

Chapter 18

Enterprise Knowledge Clouds: *Applications and Solutions*

Jeff A. Riley and Kemal A. Delic

18.1 Introduction

With the evolution of cloud computing in recent times we have proposed *Enterprise Knowledge Clouds* (EKC) as the next generation Enterprise Knowledge Management systems (Delic & Riley, 2009).

The Enterprise Knowledge Cloud abstracted architecture is shown in Fig. 18.1. This architecture interconnects business partners and suppliers to company customers and consumers, and uses future cloud technologies to harvest, process and use internal knowledge. Each of the clouds shown in Fig. 18.1 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 are independent, they will be capable of interconnection, overlap and knowledge-sharing so that, for example, customers and consumers might have access to appropriate internal enterprise knowledge, or even partner/supplier knowledge through the Enterprise Knowledge Cloud.

The emergence of these clouds 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. As an example, internal IT operations will use private clouds, Sales and Marketing would operate on public clouds, while Outsourcing businesses 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 Fig. 18.1 provides both the infrastructure for behavioural, structural and strategic adaptation in response to change as well as an environment for knowledge creation and exchange.

In the following sections we discuss Enterprise Knowledge Management (EKM) applications and solutions in terms of cloud computing and the emergence of the

J.A. Riley (✉) and K.A. Delic
Hewlett-Packard Co., New York, NY, USA
e-mails: {jeff.riley; kemal.delic}@hp.com

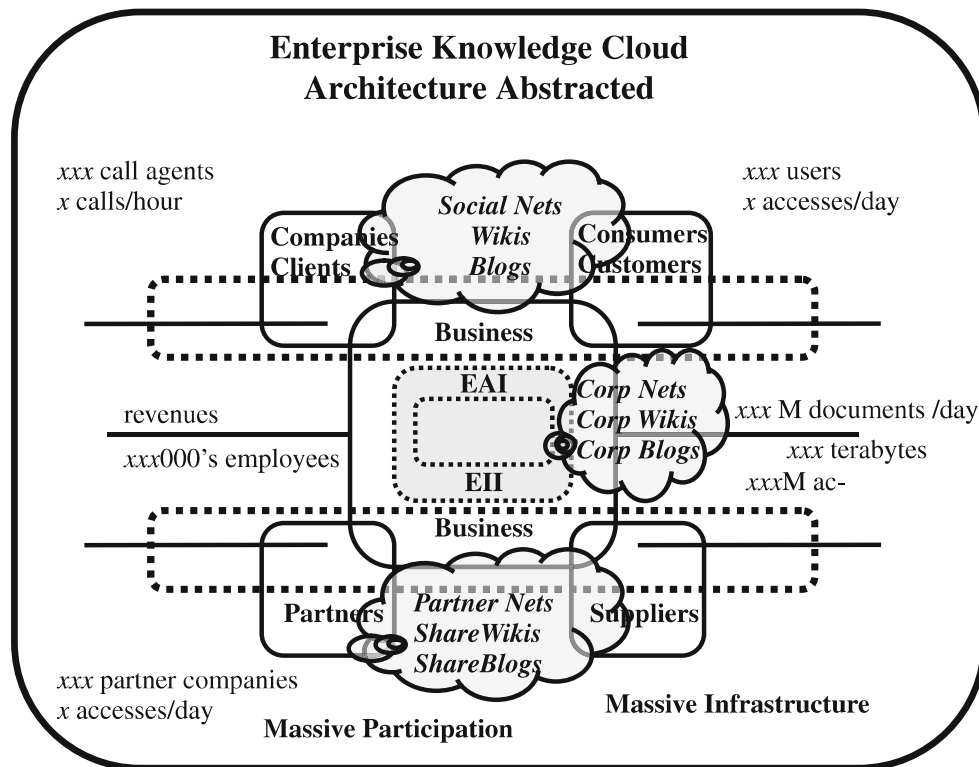


Fig. 18.1 Enterprise knowledge cloud: architectural view

Enterprise Knowledge Cloud. We present what we believe will be user expectations and requirements of cloud-based Knowledge Management Applications. We also discuss issues that will face developers and providers of KM applications as they migrate existing applications to the cloud platform, or build new applications specifically for the cloud. Finally we present our view of the future direction of Knowledge Management in a cloud computing environment.

18.2 Enterprise Knowledge Management

Unfortunately, even today there is no real consensus as to what “knowledge management” really is – ask ten people to define knowledge management and you will get twelve different answers. Worse than that – ask ten people to define “knowledge” and you will probably walk away with twice that number of definitions.

The “DIKW Hierarchy”, or “Knowledge Hierarchy”, refers to a representation of the relationships between *Data*, *Information*, *Knowledge*, and *Wisdom*. An early description of the DIKW Hierarchy was given by Milan Zeleny in 1987 (Zeleny, 1987)¹ in which he describes *Data* as “know-nothing”, *Information* as “know-what” and *Knowledge* as “know-how” (we are not concerned with Zeleny’s definition of *Wisdom* here and leave it to interested readers to refer to his paper).

¹The origin of the DIKW hierarchy is far from clear – see [12].

