

Practical Knowledge Management:

The Lotus Knowledge Discovery System



Practical Knowledge Management:
The Lotus Knowledge Discovery System

Wendi Pohs
with Gayle Thiel and Seth Earley

Edited by Marianne White



Practical Knowledge Management: The Lotus Knowledge Discovery System

Wendi Pohs with Gayle Thiel and Seth Earley, edited by Marianne White

Published by IBM Press™

Program Director, IBM Centre for Advanced Studies: Gabriel Silberman

IBM Associate Publisher: Sheila Richardson

IBM Press Alliance Publisher: David Uptmor, IIR Publications, Inc.

For information on translations or book distribution outside the United States or to arrange bulk-purchase discounts for sales promotions or premiums please contact:

IBM Press
c/o IIR Publications, Inc.
5650 El Camino Real, Suite 225
Carlsbad, CA 92008-7147
USA
760-931-8615
www.iirpubs.com/ibmpress

First edition

First printing: August 2001

ISBN: 1-931182-03-5

Printed in Canada. All rights reserved. No part of this publication may be reproduced in any form without prior permission of the copyright owner.

© Copyright 2001 by International Business Machines Corporation.
Portions © Copyright 2001 by IIR Publications Inc.

Every attempt has been made to provide correct information. However, the publisher, the editors, and the authors do not guarantee the accuracy of the book and do not assume responsibility for information included in or omitted from it.

The following are trademarks of International Business Machines Corporation in the United States, other countries, or both: IBM, the IBM logo, IBM Press, OS/2, and WebSphere. The e-business logo is a trademark of IBM and used under license.

Lotus, Domino, Freelance, Freelance Graphics, and Sametime are registered trademarks, and Discovery Server, Domino.Doc, K-station, the KM logo, Knowledge Discovery System, LotusScript, People Places and Things, and QuickPlace are trademarks of Lotus Development Corporation in the United States, other countries, or both.

Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

Microsoft, PowerPoint, Windows, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

Other company, product, and service names may be trademarks or service marks of others.

To my parents, Arnold and Constance Pohs;
to my brother Glenn, and to Jill, Lara, Jamie and Jake;
to my stepdaughters Meghan and Erin;
and especially to my husband, Mike Kilgore.



❖ Acknowledgments

They say that a dwarf standing on a giant's shoulders sees the farther of the two. In writing this book, I feel like I am standing on the shoulders of all the giants I've worked with in the information retrieval and computer industries. At the risk of sounding like an over-excited actor at the Oscars, I'd like to thank the following people:

Chapter 1: Seth Earley

Chapter 2: Scott Eliot, Stona Fitch

Chapter 3: Gayle Thiel, Seth Earley, the staff at Lotus Research, Michael Muller, Matthew Simpson, Chris Reckling, Kellie Hunter, Sara DeSouza

Chapter 4: Gayle Thiel, John Curtis

Chapter 5: Justo Perez, John McSweeney, Michael Cox, David Malone

Chapter 6: Kate Ehrlich, Lauren Wendel, Mary Carbonara, James Riel

Acknowledgments

Chapter 7: The Lotus IS Early Deployment Team: Peter Monahan, Kaye Kittredge, Barbara Brooker, Marianne Iauco, C.J. Frost, Dave Urban, Vaughn Rokosz, Steve Kelley

Overall editorial/production assistance: Marianne White, for making this happen, and Genevieve Thiers, Laura Sapelly, Judy Riessle

Iris Associates: The entire KDS development team: Dave Newbold, James Goodwin, Mim Francis, Andy Schirmer, Bob Barber, Chris Samoiloff, Chris Engstrom, Craig Wolpert, Cynthia Regnante, David Kajmo, Geno Spinoso, George Sprott, Jacob Kirk, Jaye Fitzgerald, John Curtis, Jonathan Griep, Julia Nizhegorodova, Justo Perez, Mark Keeler, Mark Rovelli, Michael Angelo, Michael Cox, Paul Wilson, Ronald Buell, Roy Goodwin, Sabrina Yee, Sheryl Jablonowski, Steve Bussell, Tom Lowell, Yaolei Hong, Jerry Wood, Dane Johnson, Paul McGrail, Susan Ricercato, Lauren Wendel, Jan Kenney

