

Privacy and knowledge management: Challenges in the design of the Lotus Discovery Server

by A. L. Schirmer

The Lotus Discovery Server™ (LDS) is a knowledge management system that lets users browse and search for information and expertise, collaborate instantly with colleagues, share knowledge, and look for resources, in a time-efficient manner. Because a knowledge management system has the power to make visible what was previously obscure, privacy issues are of particular concern. This paper discusses what the LDS product team learned about privacy issues in an enterprise knowledge management product, how it learned these things, and what steps the team took to protect users' privacy, allay their concerns, and promote the value of the product.

As long as there have been information resources, companies have needed the means to manage them, and many tools have been developed for that purpose. In recent years, companies have come to realize that the increasing complexity of their information resources encompasses more than just data. It also includes a precious asset—the human knowledge of its employees, its intellectual capital. That realization has also brought with it the need to protect that asset and to make its use within the enterprise much more effective. Knowledge that is hidden is without value. What is known needs to be connected with those who need to know. This is the chief aim of knowledge management, which extends prior information management ideas. Although knowledge management is not itself a technology, knowledge management technology solutions, such

as the Lotus Discovery Server* (LDS), have been developed to realize that aim.

As people work to achieve their goals, they constantly need to acquire new knowledge. They do so in many cases by gleaning information from documents created by others, but knowing whether such information exists and where it might be located is often very difficult. A knowledge management system can help with this problem in one way by organizing information for users and providing links to the data resources. Many times, however, people make better and faster progress by learning directly from other people—on their way to gaining expertise themselves, they benefit from someone else's expertise. How do they find someone who has the knowledge they need or who can direct them to it? Knowledge-management technology can help with that also, by assessing people's affinities for topics of information and publishing those affinities for others to see.

Having a software system make assessments about people raises a completely new concern—how will such decisions be made accurately and appropriately? Mistakes in this context can have an adverse impact on the people involved, possibly subjecting them to erroneous and obtrusive requests for information, or associating them with topics that misrepresent their skills and give an inaccurate picture of who they are. Even if the conclusions are accurate,

©Copyright 2003 by International Business Machines Corporation. Copying in printed form for private use is permitted without payment of royalty provided that (1) each reproduction is done without alteration and (2) the *Journal* reference and IBM copyright notice are included on the first page. The title and abstract, but no other portions, of this paper may be copied or distributed royalty free without further permission by computer-based and other information-service systems. Permission to *republish* any other portion of this paper must be obtained from the Editor.

people may not wish to share that information, perhaps because their involvement in a particular area has lessened, the area is only tangential to their true skills, or they simply do not wish to be bothered by information seekers. Each person's reputation, public persona, and self-image—and his or her privacy—are extremely important and must be respected.

When the LDS was first being developed, in 1998, Web commerce was flourishing, and there was a good deal of information available about the issues of privacy in business-to-customer (B2C) and business-to-business (B2B) contexts (publicized by the Platform for Privacy Preferences Project¹ and others). Public discussion about the particular issues of privacy within a business enterprise was slight, however, even when the first version of LDS was released in the spring of 2001. In the enterprise, circumstances are significantly different from B2C and B2B commerce because the users of the system are the employees of the enterprise rather than external customers. Most of the information they need to share to make the system effective has to do with their job roles and skills rather than their personal information. Rather than trying to keep that information hidden from other people as would be the case for Web commerce, the enterprise knowledge management system wants to encourage as much information exposure as possible, in order to facilitate the publication of users' information affinities and increase the flow of knowledge. It must do this while instilling confidence and trust in its users.

In recent years, Web commerce privacy has gotten even more attention (see IBM Privacy Practices on the Web² and the general references). Australia and Canada have instituted a national position of Chief Privacy Officer (CPO). A recent survey of 66 Fortune 500 middle market and emerging technology companies having CPOs indicated that "the functional orientation of the CPO is still evolving,"³ and the importance of this role is expected to grow in the future.

IBM has implemented privacy software solutions such as the Enterprise Privacy Architecture⁴ and Tivoli Privacy Manager for e-business.⁵ IBM has also selected a CPO, established the IBM Privacy Management Council, and instituted the IBM Privacy Research Institute⁶ to promote and advance research in privacy and data protection technology.

Today, there is more discussion in knowledge management forums about the importance of privacy and trust for the success of knowledge management sys-

tems. The literature for commercial knowledge management products such as LDS and products from Autonomy⁷ and Tacit Knowledge Systems⁸ point to privacy features, but differences of implementation make consistent policies and practices across enterprise knowledge management products unlikely. As knowledge management products evolve and improve, the issue of privacy promises to grow in prominence and consequence. In this paper, we present the considerations that went into the design of privacy mechanisms in the LDS and discuss the many issues and challenges involved in designing an effective knowledge management product which shares users' information while safeguarding their privacy.

The Lotus Discovery Server

The LDS⁹ provides knowledge by using a network of people, documents, and category objects (that is, concepts or topics) and their relationships, which it creates from enterprise and external information. We refer to the information that the system records and calculates to create this knowledge as the system "metrics." This process of knowledge creation is dynamic and therefore represents present conditions. It is also configurable, allowing the system administrator to change the significance given to various interactions to reflect the circumstances of a particular enterprise more accurately, as well as the threshold that determines when the system should determine affinities.