Later, Ackoff (1989) gave the following definitions for *Data*, *Information* and *Knowledge*:

- Data – data is raw. It simply exists and has no significance beyond its existence (in and of itself). It can exist in any form, usable or not. It does not have meaning of itself. In computer parlance, a spreadsheet generally starts out by holding data.
- Information – information is data that has been given meaning by way of relational connection. This “meaning” can be useful, but does not have to be. In computer parlance, a relational database makes information from the data stored within it.
- Knowledge – knowledge is the appropriate collection of information, such that its intent is to be useful. Knowledge is a deterministic process. When someone “memorizes” information (as less-aspiring test-bound students often do), then they have amassed knowledge. This knowledge has useful meaning to them, but it does not provide for, in and of itself, integration such as would infer further knowledge

We can see from these definitions that not all data is useful information, and not all information is useful knowledge. It follows then that Knowledge Management is not simply the management of data or information – just providing keyword searching capabilities over one or more information repositories is not Knowledge Management.

Referring to the definitions of knowledge proposed by Zeleny and Ackoff, we define Knowledge Management as the process by which organisations create, identify and extract, represent and store, and facilitate the use of knowledge and knowledge-based assets to which they have access.

18.2.1 EKM Applications

Knowledge can be categorized as either explicit or tacit. Explicit knowledge is knowledge that can be codified – knowledge that can be put in a form which is easily transferred from one person to another. Examples of explicit knowledge are patents, recipes, process documentation, operational plans, marketing strategies and other such knowledge-based hard assets. Tacit knowledge, on the other hand, is knowledge which hasn’t been codified and is in fact difficult to codify and transfer – knowledge that exists as “know-how” in people’s heads, often unknown or unrecognised. Examples of tacit knowledge are things such as habits, sometimes cultural and often learned by observation and imitation, and the notion of language.

Knowledge Management applications can be similarly categorised, and while traditional KM applications focused on managing explicit knowledge (knowledge base management systems, workflow management systems etc.), with the advent of *Web2.0*, *social networking* and *groupware*, *mashups*, *wikis*, *blogs* and *folksonomies*

etc., there is more opportunity for KM applications focused on managing tacit knowledge to evolve.

Typical Knowledge Management applications can be layered into three essential subsystems:

- Front-end portals that manage interactions with internal users, partners' 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 knowledge and services.
- A core layer that provides the knowledge base and access, navigation, guidance and 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.

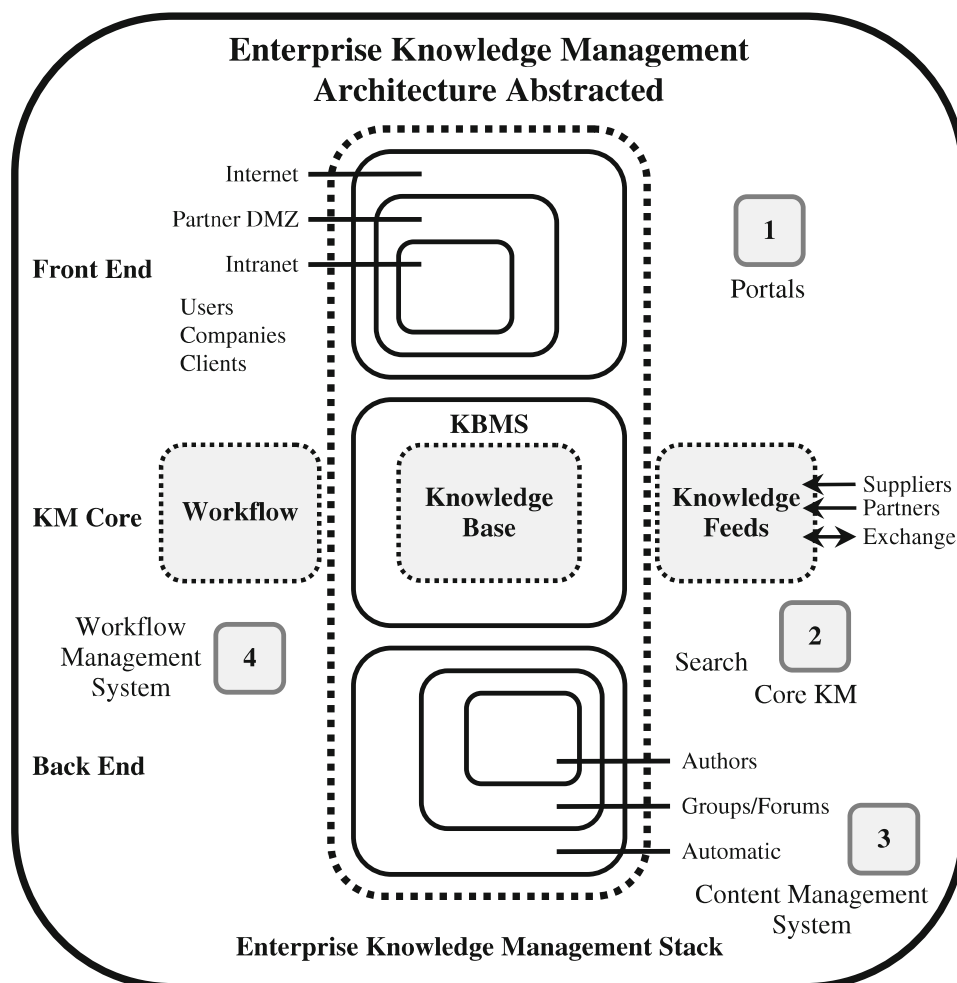


Fig. 18.2 Enterprise knowledge management: architectural view

- The back-end that supplies *Knowledge Content*, and the content management system, from various sources, authors and communities, enabling a refresh of the knowledge base.

