IBM

E-mail Archiving and Records Management Integration Solution Guide

Using DB2 CommonStore and DB2 Records Manager

- E-mail archiving type, storage model, and archiving policies
- E-mail record declaration, classification, and disposition
- End-to-end solution installation and integration

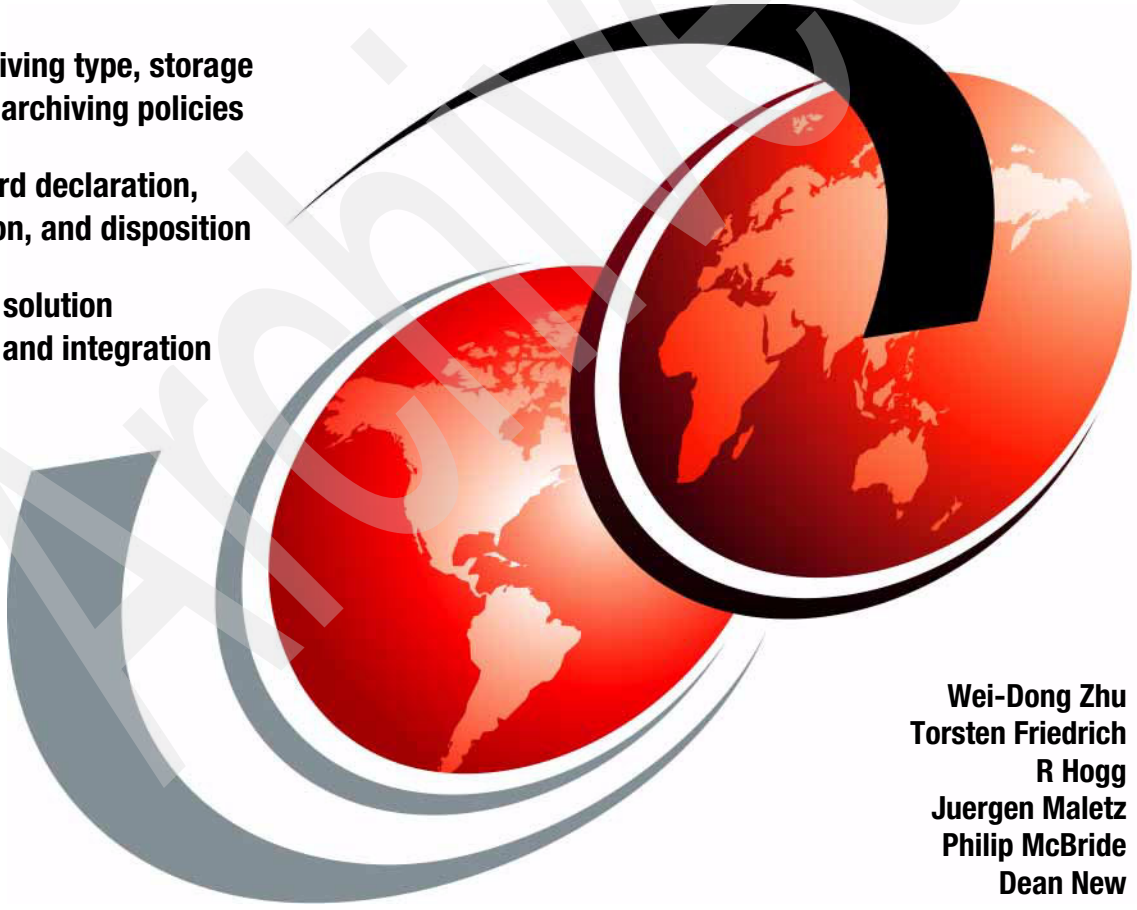

Wei-Dong Zhu
Torsten Friedrich
R Hogg
Juergen Maletz
Philip McBride
Dean New

Redbooks

ibm.com/redbooks

International Technical Support Organization

E-mail Archiving and Records Management Integrated Solution Guide Using DB2 CommonStore and DB2 Records Manager

January 2006

SG24-6795-00

Note: Before using this information and the product it supports, read the information in “Notices” on page ix.

First Edition (January 2006)

This edition applies to Version 8.3 of IBM DB2 CommonStore for Lotus Domino (program number 5724-B86). Version 8.3 of IBM DB2 CommonStore for Exchange Server (program number 5724-B85), Version 8 Release 3 of IBM DB2 Content Manager Enterprise Edition (product number 5724-B19), and Version 4 Release 1 Revision 2 of the IBM DB2 Records Manager for Windows (product number 5724-I58).

Contents

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:
IBM Director of Licensing, IBM Corporation, North Castle Drive Armonk, NY 10504-1785 U.S.A.

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:
This information contains sample application programs in source language, which illustrates programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. You may copy, modify, and distribute these sample programs in any form without payment to IBM for the purposes of developing, using, marketing, or distributing application programs conforming to IBM's application programming interfaces.

Trademarks

The following terms are trademarks of the International Business Machines Corporation in the United States, other countries, or both:

AIX®	ibm.com®	Redbooks (logo) ™
DB2®	Lotus®	Tivoli®
DB2 Universal Database™	Lotus Notes®	WebSphere®
Domino®	Notes®	z/OS®
IBM®	Redbooks™	

The following terms are trademarks of other companies:

Java, JDBC, J2EE, and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

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

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

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

Preface

Currently, under the U.S. Security and Exchange Commission rules, all banks and brokerage firms are required to keep their e-mail for three years. By July 2006, under the Sarbanes-Oxley Act, all public companies are required to comply with this rule.

Whether the companies are public or private, more companies need to manage their e-mail to support regulatory compliance, litigation, and corporate policy and to improve performance and productivity. The e-mail archiving and records management solution presented in this book will help companies to accomplish these goals.

This IBM® Redbook provides a general solution guide for the integrated e-mail archiving and records management solution. The integrated solution uses the following IBM products:

- IBM DB2® CommonStore for Lotus® Domino® V8.3 or
  IBM DB2 CommonStore for Exchange Server V8.3
- IBM DB2 Records Manager V4.1.2
- IBM DB2 Content Manager V8.3
- IBM DB2 Content Manager Records Enabler V8.3 (previously known as Records Manager Enabler or RME)

Part 1 of the book covers design and planning of the solution. In Chapter 1, we provide an overview of the integrated solution, the products involved, and their roles in the solution. In Chapter 2, we cover the basic concepts behind e-mail archiving, CommonStore features and functions, and its architecture. We address key areas to consider when planning and designing the solution from e-mail archiving perspective. In Chapter 3, we cover the basic concepts behind records management, Records Manager, and Records Enabler's features and functions. We also address key areas to consider when planning and designing the solution from an e-mail records management perspective. In Chapter 4, we focus on security of individual products and the implication on the overall integrated solution. In Chapter 5, putting everything together, we cover the design and planning of the entire integrated solution. System architecture and implementation sequence are addressed.

Part 2 of the book focuses on all major steps involved in installing and integrating all of the products for the end-to-end integrated solution. Chapter 6 covers the scenario using CommonStore for Lotus Domino in the Windows® environment. Chapter 7 covers the scenario using CommonStore for Exchange Server in a

Windows environment. Chapter 8 covers the CommonStore for Lotus Domino in the AIX® environment.

Part 3 of the book introduces some advanced topics for the solution. In Chapter 9, we examine the deployment of the integrated solution. In Chapters 10 and 11, we address Records Manager specific topics, including holding records, disposition of records, and record discovery.

This book is intended for IT architects and specialists who will be responsible in planning, designing, and implementing an e-mail archiving and records management solution. We focus on the important areas that are related to the overall end-to-end solution in this book. For product-specific information, we strongly recommend reading the existing products manuals in conjunction with this book.

The team that wrote this redbook

This book was produced by a team of specialists from around the world working at the International Technical Support Organization, San Jose Center.

Wei-Dong Zhu (Jackie) is a Content Management Project Leader with the International Technical Support Organization at the Almaden Research Center in San Jose, California. She has more than 10 years of software development experience in accounting, image workflow processing, and digital media distribution (DMD). Her development work in one of the DMD solutions contributed to a first time ever win for IBM of an Emmy award in 2005. She holds a Master's degree in Computer Science from the University of Southern California. Jackie joined IBM in 1996. She is a Certified Solution Designer for IBM DB2 Content Manager.

Torsten Friedrich is an IT specialist with IBM in Germany. He has six years of experience with IBM and holds a degree in Information Technology. Having worked on content management projects for three years, he specializes in e-mail archiving solutions. He delivers proof-of-concept and implementations for customers.

R Hogg is the Records Management solution executive expert with IBM US. With more than 10 years of experience in the field of document, information, and records management, he engages with customers and partners worldwide in analyzing, resolving, and deploying solutions.

Juergen Maletz is a senior consultant working for IBM in Germany. He joined IBM in 1986 with a degree in Information Technology. He worked on different

projects before joining the CommonStore development team. He is the team lead of the CommonStore for Exchange Server development team.

Philip McBride is a senior consultant with IBM in Australia. Originally from the UK, he has seven years of experience with IBM and holds a degree in Information Technology. He specializes in delivering solutions around document, records management, and workflow systems. Working with customers in either technical or business lead roles, he delivers proof-of-concept and implementations for customers in Australia and Asia.

Dean New is a Technical Services Associate in the IT department at Honda in Canada. He has eight years of experience in the IT field. He holds a three-year Computer Programmer/Analyst degree. His areas of expertise include Lotus Domino Administration and WebSphere® Administration. He has experience with IBM AIX, IBM DB2 Content Manager, IBM DB2 CommonStore for Lotus Domino, IBM DB2 Records Manager, and Records Manager Enabler.

Thanks to the following people for their contributions to this project:

John Dorak
Craig Kindell
Tracy Kong
Chris Lehman
Qing Lu
Jose “Hose” Martinez, Jr.
Ken Milsted
Lonnie Moore
Paul Schultes
Glen Walters
Emily White
IBM Software Group, Information Management, ECM CommonStore and Records Enabler Development, ECM eRecords Solution Development, Boca Raton, Florida, US

Bill Fuller
Neall Hards
IBM Software Group, Information Management, WW Software Support, Records Manager. L2 Support for Records Manager Enabler and Common Store, Ottawa, Canada

Become a published author

Join us for a two- to six-week residency program, and help write an IBM Redbook dealing with specific products or solutions, while getting hands-on experience with leading-edge technologies. You'll team with IBM technical professionals, Business Partners, and/or customers.

Your efforts will help increase product acceptance and customer satisfaction. As a bonus, you'll develop a network of contacts in IBM development labs, and increase your productivity and marketability.

Find out more about the residency program, browse the residency index, and apply online at:

ibm.com/redbooks/residencies.html

Comments welcome

Your comments are important to us!

We want our Redbooks™ to be as helpful as possible. Send us your comments about this or other Redbooks in one of the following ways:

- Use the online **Contact us** review redbook form found at:

  ibm.com/redbooks

- Send your comments in an e-mail to:

  redbook@us.ibm.com

- Mail your comments to:

  IBM Corporation, International Technical Support Organization
  Dept. QXXE Building 80-E2
  650 Harry Road
  San Jose, California 95120-6099

Part 1

Design and planning

This part introduces the planning and designing of the e-mail archiving and records management solution. We examine all of the products that are involved in the solution. For each product, we provide basic concepts and architecture and discuss the design options. From an integrated solution perspective, we cover areas that are specific to the integrated solution and address individual and overall planning considerations. The following chapters are included in this part:

1

Solution overview

This chapter provides an overview of the integrated solution for e-mail archiving and records management. We cover the key products involved (except IBM DB2 Content Manager) and their roles in the integrated solution.

Specifically, the topics we cover include:

- Introduction
- IBM DB2 IBM DB2 CommonStore for Lotus Domino
- IBM DB2 IBM DB2 CommonStore for Exchange Server
- IBM DB2 IBM DB2 Records Manager
- IBM DB2 DB2 Content Manager Records Enabler (CMRE)
- Integrated solution overview

1.1 Introduction

Years ago (and even today) customers dealt with their physical records without thinking of them as records—usually paper in files, boxes, and filing cabinets, or storage areas, records centers, or warehouses. Paper is not going away. Five years ago, if records management was part of a customer information management strategy or included in any Request for Proposal (RFP), it was at the bottom of the list. Today, records management is not only front and center for every customer of any size; for the majority, it leads the need. Whether a customer is mandated (for example, in the USA via SEC 17a-4, SarBox, HIPAA regulations) or simply realizes that it is a good practice, applying formal records-keeping to information — somehow — is a must. Businesses have no choice.

Current trends

Current trends demonstrate that the tide of knowledge and awareness has been raised, not only for line-of-business owners but, as is key, the IT technical community. In addition, records-centric discussions are now on the rise internally between these parties and the records and legal staff.

You cannot demonstrate compliance without applying some form of formal records-keeping, regardless of the information or system and media. For many customers, the most immediate and high-risk information area is e-mail. Either through mandate, or due to the discovery-risk exposure of the volumes and details of business decisions that e-mail now carries, the first real high-priority focus to deploy records controls across companies' e-mail has been forced.

Companies realize that e-mail backups and archiving (although a good first step) do not make a records system. Disaster recovery needs remain to ensure that daily backups occur to preserve operations, but the IT, records, and legal teams need to determine and ensure that the company's basic retention schedules are not invalidated if off-site backups are being kept for five years, for example. There are key core benefits for ongoing business efficiency and information management that e-mail archiving and de-duplication features can bring. These must not be confused with or assumed to also resolve retention requirements.

Some companies start with defining and deploying an e-mail policy for all employees. This serves both to refresh employee awareness of the importance of e-mail and records to the company and to underscore the impact and risk of their part in that responsibility. Good e-mail policies also encourage preclassification of e-mail via foldering or drop-down category options when sending e-mail. Value-add can be gained from requesting more detailed keywords to be included in the subject line of e-mail, which can aid auto-classification options. Customers can start with applying e-mail records-keeping solutions to their C-level executives to address their highest risk,

then proceed to deploy to other groups and business processes using e-mail that has been identified through their internal risk analysis.

Digital Rights Management (DRM) has also been raised in the media recently as a way to manage records and retention. To expert records staff, this is anathema. To have distributed content packaged in a way that in the future will expire and be deleted or locked up and unreadable does not facilitate modern electronic records management.

Future trends

The volume and level of business transacted via e-mail will continue to increase. Customers of all sizes will continue to leverage Web-enabled transactions facilitated by automated agents and interfaces that may slightly reduce the continued exponential reach of e-mail. Best practice–aware customers are also learning that deploying real document management systems first is key to the capture, control, and distribution of data leading to information and decisions that should never be in e-mail to begin with. For example, overuse of Exchange public folders as a pseudo-document management system creates more problems than it solves. Together with real document and process-centric workflow internally and externally, more structured control of business and a reduced record risk is a noble goal.

An integrated solution

E-mail systems from any vendor are designed to be, and are, great e-mail services. They are not effective communication, information, and content control solutions. They do not have built-in retention and disposition features, nor are they designed to be manageable long-term storage repositories. A balance is needed between maintaining the core services of an e-mail solution such as private and shared e-mail folders, drag-and-drop filing, inbound and outbound mail and attachments, and the need to apply a base retention and disposition on all of the various classifications of e-mail.

The IBM records enabled solution provides comprehensive capabilities across a range of options to empower customers to apply and deploy records-keeping across a subset or all of their e-mail, and to leverage the core e-mail mailbox-management features that aid IT in their daily challenge.

The integrated solution is key to avoid disabling the core services that e-mail provides, and ensuring that basic retention is applied and candidate records are expunged at end-of-life. To not achieve both is fundamentally a waste of time. No solution would be chosen and deployed if it takes away core e-mail features.

Note: A good practice is to take away features for end users to continue creating personal e-mail archives or PST collections.

The primary goal from a records perspective is to place the correct retention on the e-mail and to be able to appropriately delete it as soon as is possible. If that goal cannot be met, it is a waste of records-time and undermines the whole effort.

The integrated solution that we address in this book comprises the following products:

- IBM DB2 CommonStore for Lotus Domino or IBM DB2 CommonStore for Exchange Server
- IBM DB2 Records Manager
- DB2 Content Manager Records Enabler (CMRE)

We examine these products in the following sections.

1.2 IBM DB2 CommonStore for Lotus Domino

CommonStore for Lotus Domino (CSLD) is an archiving and restoring application for Lotus Notes® documents. It can archive Notes document content, including attachments, from any Lotus Notes database that is CommonStore enabled and that resides on a Lotus Domino server.

CommonStore for Lotus Domino provides the following functions:

- Move or copy document content from Notes databases to an archive.
- Search for content in an archive.
- Retrieve or restore archived content.
- Display archived content.

In the integrated e-mail archiving and records management solution, CommonStore for Lotus Domino is used to archive e-mail from Domino mail databases to Content Manager, the repository for the archived e-mail.

There are many reasons to use CommonStore for Lotus Domino to archive e-mail:

- To keep your mail database size within certain limits.
- To move old or obsolete e-mail out of the way.
- To free up space on your Lotus Domino mail servers.
- To speed up communication with your Lotus Domino servers.
- To meet legal obligations.

The archived e-mail messages are stored in the archive. The archive is a logical grouping of content stored in a repository such as IBM DB2 Content Manager, IBM DB2 Content Manager OnDemand, or IBM Tivoli® Storage Manager.

Note: Although CommonStore for Lotus Domino supports IBM DB2 Content Manager OnDemand and Tivoli Storage Manager as e-mail archives, the integrated e-mail archiving and records management solution supports only IBM DB2 Content Manager as the e-mail archive. This is because Content Manager Records Enabler does not support these archives. For more about Content Manager Records Enabler, refer to 1.5, “DB2 Content Manager Records Enabler (CMRE)” on page 15.

You can archive or delete e-mail based on *policies* that you set up. For example, you can set up a policy to archive messages that are 90 days or older if the mail database reaches certain space limits. Within this policy, you can also set it such that after an e-mail is archived, its content will be deleted from the mail database, and an e-mail stub with a link to the archived e-mail will be left in the mail database. In addition, you can archive different parts of messages, depending on your configuration, and store them differently according to your business needs.

E-mail archiving and deletion can be done manually or automatically.

Figure 1-1 shows the CommonStore for Lotus Domino system architecture.

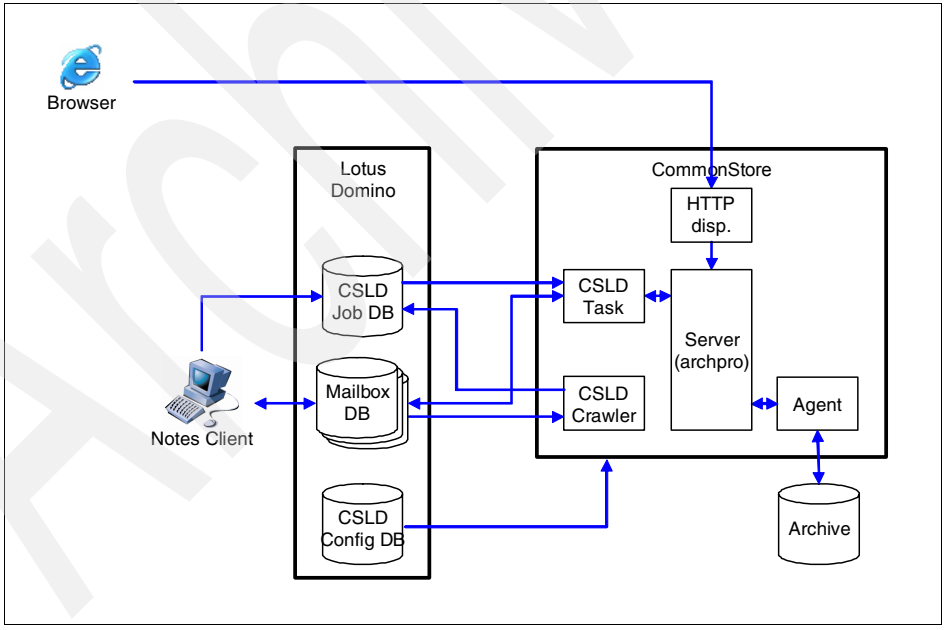

Figure 1-1 CommonStore for Lotus Domino system architecture

CommonStore for Lotus Domino uses two Notes databases:

- CSLD configuration database
- CSLD job database

CSLD configuration database

The *CSLD configuration database* is a Lotus Notes database that contains all of the configuration documents for CommonStore for Lotus Domino. The configuration database is located on a Lotus Domino server.

The configuration documents enable you to set up:

- Schedules
- Document selection criteria
- Archiving policies

CSLD job database

The *CSLD job database* is a Lotus Notes database in which all client requests, such as archiving, retrieval, search, and viewing, are collected before they are picked up by a CSLD task (see description below). The job database is located on a Lotus Domino server.

CommonStore for Lotus Domino comprises five major parts:

- CSLD crawler
- CSLD task
- CommonStore Server (archpro)
- Agent
- Web (HTTP) dispatcher

CSLD crawler

The *CSLD crawler* is the program that performs all automatic operations in CSLD. It creates automatic archiving and deletion jobs, as well as retrieval jobs that are centrally triggered by an administrator. The CSLD crawler directly accesses the databases on the Lotus Domino servers and looks for documents that match the criteria laid out in your policies.

CSLD task

The *CSLD task* is the program that directly interacts with your Lotus Notes and Domino environment. It looks for jobs in the CSLD job database, which is the place where all user requests are collected before they are processed further. The CSLD task converts the documents included in jobs or requests to files, passes the files to the CommonStore Server, and vice versa. You can run several instances of the CSLD task at the same time.

CommonStore Server (archpro)

The *CommonStore Server*, also known as *archpro*, is the heart of CommonStore. The CommonStore Server maintains a list of logical archives, and controls the flow of information to and from these archives. All input and output data is routed through the CommonStore Server, which distributes the requests among the archive agents (the interfaces to the archive).

The CommonStore Server is an application-independent archiving engine. It does not have to know where the content comes from, what format the content has, or what the content is.

Agents

The *agents* are the interfaces to the archive. Every agent is an independent archive client process. For every archive supported by CommonStore, there is a special agent. An agent calls the application programming interface (API) of the archive. These archives are supported by CommonStore for Lotus Domino:

- Content Manager Version 8 (CM8)
- Content Manager OnDemand (CMOD)
- Tivoli Storage Manager (TSM)

Agents always run on the same machine as the CommonStore Server (archpro) and are automatically started by it.

Web (HTTP) dispatcher

The *Web (HTTP) dispatcher* is a Web server for content archived by CommonStore. It enables you to view the archived content in a Web browser. The Web dispatcher is installed as part of the CommonStore Server and is automatically started by it.

1.3 IBM DB2 CommonStore for Exchange Server

CommonStore for Exchange Server (CSX), an archive-and-restore application for Microsoft® Exchange Server, can archive Exchange messages including attachments from any Exchange mailbox that resides on an Exchange server.

CommonStore for Exchange Server provides the following functions:

- Move or copy document content from Exchange mailboxes or Exchange Public Folders to an archive.
- Search for content in an archive.
- Display archived content.
- Retrieve or restore archived content.

In the integrated e-mail archiving and records management solution, CommonStore for Exchange Server is used to archive e-mail from Microsoft Exchange mail databases to Content Manager, the repository for the archived e-mail.

There are many reasons to use Exchange Server to archive e-mail:

- To keep your mail database size within a certain limit.
- To move old or obsolete messages out of the way.
- To free up space on your Microsoft Exchange Servers.
- To speed up communication with your Exchange Servers.
- To meet legal obligations.

The archived messages are stored in the archive. The archive is a logical grouping of content stored in a repository such as IBM DB2 Content Manager, IBM DB2 Content Manager OnDemand, or Tivoli Storage Manager.

Note: Although CommonStore for Exchange Server supports IBM DB2 Content Manager OnDemand and Tivoli Storage Manager as e-mail archives, the integrated e-mail archiving and records management solution supports only IBM DB2 Content Manager as the e-mail archive. This is because Content Manager Records Enabler does not support these archives. For more information about Content Manager Records Enabler, refer to 1.5, “DB2 Content Manager Records Enabler (CMRE)” on page 15.

You can archive e-mail based on *policies* you set up. For example, you can set up a policy to archive messages that are 90 days or older if the mail database reaches certain space limits. Additional options enable you to reduce the size of your mailbox by deleting the attachments, the message body, or the entire message after successful archival. For example, within the previously mentioned policy, you can set it such that when an e-mail is archived, the e-mail content will be deleted from the mail database, and an e-mail stub with a link to the archived e-mail will be left in the mail database. In addition, you can archive different parts of messages depending on your configuration and store them differently according to your business needs.

E-mail archiving and deletion can be done manually or automatically.

Figure 1-2 on page 9 shows the CommonStore for Exchange Server system architecture.

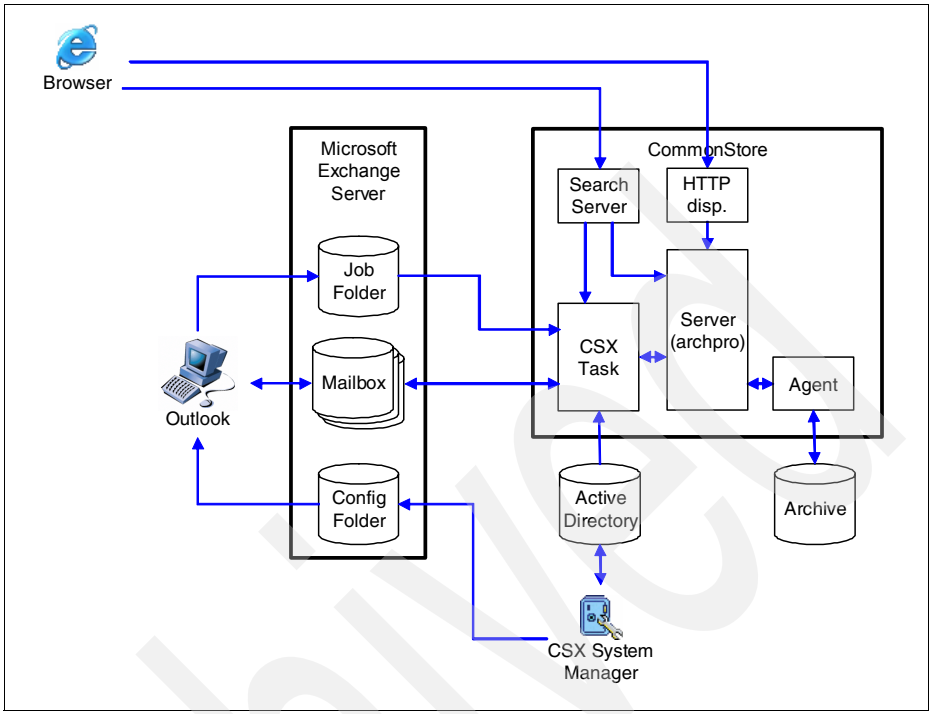

Figure 1-2 CommonStore for Exchange Server system architecture

CommonStore for Exchange Server installs a public folder named CommonStore. This folder contains:

- Configuration folder
- Job folder

Configuration folder

The *configuration folder* contains all of the information that Microsoft Outlook® clients need to execute CommonStore functionality. This includes:

- Job folder name
- Server host names and ports

This data is written during system configuration using the CSX System Manager (see description below).

Job folder

The *job folder* contains interactive client requests to archive or retrieve messages. They are collected in the folder before they are picked up by a CSX Task.

CSX System Manager

The CSX System Manager is a Microsoft Management Console (MMC) that provides a graphical user interface for the CommonStore administrator. Using the CSX System Manager, you configure the behavior of your archiving solution including the instances of the CSX Task (see description on CSX Task below). The CSX System Manager saves the settings in the Active Directory of the forest that the CSX Task instances and the connected Exchange servers belong to. When an instance of the CSX Task is started, it reads the configuration data from this Active Directory. Additionally, the CSX System Manager writes client-relevant information into the configuration folder.

CommonStore for Exchange Server is comprised of five major parts:

- CSX Task
- CommonStore Server (archpro)
- Agent
- Search server
- Web (HTTP) dispatcher

CSX Task

The *CSX Task* is the program that directly interacts with your Microsoft Exchange environment. It is responsible for performing interactive and automatic archiving and retrieval, during which it transforms Exchange messages into files and vice versa.

The CSX Task can be logically split into the following components:

- Crawler
- Poller
- Worker
- Committer

Figure 1-3 on page 11 shows the individual components of the CSX Task.

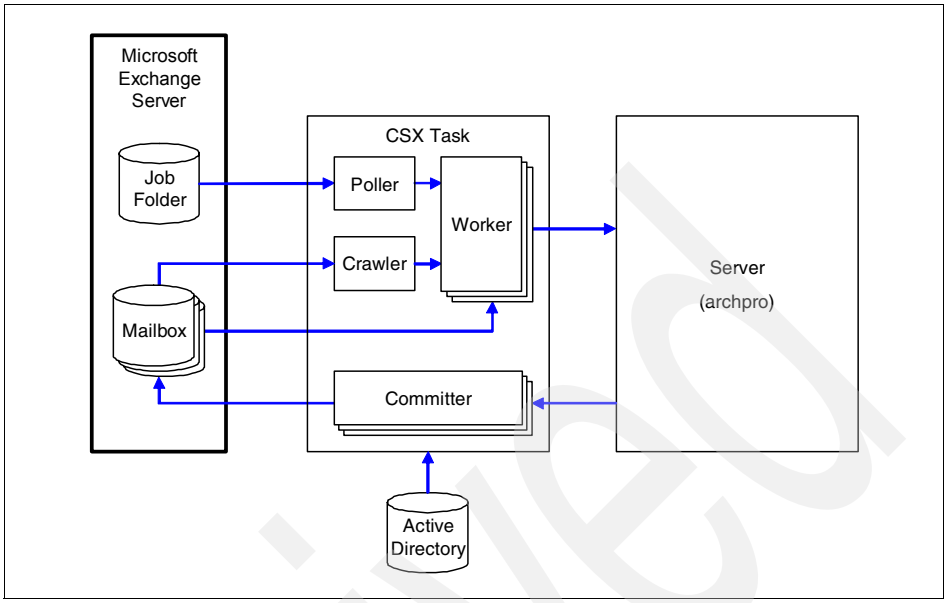

Figure 1-3 CSX Task components

CSX Task: Crawler

The *crawler* component carries out policy-driven archiving processes. It investigates mailboxes and public folders to see whether any message meets certain selection criteria. These criteria are defined in policies. If messages meet the criteria that you defined in policies, the crawler creates archiving requests for them. Each crawler does this for exactly one Exchange server.

In addition, the crawler initiates the automatic removal of retrieved content, also known as restubbing. More discussion about restubbing is provided in “CSX Task: Committer” on page 12.

CSX Task: Poller

The *poller* component continually looks for interactive job messages in the job folder. Interactive jobs are created by clients (Outlook users). The job messages contain the information needed to perform the jobs. The poller converts this information and puts it in an internal job queue.

CSX Task: Worker

The *worker* component reads the instructions from the converted job messages in the internal job queues. It then accesses the Outlook messages referred to in the job messages and reads the content to be archived or identifies the content to be retrieved. It finally creates the actual job request and passes it to the CommonStore Server (archpro).

CSX Task: Committer

The *committer* component receives the answers from the CommonStore Server and processes the instructions within it. If an operation is successful, this results in modifications of the original messages according to the instructions. For archiving jobs, this means that content may be deleted from the original messages according to the selected deletion type. Additionally, the status is set to `archived` and a few properties are added to the message. For retrieval jobs, it means restoring archived content and setting the status to `retrieved`. The committer is also responsible for restubbing.

The term *restubbing* describes the process of deleting retrieved content after a certain time and resetting the status to archived. When messages are restubbed, this is done in accordance with the deletion type that is specified at the time the message is archived. Restubbing does not involve communication with the CommonStore Server.

CommonStore Server (archpro)

The *CommonStore Server software*, also known as *archpro*, is the heart of CommonStore. The CommonStore Server maintains a list of logical archives and controls the flow of information to and from these archives. All input and output data is routed through the CommonStore Server. The CommonStore Server distributes the requests among the archive agents, which are the interfaces to the archive.

The CommonStore Server is an application-independent archiving engine. It does not have to know where the content comes from, what format the content has, or what the content is.

