

Chapter 10

Enterprise Knowledge Clouds: *Architecture and Technologies*

Kemal A. Delic and Jeff A. Riley

10.1 Introduction

This chapter outlines the architectural foundations of *Enterprise Knowledge Clouds* (EKC) (Delic & Riley, 2009), describing the underlying technological fabrics and then pointing at the key capabilities of the (hypothetical) intelligent enterprise operating in constantly evolving, dynamic market conditions. Our aim is to give readers of this chapter a better understanding of knowledge cloud architectural aims and practical insights into EKC technological components. Thanks to knowledge, the enterprise will know more, will act better and react sooner in changing environment conditions, ultimately improving its performance and enabling it to show better behaviour and measurable improvements.

The *Enterprise* is an organisational structure which may take varying forms in different domains and circumstances. For our purposes here, we consider that the enterprise is an operating business employing 5000 or more people, operating globally with revenues in excess of \$1 Billion, and supported by appropriate IT capabilities and facilities. There is a long, ongoing debate over the value and impact of IT use in business operations, but we can easily imagine what would happen if a business enterprise suddenly finds itself without any IT systems.

Knowledge gives distinctive capabilities to living creatures, with humans being at the top of the hierarchical tree of life. Tacit knowledge enables perception, reflection and action as the basic features of any intelligent behaviour. Technology, on the other hand, enables capturing and reuse of tacit knowledge in explicit form. Much of what we know as ‘knowledge management’ is about transforming tacit knowledge into explicit and vice-versa. Intuitively is clear that knowledge plays the key role in each and every part of the business enterprise. Knowledge takes various forms, has variable value and makes varying impact, and requires different technologies to

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

deal with its entire, continues life cycle. It is ‘enterprise knowledge’ which makes the difference in operational tasks (automation) and strategic situations (decision making).

Cloud computing is an emerging architectural paradigm driven by the sharp drop of technology costs followed by radically improved performance (commoditisation) (ITU, 2010). Social changes and economic advances have created a huge number of consumers and producers of various content artefacts (text, photo, music, video, etc.) representing huge user clouds and large communities (in the order of 100s of millions people). Thus we see cloud computing developing on an unprecedented scale and dynamics on a global basis (CLOUDSCAPE, 2009).

Highly abstracted, company operations are described as the interplay between people, machines and processes, providing either tangible goods or consumable services (Delic, 2003). Depending upon the business context, one component might be dominant over others, while each will contain something which we could label as ‘knowledge’. It is important to observe is that approximately 75% of the economic activity in most advanced countries is created by service industries where knowledge is the primary resource or ingredient, thus we have been hearing about ‘knowledge-based economies’ for many years now. We may conclude that the services economy is driven by the power of knowledge.

We postulate that *Knowledge Clouds* (KC) will enable the global spread of economic growth, efficient delivery of services, smoother exchange and profitable trade of goods and services.

10.2 Business Enterprise Organisation

The typical business enterprise is an hierarchical organisation which has certain characteristics of military command-control layout with executives at the top of the hierarchy (numbered in tens) senior managers and managers at the next level (numbered in hundreds) and employees at the base (numbered in thousands). Depending upon the industry branch and regional specifics, it might be that certain functions are global and others regional. This usually leads to a characteristics matrix organisation which is marked by high complexity.

At the conceptual level, we can talk about key entities as Clients and Customers, Partners and Suppliers, interconnected to the business enterprise via distribution channels and supply chains (Fig. 10.1). Internally, the enterprise will have shared functions such as Human Resources, Finance, R&D Labs, IT, Sales and Marketing synchronised with Production and Services. Specialised enterprise software (e.g. CRM, ERP, SCM etc.) enables smooth operations of the enterprise and represent typical enterprise software applications today. For large enterprises, those are key systems requiring many years to perfect and require a large effort to operate. Each system encompasses knowledge embodied either as human expertise, business processes, software algorithms or analytical models. Enterprise IT plays a special, technological role for which KM will have distinctive value and lasting importance.

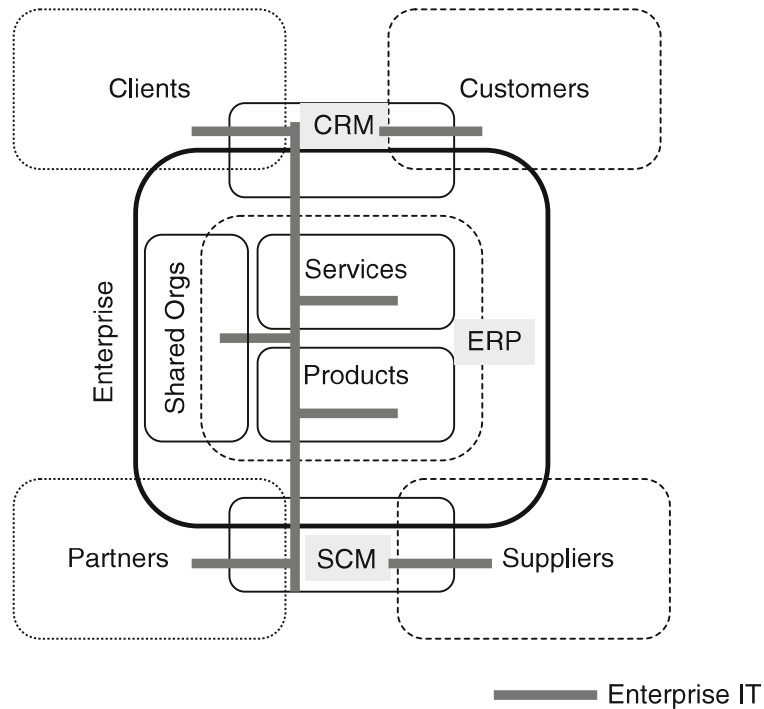


Fig. 10.1 Conceptual business enterprise organisation

From the perspective of technology components, Enterprise IT can be abstracted with the key components having their own operational indicators, such as dollars-per-call for the help-desk or cents-per-event for processing, enabling management and administrators to grasp inefficiencies and estimate the overall cost (Fig. 10.2). The ultimate objective is to minimise the cost while maximising the efficiency of the each IT unit, considering that data centres are machine intensives, help-desks are

