Enterprise Knowledge Clouds: Next Generation KM Systems?

Kemal A. Delic, Jeff A. Riley
Hewlett-Packard Co.
kemal.delic@hp.com, jeff.riley@hp.com

Abstract

We are witnessing the emergence of the global, dependable and efficient infrastructure of cloud computing. We assess the current state of the Enterprise Knowledge Management field and project the possible emergence of Enterprise Knowledge Clouds. We give some architectural views, discuss briefly the underlying technologies and describe roughly related applications. We conclude with possible developments in the next five to fifteen years.

1. Introduction

The field of Knowledge Management (KM) has passed several cycles of hype and disappointments and has created important disputes along the axis of 'knowledge' as the philosophical discourse and science and 'management' as 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 phrase 'knowledge management' remains undefined, fuzzy and disputed, while 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 kind-of 'knowledge bazaar''. In this paper we claim that the treatment of KM Systems (KMS), which represent an intricate part of many business enterprises, has yet-another chance of re-appearing in totally new technological and social circumstances.

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

After observing current deficiencies and projecting future developments, we depict a high level architecture of the Enterprise Knowledge Cloud (EKC) as a collaborative, cooperating, competing mega-structure providing computing, networking and storage services to various 'knowledge producers and consumers' – being devices, people and applications. Some architectural and design landscapes are provided for illustration. We conclude with a no-nonsense list of things we should observe happening as the sign of mega-shift from the industrial to post-industrial world of the 21st century.

1.1 Emerging Cloud Computing Infrastructures

Cloud computing denotes infrastructure which already exists today in the form of the Internet providing a wide variety of IT services. They could be bought on-demand (computing cycles, storage, and network) in a highly simplified procedure. However, we predict future growth in which we will see a huge number of common devices inter-connected, 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: as an example, a well-known social network has provided a platform for 4000 applications written by 80,000 developers in just 6 months. We observe that the cloud infrastructure is global, highly dependable and supports innovative business models and new types of social phenomena.

978-0-7695-3531-9/09 $25.00 © 2009 IEEE
DOI 10.1109/eKNOW.2009.28

1.2 Collective Intelligence

By some estimates, today there are 80 million people worldwide writing blogs. The blogs are typically topic-oriented, and some attract important readership. Authors range from large company CEO's to housewives and young children. When taken together, the cloud computing infrastructure which hosts blogospheres looks like a big social agglomeration providing a kind of "collective intelligence". It appears as an omnipresent, omniscient, cloud-like giant infrastructure as a new form of 'knowledge management'. Today this represents a massive collaboration of people only, but tomorrow we may envisage intelligent virtual objects and devices collaborating with people. 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.

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 non-physical operations versus Internet, we observe a new type of enterprise emerging: we call it the Intelligent Enterprise [1]. It is able to interact with its environment, change its behavior, structure and strategy - behaving actually as an intelligent entity. It is able to adapt to changing market circumstances, gradually change its business model and survive into the next market cycle. We postulate that the emergence of collective intelligences in the cloud computing infrastructure will influence markets and established businesses, and reshape the contemporary approach to 'Enterprise Knowledge Management'. Next, we further describe the current state of EKM.

2. Enterprise Knowledge Management: Architecture & Technologies

Constantly evolving markets exercise pressure on business enterprises to continually evolve and improve. One of the most widely used business paradigms is about Enterprise Knowledge Management – as a means to capture and express tacit human knowledge into explicit form (externalized knowledge or content) which could be later (hopefully) reused. Various schools of thought were proposed, several assistive technologies developed and an important number of successful enterprise KM stories reported. From our experience, the best domain for EKM would be in the enterprise IT domain [2][3][4]– 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, decision making, business intelligence & reporting and monitoring/tuning/automation tasks (Fig. 1). These are also the most promising areas for the future deployment of Enterprise Knowledge Clouds.

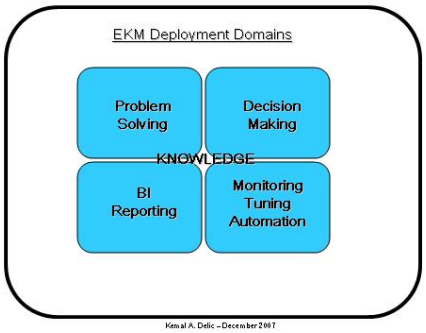

Fig. 1. Enterprise Knowledge Management: IT Deployment Domains

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 [5].

When mapped into more precise form, this conceptual drawing will evolve into the enterprise-scale KM application stack (Fig. 2).

Practically each and every KM application today can be layered into three essential subsystems:

- front-end portals that manage interactions with internal users, partner's agents and external users, while rendering 'knowledge services'.
- a core layer that provides the knowledge base and access/management services to knowledge portals and other enterprise applications.
- the back-end that supplies 'knowledge content' from various sources, authors and communities, enabling a refresh of the knowledge base.