Thus, Enterprise Knowledge Management is typically a three-tier enterprise application probably spread over several geographically dispersed data centres, and typically interconnected or integrated with enterprise portals, content and workflow management systems (Fig. 18.2).

The Enterprise Workflow System captures interactions with users and provides necessary context for the Enterprise Knowledge Management 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 knowledge management system, we presume that the knowledge base contains several million knowledge items, and users number in the hundreds of thousands. Enterprise knowledge management is considered a high-end, mission-critical corporate application which resides in the corporate data centre. High availability and dependability are necessary engineering features for such global, always-on, always-available systems. The transformation of EKM into the enterprise cloud will increase the scale and importance by orders of magnitude.

18.3 Knowledge Management in the Cloud

Knowledge Management in a cloud computing environment requires a paradigm shift, not just in technology and operational procedures and processes, but also in how providers and consumers of knowledge within the enterprise think about knowledge. The *knowledge-as-a-service*, “on-demand knowledge management” model provided by the cloud computing environment can enable several important shifts in the way knowledge is created, harvested, represented and consumed.

18.3.1 Knowledge Content

Collective intelligence is a phenomenon that emerges from the interaction – social, collaborative, competitive – of many individuals. By some estimates there are more than eighty million people worldwide writing web logs (“*blogs*”). The blogs are typically topic-oriented, and some attract important readership. Authors range from large company CEOs to administrative assistants and young children. When taken together, the cloud computing infrastructure which hosts “*blogospheres*” is 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, giant infrastructure – as a new form of knowledge management. This same paradigm applies, albeit on a smaller scale, to the enterprise cloud – the socialisation and participatory culture of today’s internet is mirrored in the microcosm of the enterprise.

Today this represents collaboration of mostly people only, but very soon in the future we may envisage intelligent virtual objects and devices collaborating with people. Indeed, 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 even hundreds of billions of real-world objects having a data representation in the virtual world is probably realistic.

It is important to note here that content will no longer be located almost solely in a central knowledge repository on a server in the enterprise data centre. Knowledge in the cloud is very much distributed throughout the cloud and not always residing in structured repositories with well-known query mechanisms. Knowledge management applications offered in the cloud need to be capable of crawling through the various structured and ad hoc repositories – some perhaps even transient – to find and extract or index knowledge, and that requires those applications be capable of recognising knowledge that might be useful to the enterprise knowledge consumers. Furthermore, we believe that over time multimedia content will become dominant over ordinary text, and that new methods for media-rich knowledge management will need to be devised.

Even in the smaller world of the enterprise, 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.

18.3.2 Knowledge Users

Leaving aside administrative tasks, there will be two categories of users within the Enterprise Knowledge Cloud: *Knowledge Providers* and *Knowledge Consumers*.

While sketching the architecture of future enterprise knowledge management applications, serious consideration needs to be given to aspects and dimensions of future users – the evolution of technology in consumer and corporate domains has created a new type of user that will be very different from contemporary knowledge consumers.

The younger generation – the so-called “*Generation Y*” or “*Millennial Generation*” – seems to have developed a way to quickly exchange information snippets, being either very short text messages or particular multimedia content. Members of this generation also typically have a much better ability to multitask naturally while not losing or intermixing communication threads – probably a natural consequence of their exposure to electronic gaming and new work and living

styles. This new generation of knowledge consumer will drive the on-demand nature of knowledge management in the cloud. Moreover, they will require their knowledge served up to them in “*BLATT*” format – with the “bottom line at the top” – and with media-rich content.

Knowledge providers will need to develop applications that recognise the salient knowledge in response to queries, and deliver a synopsis of that salient material ahead of the detail. Delivering a list of documents for the user to read through and identify and extract the knowledge themselves will no longer be acceptable. Furthermore, knowledge management applications in the cloud will need to be capable of presenting media-rich results to the consumer concurrently with more traditional text-based results. Knowledge consumers that have grown up with the internet know that knowledge is more than just text in a data base, and when they seek knowledge they will want to view the video as well as read the text. The new, internet-savvy knowledge consumers will want short, sharp, to-the-point responses to their queries – but responses which are complete and include audio and visual content where appropriate.

18.3.3 Enterprise IT

From our experience, the best domain for Enterprise Knowledge Management is in the Enterprise IT domain, as it is a domain under huge cost pressure but one which is essential for strategic development.

From a highly abstracted view, the Enterprise Knowledge Management IT domain consists of problem solving, monitoring, tuning and automation, business intelligence & reporting, and decision making tasks (Fig. 18.3).