Using one or more corporate directories, the LDS creates a profile for each person listed. It also analyzes documents of various types and formats in the enterprise's public data, analyzes the content to determine the most salient categories contained in the data, and organizes those categories into a hierarchy tree. It then assigns each document to one or more categories that best describe its content.

During document analysis, the system also records meta-data about whatever people it detects—for example, information concerning the authoring, linking, and forwarding of a document. From this, it can derive relationships between the people and the documents, and, by extension, between the people and the categories. We define a person whose connection to a category exceeds a defined threshold as having an "affinity" for that category, which to some extent connotes expertise. Users can find people with a particular expertise simply by looking at the desired category in the category tree and seeing who is listed as having an affinity for it.

The system monitors subsequent use of the documents and categories and uses that information to constantly adjust the relationship values among the people, documents, and categories. The system calculates “values” for categories and documents based on how often the LDS users interact with the categories and documents. These calculations factor in dates and times of use, and the system gives more weight to recent interactions. In this way, affinity strengths increase and decline over time.

The LDS user interfaces include:

- The Knowledge Map, which provides a means for browsing and searching the category tree and shows the relationships among the categories, documents, and people.
- The person profiles, which contain the personal information for each user.
- The Knowledge Map Editor, which provides tools for customizing and managing the category hierarchy.
- The Discovery Control Center, the locus for system administration tools.

Additional description of the LDS can be found in a previous edition of this journal.¹⁰

Planning and design. Throughout the Discovery Server project, the product design team sought to apply user-centered design techniques to understand our audience, determine user requirements through role and task analysis, survey competitive products and related domains, and repeatedly confirm our assumptions with a wide and informed set of advisors.

From the beginning, members of the product team realized that the very elements that would make the system valuable might also make its participants uneasy. We designed the product to reveal the categories that underlie an enterprise’s many data stores so that its users could see what information was available to them as well as to identify expertise in an automated way by measuring the connections between those categories and particular individuals. Though this would reveal and organize only information that was already available, the team learned from the unprecedented access to public information by Internet users how disconcerting such revelations could be. However, the parallels to the Internet were limited: the LDS would deal with enterprise information, owned by corporations rather than by individuals, and because it would provide a set of services for which there was little ex-

isting competition, the product team knew it would be breaking new ground. The team needed to start from the lessons of the Internet but also needed to consider a completely new set of special concerns.

The LDS project leader began by looking at the information acquired by existing data-mining tools, some of which could, for example, report a great deal of detail about a person’s computer use. Some of these tools had a comprehensive set of options that let users limit the system’s probing into their practices. In the best cases, the system let users opt in (that is, volunteer information) rather than asking them to opt out. The project leader realized that user preference was a key part of enabling users to control the system, and an initial privacy policy was drafted, stating:

Information can be public or private. Authorship of widely circulated documents and membership in directory groups with open access is public, and everything else is private. The default policy gives users control over the access to, and publication of, information derived from their private information, though override mechanisms are provided for system administrators.

As the product’s lead designer, I had the responsibility to research and implement this policy, plan the supporting features, design the user interface, and make sure that everything could be implemented reasonably. I began by considering my own experiences with Internet resources, concerning both public data and personal data that I had provided for activities such as online shopping. I also referred to user roles and cases developed in conjunction with my co-designer to consider the different types of privacy perspectives that we would need to accommodate. From this, I drafted a set of issues, detailed below, and discussed them with various colleagues on the project team and elsewhere.

One of the great assets available to the product team was a council of design partners, consisting of representatives from various corporations interested in knowledge management. We met with them at regular intervals during the project to get their advice on product issues and design directions. Our product managers reviewed the privacy policy at briefings with the partners and with other Lotus customers, getting their input and corroboration. They also represented LDS at an internal IBM project review of product privacy policies and confirmed them with IBM’s Chief Privacy Officer.

Usability testing for the first release of LDS ranged from early paper prototypes to videotaped laboratory sessions. During testing, the design and usability team members designed tasks to illuminate privacy issues and looked for unexpected ways such issues might manifest themselves.

From the various research processes above, it became clear that achieving a successful design would require understanding user expectations, identifying privacy issues, defining strategies for product success, and then designing the user interface.

Understanding user expectations. To gauge the scope of issues we would encounter, we assessed existing circumstances that we thought would affect perceptions of LDS.

Lack of product precedents. As an innovative product in an innovative product category, LDS would be used by users who had no experience with knowledge management applications. The product “type” itself would not help set user expectations about privacy. Understanding the sometimes complex set of relationships presented by the Knowledge Map’s network of people, documents, and categories would be challenging for them. People might well regard with puzzlement and skepticism any product claims to reveal expertise. The enormous potential of the product could tend to make users suspicious rather than confident. Throughout the project, our conversations with design partners and our usability tests confirmed that this was true.

The Internet and Web browsers. Because we designed the Knowledge Map and profile user interfaces to run in a Web browser, other experiences with browsers would understandably color user expectations, and we thought that in large measure users would conclude they were on the Internet, rather than in a traditional software application. It is likely that their previous Internet experiences would have left them concerned about “cookies,” passwords, obtrusive advertisements, unsecured data fields, and so forth. For many, such artifacts have produced the feeling that Internet interactions, even simple browsing, may tap their personal information without their knowledge.

Privacy policies and guidelines for Internet applications have been emerging for some time but are far from resolved. There are different cultural assumptions in different parts of the world, and the stringency of privacy laws varies widely from country to country. The design team also noted in usability tests

that people of different ages had different expectations about Internet privacy.

