT domain (Delic and Dayal 2000, Noël and Delic 2002, Delic and Douillet 2004)—as it is a domain under huge cost pressure, but one which is essential for strategic development.

From a highly abstracted view, the EKM IT domain consists of problem solving, monitoring, tuning and automation, business intelligence reporting, and decision-making tasks (Figure 3.2).

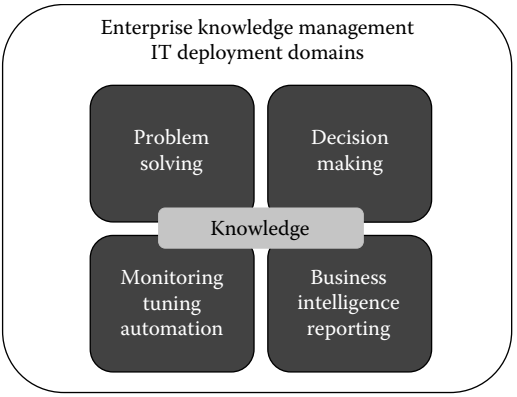

Figure 3.2 EKM: IT deployment domains. (Adapted from Delic, K.A. and Riley, J.A., Enterprise knowledge clouds: Next generation KM systems? *Proceedings of the 2009 International Conference on Information, Process, and Knowledge Management* (e*Know 2009*), Cancun, Mexico, February 2009.)

Problem solving, especially in the EKM IT domain, is the task for which knowledge management techniques and systems are most commonly deployed. The proliferation of knowledge management systems for problem analysis and solving is many and varied, spanning the gamut from knowledge capture, representation, and transformation through to recognition, extraction, and reuse. Knowledge from all sources, including human expertise, in the form of plain text, models, visual artifacts, executable modules, etc. is used by intelligent knowledge management systems to enable users to solve problems without reference to scarce, and often expensive, human experts.

In recent years, a wide variety of artificial intelligence (AI) techniques and heuristics have been deployed in knowledge management systems in an effort to make the systems smarter and more responsive. These smarter knowledge management systems are particularly well suited to automation and self-management tasks, where the goal is to provide automated monitoring of system use and predictive tuning of system parameters to achieve automatic system scale out.

Business intelligence (BI) refers to a range of methodologies, technologies, skills, competencies, and applications businesses implement and utilize in order to better understand their commercial context. Typically business intelligence systems are knowledge management systems that provide current and predictive views of the business based on historical and current data relating to the business itself and the commercial environment in which it exists. Business intelligence reporting is more than the simple reporting of data gathered—it uses a wide range of AI techniques to extract relevant knowledge from incoming data streams and repositories and provides observations, hints, and suggestions about trends and possible futures.

Decision making is most often done by humans after understanding the results of the business intelligence reporting, but with the volume of business intelligence available to analysts increasing almost exponentially, it is becoming more and more difficult for humans to make sensible, rational, and timely decisions, so this task is increasingly becoming the responsibility of AI systems tuned to the environment of their deployment.

The tasks of problem solving, monitoring, tuning and automation, business intelligence reporting, and decision making are the most promising areas for the future deployment of EKCs. These areas will have a special flavor for the development of a slew of new technologies addressing the problems that previous computing facilities could not resolve.

Currently, the majority of the indicated IT tasks include people, while we suggest that this balance will be changed in the future through automation, ultimately leading to self-managing enterprise IT systems (Delic et al. 2007). When mapped into a more precise form, this conceptual drawing (Figure 3.2) will evolve into the enterprise-scale knowledge management application stack (Figure 3.3).

