Knowledge Management on a Global Scale

Philip Klahr

Inference Corporation
100 Rowland Way
Novato, CA 94945
klahr@inference.com

Abstract

Building and using knowledge bases have become widespread in the customer support area, where the need is great and the potential benefit is high. While customer support has become a focal point for knowledge creation and use, other departments in an organization are now contributing to, and retrieving knowledge from, the knowledge bases being developed. These knowledge bases are becoming more corporate-wide repositories, with some companies even making the knowledge directly available to customers and consumers. CBR has become a useful tool to provide the representational framework and the search mechanism for these knowledge bases. As companies desire to create and distribute these knowledge bases globally, a number of key issues in authoring, distributing, localizing, maintaining and managing the whole process have been successfully addressed.

Global Customer Support

Over the past few years, there have been a number of companies developing knowledge-based systems in the area of customer support (Acorn & Walden 1992; Allen 1994; Nguyen, Czerwinski, & Lee 1993; Hislop & Pracht 1994). In this environment, customers call into a support center to ask questions about the company's products and services. Knowledge-based systems serve to provide a company's support agents with knowledge about the companies products, so the agents can answer customer questions, or more effectively solve customer problems.

Since many of these companies are global, there has been increasing interest on building and deploying knowledge bases on a global scale – both to leverage the knowledge and expertise that is distributed around the world, and also to make case bases available to all the regional support organizations to solve customer problems consistently worldwide (Borron, Morales, & Klahr 1996).

The principal AI-based technology that has been used in these customer support efforts is Case-Based Reasoning (CBR). "Cases" are used as the representational framework to capture knowledge of customer issues, problems and solutions, and general queries that customers have asked in the past. Providing a CBR search engine to end users allows them to find previous solutions that may help solve a particular customer's problem.

The financial benefits can be (and have been) enormous in terms of reduced escalations (less transferring of calls to more experienced technicians), reduced field visits by technicians (by solving more of the problems remotely on the phone), and reduced cost per call. Equally important are the benefits derived from improved customer satisfaction in having problems solved correctly on the first call. Improvements in the quality and consistency of customer support have enabled companies to improve customer loyalty, and differentiate themselves from their competitors in the marketplace.

As a result, the area of customer support has become a focal point for knowledge development and management. It is an area that is very costly to operate (many computer companies have tens of thousands of calls per day), but essential to a company's competitive positioning. There is an immediate need, and large payback in applying knowledge in customer support. Given this urgency, customer support organizations are developing a repository of knowledge, and defining mechanisms for effectively creating and managing knowledge.

Expanding Beyond Customer Support

Given the accumulation of knowledge centered in customer services, there are now companies that are broadening both the creation and usage of the knowledge beyond customer support.

Most notable in this activity are the product development

organizations which are now creating knowledge bases (again as cases for CBR reasoning and retrieval) as part of the product development lifecycle. Thus, product development is responsible for creating and releasing new products, creating and releasing documentation, and creating and releasing knowledge bases.

The main recipients for the knowledge bases are the customer support agents who, once a new product is launched, must support customers using the new product. Typically there is a transition of ownership of the knowledge bases – from product development to customer support, who then maintain and expand the knowledge. Part of the "sign-off" of the knowledge bases coming from product development is a quality check by the customer support organization to ensure the knowledge is satisfactory for their use, and for their subsequent ownership.

On the usage side, other departmental organizations can benefit from the knowledge accumulating in these case bases. Sales organizations, for example, can access the knowledge bases to better understand customer problems and questions. Marketing, as another example, can access the knowledge to better determine requirements for subsequent new releases or new products, based on the knowledge base.

As a result, these knowledge bases are being made available on internal corporate intranets and/or on the external Internet. In the latter case, there has been significant benefit achieved by companies making their knowledge directly available to customers over the World Wide Web (see, for example, Broderbund Software's web site http://www.broderbund.com in the technical support area). Customers are now accessing case bases and solving their own problems through this knowledge sharing process.

Also, there are now several products that have taken case bases and packaged them for retail sale on CD-ROMs (see for example, Symantec's PC Handyman product and CyberMedia's First Aid 97 product) for direct use by PC users.