The tasks of problem solving, monitoring, tuning and automation, business intelligence and reporting, and decision making are the most promising areas for the future deployment of Enterprise Knowledge Clouds. The knowledge available to both IT administrators and automated management agents via the Enterprise

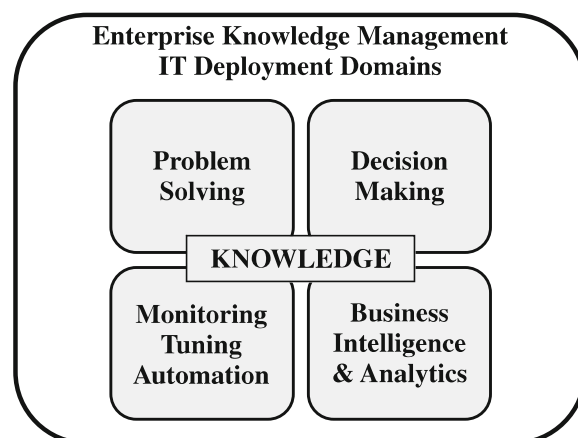


Fig. 18.3 Enterprise knowledge management: IT deployment domains

Knowledge Cloud will help drive the development of a slew of new technologies addressing the problems which previous computing facilities couldn't 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 & Faihe, 2007). When mapped into more precise form, this conceptual drawing (Fig. 18.3) will evolve into the enterprise-scale knowledge management application stack discussed earlier (Fig. 18.2).

18.3.3.1 Problem Solving

Problem solving, especially in the Enterprise Knowledge Management 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 artefacts, 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.

With the advent of the Enterprise Knowledge Cloud and associated interfaces to knowledge bases and knowledge feeds, software agent-based problem solving and guidance becomes much more of a realistic proposition. Intelligent software agents, able to be trained by “watching” and learning from experienced and expert human support engineers, can be deployed across the cloud to assist less experienced, sometimes novice, engineers and end users. These intelligent software agents will have knowledge of and access to all knowledge repositories and knowledge feeds across the Enterprise Knowledge Cloud.

The potential for collaborative problem solving is expanded in the Enterprise Knowledge Cloud, with the social networking aspect of the cloud environment facilitating both greater interaction between the end user and support engineer and amongst support engineers and others with relevant knowledge – including software agents, and eventually intelligent software agents collaborating with other software agents.

The Enterprise Knowledge Cloud provides a platform not only for collaborative problem solving, but also for distributed problem solving. Distributed problem solving is not new, and there are established grid middleware and resource brokering technologies that facilitate the development of distributed applications (e.g. Foster & Kesselman, 1999; Frey, Tannenbaum, Livny, Foster, & Tuecke, 2004; Venugopal, Buyya, & Winton, 2006), but more recently distributed problem solving applications developed specifically for the Enterprise Cloud have been described. In Vecchiola, Kirley, and Buyya (2009), Vecchiola, Kirley and Buyya describe a network-based, multi-objective evolutionary algorithm to be deployed on an Enterprise Cloud for solving large and complex optimisation problems – for example, the tuning and management of computer systems and networks.

18.3.3.2 Monitoring, Tuning and Automation

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.

Delic et al. describe a system for the self-management of Apache Web Servers using a hierarchical control system well-suited to implementation on an Enterprise Knowledge Cloud platform (Delic, Riley, Bartolini, & Salihbegovic, 2007a). Their hybrid approach harvests and combines knowledge from different sources and utilises machine learning techniques to learn and implement heuristics to manage the web server (Fig. 18.4).