All knowledge management applications today can be layered into three essential subsystems:

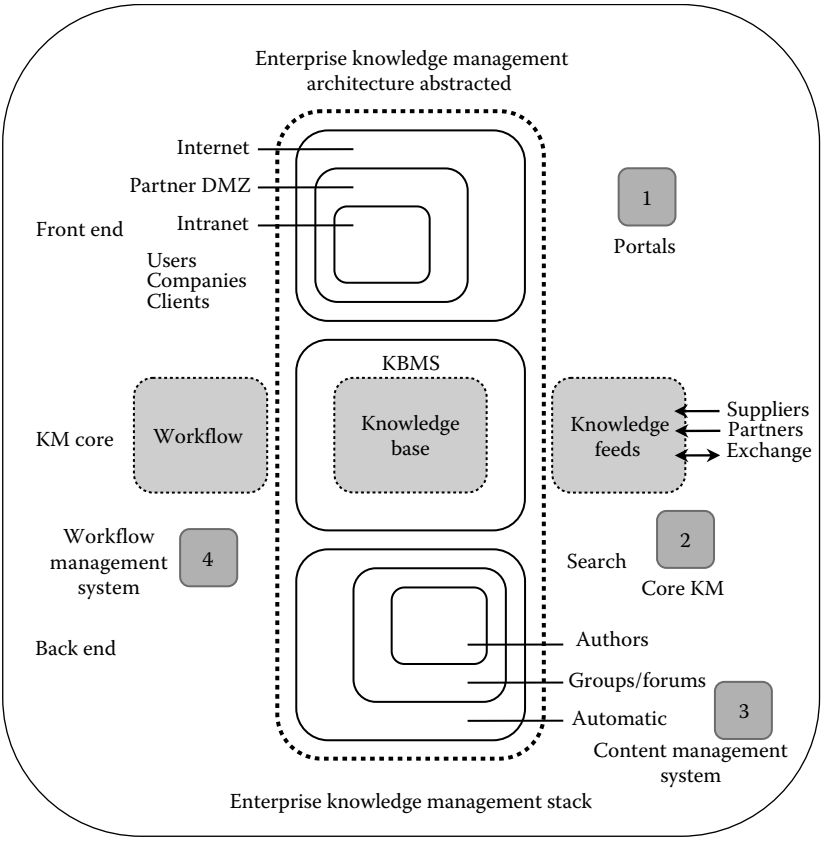

Figure 3.3 EKM: Architectural view. (Adapted from Delic, K.A. and Riley, J.A., Enterprise knowledge clouds: Next generation KM systems? *Proceedings of the 2009 International Conference on Information, Process, and Knowledge Management* (e*Know 2009*), Cancun, Mexico, February 2009.)

- Front-end portals that manage interactions with internal users, partner's agents, and external users while rendering various "knowledge services." Different classes of users—e.g., internal vs. external—are often presented with slightly different portals allowing access to different types of knowledge and services.
- A core layer that provides the knowledge base and access/navigation/guidance/management services to knowledge portals and other enterprise applications. The core layer provides the Knowledge Base Management System (KBMS), the knowledge feeds—the means by which knowledge is added to the knowledge base or exchanged with other knowledge management systems or users—as well as the mechanism to distribute and inject appropriate knowledge into business processes throughout the enterprise.

- The back-end that supplies "knowledge content" and the content management system from various sources, authors, and communities that enables a refresh of the knowledge base.

The Enterprise Workflow System captures interactions with users and provides necessary context for the EKM system. Various feeds enable the flow and exchange of knowledge with partners and suppliers. Today these feeds are mainly proprietary, while we expect that they will evolve into standards-based solutions for large-scale content flows (RESTful services, RSS, ATOM, SFTP, JSON, etc.). To indicate the scale and size of the typical corporate knowledge management system, we presume that the knowledge base contains several million knowledge items, and users number in the hundreds of thousands. EKM is considered a high-end, mission-critical corporate application, which resides in the corporate data center. High availability and dependability are necessary engineering features for such global, always-on, always-available systems.

Thus, EKM is typically a three-tier enterprise application probably spread over several geographically dispersed data centers and typically interconnected or integrated with enterprise portals, content, and workflow management systems. In essence, EKM consists of the enterprise knowledge base (KB) with appropriate knowledge management routines (add/remove/modify KB), whose content is usually accessed via search access routines.

The ultimate result is that we are witnessing emerging social phenomena (writing blogs, participating in social networks, collaborating in wikis) enabled by an always-available, globally accessible, and secure infrastructure that can be used for free, or at a very low-cost, and running a mushrooming number of user-created applications. Some major companies are already announcing their intention to enter, drive, and dominate this field (*The Economist* 2008).