In all, people’s expectations for products that they experience through a browser are diverse and fluctuating. LDS is not a pure Internet application, however, and the experience of using LDS is unlike the experience of using the Internet in several ways that may contribute to a user’s anxiety. One factor is that the likelihood of LDS users finding documents they themselves authored is much greater than when browsing the Internet, because LDS specifically targets data from the user’s own corporation. Another factor is that LDS introduces derived information, such as document values and affinity strengths, which are nonexistent or invisible when browsing the Internet. The novelty of these two factors, while perhaps contributing to a user’s curiosity, may also increase the sense of uncertainty.

Change in users’ expectations. As with any complex product, we expected users’ expectations to change over time as they became more familiar with the product. We believed that their reactions would follow a logical progression:

1. Confusion.
At first, people would face new concepts—asking for example, “What are affinities?”
2. Suspicion.
After they understood the concepts, they would wonder how the product works. “How does the system know that? How accurate is its assessment? Does it reveal things in appropriate ways?”
3. Desire for control.
Based on their conclusions about accuracy, users would want some measure of control over the system’s evaluations. “Am I forced to participate? Do I have a say in what the system concludes?” System administrators would want the ability to tune the system to get the best accuracy for their particular sites and circumstances, and would want a say in the way privacy policies are implemented.
4. Trust.
If provided with a product that solved these problems, people would feel comfortable and discover and appreciate the benefits of the system.

Identifying privacy issues. With these expectations in mind, we needed to draw on our various inter-

views and analyses to identify specific privacy issues and form a strategy for dealing with them. We found that after potential users understood how the product would analyze and organize information and how it would derive affinity strengths and document and category values, they often had the following questions:

- Would the system display only appropriate information?
- Would the system use only appropriate information to reach its conclusions?
- Even if it used appropriate information, would the system reach conclusions accurately?
- Would system users have control over sources of information about them?
- Would system users have control over conclusions made about them?
- Would the system allow access to information that was intentionally limited?

Our review of the role of product system administrator revealed an additional question.

- Recognizing that enterprise requirements in different domains and regions might differ, would the system provide flexibility in controlling privacy policies?

Additionally, we knew that the most critical problem was whether uncertainty about these issues would prevent people from appreciating the value of the tools we were giving them. If it did, all the advantages of the solutions would be useless, and the product would fail.

Strategies for product success. To ensure that the product would succeed, we devised a set of strategies to cope with the uncertainties we had identified, as outlined in the following sections.

Displaying appropriate information. Our research substantiated the early decision to prevent the display of private information, so we designed the product to follow certain guidelines: (1) The system would use only public information from enterprise data sources about people, selected by the system administrator, to populate the person profiles. (2) Though it might have access to private data, such as e-mail, the system would create categories only from public data. (3) By default, the system would notify users of any affinities that it had detected for them and require their review and approval before publishing those affinities. It would hide any information about

rejected affinities. (4) The system would respect any information security policies already in place for the public data sources. (For related issues, see the section “Respecting existing security policies” that follows.)

Deriving knowledge using appropriate information. As mentioned earlier, people were concerned about how the system might use their private information, even if not visible to others, to infer knowledge about them. To alleviate those concerns we designed the system to create affinities between users and public categories only. It would not generate affinities based on any information discovered in private data. Although the system might use information from private e-mail data to modify affinity scores, by default it would need the permission of the e-mail data’s owner (that is, any person—sender or recipient—who has a copy of the e-mail) to do so. The system would not reveal any information about the e-mail data itself.

Anticipating uncertainty and explaining policies. It was clear to us that we wanted to answer the “uncertainty” questions we had identified through our research and detailed in the earlier section “Identifying privacy issues.” We chose to design the product so that, as much as possible, it addressed these questions directly in the user interface. First, we wrote extensive explanations for the anticipated issues and referenced them in the product’s help system. Then we designed the user interface to include help links and in-place documentation to clarify the features where they appear. More information about the interface details is given in the section “Designing the user interface.”

Setting reasonable expectations. We designed IDS to be dynamic. By opening, forwarding, and linking documents, users express their opinions about the usefulness of those documents. As the system senses these interactions, it recalculates the metrics used for affinity strengths and document and category values. Over time, this tends to smooth out outlying values and increase the accuracy of the scores.

The metrics system does not ensure the precision of any particular instance of a document or category value or of affinity strength. Such values are, of necessity, determined by the particular circumstances of the data and the domain of the enterprise.

To give users a better perspective on the important part their judgment plays when using the system, the

product design team wanted not only to react to their expectations, but also to anticipate and clarify their questions by choosing our terminology appropriately.

Rather than claiming that it will come to particular conclusions, LDS documentation and instructions attempt to quantify the characteristics of interactions between users and categories of information. For ex-

**We designed ways
to help users control
the creation of affinities
to suit their needs.**

ample, we use the term “affinity” rather than “expertise” when talking about a user’s relationship to a certain subject, in the belief that, because the system creates an affinity from users’ real interactions with real data, it is a good indicator of expertise. Systems that rely on their users to manually update records of their skills or to “vote” on the skills of others are arbitrary, static, or both, and find their precision to be transient.

Understanding these concepts helps users realize that an assessment of expertise is an educated suggestion, and that they and others still play a large role in interpreting the information conveyed by the system. Expertise is still in the eye of the beholder.