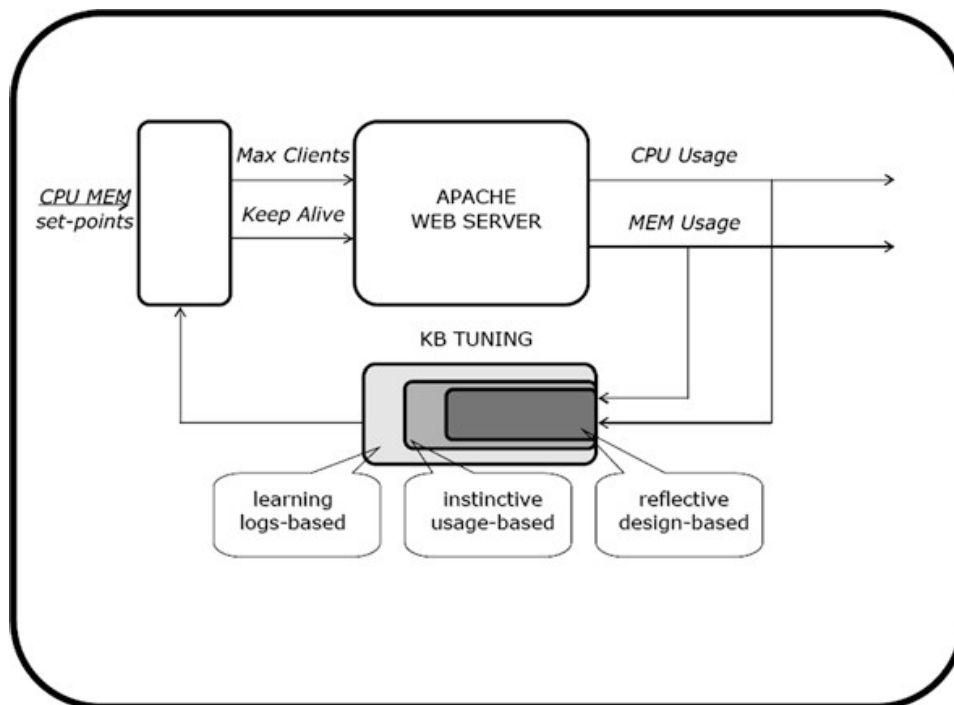


Fig. 18.4 Control system for self-management of apache web servers

The self-managing Enterprise IT system is becoming more achievable as knowledge management systems become smarter and knowledge more accessible. The automation of system monitoring, problem recognition, analysis and, when necessary, routing to a human expert is a critical step in architecting self-managing systems. Figure 18.5 shows an architecture for a state-based, adaptive, self-managing Enterprise IT system (Delic, Riley, & Faihe, 2007b). The system shown in Fig. 18.5 recognises the system state and applies solutions to known-problem where applicable, or routes the problem to an expert to solve as necessary. Problems routed to experts are solved, and knowledge of the solutions is fed back into the system to allow that knowledge to be applied as “known-problem solutions” in the future.

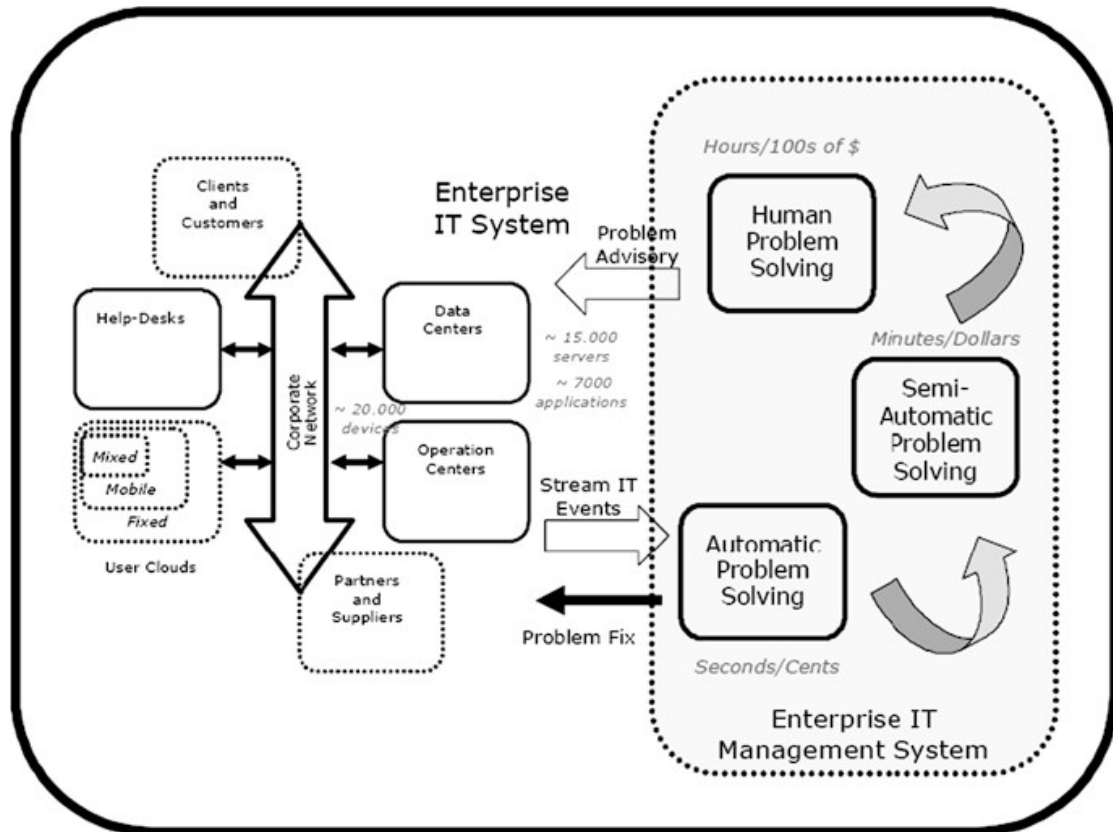


Fig. 18.5 Self-managed enterprise IT system

System state and known-problem recognition are learned in an “offline” mode by the system – facilitated by the integrated knowledge management systems of the Enterprise Knowledge Cloud in a way that has not been possible, or is vastly more difficult, in a non-cloud environment.

18.3.3.3 Business Intelligence and Analytics

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 – BI reporting uses a wide range of AI techniques to extract relevant knowledge from incoming data streams and repositories and provide observations, hints and suggestions about trends and possible futures.