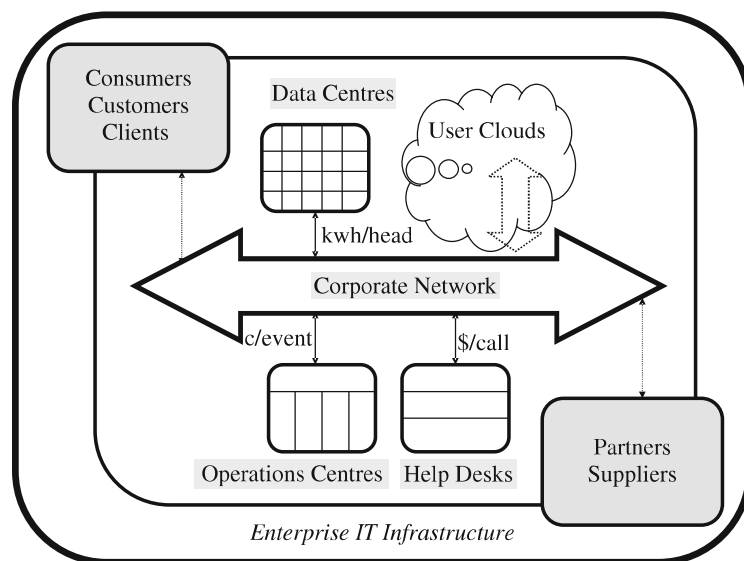


Fig. 10.2 Conceptual enterprise IT organisation

labour intensive, and operational centres and corporate networks are event intensive units.

To enable synchronised and orchestrated changes, enterprise architecture captures the overall state of the enterprise business and IT infrastructure and provides guidance expressed with a set of architecting principles.

Enterprise infrastructure used for business operations represents interconnecting, mediation fabrics which improve operational behaviour captured and indicated via key performance indicators. Removal of the IT fabrics will cripple business operations, demonstrating that today's business operations are not imaginable without the deployment of IT. In fact, the majority of businesses consider IT as a business-critical component which must be cleverly architected and well designed as a highly dependent part of the business.

10.3 Enterprise Architecture

Enterprise architecture is a strategic framework that captures the current state of the enterprise business and supporting IT, and outlines an evolutionary path towards the future state of the business and IT (Delic, 2002b). It is a very hard challenge to provide synchronised development of business and IT in dynamic and unpredictable market conditions. Thus, having a sound enterprise architecture that charts evolutionary change over 3–5 years is an important, competitive advantage. In simplistic terms, enterprise architecture is a model depicting the evolution of business, infrastructure, applications and data landscapes over 3–5 years. Each artefact of the enterprise architecture has a very high monetary value, a strong proprietary nature, and a vital importance for the future of the enterprise.

In a simplistic fashion, these artefacts can be represented in an hierarchical manner, implying the type of models appropriate for each layer, characteristic entities and key metrics (Fig. 10.3). Enterprise architecture can be viewed as a global strategic plan which will synchronise business evolution with IT development and ensure that future needs are properly addressed. As such, enterprise architecture represents the most valuable strategic planning item for enterprise executives and management: business can plan underlying technology changes after observing forthcoming technology shifts (Zachman, 1999). Such technology changes will critically improve or cripple business performance, and mature businesses should keep those plans private, current and sound.

The expanded landscape shown in Fig. 10.4 indicates some characteristic operational figures for the very large, global business enterprise, and illustrates the scope, scale and complexity of operations (Delic, 2002a). It is intuitively clear that such a complex environment contains several points of inefficiency and structural weakness which could be best dealt with via deployment of KM techniques. A more developed enterprise landscape (e.g. Delic, Greene, & Kirsch, 2010) will also contain data points, indicate the dynamics and spread of data flows, and show the key technological, business and market indicators.