Lotus Development Corporation: Kathleen Murphy, Peter Orbeton, Dick McCarrick, Karen Bjorkman, Robin Rouse, Christopher Dougherty, Gary Pinder, Glen Kelley, Akiba Saeedi

Digital Equipment Corporation: Lois Lyman, Joyce Ward, Nancy James, the original Source Managers: Lydia Dehn, Gail Bosworth, Anne Stoma, Shirley Barron, William Whalen, Marianne White, and Annmarie McDonough, Laura Hunt, Denice Brait, Pam Morrison, Marlene Rockmore

Mathematical Reviews: William B. Woolf, Sandy Barth, Jean Gibney, Juan Rodriguez-Padro, Peggy Gerber, Fran Bryant, and SBW, wherever you are

Wendi Pohs
Westford, Massachusetts
May, 2001

❖ Contents

Foreword by Alan D. Marwick	xiii
Introduction by Dave Newbold	xv
About this book	xvii
Chapter 1: Knowledge discovery	1
Knowledge at Web speed	2
One small problem	3
Historical perspective	4
The modern problem	5
A confluence of developments	6
Search comes of age	7
Measurement and metrics	9
Digital breadcrumbs and micro-decisions	11
Putting it all together	11

Contents

Chapter 2: Managing knowledge with the Lotus K-station and the Discovery Server	13
The importance of knowledge in today's organizations	14
The quest for knowledge	15
The value of experience	15
Creating context.	16
Managing collective experience	17
A different approach to content management	17
Focusing on people	19
Putting the pieces together	20
Enabling effective knowledge management	20
The knowledge organization	21
Who works like this now?	25
Chapter 3: Communities and K-station places.	27
Communities: repositories of knowledge	28
The role of communities.	29
Communities of practice	30
Supporting successful communities	30
The Lotus K-station portal	31
Using the K-station portal	31
Creating shared places for CoPs	32
Making good use of shared places	32
Chapter 4: Capitalizing on full-text search.	35
From information retrieval to knowledge management	36
Reference and publishing	36
The Internet.	36
Spiders, indexes, and search engines	37
The Web user experience.	38
The query	38
The results	39
Directories.	39
The missing link	40
The Discovery Server: the user experience.	41
Benefits of the K-map	42
Narrowing a search	44
Simple queries, sophisticated results.	45

The Discovery Server: the search engine	45
Multipurpose spiders	45
The full-text index	46
The search engine	46
Chapter 5: Creating and managing meaningful taxonomies.	49
Benefits of a business taxonomy.	50
Effective information reuse	51
Thought leadership	51
Organization and cleanup.	52
Economies of scale	53
Full-text search enhancement.	53
Information discovery	53
The library model.	53
Controlled vocabularies	53
Taxonomies	54
Classification schemes	55
Fields and tagging	57
Existing categorical evidence	58
Creating taxonomies using Lotus Discovery Server tools	59
Using the K-map Builder	59
Using the K-map Editor	62
Using the K-map	63
Performing a knowledge audit	63
Selecting data repositories	64
Gathering external references.	65
Organizing the taxonomy	66
Testing the taxonomy	67
Training the taxonomy	67
Summary.	67
Chapter 6: Collecting and managing information about people	69
Finding experts	70
Mapping expertise	73
Components of the Discovery Server's people-finder	73
Mining e-mail.	75
Expertise and affinities	75
Roles and responsibilities	76
Communicating benefits	78

Contents

Privacy and trust	78
Sample privacy guidelines	79
Chapter 7: A Lotus case study	81
Planning	82
Gathering content	83
Building the taxonomy	84
Deployment	86
An impetus for content improvement	87
Large-scale deployment	88
Appendix: Sample taxonomies.	91
A software development taxonomy	91
An industry taxonomy	93
Glossary	97
Bibliography	103
Index	107
About the Authors	109



Foreword

Reviewing the draft of this book before writing this Foreword, I was struck by the way the Lotus Knowledge Discovery System integrates features that address some of the key challenges of knowledge management. The expertise location function helps in finding colleagues with tacit knowledge to tap. The integrated synchronous collaboration and virtual place creation will help you work with other people in a team (and the architecture of the K-station supports communities of practice). These important people-centric attributes balance the information-centric approach that is too often a surrogate for effective knowledge management.