Finally, we wanted to justify users’ expectations by making affinity judgments flexible. To that end, we designed ways to help users control the creation of affinities to suit their needs. This is described further in the section “Providing choice, flexibility, and user control.”

Respecting existing security policies. Through our research on Internet privacy issues and our prior experiences in developing other applications, we understood the importance of security in various forms.

To properly gauge the interactions between a particular user and documents and categories, LDS must know who the user is. For this reason, the system requires each user to verify his or her identity by logging in.

The system must also respect existing data access rights. For example, LDS must not override the doc-

ument security provided by the operating system or other applications. Knowing the identity of the user, LDS does not show in the Knowledge Map any secure document for which the user does not have access rights.

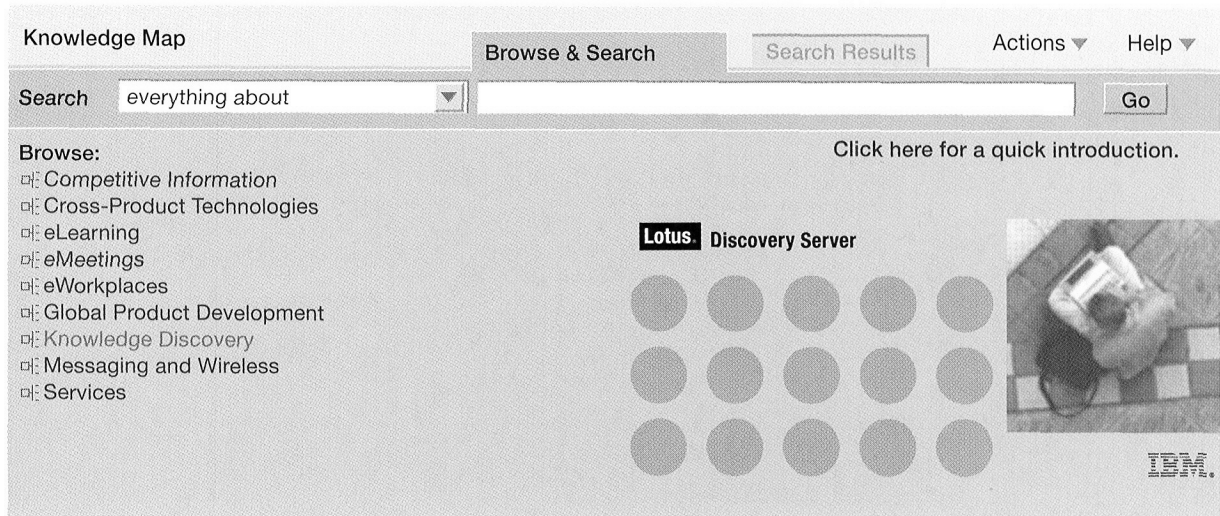
The system does not expose concepts that are not germane to the enterprise. LDS creates its category tree by using the concepts found in public data from sources provided by the system administrator. It also creates the affinities between categories and people by using that public data. Optionally, users can choose to let LDS look for affinities by using their e-mail data, but the system bases those affinities only on concepts derived from the public data, not on categories detected in the e-mail.

As an additional protection, LDS provides the Knowledge Map Editor user interface for the renaming, reconstituting, or removing of any categories they deem inappropriate or erroneous. See Figure 1 for the Knowledge Map home page and Figure 2 for an example of the Knowledge Map, second level.

Providing choice, flexibility, and user control. The features we created to give users control over the use of their private information are affected by the system administrator’s choices as set in the Discovery Control Center. By default, LDS provides settings that give users a maximum of choice, and we encourage customers to retain them. The default and alternate settings are described in the section “Administration.”

Affinities. When the system calculates that a user’s connection to a particular category exceeds the threshold value, it sends an e-mail notification to that user. The message explains why affinities are useful and contains a hyperlink to his person profile. The profile contains a message explaining the proposed affinities for the user to review. The user clicks a button to edit the profile and sees a list of one or more categories representing proposed affinities. Each category is also hyperlinked. If they wish, users can click the link to open the Knowledge Map browsing interface at the specified category. This allows a user to review the category in its context—its place in the category hierarchy, documents conceptually connected to it, other users who already have an affinity for it, and so forth. The list of proposed affinities in the profile provides choices that let the user decide whether to approve each proposed affinity, reject it, or defer the decision. Figure 3 is an example of a

Figure 1 Knowledge Map home page



person profile. Figure 4 shows the person profile in edit mode.

When the user approves an affinity, the system adds the user's name to the Knowledge Map list of people with affinities for that category. It also adds the category to the list of affinities in the user's profile. If the user rejects the affinity, the system does not publish it and never proposes it again.

If a user later decides to publish a rejected affinity, or to publish an affinity that the system has not proposed, he or she can use a menu choice in the Knowledge Map to declare an affinity for a category, and the system will publish it.

If a user decides at any time to stop the publication of an affinity, he or she can edit the profile and reject the affinity, after which the system will remove that name from the Knowledge Map list of people with affinities for the category. It will also remove the category from the list of affinities in the user's profile.