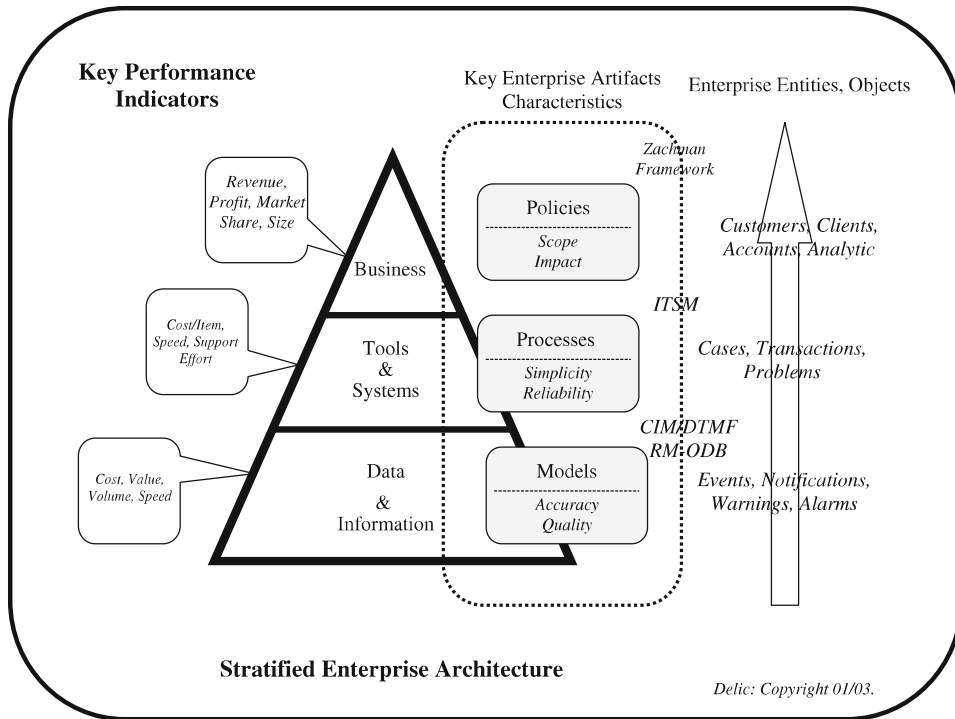


Fig. 10.3 Simplified enterprise/business/IT artefacts

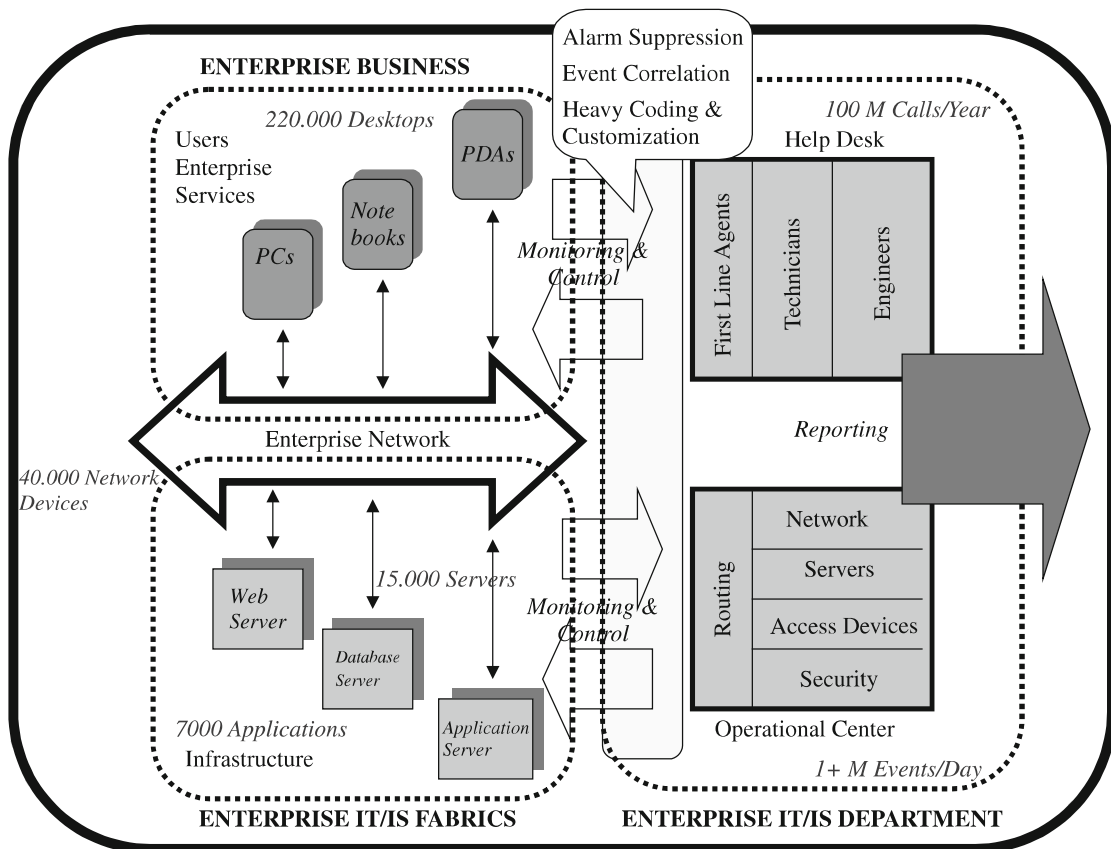


Fig. 10.4 Simplified enterprise/business/IT architecture landscape

As an example, for an enterprise business to efficiently handle in the order of 100 million calls per year, it should capture and deploy knowledge about callers, encountered problems and solution procedures. Similarly, to configure, manage and maintain in the order of 40,000 network devices, an enterprise business needs deep and reliable knowledge about network topology, fault behaviours and overall traffic flows. For each domain, knowledge will have different capture paradigms and technology, and will have a different impact on internal IT performances and cost.

10.4 Enterprise Knowledge Management

Enterprises as large, distributed and complex entities have several points of inter-operations with their environment which could be improved via deployment of KM applications. Knowledge about clients and customers will improve financial results and customer satisfaction. Knowing partners and suppliers better will help improve cooperation. Internal systems may help harvest employees ideas which then might be transformed into valuable intellectual property. As previously indicated, the IT domain is especially suitable for the deployment of KM systems, and this is an area in which the authors both have many years experience. We describe in detail three examples of KM deployment for internal IT operations, decision support, and knowledge harvesting.

An example of KM deployment in the IT domain is in the use of various knowledge repositories and systems to resolve a range of IT problems (Delic et al., 2010). Following a problem event from the IT infrastructure we see from Fig. 10.5 that

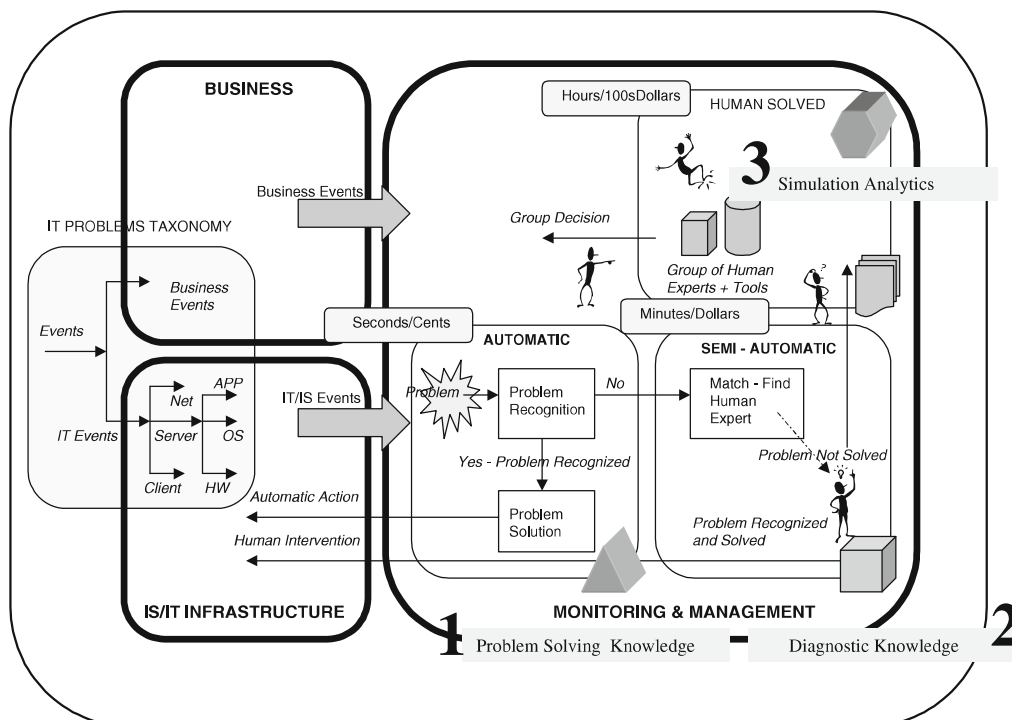


Fig. 10.5 Knowledge for IT problem solving

(1) problem recognition software will search a knowledge base containing problem solving knowledge, and if it recognises and identifies the problem it will deploy the solution found in the knowledge base; otherwise, (2) a knowledgeable human expert will be identified and, after deploying the diagnostic procedure, knowledge will be forwarded into automatic problem solving layer. For more complex, intricate or inter-dependent problems, (3) a group of human experts will be engaged to use knowledge captured in simulation analytics to resolve the problem via group decision making.

As millions of problems are solved daily, it is clear that the cost and speed of problem resolution are important parameters that illustrate the value of KM deployment for IT operations. At the very high level of abstraction, we see transformation of the raw data into information and then into knowledge and problem-solving acts, having measurable business impact and monetary value. It is important to note that KM techniques serve important roles in support and services, and that technologies deployed mainly originated from the field of Artificial Intelligence.