But of course information is important. All of us know that the information we use in business is moving online. Actually, my profession, scientific research, was a leader in moving bibliographic and other scientific databases online in the days when 110 bits per second teletype was considered to be a hot box. Today, all companies are in the process of moving their vital information resources online.

At IBM Research, we have a focus on a fast-growing category of that information, “unstructured information,” that is designed for people to read but which technology can and should help us to use more effectively. For example, we

Foreword

would like people to be able to use text search to find helpful documents when they have a need for specific information.

One of the challenges of effective information use in knowledge management is to improve the usability of text search. The approach in the Knowledge Discovery System provides an innovative way to expand the notion of “relevance.” In other words, to include an element of quality (relevance) is the classic information retrieval measure of goodness. This measure can be made to influence a document’s position in the search results list.

It is worth pointing out that the quality measure used by the Knowledge Discovery System is an empirical one. In the absence of reliable, absolute measures of quality (the Greeks debated this), “usefulness” is an acceptable replacement, for which usage is a measurable surrogate. In effect, people (by choosing to use certain documents) influence the outcome of searches that their colleagues later perform.

As a researcher, I know that many other challenges are ahead as we invent and adapt technology to the human need to share and use knowledge. We know from the analysis in Nonaka and Takeuchi’s seminal 1995 book, *The Knowledge-Creating Company*, that organizational learning is supported by transformation of knowledge from tacit to explicit form, and back again in a spiral of sharing, reconceptualization, and use. Thus, for information technology to best support knowledge management, we need the balanced combination of information-centric and people-centric approaches described in this book.

Alan D. Marwick
IBM T.J. Watson Research Center
May 2001



Introduction

The Lotus Discovery Server was conceived as a tool for helping individuals discover content and expertise within their organizations. As with any tool, the results derived are determined in large part by the skill of execution. At Lotus, we have always stressed that knowledge-management (KM) technology plays an important, but partial, role in a successful KM solution. This book is part of our commitment to the human side of the solution. I hope that everyone involved in testing, deploying, and supporting a Discovery Server solution has the opportunity to read and understand this material.

As we all move forward into a world where there is more information than we can possibly process, and our attention becomes overwhelmed, catalogs can provide some needed context. I believe that content catalogs are the best way to give users context for evaluating and understanding a subject area. While creating a well-crafted taxonomy is a difficult task, the Discovery Server provides a powerful set of tools to assist you in understanding your organization's collected knowledge assets and to help you rationalize the content into an organized system. The weaving of content into such a taxonomy requires the coordination of the content owners and experts, the support of your IT department, and the consent of management. This is a challenging and ultimately very rewarding task.

Introduction

One of the most compelling aspects of the Discovery Server is that it can be useful in identifying expertise within an organization. Automating affinity generation and leveraging the content catalog categories resolves what has been an intractable goal for many organizations. The usefulness of the Expertise Locator is a function of the clarity and structure of the content catalog. Hence, our focus on giving you the tools and methodology to be successful.

Welcome to the community of Discovery Server users. So that we can make the product more responsive to your needs and so that we can continue to share our best practices, I hope you will share your experiences.

Thank you!

Dave Newbold
For the Discovery Server Development team
dnewbold@iris.com
<http://www.lotus.com/discoveryserver>

❖ About this book

Practical Knowledge Management: The Lotus Knowledge Discovery System describes a methodology for using a new, knowledge discovery software product from Lotus Development Corporation called the Lotus Knowledge Discovery System. Using several leading-edge technologies, the product systematically catalogs corporate expertise and information resources, personalizes and organizes knowledge for individuals and communities, and provides a place for teams to work, make decisions, and take action. It also creates a search-and-browse user interface, called a *K-map*, computes document value, and includes a Control Center and a K-map Editor for system administrators and taxonomy editors.