Recognizing that user control is important, the LDS design includes two manual mechanisms for creating affinities. The first allows any user to declare a personal affinity, and the second allows an approved manager to declare an affinity for someone else. For a self-declared affinity, the system initially assigns the lowest affinity strength because an "unqualified"

opinion, rather than system "evidence," is the source of the affinity. For a manager-declared affinity, the system assigns the highest affinity strength because it is based on a "qualified" opinion. This provides a way for managers to "declare" an expert for a category, which is especially important in the early stages of product deployment, when little evidence may have been accumulated. In both cases, the LDS metrics system adjusts the affinity strengths over time according to the real interactions between users and information, just as it does with affinities that it has discovered.

Affinity assignment policy. To provide the means to change system privacy policies to suit particular requirements, the Discovery Control Center contains a mechanism that lets administrators override the default settings for user approval and e-mail analysis. As indicated earlier, the default policy dictates that users be notified of, and approve, any proposed affinities, before such affinities are published.

Changing the default settings is illegal in some countries where users must be notified of information about them that a software system has collected, and must approve the sharing of that information. The LDS white papers, administration guide, and FAQs (frequently asked questions) all note the presence of the system options and require that a corporate privacy authority ensure compliance of the LDS af-

Figure 2 Example of Knowledge Map, second level

The screenshot shows a web interface for a Knowledge Map. At the top, there are tabs for 'Browse & Search' and 'Search Results'. A search bar contains the text 'everything about' and a 'Go' button. Below the search bar, there is a breadcrumb trail: 'Browse: Home>Knowledge Discovery'. On the left side, there is a list of subcategories under 'Domino Extended Search', including 'Domino.Doc', 'Lotus Discovery Server', 'Lotus Web Content Management System', 'Lotus Workflow Project Information Sage', and 'KM Competitive'. The main content area is divided into two sections. The first section is titled 'Documents About (25+)' and contains a table with columns for 'Value' and 'File Type'. The second section is titled 'People Who Know About (11)' and contains a table with columns for 'Affinity' and 'Job Title'.

Documents About (25+)	Value	File Type
LDS API Test Plan	3	Lotus Notes
(press release) New Enterprise Software Harnesses the Power of Social Networks for Corporate Portals and Intranets	2	Lotus Notes
Discovery Server Customer Council Milestones IV	2	Lotus Notes
Skills Inventory - Lotus Consulting/LPS	1	Lotus Notes

Documents Displayed: 1-25 of 25+ Next

People Who Know About (11)	Affinity	Job Title
■ John Ritsko/Watson/IBM	52	Distinguished Engineer
■ David Seidman/Watson/IBM	24	Senior Program Manager
■ Andy Schirmer/Westford/IBM	23	Product Design
● Shirley Ackerman/Watson/IBM	20	Chief Architect

finity policy settings with applicable governmental and industrial rules.

In addition, by default, there is no limit on the time a user can take to decide whether to publish an affinity. Because the publishing of affinities is an important part of the knowledge made visible by LDS, it is undesirable to let potential affinities languish. To prevent this, there is an optional time-out setting that lets administrators automatically publish pending proposed affinities that users have not approved or rejected by the end of a waiting period. The administrator determines the length of the waiting period, and the system automatically generates reminder e-mails to warn users when they are close to the deadline.

For various reasons, administrators may even want to override the user approval process entirely, and

there is a setting to force the publishing of affinities as soon as the system discovers them. This is particularly useful during the initial deployment of LDS, perhaps using test or limited data, to allow the administrator to test and alter affinity threshold settings to best effect.

It is important to note that the administrative policy settings never remove a user's ability to declare a new affinity or to reject an existing one.

E-mail privacy. While editing their profiles, users can also choose whether to allow the system to compare the concepts in their private e-mail data with the categories derived from public data. This setting is off by default, but if a user turns it on, the system can use this additional important information to propose new affinities.

While this setting is on, the system will periodically reevaluate the user's e-mail to look for additional affinities in new messages. Users can turn the setting off or on again at any time.

E-mail mining policy. During our research, some design partners told us that their users' e-mail, rather than the enterprise's public data, represented most of their enterprise's knowledge. For this reason, the option to analyze users' e-mail for affinities was especially important. When e-mail analysis is essential, the administrator can set an override that allows it by default.

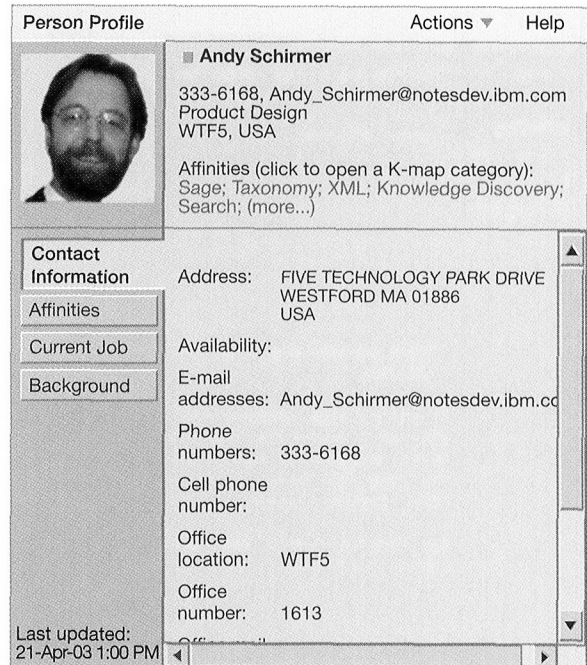
Limiting information. To protect the enterprise as a whole, we designed the product so that the enterprise maintains control of the metrics data created by the system. LDS transfers no information to any entity outside the enterprise.