Agents

The *agents* are the interfaces to the archive. Every agent is an independent archive client process. For every archive supported by CommonStore, there is a special agent. An agent calls the application programming interface (API) of the archive. The following archives are supported by CommonStore for Exchange Server:

- Content Manager Version 8 (CM8)
- Content Manager OnDemand (CMOD)
- Tivoli Storage Manager (TSM)

Agents always run on the same machine as the CommonStore Server (archpro) and are automatically started by it.

Search server

The *search server* is a separate Web application server process for the search function. It is automatically started by the CommonStore Server. It routes the

search requests via the CommonStore Server to the archives that support this function (Tivoli Storage Manager does not support it). After receiving the results from the CommonStore Server, it arranges them into a list and passes that list to the requesting Web application that displays it in the browser window. The result list enables the requesting user to view the messages and attachments that are found in a browser. In addition, the user can restore the archived content to a special folder in the user's mailbox. Messages are restored to their original form, and attachments are restored to container messages.

Web (HTTP) dispatcher

The *Web (HTTP) dispatcher* is a Web server for content archived by CommonStore. It enables you to view the archived content in a Web browser. The Web dispatcher is installed as part of the CommonStore Server and is automatically started by it.

1.4 IBM DB2 Records Manager

Records Manager (IRM) is a records management engine and infrastructure tool to records enable business applications. It adds the benefits of records management to business applications, with no software to install on workstations. It provides a single and consistent records management platform with extensive record-keeping capabilities for both electronic and physical information assets. It helps meet government and industry requirements for formal records management.

Records Manager is DoD 5015.2-STD Chapter 4 certified with DB2 Content Manager and is Chapter 2 certified with IBM DB2 Content Manager and IBM DB2 Document Manager.

Records Manager can records enable content, documents, and e-mail using one central repository for all organizational records. It is integrated with:

- IBM DB2 Content Manager
- IBM DB2 Document Manager
- IBM DB2 CommonStore for Lotus Domino
- IBM DB2 CommonStore for Exchange Server

It supports IBM DB2 Universal Database™ V8.2 FP 1, Microsoft SQL Server 2000 SP3a, and Oracle 9i v9.2.0.5.

Records Manager provides the following functions:

- Declare and classify corporate records.
- Apply life cycle management of the declared records.
- Embedded engine technology that provides record keeping capability.

- Search, retrieve, view and print record metadata.
- Retrieve record contents for viewing.
- Delete and un-declare corporate records.
- Create various reports for auditing and records management purposes.
- Manage user and group accounts, administer life cycle of records, and perform other records management functions.
- Import of host users and groups.

In the integrated e-mail archiving and records management solution, Records Manager is integrated with the CommonStore and Content Manager products to records enable e-mail.

There are many reasons to use Records Manager with your business applications:

- To reduce litigation risk via structured document destruction.
- To reduce discovery costs during litigation via improved evidence discovery.
- To demonstrate compliance with regulations.
- To improve decision-making.
- To reduce operational costs.
- To securely maintain corporate records.
- To support key standards.

Features include:

- Embedded engine technology; no new application to learn and maintain.
- Web-based client for records administration.
- Quick integration using multiple client technologies, such as Java™, C++, and .Net.
- Scalable architecture.
- Content maintained in host repository; no redundant data.

You can customize your business application to declare and classify records manually, automatically, or somewhere in between.

Records Manager is composed of four major components:

- Records Manager engine
- Records Manager database
- Records Manager APIs
- Records Manager administrator Web client.

Records Manager engine

The Records Manager engine provides all of the business logic that is required to enable life cycle management. You can embed the engine into the e-mail, document management, or other applications to records enable your applications.

Records Manager database

The Records Manager database contains all the configuration data of the Records Manager system. This includes file plan, life cycle configuration, retention rules, users and user groups, security setup, and other system information. It also contains the records-related metadata of the documents that have been declared as records; however, the declared documents are still stored in the same repository of the original business application.

Records Manager APIs

The Records Manager application programming interface is used to embed Records Manager into any business application. The APIs enable an application to access Records Manager data. You can also use the APIs to modify, enhance, customize, or rewrite the existing Records Manager administrator Web client interface.

Records Manager administrator Web client

The Records Manager administrator Web client is a browser-based application that enables administrators to design, build, and maintain corporate file plans and retention rules, and manage records life cycles, users, security, and other administrative tasks.

1.5 DB2 Content Manager Records Enabler (CMRE)

Content Manager Records Enabler (CMRE) adds record management function to the Content Manager product by integrating it with Records Manager. Together, Records Manager provides the record management functions, and Content Manager provides the content repository. Records Enabler links the two and provides records management capability to existing e-mail, content management, and document management clients such as the Content Manager Client for Windows, eClient, Document Manager, Lotus Notes client, and Microsoft Outlook client. Working with CommonStore for Lotus Domino or CommonStore for Exchange Server, the solution provides e-mail archiving and records management capability.

With Records Enabler, you can declare corporate documents or content as corporate records. When a document is declared as a record, the contents and associated Content Manager metadata of the record remain in the Content Manager repository, and the document can no longer be edited or deleted. The records-related metadata is stored in the Records Manager database. The access permissions and life cycle of the record and its content are governed by the access permissions and life cycle rules that are defined for the record in the Records Manager system. Only authorized users, such as records administrators, can process or manage the life cycle of the records.

Records Enabler provides the following functions and features:

- Integrates with CommonStore for Lotus Domino and CommonStore for Exchange Server to enable users to declare e-mail messages and attachments as records.
- Integrates with Content Manager and Document Manager to enable users to declare corporate documents as records.
- Keeps the content of the declared records in the same Content Manager repository.
- Performs scheduled and on demand permission synchronization tasks.
- Enables users to import Content Manager user and group accounts into Records Manager.
- Enables users to view declared records using the Content Manager clients.
- Enables users to perform text searches on items stored in the Content Manager repository.

Figure 1-4 shows the Records Enabler subsystems.

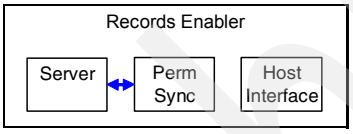

Figure 1-4 Records Enabler subsystems

Records Enabler is comprised of the following subsystems:

- Records Enabler Server
- Records Enabler Host Interface Server
- Records Enabler Permission Synchronization Server
- Records Manager Extensions

Records Enabler Server

Records Enabler Server provides a centralized base for handling Records Enabler user requests and their appropriate actions.

Records Enabler Host Interface Server

Records Enabler Host Interface Server implements the Records Manager Host Interface application program interface (API) for Content Manager to enable communication between the Records Manager and the Content Manager servers.

Records Enabler Permission Synchronization Server

Records Enabler Permission Synchronization Server performs on-demand and scheduled permission synchronization tasks from the latest file plan components

that have their permission policies changed since the previous synchronization and the corresponding items in Content Manager.

Records Manager Extensions

Records Manager Extensions extends the Records Manager installation to support Records Enabler. It provides the notification service for permissions synchronization between Records Manager and Content Manager. This subsystem must be installed on the Records Manager server.

1.6 Integrated solution overview

The integrated e-mail archiving and records management solution that this book focuses on is comprised of the products we introduced in the previous sections.

Figure 1-5 shows the system architecture of the integrated system.

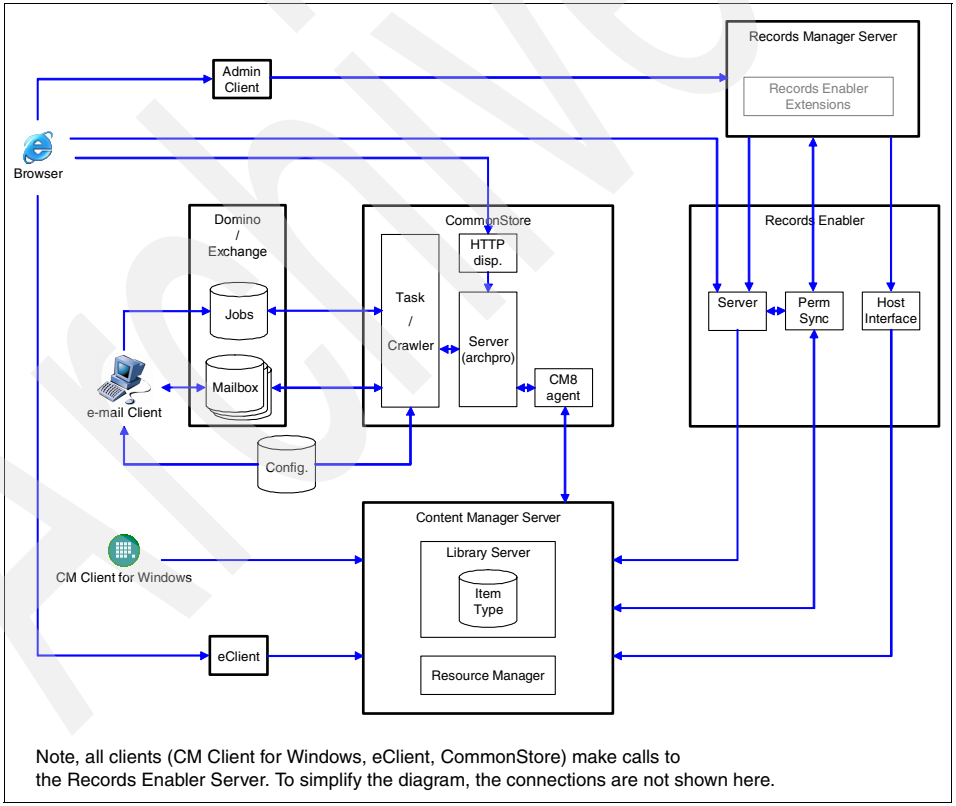

Figure 1-5 Integrated solution system architecture

The integrated solution is an integration of the following products:

- Lotus Domino server/Microsoft Exchange Server (Domino/Exchange)
- Mailboxes (Lotus Notes databases or Outlook mailboxes)
- E-mail client (Lotus Notes/Microsoft Outlook)
- CommonStore Server (for Lotus Domino or Exchange)
- Records Manager server (with Records Enabler extension)
- Records Enabler
- Content Manager server

This solution provides e-mail archiving capability and e-mail records management function. Specifically, the solution offers the following functions:

- Using Lotus Notes client with CommonStore for Lotus Domino, or Microsoft Outlook client with CommonStore for Microsoft Exchange Server, you can:
  - Archive e-mail messages, attachments, and documents.
  - Declare and classify e-mail messages and attachments as records.
  - View archived content.
  - View record information.
  - Search and retrieve archived content and records.
- Using the CommonStore for Lotus Domino configuration database or the CommonStore for Exchange Server system manager, you can:
  - Perform user administration.
  - Set up archiving rules and policies.
- Using Records Manager administrator client, you can:
  - Declare electronic and physical documents as records.
  - Delete or undeclare records.
  - Search, import, retrieve, and view record contents.
  - Retrieve record contents from Content Manager repository.
  - Query records or generate record reports.
  - Import into Content Manager and declare records.
  - Perform records management functions on records.
  - Perform user administration, records life cycle management, records management, and other system management functions.
- Using eClient or Client for Windows, you can:
  - Declare and classify Content Manager items as records.
  - View record information.

2

Design and planning for e-mail archiving

This chapter covers e-mail archiving solution design and planning.

In the integrated solution discussed in this book, e-mail archiving is achieved with IBM DB2 CommonStore for Lotus Domino or IBM DB2 CommonStore for Exchange Server. In this chapter, we introduce the basic concepts behind e-mail archiving, and describe CommonStore features and functions. We address key areas to consider when planning and designing the e-mail archiving portion of a solution.

We cover the following topics in the chapter:

- CommonStore e-mail archiving design options
- CommonStore e-mail archiving solution planning

This chapter mainly focuses on areas that are related to the integrated solution and is intended to be used in conjunction with the existing product manuals. To have an in-depth understanding of the individual product, we recommend reference to the existing product manuals.

2.1 CommonStore e-mail archiving design options

CommonStore (for Lotus Domino and for Exchange) offers a variety of methods for archiving e-mail. Different methods determine different e-mail archiving behavior. It is important to understand what happens to an e-mail message when it is archived by CommonStore using different methods. Additionally, it is useful to understand what is left in the e-mail entry within the mail database after it is archived using different methods.

In this section, we examine:

- E-mail message layout
- Archiving types
- Document storage model
- Deletion types
- Retrieving archived content
- Viewing archived content
- Archiving options and policy

The archiving type controls which parts of a message are archived in the archive repository and which hyperlinks make sense to be included in the message stub, and the document storage model controls how the archived content is stored in the Content Manager archive. The deletion type determines which parts of the original e-mail message are removed after it is archived successfully. Including retrieval and viewing hyperlinks in the message stub and the options to configure this are different for the two CommonStore products.

This section provides an overview of the main concepts that are important in designing and planning the integrated e-mail archiving and records management solution. For more detailed information, refer to the following publications:

- *IBM DB2 CommonStore for Lotus Domino: Administrator's and Programmer's Guide* Version 8.3, SH12-6742
- *IBM DB2 CommonStore for Exchange Server: Administration and User's Guide Version 8.3*, SH12-6741

2.1.1 E-mail message layout

To work with e-mail archiving, you must first understand the layout of an e-mail message. An e-mail message consists of the information in the message *body*, information stored in the message *properties* (visible or invisible), and *attachments*. See Figure 2-1 on page 21.

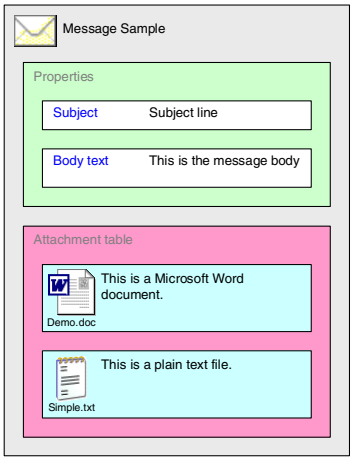

Figure 2-1 E-mail message layout

2.1.2 Archiving types

The e-mail *archiving type* controls *which parts of a message* are stored in a repository.

The following archiving types are available:

- Entire
- Component
- Attachment

Note: The archiving type does *not* control how the archived data is stored in the Content Manager archive. This is controlled by the storage model selected. (Refer to 2.1.3, “Document storage model” on page 24.)

With entire and component archiving types, Microsoft Exchange e-mail messages are stored in the Microsoft Message Format (using extension .msg), and Lotus Notes e-mail messages can be stored in Notes native or Domino XML (DXL) format. The DXL format has the added benefit of a message being able to be displayed in a browser. The Notes native format requires a Notes client to view it.

Archiving type: Entire

Using archiving type *Entire*, the complete e-mail message content is archived. That is, the body text, the attachments, and all other message properties are archived. See Figure 2-2.

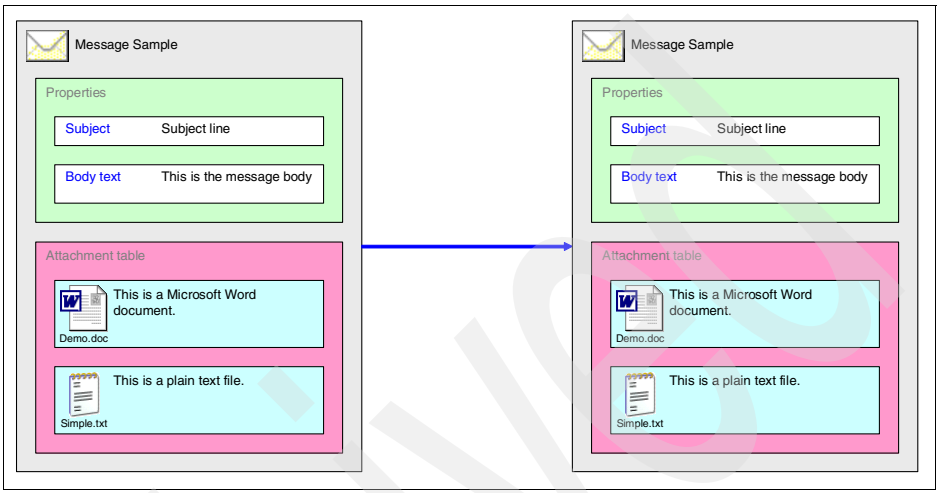

Figure 2-2 Archiving type: Entire

Note: Using archiving type Entire does *not* allow you to declare individual attachments as records.

Archiving type: Component

Using archiving type *Component*, the entire message is archived. However, component archiving does not store the entire content in a single file, but decomposes the message into separate attachment files and the rest into a single file containing the message remainder. For example, if a message has two attachments, the message is stored in three parts. See Figure 2-3.

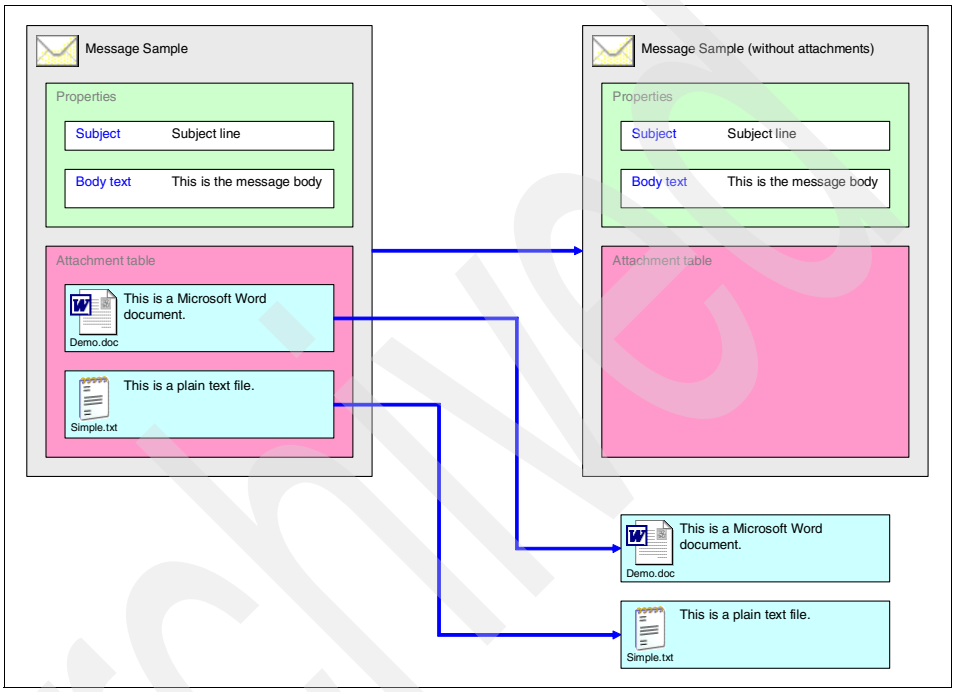

Figure 2-3 Archiving type: Component

Note: Component archiving in combination with the *GENERIC_MULTIDOC* document model is one of the requirements for DoD or PRO2 compliance.

Refer to 2.1.3, “Document storage model” on page 24 for discussion about GENERIC_MULTIDOC.

Archiving type: Attachment

Using archiving type *Attachment*, only the attachments are archived. The format of the attachments remains unchanged. See Figure 2-4.

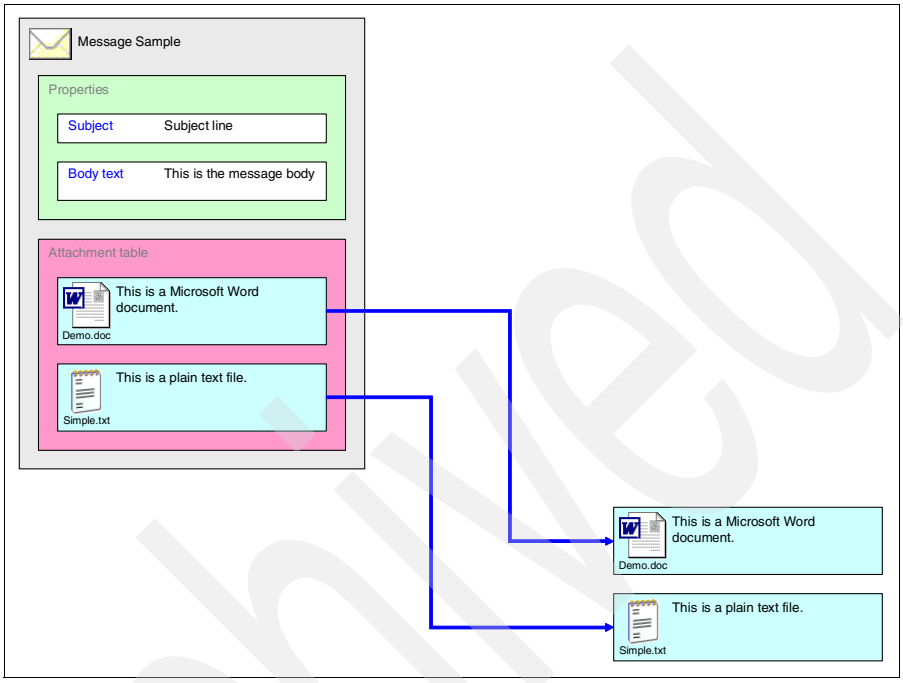

Figure 2-4 Archiving type: Attachment

Note: Archiving type Attachment is an *incomplete* archiving type, because the message body is *not* archived.

2.1.3 Document storage model

Compared to the archiving type, which controls what parts of a message are archived in the Content Manager archive, *document storage model* controls *how* the archived data is stored in the Content Manager archive.

Document storage models are repository specific.

The models provided with Content Manager V8.3 repository are:

- GENERIC_MULTIPART
- GENERIC_MULTIDOC
- BUNDLED

These storage models can be used in combination with any archiving type, which allows nine different combinations.

Storage model: GENERIC_MULTIPART

All e-mail components of a single e-mail, such as the message body and the attachments, are stored as *one* Content Manager V8 document with *one* or *more* parts. See Figure 2-5.

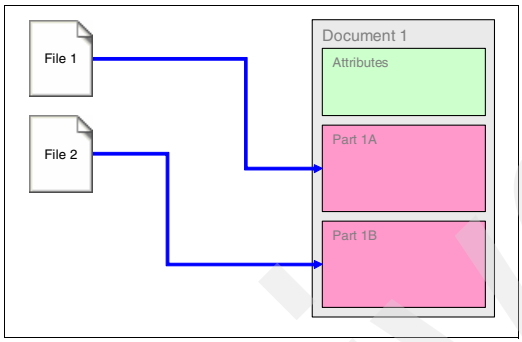

Figure 2-5 Storage model: GENERIC_MULTIPART

As this storage model makes use of the generic document model of Content Manager, the individual message components (entire message, attachments or message remainder) can be accessed from Content Manager clients unless the CSN format for CSLD is used to store messages.

When using the text-search user exit, there are some limitations with the archiving types Attachment and Component. See the CommonStore document *Text Search Configuration for IBM DB2 Content Manager V8* for more information.

Note: Using storage model GENERIC_MULTIPART does *not* allow you to declare individual attachments as records.

GENERIC_MULTIPART with archiving type Entire

Using archiving type Entire with storage model GENERIC_MULTIPART, the complete message is archived as *one* Content Manager V8 document with *one* part. See Figure 2-6.

Based on the property mappings you defined in your configuration, the message properties are stored as attributes in Content Manager.

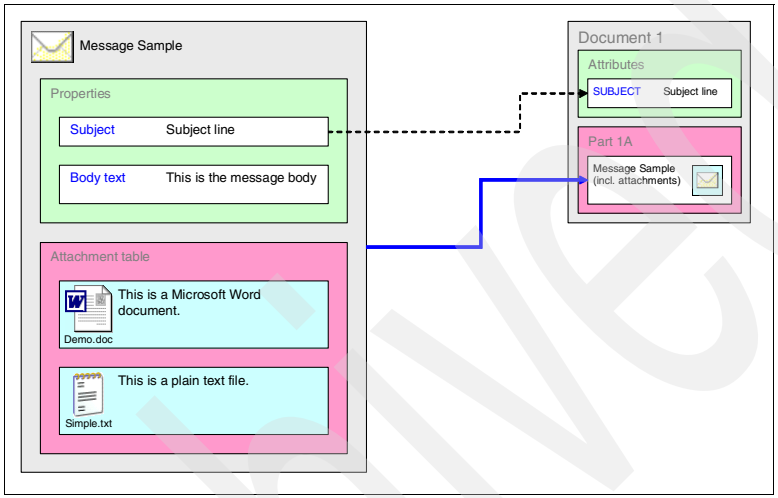

Figure 2-6 Storage model GENERIC_MULTIPART with archiving type Entire

GENERIC_MULTIPART with archiving type Component

Using archiving type Component with storage model GENERIC_MULTIPART, the message is decomposed into the attachments and the message remainder. These message components are archived as *one* Content Manager V8 document with *one or more* parts. See Figure 2-7 on page 27.

Based on the property mappings you defined in your configuration, the message properties are stored as attributes in Content Manager.

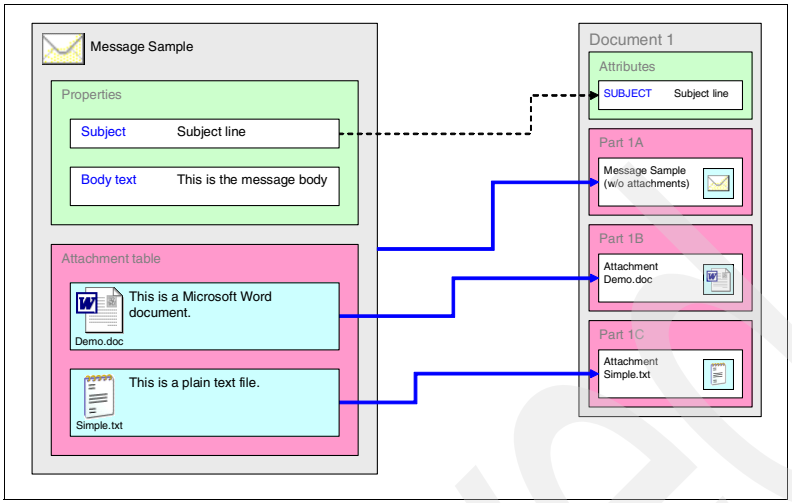

Figure 2-7 Storage model GENERIC_MULTIPART with archiving type Component

GENERIC_MULTIPART with archiving type Attachment

Using archiving type *Attachment* with storage model *GENERIC_MULTIPART*, the attachments are detached from the message and archived in Content Manager. Based on the property mappings you defined in your configuration, the message properties are stored as attributes in Content Manager.

Using CommonStore for Exchange Server, the attachments are archived as *one* Content Manager V8 document with *one* or *more* parts. See Figure 2-8.

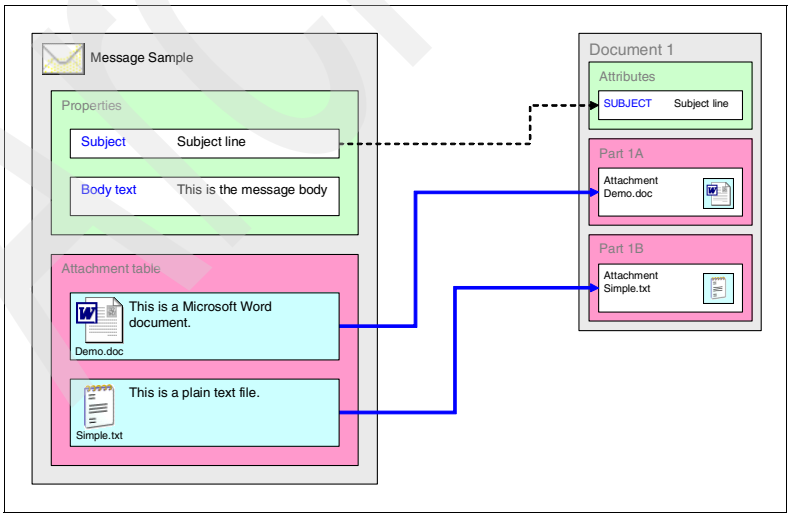

Figure 2-8 Storage model GENERIC_MULTIPART with archiving type Attachment (CSX)

Using CommonStore for Lotus Domino, the attachments are archived as *individual* Content Manager V8 documents with *one* part each. See Figure 2-9.

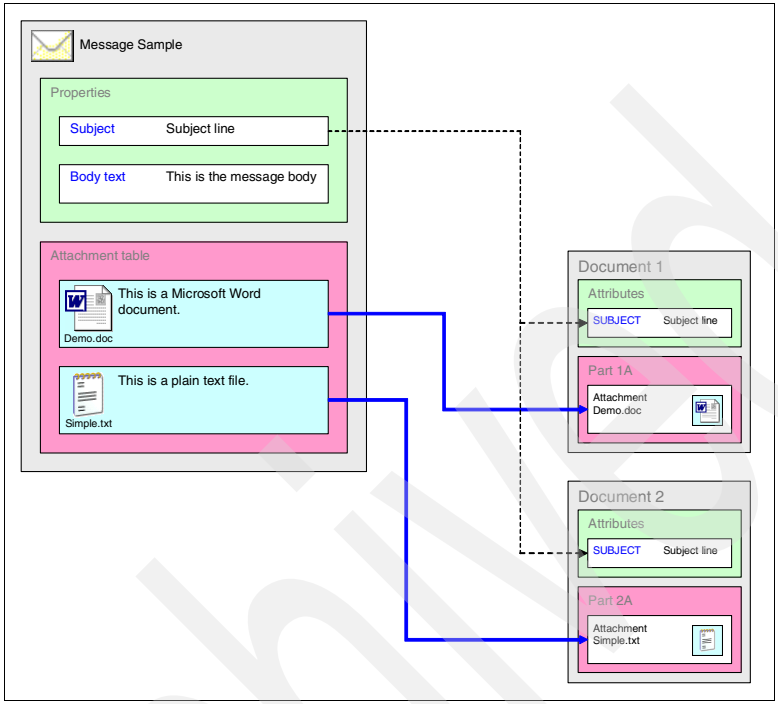

Figure 2-9 Storage model GENERIC_MULTIPART with archiving type Attachment (CSLD)

Storage model: GENERIC_MULTIDOC

All e-mail components of a single e-mail (such as the body and the attachments) are stored as *one or more individual* Content Manager V8 documents with *one* part each. See Figure 2-10.

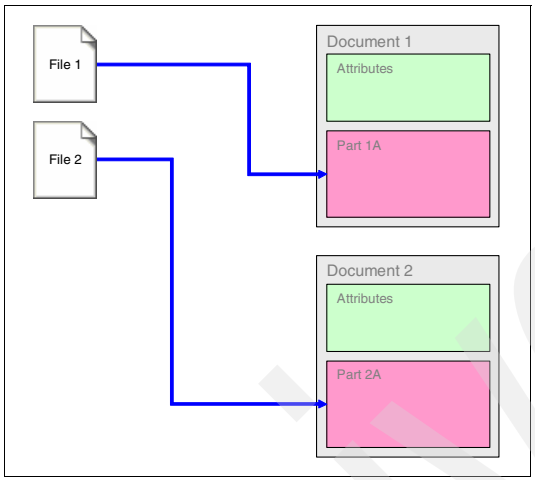

Figure 2-10 Storage model GENERIC_MULTIDOC

As this storage model makes use of the generic document model of Content Manager, the individual message components (entire message, attachments, or message remainder) can be accessed from Content Manager clients unless the CSN format for CSLD is used to store messages.

When using the text-search user exit, there are some limitations with the archiving types Attachment and Component. See the CommonStore document *Text Search Configuration for IBM DB2 Content Manager V8* for more information.

GENERIC_MULTIDOC with archiving type Entire

Using archiving type Entire with storage model GENERIC_MULTIDOC, the complete message is archived as *one* Content Manager V8 document with *one* part. See Figure 2-11.

Based on the property mappings you defined in your configuration, the message properties are stored as attributes in Content Manager.

