Waiting for Gadgets

Lord of the Files—Social Networks, Portals and Knowledge Management

DR. PRABHAKAR RAGHAVAN AND CHRIS BIOW
Vice President and Chief Technology Officer, Verity, Inc.   *Professional Services, Verity, Inc.*

Government and business enterprises recognize the need to make their employees more productive through effective knowledge management. However, some of the most valuable knowledge in an enterprise resides in the minds of its employees.[1] Technology is available to capture, analyze and retrieve information from the documents and databases in the enterprise, but enterprises face a challenge in combining their digital information with the knowledge and experience of employees. A promising answer to this challenge comes from social network theory. This article describes the elements of a modern knowledge management platform, building up to the role of social networks.

Dr. Prabhakar Raghavan

Chris Biow

Social network theory is rooted in the work of Stanley Milgram, a Harvard social psychologist.[2] Milgram reported results of experiments in which he asked each of several subjects—in Omaha, Nebraska and in Wichita, Kansas—to convey a letter to an associate of his in Boston by forwarding the letter to someone they knew on a first-name basis. The objective was to get the letter to Milgram's associate as quickly as possible, with each forwarding being only to someone known on a first-name basis. By analyzing the patterns through which the letters finally reached their destination, Milgram found that most paths took six or fewer hops to arrive. This has led to the folklore that any two people in the United States are linked in a social network by a mere "six degrees of separation".

These social networks can be exploited for enterprise knowledge management. But first, we'll review modern knowledge management infrastructure. Although still rela-

tively new, portals have emerged as one of the key tools for leveraging knowledge assets. However, as many industry analysts agree, no single "portal-in-box" approach can fulfill the unique needs of an enterprise:

"Although leading portal frameworks are increasingly function-rich, they are typically incomplete in terms of a total combined knowledge management/integration solution."[3]

Still, there are common elements to effective portal and knowledge management systems. One recent survey referred to these as "key enabling features;" when polled, executives at Global 3, 500 companies said the top three were *search, taxonomy* and *personalization*.[4] Each of these can be thought of as a tier in the system infrastructure.

The first tier takes the classic notion of search, and extends it with a range of methods for information access. When most people think of search, consumer-oriented web search is the first thing that comes to mind. While there are similarities, enterprise search is unique in a number of ways. Most information on the web is found on public HTML pages stored on Web servers, but enterprise information typically resides in various different repositories, systems and document formats. In addition, enterprise data carries with it fine-grained security requirements that are absent in web search.

Let us briefly examine these issues. By "different repositories" we are referring to the fact that in a typical enterprise IT environment, valuable information is scattered across many systems—email, groupware, relational data-bases, websites and content management systems. However, a user looking for information cannot predict in advance which of these sources contains relevant content. It is therefore necessary to index documents from all of these different sources to make them searchable with a single query. The documents can exist in a variety of formats ranging from simple web pages to slide presentations to PDF—the indexing engine must decipher and extract information from all of these formats.

These documents have various sorts of *structure,* some consisting of simple text, others being organized or for-

[1] Beth Serepca, "Why Knowledge Management," *The Journal of Public Inquiry*, Spring/Summer 2001.

[2] Stanley Milgram. "The small world problem," *Psychology Today*, May 1967.

[3] David Yockelson, "Portal Arms Dealers: Part 1—Filling Out the Framework," META Group Inc., September 20, 2001.

[4] Frank E. Gillet et al, "Making Enterprise Portals Pay," Forrester Research, Inc., August 2001.

matted into recurring sections that have a specific meaning, or *semantics*.

For instance, a data sheet on lumber may show three fields containing the dimensions of the lumber (e.g. width, height, and length in inches), one field specifying its grade, and one listing the unit price in dollars. The sheet may also give a textual description of the lumber. This would be typical of the *semi-structured* documents most common in enterprises—there is a structured component (the dimensions, grade and price) as well as an unstructured component (the free-text description). The significance of this distinction is that whereas computers can readily extract the semantics of the structured portion, they typically view the free-text description as a bag of words. By specifying the structured components as parameters, end-users can readily locate information. For example, the user can enter a query specifying *price less than $50* and *grade "construction."* The user may also search the unstructured portion, such as by specifying that text should contain the word *cascades*. This represents a significant step beyond traditional keyword search, giving the end user a range of options with which to more effectively find information. This efficient retrieval of documents matching such "parametric" queries is a key challenge for state-of-the-art knowledge management systems.