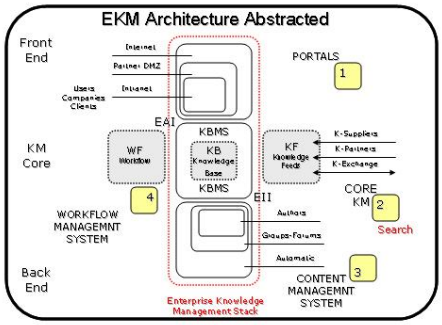

Fig. 2. Enterprise Knowledge Management : Architectural View

The Enterprise Workflow System captures interactions with users and provides necessary context for the EKM system. Various feeds enable 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 KM 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 3-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), and 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 which 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 [10].

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, non-stop application and would be supported 24x7 and always available. Some less critical applications will have yet another label, be supported 24x5 and non considered non-stop 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.

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 necessary features of a mission-critical application. It is out of scope of this paper 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.

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 would assume that multimedia content will become dominant in the future and that new methods for knowledge capture and rendering would be devised.

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. Thus 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 of future users of EKM systems.

3. Enterprise Knowledge Cloud (EKC)

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.

An abstracted business enterprise architecture is shown in Fig. 3. This architecture interconnects business partners and suppliers to company customers and consumers, and uses future cloud technologies to harvest, process and use internal knowledge (Corp 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 the wide variety of purposes.

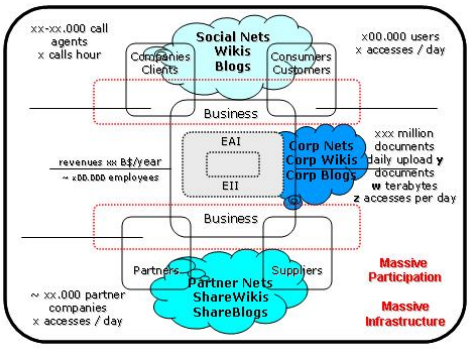

Fig. 3. Enterprise Knowledge Cloud : Architectural View

Each of the clouds shown in Fig. 3 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 enterprise.

The emergence of these clouds (Private, Partner, Public) and their coalescence into the Enterprise Knowledge Cloud 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. The interaction and cooperation of the user groups, their knowledge, and the collective intelligences across the three clouds shown in Fig. 3 provides the infrastructure for behavioral, structural and strategic adaptation in response to changes in the (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 – the Knowledge Exchange for example, enabling trade, exchange and monetizing of knowledge assets. All this is in a very early stage, but one should sense that developments may go in this direction.

4. Conclusion: The Next Five to Fifteen 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' [6]. 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 moving to decoupled, completely pluggable, intelligent KM appliances capable of adapting to interface with Enterprise Knowledge Clouds as required (Table 1).

Table 1. Evolution of EKM Systems

EKM Systems	Today	Tomorrow	Beyond
Architecture	Enterprise Stack	Distributed	Decoupled/Pluggable
Infrastructure	DC	Grid	Cloud
Application	IT Controlled	User Produced	Mix of X ?
Content	Mainly Text	Multimedia	Active
Users	Fixed/Mobile	Virtual	Avatars
Standards	3W.org	Web 2.0	Web 3.0

We are in the midst of important social, technological and market changes where we see some major companies announcing their intention to enter, drive and dominate the field of cloud computing [7][8][9]. We see this as a precondition for the emergence of the intelligent, adaptive enterprise which 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 Enterprise Knowledge Clouds embedded in the global, dependable fabrics consisting of subjects, objects and devices. 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 Enterprise Knowledge Management will be rescaled by an order of magnitude and will spawn the creation of a new kind of KM system.

5. Acknowledgements

The authors would like to thank Jamie El Fattal for his thoughtful review comments.

6. References

[1] Kemal A. Delic and Umeshwar Dayal: The Rise of the Intelligent Enterprise, ACM Ubiquity, Issue 45 (December 31, 2002 - January 6, 2003)

[2] Kemal A. Delic, Umeshwar Dayal: Knowledge-Based Support Services: Monitoring and Adaptation. DEXA Workshop 2000: : 1097-1101

[3] Franck Noël, Kemal A. Delic: Knowledge-Based Self-Support System for Corporate Users. DEXA Workshops 2002: 149-156

[4] Kemal A. Delic, Laurent Douillet: Corporate IT Knowledge Workbench: Case Study. DEXA Workshops 2004: 494-497

[5] Kemal A. Delic , Jeff A. Riley, Yassine Faihe: Architecting Principles for Self-Managing Enterprise IT Systems, IEEE ICAS archive : Proceedings of the Third International Conference on Autonomic and Autonomous Systems, August 2007

[6] Mario Cannataro , Domenico Talia, The knowledge grid, CACM Volume 46 , Issue 1 (January 2003)

[7] Is Cloud Computing Ready For The Enterprise, Forrester Research, March 7, 2008

[8] Aaron Weiss, Computing in the clouds, netWorker, December 2007

[9] Brian Hayes, Cloud Computing, Communication of ACM, Vol. 51., No 7, July 2008. pp. 9-11

[10] The Economist: When Clouds Collide, February 7, 2008.