Another example of knowledge deployment is in decision support systems for enterprise operations, based on Enterprise Management Analytics (Casati et al., 2004; Cerri et al., 2008). We depict a layered IT architecture serving the business to orchestrate operations with clients and customers while being supported by suppliers and partners (Fig. 10.6). Those layers have distinctive architectures dictated by the general intent, so that all events from the instrumentation layer are served in

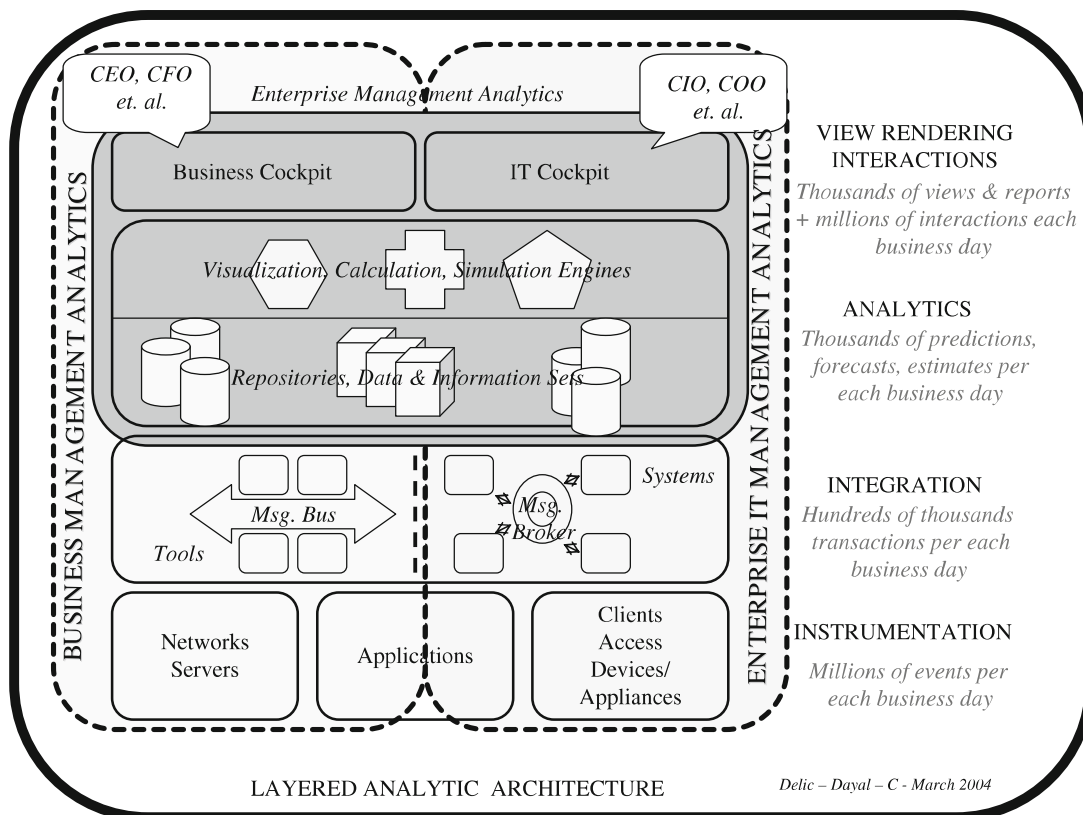
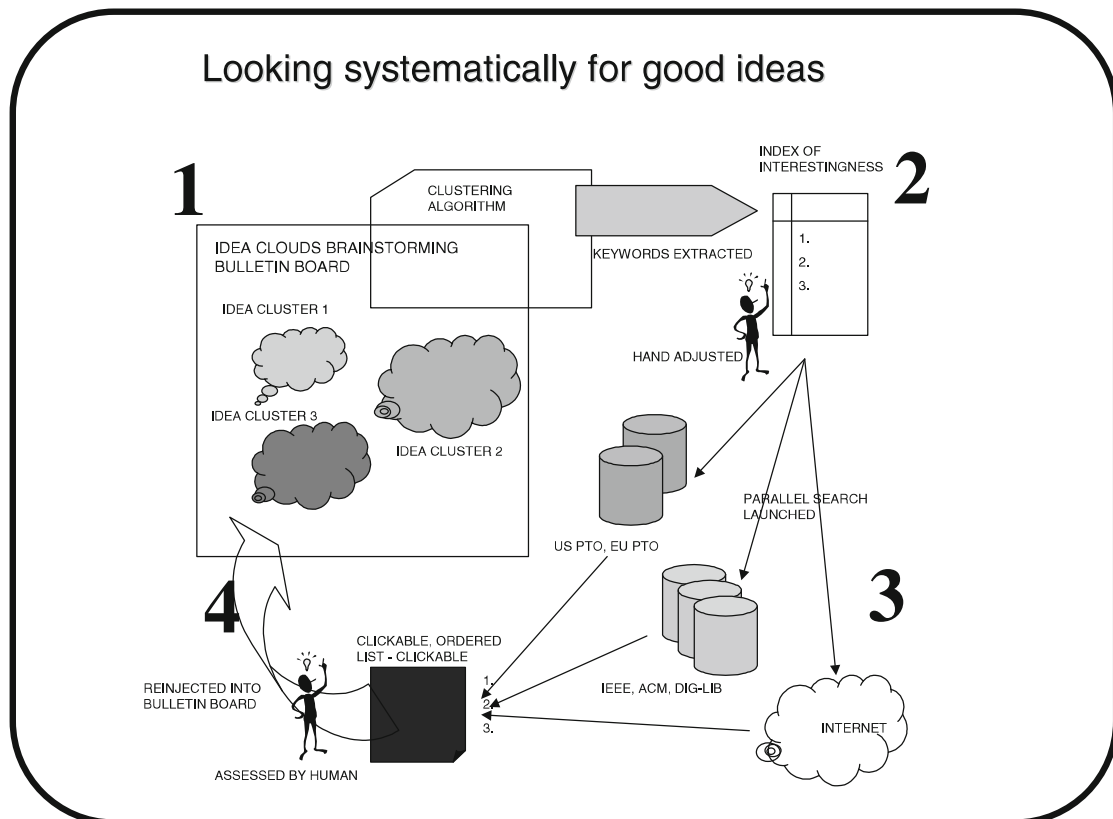


Fig. 10.6 Enterprise knowledge for decision support via analytics

timely manner and never missed; transactions in the integration layer are captured and never lost; and analytics in the interaction layer are always delivered and never inaccurate.