3.2.1 Enterprise Knowledge Management Infrastructure

Enterprise data centers are the key computational, storage, and network resources arranged around an efficient corporate network as the backbone of the enterprise IT infrastructure. Consequently, they are designed in such a way that the enterprise applications are categorized according to their criticality and provided with adequate infrastructural support. Thus, if many millions of users are critically dependent on an application, it would be categorized as a mission-critical, nonstop application and would be supported 24×7 and be always available. Some less critical applications will have yet another label, be supported 24×5, will not be considered nonstop, and would be something less than always available.

Thus, for EKM, if the risk of monetary and/or reputation loss is high, we will provide the infrastructure (clusters or high-end servers with some distinctive disaster recovery capabilities) and support, which will fulfill expectations and fit into dependability requirements—with appropriate trade-offs between cost and features.

3.2.2 Enterprise Knowledge Management Applications

Once we have categorized our EKM needs and provided the appropriate infrastructure, we should architect, design, and engineer EKM applications so that they fit into the entire EKM criticality. Thus, if the infrastructure is mission-critical, then EKM should have all the necessary features of a mission-critical application. It is out of the scope of this chapter to discuss this in more depth, but one should be well aware of this requirement as it will have implications for the software architecture, choice of operating system, platform, and programming environment: they should all respect the criticality label of the EKM system.

3.2.3 Enterprise Knowledge Management Content

Having briefly described the EKM infrastructure and applications, we should consider how enterprise knowledge will be represented, captured, processed, and delivered. Problem solving documents (Problem Description-Problem Solution) are the most simple and widely used way of capturing problem solving tasks. Some early EKM systems used a rule-based representation of knowledge; executable models (decision trees, case-based reasoning systems, neural networks) are more recent knowledge capturing paradigms. We believe that multimedia content will become dominant in the future and that new methods for knowledge capture and rendering will be devised.

3.2.4 Enterprise Knowledge Management Users

The evolution of technology in consumer and corporate domains has created a new type of user who will be very different from contemporary users. While sketching the architecture of future EKM systems, one should seriously analyze and consider several aspects and dimensions of future users. The best way would be to look at our children: they seem to have developed a way to quickly exchange information snippets, being either very short text messages or particular multimedia content. Also, it seems that they have a much better ability to multitask naturally while not losing or intermixing communication threads. This is the natural consequence of their exposure to gaming and new work and living styles. The so-called Millennium Generation will be the model for future users of EKM systems.

3.3 Enterprise Knowledge Cloud

Following social developments in the Internet world, it will be in the interest of business enterprises to deploy some of these new paradigms (social networks, blogging, open source) within their environments and with business intentions. Extrapolating what's going on in the open Internet, we project that enterprises will create several clouds for various purposes.