Fine-grained security is another difficult issue in enterprise search. On the web we are all "equal citizens," allowed to access the same set of web pages. The situation is considerably different in an enterprise. The role that each employee plays translates into the documents that an individual may access. Access control implies that the only documents presented to a user in response to a search should be ones that user has permission to access. Even listing (say, the title and source) of other documents in result lists may lead to the leakage of secure information that is unacceptable in an enterprise setting. Thus an engineer not privy to M&A information should not, on issuing a query for "IBM" be shown a document title "M&A discussions with IBM" in the search results. Even if he can't access and view the document, secure information has already been compromised. The best approach to addressing this issue is to filter the results of searches by the user's access privileges—at the time of the search. This issue is compounded by the different security models used by various enterprise information sources, all of which need to be supported by the retrieval system.

Robust, scalable powerful and secure search thus forms the basis of modern knowledge management infrastructure. The next tier exposes the *organization* of all of the content in the enterprise, to make it easier for the user to navigate. Libraries have done this for centuries: their books are not randomly laid out on shelves, but rather grouped together by subject matter and housed on logically labeled shelves. Similarly, products in stores are organized into aisles and shelves.

How does one go from the set of all documents in an enterprise to an organized, navigable *taxonomy* of labeled concepts (think of a directory tree similar to the one you see in Windows Explorer), with each document assigned to its category in the taxonomy? Unlike a physical store or library, a document in an enterprise taxonomy may in fact be assigned to multiple positions; a list of IBM servers might be assigned to a subcategory called "Hardware" under the *Engineering* category, as well as to a subcategory called "Duplicate Product Lines" in an "IBM" category found under *Business Development*. There are generally three steps to achieving this

- Build the taxonomy
- Assign classification rules or models to the categories in the taxonomy
- Populate the taxonomy.

Let's expand on these three steps:

- **Build and label the taxonomy.** A worker decides on the skeleton of the navigation structure, without necessarily specifying how to decide which documents go where. The worker might decide that at top level the user will see *Arts, Business, Government* and *Science* categories; under *Business,* the subcategories could be industry groups such as *Construction, Services, Publishing,* etc. This is equivalent to a librarian laying out the library's shelves and affixing subject labels. This step is usually performed by *domain experts* who understand the enterprise, its content and its user base.

  Technology can help these domain experts, using ideas from the discipline of *unsupervised machine learning.*[5] Here we use a computer program to analyze the content in an enterprise, then extract and present the dominant themes in this content to the domain expert as candidate categories. The program may also suggest a hierarchy by which to organize categories and subcategories for end-user navigation. Although this technology has made significant improvements in recent years, it should not be expected to completely replace human domain expertise. However, such techniques can be a significant aide in enhancing the productivity of the expert building the taxonomy. It is also possible to use a pre-defined taxonomy (much as librarians use a standard classification scheme, and stores may use standard layouts) for the particular industry or knowledge domain.

- **Assign rules.** We now essentially have labeled shelves and need to determine how to decide what

---

[5] Ryszard Michalski, Jaime Carbonell, Tom Mitchell, Machine Learning: An Artificial Intelligence Approach, Volume II, Morgan Kaufmann Publishers, 1986.

shelf a book goes onto. Our second step is then to assign to each category in the taxonomy a set of rules, or a *model*, that determine whether a document should be a member of that category. For instance if we had a category labeled *Scotland* in our taxonomy, the rule might specify that if the document had many occurrences of the words *tartan, loch, skirling* or *malt*, it should belong to this category.

The alternative, supervised machine learning, begins with a set of *training documents* for each category; in the above example, we would need a number of documents about Scotland. These are fed to the system in a "training phase," during which it learns a rule for the category *Scotland* from the content of the training documents. The more training data, the more accurate the rules produced. This method reduces the need for domain experts but still requires pre-classified training documents, which are classified by domain experts. Later, maintaining such machine-built rules can be a challenge: as the vocabulary changes, one has to retrain the system with new sets of classified data. Moreover, the best research in supervised machine learning suggests that rules obtained using this method are rarely of sufficient quality to produce the level of classification accuracy required in enterprise settings.[6] On the other hand, this method can be used to make domain experts significantly more efficient; it is this combination that, in our experience, tends to work best.

■ **Populate the taxonomy.** The first two steps have created a skeletal taxonomy structure with rules defining each of the categories. But there are no documents attached to the category. It is as if our library has empty, labeled shelves with rules that tell us whether a book is about *medieval music.* So we must populate the taxonomy, determining which category rules a document matches, and placing the document in the appropriate categories.