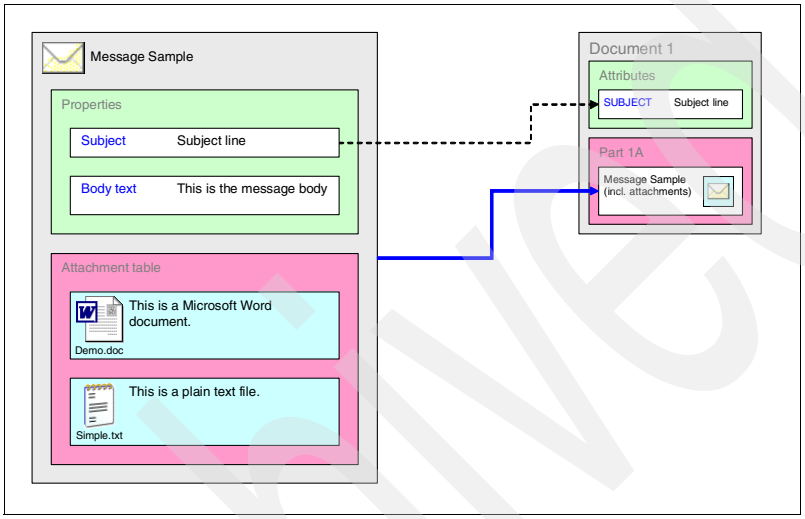

Figure 2-11 Storage model GENERIC_MULTIDOC with archiving type Entire

GENERIC_MULTIDOC with archiving type Component

Using archiving type Component with storage model GENERIC_MULTIDOC, the message is decomposed into the attachments and the message remainder. These message components are archived as *one or more individual* Content Manager V8 documents with *one* part each. See Figure 2-12.

Based on the property mappings you defined in your configuration, the message properties are stored as attributes in Content Manager with each document.

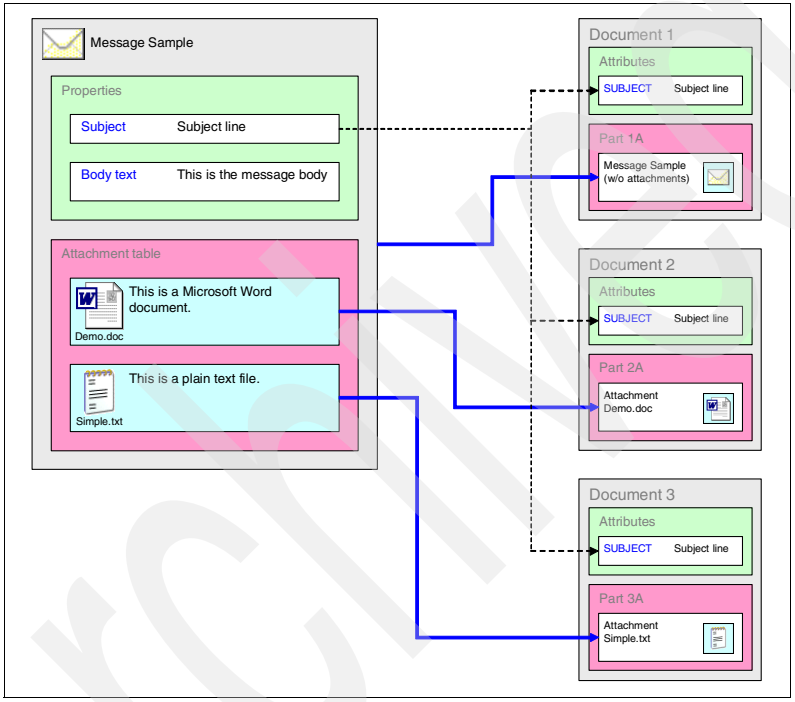

Figure 2-12 Storage model GENERIC_MULTIDOC with archiving type Component

GENERIC_MULTIDOC with archiving type Attachment

Using archiving type Attachment with storage model GENERIC_MULTIDOC, the attachments are detached from the message. The attachments are archived as *individual* Content Manager V8 documents with *one* part each. See Figure 2-13.

Based on the property mappings you defined in your configuration, the message properties are stored as attributes in Content Manager with each document.

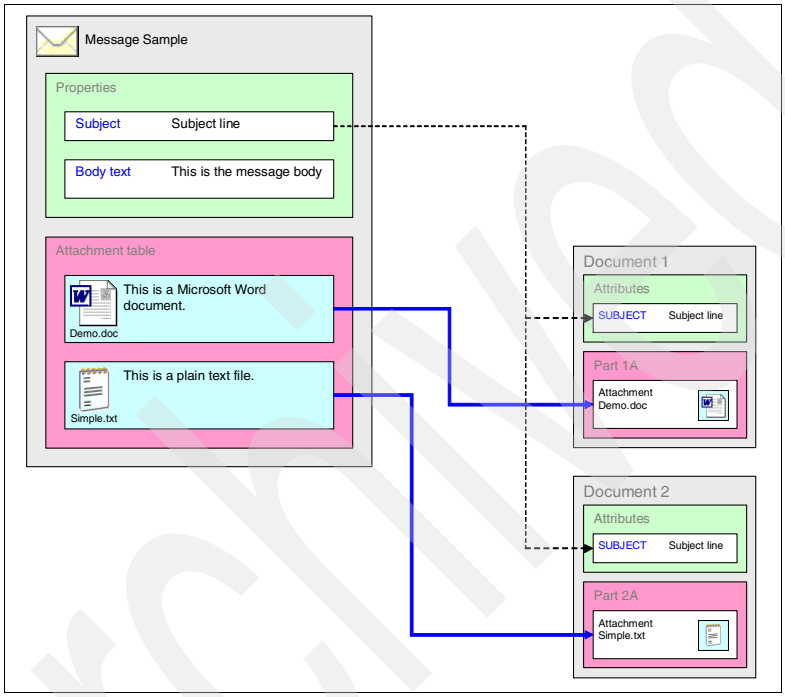

Figure 2-13 Storage model GENERIC_MULTIDOC with archiving type Attachment

Storage model: BUNDLED

All e-mail components of a single e-mail (such as the body and the attachments) are stored in *a single* Content Manager V8 *resource item*. See Figure 2-14.

Independent of the property mappings you defined in your configuration, the basic message properties Subject, From, To, Cc, and Bcc are stored in the resource item if available.

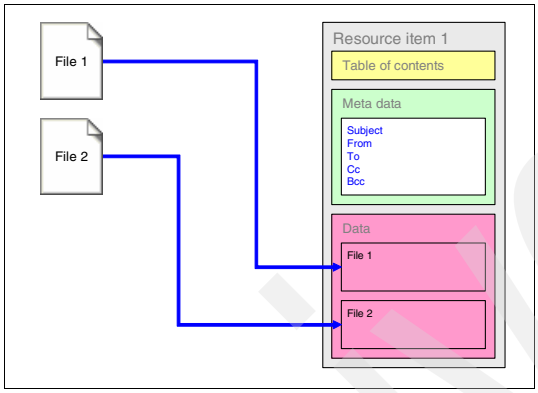

Figure 2-14 Storage model BUNDLED

Note: Using storage model BUNDLED does *not* allow you to declare individual attachments as records.

This storage model does not make use of the generic document model of Content Manager. As the format is a CommonStore native format, archived data *cannot* be accessed from Content Manager clients.

When using the text-search user exit, there are no limitations. See the CommonStore document *Text Search Configuration for IBM DB2 Content Manager V8* for more information.

BUNDLED with archiving type Entire

Using archiving type Entire with storage model BUNDLED, the complete message is archived in *one* Content Manager V8 *resource item*. See Figure 2-15.

Based on the property mappings you defined in your configuration, the message properties are stored as attributes in Content Manager. In addition to that, the basic message properties Subject, From, To, Cc, and Bcc are stored in the resource item.

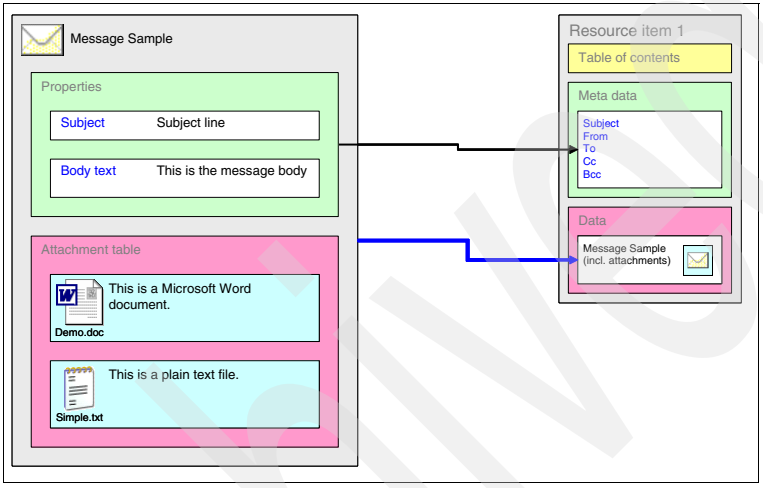

Figure 2-15 Storage model BUNDLED with archiving type Entire

BUNDLED with archiving type Component

Using archiving type Component with storage model BUNDLED, the message is decomposed into the attachments and the message remainder. These message components are archived in *one* Content Manager V8 *resource item*. See Figure 2-16 on page 35.

Based on the property mappings you defined in your configuration, the message properties are stored as attributes in Content Manager. In addition to that, the basic message properties Subject, From, To, Cc, and Bcc are stored in the resource item.

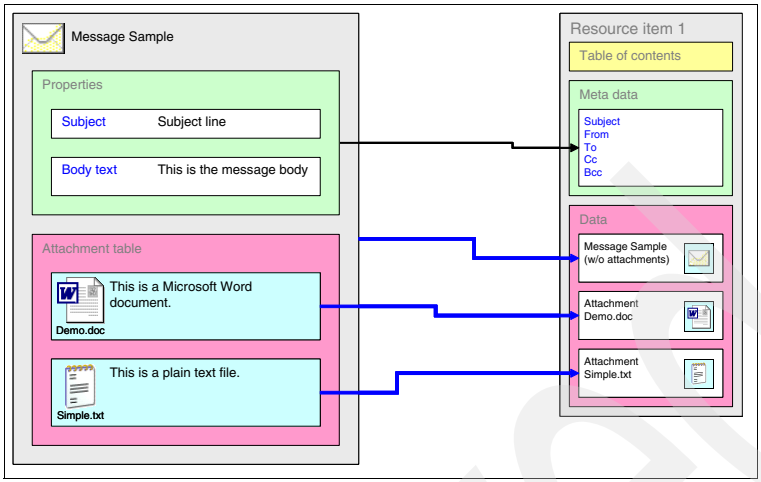

Figure 2-16 Storage model BUNDLED with archiving type Component

BUNDLED with archiving type Attachment

Using archiving type Attachment with storage model BUNDLED, the attachments are detached from the message and archived in Content Manager. Based on the property mappings you defined in your configuration, the message properties are stored as attributes in Content Manager.

Using CommonStore for Exchange Server, all attachments are archived in *one* Content Manager V8 *resource item*. See Figure 2-17.

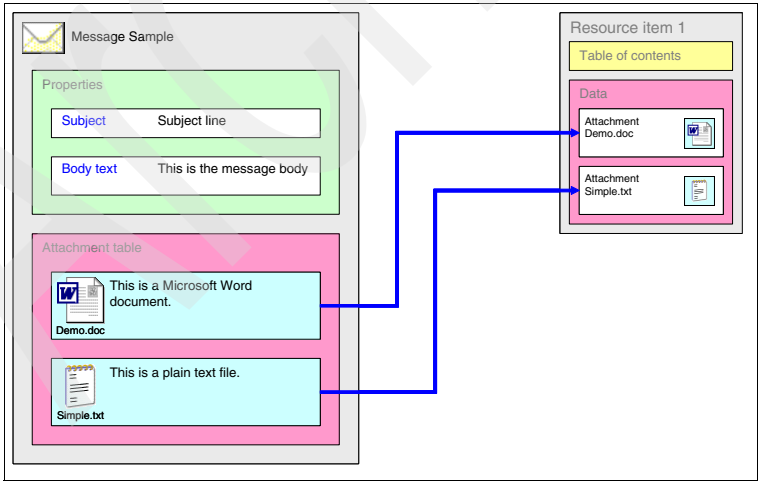

Figure 2-17 Storage model BUNDLED with archiving type Attachment (CSX)

Using CommonStore for Exchange Server, each attachment is archived in an *individual* Content Manager V8 *resource item*. See Figure 2-18.

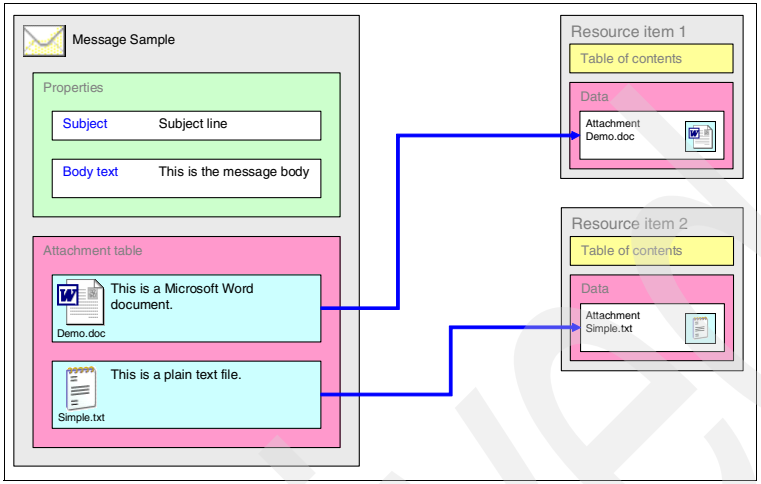

Figure 2-18 Storage model BUNDLED with archiving type Attachment (CSLD)

2.1.4 Deletion types

The e-mail *deletion type* determines which part of the original e-mail message is removed after it is archived successfully.

Selection of the deletion type to be used is performed in either the Lotus Notes CSLDConfig database (using CommonStore for Lotus Domino) or the CSX System Manager (using CommonStore for Exchange).

For each archiving policy or archiving rule, you can select one of the following deletion types:

- Nothing
- Attachments
- Body
- Message

Deletion type: Nothing

The original e-mail messages are left untouched after they are archived into the Content Manager repository. This is similar to creating a copy of the messages in the archive. This option does not reduce the size of the mail boxes or public folders. It slightly increases their sizes because archive information is added to the original messages. See Figure 2-19 on page 37.

You can select this deletion type if the archiving type is Entire, Component, or Attachment.

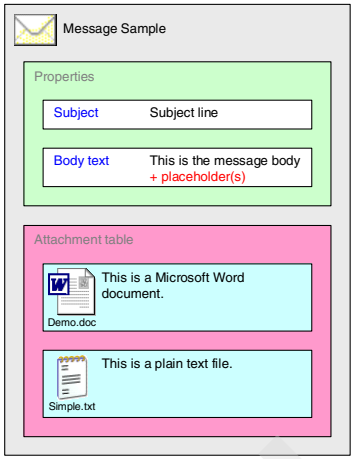

Figure 2-19 Deletion Type: Nothing

Deletion type: Attachments

After the attachments are archived, the attachments within the original e-mail message are removed. The body texts remain in the mail boxes and the public folders. The message stubs contain the body text and the links of the archived attachments. See Figure 2-20.

You can select this type if the archiving type is Entire, Component, or Attachment.

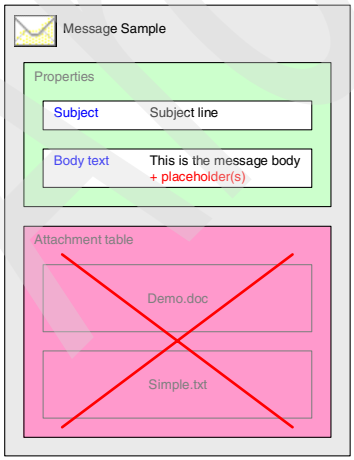

Figure 2-20 Deletion type: Attachment

Deletion type: Body

The message bodies are removed from the archived e-mail messages. This includes the attachments that a body might contain. Message stubs containing the header information (that is, the message properties such as Sender and Subject) are left in the mail boxes and public folders. Links to the archived messages are appended to the message. See Figure 2-21.

You can select this type if the archiving type is Entire or Component.

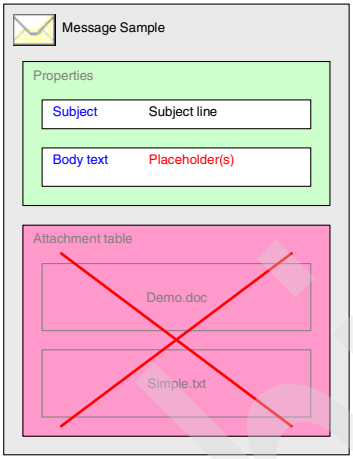

Figure 2-21 Deletion type: Body

Deletion type: Message

With this deletion type, all archived messages are removed. No message stubs remain in the mail boxes and public folders. See Figure 2-22 on page 39. Consequently, links to the archived content are not available.

Selecting this type prevents users from retrieving archived content directly. To retrieve the archived messages, an authorized user must first search the archive for the messages by using a search application.

You can select this type if the archiving type is Entire, Component, or Attachment.

Figure 2-22 Deletion type: Message

2.1.5 Retrieving archived content

When retrieving archived content, the data will be fetched from the archive and restored to your mail database. You can trigger this by selecting the e-mail stub and clicking the Retrieve button. Additionally, if you have configured CommonStore for Lotus Domino to insert retrieval hyperlinks in the document stub, you can retrieve archived content using that hyperlink. A click on a hyperlink shows the archived attachment in a viewer. This can be a Web browser or the viewer of the Lotus Notes client. You also can trigger retrieve requests from search result lists.

Note: Retrieving archived content involves the CommonStore Task to interact with and manipulate the mail database.

The following parameters have an impact on content retrieval:

- Deletion type = Message
- Archiving type = Entire
- Archiving type = Component
- Archiving type = Attachment

Deletion type = Message

When an e-mail is archived and the deletion type Message was used, there is no message stub to trigger retrieval from. In this case, the only way to restore archived content is to trigger retrieval from a search result list.

Archiving type = Entire

When an e-mail is archived using archiving type Entire and you trigger retrieval using the message stub, the e-mail is fetched from the archive and replaces the stub, thus restoring the original e-mail. If you trigger retrieval from a search result list, the behavior differs for the two CommonStore products. When using CommonStore for Lotus Domino, you can select whether you want to replace the stub by the restored e-mail or create a new e-mail. When using CommonStore for Exchange Server, the e-mail is restored as a new e-mail in the search restore folder.

Archiving type = Component

When an e-mail is archived using archiving type Component and you trigger retrieval using the message stub, all e-mail components are fetched from the archive. An e-mail is composed from these components and replaces the stub, thus restoring the original e-mail. If you trigger retrieval from a search result list, the e-mail is restored as a new e-mail in the search restore folder.

Archiving type = Attachment

When an e-mail is archived using archiving type Attachment and you trigger retrieval using the message stub, the attachments are fetched from the archive and reattached to the message stub. If you trigger retrieval from a search result list, a container document is created and the attachment is added to it.

2.1.6 Viewing archived content

When viewing archived content, the data will be fetched from the archive and displayed in your Web browser. You can trigger this by selecting the e-mail stub and clicking a hyperlink that CommonStore has added after successful archival. You also can trigger viewing requests from search result lists.

Note: Viewing archived content does not involve the CommonStore Task. The request is handled by the Web (HTTP) dispatcher.

The following parameters have an impact on viewing archived content:

- Deletion type = Message
- Archiving type = Entire
- Archiving type = Component
- Archiving type = Attachment

Deletion type = Message

When an e-mail is archived and the deletion type Message was used, there is no message stub to trigger viewing from. In this case, the only way to view archived content is to trigger viewing from a search result list.

Archiving type = Entire

When using CommonStore for Lotus Domino, viewing of an e-mail archived with archiving type Entire is not possible. When using CommonStore for Exchange Server, for e-mail archived using archiving type Entire, the e-mail is fetched from the archive and sent to your browser for display. As the browser cannot display the content, it will pass the request to Outlook, which has to be installed to display the e-mail.

Archiving type = Component

When using CommonStore for Lotus Domino, viewing from the message stub is not possible. When using CommonStore for Exchange Server, for e-mail archived using archiving type Component, you can view the individual components of an e-mail:

- The message body (without the attachments)
- The individual attachments

The results returned in a search result list for archiving type Component depend on the document storage model in use, because search requests in Content Manager always return documents. Hence a search request when using GENERIC_MULTIDOC can create one hit per e-mail component, but a search request when using GENERIC_MULTIPART will always create one hit per e-mail. Therefore, when using GENERIC_MULTIPART, only the first document part can be viewed. The first component of an archived e-mail usually is the body (without attachments).

When you select to view the message body, this component is fetched from the archive and sent to your browser for display.

Limitation: When using CommonStore for Lotus Domino, messages archived with archiving format Notes cannot be displayed. When using CommonStore for Exchange Server, the browser cannot display the content and will pass the request to Outlook, which must be installed to display the e-mail. Note that only the message body is displayed. Clicking an attachment icon will not have the desired effect; it results in an error message.

When you select to view an attachment, the associated content is fetched from the archive and sent to your browser for display.

Archiving type = Attachment

When you selected to view an attachment, the associated content is fetched from the archive and sent to your browser for display.

When using CommonStore for Exchange Server, the results returned in a search result list for archiving type Attachment depend on the document storage model in use, because search requests in Content Manager always return documents. Hence a search request when using GENERIC_MULTIDOC can create one hit per attachment, but a search request when using GENERIC_MULTIPART will always create one hit per e-mail (regardless of the number of attachments). Therefore, when using GENERIC_MULTIPART, only the first document part can be viewed.

2.1.7 Archiving options and policy

CommonStore offers *automatic* and *manual* archiving options. If using the automatic archiving option, CommonStore automatically archives e-mail based on predefined archiving policies. With the manual archiving option, users decide when and how to archive their e-mail.

Archiving policies contain archiving rules for e-mail message archiving. You can assign different archiving policies to different users' mail boxes. A default archiving policy applies to users' mail boxes that have not assigned a specific archiving policy yet.

For automatic archiving, the archive rules contain criteria that an e-mail message or its associated database must meet before CommonStore archives the e-mail. The criteria can be based on the size of the mail database, the size of a particular e-mail, how long the e-mail has been created, and e-mail message properties. Archive rules also include the archiving type of the e-mail, the deletion type, and the logical archive ID. You can combine multiple rules together and prioritize the rules to create an archive policy.

For manual archiving, the archiving rules define the archiving type, storage model, and deletion type.

2.2 CommonStore e-mail archiving solution planning

As part of the overall design and planning, you should review and understand the CommonStore e-mail archiving options as discussed in the previous section. These include:

- Archiving types
- Document storage model

- Deletion types
- Retrieval method
- Archiving options and policy

To plan for the solution, you must understand the business requirements and the purpose of the e-mail archiving solution.

Purpose of the e-mail archiving solution

Some of the key uses for an e-mail archiving system are to reduce the size of e-mail files, make it manageable, and improve system performance.

Without archiving, this can be achieved from a system management perspective by:

- Deleting mail past an arbitrary age
- Strictly limiting the storage size per user
- Limiting the size of attachments
- Utilizing individual archives

These measures can limit the ability of key users to quickly access critical documents that may be days, weeks, months, or even several years old.

Using CommonStore to archive e-mail provides a better way to achieve these goals.

Note: In addition to user e-mail files, IBM CommonStore offers an ability to effectively control the size of potentially any Lotus Domino database or Exchange message store.

There are many ways to implement an e-mail archive solution. As part of the planning and designing exercises, we provide a list of implementation options to get you started:

- Allow users to manually select which e-mail to archive. Users have the option to select which parts of the e-mail will be archived (the entire e-mail, the attachments only, or the body of the e-mail and the attachments).
- Establish an automated e-mail archive policy based on one or more aspects of a user's e-mail information, or the associated mail database.
- Decide whether to keep or delete any of the e-mail after it is archived.

When adding records management function into the e-mail archive solution, consider the following questions and issues (refer to Chapter 3, "Design and

planning for e-mail records enabling" on page 45 for records management related concept):

- Is there a requirement to retain all e-mail messages for a specified number of years? If so, all messages must be archived automatically.
- If a user is allowed to manually archive an e-mail, will the user later be allowed to declare the same e-mail as a record?
- If the legislation governing your organization specifies that messages and attachments must be declared as separate records, what document model should you use?
- If you specify an archive policy that deletes the e-mail after archiving, how will a user declare the e-mail or its attachments as a record?
- Before an e-mail can be declared as a record, it must exist in the archive repository, so either there must be a mechanism for users to manually archive e-mail or there must be an automatic archive policy in place.

Important: It is very important to understand your business requirement or legal requirement before setting up the e-mail archiving solution.

Much of the legislation governing how organizations set up their records management systems dictates that *all records are stored as separate objects*.

To meet this requirement, the combination of document model and archive type must be:

Archiving type = Component; document model = GENERIC_MULTIDOC.

This combination results in *every* component (e-mail body and each e-mail attachment) of an e-mail being stored separately within the Content manager repository.

To implement records management function into the solution, we examine Records Manager in the next chapter.

3

Design and planning for e-mail records enabling

This chapter covers design and planning for records enabling e-mail.

In the integrated solution discussed in this book, the e-mail records enabling is achieved with IBM DB2 Records Manager. In this chapter, we introduce the basic concepts behind records management and describe the basic concepts behind records management in the context of Records Manager. We address key areas to consider when planning and designing the e-mail records enabling (management) portion of a solution.

We cover the following topics in the chapter:

- Records Manager design options
- Declaration and classification: what is involved
- Records Manager design and planning considerations

This chapter mainly focuses on areas that are related to the integrated solution and is intended to be used in conjunction with the existing product manuals. To have an in-depth understanding of the individual product, we recommend referring to the existing product manuals.

3.1 Records Manager design options

Driven by compliance needs or risk management, organizations have a need to add records management capability to their business applications. In the e-mail archiving and records management solution we address in this book, Records Manager, working with Content Manager Records Enabler, Content Manager, and CommonStore for Lotus Domino or CommonStore for Exchange Server, adds the records management function to the e-mail archiving solution.

Note: The IBM DB2 Records Manager system has been used to records enable more than just e-mail systems. It contains the records management engine that you can use in conjunction with IBM DB2 Content Manager solution alone or with an IBM DB2 Document Manager solution. In this IBM Redbook, we focus our discussion on records enabling Microsoft Exchange and Lotus Domino mail databases using IBM DB2 Content Manager as the repository.

When used in the context of clear, consistent organizational records management policies, the records management controls are required to achieve one or more of the following:

- Reduced litigation costs and risks through structured document destruction
- Minimized discovery costs during litigation
- Ability to demonstrate compliance with organization and legislative regulations

Application of these controls is through an effective configuration of Records Manager. Records staff and administrators must ensure that the correct rules and policies are put in place and are carried out.

Records Manager offers a number of records management features and functions. These features and functions should be familiar to records administrators whether they use an existing electronic system or more traditional, physical records management system.

Like many electronic systems, the extent of the functionality exceeds what is required by any single organization. The capabilities of Records Manager must be reviewed carefully before implementing the system. When you understand what Records Manager has to offer, decide which functions must be implemented and how they should be implemented to meet your organization's records management requirements.

Many of these functions are driven by legislative requirements such as the rules that govern retention of all trading records and documents to satisfy SEC 17a-4

legislation. Here, all documents referencing share-trading actions must be kept as records. This extends to all e-mail correspondence and instant message traffic. Our solution is applicable in the case of e-mail retention.

Other functions may be needed to meet your organization's own procedures, such as the need to be able to quickly and efficiently locate information triggered by legal actions. An example of how Records Manager meets this requirement is in the way it allows the application and capture of metadata on what would otherwise be unstructured data such as e-mail.

To implement the records management function into your solution, it is important to understand the basic concepts of records management within the context of Records Manager.

In this section, we examine:

- File plan
- Life cycle (including retention schedule)
- Declaration
- Classification
- Disposition
- Security
- Physical records management
- Legal hold

This section provides an overview of the main concepts that are important in designing and planning the integrated e-mail records management solution. For more detailed information, refer to the following publications:

- *IBM DB2 Records Manager: Concepts Guide*, SC18-9182
- *IBM DB2 Records Manager: Administrator's Guide,* SC18-9180
- *IBM DB2 Records Manager: Technical Reference Guide*, SC18-9181

3.1.1 File plan

A *file plan* specifies how records are organized hierarchically in a records management environment. A file plan is similar to a collection of containers; a container represents a holding place into which you place records related to a common subject or theme, or another container together.

In Records Manager, you can create and manage flexible file plans according to your organization's business requirement. A records administrator can design any file plan based on unary or hierarchical object (records and container) relationships, define different user views and security policies, and establish relationships between objects. Virtually any record-keeping process can be implemented.

For many companies, the Records Manager's file plan maps to the companies' existing physical or electronic file organization scheme.

A file plan can support both the electronic and physical record-keeping of a company.

Figure 3-1 shows the file plan design of the scenario we use in this book. A file plan always starts with a Root component. Starting from the Root, we organize records based on Department, Region, and Division respectively. Within the Division, we organize records into two categories: eDocument or email. The declared e-mail falls into the email category (also called records component).

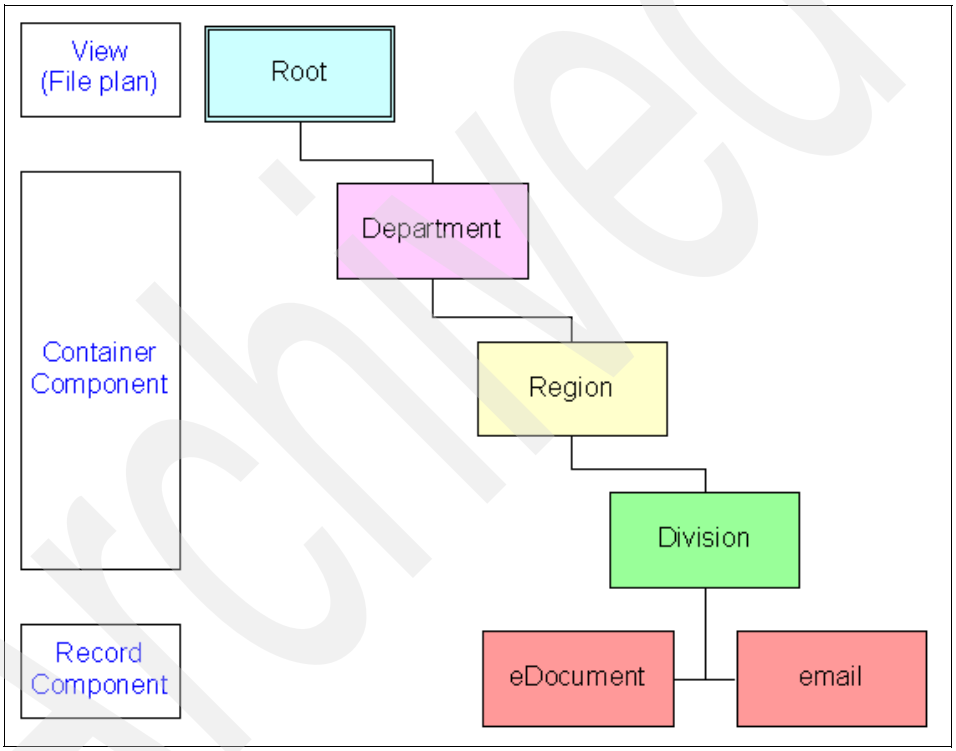

Figure 3-1 Sample file plan design

Important: If you plan to do policy-driven record declaration and are using CSX, the CSX Task requires that the record component be called *email*, all lower case and no hyphen.

File plan component definitions

A file plan is comprised of components. These components are instances of *component type definitions*.

Using an analogy from Content Manager, you can think of component definitions as item types, and components as items. From a relational database perspective, you can think of component definitions as table definitions, and the actual instantiated tables are components.

There are two file plan component definitions:

- System component definitions
- Custom component definitions

System component definitions come with Records Manager. They are used to instantiate system components that are needed for a Records Manager's file plan. All file plans come with a default system component definition, *Root*. It is at the start of your file plan, where you define the rest of your file plan components. See Figure 3-1 on page 48.

Custom component definitions are those you define according to your business rules. In Figure 3-1 on page 48, custom component definitions are Department, Region, Division, eDocument, and email.

Container and record component definitions

There are two types of custom component definitions:

- Container component
- Record component

A *container component* is one of the containers for records or other containers. For example, Department, Region, and Division from Figure 3-1 on page 48 are container component definitions.