CBR as a Knowledge Management Tool

CBR, as mentioned above, has become a common tool for implementing knowledge systems for customer support. There are a number of reasons CBR has become popular:

1. CBR provides a representational framework that maps nicely to the reasoning process technical engineers go

through to troubleshoot customer problems. This process is one of referring to past similar situations.

2. CBR provides inexact matching and inexact reasoning to locate and retrieve knowledge that may be relevant. Matching can take into account missing information, and, in many situations, can operate with some degree of incorrect or conflicting information.

3. Cases typically contain a description of the case and a set of features (e.g., symptoms) that define the case, as well as the appropriate action to take in that situation. Based on the initial problem/query description, the system can retrieve similar cases and provide users with appropriate questions to ask (the features) to help further narrow the search. This interactive dialog provides a user with guidelines on what questions to next ask. CBR can be proactive.

This latter point is particular important when comparing CBR to text retrieval tools in the area of customer support. With text retrieval, users are provided little guidance as to how to formulate a query, or how to refine it. Users typically need to know what they are looking for, and how to access the knowledge (e.g., what keywords to use). CBR can alleviate this problem through the question-asking dialog process. Often CBR and text retrieval approaches can be complimentary. For example, when specific solutions are known, they can be entered in cases; where there is research required, text retrieval can provide a vehicle to access technical notes or documents.

The advantages of a proactive dialog to help define the specific problem are even more pronounced when the user is a customer, someone typically unfamiliar with the problem domain and with the knowledge base. They often cannot even describe their problem correctly. Providing a simple dialog to elicit problem symptoms and features, makes the diagnostic process quicker and much easier for the user.

Global Knowledge Management

I define knowledge management as

> the process of creating, refining, maintaining, monitoring and extending knowledge, and providing delivery mechanisms for effectively accessing and retrieving knowledge.

Knowledge management becomes even more complex within a global framework. There are a number of key issues that companies need to address in global knowledge

capture and management (Klahr 1996). I summarize them as follows:

- *distributed authoring* – capturing expertise and experience from customer support organizations distributed around the world; embedding that experience in case bases.

- *knowledge distribution* – distributing knowledge bases, and updates, and having the local sites incorporate changes and additions.

- *localization* – being able to modify global knowledge and add local knowledge; customizing to local needs (e.g., language translation, local actions, local questions, local features).

- *maintenance* – updating and adding both global and local knowledge.

- *management and organization* – managing a global project and establishing (and enforcing) processes for the above activities; ensuring funding, appropriate staffing, accurate scheduling and other project management activities – on a global scale.

Each global company addresses these issues in different ways depending on their business objectives, operational infrastructures, and technical requirements.

As an example of one particular effort, consider the global case base developed by Reuters (Borron, Morales, & Klahr 1996). The principal roles and responsibilities that were defined and used in that project (as shown in Figure 1) include the following:
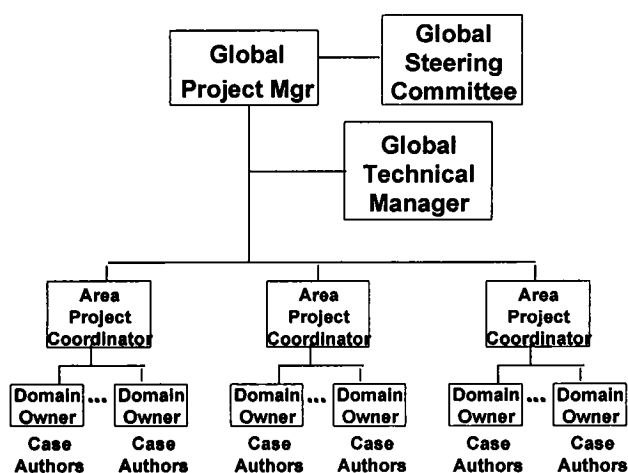

Figure 1. Example Global Management Structure

The *Global Steering Committee* consists of the heads of customer support from each of the global areas, plus a senior customer support representative from the corporate headquarters. The Global Steering Committee meets periodically (typically quarterly) to review project milestones, issues and directions. Biweekly phone conference calls provide interim status updates and address immediate issues and needs.