Once we have gone through this three-step process an end-user may browse through the taxonomy or issue a *scoped search*—a text search restricted to a specific category or subcategory. Studies show that users tend to employ a combination of searching and browsing to cull the information of value to them.

Thus far we have discussed knowledge management technologies that rely entirely on the content of the documents in an enterprise. The logical next step is to bring in the interactions between users and documents, and between

users themselves.[7] It is here that ideas from social networks come in to play. Milgram's classical work dwelt on the network of acquaintances. We extend this to consider the broader social network involving people and the documents they access and author, as well as the organizational links between people in the same departments and job functions. How can one exploit this network? For example, if several of a lawyer's colleagues were to find a document useful, it is likely that the lawyer would also find the same document useful. Conversely, if the lawyer were to access many of the same documents that a colleague accessed, the system could recommend that the lawyer talk to that colleague—presumably about the subject matter of the documents of common interest. Technologies based on such ideas are already in everyday use. Analysis of the network of hyperlinks is used by most web search engines to boost the accuracy of search. This network does not explicitly involve links between humans, but it does implicitly link web content creators and the topics of interest to them.[8] In e-commerce, analysis of the affinity between groups of customers and products has led to *collaborative filtering* systems that recommend products likely to be of interest to a user.

These examples demonstrate that analysis of the social network has the potential for mining the knowledge in the minds of people within enterprises. However, several technical challenges present themselves: (1) how does one relate the various entities of interest—people, documents, queries—in a uniform manner? (2) how does one weight the various cues derived from linkages in our network—say between people and the documents they read, versus people and their organizational links to colleagues? (3) how do we decide which linkages are meaningful and which ones coincidental? These and other challenges are raised and addressed in a variety of research projects such as Referral Web[9] and Yenta[10], with the results finding their way back to commercial enterprise applications

Let us now examine the user experience and benefits resulting from such analysis. The system maintains, over time, a profile of the user's role and interests in the enterprise. When the user issues a search, the system cannot only present documents matching the search criteria as in a traditional retrieval system, it can also recommend documents that take into account the user's profile. Thus when a user whose profile incorporates an interest in cars

[6] Kathleen Hall and Daniel Rasmus, "Criteria for Selection: Classification Technologies," Giga Information Group, August 21, 2001.

[7] Charles Canning-Smith, "Expertise Software: KM's 'Killer Application'?", Basex, Inc., May 2001.

[8] Prabhakar Raghavan et al, "Hypersearching the web," *Scientific American*, June 1999

[9] H. Kautz, B. Selman, M. Shah, "Referral Web: Combining Social Networks and Collaborative Filtering," *Communications of the ACM, Vol. 40, no. 3,* March 1997.

[10] Leonard Foner, "Yenta: A Multi-Agent Referral-Based Matchmaking System," First International Conference on Autonomous Agents (Agents '97) Proceedings, February 1997.

searches for "Jaguar," the system recommends documents on the Jaguar automobile rather than the animal, even though the search results may include a mixture of both kinds of documents. When the user selects a document to view, the system can recommend documents related to the selected document, taking into account the context set by the query, the user's profile, and the preferences and profiles of other users who chose to view the same document. A final scenario—consider several users choosing to view the same set of documents. A new user viewing some of the same documents is quickly presented with the docu- ments implicitly deemed most useful (by the collective wisdom of previous users), as well as the names, contact information and instant messaging links of the users who previously trod similar paths of searching and viewing.

The net result is that users can quickly locate both the information and expertise relevant to their current task, aided by a system that incorporates analysis of their context with the more traditional analysis of document content. While the science and technology in this area are still relatively young, social networks already offer great promise in tapping the tacit knowledge within an enterprise. ☎

Devices and Desires

The Workplace of the Future: Your Office Is Anywhere You Are

P R A N S H U   T E W A R I
Workstation Brand Manager for Government, Dell Computer Corporation

The workplace of America is transforming year by year. From typewriters to word processors, from snail mail to e-mail, and from neighborhood store to online store, technology has been a key enabler in helping the workforce become more accessible, efficient and productive with fewer resources.

Government is no exception. Technology has changed the way government conducts its daily operations. Moving more operations online and consolidating networks are two popular information technology initiatives government is conducting.