Enterprise analytic systems provide different classes of users a wide range of analytics from holistic views of the enterprise state and subsystems to the tasks of optimization and forecasting (Delic & Dayal, 2003). These systems cover different

domains, address varying value-at-risk entities, and require different analytic artefacts to be provided to users in order to improve and accelerate decisions. Highly abstracted, the enterprise architecture can be dissected into event, transaction and analytic layers, each having different purposes, objectives and design constraints (Fig. 18.6).

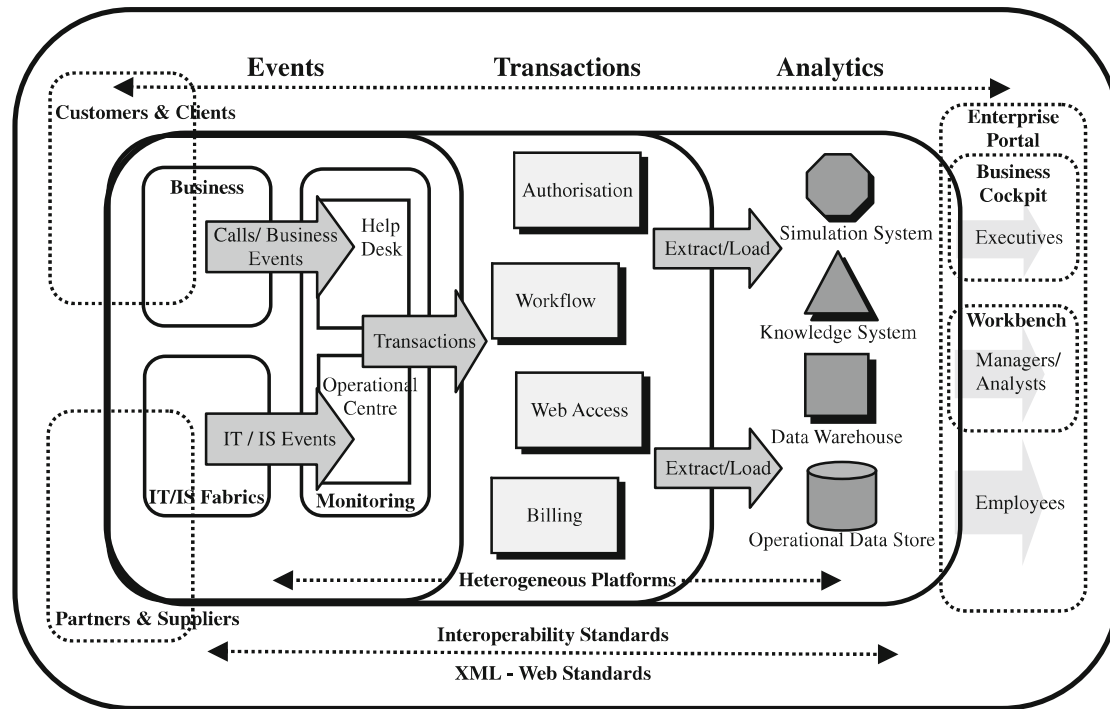


Fig. 18.6 Abstracted enterprise architecture

Different types of users of enterprise analytics use different artefacts: senior executives and analysts typically require interactive/active content, whereas passive/static content is used by larger communities of employees. Each generic type of user also exhibits different usage patterns: analysts typically need a powerful client/server setup; managers need browser access to analytics from a portable system; lower-level employees might need analytics right on their desktops or accessible via temporarily connected PDAs, mobile phones, or other portable devices.

The Enterprise Knowledge Cloud brings together the technologies and knowledge repositories underpinning today's business intelligence and analytics systems.

18.3.3.4 Decision Making

Decision making is most often done by humans after consuming the results of the business intelligence and analytics 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. For this reason more responsibility for the decision making task is being given to

knowledge-driven Decision Support Systems (DSS) which typically employ some sort of artificial intelligence tuned to the environment of the DSS deployment.

Delic et al. describe a real-time Decision Support System for a typical IT situation: help-desk support for enterprise customers (Delic, Douillet, & Dayal, 2001). Such help-desk support deals with IT equipment usage problems (usually desktop workstations), provides operations centre support to manage the network, servers and applications. Figure 18.7 depicts a help-desk system which uses a workflow system to capture “*problem cases*” (phone call records) and “*trouble tickets*” (events initiated from the equipment or/and applications). Two kinds of knowledge bases support IT help-desk operations: case knowledge bases containing problem-solving knowledge (often in the form of documents and search tools, diagnostic tools such as Bayesian networks, case-based reasoning tools, etc.); and real-time knowledge bases containing events, event management knowledge (e.g., *event-condition-action* rules) and enterprise topology.