The *Global Project Manager* leads the project and coordinates the development and deployment efforts in the various global regions. Specific responsibilities include;

- completion of all milestones and deliverables
- project scheduling and monitoring
- communication, both up to the Global Steering Committee and down to all project team members
- planning and implementing a technology transfer program for all global regions.

The *Global Technical Manager* is responsible for supervising all technical aspects of the project, including:

- developing a single, standard global style guide for the knowledge base
- leading all software development efforts, e.g., utilities to support global procedures
- supervising technical deliverables of external consultants
- approving case bases for global distribution
- maintaining a central library of global cases
- distributing global case bases and/or updates to the various regions worldwide
- technically supporting the regions.

The *Area Project Coordinators* are responsible for the overall operation of their local region, including:

- managing the hardware and software infrastructure locally
- authoring knowledge bases (building case bases and supervising the domain owners and case authors)
- training authors and end users
- testing and approving locally-built knowledge bases
- transmitting knowledge bases to the central knowledge repository
- receiving global case bases and updates from the central repository and implementing them locally.

The *Domain Owner* is the person responsible for a particular case base. Each case base contains cases relative to a particular product. This segmentation of knowledge based on product is one used by many companies. Customer calls usually focus around the particular product

84

the customer is using, and problems or issues associated with that product. A particular product case base is assigned for development to the particular region that has the most expertise in the product. The Domain Owner is the individual assuming ownership (content, delivery, maintenance) of a product case base.

In some cases the Area Project Coordinator (APC) can also be a Domain Owner for one or more product domains. These added responsibilities for the APC can vary, and are based on the APC's workload and domain expertise.

A Domain Owner can use multiple *Case Authors* to help author the knowledge in a particular case base. Again this is dependent on workload responsibilities and product expertise. Thus, many configurations are possible: an APC can fully author a small product case base, or can supervise a Domain Owner who has several Case Authors to author the knowledge. The Domain Owner is ultimately responsible for the knowledge content and organization within the assigned case base. That individual typically accumulates cases from the Case Authors, looks for redundancies, and ensures consistency and style.

Summary

The Reuters project is one example of a company leveraging expertise that exists around the world and putting that expertise into a single knowledge base, which is then distributed to all its global sites. They have established a global organization and infrastructure to support the project and technical requirements. CBR is the technology used to represent the knowledge, and the mechanism for search and retrieval. The customer support area has become a focal point for knowledge creation and sharing.

References

Acorn, T. L., and Walden, S. 1992. SMART: Support Management Automated Reasoning Technology for Compaq Customer Service. In Innovative Applications of Artificial Intelligence 4, Proceedings of IAAI-92 (A. Scott & P. Klahr, eds.), 3-18. Menlo Park, California: AAAI Press.

Allen, B. 1994. Case-Based Reasoning: Business Applications. *Communications of the ACM* 37(3): 40-42.

Borron, J., Morales, D., and Klahr, P. 1996. Developing and Deploying Knowledge on a Global Scale. In Proceedings of the Thirteenth National Conference on Artificial Intelligence and the Eighth Conference on Innovative Applications of Artificial Intelligence, 1443-1454. Menlo Park, California: AAAI Press.

Hislop, C., and Pracht, D. 1994. Integrated Problem Resolution for Business Communications. In Proceedings of the Sixth Innovative Applications of Artificial Intelligence Conference (E. Byrnes & J. Aikins, eds.), 63-74. Menlo Park, California: AAAI Press.

Klahr, P. 1996. Global Case-Base Development and Deployment. In Advances in Case-Based Reasoning, Proceedings of the Third European Workshop, EWCBR-96 (I. Smith & B. Faltings, Eds.), Lecture Notes in Artificial Intelligence 1168, Springer, Berlin, 1996, 519-530.

Nguyen, T., Czerwinski, M., and Lee, D. 1993. Compaq QuickSource: Providing the Consumer with the Power of Artificial Intelligence. In Proceedings of the Fifth Innovative Applications of Artificial Intelligence (P. Klahr & E. Byrnes, eds.), 142-151. Menlo Park, California: AAAI Press.