The ability to work remotely has emerged as another important trend of the 21st century. As notebook computers, personal digital assistants and wireless technology price barriers wear down, more organizations have the means to provide mobile technology to employees, allowing them to have remote access to their data as well as colleagues.

One drawback to mobile computing has been around the lack of processing speed to operate powerful applications such as geographic information systems (GIS) and computer-aided design (CAD)—software that requires graphic and processing performance previously only found in desktop workstations. Until now.

In conjunction with new Intel mobile processors running at high speeds, Dell recently introduced one of the industry's first mobile workstations in the form of a notebook computer, the Dell Precision™ Workstation M50.

The mobile workstation was specifically designed for professional users who require the freedom and flexibility to be productive no matter where they are located.

Mobile workstation capabilities can help government agencies and departments with a variety of functions. Take GIS for example. GIS allows users to transform geographic data into mapped images. Users from many federal government agencies can now use GIS tools in the field to map town boundaries, property data, weather patterns, population forecasts and land-usage. Law enforcement also can use GIS for neighborhood crime tracking and mapping.

CAD is another popular workstation-class application for government. CAD allows architects, engineers, drafters and artists to create 2-D drawings and 3-D models. Users can create technical government buildings, including office space and highly complex structures like nuclear reactors, as well as other mechanical objects such as satellites, fighter jets and vehicles. With a mobile workstation, these users are no longer tied to their office and can modify and refine their models and illustrations anywhere they need.

Another requirement for high-end mobile computing is a notebook screen that looks as sharp as a desktop monitor. Some users require immediate access to crisp and clear video footage and a contrast ratio that renders rich and sharp colors. Some may need wider angles that allow users to view the screen from various positions without compromising image quality. For example, some situations call for a group viewing of data. With access to only one notebook in the field, the ability for a small group to see the screen is critical when quick decisions must be made.

As technology advances, mobile computing will continue to evolve and will expand to uses beyond the imagination. With the introduction of the mobile workstation, the industry is one step closer to breaking the mold of the traditional workplace, accomplishing goals that could never have been attained before. ☎

Riding in Cars with Toys
Wireless

AL STERN

Director of Marketing, Go America

In a business world that demands ever-increasing levels of speed and productivity on the one hand . . . while demanding you limit and control costs on the other . . . how do you meet the challenge?

In a business world that says you should travel far and wide for every opportunity . . . but must not let go of existing responsibilities . . . how can you get it all done?

In a business world that insists you stay in touch . . . even when you are out of reach . . . how should you communicate?

In a world where emergency response is a day to day necessity . . . not a training drill . . . how can you rise to the occasion?

An important and effective part of the answer to all of these issues is *Wireless Data.*

What Is Wireless Data?

Wireless data provides the ability to gain access to e-mail, intranet, corporate database information and the Internet using a mobile device, eliminating the need for business travelers to locate an electric outlet or telephone jack.

With the right wireless data capabilities you can not only send and receive information, create and edit documents and exchange messages; you can access critical corporate data in a secure fashion.

What Can It Really Do?

In New York City on 9/11 there was virtually no wired or wireless phone service. When the first airplane hit, the wife of a wireless data company president, who worked near the Trade Center, sent a wireless data message to her husband saying that she heard a crash, was afraid, and was heading home. Moments later the second plane hit. Throughout her trip she was able to communicate with her husband. He knew she was safe . . . and she knew what was happening at all times.

A CEO traveling for business faced an impending disaster as a tornado headed for his corporate headquarters—just the sales team was scheduled to hit the streets. He quickly posted a notice on the company intranet using his wireless device. Within minutes, sales staff both in the office and in the field, received notification on their wireless devices with the day's emergency protocol.

A financial services professional was meeting with clients in D.C., when his colleague published an important research report regarding a product from a company competitor. During the meeting, she accessed the report, sent it to a local fax for printing and checked her Bloomberg Phon to see the stock reaction—all from her wireless device—all while sitting with the client.

A scientist for a major pharmaceutical company was outlining benefits of a new, but controversial drug to a group of physicians when the company community-relations expert learned of a spontaneous and dangerous protest outside the convention hall. She was able to email the executive from her seat in the audience, informing him of the danger, and was able to call additional security guards to the scene using a customized wireless application developed by company's internal IT staff.

Al Gore was traveling in a motorcade on his way to make a concession speech, when he received an urgent, secure email on his wireless device from an aide telling him to turn around.

Wireless data has already changed the way Americans do business and live their lives. For an industry in its infancy—the potential impact is already clear. Wireless data is no longer a "nice to have" for American businessmen and women . . . it is a "must have."