Decision support is provided via portals embodied as business and IT cockpits for executives, an operational workbench for managers, and working spaces for employees. Knowledge is captured in enterprise management analytics. This is yet another enterprise architecture landscape which combines the three-layers stratification principle with analytics technologies to illustrate the current state of enterprise KM systems for dependable and effective decision support.

Harvesting of employees' ideas represents an important activity as it may spawn the seeds of valuable new processes, inventive technologies, or innovative solutions. After initial triage and assessment, ideas could be suitable for transformation into valuable intellectual property – as patents for example. Figure 10.7 depicts a hypothetical example, illustrating that a large brainstorming exercise, or grand challenge, can create big idea clouds which could be harvested, transformed and potentially monetised. It is an illustration of KM deployed for innovation on a mass scale, where emerging cloud computing facilities may enable rescaling of these processes by orders of magnitude (Delic & Fulgham, 2004).



Delic & Fulgham – Dec 2004

Fig. 10.7 Knowledge harvesting and transformation: conceptual IP architecture

Figure 10.7 depicts, in stylised form, that (1) ideas are created, spawned and enriched; (2) the ideas are then organized and ordered by potential value or interest-iness; then (3) the ideas are compared and measured against similar ideas in patent repositories, document libraries or internet documents; and finally, (4) the ideas are either refined and formalised or re-injected into yet another round of brainstorming.

Some large companies have arranged intensive sessions or grand challenges creating more than 100,000 ideas in a very short period of time – so the next possible challenge will be in the automation of processes related to triage, evaluation, valuation and formalization of the assessments of ideas. Due to this automation, the amount of innovation knowledge captured will be extremely large, and the power of scope and scale of such a system and its potential monetary value can only be imagined.

10.5 Enterprise Knowledge Architecture

Contemporary enterprise applications usually reside in data centres and have a typical stacked architecture (Fig. 10.8). Web servers manage interactions, deliver content and capture traces (front-end system) for enterprise applications residing in application servers (middleware). It is common for the databases capture events, transactions and analytics in the back-end system. To deal with high load and transient peaks, load-balancers are installed on the front-end and SAN (Storage Area Networks) for archiving in the back-end.

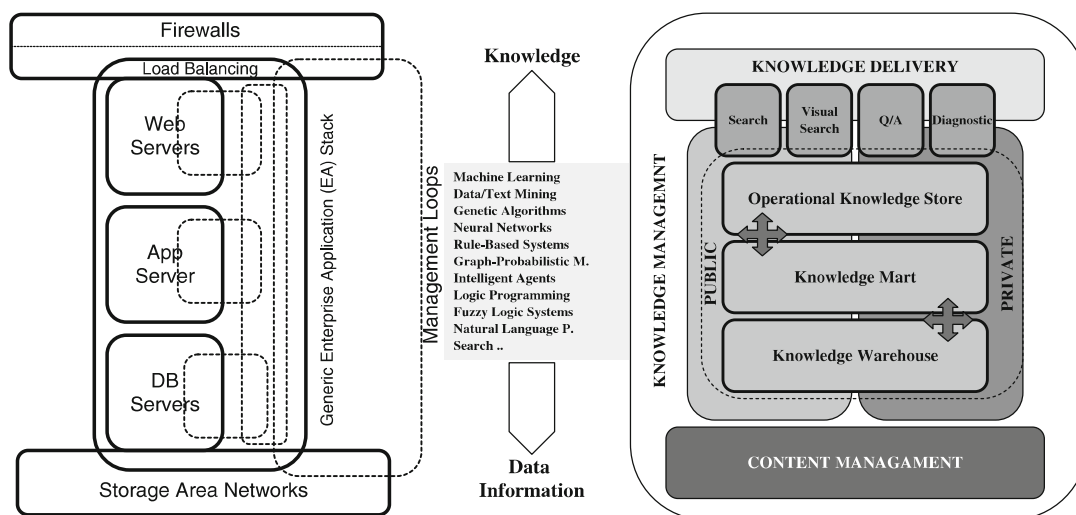


Fig. 10.8 Enterprise knowledge management stack

The abstracted enterprise knowledge management system can also be shown with a stacked architecture (Fig. 10.8), and here we recognise three characteristic layers: front-end, middle/application layer and back-end. We indicate a whole slew of technologies (middle/application layer) originating from Artificial Intelligence research which represent the essence of many KM applications (Cannataro & Talia, 2003;

Delic, Riley, Bartolini, & Salihbegovic, 2007). Knowledge is delivered via various types of portals, either to registered internal or external users or anonymous web consumers. It is typically the case that knowledge users can become knowledge producers via various discussion forums. Another channel of delivery and exchange of knowledge is via machine-to-machine exchanges.

The knowledge (idea) harvesting system described previously can be implemented as an enterprise application with three layers in which content (as externalised knowledge) is being processed and stored within three logical knowledge layers. The Operational Knowledge Store provides rapid access; the Knowledge Mart is an intermediate knowledge repository; and archived knowledge is stored in the Knowledge Warehouse.

We should stress that all conceptual drawings show an architecture which can be materialised with different logical and physical architectures, depending on the deployment domain and choice of key technologies – KM technologies such as content management, enterprise search, delivery portals, discussion forums as key enterprise components glued together via Service Oriented Architecture (SOA) into service delivery fabrics.

10.6 Enterprise Computing Clouds

Cloud computing is the next evolutionary step in the distributed computing field enabled by:

- radical price/performance improvement leading to commoditization
- technology advances with multi-core and energy-aware chip designs
- architectural interplay of warehouse-scale computing and huge number of intelligent edge devices

Large business enterprises have strong incentives to consider their architecture plans in light of developments in cloud computing (Sun Microsystems, 2009).

A possible instance of cloud computing serving billions of users can be depicted as the next wave internet in which the number of devices, gadgets and things can easily surpass 10 billion items (Delic, 2005; Delic & Walker, 2008) (Fig. 10.9). It will be served by strategically placed data centres federated into grids via efficient communication fabrics. In data centres clusters of various sizes (hundred to thousands of machines) will be dynamically allocated to handle varying enterprise workloads. At the chip level, programming of multi-core will become the principal preoccupation of designers aiming at energy efficient designs.