A *record component* is the component type used for instances of a record. It differs from the container component in that it has content, such as the e-mail message and its attachments. For example, eDocument and email from Figure 3-1 on page 48 are record component definitions.

File plan components

After component type definitions are defined, you can create components from the component type definitions.

Figure 3-2 on page 50 shows the instantiated components that comprised the file plan we use in this book. From the Department component definition, we instantiate two components: Finance and Sales. From the Region component

definition, under the Finance component, we instantiate Europe, Asia Pacific, and America. Additionally, under the Europe component, using the Division component type definition, we instantiate Accounts Receivable and Accounts Payable. Actual e-mail messages and other documents will be placed under eDocument and email buckets.

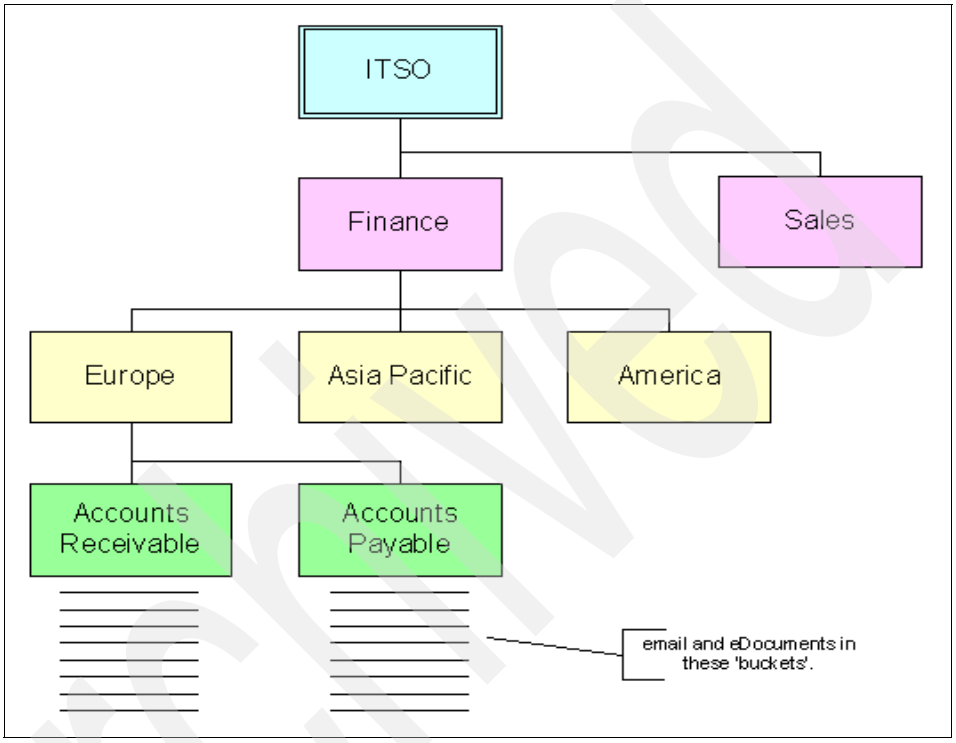

Figure 3-2 Sample file plan components

File plan view

A *file plan view* represents a grouping of component relationships. It helps you to navigate within a file plan.

Using an analogy from a relational database perspective, a file plan view is similar to a view in a database. All file plan components must have at least one view.

There are three types of file plan views:

- Hierarchical
- Link
- Set

To simplify the basic Records Manager concept we present here, we discuss only the hierarchical view in this book.

A *hierarchical view* represents a tree-like structure in a parent and child relationship. All file plans must have a hierarchical view. The file plan example that we introduced in this section (Figure 3-1 on page 48) uses the hierarchical view. The view can also represent a containment relationship. For example, a Department can contain multiple Divisions. A Division can contain multiple Regions.

3.1.2 Life cycle

Life cycle is a collection of phases a record must go through from the time it is declared as a record to the time it is disposed.

A life cycle can consist of one or more phases. Each *phase* specifies a certain duration and denotes a specific records management activity that must be performed at the beginning or end of the phase.

Retention schedule, also known as *retention rules*, specifies how long a record stays (is retained) in a phase and when the record transitions to the next phase.

A retention schedule is based on one of the following:

- Time
- Event
- Event time

If the retention schedule is based on *time*, then, after a specified time, a record will be moved out of the current phase and into the next phase. The time can be calculated from the time the record is added to the phase or from a life cycle date entered by a user or a program. Depending on the configuration and the setup, the time can also be calculated from the record's creation date.

If the retention schedule is based on *event*, then after the particular event has happened, a record will be moved out of the current phase and into the next phase.

If the retention schedule is based on *event time*, then, when a specified time has taken place after a particular event has happened, a record will be moved out of the current phase and into the next phase. Note, the time does not start calculating until the moment the particular event has taken place.

An example of a retention rule is to keep a record for three years from the time the record enters a particular phase. This rule can be applied to e-mail records to ensure that all e-mail will be retained in the system for at least three years.

The total length of each phase in a life cycle comprises the life cycle *duration*. The life cycle duration can be as short as a day or as long as hundreds of years.

In Records Manager, you can design and configure different life cycles for different classes of records, with multiple life cycle phases and phase transition parameters, including manual or automatic inter-phase transfer. The ownership and security of records can also change upon phase transition. This information should be derived from your organization's retention schedule.

We highly recommend that the records administrator *has a complete understanding* of the internal and external records rules that apply to various documents or e-mail.

3.1.3 Declaration

Declaration is designating a document, such as an e-mail, to be a corporate record. In our scenario, the declaration process is performed using the user's e-mail client, such as Lotus Notes client, Microsoft Outlook, or a CommonStore agent. You can declare e-mail including attachments or only the attachments as records.

Records can be declared by one of the following processes:

- Manual process
- Automated process

With *manual process*, the user decides when to declare a document as a record. The user sets a property or selects a menu option to declare the document as a record. In our e-mail archiving and records management solution, this is a user-triggered declaration of an e-mail, regardless whether the e-mail has been archived.

Note: In the e-mail archiving and the records management solution we address in this book, an e-mail or its component must be archived before the system can declare it as a record.

At the time of record declaration, the Records Enabler for Content Manager (CRME) checks whether the e-mail or its component has been archived. If it has not been archived, the system will automatically archive it to the content repository, waits for a document ID to be returned from Content Manager, and then declares the e-mail component as a record. If it is already archived, the system will declare it as a record immediately.

With *automated process*, a certain property triggers the automatic declaration of the record. In the e-mail archiving and records management solution, this can be

triggered with an e-mail component's *archive action* that is managed through CommonStore, such that the e-mail component is automatically archived and then declared as a record.

After a document is declared as a record, Records Manager removes any deletion or editing rights users previously have on the document. The edit and delete control of the document is now passed from the user to the record-keeping process, as administered by the corporate records management professionals. View access will only be allowed for users who are listed as having View access (in Records Manager) to the record. The declared record can only be modified or deleted in accordance with the organization's formal records management process now and not by the user.

Automated declaration (and classification) should be the *goal* of the organization. It should only be the exception that a small subset of records is manually declared. Organizations need to aim for *touchless records capture and keeping* and aim wherever practical to deploy automatic declaration features for various candidate record content.

We explain more about declaration in 3.2, "Declaration and classification: what is involved" on page 56.

3.1.4 Classification

Declaration and classification work together: You first declare a document as a record, and then you immediately classify the record.

Classification is assigning where in the file plan to classify a record, which then defines, usually via inheritance, the retention and disposition rules on the declared records. This is done by assigning the records to a particular file plan component (or bucket). Similar to declaration, classification can be completely manual or process-driven and can be automated, depending on your business requirement. By classifying an e-mail correctly, the correct retention rule is applied. One aspect of classification is capturing of metadata relevant to the document. This data may be captured automatically or it may be manually annotated by a user.

ISO 15489 (paragraph 7.2) gives the general characteristics of a record as: "a record should correctly reflect what was communicated or decided or what action was taken. It should be able to support the needs of the business to which it relates and be used for accountability purposes." The consequent definition of metadata given in ISO 15489 states: "data describing context, content and structure of records and their management through time."

Essentially, the metadata should present sufficient information about the document to classify it appropriately.

Figure 3-3 shows what happens to the document's metadata when a user declares and classifies it as a record: Its content remains within the host application's content repository, and the related metadata is stored in the Records Manager database. This metadata stored in the Records Manager database helps Records Manager in locating and managing the record.

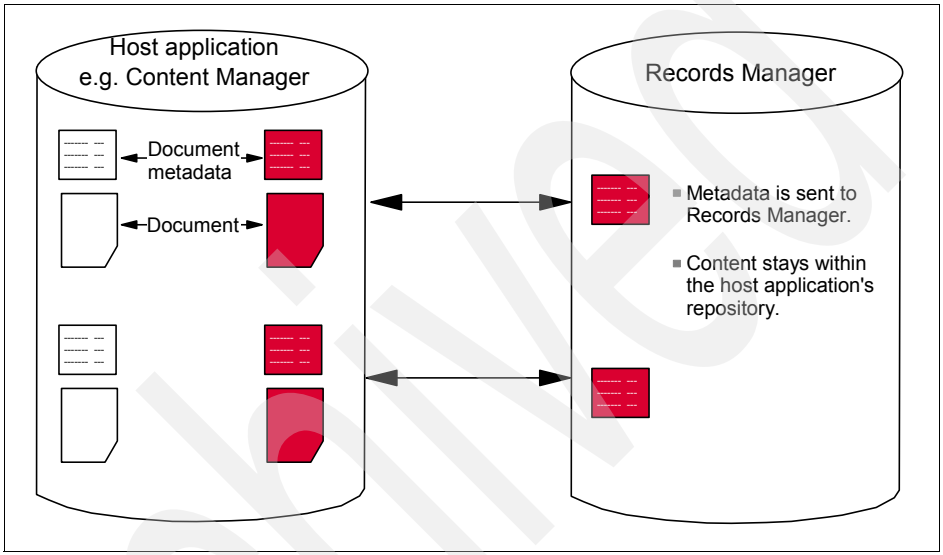

Figure 3-3 What happens when an e-mail is declared and classified as a record

One highly desirable capability of Records Manager is automatic classification.

Automatic classification

Automatic classification eliminates the need for users to manually assign retention rules (or more accurately select a record component from the file plan hierarchy the user is permitted to use). Records Manager offers metadata-based automatic classification, whereby the records administrator can define classification rules based on metadata of the record. If you plan on using automatic classification for your e-mail archiving solution, you should understand the lack (or potential lack) of suitable metadata. In the case of e-mail, it usually contains the From, To, and Subject fields. To have successful automatic classification, these fields or other metadata of the e-mail must contain sufficient information for the system to auto-classify.

3.1.5 Disposition

Disposition is the last stage in the record life cycle where records are disposed. Disposing a file plan component also disposes its descendants.

There are four ways to dispose a file plan component:

- Accession
- Destroy
- Export
- Review

Accession results in deletion of the record's metadata from the Records Manager database. In addition, it involves permanently transferring the record and its metadata to another authority that assumes responsibility and ownership of the record. This is done by the business application (in this case, Content Manager) making a copy of the record and its metadata to a specified directory and then deleting the record from its repository. The copied record content is kept elsewhere by the new record owner.

Destroy results in deletion of the record's metadata from the Records Manager database. If the record is in electronic format, Record Manager will inform the business application (for example, Content Manager) to delete the record. It is the responsibility of Content Manager to ensure this task carries through. If the record is a physical document, the Records administrator is responsible in overseeing the shredding or burning of the physical document.

If the disposition of a record is *export*, Records Manager translates the record into XML. When the record is exported, you must return to the task and proceed with the transition (clear the task to abort the destruction of the record).

If the disposition is set to *review*, then a records administrator must examine its metadata, its history, before deciding the record's final disposal schedule. If the final disposal schedule cannot be determined, the administrator can retain the record for further review and specify details about the review decision.

3.1.6 Security

Security is important to ensure that the records that are required to be kept are not deleted and the records that must be disposed are disposed properly by the right personnel. When a document has been declared as a record, only those users who have been assigned records permissions will have any access to the documents. Depending on the business needs, you can set up the system such that the original authors or receivers of the e-mail can no longer access their e-mail messages after they have been declared as records.

Records Manager enables you to control user access to the Records Manager functions and features, and user access to individual objects in a file plan. Access to Records Manager functions and features is called *function access rights*. Access to individual objects in a file plan is called *permissions*.

For example, the permissions that can be applied to an object in a selected file plan component include: add, delete, update, and dispose.

Refer to Chapter 4, “Security and user IDs” on page 67 for more information about the security of different components that make up this integrated e-mail archiving and records management solution.

3.1.7 Physical records management

The same underlying record-keeping infrastructure and processes that are applied to electronic e-mail also apply to manage the business’ traditional physical (paper) records. Depending on business requirements, you may need to track individual records, folders, and boxes; apply barcode technology; and manage physical storage space. Additionally, disposition schedules that are applied to an electronic record may also be applied to physical records, or at the same time as its associated electronic records.

3.1.8 Legal hold

Records Manager supports *legal hold* (sometimes referred to as a *suspension*, tax hold, or record hold) to designated records. Records under legal hold are protected from any possible destruction until the hold is lifted. This is usually driven by legal discovery litigation needs.

For more information, refer to 10.2, “Holds” on page 402.

3.2 Declaration and classification: what is involved

How declaration and classification are carried out is one of the key design decisions you have to make for the integrated records management solution. In this section, we take a closer look at what is involved in this process, and provide some guidelines to assist in designing your solution.

Within any records management system, there are two types of users: Those users who fully understand the concepts and those users for whom records management is a task that someone else performs. It remains a goal of many records management systems to streamline the whole process, particularly for users. This is because many users have little or no understanding of what constitutes a record nor do they wish to take time away from their work to

understand how to classify a document as a record. For these users, the greater the automation involved in records declaration, the better. This applies particularly to senior staff for whom a very real requirement exists that e-mail sent or received by these staff members must all be captured as records (depending on which legislation applies).

At the other end of the records management spectrum lay the professionals whose job it is to establish, own, and maintain an effective records management system whether they are paper-based or electronic-based. This staff fully understands what constitutes a record as well as the importance of why records need to be effectively managed. Classification of documents as records is a straightforward process.

A major challenge in the adoption of effective records management is how to manage the declaration and classification of records for all users, not just records professionals. The Records Manager–based solution enables declaration and classification techniques from auto-declare and auto-classify to a very manual declare and classification process. This *records management continuum* can be effectively applied to an e-mail archiving system such as CommonStore. Whatever level of automation is applied, *the goal is the accurate, repeatable, consistent application of retention rules to captured records*. Figure 3-4 shows this records management continuum.

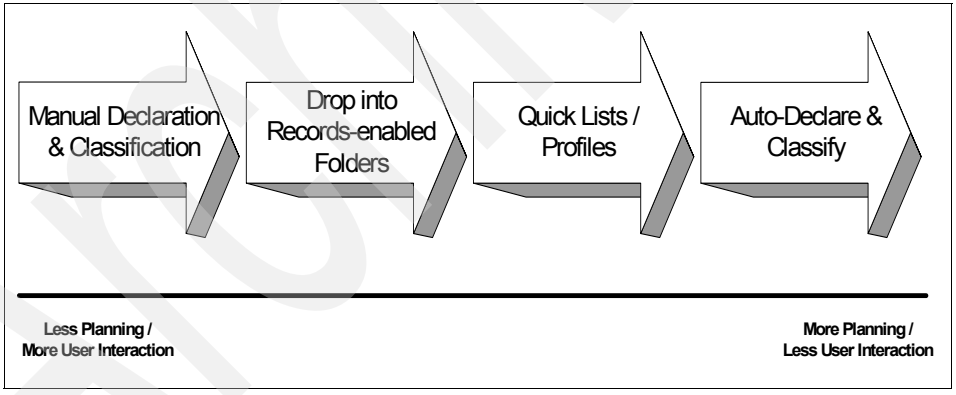

Figure 3-4 Records management continuum

The more manual the process of record declaration and classification, the less planning is required, yet more user interaction is involved and more record knowledge is required by the user. The more automated the process of record declaration and classification, the more planning is required beforehand for the entire system, yet less user interaction is involved.

Let us examine and compare the declaration and classification options:

- Automatic everything
- Manual everything
- Middle ground (quick list and foldering)

3.2.1 Automatic everything

To some organizations, the capture of every e-mail is vital to assist them in their legislative challenges. This process is often referred to as *journalling*. If this requirement is placed within the context of records management, all of this e-mail must be declared automatically as records, with no user intervention. The classification of these records must also be as automatic as possible as there would be no opportunity for classification by users. Otherwise, the administrative load on each user (to have to classify each e-mail) would soon lead to complaints from those users; others will attempt to shortcut the process. Notwithstanding the configuration required from the journalling component, it is not a good solution to place this classification load on the records administration staff. Your organization may send and receive hundreds of thousands of e-mail messages each day, so Manual classification of these documents is not an option.

Functions within Records Manager to apply auto-classification require planning and configuration ahead of time. You must decide:

- What auto-classification rules should be established.
- What metadata will be used in the classification process.
- How to handle documents that cannot be classified due to either no single rule satisfying a match or insufficient metadata available.

After auto-classify is configured, you can allow Records Manager to automatically locate the correct place in a file plan to put a new document. When you enable this feature, a user can automatically classify a document by clicking an Auto Classify button. The auto-classify feature is designed for use with business applications, and it is available for any record. When a user chooses to *auto-classify* a document, Records Manager reviews the previously configured auto-classification rules, locates all the rules that apply, parses through each one, determines whether the data matches the criteria, and then returns only those file plan components that match the criteria.

3.2.2 Manual everything

The other end of the scale is where the declaration and classification are both manual operations. An assumption here is that the classification of a record is made by the user rather than the records administrator.

Note: This is an example of a manual e-mail declaration policy designed to assist users to understand *when* to declare a record.

An e-mail should be declared as a record by you in these circumstances:

- When you are the author or creator of the e-mail.
- When you are the primary or the only recipient of an e-mail issued by someone external to the organization.

If the e-mail does not fall into one of these circumstances, then it should be deleted as soon as it is no longer required unless there is any other legal requirement to hold the e-mail.

If your organization allows users to perform manual declaration and classification, you must understand that the responsibilities placed on users must be carefully considered. Users may have no previous experience or skills in this area. Your organization must define and publish retention policies relating to the importance of e-mail and records within your organization.

There are two ways in which these tasks can be streamlined within an organization. Both of these options must be considered as they are not mutually exclusive:

- Presentation of a user or a group-specific minimum subset of the file plan to the user at the declare and classify time.
- Adequate and relevant user training.

Another method is to establish profiles.

Profiles

In Records Manager, a *profile* is a data entry form. Profiles are what users see when they add or edit an item in Records Manager. A profile contains one or more of the fields (attributes) that comprise the file plan component definition. You can create numerous profiles for almost any type of file plan component. This includes file plan components such as files, folders, and documents; it also includes system components such as users and groups. After you create a profile, you can assign it to specific users and groups or assign it to the Public group. Profiles let you filter or restrict user access to certain fields as well as limit the actions a user can perform on the fields that are included in the profile. When you create a profile, you can select the fields you want to include in the profile. This lets you control what a user can do in these fields. You can make fields read-only or mandatory.

3.2.3 Middle ground (quick list and foldering)

When auto-classification or manual classification are not suitable for an organization, Records Manager allows an alternative where a user wishing to declare a document as a record can do so using a combination of manual- and auto-classify.

This function is provided through quick lists or pre-configured folder on a user e-mail file.

Quick list

A *quick list* is a subset of file plan components presented to users to help them to quickly classify a document during the declaration of a record. It lets authorized users navigate to a subset of components without having to navigate through the entire file plan.

Foldering

The *foldering* feature enables you to preconfigure (or records enable) a folder in the user mail database to connect to a particular file plan component. When a user drops a particular e-mail into one of these folders, the e-mail is archived and declared as a record and is automatically classified.

The limitation of foldering in this way is that only a small set of file plan components can be exposed. If a large number of file plan components are exposed in this way, the list of folders may become unmanageable.

3.2.4 Comparison

There are advantages and disadvantages to using auto-declare and classification, manual-declare and classification, or somewhere in between. Knowing the trade-off and what is suitable for your organization helps you make design decisions.

Table 3-1 on page 61 summarizes the advantages and disadvantages of the two key methods of declaration and classification.

Table 3-1 Auto and manual declaration and classification comparison

Auto-declare and classify	Manual-declare and classify
Little or no user input. Little or no training is required for the users.	Require user input and training. More disruption to the normal daily workload.
For the documents covered by the subset of the file plan used, an organization can be assured that records management for these documents is taking place (assuming the correct configuration has been applied).	An organization cannot guarantee that users are declaring and classifying appropriate records, nor can they guarantee that the documents are declared and classified correctly.
Little opportunity for users to abuse the records management procedures. Users may not even be aware that records are being captured.	Users may take shortcuts and this will result in misclassified records or they may choose to not classify records at all. Re-declaration and classification may be required by records administrators.
Modifications to the file plan can be implemented with little or no user involvement, other than in the initial planning.	Whenever modification is done to the file plan, users must be informed and trained about the changes.
More up-front planning. More up-front configuration	Less up-front planning. Very little up-front configuration is needed.
Depending on how auto-classification is set up, these policies must be regularly reviewed by records administrators.	Records administrators do not need to review the setup. Less load on records administrators to reclassify incorrectly classified documents if users are well trained. Requires relevant training for users.
Some documents may not fall into any auto-classification rules and thus require manual classification.	Flexible declaration and classifying of records.

3.3 Records Manager design and planning considerations

This section is designed to present a summary of the information areas needing research required to set up a basic records management configuration.

Records Manager design and planning involves the following main steps:

- File plan design considerations
- Life cycle design considerations
- Users and security planning
- Records destruction planning

3.3.1 File plan design considerations

Establishing an effective file plan is crucial to both the ease of use by users as well as to any retention rules that need to be applied to records managed by Records Manager.

Before you begin the process of building a file plan in Records Manager, you must have an understanding of how your organization is currently organizing documents. An important step in file plan design and planning is reviewing your current file plan or how your files are currently organized within the organization.

Reviewing your current file plan or file organization

Every organization uses certain ways to organize its information. You may organize your information by files, by prefix, file, section, folder, and volume.

If your company has not yet designed a file plan, this task must be undertaken before a records management system is introduced.

Reviewing the existing file plan or existing file organization ensures that a records administrator:

- Is familiar with what records are being generated within the organization.
- Understands the relationship or hierarchy between groups of records and where records are folders (or containers) of documents.
- Is confident that all records generated within the organization are represented in the file plan.
- Understands who generates (or who will generate) records within the organization.
- Knows what metadata is required to be captured when records are declared and understands the technical capabilities of the systems in capturing the required metadata.

3.3.2 Life cycle design considerations

After you design the file plan, figure out the life cycle of records or file plan components. Legal staff in your organization should draft the retention schedules for the company's records according to company rules and legal obligations. You should review the current retention schedules.

Reviewing the current retention schedules

The *record retention (and disposition) schedule* for an organization should cover both physical and electronic records, and it should define:

- How long the organization needs to retain the information.

- How that information is disposed of when you no longer have to retain the information: destroy or accession to official archives? See 10.3, “Records disposition” on page 405 for more information about destroy and accession.
- The life cycle for each records series. For example:
  - How long is the record active?
  - How long is it dormant?
  - Where does it reside at each stage?
  - How does it transition between stages?
- Metadata associated with the records retention schedule. This includes:
  - Title
  - Item number
  - Description
  - Retention period
  - Method of destruction

Retention rules are applied to records through a formal, multi-stage process. With this process, the pre-defined retention and disposition rules and policies are applied to all of the declared records, so that only the relevant records are deleted at the appropriate time.

3.3.3 Users and security planning

In summary, the following reviews should be carried out to establish a more effective security policy for electronic documents:

- Establish user groups. It is easier to manage security with groups.
- Decide what groups need access to which files. This can be achieved through applying security permissions to components of the file plan.
- Decide what permissions user groups need.
- Decide what privileges records administrator staff needs. If an organization has a large records administration function, Records Manager can allow a granular application of security to the system. For example, only certain staff can modify the file plan while others may only have access to them.
- Decide what files need more security. Individual files can have their security set explicitly.

Some of the questions you need to address while planning for security include:

- What is your user population?
- Where are your user credentials managed?

What is your user population?

The drivers for implementing a records management system may determine what part of your user population requires access to the system. For example, your Human Resource department produces electronic documents that relate to employees. This type of document has to be kept for certain number of years. Typically, these documents would have been printed as paper copies and managed appropriately according to records policies. More often, these documents should be kept in their electronic form and managed by an electronic records-keeping system.

The legislation may affect multiple groups of users or, in the case of document discovery mitigation, all users. Identification of these groups will then lead to the following planning issue. (For more information about discovery, see Chapter 11, “Discovery” on page 409.)

Where are your user credentials managed?

Almost every organization considering an electronic records management system uses e-mail systems. Ideally, the e-mail messages are from a single vendor with no interoperability issues or multiple-user repositories. If this is the case, your user credentials should be located in a single directory system. We recommend that Records Manager takes its users from Content Manager. It is possible to create and manage users directly in Records Manager, but this would only be used for a few users, such as records administrators, as the load of maintaining additional users in multiple resource directories would be inconvenient. If Content Manager is the user resource for most Records Manager users, where are those user’s credentials managed? This depends on your environment. If your e-mail system uses an LDAP resource, this could be used by Content Manager, importing users and groups into Content Manager on a scheduled basis.

3.3.4 Records destruction planning

Most e-records standards require that the destruction of electronic records be non-recoverable or expunged. Records Manager relies on the host business application to perform the actual destruction, which Records Manager triggers and commonly carries out in non-recoverable form.

As an example of a records-keeping standard (the United Kingdom’s The National Archive), control of document destruction still requires some manual intervention as illustrated by this section from The National Archive (TNA) Requirements for Electronic Records Management Systems (2002):

```
A.4.64 (M) The ERMS must seek confirmation of destruction from an
authorized user as a mandatory step in the disposal process, before any
```

```
action is taken on folders, parts or records; and enable cancellation of
the disposal process at this point if confirmation is not given.
```

Part of your planning consideration must include what happened at the end of records life cycle: when and how the records should be destroyed.

This concludes the basic introduction to records management's basic concept, design, and planning considerations.

4

Security and user IDs

This chapter introduces the security models of the different components used within an e-mail archiving and records management solution. We provide recommendations on how an overall security model should be set up.

We cover the following topics:

- Content Manager security
- CommonStore for Lotus Domino security
- CommonStore for Exchange Server security
- Records Manager security
- Content Manager Records Enabler security
- Integrated solution security overview
- Important user IDs summary

This chapter does not discuss all of the security features of every component involved in the solution, but it introduces the basic concepts. For more detailed descriptions of each product security, read the product-specific administration guide.

4.1 Content Manager security

Content Manager functions as the backend content repository for the archived e-mail in the e-mail archiving and records management solution. When the archived e-mail messages are declared as records, they remain within the Content Manager repository.

To ensure proper user access to the repository, we examine the Content Manager security in this section.

4.1.1 Overview

Content Manager security consists of the following configurations and concepts:

- Privilege and privilege group
- Privilege set
- User and user group
- Item type and item
- Access Control List (ACL)

Privilege and privilege group

Privilege is the permission to perform an action within a Content Manager system. There are predefined privileges in Content Manager such as *ItemAdd* and *ClientImport*. A user with these privileges has the permission to insert an object into Content Manager (with the ItemAdd privilege) and import an object using a Content Manager client (with ClientImport privilege).

Privilege group is a logical grouping of privileges that makes the administration process easier. There are predefined privilege groups in Content Manager such as *ClientTaskCreate*. This privilege group contains privileges such as ItemAdd and ClientImport that gives a user the permission to add and import an object to Content Manager.

Privilege set

Privilege set is a collection of privileges and privilege groups. Only privilege sets can be assigned to users or ACLs. (Refer to the description of ACL that follows.) A privilege set describes the maximum set of privileges a user can have. Even if an ACL grants more rights, the user cannot use them, as long as the user's assigned privilege set does not have the privilege. There are predefined privilege sets in Content Manager, such as *ClientUserReadOnly*. This privilege set contains all privileges to grant read-only access to users.

User and user group

Every *user* is assigned a maximum privilege set. This privilege set describes the maximum rights a user can be granted in the system. For example, the assigned privilege set may specify that the user is an administrator or a normal user with restricted rights to create and view objects.

User group is a logical grouping of users. Privilege sets cannot be assigned to groups directly, but only to users. Privilege sets are only associated to groups within an ACL. This privilege set describes the maximum rights a user of this group can have when the ACL is used.

Item type and item

An *item type* is a template for defining and locating items (objects). It is the basic entity in Content Manager. By using the same template, items of the same type are consistently constructed, which helps you to locate them and quickly define new ones.

An *item* is an instance of any item type.

Using relational database as an analogy, you can think of item type as a database table, and items as entries in the table. For example, you can have *DominoMail* as an item type that is used as a template to store Domino e-mail messages. The users' e-mail messages are then stored in Content Manager as items, using the item type DominoMail.

Access Control List (ACL)

Access Control List (ACL) is a list consisting of one or more individual user IDs or user groups and their associated privileges. ACLs are used to control user access to objects in a Content Manager system (such as item and item types). A privilege set assigned to a user (during the user creation phase) defines the maximum ability that a user can have to the system. A privilege set associated with the user in an ACL further defines the exact user's access rights wherever this ACL is used.

An ACL does not grant additional privileges to a user. Even if the privilege set associated with a user in an ACL contains more privileges than the user has directly assigned to it, the ACL *limits* user access; it does not grant more access.

Figure 4-1 on page 70 gives an example of how the privilege set assigned directly to a user and a privilege set used in an ACL work together with a user ID that result in the rights a user has on a specific object (item type or item) that uses this ACL.

In this example, the user User 1 is assigned the maximum privilege set of ClientUserEdit. This means that in general, this user has the possibility to create,

view, and edit objects. For a specific item or item type's ACL in the figure, User 1 is assigned no privilege (NoPrivs). The ACL therefore limits the user to have no rights on the particular item or item type.

For User 2, the person is also assigned the maximum privilege set of ClientUserEdit. In addition, within the ACL of the specific item or item type, User 2 also has ClientUserEdit privilege. In this case, User 2 can perform create, view, and edit on the particular item or item type.

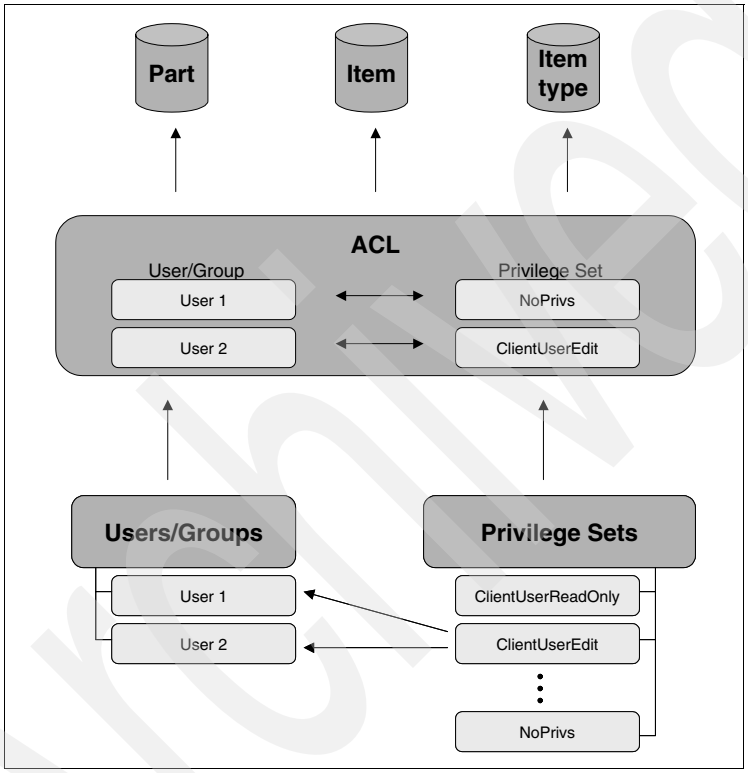

Figure 4-1 Users, privilege set, and ACL example

Important: Every item type has an ACL assigned to it. During the creation of an item type, you can choose whether the ACL works on the item type level or on the item level.