How Does All This Work?

Mobile Devices

Workers within the enterprise vary greatly on the mobile device they want to use. For many years, the laptop was considered the mobile device of choice. It provided the mobile user the connectivity needed and the applications wanted. It was an easy favorite of the IT department as well. Standardization of both hardware and software allowed the IT staff to keep costs low, maintenance in check, and security high. The introduction of the PDA (Personal Digital Assistant) has since changed that.

PDA/Mobile Devices

RIM 950,957	Compaq iPaq
RIM 850,857	Casio Casiopeioa
Palm III, V, M500 series	HP Journada
Handspring	Laptops with Type II PMC Slots
Sony Clie	

Conservative estimates approximate the current number of PDA's being used by consumers to be in the 20–30 mil-

lion range. That range has the expectation of doubling within the next two years. It is believed that half of these devices will have some sort of wireless connectivity, whether it is strictly email, or email and Internet access. The choice of the PDA has previously been personal and diversified by many factors including position within the enterprise, demands of the position, and geographic location.

In order for the PDA to be embraced and supported by the enterprise, it has to have wide connectivity capabilities including email, access to intranets and the internet, and corporate applications. Offline functionality, easy access to local data stores, signature capture, and multi-threaded operating systems will be the required functionality, not the exception. Enterprises use these devices as a way to improve quality of work life, while staying connected to the user.

Connectivity allows the mobile employee the access to the information they need and want, when and where they want it. For the IT staff, connectivity is not measured as knowing where the employee is or what data they are accessing, but rather:

- What applications are on the device?
- What version of operating system is the device running?
- When was the last backup of the device completed?
- What upgrades to the device have been implemented?

Application Language

The allure of the Internet for the enterprise was the ability to move to a more conventional way of providing access to corporate data. Client/Server architecture initially enhanced users capabilities, but also increased IT support of the application, especially for remote users. Using "visual tools" the enterprise could rapidly develop applications, scaling user access to data quickly, but often not effectively. Too often, "off the shelf" software loaded on the remote computer conflicted with the enterprise-developed applications, causing frustration for the user and the IT staff.

Wireless Enterprise Applications

Email	FFA (Field Force Automation)
SFA (Sales Force Automation)	Time and Billing
CRM Customer Relationship Management	Intranets

The Internet spurned a new set of application development for the enterprise. The deployment was as rapid, but the standards were easier to support. The architecture, while still client/server, was standardized on the client side. The server did most of the work while causing few, if any, conflicts on the client. Employee and IT frustration were curtailed dramatically, and the application usage increased.

The Internet standard of HTML (Hyper Text Markup Language) has been trained, accepted, and deployed within the enterprise. As demands increased, so did application language capabilities. Introduction of Java, Java Scripting, ASP (Active Server Pages), and XML (Extensible Markup Language) has allowed the enterprise not only to extend the application, but provide extensive scalability and capability in a singularly controlled environment, specifically for the wired world.

Wireless access has changed the application access again. New languages, protocols and standards have emerged. HDML (Handheld Device Markup Language), WML (Wireless Markup Language), Slim HTML, WAP (Wireless Access Protocol), and J2ME have affected the extension of the application. They have also affected application deployment and scalability.

Security

When wireless mobile data access is required, the architecture of application deployment, while not compromised, must change. This is very similar to the way in which networks had to adapt to the Internet when first implemented. The enterprise is required to open its architecture to third parties, but this accessibility does not have to compromise the existing security strategy.

Common Encryption Technologies

SSL, SSL128	DES, Triple DES	PGP
PPTP	ECC	RSA

Security solutions like firewalls, VPN's, encryption schemas, intrusion detection, anti-virus software, and other tools, build a protective wall around important data.

There are two models of wireless access can be deployed. The first is a hosted model by Application Service Providers (ASP). The second is a server solution instituted behind the corporate firewall.

By selecting the security model that meets your organization's needs Wireless Data can deliver a level of secure communication that can be made acceptable to even the most security conscious organizations.

In Summary

In a world where smarter, faster, cheaper and more effective is cost of entry in a world where emergency response has become part of our daily lives and in a world where "I didn't get the message" is unacceptable . . . safe and secure wireless data is a must have.

The industry is abuzz with visions of the *killer app.* But in truth, only you know what that is. For you, the *killer app* is whatever you need to access every day, any way, in a secure fashion.

Some of the toughest have already gone wireless— should you? ☎