The entire chain from the chips, via racks, clusters and data centres should be designed with cloud computing in mind. The same should apply for the software design. At the level of large aggregation of grids, entirely new economics and legal concerns can govern traffic flows, data storage and choice of application execution

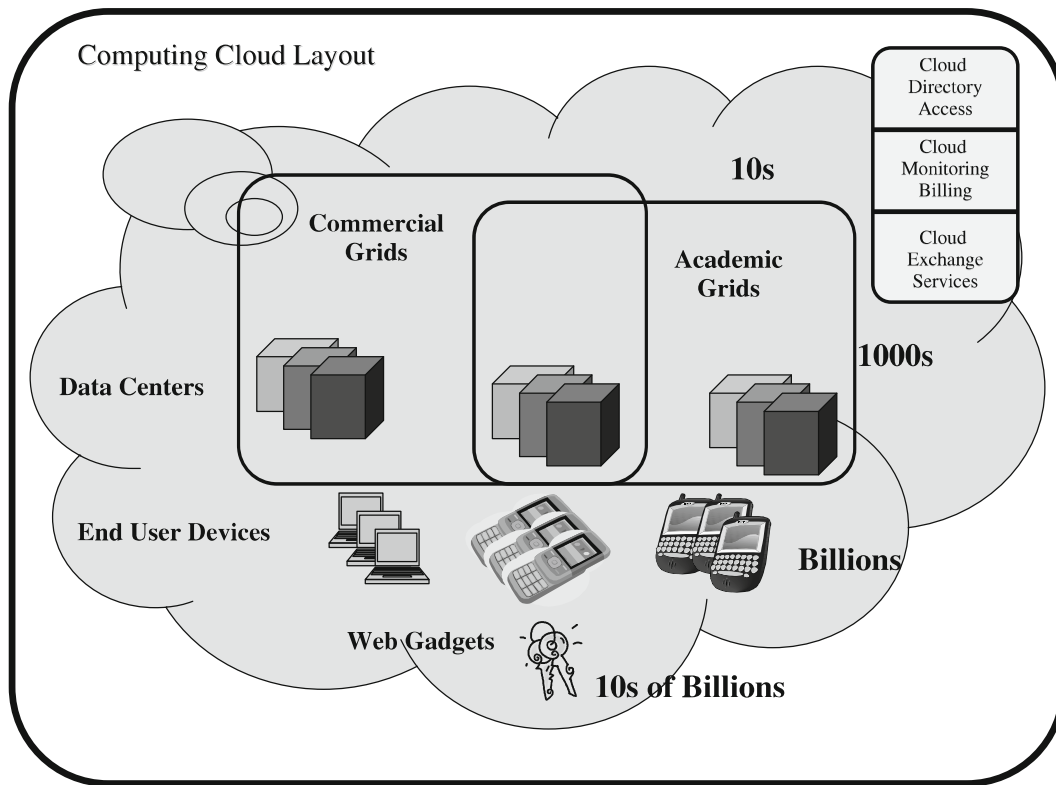


Fig. 10.9 Computing cloud layout

location. This represents a wide number of unresolved issues and particularly hard challenges.

The business reality of the enterprise and constantly changing market conditions will dictate specific choices of the Enterprise Knowledge Clouds, which we describe in the next section.

10.7 Enterprise Knowledge Clouds

Taking into account the current organisational layout of the typical enterprise, we can outline a generic architecture for enterprise clouds. This generic architecture has three principal architectural layers: *private*, *partner* and *public* cloud. We postulate that knowledge management techniques will be appropriately spread over the each and every enterprise cloud. This natural separation is dictated by the required capability of each cloud: security and privacy is a must for the private cloud; availability and reliability is a precondition for the partner cloud; and rescaling and coverage is important for the public cloud (Fig. 10.10). These requirements will be not only guiding principles but also design criteria for enterprise clouds.

We can easily imagine that Finance, Human Resources and R&D Labs will be the prime candidates for the *private* enterprise KM cloud. The Supply Chain and Delivery organisations will naturally fall into the *partner* cloud; while Sales, Marketing, Public Relations and Publicity would be natural fits for the *public* cloud.

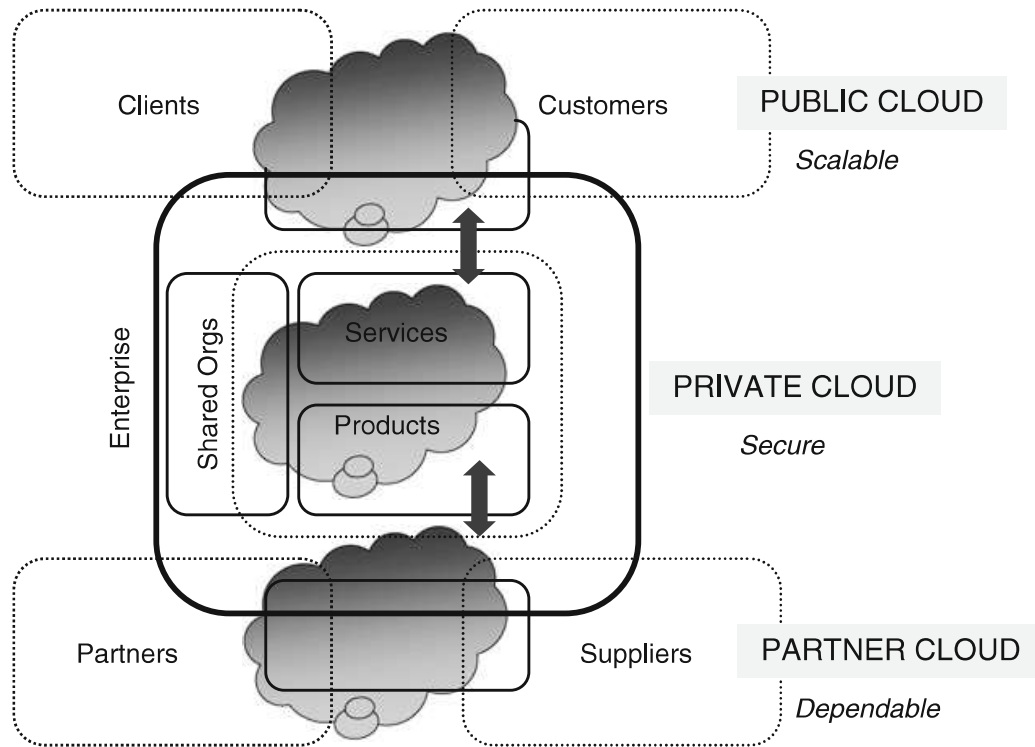


Fig. 10.10 Enterprise clouds stratified

The private, partner and public enterprise clouds should have the facility to inter-operate and exchange data, information and knowledge on a regular and intensive basis. The choice of technologies for the enterprise clouds will be critical, and the emergence of suitable standards is keenly anticipated.

It is not expected that the large enterprises will switch overnight onto cloud computing fabrics, but we expect they will start to gradually deploy cloud-based applications for a few, carefully selected domains. Each previous wave of enterprise technologies has gone through the *prototype-test-deploy* cycle, and cloud technology will not be different. It is also during this time that choices of key cloud technologies respecting the ultimate capability for each cloud type will be made.