For a record enabled environment, with Content Manager V8.3, Fix Pack 1, ACLs should set on the item level.

4.1.2 Recommendations

To properly set up the security of Content Manager, you must understand the concepts behind Content Manager security. In addition, we provide some recommendations for the setup of the privilege sets, groups, and ACLs.

Note, these setup recommendations also reflect how the sample scenario we used in this book is set up.

Recommendations for defining privilege sets

Create privilege sets with reasonable names, so that it is easy to administer and use them. *Make sure that there are no super user privileges in it*; otherwise, ACLs will be ignored.

Examples of meaningful privilege set names are: CreateUpdateDelete and CreateUpdate. Privilege sets should be item type independent because they only describe a certain number of actions you can perform in a Content Manager system.

Recommendations for defining groups

One simple way to administrate a Content Manager system is to create one group per item type. This group is used within the item type specific ACL. By doing so, changes made to that group or to the ACL only affect the associated item type. This may dramatically reduce the risk of side effects when ACLs are changed.

In a large environment where there are many item types, however, this may not be the best administrative solution.

In general, the ACL should only include the group created for that item type and associate a certain privilege set that you want to allow for the item type.

Recommendations for setting up ACLs

Following the idea mentioned above, we recommend creating one ACL per item type. In this ACL, the item type specific group is associated with a certain privilege set. To change the access rights of an item type, it is recommended not to assign a new ACL but to change the item type specific ACL. Particularly if the access control is on item level, assigning a different ACL will not effect the items already stored in the item type. Instead, modifying the ACL of the existing item type will affect items stored in the future and items already stored. If the items have already been declared as records, they will not be affected by the ACL modification on the item type, as their ACLs would have already been reassigned.

4.2 CommonStore for Lotus Domino security

In this section, we examine CommonStore for Lotus Domino security. We cover the components involved from CommonStore for Lotus Domino and Content Manager. We discuss the communicators between these components, and how security is set up between Domino, CommonStore for Lotus Domino, and Content Manager.

To configure the security for a records-enabled solution, refer to 4.6, “Integrated solution security overview” on page 93.

4.2.1 CommonStore components

CommonStore consists of different components that are necessary to communicate with Domino on the one side and Content Manager on the other side. Figure 4-2 shows the key CommonStore components and specific details on both the Domino and Content Manager sides.

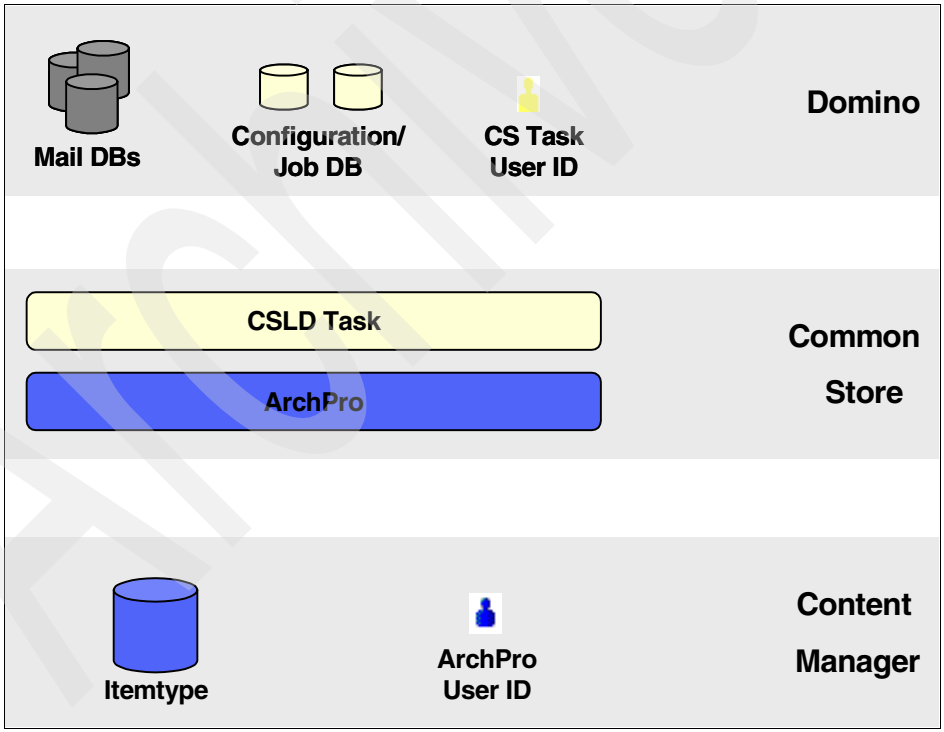

Figure 4-2 CommonStore for Lotus Domino components

We examine these components and the users involved in the following sections.

4.2.2 CommonStore and the Domino security

The *CommonStore Task* is a program that interfaces with Domino. It accesses a Domino system using the Domino APIs. It needs a Domino user ID to log on to the system. This user ID is usually set as *CSLD Task* or *CS Task*. It can be named anything else to suit your naming convention.

From the Domino side, as shown in Figure 4-2 on page 72, we deal with the access rights of this Domino user ID and the security on the Domino's mail databases, configuration database, and the job database.

Mail database security

The CommonStore Task needs to search through all mail databases that have e-mail to be archived. The CommonStore Task user (CS Task, for example) must be able to log on to the Domino system and have access to all mail databases that need to have e-mail archived.

Configuration database security

The CommonStore Task reads its configuration from the Domino's configuration database. The CommonStore Task user (CS Task, for example) must therefore have read access to the database.

A Domino administrator should have full access to the configuration database. There is no need for a normal mail user to have access to the configuration database.

Job database security

The CommonStore Task processes jobs in the job database, so the CommonStore Task user (CS Task) should have access to all jobs in the job database. The role [CSLDUser] is predefined in the job database and it can be assigned to the CommonStore Task user (CS Task).

Every e-mail user has to have at least the right to create a job (a Note document) in the job database. The security should be set up in the job database such that users can see jobs they themselves have created.

4.2.3 CommonStore and the Content Manager security

The *CommonStore ArchPro* is a program that interfaces with Content Manager. It accesses the Content Manager system using the Content Manager APIs and needs a Content Manager user ID to log on to the system. The ArchPro configuration file (archint.ini) specifies the Content Manager user ID that is used by ArchPro. The user ID has to have the privileges to store, update, and delete documents in the Content Manager system. This user ID is usually set as *ArchPro*. It can be named anything else to suit your naming convention.

From the Content Manager side, as shown in Figure 4-2 on page 72, we deal with the access rights of ArchPro and the security on the Content Manager's item types.

Item type security

At least one item type has to be defined for CommonStore to be able to archive e-mail into Content Manager. The ACL for this item type should include the Content Manager user ID used by ArchPro. Because e-mail and attachments from all Domino users are stored in this one item type (for example DominoMail), access to this item type should be limited.

CommonStore relies on its own security model based on Content Manager attributes (see 4.2.4, "CommonStore security model" on page 74). If a user searches e-mail from a Notes client, the CommonStore security ensures that only e-mail messages that belong to that user are visible.

If a user has access to that item type (DominoMail), searches using a Content Manager client returns all e-mail messages even if they do not belong to the particular user. This is because the Content Manager security model relies only on the ACL of the item or item type and the user's maximum privilege set to determine security; it is not aware of the CommonStore security restrictions.

Important: In an e-mail archiving *only* solution (without the records management functionality), it is advisable to only give the Content Manager user ID used by CommonStore access to the item type. No other Content Manager users should have access to the item type unless these users have a business need to access documents of the item type.

When configuring a solution to be records enabled, the security configuration should be set up differently. Refer to 4.6, "Integrated solution security overview" on page 93 for more details.

4.2.4 CommonStore security model

CommonStore uses its own security model to protect documents (e-mail and the attachments) from unauthorized access. The Content Manager security model is not used by CommonStore to determine which Content Manager user can access which document. Instead, the item type (that is used to store archived e-mail) has to have the attribute CSLDOrigUser or CSLDOrigDB. These attributes are filled during archiving (depending on the chosen security level) and interpreted during retrieval. The attribute CSLDOrigUser is used to store the Domino user ID that requested the archive. This is with limited use, as a user name can change. We recommend use of the attribute CSLDOrigDB that stores the Replica ID of the database a document originates from. No matter which of

these two attributes is used, it is important to understand that the attributes are not part of the Content Manager security model and therefore are not included during the authorization process to determine whether a Content Manager user has access or not to a specific item. The attributes are only interpreted by CommonStore to ensure that requests from Lotus Domino are served correctly.

The following paragraph explains the CommonStore security model based on a small example, shown in Figure 4-3.

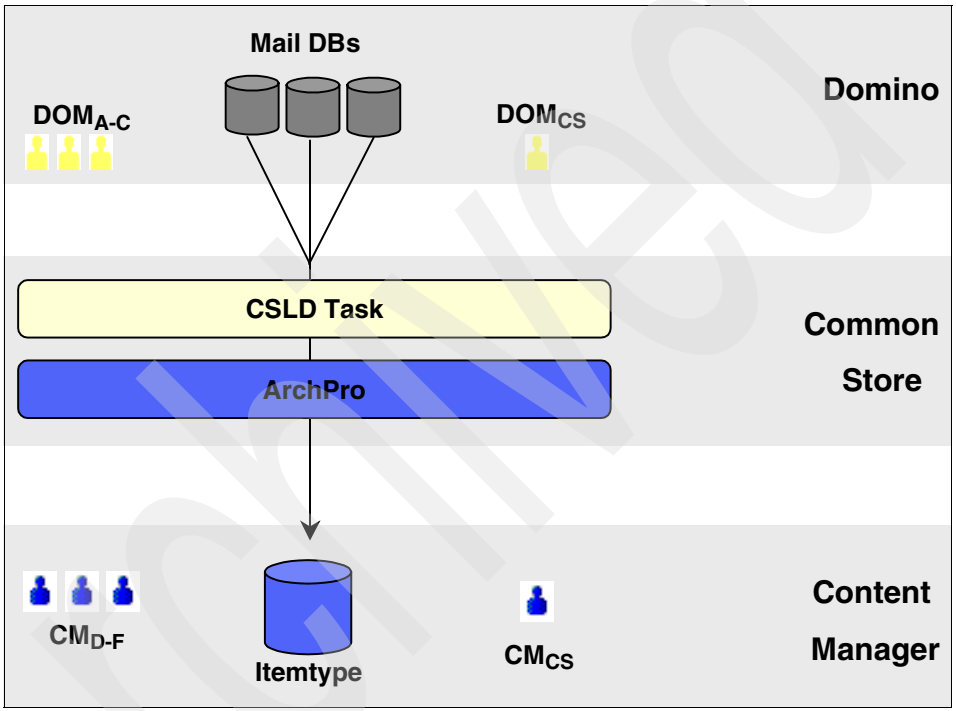

Figure 4-3 CommonStore security example

Assume that on the Domino side, there are three e-mail users (Dom_A, Dom_B, Dom_C) with one e-mail database per user and one Domino user ID used by CommonStore (Dom_{CS}). The Domino user ID (Dom_{CS}) has access to all three e-mail databases.

An item type (for example, DominoMail) is created to store archived e-mail from the three e-mail users. This item type has some standard attributes such as Sender, PostedDate, Subject, and the attribute CSLDOrigDB. The ACL assigned to the item type grants appropriate access to the Content Manager user ID used by CommonStore (CM_{CS}). This enables the Content Manager user to archive everyone's e-mail.

If an e-mail is archived from the user D_A's mail database, it is stored as an item of the item type (DominoMail), and the attribute CSLDOrigDB for this item is filled with the Replica ID from the user D_A's mail database. If e-mail messages from the users D_B and D_C are archived, they are stored with the same item type, and the attribute CSLDOrigDB is filled accordingly. In this example, three items are stored in the same item type, and there is one archived e-mail from each mail database.

If a mail user performs a search from its database against the archive (using the CommonStore function that is added to the Notes database), CommonStore searches the database using its Content Manager user ID (CM_{CS}), which has full access to all items of the item type. The ArchPro interprets the CSLDOrigDB attribute based on the Replica ID of the database that issued the search. Only documents that originate from the mail user's database are included in the hit list.

From the Content Manager side, only CM_{CS} can view the e-mail in Content Manager.

Now, for this example, we add three Content Manager users (CM_D, CM_E, CM_F) and assign them to the item type's ACL. In this scenario, these Content Manager users can see all e-mail of the item type using a Content Manager client. The Content Manager security model does not interpreted the attribute CSLDOrigDB and therefore does not limit access to the e-mail (documents). To avoid this access, the ACL for the item type should only be granted to the Content Manager user ID used by CommonStore (CM_{CS}).

Important: Again, in an e-mail archiving *only* solution (without the records management functionality), it is advisable to only give the Content Manager user ID used by CommonStore access to the item type. No other Content Manager users should have access to the item type unless these users have a business need to access documents of the item type.

When configuring a records-enabled solution, the security configuration should be set up differently. Refer to 4.6, "Integrated solution security overview" on page 93 for more details.

4.2.5 Recommendations

To properly set up the security of the integrated system, you must understand the concepts behind CommonStore security, and the interrelationship with Domino and Content Manager as described in the previous section. In addition, we provide some recommendation for the security setup of Domino, Content Manager, and CommonStore.

Note, the setup recommendations also reflect how our sample scenario used in this book is set up.

Recommendation for Domino security setup

Create a Domino user ID (CS Task) and grant access to all mail databases. This user ID requires the rights to create, update, and delete documents. Furthermore, this user ID requires read access to the configuration database and full access to the job database as well as the role CSLDUser on the job database.

Recommendation for Content Manager security setup

Create a Content Manager user ID (ArchPro) with the appropriate privilege set to create, update, and delete items (ClientUserAllPrivs).

Create an ACL specific for an item type (DominoMail), add the newly created Content Manager user (ArchPro) to it, and associate the appropriate privilege set to it (ClientUserAllPrivs).

Create the item type (DominoMail) that includes the attribute CSLDOrigDB and assign the newly created ACL to it.

This recommendation is valid for an e-mail archiving *only* solution. For an integrated records-enabled solution, you will set up the security differently. Refer to 4.6, “Integrated solution security overview” on page 93 for more details.

4.3 CommonStore for Exchange Server security

In this section, we examine CommonStore for Exchange Server security. We cover the components involved from CommonStore for Exchange Server and Content Manager. We discuss the communicators between these components, and how security is set up between Active Directory, Exchange, CommonStore for Exchange Server, and Content Manager.

To configure security for an integrated e-mail archiving and records management system, refer to 4.6, “Integrated solution security overview” on page 93.

4.3.1 CommonStore setup

During the setup of CommonStore for Exchange Server, the infrastructure is created. This includes:

- Extending the Active Directory Schema with objects defining the layout of CSX configuration data.
- Importing base configuration data.
- Publishing Outlook forms to your organizational forms library.

Extending the Active Directory Schema with objects defining the layout of CSX configuration data

The script CsxImportSchema.vbs is part of the Directory Extension Tool component of the installation package. It uses the program ldifde to import data defined in ldf files. Execute this script on the Domain Controller workstation using a user ID with Schema Administrator authority. See Figure 4-4.

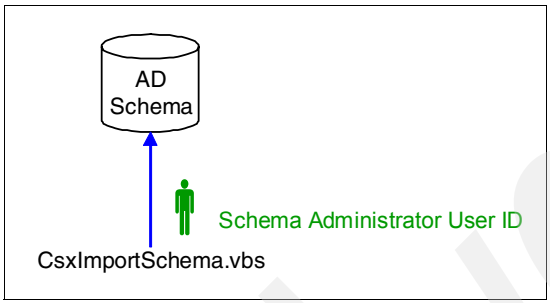

Figure 4-4 Extending Active Directory Schema for defining CSX config data layout

In our scenario we used the user ID CSX Admin, which we added to the Windows group Schema Administrators.

Importing base configuration data

The script CsxImportEmptyDirectory.vbs is part of the Directory Extension Tool component of the installation package. It uses the program ldifde to import data defined in ldf files. Execute this script on the Domain Controller workstation using a user ID with Domain Administrator authority. See Figure 4-5.

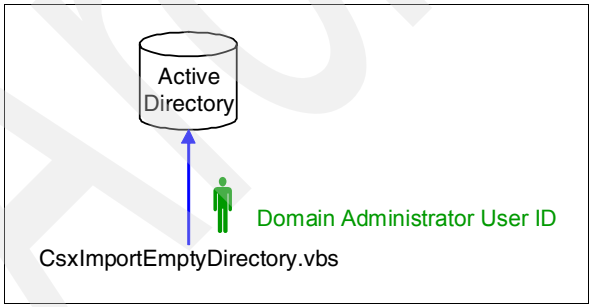

Figure 4-5 Importing base configuration data

In our scenario, we used the user ID CSX Admin, which we added to the Windows group Domain Administrators.

Publishing of Outlook forms to your organizational forms library

The InstallFormsAndFolders.vbs script is part of the Forms & Public Folders component of the installation package. It adds CommonStore Forms to the language-dependent Organizational Forms Library and creates CommonStore related public folders. See Figure 4-6, “Publishing of Outlook forms to your organizational forms library” on page 79.

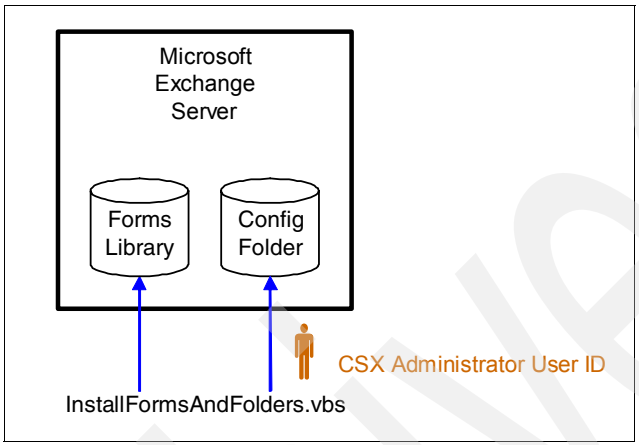

Figure 4-6 Publishing of Outlook forms to your organizational forms library

In our scenario we used the user ID CSX Admin, which we gave the Folder visible, Create items, and Read items permissions to the forms library.

4.3.2 CSX System Manager

Figure 4-7 illustrates the objects accessed by the CSX System Manager.

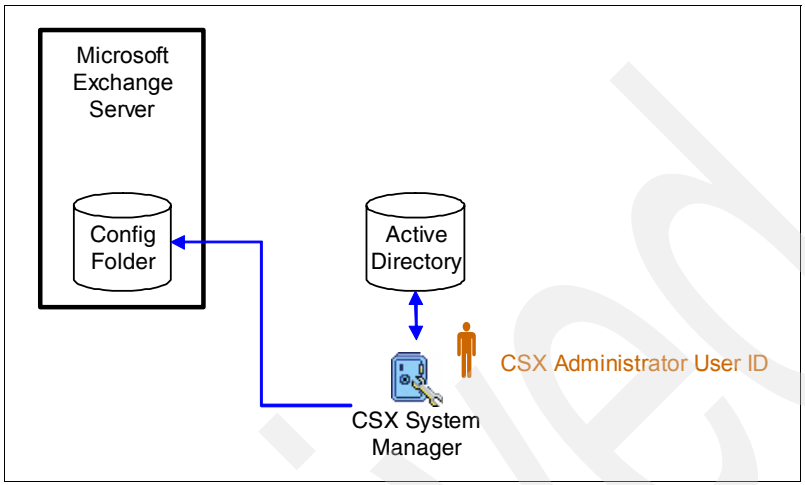

Figure 4-7 CommonStore for Exchange Server - System Manager

The CSX System Manager uses the Active Directory to store CSX configuration data and to provide it to the CSX Task. The workstation running the CSX System Manager must be in the same Windows forest as the Exchange organization. The user starting the CSX System Manager requires write access to the CommonStore node in Active Directory.

As the CSX System Manager also writes client relevant information to the CommonStore folders, the user starting it needs special permissions on the public folders CommonStore\Configuration and CommonStore\Job Folders.

In our scenario we used the user ID CSX Admin to run the CSX System Manager. As this user ID was also used to create the CommonStore infrastructure, no additional permission is required.

4.3.3 CSX Task

Figure 4-8 illustrates the objects accessed by the CSX Task.

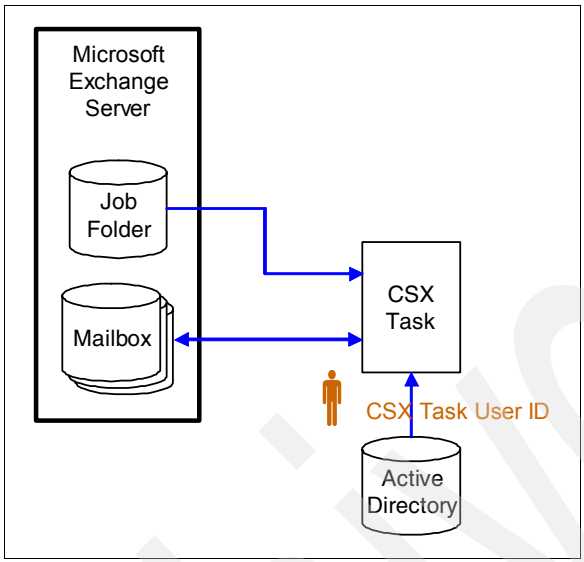

Figure 4-8 CommonStore for Exchange Server - Task

The CSX Task reads the CSX configuration data from the Active Directory. The workstation running the CSX Task must be in the same Windows forest as the Exchange organization. The user starting the CSX Task requires read access to the CommonStore node in Active Directory.

Note: The CSX Task does not read configuration data from the configuration folder in Exchange.

The CSX Task communicates with the Exchange system using the Microsoft messaging API (MAPI). It must have an Exchange user ID and a mailbox to log on to the system. This user ID must have write access to the job folder to read and update interactive archiving and retrieval requests.

The CSX Task must be able to search through all mailboxes and public folders that have e-mail to be archived. Hence it must be able to log on to the Exchange system and have access to all mailboxes and public folders without even knowing the credentials of individual users. To gain this authority, the CSX Task user ID must have the Exchange Administrator role on either the entire Exchange Organization or on all Administrative Groups that contain Exchange Servers to be archived. This role is assigned using the Delegate control context menu in the Exchange 2000/2003 System Manager.

In our scenario we used the user ID CSX Admin to run the CSX Task. As this user ID was also used to create the CommonStore infrastructure, no additional permission was required on the job folders. We gave the user ID CSX Admin the role Exchange Administrator on the entire Exchange Organization.

4.3.4 CommonStore Server

Figure 4-9 illustrates the objects accessed by the CommonStore Server.

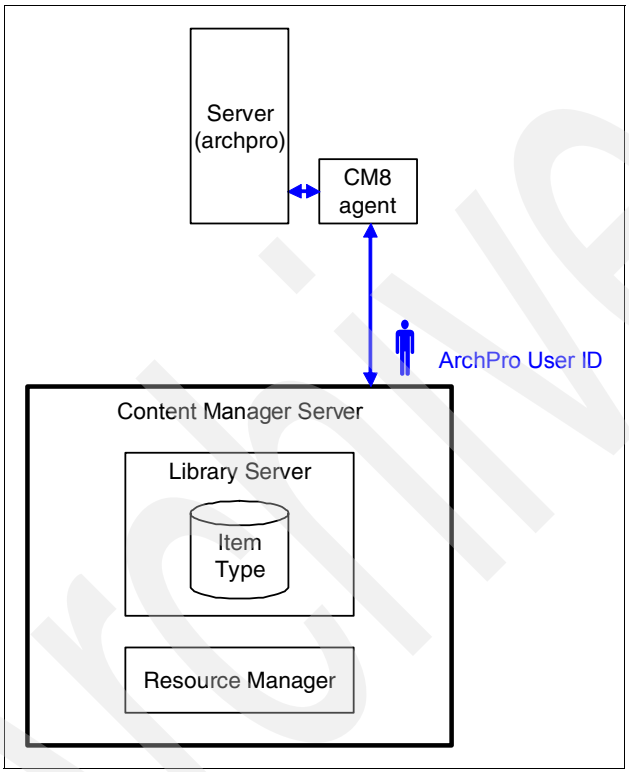

Figure 4-9 CommonStore for Exchange Server - CommonStore Server

The CommonStore Server communicates with Content Manager using the Content Manager APIs. It needs a Content Manager user ID to log on to the system. The ArchPro configuration file (archint.ini) specifies the Content Manager user ID that is used by ArchPro. The user ID has to have the privileges to store, update, and delete documents in the Content Manager system. This user ID is usually set as ArchPro. It can be named anything else to suit your naming convention.

From the Content Manager side, as shown in Figure 4-9 on page 82, we need to deal with the access rights of ArchPro and the security on the Content Manager's item type.

At least one item type has to be defined for CommonStore to be able to archive e-mail into the Content Manager. The ACL for this item type should include the Content Manager user ID used by ArchPro. E-mail messages from all Exchange users are stored in this item type, so access to this item type should be limited.

CommonStore relies on its own security model based on Content Manager attributes. (See 4.3.6, "CommonStore security model" on page 85.) If a user searches e-mail from an Outlook client, the CommonStore security ensures that only data that belongs to that user is visible.

If a user has the right access to the item type, searches using a Content Manager client returns all e-mail messages even if they do not belong to the particular user. This is because the Content Manager security model relies only on the ACL of the item or item type and the user's maximum privilege set to determine security; it is not aware of the CommonStore security restrictions.

Important: In an e-mail archiving *only* solution (without the records management functionality), it is advisable to only give the Content Manager user ID used by CommonStore access to the item type. No other Content Manager users should have access to the item type unless these users have a business need to access documents of the item type.

When configuring a solution that must be records enabled, the security configuration should be set up differently. Refer to 4.6, "Integrated solution security overview" on page 93 for more details.

4.3.5 Outlook clients

Figure 4-10 illustrates the objects accessed by Outlook clients.

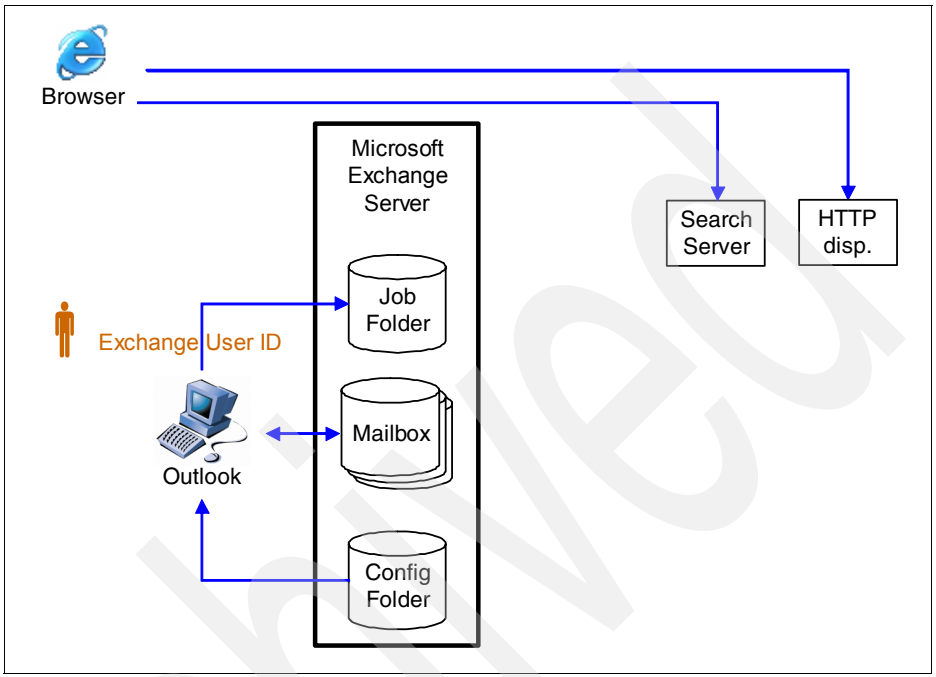

Figure 4-10 CommonStore for Exchange Server - Outlook Clients

Outlook clients read the CSX configuration data from the CommonStore configuration folder. They need read access to this folder to identify the name of the job folder as well as the host name and port number of the CSX Search Server. The job folder name and write access to it is required to trigger interactive archiving and retrieval requests. The CSX Search Server parameters are required to start a search session in the browser.

Note: Access to configuration folder and job folder is not required to view archived items from a Web browser. The URL included in the e-mail stub is sufficient to contact the CommonStore Server and display the archived e-mail or attachment. Moreover, no additional authentication is done. This allows access to forwarded archived e-mail messages and eliminates the need to restore them before forwarding.

4.3.6 CommonStore security model

CommonStore uses its own security model to protect documents (e-mail and the attachments) from unauthorized access. The Content Manager security model is not used by CommonStore to determine which Content Manager user can access which document. Instead, the used item type has to have the attribute CSORIGINATOR. This attribute is filled during archiving and added to search requests by the CSX Search interface, and is used to store the distinguished name (DN) of the user mailbox the e-mail resided in. It is important to understand that this attribute is not part of the Content Manager security model and therefore is not included during the authorization process to determine whether a Content Manager user has access to a specific item when working with Content Manager clients. The attribute is only used by CommonStore to ensure that requests from an Outlook client are served correctly.

Figure 4-11 on page 86 shows the CommonStore security model based on a small example, which is explained in the paragraph that follows it.

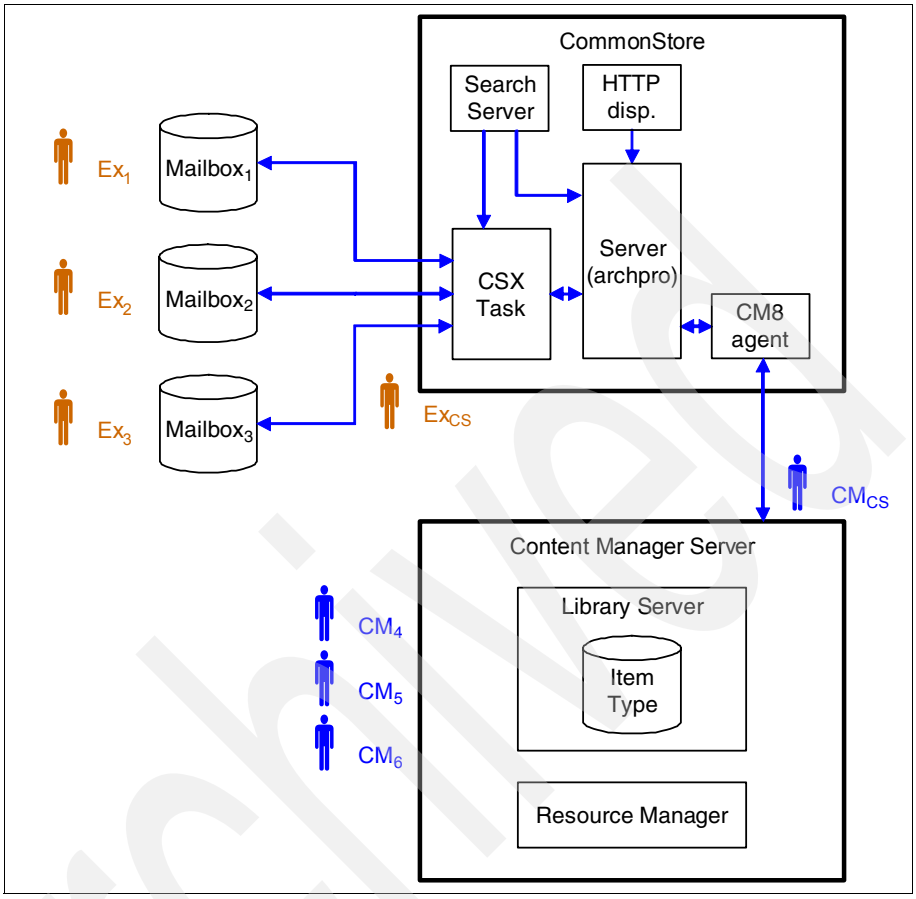

Figure 4-11 CommonStore security example

Assume that on the Exchange side, there are three e-mail users (Ex_1, Ex_2, and Ex3) with one mailbox per user and one user ID used by CommonStore (Ex_{CS}). The Exchange user ID (Ex_{CS}) has access to all three mailboxes.