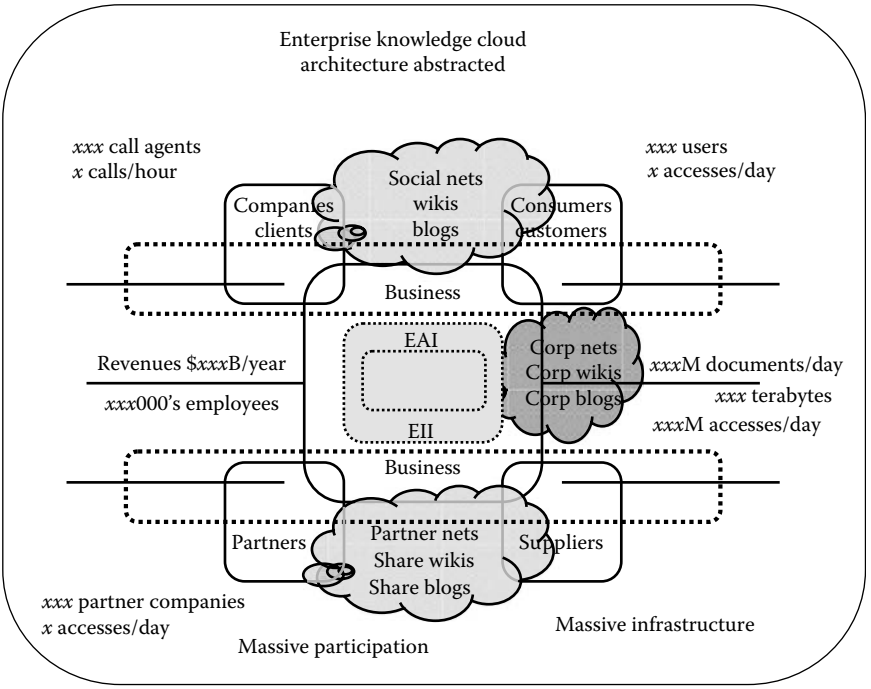

Figure 3.4 EKC: Architectural view. (Adapted from Delic, K.A. and Riley, J.A., Enterprise knowledge clouds: Next generation KM systems? *Proceedings of the 2009 International Conference on Information, Process, and Knowledge Management* (e*Know 2009*), Cancun, Mexico, February 2009.)

An abstracted business enterprise architecture is shown in Figure 3.4. This architecture interconnects business partners and suppliers to company customers and consumers and uses future cloud technologies to harvest, process, and use internal knowledge (corporate nets, wikis, blogs). Furthermore, similar partner/supplier clouds will be developed to harvest, enrich, and deploy yet another knowledge cloud. Finally, the largest enterprise cloud will cover clients and consumers, which could be used for a wide variety of purposes.

Each of the clouds shown in Figure 3.4 is an autonomous entity, existing for its own purpose and capable of collecting, warehousing, managing, and serving knowledge to its own group of users. However, while the clouds discussed are independent, they should be capable of interconnection, overlap, and knowledge-sharing with appropriate rules and safeguards, so that, for example, customers and consumers might have access to appropriate internal enterprise knowledge or even partner/supplier knowledge through the cloud.

The emergence of these clouds (Private, Partner, Public) and their coalescence into the EKC allows, indeed encourages, the collective intelligences formed within

each cloud to emerge and cooperate with each other, thus becoming the driving force for the true Intelligent Enterprise. As an example, internal IT operations will use Private Clouds, Sales, and Marketing and would operate on Public Clouds, while the Outsourcing business may reside on the Partner Clouds—each having different types of users and customers. The interaction and cooperation of the user groups, their knowledge, and the collective intelligences across the three clouds shown in Figure 3.4 provides the infrastructure for behavioral, structural, and strategic adaptation in response to changes in the respective (business) environment.

To see this happening in the future, we would expect to see the development of some major cloud computing technologies and adoption of common standards. This will enable yet another type of mega-application—Knowledge Exchange, for example, enabling the trade, exchange, and monetizing of knowledge assets. However, one should not underestimate the huge obstacles in the security, privacy, performance, and dependability of those clouds as the clear precondition for real-world deployments. One intricate problem to address will be the interoperability of clouds, leading to enabling technical standards but also aiming to establish monetary/value ground (accounting and billing systems) for the exchange of various cloud contents. All this is in a very early stage, but one should sense that developments may go in this direction.

3.4 The Next 5 to 15 Years

Today's enterprise applications are developed by IT departments, but for the future we predict a shift towards user-developed applications: mash-ups written in high-level mash-up languages. Content today is mainly text-based, but for the future we see an evolution towards multimedia context and active content (later).

Users today are either fixed or mobile—tomorrow we expect they will be virtual and later will take personalities of "avatars" to protect privacy and integrity.

Standards will evolve with the current Web 2.0 and will eventually evolve into something like Web 3.0—which we assume to be cloud computing.

Current EKM systems are enterprise applications in data centers, while we expect them to evolve into "enterprise grids" on which others envisage the development of "KM grids" (Cannataro and Talia 2003). Once the technology is stable and markets grow, we predict the development of clouds as the super-structure of enterprise grids, interconnecting enterprise data centers providing various functionalities.

Thus, while the architecture of today's EKM systems is built around the enterprise stack, tomorrow's EKM architecture will be distributed and loosely coupled and later will move to decoupled, completely pluggable, intelligent knowledge management appliances capable of adapting to interface with EKCs as required (Table 3.1).