Ensuring accuracy. Because the accuracy of inference is a primary concern of the users of a knowledge management system, we made the LDS metrics as precise as we could. However, knowing that there would be many circumstances we could not anticipate, and that circumstances would vary among enterprises, we planned the LDS metrics functions so that the algorithms are configurable. The complex metrics that compute activity levels, affinity strengths, values, and so forth, use a number of individual component metrics, each of which has a weighting factor. The characteristics of an enterprise's particular knowledge domain or culture may require that the system administrator optimize the algorithms by adjusting the component weighting. In addition, it is possible to plug in custom algorithms to further modify the complex metrics.

Administration. System administrators can control many elements of LDS behavior by using the Discovery Control Center (see Figure 5). Our research showed that, in addition to the user protections described earlier, administrators would need the ability to tailor the privacy policies of the system to fit the needs of their enterprise. They would also want certain general protections for their enterprise. As mentioned earlier in the section "Affinity assignment policy," one such protection is provided by the Discovery Control Center mechanism that lets administrators override the default settings for user approval and e-mail analysis.

Another Discovery Control Center service is the Metrics Report, which we designed to help admin-

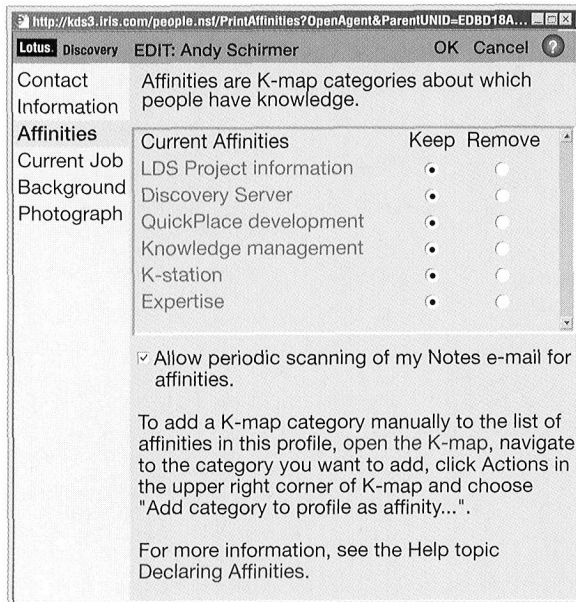
Figure 3 Example of Knowledge Map person profile



istrators assess and manage system use. It provides a level of detail about categories, documents, and people beyond what the user interface, the Knowledge Map, shows. Because of the greater sensitivity of the information presented in those reports, we restricted their use to administrators. Beyond that, we made some purposeful decisions to protect individual privacy by limiting the type of information that the reports produce, primarily by using aggregate values rather than low-level metrics reflecting information on individuals.

Demonstrating the value of LDS. In great part, LDS is exciting because of its ability to evaluate the strength of people's connections to concepts by analyzing their interactions with information. It is also dynamic—it continues to reveal new categories and affinities as new information and people come into it. Affinities increase or decline over time, based on users' interactions with categories and documents. The system keeps things up to date without requiring users to manually adjust the relevance values of their skills. Such benefits make it much more valuable than other more manual, less thoughtful systems, which pepper users with keywords and nag them to revise the records of their skills and vote on

Figure 4 Knowledge Map person profile in edit mode



the perceived value of documents. Such techniques interfere with the users' true goal—to get their work done—and quickly overwhelm users' tolerance and willingness to take part.

Because of the ways LDS infers knowledge (described in "The Lotus Discovery Server" and "Setting reasonable expectations"), LDS works better as more people use it. To encourage user participation, it is important to properly convey the depth of the benefits of LDS to its potential user audience. One way to do that, of course, is to document the features, which we did by writing a product methodology guide containing recommendations, process guidelines, and deployment scenarios. These materials, and the extensive help system, convey a clear picture of the system's potential benefits.

More compelling evidence of the system's value comes when users actually find it helping them with their work. The ability of LDS to display associated documents and people for any category that the user is browsing is a tremendous advantage over systems that require users to search for information and to understand the rules of searching before they can do so. Using affinities to find an authority on a subject for which a user needs help is a great way to demonstrate the benefits of publishing affinities.

The more the system helps people, the more they use it. The more people use it, the more accurate, and thus helpful, are its metrics.

Designing the user interface. Working through the issues detailed earlier guided our user interface design. Aside from simply providing places to expose the features, we wanted to make the user interface design part of our strategy for setting users' expectations and answering their questions about privacy issues.

To achieve this, we included in our user interface design plan additional goals that would support our strategies to address those privacy issues:

- Explain what the product does, and why, to increase understanding and set expectations.
- Make privacy safeguards and preferences salient, to inform users about the choices they can make to control information concerning them.
- Demonstrate the value of product features to show what problems they solve and why participation is desirable.

Following are some examples of aspects of the LDS user interface design intended to meet those goals.

Knowledge Map. In the Knowledge Map browsing interface (see Figures 1 and 2), users can traverse the category tree. For any particular category, the Knowledge Map displays a list of documents classified in the category and a list of people who have an affinity for the category. The list of people is sorted in descending affinity strength, so that the users with the strongest indications of expertise appear first. This not only helps users find likely contacts more quickly, but also reinforces their understanding of the affinity concept. In turn, this demonstrates the value of the feature and, by extension, the value of users approving their affinities' publication.