An item type is created for CommonStore to store documents from the three mailboxes. This item type has some standard attributes such as Sender, Sent Date, Subject, and the attribute CSORIGINATOR. The ACL assigned to the item type grants appropriate access to the Content Manager user ID used by CommonStore (CM_{CS}). It enables this Content Manager user to archive everyone's e-mail messages.

If an e-mail is archived from user Ex_1's mailbox, it is stored as an item of the item type, and the attribute CSORIGINATOR for this item is filled with DN from user Ex_1's mailbox. If e-mail from the users Ex_2 and Ex_3 is archived, it is stored with

the same item type, and the attribute CSORIGINATOR is filled accordingly. In this example, three items are stored in the item type: one archived e-mail from each mailbox.

If a mail user performs a search from its database against the archive (using the CommonStore function that is added to the Outlook client), CommonStore searches the database using its Content Manager user ID (CM_{CS}), which has full access to the item type. The CSX Search Server modifies the query by adding the mailbox DN of the user that issued the search. Only documents that originate from this mailbox are included in the hit list. Privileged users can be defined to be able to search for e-mail originated from all mailboxes.

From the Content Manager side, only CM_{CS} can view the e-mail in Content Manager.

Now, for this example, we add three Content Manager users (CM_4, CM_5, CM_6) and assign them to the item type's ACL. In this scenario, CM_4, CM_5, and CM_6 can see all e-mail messages of the item type using a Content Manager client. The Content Manager security model does not interpret the attribute CSLDOrigDB and therefore does not limit access to the e-mail (documents). To avoid this access, the ACL for the item type should be granted only to the Content Manager user ID used by CommonStore (CM_{CS}).

Attention: Again, in an e-mail archiving *only* solution (without the records management functionality), it is advisable to only give the Content Manager user ID used by CommonStore access to the item type. No other Content Manager users should have access to the item type unless these users have a business need to access documents of the item type.

When configuring a records-enabled solution, the security configuration should be set up differently. Refer to 4.6, "Integrated solution security overview" on page 93 for more details.

This security model applies to searching for archived e-mail and attachments, so authorization for Web viewing is not checked by CommonStore. For example, if an e-mail is archived from the mailbox of user Ex_1, a URL is included in the message stub to allow user Ex_1 to view the archived message from the Web by clicking on the URL. If user Ex_1 forwards the message stub to user Ex_2, user Ex_2 can also use the URL in the forwarded e-mail to view the archived message from the Web.

4.3.7 Recommendations

To properly set up the security of the integrated system, you must understand the concepts behind CommonStore security, and the interrelationship with Active Directory, Exchange, and Content Manager as described in the previous section. In addition, we provide some recommendations for the security setup of Exchange, Content Manager, and CommonStore. The setup recommendations also reflect how our sample scenario used in this book is set up.

Recommendation for Exchange security setup

Create an Exchange user ID (CSX Admin) and use this ID for all configuration steps, to run the CSX System Manager, and to run the CSX Task. This avoids the need for additional permission changes on public folders created during setup.

Add the user ID to the Windows groups Schema Administrators and Domain Administrators. Grant the user ID permissions Folder visible, Create items, and Read items to the Organizational Forms libraries for the languages supported by your Exchange environment. Give the user ID the role Exchange Administrator on the entire Exchange Organization.

Recommendation for Content Manager security setup

Create a Content Manager user ID (CSX) with the appropriate privilege set to create, update, and delete items (ClientUserAllPrivs). Create an ACL specific for the item type (CSXMAIL), add the newly created user to it, and associate the appropriate privilege set to it (ClientUserAllPrivs).

This recommendation is valid for an e-mail archiving *only* solution. An integrated records-enabled solution requires a different security setup. Refer to 4.6, “Integrated solution security overview” on page 93 for more information.

4.4 Records Manager security

In this section, we examine Records Manager security in relation to the entire e-mail archiving and records management solution.

Some of the basic concepts were introduced in Chapter 3, “Design and planning for e-mail records enabling” on page 45. We repeat their definitions here for the continuity of the discussion and for your review.

4.4.1 Overview

Records Manager security is based on:

- File plan and its components
- Function access rights and permission
- User and user group

File plan and its components

A *file plan* specifies how records are organized hierarchically in a records management environment. A file plan is similar to a collection of containers, where a container represents a subject or a container into which records related to a common subject or theme are placed.

A *component* is a logical container that comprises the file plan and denotes a class of information. There are two types of components:

- *Component (container) component*: These are the physical or logical containers in the file plan. Examples of container-type components are: files, folders, departments, boxes, and floors in a building.
- *Record component*: These represent the actual records. They differ from container components in that they can have content. Examples of record components are: documents, e-mail messages, and illustrations.

In Records Manager, you can set security on any of the file plan components and file plan component definitions.

Function access rights and permission

You can control user access to Records Manager's application features, and user access to individual objects of a file plan. Access to Records Manager's application features is called *function access rights*. Access to individual objects in the file plan is called *permissions*.

Users are assigned to function access rights. The *function access rights* give the users permission to perform actions within a Records Manager system. There are predefined function access rights such as File Plan Administration. A user with this function access right has the *general* permission to insert records into a file plan. Whether the user can actually insert a specific record into a file plan depends on the system permissions (see description below) the user has to the specific file plan or a specific component of the file plan. For example, a user with the function access right File Plan Administration who does not have the appropriate system permissions to the component Division (our case study example) cannot create any records in any instance of this specific component.

Permission is the access rule that defines what a user can do to a file plan component and file plan component definition. There are predefined permissions such as Add. A user with the Add permission can add a file plan component.

Permissions can be set in three ways:

- Systemwide: Permissions for file plan component types. Using our case study example, you can set systemwide permission to component types Department, Region, and Division.
- Component level: Permissions for individual component instances. Using our case study example, you can set component-level permission on component instances such as Finance, Europe, and AccountsRec.
- Field level: Permissions for individual file plan component definition fields.

Figure 4-12 shows the permission list that can be applied to selected file plan components for a specific set of users.

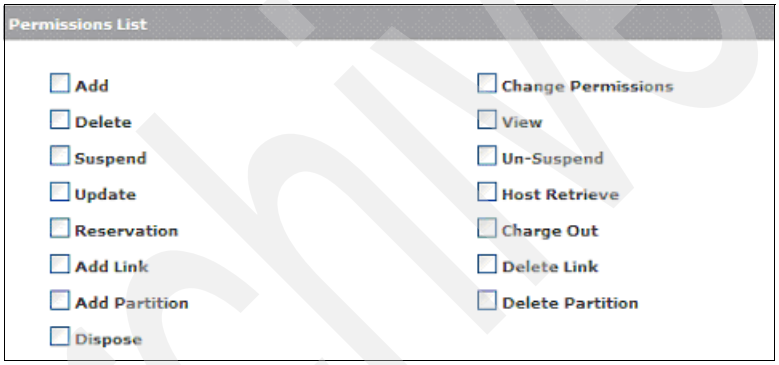

Figure 4-12 Permission available in Records Manager

User and user group

Records Manager differentiates between local and host users. *Local users* are defined within the Records Manager system. *Host users* are defined within a repository that is records enabled, such as Content Manager. In order to give Content Manager users the possibility to declare records, they have to be imported into Records Manager. During the import, function access rights are assigned to the users. After the import, the permissions have to be set for each user; otherwise the user will not be able to access the file plan.

4.4.2 Recommendations

To properly set up the security of the Records Manager system, you must understand the concepts behind Records Manager security as discussed earlier. In addition, we provide some recommendations for the Records Manager security setup.

These setup recommendations also reflect how our sample scenario used in this book is set up.

Recommendation for user and user group setup

For every Content Manager user who has to declare records, you must import a Content Manager user ID into Records Manager. This user needs the File Plan Administration function access right.

Recommendation for permissions setup

Every user who has to declare records into a specific component needs View access to all components mentioned above in the hierarchical file plan, and View, Add, and HostRetrieve access permissions to the component itself.

4.5 Content Manager Records Enabler security

Records Manager provides the records management engine that you can embed into your business application. Content Manager Records Enabler (CMRE) works between Content Manager and Records Manager to bring records management capability into Content Manager.

In this section, we discuss Content Manager Records Enabler security.

4.5.1 Overview

The Records Enabler does not have its own security model. It uses the Content Manager and Records Manager's security. One of the main functions of the Records Enabler is to transform Records Manager security into Content Manager ACLs. Content Manager Records Enabler has to have a Content Manager user ID with the appropriate rights to create new ACLs and item types. Figure 4-13 on page 92 shows the user IDs (CMREID and administrator) used by Records Manager to communicate with the connected systems.

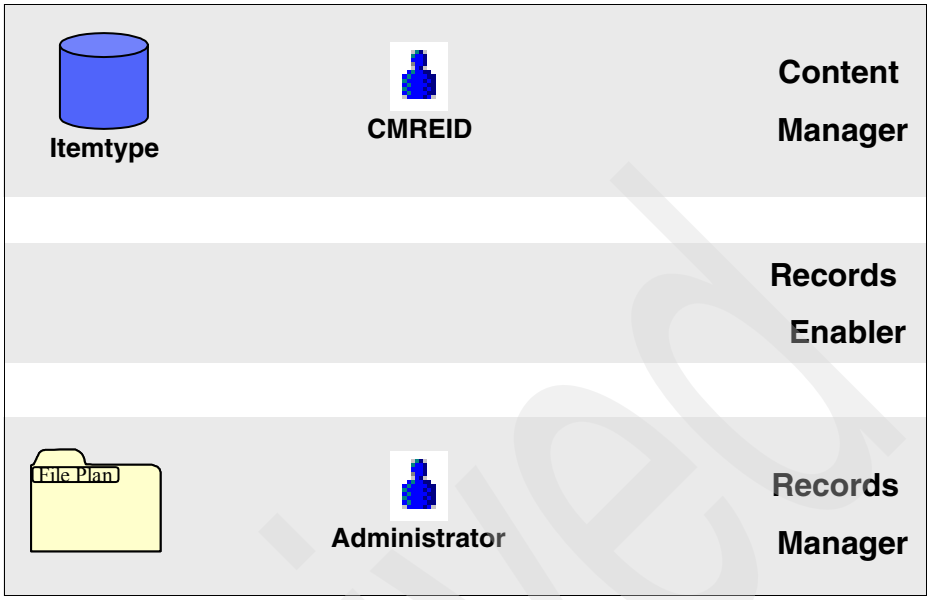

Figure 4-13 Recommended user IDs that work with Records Enabler

4.5.2 Recommendations

We recommend that you create the two user IDs as specified in Figure 4-13. For each ID, we describe what it is used for and provide some recommendations.

The setup of these IDs also reflects how our sample scenario used in this book is set up.

CMREID

The CMREID user ID is a Content Manager system administrator ID. The CMREID user ID is created in Content Manger during Records Enabler installation. Records Enabler saves this user ID in the host configuration of Records Manager, and uses it to configure Content Manager. For example, Records Enabler uses this user ID to create the needed item types, privilege sets, and ACLs when a Content Manager system is records enabled. This is also the ID that is stored in the WebSphere Data Source and later used to add (and remove) triggers to eRecord enabled item types. This user should not be imported into Records Manager. It is meant to be used under the covers by the Record Enabler code.

RMEADMIN

The RMEADMIN user ID should be used by the Records Enabler administrator to log on to the Records Enabler Administration client. Records Enabler does not

create this user ID automatically. This user should be created in Content Manager before the Records Enabler installation. The user does not need to have any explicit Content Manager permissions. As a minimum, this user can be assigned the privilege set NoPrivs and the ACL NoAccessACL. When this user logs on to the Records Enabler Administration client, the CMREID user will be used to perform any necessary Content Manager configurations.

After creating the RMEADMIN user ID in Content Manager, this user should be imported into Records Manager and given the System Configuration Management function access. This user does not require any other Records Manager permissions to use the Records Enabler Administration client.

4.6 Integrated solution security overview

The integrated e-mail archiving and records management solution is comprised of Content Manager, CommonStore for Lotus Domino or CommonStore for Exchange Server, Records Manager, and Content Manager Records Enabler.

In this section, we put all of the products together and take a look at the security for the entire integrated solution.

4.6.1 Overview

In the integrated solution, Content Manager is a records-enabled system. This means that content stored in Content Manager can be declared as records. After the content (e-mail) is declared as records, the data is under the control of Records Manager.

In the solution, the Lotus Domino and Exchange Server themselves are technically not records enabled. Every e-mail that has to be declared as a record must be archived (moved or copied) from Domino or Exchange Server into a records-enabled content repository system first (in this case, the Content Manager system). The Records Manager system communicates with the Content Manager system and not Domino or Exchange Server.

The e-mail users must have unique credentials in Content Manager and Records Manager to properly manage access to the e-mail messages that have been declared as records. This is achieved by associating a unique Content Manager user ID with each e-mail user. This unique Content Manager user ID has the proper access to the item type that is used to store archived e-mail in Content Manager. The Content Manager user ID also has to be imported into Records Manager so that the Records administrator can assign appropriate access rights to it.

The mapping of an e-mail user to the unique Content Manager user ID is managed via a user exit in the CommonStore server, which is installed as part of the Records Enabler configuration.

After installation of the user mapper support, an e-mail user is prompted for the associated Content Manager user ID when first opening his or her mail database (*if* the mapping is not already created for the user). This process adds the Content Manager user ID and password into the user mapping tables at the CommonStore Server and maps the Content Manager user ID to the e-mail user ID. Each subsequent access to archived e-mail (documents), whether a record or not, is authenticated based on the access control rules assigned to the archive. These access control rules are updated when an archived e-mail becomes a record to further restrict access based on the access permissions assigned to the record in Records Manager.

To retain access to archived e-mail, each Content Manager user must have the appropriate query privileges in the Content Manager's ACL assigned to the item type that is used for the archived e-mail messages. This, however, gives the user "query" access to all archived e-mail stored in this item type until the e-mail is declared as a record and the access privileges are modified for that individual e-mail. Although the mail client restricts access to a user's own personal e-mail (as described earlier when discussing CommonStore security models), Content Manager clients are governed by the Content Manager access rules.

Thus in an integrated e-mail archiving and records management solution environment, where both Notes (or Exchange Server) and Content Manager clients are used, precautions should be taken when assigning access rights to the item type.

One method to achieve greater privacy of archived e-mail and still allow for unique access privileges to records is to use a unique yet *private* Content Manager user ID for each e-mail user, known only to the CommonStore and Records administrators. You can create the private Content Manager user IDs, add the IDs to the user mapping tables at the CommonStore server, and associate these private IDs with the e-mail users. These private user IDs will be used for all authentication functions for the e-mail users.

Systematically creating these Content Manager user IDs based on the existing e-mail users and pre-populating the user mapper table simplifies the process for e-mail users as they no longer will be prompted for a Content Manager user ID when first accessing their mail database using the mail client. These private user IDs (and passwords) are not known to the mail users, so the mail users would not use the private Content Manager IDs to access the archived e-mail or records using Content Manager clients (if access to them is restricted to only these private users).

Figure 4-14 shows an example of mapping of the Domino or Exchange users to the Content Manager users in the integrated solution.

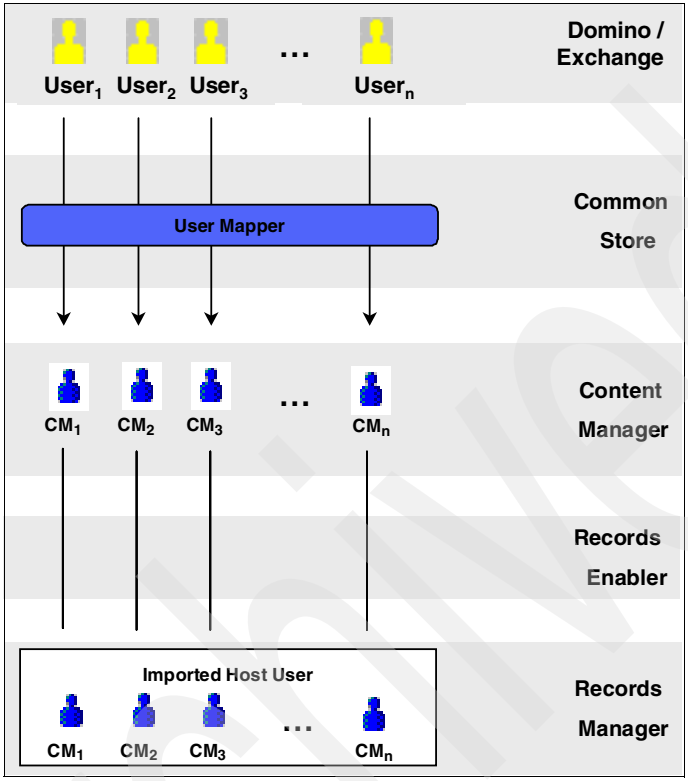

Figure 4-14 Domino/Exchange and CM user mapping in the integrated solution

In this example, $User_1$ is mapped to the Content Manager user CM_1. CM_1 is imported into Records Manager and has the appropriate rights to declare records in the file plan.

Note: A Domino or Exchange user who *is not mapped to a Content Manager user* cannot declare e-mail as records. Additionally, a Domino or Exchange user using a Content Manager ID *that is not imported into Records Manager* cannot declare e-mail as records.

4.6.2 Recommended configuration

In a non-Records configuration, it is assumed that a user always has access to the e-mail that he has archived. Thus access to the archived e-mail can be limited to the one specific CommonStore administrator ID in Content Manager.

CommonStore can then simply use the CSLDOrigDB attribute to filter the list of archived items to a specific user. In a Records environment, however, users may or may not have access to their archived e-mail based on the rules established in the Records Management system. Archived e-mail classified to certain areas of the file plan may be restricted to only certain users, or access to e-mail in a specific life cycle stage may be restricted. Because of the complexity of the access rules for “records,” filtering of user access must be done outside of CommonStore.

This is accomplished by setting the specific access to each item in the archive to specific users based on where in the file plan the e-mail is classified. Access is also appropriately updated in the archive as necessary when access rights change in Records Manager for each user or file plan Component. The access for each user is controlled by the ACL assigned to the archived e-mail.

Based on the previous discussion, every Domino user or Exchange Server user must have a Content Manager user ID mapped to it. The Content Manager user ID has to be imported into Records Manager and has to have the appropriate function access right and permissions to declare and retrieve records. This Content Manager user ID is used to assign user access to the file plan in Records Manager and is used in the Content Manager ACL on the archived e-mail to grant or deny access.

After an archived e-mail becomes a record, based on the rules configured in Records Manager, a user could lose access to it immediately or at some later phase in the life of the record. The ACL assigned to the e-mail item in the archive item type governs this access. Prior to becoming a record, the ACL must grant access to the archived e-mail to the user who archived it because there are no restrictions on its access yet. When CommonStore archives an e-mail, it simply imports the item into the archive item type. This results in the ACL assigned to the item type being assigned to the item imported. In the recommended non-Records configuration, this would mean that the CommonStore administrator would be the only CommonStore user with access to the item. However, this will not achieve the desired results in a records environment.

When an item becomes a record, user access to the item must map to the access defined in Records Manager for the file plan component to which the record is classified. New users added to the Content Manager ACL at this point would not be granted access to the item by Content Manager unless they also had access to the item type containing the item. This means that a new ACL assigned to the item now based on the access defined in Records Manager would not grant the necessary access unless these users were also listed with access in the archive item type’s ACL. Therefore, all mapped Content Manager user IDs would have to have read access in the item type’s ACL to enable further restriction of access to each declared item.

This can most easily be accomplished by assigning the Content Manager PublicReadACL to the archive item type; however, this would give all Content Manager users access to the archived e-mail items that have not yet been declared as records, as their ACLs would still reflect the item type's public access ACL. This is fine if Content Manager client applications are not used in your environment or are restricted to only those who should have access to all undeclared items in the e-mail archive anyway. Thus use of the PublicReadACL configuration should not be used if non-administrative users have access to Content Manager client applications in your environment.

Important: If using the PublicReadACL to control access, the CommonStore user ID has to be a super user; otherwise CommonStore cannot write any items into the item type.

An alternative configuration assigns every mapped user ID read access in the ACL assigned to the item type that uses private Content Manager user IDs for each e-mail user, as discussed previously. Thus every mapped user would have access to all archived e-mail, but e-mail users would not be aware of their mapped user ID and password and thus could not access the archive via Content Manager client applications. CommonStore would still access the archive on behalf of a CommonStore client (Notes/CSX), but again filters the results by the CSLDOrigDB attribute so that the user would only see e-mail he or she archived. As items in the archive are declared as records, the item's ACL is changed to grant access only to those permitted in Records Manager settings, thus further restricting the archived items returned to an e-mail user in a CommonStore client.

Note: A sample utility is available that simplifies the process of automatically creating private Content Manager user IDs for every e-mail user and pre-populating a mapping table with these IDs. It helps you to perform the following tasks:

1. Create Content Manager user IDs and passwords for each e-mail user. These private IDs are known only to the CommonStore and Content Manager administrators.
2. Add these users to a predefined Content Manager user group (a special group for this purpose). This group can then be assigned read access in the archive item type's ACL.
3. Add these users to the CommonStore user mapper table and associate them with the existing e-mail users.
4. Reset passwords for mapped users.

For more details, download the additional material associated with this book.

4.7 Important user IDs summary

In this section, we provide an overview of all of the important user IDs that are used for different components. You should understand the importance of these IDs with the affected components, especially when you decide to change the passwords of the IDs in the future.

User IDs for CommonStore for Lotus Domino

CommonStore for Lotus Domino uses the following IDs as described in Table 4-1 and represented in Figure 4-15:

- CSLD Task (Domino user ID)
- ArchPro (Content Manager user ID)

These names can be changed to suit your naming convention. We use these names in our scenario setup.

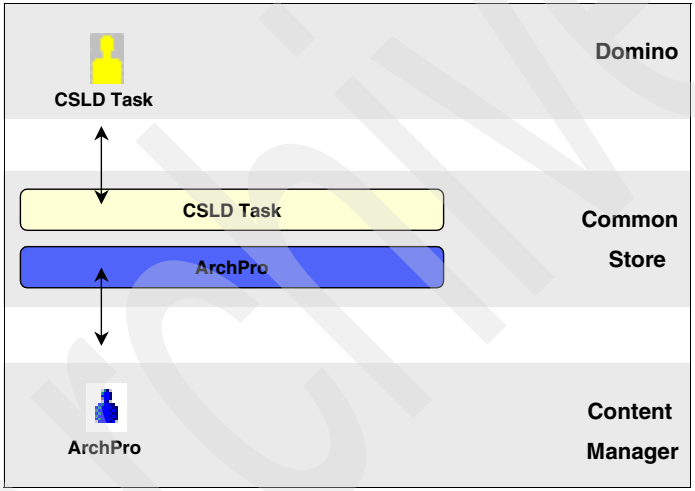

Figure 4-15 CommonStore for Lotus Domino user IDs

Table 4-1 CommonStore for Lotus Domino user IDs

System	Sample name	How to save new password
Domino	CSLD Task	From the command line, issue this command: `csld -f serverpasswd -i <notes.ini>`
Content Manager	ArchPro	From the command line, issue this command: `archpro -f serverpasswd -i <archint.ini>`

User IDs for CommonStore for Exchange Server

CommonStore for Exchange Server uses the IDs described in Table 4-2 and represented in Figure 4-16:

- CSX Admin (Exchange Server user ID)
- CSX (Content Manager user ID)

These names can be changed to suit your naming convention. We use these names in our scenario setup.

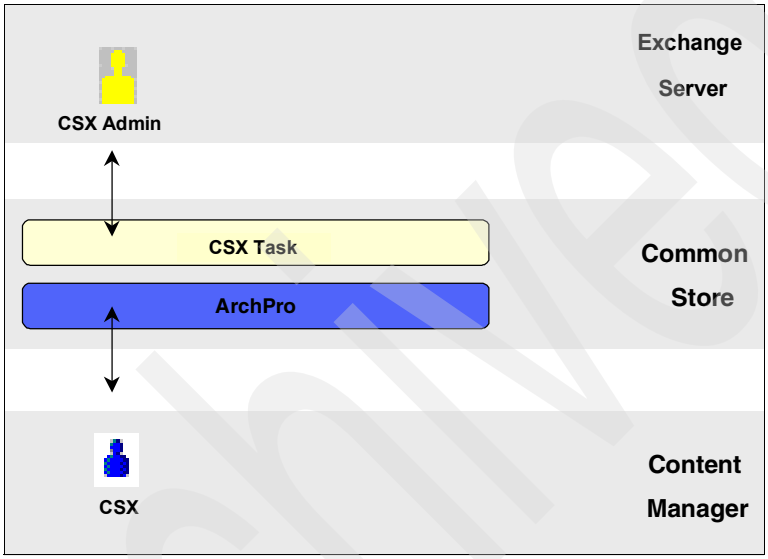

Figure 4-16 CommonStore for Exchange Server

Table 4-2 CommonStore for Exchange Server user IDs

System	Sample name	How to save new password
Exchange Server	CSX Admin	For the user CSX Admin used for CSX System Manager, password cannot be saved. For the user CSX Admin used for CSX Task, from the command line, issue this command: `csx <task_name> -p`
Content Manager	CSX	From the command line, issue this command: `archpro -f serverpasswd -i <archint.ini>`

User IDs for Content Manager

Content Manager uses the following IDs as described in Table 4-3:

- icmadmin
- rmadmin

These names can be changed to suit your system needs. We use these names in our scenario setup.

Table 4-3 Content Manager user IDs

System	Sample name
OS / DB2	icmadmin
OS / DB2	rmadmin

User IDs for Content Manager Records Enabler

Content Manager Records Enabler uses the IDs represented in Figure 4-17 and described in Table 4-4 on page 101:

- CMREID (Content Manager user ID)
- Administrator (Records Manager user ID)

These names can be changed to suit your system needs. We use these names in our scenario setup.

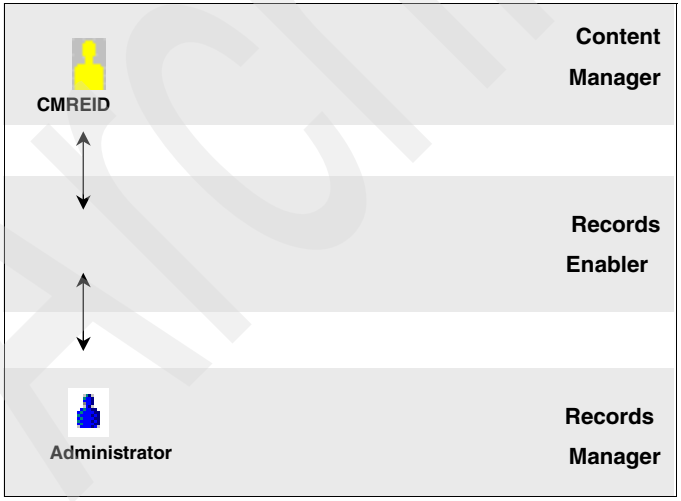

Figure 4-17 Content Manager Records Enabler user IDs

Table 4-4 Content Manager Records Enabler user IDs

System	Sample name
Content Manager	CMREID
Records Manager	Administrator

User IDs for Records Manager

Content Manager Records Enabler uses the following IDs:

- CMREID (Content Manager user ID)
- Administrator (Records Manager user ID)

These names can be changed to suit your system needs.

Table 4-5 lists the WebSphere administration user ID we used.

Table 4-5 WebSphere administration ID

System	Sample name
WebSphere	wasadmin

5

Integrated solution design and planning

This chapter discusses the overall design and planning of the integrated e-mail archiving and records management solution. To guide you through the design and planning process, we present a list of the questions and considerations that you should address during the process. In addition, we provide the end-to-end solution architecture, system configuration options, and the approach you may take to systematically implement and deploy the integrated solution.

We cover the following topics:

- Solution integration overview
- Planning considerations for the integrated solution
- System configuration
- Solution implementation and deploy sequence

The information covered in this chapter also includes the design and planning considerations presented from earlier chapters when dealing with only the e-mail archiving or the records management portion of the solution.

5.1 Solution integration overview

The e-mail archiving and records management solution architecture is open, portable, and extensible. Four main components are involved in the integrated system:

- Content Manager: Used as the backend server for content repository.
- CommonStore for Lotus Domino (CSLD) or CommonStore for Exchange Server (CSX): Used as the archiving engine that archives e-mail messages, including attachments to Content Manager.
- Records Manager: Used as the engine to provide records management function.
- Content Manager Records Enabler: Works with both Content Manager and Records Manager to provide records management capability to Content Manager, thus records enabling the entire integrated solution.

5.2 Planning considerations for the integrated solution

In this section, we look at the general design and planning considerations for the end-to-end, integrated e-mail archiving and records management solution. Because each organization has a different set of requirements and may have different approaches for this process, we recommend using the information provided here as a guide to help you start your design and planning process. In addition, we recommend using the existing product manuals in conjunction with this book during your design and planning process.

The considerations we cover include:

- General considerations for an integrated solution
- Security
- Other planning areas
- Key departments involved in planning process

5.2.1 General considerations for an integrated solution

There are many areas to consider when planning and designing an integrated e-mail archive and records management solution. You should always review the existing business requirements and legal requirements, working with legal staff and records personnel, to see how you should implement the solution.

We divide general considerations into these topics:

- “Archiving considerations” on page 105

- "Records management considerations" on page 109
- "Additional considerations" on page 114

Archiving considerations

From the e-mail archiving side, we examine the following questions:

- Why implement an e-mail archiving solution?
- What are the goals of e-mail archiving?
- Who are the target users or what are the target e-mail databases?
- How will e-mail be archived automatically?
- How will e-mail be archived manually?
- When will e-mail be archived?
- How to retrieve or locate archived e-mail or their attachments?

Why implement an e-mail archiving solution?

Many organizations implement an e-mail archiving solution before they implement an electronic records management system. One of the main reasons may be that historically, records management on the scale we are beginning to see these days was not a major driver for this technology. Many organizations are not covered (at least not extensively) by legislation requiring them to implement a records management system.

A common driver for e-mail archiving is size reduction on mail databases, regardless of whether records management is being considered.

The first thing to answer here is why you need to implement an e-mail archiving solution. Some reasons include:

- Improve mail database and server management.
- Improve system performance and user experience.
- Reduce mail database size, server size, or individual e-mail file size.
- Legal obligation; for example, need to retain e-mail for up to three years.
- Legal obligation; for example, need to declare e-mail as records.
- Some or all of these reasons.

What are the goals of e-mail archiving?

If the purpose of implementing an e-mail archiving solution is to reduce mail database size, improve system performance, or improve mail database and server management, or all of the above, then what are the specific quantitative goals you need to achieve?

To answer that, review your business requirement, and assess and compare that to the current mail database usage. Some of the statistics and behavior to investigate include:

- How many mail database users do you currently have?

- How many mail servers do you have?
- What is the average database size?
- What is the size of the mail database servers?
- What is the average number of mail messages in the mail database?
- How long do people usually keep their mail?
- What criteria do people currently use to clean up their e-mail?
- How often do users clean up their e-mail?
- What is the current response time? Are users using Outlook PST files that must be imported?

With this assessment, you can set specific e-mail archiving goals. Examples of specific goals include:

- Reduce the average mail database size to a certain number.
- E-mail older than a particular number of days should be archived and removed from the system.
- E-mail larger than a particular size should be archived, or its attachment should be archived.

If the purpose of implementing an e-mail archiving solution is to implement a records management solution, then you do not need to consider the specifics such as mail database size and server size.

Who are the target users or what are the target e-mail databases?

If the goal to implement an e-mail archiving solution is to reduce database size, improve system performance, and manage e-mail databases, you should investigate and decide who are the target users or what are the target e-mail databases:

- How extensive will the archiving system be?

  Will archiving apply to everyone's mail database or selective ones? What are the hardware and software implications?

- If e-mail archiving applies to only selected users, what is the criteria?

  You may have already answered this question from the previous discussion. For example, the criteria can be: for any e-mail database size that is greater than a certain number, its messages will be archived (based on a finer set of criteria). Certain users, due to the nature of their work, may have large volumes of incoming or outgoing e-mail, or large e-mail files; especially with large attachments, you may target these users for archiving activities.