Monitoring, measurement and calculation of key performance parameters for each enterprise cloud should be undertaken in order to measure the impact of the new cloud architecture on enterprise performance, and justify investment in the new technologies.

10.8 Enterprise Knowledge Cloud Technologies

Figure 10.11 depicts an abstracted cloud architecture and shows three principal groups of technologies that provide *virtualisation*, *automation* and *scheduling*. *Virtualisation* (of hardware and software) will provide better use of resources; *Automation* will lower support costs and improve dependability of the clouds; and *Scheduling* will enable economics-based reasoning about the use of resources and

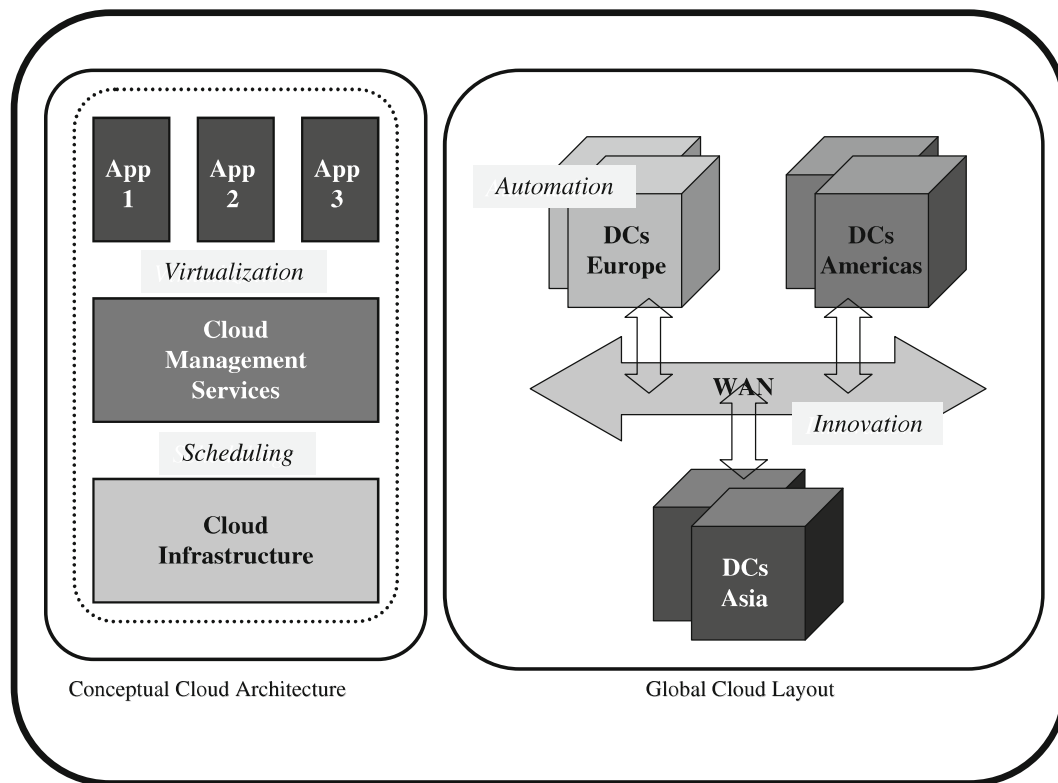


Fig. 10.11 Cloud computing technologies

dispatching of the enterprise workloads. Finally, we assume that the innovation in communication part of the clouds should lead to breakthrough improvements.

All these technologies could be categorised as open-source, proprietary or hybrid. It is beyond the scope of this discussion to delve deeper into detail, but we point out that these choices are critical for real-world deployment. Intuitively we would suggest open-source technology for the public cloud, proprietary for the private cloud, and hybrid for the partner cloud, but we also recognise that this is a difficult problem and once put into a real-world application context, the choices might not be right nor will alternatives be obvious.

The choice of technologies for enterprise clouds will be the difference between success and failure. We expect that the private, public and partner clouds will inter-operate, so the choice of technology should take into account existing or emerging standards which will enable future flows among clouds. All interested parties in this domain have reasons to participate in establishing standards: some self-serving, some altruistic. We expect to see some robust negotiations between the “proprietary” and “open” camps.

Looking back over the short history of cloud computing, we can identify some early platforms on which a very high number of users have developed and deployed a large number of applications. Thus we expect that the leading vendors will try to create very large scale platforms that will attract millions of developers. Platform development usually means the choice of programming language and associated framework. This has advantages for developers and consumers, but also for the

vendors in that it tends to lock developers into a single platform. Interoperability of platforms will pose some of the greatest challenges for cloud computing in the future.

Ongoing technology developments are especially noticeable, and sometimes targeted, by small companies which aim to exit the business via a sale to an established, global vendor. We have seen this happen already with some small companies having been sold for (up to) several hundreds of million of dollars. This is an area for future and exciting developments.

Automation, especially of data centres, will represent the most intricate part of the cloud, as it must address multiple engineering issues and big challenges in which the ultimate goal is multi-objective: the maximisation of utilisation and monetary benefits, and the minimisation of energy cost. All this is to be done whilst guaranteeing dependability and achieving performance objectives. We see here a long road of future research, engineering development and technology innovation (Armbrust et al., 2009). All this will be a critical path to see enterprises running the majority of their business in the cloud.

10.9 Conclusion: Future Intelligent Enterprise

If we take a longer perspective look into past technology developments and business evolution, we observe some distinct phases characterised by a single word to describe an entire technology epoch (Delic & Dayal, 2002). For the automotive industry *automation* of production was the key technological advance; *integration* for the aviation industry and aircraft production; *optimisation* for e-commerce; and for the forthcoming service industry, it is *adaptation* (Delic & Greene, 2006) (Fig. 10.12).

Adaptive behaviour is a characteristic of living systems, while businesses are hybrid systems combining people, technology and processes into orchestrated whole. We believe that the injection of technologies will improve interconnectivity, reduce latencies and increase speeds while improving the problem solving capabilities based on a higher ‘knowledge density’. We call such an enterprise an *Intelligent Enterprise* to denote improved behaviour and the ability to adapt and survive in changing circumstances. While we do not (yet) compare this artefact to intelligent living creatures, the analogy is clear.

We postulate that the synergy among big data, big mobile crowds and large infrastructures will lead to unprecedented improvements in the key indicators above (Fig. 10.12).