On the home page of the Knowledge Map, there is a hyperlink that brings up an introduction to LDS concepts and Knowledge Map features and behaviors.

A help button is displayed in the Knowledge Map at all times, which users can click to bring up the comprehensive, hyperlinked LDS help system described below in the section "Help."

Person profiles. Our user assistance writers also collaborated with the interface designers to craft in-place documentation to guide users through the pro-

cess of reviewing and approving affinities during the editing of their person profiles.

When users edit the affinities section of their profiles, in-line help text explains what affinities are and why they are useful. It shows the choices available for approving or rejecting affinities and explains the consequences of each. Also visible during profile editing is the e-mail analysis option. Again the text explains the purpose of this feature, as well as why it is safe and desirable. A help button is displayed in the profile, offering contextual help about working with and editing a profile.

Discovery Control Center. In the Discovery Control Center (as shown in Figure 5), the system administrator can adjust the system privacy policies as described above. The controls help to clarify the available privacy policy choices. The administrator can also modify the content of the three types of automatic e-mails that the system sends to notify users about affinities.

Help. To increase users' understanding of the product and set their expectations, our user assistance writers created comprehensive explanations and examples in the product help system. Figure 6 shows the help "overview" screen, and Figure 7 shows an example of help on specific aspects of privacy protection. They provided not only the usual descriptions of product features, but also included conceptual overviews and concise descriptions of how various features touch on issues of privacy, such as the use of e-mail in helping to determine affinities, as shown in Figure 7.

Summary

The Lotus Discovery Server is a knowledge management product with tremendous potential to enlighten users about the scale and scope of their information and to help them locate and assess people of authority in areas of their interest. Our work showed that a successful knowledge-management system must reveal relationships between people and other entities without sacrificing trust. We learned about a number of privacy issues particular to a knowledge management product. These led us to a set of strategies, which we effected in the design and implementation of the product's features and user interface. Among the essential strategies were displaying and deriving knowledge from appropriate information, explaining privacy policies to users and administrators and

Figure 5 Discovery Control Center

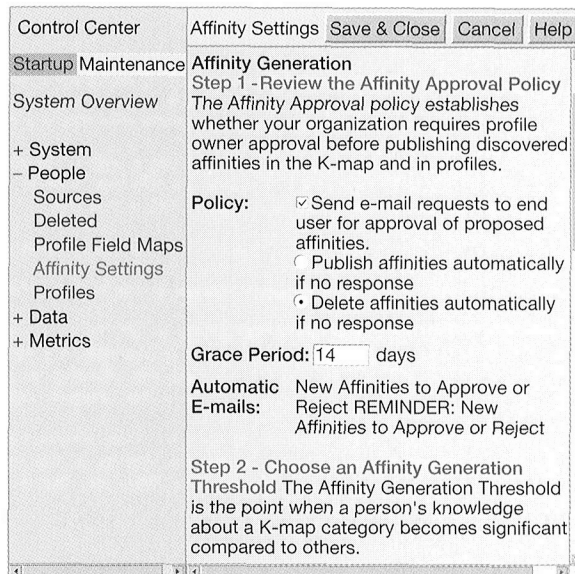
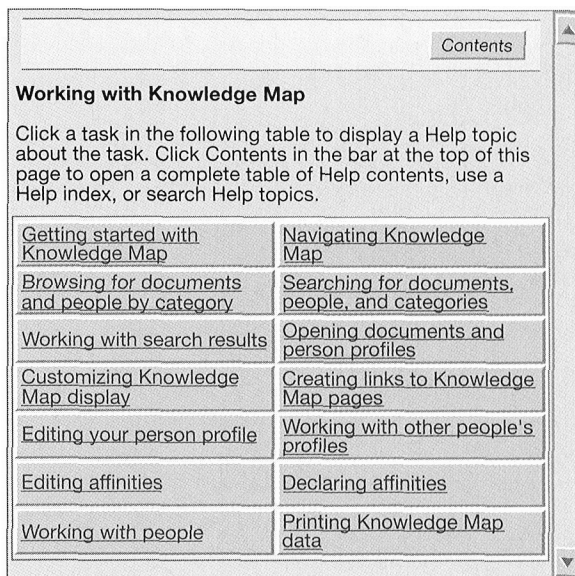