This book defines the Lotus Knowledge Discovery System concepts at a high level and then shows how each of the Lotus Knowledge Discovery System components can be used. It introduces the Lotus Knowledge Discovery System's K-map Builder and presents a corresponding methodology for classifying large collections of information. It also includes a case study with details about how the system was deployed internally at Lotus.

Developers will find more detailed technical information in the IBM Redbook *Inside the Discovery Server* (ISBN: 0738421901). Also, readers can refer to the Bibliography for additional resources.

Who should read this book

This book can be used by chief technology officers (CTOs), consultants, planners, practitioners, and administrators of various types interested in developing and supporting knowledge management solutions within an organization.

Chief technology officers will find information on the underlying Knowledge Discovery System technologies. Chapters 1 and 2 identify the generic building blocks required to deploy a knowledge management solution and subsequent chapters detail how to implement policies to support each component.

Consultants can adapt the book's methodology to various customer environments. Chapter 5 offers some basic principles about building corporate taxonomies. Chapter 6 explains the advantages of deploying effective people-finder systems and includes information about privacy issues.

Planners can consult the book to get an overview of all the Knowledge Discovery System components. Chapter 2 offers guidelines on how to staff a knowledge-management organization. Chapter 7 presents a case study that can be used to plan deployments.

Knowledge management practitioners and administrators should read the entire book before they deploy the Knowledge Discovery System. Because knowledge management solutions are larger and often more complex than more traditional applications, practitioners can use the book to set appropriate expectations among their users.

Hands-on users of the Lotus Knowledge Discovery System components, including corporate librarians and knowledge management evangelists, will also find practical information in the book. Chapter 6 describes how librarians can use their existing skills to create corporate knowledge management applications.

Features of the book include:

- Clear instruction for early adopters of knowledge management components
- Transfer of expertise to customers from the Lotus Knowledge Discovery System product team
- Fast and efficient deployment hints

Chapter summaries

Chapter 1, "Knowledge discovery," provides the big-picture view and describes the development team's vision. It introduces the Lotus definition of knowledge management and discusses problems you can solve using the Lotus Knowledge Discovery System. This chapter explains why the technology and the methodology are coming together now.

Chapter 2, "Managing knowledge with the Lotus K-station and the Discovery Server," defines the Lotus Knowledge Discovery System concepts at a high level and makes the business case for deployment. It is based on the user roles Lotus defined to scope the project. The chapter describes the concepts behind the Lotus Knowledge Discovery System components (the K-station and the Lotus Discovery Server) and proposes a knowledge management organization. Read this chapter to learn about the building blocks of a knowledge management infrastructure.

Chapter 3, "Communities and K-station places," covers communities of practice. It explains how a Lotus Knowledge Discovery System K-station feature called places allow these communities to work effectively. It also provides an overview of the K-station portlet components.

Chapter 4, "Capitalizing on full-text search," briefly describes a few of the Internet search engines, and makes the case for information discovery using the Lotus Discovery Server's full-text search model. It explains how the Discovery Server's search engine works to access information about people, places, and things in an organization.

Chapter 5, "Creating and managing meaningful taxonomies," describes the Discovery Server K-map Builder and the organizational skills required for effective taxonomy creation and maintenance using the K-map Editor. Corporate librarians, content managers, and information professionals should read this chapter.

Chapter 6, "Collecting and managing information about people," provides guidelines for successfully introducing a "people-finder" system into an organization. The chapter introduces the concepts of affinities and metrics processing, and discusses privacy and security issues.

Chapter 7, "A Lotus case study," is the story of how the Early Deployment group put it all together to deploy a working Knowledge Discovery System for the

About this book

IBM/Lotus sales force. This chapter presents the actual taxonomy editing and deployment recommendations made by the group.

In addition, the book includes a knowledge management glossary and an appendix containing two sample taxonomies that can be customized.