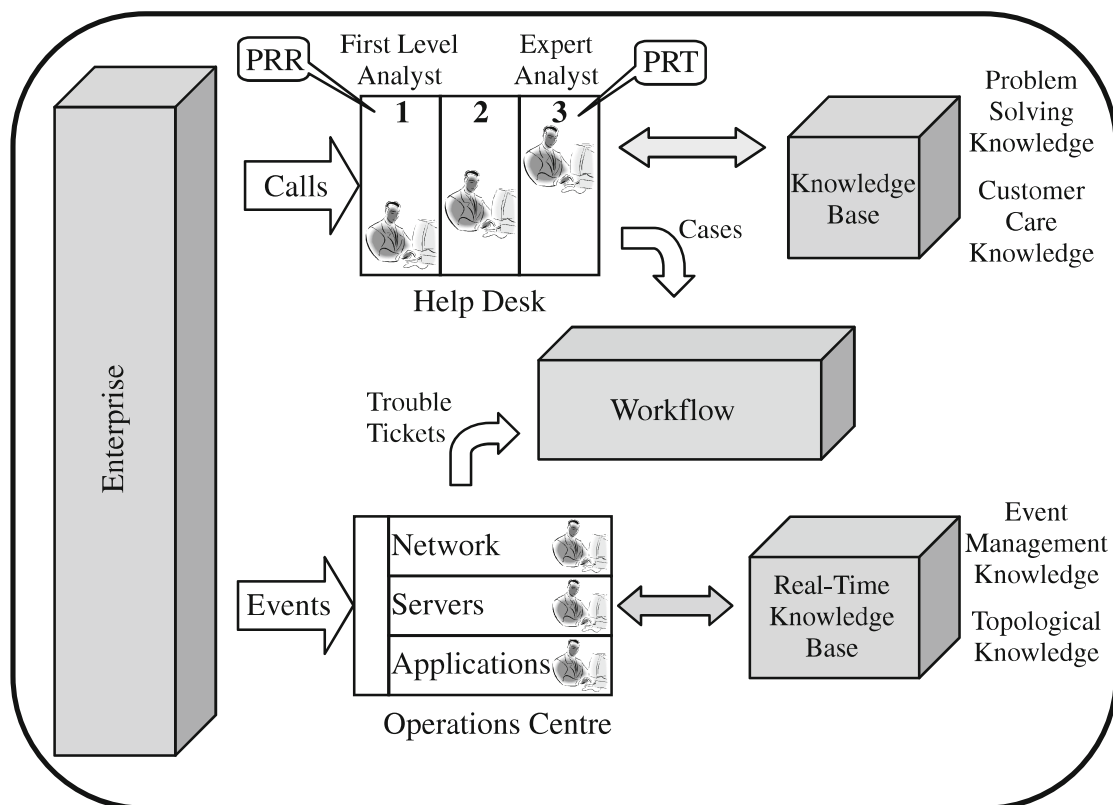


Fig. 18.7 Real-time decision support system for IT help desk

A help-desk is usually staffed by first call contact analysts who use the knowledge bases and tools available to solve less complex problems, and dispatch the more complex problems to second and third line analysts for resolution.

A manager responsible for IT help-desk support operations typically wants to see a variety of performance metrics, including Problem Resolution Rate [PRR] for the front-line analysts, and Problem Resolution Time [PRT] for the second and third line of support. Traditionally the manager would typically get monthly or weekly

reports from the workflow system, but with more current information the manager could react more quickly to changes in operating conditions. For example, if the help-desk manager was able to observe *in real time* that PRR is going down and PRT is going up, he or she could decide quickly what corrective action to take. The causes of the problem could be varied and interdependent, masking the prevailing cause or disabling proper diagnosis and decision-making. The manager may need additional information to determine the actual causes, so a real-time Decision Support System is critical in this case.

The Enterprise Knowledge Cloud, with its integrated knowledge repositories and Knowledge Management Systems is an ideal environment for a knowledge-based DSS. A knowledge-driven Decision Support System can be viewed in general terms as the combination of a problem solving system and business intelligence and analytics reporting. The knowledge-driven DSS applies specialist problem solving capabilities to the analytics reported by the enterprise analytics and business intelligence systems and recommends decisions and actions to users of the system. In some cases more advanced and intelligent systems will actually perform recommended actions or implement recommended decisions.

18.3.4 The 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 the internet, we observe a new type of enterprise emerging: the *Intelligent Enterprise* (Delic & Dayal, 2010).

The Intelligent Enterprise is able to interact with its environment, change its behaviour, 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 intelligences 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.

18.4 Moving KM Applications to the Cloud

Not all applications are suited to cloud computing, so it follows that not all Knowledge Management applications will be good candidates for migration to a cloud platform. A number of issues need to be considered when deciding whether any application should run on a cloud platform, all of which apply to KM applications.

Some considerations when either migrating applications to a cloud environment or designing for the cloud are:

- *Security and Privacy*: while less of a problem in an enterprise cloud than a public cloud or the open internet, security and privacy issues need to be considered when migrating KM applications to the cloud, or when designing new KM applications for the cloud.
- *Latency*: if the speed with which knowledge queries are answered is critical, then careful consideration should be made as to whether the knowledge application serving those queries should be moved to, or developed on, the enterprise cloud. Individual business, even within larger enterprise business, may be able to provide faster access to dedicated knowledge bases than is possible in a cloud environment.
- *Transaction Management*: data integrity, concurrency and latency are all issues for transaction management in the cloud. There are many strategies for coping with issues in this area – the right ones need to be chosen for each application.
- *Criticality and Availability*: the criticality of the underlying cloud infrastructure should match the criticality of the applications to be run on that infrastructure. Mission critical KM applications, for example, should not be deployed in a non-mission critical cloud environment.