We are in the midst of important social, technological, and market changes where we see some major companies announcing their intention to enter, drive,

Table 3.1 Evolution of EKM Systems

EKM Systems	*Today*	*Tomorrow*	*Beyond*
Architecture	Enterprise stack	Distributed	Decoupled/pluggable
Infrastructure	Datacenter	Grid	Cloud
Application	IT controlled	User produced	On demand
Content	Mainly text	Multimedia	Active
Users	Fixed/mobile	Virtual	Avatars
Standards	3W.org	Web 2.0	Web 3.0

and dominate the field of cloud computing (Weiss 2007, Forrester Research 2008, Hayes 2008). We see this as a precondition for the emergence of the intelligent, adaptive enterprise that was announced in the previous century, but can be created only in the right technological circumstances.

We believe that enterprise intelligence will draw its capacities from the EKCs embedded in the global, dependable fabrics consisting of subjects, objects, and devices. Cloud computing will enable massive and rapid rescaling of the content production, consumption, and participation of the various groups of cloud users at an unprecedented scale. This may yet evolve into a "social computing" paradigm as the likely advanced form of future society.

Massive collaboration (on content tagging, for example) followed by the emergence of ontologies based on the Semantic Web, and adjusted by the *folksonomies* developed as user-oriented Web 2.0 applications, will embody "collective intelligence" as the new source of knowledge. To see this happen, we postulate the necessity of massive, global, mega-scale infrastructure in the form of "cloud computing" (interconnected grids and data centers). We are at the very beginning of important new developments where we expect that the field of EKM will be rescaled by an order of magnitude and will spawn the creation of "a new kind of EKM system." We expect that the monetary value of the enterprise knowledge exchanges will largely surpass the cost of the use of the cloud infrastructure based on commodity components. This will fulfill an old predicament of the "content as the king" of commerce.

References

Cannataro, M. and Talia, D. 2003. The knowledge grid. *Communications of the ACM* 46(1): 89–93, January 2003.

Delic, K.A. and Dayal, U. 2000. Knowledge-based support services: Monitoring and adaptation. *DEXA Workshops 2000*, London, U.K., pp. 1097–1101.

Delic, K.A. and Dayal, U. 2002. The rise of the intelligent enterprise. *ACM Ubiquity*, 2002(45) (December 1–31, 2002) (accessed August 1, 2009).

Delic, K.A. and Douillet, L. 2004. Corporate IT knowledge workbench: Case study. *DEXA Workshops 2004*, Zaragoza, Spain, pp. 494–497.

Delic, K.A. and Riley, J.A. 2009. Enterprise knowledge clouds: Next generation KM systems? *Proceedings of the 2009 International Conference on Information*, *Process*, *and Knowledge Management* (*eKnow 2009*), Cancun, Mexico, February 2009.

Delic, K.A. and Walker, M.A. 2009. Emergence of the academic computing cloud. *ACM Ubiquity*, 9(31) (August 5–11, 2008) (accessed August 1, 2009).

Delic, K.A., Riley, J., and Faihe, Y. 2007. Architecting principles for self-managing enterprise IT systems. *Proceedings of the Third International Conference on Autonomic and Autonomous Systems*, Athens, Greece, August 2007.

Forrester Research. 2008. Is cloud computing ready for the enterprise? *Forrester Research*, March 7, 2008.

Hayes, B. 2008. Cloud computing. *Communications of the ACM*, 51(7): 9–11, July 2008.

Mell, P. and Grance, T. 2009. Draft NIST working definition of cloud computing. http://csrc.nist.gov/groups/SNS/cloud-computing/index.html (accessed August 1, 2009).

Noël, F. and Delic, K.A. 2002. Knowledge-based self-support system for corporate users. *DEXA Workshops 2002*, Aix-en-Provence, France, pp 149–156.

The Economist. 2008. When clouds collide, February 7, 2008.

Weiss, A. 2007. Computing in the clouds. *netWorker*, 11(4): 16–25, December 2007.