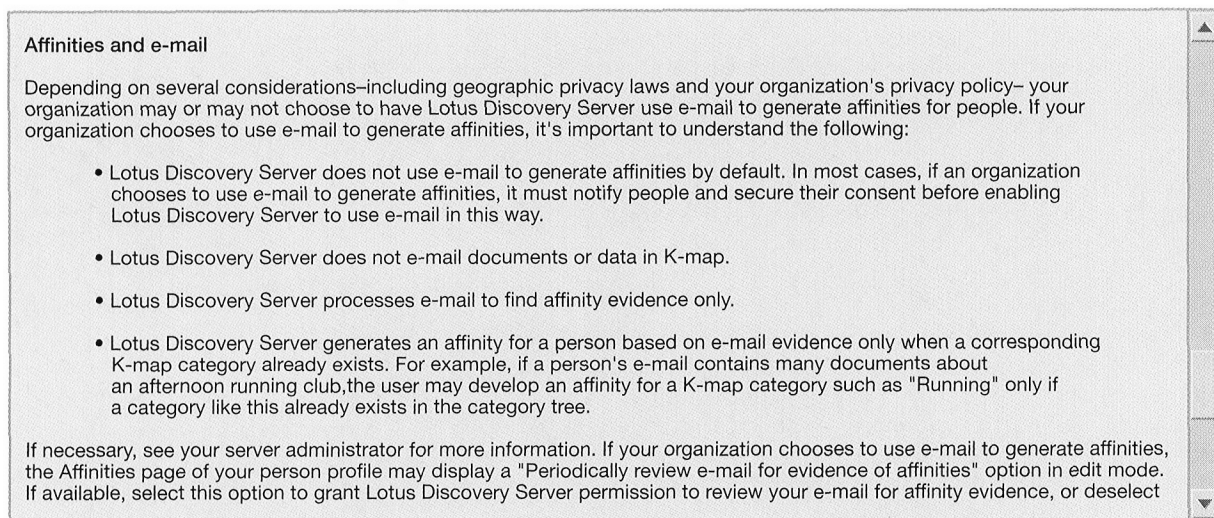
Figure 6 Knowledge Map help: Overview



giving them control over their information, and respecting existing security policies.

We believe these strategies helped us to protect the product's users, to allay their concerns, and to ad-

Figure 7 Example of help on privacy specifics



vance the value of the product in a way that encourages them to use it effectively. Since the product first shipped, we have continued to learn about users' experiences with it in the real world, exploring ideas for improvements and enhancements that better achieve knowledge management goals while still respecting our customers' privacy concerns.

Acknowledgements

For their generous help reviewing this article, I would like to thank Karen Bjorkman, James Goodwin, Scott Eliot, Lauren Wendel, and Majie Zeller.

I would also like to acknowledge the contributions of Dave Newbold (Lotus Discovery Server project leader), Cindy Regnante (co-product designer of the Discovery Server), Mary Carbonara (knowledge-management user-assistance writer), Dave Kajmo and Lauren Wendel (Lotus Discovery Server product managers), and Harriet Pearson (Chief Privacy Officer for IBM).

*Trademark or general trademark of International Business Machines Corporation.

Cited references

1. Platform for Privacy Preferences (P3P). See <http://www.w3.org/P3P>.
2. IBM Privacy Practices on the Web. See <http://www.ibm.com/privacy>.
3. R. Nelson and A. Ladd, *The Chief Privacy Officer: Pricewa-*

terhouseCoopers Surveys an Emerging Role, PricewaterhouseCoopers, Inc. (2000). Available at <http://www.pwcglobal.com/extweb/manissue.nsf/DocID/BD73BCD4420B713E85256A640048E4C0>.

4. Enterprise Privacy Architecture. See <http://www.ibm.com/services/security/epa.html>.
5. Tivoli Privacy Manager for e-business. See <http://www.ibm.com/software/tivoli/products/privacy-mgr-e-bus>.
6. IBM Privacy Institute. See <http://www.ibm.com/research/privacy>.
7. Autonomy Corporation plc. See <http://www.autonomy.com/Content/Legal/PrivacyPolicy>.
8. Tacit Knowledge Systems, Inc. See http://www.tacit.com/solutions/what_can_tacit_do/discover_activity.html.
9. Lotus Discovery Server. See <http://www.lotus.com/products/discserver.nsf>.
10. W. Pohs, G. Pinder, C. Dougherty, and M. White, "The Lotus Knowledge Discovery System: Tools and Experiences," *IBM Systems Journal* **40**, No. 4, pp. 956–966 (2001).

General references

BBBOnline Privacy Program. See <http://www.bbbonline.org/privacy/index.asp>.

Code of Fair Information Practices. See http://www.epic.org/privacy/consumer/code_fair_info.html.

Electronic Privacy Information Center (EPIC). See <http://www.epic.org>.

Georgetown Internet Privacy Policy Survey (1999). Available at <http://www.msb.georgetown.edu/faculty/culnanm/gippshome.html>.

Junkbusters Corporation. See <http://www.junkbusters.com>.

Organization for Economic Cooperation and Development, Information Security and Privacy. See <http://www.oecd.org/EN/home/0,,EN-home-43-nodirectorate-no-no-no-13,00.html>.

Privacy & Human Rights 2002: An International Survey of Pri-

vacancy Laws and Developments. See <http://www.privacyinternational.org/survey>.

Privacy International. See <http://www.privacyinternational.org>.

Privacy.org. See <http://www.privacy.org>.

Privacy Rights Clearinghouse. See <http://www.privacyrights.org>.

T. H. Davenport and L. Prusak, *Working Knowledge: How Organizations Manage What They Know*, Harvard Business School Press, Boston, MA (1998).

TRUSTe. See <http://www.truste.org>.

Accepted for publication April 22, 2003

Andrew L. Schirmer IBM Software Group, 4 Technology Park Drive, Westford, MA 01886 (Andy_Schirmer@notesdev.ibm.com).

Mr. Schirmer is a senior software designer and team leader for the User Experience Roadmap at Lotus. He earned a B.A. degree in history from Carleton College in 1972. He joined Lotus in 1987 and was instrumental in the development of several products as a senior software engineer. Mr. Schirmer has been a leader in product design for 13 years, working on projects in knowledge management, social networking, Web services, and office applications. He has filed over 30 patent applications for user experience and knowledge-management innovations. He is a member of the Association for Computing Machinery.