The emerging cloud computing paradigm embodied in useful applications for the enterprise knowledge management offers:

- radical cost reduction
- a great ability to scale
- much improved agility

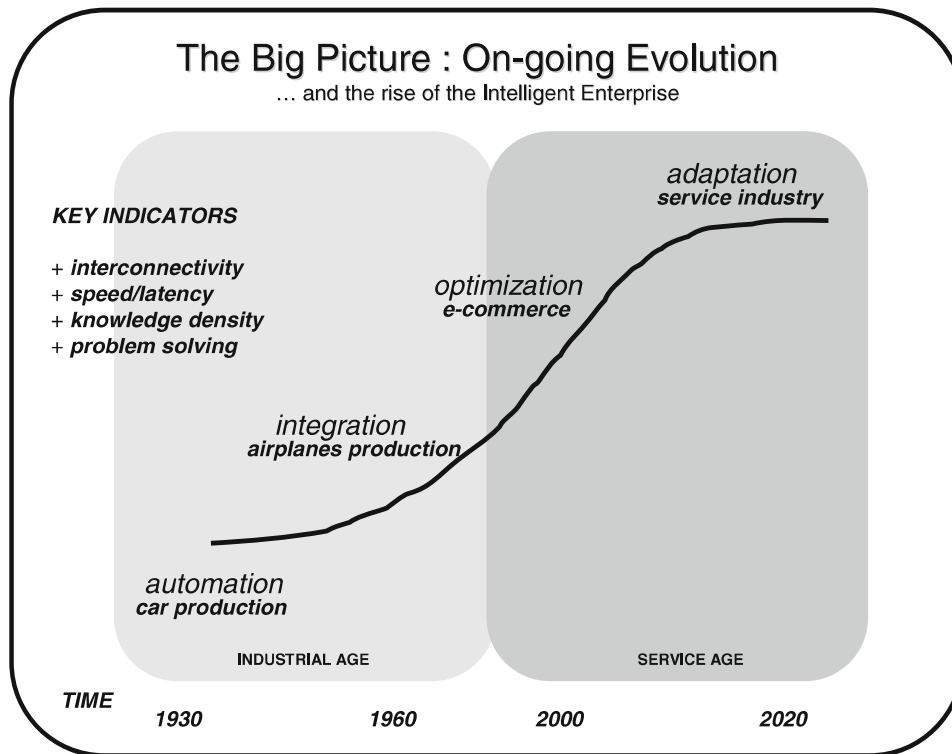


Fig. 10.12 Ninety years of technology evolution – market view

As such, it might be a good showcase for cloud computing success. In practical terms, enterprise architecture will evolve towards cloud architecture, and all architectural components and layers will be impacted and adapted accordingly.

New technologies will enable the next wave of business models, impact market developments, and see the rise of much changed business enterprise.

References

- Armbrust, M., Fox, A., Griffith, R., Joseph, A. D., Katz, R. H., Konwinski, A., Lee, G., Patterson, D. A., Rabkin, A., Stoica, I., & Zaharia, M. (2009). *Above the clouds: A Berkeley view of cloud computing* (University of California, Berkeley, Tech. Rep. UCB/EECS-2009-28).
- Cannataro, M., & Talia, D. (2003). The knowledge grids. *Communications of the ACM*, 46(1), 89–93.
- Casati, F., Castellanos, M., Chan, P. K., Dayal, U., Delic, K. A., Greene, W. M., et al. (2004). Enterprise management analytics. *Proceedings of the Eleventh OpenView University Association Workshop (OVUA04), Paris, France*.
- Cerri, D., Valle, E. D., Marcos, D. D. F., Giunchiglia, F., Naor, D., Nixon, L., Teymourian, K., Obermeier, P., Rebholz-Schuhmann, D., Krummenacher, R., & Simperl, E. (August 2008). *Towards knowledge in the cloud* (Tech. Rep. DISI-08-30, University of Trento, Italy).
- CLOUDSCAPE (October 2009). The 451 group and Tier1 research. *The cloud codex*, CLOUDSCAPE. Available from: http://www.the451group.com/cloudscape/451_cloudscape.php.
- Delic, K. A. (2002a). Enterprise IT complexity. *ACM Ubiquity*, 2002(Issue January). Available from: http://www.acm.org/ubiquity/views/k_delic_1.html.

- Delic, K. A. (2002b). Enterprise models, strategic transformations and possible solutions. *ACM Ubiquity*, 2002(Issue July). Available from: http://www.acm.org/ubiquity/views/k_delic_2.html.
- Delic, K. A. (2003). IT Services = People + Tools + Processes. *ACM Ubiquity*, 4(37). Available from: http://www.acm.org/ubiquity/views/v4i37_delic.html.
- Delic, K. A. (2005). Science and engineering of large-scale complex systems. *ACM Ubiquity*, 2005(Issue March). Available from: http://www.acm.org/ubiquity/views/v6i8_delic.html.
- Delic, K. A., & Dayal, U. (2002). The rise of the intelligent enterprise. *ACM Ubiquity*, 2002(Issue December). Available from: http://www.acm.org/ubiquity/views/k_delic_4.pdf.
- Delic, K. A., & Dayal, U. (2004). Adaptation in large-scale enterprise systems: Research and challenges. *ACM Ubiquity*, 2004(Issue August). Available from: http://www.acm.org/ubiquity/views/v5i23_delic.html.
- Delic, K. A., & Fulgham, M. T. (2004). Corporate innovation engines: Tools and processes. *Proceedings of the 5th International Conference on Practical Aspects of Knowledge Management (PAKM2004)*, Vienna, Austria, 220–226 .
- Delic, K. A., & Greene, W. M. (2006). Emerging sciences of service systems. *Proceedings of the R&D Management Conference, New Delhi*, 140–144.
- Delic, K. A., & Riley, J. A. (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., & Walker, M. A. (2008). Emergence of the academic computing clouds. *ACM Ubiquity*, 2008(Issue August). Available from: http://www.acm.org/ubiquity/volume_9/v9i31_delic.html.
- Delic, K. A., Greene, W. M., & Kirsch, H.-J. (2004). Intelligent routing of enterprise IT Events. *InfoManagement direct*, February 2004, Retrieved March 10, 2010, from <http://www.information-management.com/infodirect/20040227/8166-1.html>.
- Delic, K. A., Riley, J. A., Bartolini, C., & Salihbegovic, A. (2007). 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 Herzegovena.
- International Communication Union ITU (2009). *Distributed computing: utilities, grids & clouds* (ITU-T Tech. Watch Rep. 9, 2009). Retrieved 10 March 2010, from http://www.itu.int/dms_pub/itu-t/oth/23/01/T23010000090001PDFE.pdf.
- Sun Microsystems (2009). *Introduction to cloud computing architecture*. White Paper.
- Zachman, J. (December 1999). Enterprise architecture: The past and the future. *Information Management Magazine*. Available from: <http://www.information-management.com/issues/19991201/1702-1.html>.