18.5 Conclusions and Future Directions

Cloud technology is driving a new paradigm in knowledge management. The participatory and collaborative nature of the internet and cloud computing is both creating more knowledge and providing access to knowledge that was hitherto not generally accessible. There is much more data and information to sort through to find the gems of knowledge – some of the data and knowledge is transient, and some of it not recognisable as knowledge. New technologies will be developed to cope with an almost overwhelming volume of data, information and knowledge.

Knowledge content today is mostly text-based, but for the future we see an evolution towards multimedia and active content. 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. And while today’s enterprise applications are developed by IT departments, with users becoming more tech and internet savvy we predict a shift towards user-developed applications: mash-ups written in high-level mash-up languages.

Current Enterprise Knowledge Management systems are enterprise applications in data centres, while we expect them to evolve into “Enterprise Grids” on which others envisage the development of “Knowledge Management Grids” (Cannataro & 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 centres providing various functionalities.

Thus, while the architecture of today’s Enterprise Knowledge Management systems is built around the enterprise stack, tomorrow’s Enterprise Knowledge Management architecture will be distributed and loosely-coupled, and later moving to decoupled, completely pluggable, intelligent knowledge management appliances capable of adapting to interface with Enterprise Knowledge Clouds as required (Table 18.1).

Table 18.1 Evolution of EKM systems

EKM systems	Today	Tomorrow	Beyond
Architecture	Enterprise Stack	Distributed	Decoupled/Pluggable
Infrastructure	Datacentre	Grid	Cloud
Application	IT Controlled	User Produced	On Demand
Content	Mainly Text	Multimedia	Active
Users	Fixed/Mobile	Virtual	Avatars
Standards	3 W.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 (Hayes, 2008; Staten, 2008; Weiss, 1987). 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. 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.

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 centres). 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 EKM system – the *Enterprise Knowledge Exchange*, enabling trade, exchange and monetisation of knowledge assets.

References

- Ackoff, R. L. (1989). From data to wisdom. *Journal of Applied Systems Analysis*, 16, 3–9.
- Cannataro, M., & Talia, D. (2003). The knowledge grid. *Communications of the ACM*, 46(1), 89–93.
- Delic, K. A., & Dayal, U. (2003). A new analytic perspective. *Intelligent Enterprise Magazine*, 6(11). Available from: http://intelligent-enterprise.informationweek.com/030630/611feat2_1.jhtml.
- Delic, K. A., & Dayal, U. (2010). The rise of the intelligent enterprise. *ACM Ubiquity*, (45). Retrieved February 1, 2010, from <http://www.acm.org/ubiquity>.
- Delic, K. A., & Riley, J. (2009). Enterprise knowledge clouds: Next generation KM systems? *Proceedings of the 2009 International Conference Information, Process, and Knowledge Management (eKnow 2009), Cancun, Mexico*, 49–53.
- Delic, K. A., Douillet, L., & Dayal, U. (2001). Towards an architecture for real-time decision support systems: Challenges and solutions. *Proceedings of the 2001 International Database Engineering & Applications Symposium (IDEAS'01), Grenoble, France*, 303–311.
- Delic, K. A., Riley, J., Bartolini, C., & Salihbegovic, A. (October 2007a). Knowledge-based self-management of apache web servers. *Proceedings of the 21st International Symposium on Information, Communication and Automation Technologies (ICAT'07), Sarajevo, Bosnia and Herzegovina*.
- Delic, K. A., Riley, J., & Faihe, Y. (June 2007b). Architecting principles for self-managing enterprise IT systems. *Proceedings of the 3rd International Conference on Autonomic and Autonomous Systems, Athens, Greece*, 60–65.
- Foster, I., & Kesselman, C. (1999). Globus: A toolkit-based grid architecture. In I. Foster, & C. Kesselman (Eds.), *The grid: Blueprint for a new computing infrastructure* (pp. 259–278). Los Altos, CA: Morgan Kaufmann.
- Frey, J., Tannenbaum, T., Livny, M., Foster, I., & Tuecke, S. (2004). Condor-G: A computation management agent for multi-institutional grids. *Cluster Computing*, 5(3), 237–246.
- Hayes, B. (2008). Cloud computing. *Communications of the ACM*, 51(7), 9–11.
- Staten, J. (March 2008). Is cloud computing ready for the enterprise?. *Forrester Research*, March 7.
- Vecchiola, C., Kirley, M., & Buyya, R. (2009). Multi-objective problem solving with offspring on enterprise clouds. *Proceedings of the Tenth International Conference on High-Performance Computing in Asia-Pacific Region (HPC Asia 2009), Hongkong*.
- Venugopal, S., Buyya, R., & Winton, L. (2006). A grid service broker for scheduling e-science applications on global data grids. *Concurrency and Computation: Practice and Experience*, 18(6), 685–699.
- Weiss, A. (2007). Computing in the clouds. *netWorker*, 11(4), 16–25.
- Zeleny, M. (1987). Management support systems: Towards integrated knowledge management. *Human Systems Management*, 7(1), 59–70.