- Will senior staff be included in e-mail archive policies?

  Some organizations choose to ignore C-level (in this example, CEO) staff mail databases for security reasons. However, this must be considered carefully in a record-enabled archiving environment as some legislation requires that *all* e-mail must be managed.

How will e-mail be archived?

Consider whether you want users to manually archive e-mail or let the system automatically archive it.

You might want some users to have the authority to manually archive their e-mail, and for other users or e-mail, you want to configure the system to automatically archive the e-mail.

There are different considerations and implications when you decide to set up automatic or manual archiving:

- How will e-mail be archived automatically?
- How will e-mail be archived manually?

How will e-mail be archived automatically?

If using automatic e-mail archiving, consider:

- Whose e-mail will be archived automatically?
- What rules will be implemented to archive e-mail automatically?

  Will these rules be based on the size of the mail database, the age of e-mail, or both?
- How will e-mail be archived automatically?

  What archiving types and storage models are being used? Will attachments be archived separately from messages or will they be stored as a single document? This has implications on how these documents are to be record enabled, compliant to regulations.
- When auto e-mail archiving, what will be left in the e-mail databases?

  Will there be message stubs for the entire e-mail, or just for its attachment? Will there be cases when the entire e-mail message be deleted from the system? For example, in the case where an e-mail is 365 days or older, the entire e-mail will be archived and no stubs will remain.
- What metadata will be collected at the time an e-mail is archived?

  This has implications on how e-mail can be found if there is no stub left in the original mail database.

Automatic e-mail archiving policy examples:

- Example #1: Perform automated e-mail archiving after a mail database reaches 200 GB. When this criteria is met, the system will automatically archive all e-mail older than 90 days.
- Example #2: E-mail 120 days or older will be archived. Only the stubs will be left in the original e-mail message.

- Example #3: If an e-mail is larger than 10 GB, its attachments will be archived and removed from the e-mail message automatically, and hyperlinks will be inserted that enable the user to view the archived attachment.

How will e-mail be archived manually?

If you will allow some or all of the users to archive e-mail manually, consider:

- Who are the users who can archive e-mail manually?
- What recommendations do you provide to users about how they should archive their e-mail?

  You may consider setting some guidelines for users to follow. For example, based on users' mail database size, the age of the e-mail, and the file size of a particular e-mail and its attachment, users should archive some or individual e-mail messages.
- How will e-mail be archived manually and what will be left in the e-mail database?

  When using CommonStore for Lotus Domino, how do you present the archive options for users? What do you allow and not allow users to do when performing the task? What archiving type will be used or offered to users to select from? Will you force users to use message stubbing for the entire e-mail, or just for its attachment? Will you allow users to or force users to delete entire messages from the system after they are archived?

  When using CommonStore for Exchange Server, will you provide users the choice to select the archiving type and deletion type?
- What metadata will be collected when users archive their e-mail?

  This has implications on how e-mail can be found if there is no stub left in the original mail database.
- What recommendations or guidelines do you provide to users who will perform manual archiving?

When will e-mail be archived?

For automatic operation, when do you run the automatic e-mail archiving task and how often do you run it?

For manual operation, this depends on users. As mentioned earlier, you should set some recommendations to users as to when they should manually archive their e-mail.

How to retrieve or locate archived e-mail or their attachments?

How will e-mail be retrieved, via hyperlink within the original e-mail message (only available when using CommonStore for Lotus Domino), the retrieve button, or a search application?

If you allow users or the system to delete the entire message after it is archived, you must provide a search function within the e-mail database.

If one of the reasons to archive is to control mailbox size, and whenever a document exceeds a certain size it will be archived, then you should design the system so that content will be removed after archiving. In most cases, attachments are the cause of large document sizes.

If e-mail archiving is based on a set criteria for a group of e-mail messages, and users have a requirement to be able to retrieve all of this type of e-mail at once, then a solution that enables users to select multiple archived e-mail message stubs and then click the retrieve button to retrieve all e-mail content may be a good option.

Another consideration is whether you want users to retrieve any archived e-mail. There may be circumstances when you do not want them to do so. You should plan this ahead of time.

Records management considerations

From a records management perspective, we examine the following questions:

- Why implement records management function to the solution?
- What are the goals of an e-mail records management solution?
- What is the target user mail database or e-mail?
- What file plan and life cycle to use?
- How will e-mail be declared and classified?
- What are the considerations associated with automatic process?
- What considerations are related to the manual process?
- When will e-mail be declared and classified?
- What happened to the e-mail when they are declared as records?
- How to retrieve or locate e-mail records

Why implement records management function to the solution?

Records-managing e-mail within an organization is becoming a more common requirement. Organizations may or may not need to implement an e-mail archiving system at the same time.

The question here is why you need to implement the records management solution. Some reasons include:

- Reduce litigation costs and risks through structured document destruction.
- Minimize discovery costs during litigation.
- Demonstrate compliance with organization and legislative regulations.

What are the goals of an e-mail records management solution?

What goals do you need to achieve when implementing a records management solution for e-mail?

The records professionals in your organization should review the organization's business requirements and legislative rules, and set specific goals for e-mail records management. Examples include:

- Records-manage and retain all e-mail for seven years.
- Records-manage e-mail from the Human Resource department.
- Records-manage any e-mail related to finance.

This will be used when designing file plan and life cycle of the records management system.

What is the target user mail database or e-mail?

Investigate and decide the target user mail databases or whose e-mail should be under records control. This should tie closely to the goals specified earlier by the records staff.

The following should be considered:

- How extensive will the records management system be?

  Will records management apply to everyone's mail database or selective ones? What are the hardware and software implications?

- If it applies only to selected users, what are the criteria?

  You may have answered this question already in the previous discussion. For example, if you need to records-manage the Human Resource department's e-mail, then your target e-mail or mail databases that should apply records control are ones from those who work in the Human Resource department.

- Will senior staff be included in records management?

  Some legislation requires that all C-level (in this example, CEO) staff e-mail must be under records control.

Make sure the target users's mail databases or e-mail that should be under records control match that of the goals set earlier. This should be reviewed by the records staff.

What file plan and life cycle to use?

Again, review the business requirements and work closely with the records staff in designing a file plan and life cycle for the e-mail records management solution.

Refer to "File plan design considerations" on page 62 and "Life cycle design considerations" on page 62 for more information.

The file plan design, life cycle design, and targeted e-mail (based on person, or e-mail type) need to comply to the legal rules or business requirement. Make sure the rules and design are documented, and reviewed and approved by the legal and records professionals.

How will e-mail be declared and classified?

Consider whether you want users to manually declare and classify e-mail as records or let the system automatically do the job.

It is possible that for certain types of e-mail, or certain people's e-mail, you want an automated process, and for others, it is a manual process.

There are different considerations and implications for automatic or manual records declaration and classification.

What are the considerations associated with automatic process?

If using automatic e-mail declaration and classification, consider:

- Whose e-mail will be declared and classified automatically?
- What rules will be applied to classify this e-mail?

  Auto-classification in Records Manager can be based on an e-mail's metadata. What are the rules or criteria used to match an e-mail's metadata with the specific file plan components?
- What if none of the rules apply to the e-mail?

  The metadata within an e-mail may not provide sufficient information for the auto-classification process. If auto-classification fails, the e-mail will reside in a collection of unclassified records. The records administrator staff has to classify them manually. What policy is established to handle this manual operation?
- How will e-mail be declared and classified?

  Will an entire e-mail be declared as one record? Or will the e-mail message be declared as one record, and the attachments be declared as separate records? The business requirement and legislation directly control what should be decided here. This affects the way e-mail should be archived. For example, if legislation requires that the e-mail message and its attachments should be declared as separate records, then you should choose the appropriate archive type and storage model when archiving the e-mail, and appropriate options to declare each component as a record.
- When using automatic e-mail declaration and classification, what will be left in the e-mail databases and who will have access to it?

  Will there be a message stub for the entire e-mail, or just for its attachment? Will there be cases when entire e-mail messages are deleted from the

system? You should consider both the archiving requirements and the records management requirement to see how to design the system that satisfies both areas.

It is possible that after the e-mail becomes a record, the rules will require that the original user can no longer access the e-mail. We provide more discussion about this in upcoming sections, but you should consider the implication here nevertheless.

- What metadata will be collected at the time an e-mail is classified?

  There are metadata associated with e-mail. Records Manager needs to store a record's metadata in its database. This metadata is comprised of the Records Manager's system data and some of the e-mail's original metadata. For automatic declaration and classification, which metadata information will be stored automatically to Records Manager? Will there be any metadata translation before they are stored in Records Manager? This has implications for how records can be found later.

What considerations are related to the manual process?

If you decide to allow some or all of the users to manually declare and classify e-mail, consider:

- Who are the users who can perform this task manually?
- When do you allow users to manually declare and classify records?

  You may need to automatically declare and classify records, and you may also allow people to manually declare and classify some other type of records. What rules govern this decision?

- What recommendations do you provide to users and under what circumstances should they declare and classify their e-mail?

  You may consider setting some guidelines for users to follow. For example, anything to do with finance or a particular project should be declared as records.

- What methods do you provide for users to perform e-mail declaration and classification? Will foldering, quick list, profile be used, or manual everything?

  How do you present the options to users? What do you allow and not allow users to do when performing the task?

  Will e-mail folders be configured to allow users to drag and drop mail into folders? Each folder can be configured to represent a particular bucket in the file plan.

- How much of the file plan will be exposed to each user if the foldering option is not set?

  Will the Records Manager quick list be used? Will Records Manager profiles be used?

- When manually declaring and classifying e-mail, what will be left in the e-mail databases?

  Again, what options will you allow or not allow users to have? Will you force users to use message stubbing for the entire e-mail, or just for its attachment?

- What metadata will be collected at the time users declare and classify their e-mail?

  This has implications for how e-mail records can be located later.

- What recommendations or guidelines, in general, do you provide to users who will perform manual e-mail records declaration and classification?

When will e-mail be declared and classified?

For the automatic operation, when do you run the automatic records declaration and classification? Will the e-mail have been archived already by an automatic process, or will the e-mail have not been archived yet?

For the manual operation, this depends on users. As mentioned earlier, you should set some recommendations for users as to when they should manually declare and classify their e-mail. Do you expect some users may have already archived the e-mail and some have not? Are there any implications? What are the rules you want to establish or guidelines users should follow as when to declare e-mail as records?

If using foldering, when e-mail is dropped to a folder, it will become records later. When should the process kick in: immediately or at a later time? What is the impact on the system? How should the design be carried out?

What happened to the e-mail when they are declared as records?

Will the original owners of the e-mail still have rights to view them? This must be configured in Records Manager, not CommonStore. The decision may depend on the type of e-mail. For some e-mail, users may maintain their access rights; for other types of e-mail, users may no longer be able to view them. This also may be determined by the type of users. Whatever is decided by the records professionals in your organization, you must be clear up front and set up security properly.

How to retrieve or locate e-mail records

How will e-mail records be retrieved? This depends on how the e-mail is archived. You may use the hyperlink within the e-mail, the retrieve button, or the search function within the user interface.

If users are allowed to search the e-mail records, how will they be searched?

Additional considerations

We provide additional considerations for your planning and designing process. They are:

- Will users need to archive e-mail without declaring it as records?
- Need to synchronize archive policies with records management policies.
- How will users locate documents and records?
- How to manage folders for e-mail archiving or records declaration?

The need to archive e-mail without declaring it as records

Will every e-mail have to be declared as a record? The answer to this depends on which legislation governs your organization's record-keeping requirements.

For those organizations required to declare every e-mail as a record, declaration of e-mail may take place from the journalling component of your messaging system (assuming that every user is covered by the relevant legislation as opposed to just senior staff). If this is case, you still need to consider whether to allow users to declare and classify e-mail or their attachments.

Deleting e-mail before it becomes records

Ensure that *all* e-mail that must become records will actually be declared as records whether you use the automatic or manual declaration process.

Is it possible that e-mail is deleted by users before it must be declared as records? CommonStore for Lotus Domino may be configured to automatically delete e-mail and its attachments whether or not the originals have been archived. This function is not available in CommonStore for Exchange Server. The possible early destruction of e-mail may conflict with organizational or legislative policies. If this is the case, configuration of automatic policies that delete the documents prior to their possible declaration as records or users that are allowed to manually delete e-mail prior to declare must be reviewed. You must identify the situations and decide how to prevent them.

Synchronize archive policies with records management policies

Depending on how e-mail archiving is set up, there should not be any conflict between archiving and records management policies.

For example, you establish an archive policy that archives the entire e-mail message and completely removes the message from the mail database after it is archived. If in Records Manager you have not set up any auto-declare and classification policies, then all e-mail archived under the archive policy above will not be declared as records (unless this task is performed manually by the user before this archive policy removes the message). After these messages are archived and removed from the mail system, the records administrator must manually declare and classify these records. Is this what you want? How can you avoid that?

Review all of your archiving policies and make sure they work with the policies you set up in Records Manager.

How will users locate documents and records?

After e-mail is declared by users, they may still need to view the e-mail content, its metadata, or its record state. This is an issue particularly if the archive policy includes removal of the e-mail and its attachments from the user's mail database. Unless the search function is retained in the e-mail client (Records Manager provides this function for both Lotus Notes and Exchange systems), the client software a user can use to search for and view the record are:

- The Records administrator client: This browser-based client can be accessed by users only if they are given the required functional access by the Records administrator.
- One of the Content Manager clients (Windows client or eClient): They must be distributed to users in an e-mail archive and records management solution. Using this client requires Records Enabler for Content Manager software to be configured.
- A federated search client such as Information Integrator for Content.

Note that any access to e-mail in Content Manager that has been declared as records will be determined by the permissions granted by the records administrator using the records administrator client.

How to manage folders for e-mail archiving or records declaration?

How do you manage folders? How do you manage the changes in folders? How do you propagate the changes? What happens when there are too many folders to manage?

Other considerations

We add some other questions and areas to consider during the design and planning process. They may repeat some of the questions asked earlier in this chapter or other chapters, but we want to make sure that you are aware of them and recapture them here.

Other considerations include:

- When using CommonStore for Lotus Domino, e-mail can be archived, retrieved, updated, and re-archived. Records can also be declared multiple times. Do you want to offer these options and how will you offer the options? Note that re-archiving is not available in CommonStore for Exchange Server.
- Does all e-mail have to be retained and declared as records? If so, do you need to worry about records management side only, or also mail database size and system performance? If you only have to implement records control, you may want to keep e-mail as is after archiving (that is, without removing any parts of the e-mail from the original database). Otherwise, decide what and how to remove certain parts of the archived e-mail, and how retrieval will be provided.
- What about e-mail journalling? Should you use it to ensure that all e-mail is captured? What is the limitation? Find out before using it.
- Be aware that some systems cannot archive an e-mail twice; for example, this is the current limitation of the CommonStore for Exchange Server. Because of this limitation, at the time of the writing, we do not recommend using the journalling feature when using CommonStore for Exchange Server.
- Do you need single instance store (SIS)? What is the current limitation when using single instance store in your environment? What is the impact or limitation when using it in conjunction with Records Manager?
- How can you ensure that the necessary e-mail is made into records? What criteria do you use to ensure no loop holes exist for the e-mail?
- Will people have rights to all e-mail, or should someone have rights to search on all users' e-mail databases? Records staff may need that authority.

5.2.2 Security

When implementing Records Manager into an e-mail archiving system such as CommonStore, the security of archived documents changes as soon as the e-mail is declared as records. The access users have to their e-mail after declaration depends on the permissions set by the records administrator.

When an e-mail is archived in a non-Records Manager environment, the user generally maintains access to that e-mail. In the Records Manager environment, access to the declared e-mail record is modified at declaration time. The modified access may (but not always) have the effect of denying the original owner of the e-mail any access to the e-mail after it is declared as a record. It is therefore important to review how security affects user access to e-mail before and after it is declared as records.

The questions below should be considered before implementing the solution:

- Should a user have access to an e-mail after it is archived?
- Will any element of an e-mail remain in the user's mail database post-archive?
- Can a user perform a manual record declaration on an e-mail?
- Do users have a minimum of View access to relevant sections of the file plan components in Records Manager?

> **Note:** If the user is not explicitly listed or is not a member of a group in a Records Manager system with a minimum of View access to relevant portions of the file plan, e-mail archived in CommonStore and subsequently declared as records may not be available to the user post-declaration.

A user with View access to the declared e-mail may access the it through the e-mail client, as with a normal e-mail.

For more considerations on security, refer to Chapter 4, "Security and user IDs" on page 67.

5.2.3 Other planning areas

Other areas related to planning should be addressed before you implement the entire solution:

- Performing audits on your mail system
- Planning for discovery
- Planning for document deletion

Performing audits on your mail system

You should audit your existing mail system prior to implementing the integrated e-mail archiving and records management system. The data captured at this time enables you to understand the magnitude of any problem and lay down base data that can be used to compare with post-implementation data to see what improvements are needed.

It is possible that your reasons for implementing an integrated e-mail archiving and records management solution is not for system improvements. It may simply be to comply with relevant legislation, whether or not that compliance brings any other system benefits.

Some of the audits you perform can be automated and can be accomplished using a variety of software tools that exist in the market today. Some of the audits can be in the form of staff and user interviews.

The audits you perform can include identifying:

- Average mail database size.
- Volume of inbound and outbound e-mail for a given time period.
- The largest users of the system in terms of mail database size and usage.
- Age of e-mail that exists in people's mail database. This is useful to understand how relevant older e-mail may be to user's current work.
- Average time users spend managing their mail files and if this is due to any mail database quotas imposed on their databases or in searching for e-mail.
- Why e-mail is kept for long periods of time. Can it be because a user is keen to cover themselves in case of subsequent inquiries?

Planning for discovery

Understanding how to manage a litigation-driven document discovery process before it happens saves a great deal of time *and* money. The negative effect of a document discovery process can be mitigated by ensuring that your organization has effective procedures around the way documents, e-mail, and data are stored.

Document discovery events can be very costly exercises that potentially fail for any of a number of reasons. Usually the failure occurs due to a lack of *advanced* planning — legal staff descending on the IT/IS department with demands for data and documents to be produced quickly around key terms causes disruption. A large amount of business may be transacted through e-mail systems.

Planning for document discovery can take the following form:

- Understand the range of data your organization holds.

  This can be both electronic and paper-based.
- Locate the repositories of data.

  This can include mail databases, file directories, other databases, backup systems (including all media types such as tape and optical). Data repositories may be geographically dispersed (as in the case of regional or head offices) as well as under the control of users on their local workstations.
- Understand how to retrieve the data quickly.

  The longer it takes to locate and retrieve data, the more costly the discovery task will be.
- Understand in what format the data may exist.

  Retrieved data is not of much use if it is in a format that you can no longer support. This could include backups that may have been taken using older backup technology.
- After the data has been retrieved, understand how it can be searched.

For more information about discovery, see Chapter 11, “Discovery” on page 409.

Planning for document deletion

How will records be disposed of when the time arrives? What is the legal or business requirement for it? What and how do you delete the records?

Refer to 10.3, “Records disposition” on page 405 for more information.

5.2.4 Key departments involved in planning process

Implementing an integrated system such as CommonStore, Records Manager, and Content Manager requires a large amount of planning. Communication must be established between a range of departments and disciplines.

The following departments should work together on the project:

- IT/IS department
- Legal department
- Records management department
- Help desk and support department

IT/IS department

The Information Technology/Information System (IT/IS) department is in charge of installing and managing the system. It is responsible for consolidating a number of servers and for refining processes across storage, backup, and e-mail systems.

Legal department

The Legal department is responsible for guidance on legal requirements with document retention and discovery mitigation.

Unless carefully planned and implemented, the system may not return the results as anticipated. You need to understand the key legal requirements that lead to the implementation of this system and prioritize them. For example, if the highest priority is to mitigate any document discovery litigation, ensure that you have the systems in place to allow fast searching of all of your data.

Usually, any litigation specifies the class of data needed and may even specify keywords. Even if your organization has an effective backup strategy for all of your data, this does not provide an efficient solution if data has to be fully restored first before searching can commence. Furthermore, you may not possess the software tools to rapidly search across disparate data systems.

Working closely with the legal department is important to ensure the understanding of the legal requirement and anticipate future system needs.

Records management department

The records management department is responsible for establishing and managing an appropriate file plan and life cycle, and implementing the organization's records-keeping policies.

Your organization may be governed by key legislative requirements and in most cases more than just one piece of legislation. You must understand the range of of internal and external legislative and procedural rules that govern the organization so that you can distill them into a representative set of file plan components. Many of these rules specify what documents should be kept, how long they should be kept and what should happen to the document at the end of its life. Some of these rules may contradict each other, in which case rulings may be required. It is up to the Records Management staff to translate and apply these rules effectively.

In addition, not all documents and records produced by businesses fall under any particular legislation, so it is up to your organization to adequately manage this data. Having policies and procedures in place will at least begin to mitigate any broad requirements in cases where document control is being tested.

Help desk and support department

The help desk and support department is responsible for supporting the introduction of any new technology into an organization.

This group must also be involved early in the planning process to assist in the formulation of support policies regarding the transition to a new records and archive system.

If you initially implement the system in a limited form, it may serve many purposes. During this limited implementation or pilot system, the support team can gear up to provide at least limited support during the pilot phase. The first-line support provision may only extend to the forwarding of support calls to the implementation team for resolution. At minimum, the support team should be able to identify that the issue is related to Content Manager, CommonStore, or Records Manager.

Make sure that one of the first groups to get any technical training is the help desk and support department.

5.3 System configuration

The e-mail archiving and records management solution uses CommonStore for Lotus Domino or CommonStore for Exchange Server, Content Manager, Records Manager, and Records Enabler.

Figure 5-1 illustrates the basic system architecture for the end-to-end e-mail archiving and records management solution.

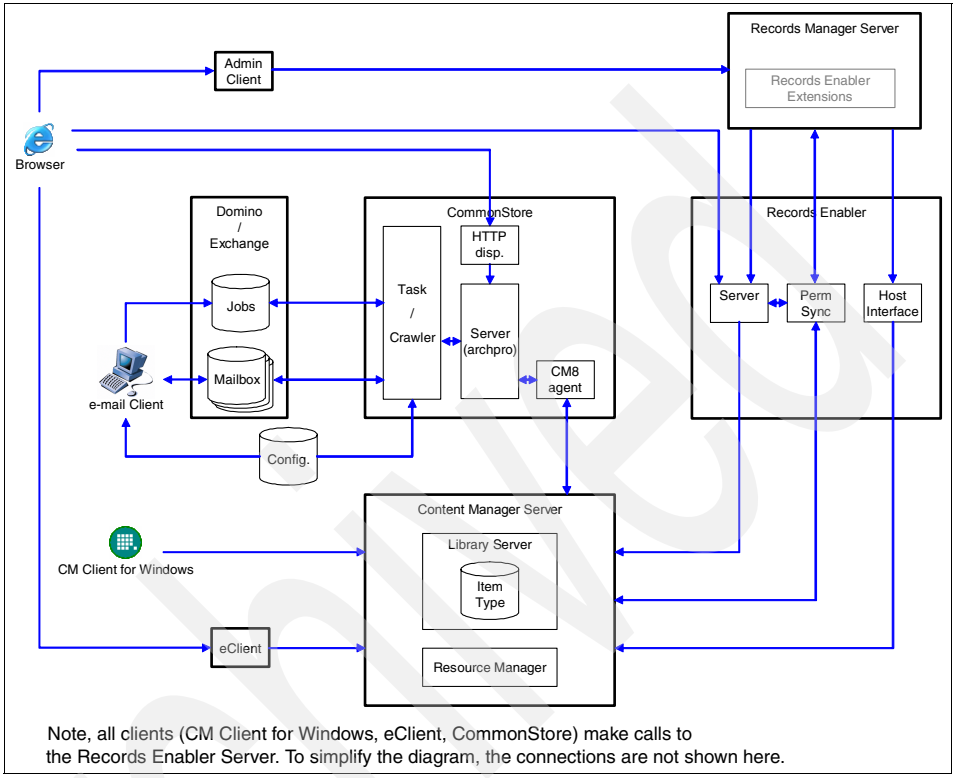

Figure 5-1 Basic system architecture

For your system configuration, decide how many servers you need, what to install on each server, where to install them if your organization has multiple sites, and how they will communicate with each other.

Review the existing e-mail infrastructure (Domino server or Microsoft Exchange Server). How many of these servers do you currently have, where are they allocated, and how are they being used?

In addition, you can start sizing your system with the following assessment:

- The total number of concurrent users during normal operations.
- The average amount of e-mail the system has to handle on daily basis.
- The average file size of individual e-mail databases and total mail databases.
- The total possible number of e-mail messages the system has to handle.
- The total system storage size for the archived e-mail.
- The system performance capacity.

- The number of records that will be declared and classified on a daily basis.
- The total number of records to be managed.
- Frequency of records auditing and reporting tasks.

If you performed the mail system audit as discussed in “Performing audits on your mail system” on page 117, you already have some of the answers, The assessment helps you decide how many servers you need, how fast the servers should be, and how much memory you may need.

Depending on the business and legal requirement, if you need to mitigate any document discovery litigation, make sure your configuration and resource allocation provide a fast search of all your data.

You should also consider the purpose and usage of all of the components involved in the solution. Will you be using Content Manager for purposes other than e-mail archiving? Will you be using Records Manager for purposes other than e-mail records control? What about your CommonStore Server? In a Lotus Domino environment, will you be using it to manage Notes databases other than mail databases? This should help you to determine whether you want to install them on their own servers or group components together on one machine.

5.3.1 Configuration options

We introduce configuration options based on three system sizes: small, medium, and large. Although it is hard to define exactly what constitutes a system size, we use the size in a general sense to give you an idea of how you should set up your servers for your solution.

Other factors may influence your decision. For example, if you have a very powerful server, you may not need to separate components onto different servers.

All in one server

In this configuration, you install everything on a single machine as shown in Figure 5-2. We recommend this configuration for prototype only. You may also use this configuration for testing and development.

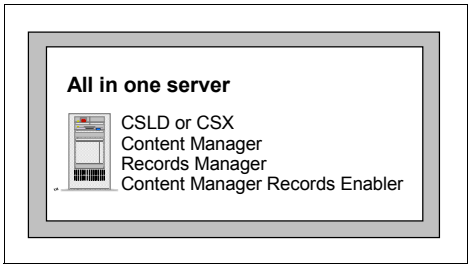

Figure 5-2 All in one server

Small-size system

Because the Records Manager engine usually reaches its hardware limit first, if the resource is available, we recommend moving the Records Manager engine to another server. See Figure 5-3.

Content Manager Records Enabler consists of three server programs. They can also be moved to the second server where the Records Manager engine resides.

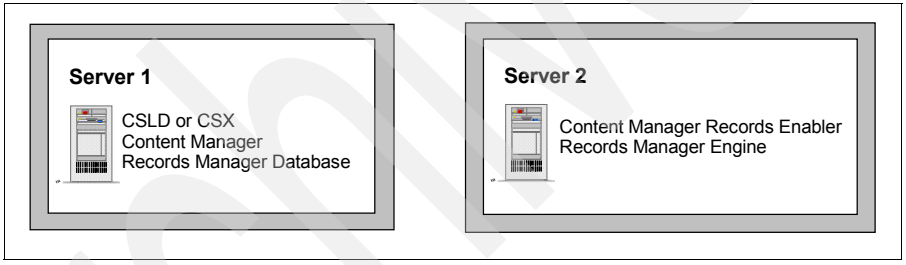

Figure 5-3 Small system configuration suggestion

Medium-size system

For a medium-size system, we recommend installing Content Manager on its own server and the Records Manager engine on its own server. You can further isolate CommonStore for Lotus Domino or CommonStore for Exchange Server on its own server. See Figure 5-4 on page 124.

CommonStore for Exchange Server also supports (and encourages) separating the archpro and the CSX Task, each to their own server hardware. CommonStore for Lotus Domino does not support this additional separation.

If you use Content Manager for purposes other than e-mail archiving and records management, you may want to install it on a server of its own regardless of the size of your current mail archiving system.

Depending on how large your e-mail volumes are and how much archive work the system has to perform, you can separate the CommonStore Server (CSX or CSLD) on its own box. In addition, depending on how many mail servers you have and how they are located, you may have multiple CommonStore Servers, one working with one or a group of mail servers.

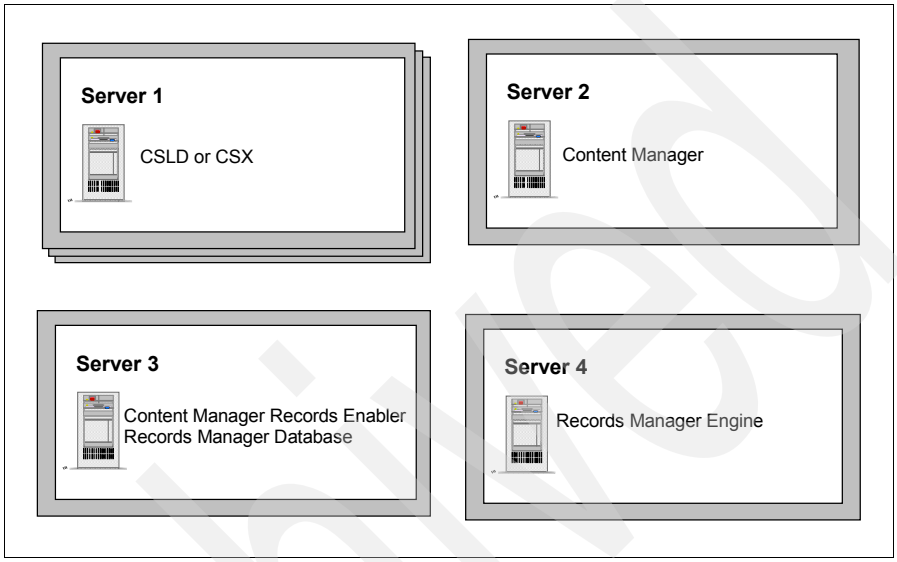

Figure 5-4 Medium-size system configuration suggestion

Large-size system

For a large system implementation, you may need to separate Content Manager's Library Server and Resource Manager onto two different servers. For example, if your organization is worldwide with mail servers located in geographically different locations, you may want to have a Resource Manager that is close to the mail servers for faster e-mail archiving and retrieving time.

If additional hardware resources are available, you may also opt to install Content Manager Records Enabler onto a separate server. See Figure 5-5 on page 125.

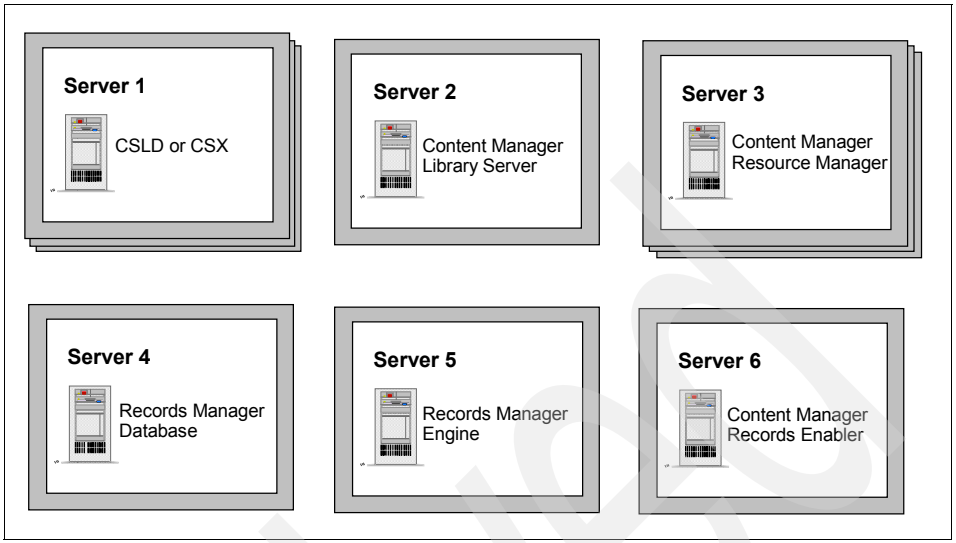

Figure 5-5 Large-size system configuration suggestion

In summary, there are many ways to configure your system for an integrated solution. You can set up everything in one server, or separate components on multiple servers, starting with the Records Manager engine. When needed, you can have multiple CommonStore Servers, separate out Content Manager Resource Manager, have multiple Content Manager Resource Managers, and separate the Content Manager Records Enabler to its own server.

System configuration directly affects system performance. You must balance your budget, your system needs, and your business requirement in finally deciding which system configuration to use.

5.3.2 Configuration consideration

As you design your system configuration, keep the following points in mind:

- If hardware resources are available, always separate the installation of the Records Manager engine and Records Manager database onto two different servers.
- If possible, it is best to keep Content Manager Library Server and Resource Manager together in one machine.
- The Records Manager engine usually reaches its hardware limit first when:
  - The Records Manager engine reaches 90% CPU usage.
  - The Records Manager database reaches 15%.
  - Content Manager reaches 30%.

- Content Manager Records Enabler reaches 30%.
- Use RAM Disc for Temp-Directory for CommonStore.

5.4 Solution implementation and deploy sequence

If both e-mail archiving and records management requirements are being considered but are not being implemented at the same time, an organization should understand the implications in selecting a particular implementation sequence.

There are three approaches you could take to implement the solution:

- Implement and deploy records management of e-mail first and then implement and deploy the e-mail archiving capability.
- Implement and deploy the e-mail archiving solution first and then add the records control function into the solution.
- Implement and deploy the complete, end-to-end e-mail archiving and records management solution at the same time.

5.4.1 Implement e-mail records control, then e-mail archiving

As part of implementing records control into an e-mail system, e-mail archiving may not have been considered as a prelude. If this is the case, some planning is required in the design of user mail templates. If you allow users to manually declare an e-mail as a record, the user interface must be modified. Records Manager provides the necessary tools and templates to accomplish these modifications.

When e-mail archiving is added after the records management solution has been implemented, further modifications must be made to the user e-mail database. Careful planning and management of these changes is required to ensure a smooth transition from a single solution that allows records declaration through an e-mail client, to an integrated solution where both records declaration and archiving can be performed.

Assuming that the records declaration (and classification) and e-mail archiving are performed manually by the user through their e-mail databases, we examine some advantages and disadvantages of implementing the system using this path.

Advantages

Reasons why you may want to implement the records management system prior to implementing the e-mail archiving capability include:

- The need to do records management and the imposition on businesses to legally comply is ever increasing. Not being compliant is no longer an option. The sooner you implement the records management capability to the e-mail system, the better. There is no legislative penalty for not archiving e-mail.
- Nothing is more admissible or can help an internal discovery requirement than to be able to demonstrate that an effective, internal records policy is in place.
- Legal discovery is not restricted in scope to only information within the records management program. Therefore, having an effective and auditable disposition process can reduce risk and exposure to possible legal actions.

Disadvantages

Some of the reasons why you may not want to implement the system using this path are as follows:

- The need to be records compliant is a relatively new imposition on businesses that are still grappling with the fundamentals.
- Organizations have tended, for a records-defensive position, to use a disposal-is-suspended-indefinitely driver for information handling and information preservation. While this has served as a legally defensive position historically, it is an unknown and is a risk for future discovery.

5.4.2 Implement e-mail archiving, then records management

It is common for organizations to implement e-mail archiving before records management. Typically this is because of the relative infancy of wide-scale adoption of electronic records management. As a result, large volumes of data (e-mail and their attachments) may already reside in the archive repository.

Depending on how archiving is implemented, these already archived documents may not have stubs remaining in users' e-mail databases. If this is the case, the options available for *locating* these documents and subsequently declaring and classifying them as records will be limited and affected by:

- What software clients are available to be used for searching.
- What security is applied to the archived e-mail.

  This may limit who can access the documents.
- Who has access to the records administrator client (RAC).

If there are many archived e-mail messages stored in the repository, they may have to be declared as records as well. This declaration and classification load may be prohibitive, particularly if there are no stubs left in the users' e-mail databases. The options for declaring and classifying records are:

- Have each user manually declare each of the documents that have previously been archived (assuming that there are stubs left in the e-mail file).
- Have the records administrator staff search for each archived document and declare and classify the document. The records administrator client can be used for this purpose but suitable search techniques will have to be employed.
- Write an application to automatically declare them as records based on the metadata of the existing e-mail. This is probably the best way to go, as it will require minimum effort from users or records administrators and can ensure that all of the e-mail is declared as records.

Advantages

Some of the reasons why you may want to go with this sequence are as follows:

- This path can be deployed quickly and deployed in stages if required.
- From a transformational management perspective, archiving is more easily understood by users.
- This adds immediate time, storage, and cost savings to the organization.
- Some of the server software and infrastructure required for both archiving and records management can be implemented and tested in advance of a records management system (assuming common use of the backend repository by both systems).

Disadvantages

Some of the reasons why you may not want to go with this path are:

- It is only 50% of the solution where appropriate records control is required.
- Already archived e-mail and documents may have to be revisited for records-declaration and classification purposes.
- You may later alter core mandatory corporate metadata needs.
- You may add to ongoing change-management procedures.

5.4.3 Implement the end-to-end solution at the same time

Although planning for a joint e-mail archiving and records management system can be complex, the results may be worth the effort.

Advantages

Some of the reasons why you may want to implement the complete integrated system at the same time are as follows:

- Jointly piloting and deploying an integrated e-mail archiving and records solution can reduce the impact, change, and pain for processes and users.
- Prudently reuses much of the same information, which may save effort and cost overall.
- An organization gets immediate benefit from archived e-mail that is also records-compliant.
- Sets an integrated functional infrastructure that is less likely to need change and is extensible across the rest of the organization.
- Utilizes a single repository for both archived e-mail and records.

Disadvantages

Disadvantages include:

- This path can add to overall project risk as the combined effort can appear to have a longer startup duration.
- Planning for an integrated solution can take longer and involves input from multiple disciplines.

Which implementation sequence you use to implement the solution depends on the intermediate needs of an organization. We recommend reviewing the business requirements and regulatory obligations, understanding all aspects of the e-mail archiving and records management solution planning and design, and then deciding which way to approach the implementation.

Part 2

Installation and configuration

This part introduces the sequence of major steps involved in installing and configuring the integrated e-mail archiving and records management solution. We focus on the overall end-to-end solution installation and configuration, and break the coverage into three chapters:

- Chapter 6, “Installation and configuration in a Lotus Domino and Windows environment” on page 133
- Chapter 7, “Installation and configuration in a Microsoft Exchange environment” on page 205
- Chapter 8, “Installation and configuration in a Lotus Domino and AIX environment” on page 289

6

Installation and configuration in a Lotus Domino and Windows environment

This chapter describes the installation and configuration of an e-mail archiving and records management solution using CommonStore for Lotus Domino, Records Manager, Content Manager, and Content Manager Records Enabler. Using a sample environment, we describe the major steps involved in installing and configuring the various components in a Windows environment. For more detailed information, see the appropriate product documentation.

We cover the following topics in this chapter:

- Overview
- Introduction to the sample environment
- Prerequisites and prerequisite software installation
- Content Manager installation and configuration
- CommonStore (CSLD) installation and configuration
- Records Manager installation and configuration
- CRME installation and configuration
- Configuring the CommonStore Server and Notes

6.1 Overview

In this section, we provide an overview for the e-mail archiving and records management integrated solution installation and configuration.

We cover:

- Software used for the integrated solution
- Installation and configuration steps and recommendation

6.1.1 Software used for the integrated solution

Several products are used in this end-to-end integrated solution. We list the software and purpose in the solution in Table 6-1.

For clarity, fix pack details have been omitted. These are referenced later in solution installation and configuration.

Table 6-1 Software used in the integrated solution and its purpose in the solution

Product	Purpose
IBM DB2 Content Manager (CM)	Repository used to store the documents and metadata for both the archive and records management systems
IBM CommonStore for Lotus Domino (CSLD)	E-mail archive system for Lotus Notes
IBM DB2 Records Manager (IRM)	Engine and administration for Records Manager
Records Enabler for Content Manager (CMRE)	Records enables the Content Manager repository. Also provides records management functions for Lotus Notes or Outlook users.
IBM DB2 ESE UDB (DB2)	Enterprise-class database used to hold both system configuration and objects metadata
IBM DB2 Net Search Extender	Extension to DB2 that adds full text search capabilities for both object metadata and documents including attachments.
IBM WebSphere Application Server with Embedded Messaging	Web application server that hosts the Content Manager Resource Manager, Content Manager Records Enabler servers, the Records Manager applications, and Content Manager eClient.
Lotus Domino server	E-mail system

6.1.2 Installation and configuration steps and recommendation

Four main products are involved in the end-to-end solution:

- IBM DB2 Content Manager
- IBM DB2 CommonStore for Lotus Domino
- IBM DB2 Records Manager
- IBM DB2 Records Enabler for Content Manager

Plan your system configuration first. (See 5.3, “System configuration” on page 120.) Decide where you want to install each main product and review the prerequisites for each product (especially if you decide to separate some components onto different machines). Before you start, make sure you know exactly what should be installed on each machine and the sequence of the installation.

Some of the product installation and configuration can be done at the same time, and others are better done sequentially. We recommend the following sequence of steps for installation and configuration:

1. Install a working Content Manager system on a designated machine.

   This includes the installation of the prerequisites first (DB2, Net Search Extender, WebSphere Application Server), then the Content Manager product.

   Validate that the Content Manager system is working by importing some documents, retrieving them, and viewing them.

   If you decide to separate the installation of the Content Manager’s Library Server and the Resource Manager onto different servers, you may have to install different prerequisites onto the machines. Refer to the product documentation to install the proper prerequisites for each server. You also need to validate the system after the installation and configuration as mentioned above.

2. Install a working CommonStore system on a designated machine.

   This includes the installation of the prerequisites (Content Manager V8 Connector from the Information Integration for Content installation, DB2 Runtime Client or DB2 Administration Client, and Notes client), and then the CommonStore for Lotus Domino product.

   Validate that CommonStore is working properly with your mail database and Content Manager by setting up some policies, archiving some e-mail, retrieving the archived e-mail, and viewing it. Also look into the Content Manager repository to ensure that the archived e-mail is there as expected.

3. Install a working Records Manager on a designated machine.

   This step can be done in conjunction with step 2. If multiple people are doing the installation, you can work together in parallel. Otherwise, we recommend performing this task after the CommonStore installation.

   The step includes the installation of the prerequisites, then the Records Manager product.

   We recommend installing the Records Manager engine and its database on separate machines. The prerequisites for the engine include WebSphere Application Server and DB2 Runtime or DB2 Administration Client. The prerequisites for the Records Manager database is DB2 server. Per system configuration option discussed in 5.3, “System configuration” on page 120, you can optionally put the Records Manager database where Content Manager is installed.

   Validate that Records Manager is working properly by using the Records Manager Administration client to create a default file plan, add a record to the system, and view the added record.

4. Install and configure Content Manager Records Enabler on a designated machine. Configure Content Manager and Records Manager. Install and configure Records Manager Extension.

   Records Manager Extension must be deployed on the same WebSphere Application Server as Records Manager.

   Validate the system: Import an item of the record enabled item type, declare it as a record, and make sure that the record is marked as declared in Content Manager and the record’s metadata is in Records Manager. After this, archive an e-mail and declare the e-mail as a record. Check both Content Manager and Records Manager to ensure that proper information is stored there. (Content Manager should have the e-mail information, based on your archiving method, and Records Manager should have the record’s related metadata.)

Tip: For the machine that will run *Permission Synchronization Server* of the Records Enabler for Content Manager product, make sure that *WebSphere Application Server with Embedded Messaging* is installed on it. This is especially important if you start your installation from an existing system such as a working Content Manager system.

If the Embedded Messaging feature was not installed, you must deinstall Content Manager, uninstall WebSphere Application Server, then reinstall WebSphere Application Server, including the Embedded Messaging feature (installed by default in V5.1.1 and later), and then Content Manager again. Do not take shortcuts or you may experience strange results.

If the existing Content Manager is currently being used and cannot be deinstalled, we recommend installing the Permission Synchronization Server on another machine.

6.2 Introduction to the sample environment

We use a sample environment to show you how to install and configure an e-mail archiving and records management solution.

Figure 6-1 shows the sample environment before any e-mail archiving and records management software are installed.

There is one client machine, Watson, and three servers, Brighton, Charger, and London, in this environment. A Domino server is running on Brighton. A Notes client is running on Watson. This should be your starting point.

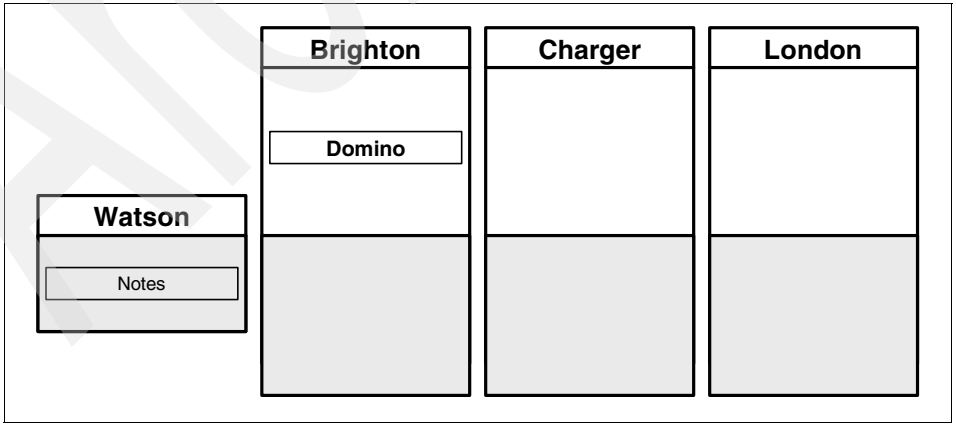

Figure 6-1 Sample environment before any software installation

Figure 6-2 shows the sample environment after all components are installed including the necessary prerequisites on each server.

On Charger, we install prerequisite software including a Notes client (Notes in Figure 6-2), a WebSphere Application Server (WAS), DB2 server with Net Search Extender (DB2 + NSE). In addition, we install CommonStore for Lotus Domino (CSLD), Content Manager (CM), and IBM Records Manager database (IRM DB).

On London, we install prerequisite software including the Content Manager V8 connector (CM connector) from Information Integrator for Content, WebSphere Application Server with Embedded Messaging (WAS*), and a DB2 client. In addition, we install Content Manager Records Enabler (CMRE) and IBM Records Manager Engine (IRM Engine).

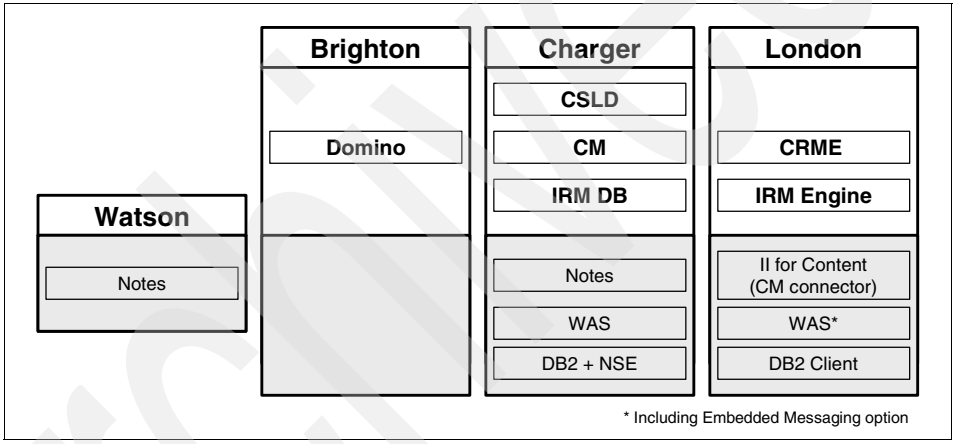

Figure 6-2 Sample environment after all software installation

6.3 Prerequisites

Four core components are involved in the e-mail archiving and records management solution: Content Manager for Multiplatforms, CommonStore for Lotus Domino, IBM Records Manager, and Records Enabler for Content Manager.

Each core component has different prerequisites. We list what they are and their version requirements. In addition, we explain why you should install the prerequisite. Understanding this should help you when you build your system that fulfills your business needs.

Prerequisite for Content Manager V8.3

Table 6-2 describes the prerequisites for Content Manager V8.3.

Table 6-2 Prerequisites for Content Manager V8.3

Product	Version	Reason why we need it
WebSphere Application Server	5.1.1.2	Resource Manager is a J2EE™ application and thus needs WebSphere Application Server.
DB2	8.2	Library Server uses stored procedures and needs a relational database (icmnlsdb). Resource Manager stores metadata in a relational database (rmdb).
DB2 Net Search Extender	8.1	For full text search within Content Manager.

Prerequisite for CommonStore for Lotus Domino V8.3

Table 6-3 describes the prerequisites for CommonStore for Lotus Domino V8.3.

Table 6-3 Prerequisites for CommonStore for Lotus Domino V8.3

Product	Version	Reason why we need it
Information Integrator for Content (II for Content) - CM V8 connector	8.3.0	The CommonStore agent needs the APIs (connector) to communicate with the Content Manager Library Server.
DB2 Runtime Client	8.2	If CommonStore for Lotus Domino is not on the same machine as the Content Manager Library Server, the Content Manager V8 connector (which is needed by the CommonStore agent) needs the Library Server database to be cataloged on the machine where CommonStore is installed. DB2 Runtime Client is thus needed. For ease of installation, DB2 Administration client is recommended.
Notes client	6.5	The CommonStore Task communicates with the Domino server; therefore it needs the Notes API, which will come when you install Notes client.

Prerequisite for IBM Records Manager V4.1.2

Table 6-4 describes the prerequisites for the Records Manager engine for IBM Records Manager V4.1.2.

Table 6-4 Prerequisites for IBM Records Manager Engine V4.1.2

Product	Version	Reason why we need it
WebSphere Application Server	5.1.1.2	Resource Manager is a J2EE application, thus needs WebSphere Application Server.
DB2 Runtime Client	8.2	The Records Manager Engine needs to communicate with the Records Manager database. If the Records Manager database is installed on another machine, then the machine with the Records Manager Engine installed needs the database to be cataloged on it and thus it needs DB2 Runtime Client. For ease of installation, DB2 Administration client is recommended.

Table 6-5 describes the prerequisite for the Records Manager database for IBM Records Manager V4.1.2.

Table 6-5 Prerequisites for IBM Records Manager database V4.1.2

Product	Version	Reason why we
DB2	8.2	The Records Manager database is a relational database. The machine that installs the Records Manager database needs to install DB2 server.

Prerequisite for Records Enabler V8.3

Table 6-6 describes the prerequisites for the Records Enabler server (CMRE server), Permission Synchronization server (PermSync server), and Host Interface server.

Table 6-6 Prerequisites for CMRE server, PermSync server, and Host Interface

Product	Version	Reason why we need it
WebSphere Application Server with Embedded Messaging	5.1.1.2	CMRE server, PermSync server ,and the Host Interface are WebSphere Application Server applications that require WebSphere Application Server. PermSync server also needs the Embedded Messaging feature.

Product	Version	Reason why we need it
Information Integrator for Content (II for Content) - CM V8 connector	8.3.0	All three servers communicate with Content Manager Library Server and therefore need Content Manager APIs.
DB2 Runtime Client	8.2	The CM V8 connector must communicate with the Content Manager Library Server database. If the Content Manager Library Server is installed on another machine, then the machine with these servers installed needs the database to be cataloged on it and thus requires DB2 Runtime Client. For ease of installation, DB2 Administration client is recommended.

Table 6-7 describes the prerequisite requirement of the Records Manager extension.

Table 6-7 Prerequisite for Records Manager Extension V8.3.0

Product	Version	Reason why we need it
WebSphere Application Server	5.1.1.2	The CMRE server, PermSync server, and the Host Interface are WebSphere Application Server applications. They require WebSphere Application Server V5.1 with Fix Pack 1 and cumulative Fix 2.

6.4 Prerequisite software installation

In this section, we describe the main steps of installing the basic software that is necessary prior to the installation of the main components of the e-mail archiving and records management solution.

The prerequisite installations include:

- “DB2 server installation” on page 143
- “DB2 Administration client installation” on page 144
- “WebSphere Application Server installation” on page 145
- “Information Integrator for Content (CM connector) installation” on page 147

Table 6-8 summarizes the prerequisites for each software product from 6.3, “Prerequisites” on page 138. Note the following:

- DB2 client (runtime or administration) is needed for Records Manager (IRM) engine, and DB2 server is required for Records Manager database if they are installed on separate machines; otherwise, they need only DB2 server.
- WebSphere Application Server (WAS in Table 6-8) with Embedded Messaging (WAS*) is required for Permission Synchronization server from CMRE.
- DB2 Net Search Extender (NSE) is required if you need full text search capabilities.

Table 6-8 Prerequisites for each software product

CM	CSLD	IRM	CMRE
WAS		WAS	WAS* (PermSync server)
DB2 +NSE	DB2 client	DB2 client (engine) DB2 (database)	
	II for Content - CM V8 connector		II for Content - CM V8 connector
	Notes client		

We install the following prerequisites on two servers as we described in 6.2, “Introduction to the sample environment” on page 137:

- London (this is where CMRE and IRM engine will be installed):
  - DB2 Runtime Client V8.2
  - Information Integrator for Content - Content Manager V8 connector
  - WebSphere Application Server (including Embedded Messaging) V5.1.1.2
- Charger (this is where CSLD, CM, and IRM database will be installed):
  - DB2 server V8.2
  - DB2 Net Search Extender V8.2
  - WebSphere Application Server V5.1.1.2 (Embedded Messaging not necessary)

Note: We do not include detailed steps of the installation in this section. We recommend using the existing product manuals in conjunction with the materials we present here for successful installations and configurations.

6.4.1 DB2 server installation

Content Manager and Records Manager use relational databases to store content (objects and records) metadata and system configuration information.

Install DB2 server to the machine (or machines) that will store Content Manager databases and Records Manager databases:

1. Install DB2 server software.
2. Verify DB2 server installation.
3. Install DB2 Net Search Extender.
4. Install DB2 Fix Pack 8.

In our sample environment, both the Content Manager database and Records Manager database are located on one server, Charger. We install the DB2 server on this machine.

Installing DB2 server

Table 6-9 shows the input values that we used during the DB2 server installation. Replace our sample input values according to your environment setup.

Table 6-9 DB2 server installation input summary for the sample environment

Required input field	Sample input value	Description
DB2 Version	8.1.7	Content Manager V8.3 requires at least DB2 V8.1.7.
Installation type	Typical	
Drive	C:\	
Installation directory	C:\IBM\DB2\SQLLIB	
DB2 Administration server user ID	db2admin	This Windows user ID is used by the DB2 Administration server to log on to the system as a service.
DB2 administration group	DB2ADMNS	This group is created automatically. The name cannot be changed. Every user ID that should be a DB2 administrator has to be in this group.
DB2 instance	DB2	

Verifying installation

To verify the DB2 server installation, create the sample database from the First Step menu. Make sure that the sample database is created successfully and that you can view the data.

Installing Net Search Extender

This is an optional step. If you need to use the full text search feature of Content Manager, install DB2 Net Search Extender before installing Content Manager. When Content Manager is installed, you can activate the full text search feature.

Use the database administrative user ID (db2admin) to run the Net Search Extender as a service.

In our sample environment, we choose not to configure for full text search.

Installing DB2 Fix Pack 8

After installing the DB2 server and the Net Search Extender, install the DB2 Fix Pack 8. The Records Manager database requires DB2 V8.2 (which is equivalent to V8.1.8) as a prerequisite.

See appropriate documentation for detailed steps of installing the fix pack.

Key information to remember

At this point, as part of your overall configuration setup documentation, write down the key configuration information that will be used during the rest of the software installation and configuration and future system operation.

Table 6-10 contains the key configuration input values to remember after the DB2 server installation.

Table 6-10 Key information to remember after the DB2 server installation

Configuration data	Sample input value
DB2 Administration server user ID	db2admin
DB2 administration group	DB2ADMNS
DB2 database server host name	charger.redbook.bocaraton.ibm.com

6.4.2 DB2 Administration client installation

We recommend installing the Records Manager database and Records Manager engine on two separate servers.

The Records Manager engine needs access to the Records Manager database. If they are not installed on the same machine, the Records Manager database has to be cataloged on the Records Manager engine machine. To catalog a database, a DB2 Runtime Client is required.

In addition to installing DB2 client to the machine where the Records Manager engine will run (if it is installed on a separate machine other than the Records Manager database), the DB2 client should also be installed on the machine where you will run CommonStore.

Although the DB2 Runtime Client is the minimum requirement for both instances, we recommend installing the DB2 Administration client because it provides a simple graphical user interface to catalog a database and other powerful DB2 tools to administer a remote database.

The installation process is straightforward, and we do not cover it here.

In the sample environment, we install the Records Manager database on Charger and the Records Manager engine on London, so we must install DB2 client on London.

6.4.3 WebSphere Application Server installation

The three core components of the solution (Content Manager, Records Manager, and Records Enabler), rely on WebSphere Application Server, so it should be installed on the machines where these products are installed.

Important: Before you continue, make sure that WebSphere Application Server with Embedded Messaging is installed on the server that runs the Records Enabler; otherwise, you may encounter problems in future.

If, after you installed everything, you discover that Embedded Messaging is not installed, you must deinstall everything including Content Manager and WebSphere Application Server, then reinstall all of the software from this point on. Otherwise, you may experience strange behavior. Do not take a shortcut here.

By default, Embedded Messaging should be installed with WebSphere Application Server installation.

The steps involved in WebSphere Application Server installation are:

1. Install WebSphere Application Server software.
2. Verify installation.
3. Install Fix Pack 1 and any accumulative fix packs.

In the sample environment, it is necessary to install WebSphere Application Server on both Charger and London. We also have to install Embedded Messaging on London because Records Enabler will be installed on the machine.

Installing WebSphere Application Server software

To help your WebSphere Application Server installation process, Table 6-11 shows the input values that we used during our installation on both servers. Replace our sample input values according to your environment setup.

Table 6-11 Installation input summary for the sample environment

Required input field	Sample input value	Description
WebSphere Application Server version	5.1.0	
Installation type	Full	This includes WebSphere Application Server with Embedded Messaging. Again, this must be installed on the machine that will run the Records Manager engine. **Note:** When choosing full installation, the installation will include all sample applications that may increase the install time.
Installation directory for WebSphere Application Server	C:\IBM\WS\WAS	
Installation directory for IBM HTTP Server	C:\IBM\IHS\	
Installation directory for Embedded Messaging server and client	C:\IBM\WS\WSMQ	
Node name	► for Charger: charger ► for London: london	We install WebSphere Application Server on both Charger and London servers.
Host name	► for Charger: charger.redbook.bocaraton.ibm.com ► for London: london.redbook.bocaraton.ibm.com	
WebSphere Administrator user ID	wsadmin	This is the Windows user ID used to run the WebSphere services.

Verifying installation

To verify successful WebSphere Application Server installation, at the WebSphere Application Server - First Steps window, click **Verify Installation**. The message `Installation Verification is complete` should show up.

Installing Fix Pack 1 and cumulative fixes

In order to fulfill the prerequisite, install Fix Pack 1 and the cumulative Fix 3.

Important: Before the installation, stop all WebSphere services (server1) as well as the IBM HTTP Server.

Key information to remember

At this point, as part of your overall configuration setup documentation, write down the key configuration information that will be used during the rest of the software installation and configuration and future system operation.

Table 6-12 contains the key configuration input values to remember after the WebSphere Application Server installation.

Table 6-12 Key information to remember after WebSphere Application Server installation

Configuration data	Sample input value
WebSphere cell	for Charger: charger for London: london
WebSphere node	for Charger: charger for London: london

6.4.4 Information Integrator for Content (CM connector) installation

Content Manager V8 connector (which comes with the Information Integrator for Content) is a prerequisite for Content Manager Records Enabler (CMRE) and CommonStore. Install the Content Manager V8 connector on both the CommonStore machine and the Records Enabler machine if these products are not installed on the same system as Content Manager.

Content Manager V8 connector must be installed before CMRE installation.

We recommend installing Content Manager before you install the connector to any of the machines in question because, during the Content Manager installation, many values necessary to configure the connector will be defined.

Although we describe the connector installation in this section, defer the installation until the Content Manager installation is done as described in 6.5, “Content Manager installation and configuration” on page 148.

Installing Content Manager V8 connector

To install Content Manager V8 connector, launch the Information Integrator for Content installation process, select **Connector** and then select **Content Manager V8 Connector** from the appropriate installation windows. Refer to product manual for specific installation instructions.

Table 6-13 shows the input values that we used during installation in our sample environment. Replace our sample input values according to your environment setup.

Table 6-13 CM V8 connector installation input for sample environment

Required input field	Sample input value	Description
Database server type	DB2 Universal Database	This is the database used by the Content Manager System.
Library Server database	icmnlsdb	Library Server stores metadata in a relational database. This value specifies the name of the database.
Library server schema name	icmadmin	
Authentication type	Server	
icmcont	icmconct	A connection user ID that is used for clients to connect to Library Server database if they do not have a valid database user ID.
icmconct password		
Local		

6.5 Content Manager installation and configuration

In the e-mail archiving and records management solution, Content Manager is used as a repository for archived e-mail. When e-mail becomes records, they are still stored in the Content Manager repository.

In this section, we describe the main steps involved in installing and configuring Content Manager, a major component in the integrated solution.

The main steps involved include:

1. Install Content Manager system.
2. Implement Windows service.
3. Verify installation.

Note: It is not our intention to include all detailed steps of the installation in this section. We recommend using the existing product manual in conjunction with the materials we present here for a successful installation and configuration.

In our sample environment, we install Content Manager on Charger.

Installing Content Manager system

We choose a typical Content Manager system installation here.

Attention: During a typical installation, full text search is *not* configured. If you need the full text search feature, choose a custom installation.

Important: The installation directory that we used in the sample environment, C:\IBM, is *not* the default installation directory. We used it to keep the PATH system variable short and avoid potential problems with long PATH value.

Follow the product manual for detail installation steps.

To help your installation process, Table 6-14 shows the input values we used during our installation for our sample environment. Input fields are grouped by the input window. Replace our sample input values according to your environment setup.

Table 6-14 Content Manager installation input values

Input window / field	Sample input value	Description
Installation destination input window		
Installation directory	C:\IBM\CM	Root directory of the Content Manager installation.
Installation type input window		
Installation type	typical	This includes the HTTPs configuration for the Content Manager internal communication. If you need to configure full text search, use custom installation type.

Input window / field	Sample input value	Description
System information input window		
host name	charger.redbook.bocaraton.ibm.com	The server's fully qualified network name.
Library Server database input window		
Library Server database name	icmnlsdb	Library Server stores metadata in a relational database. This value specifies the name of the database.
Library Server administration ID	icmadmin	This is the Content Manager administrator ID. If it does not exist as a Windows user ID, it will be created during the installation and will be added to the Windows' Administrator group with appropriate system rights.
Resource Manager database input window		
Resource Manager database name	rmdb	Resource Manager stores information about the saved documents (such as location on a disk) in a relational database. This value specifies the name of that database.
Resource Manager database administrator	rmadmin	This is the Resource Manager administrator ID. If it does not exist as a Windows user ID, it will be created during the installation and will be added to the Windows' Administrator group with appropriate system rights.
Resource Manager volume mounting point	C:\	This is the partition on which Resource Manager stores the archived documents. During a typical installation, the directory *staging* is used as a cache area for documents retrieved from a connected Tivoli Storage Manager system.
Resource Manager application input window		
Application server node name	Charger	Resource Manager is a J2EE application. This value specifies on which node the application is deployed and a new WebSphere Application Server server is created. During a typical installation, a new server called icmrm is installed. The application running in this server is deployed as icmrm.

Implementing Windows service

We recommend setting the Resource Manager application, a WebSphere Application Server, as a Windows service for ease of management.

To install a WebSphere Application Server as a Windows Service, use the following command:

```
wasservice -add <Windows Service Name> -serverName <WebSphere Server>
```

For our sample environment, we use the following command to implement the Windows service:

```
wasservice -add ResourceManager -serverName icmrm
```

6.5.1 Installation summary and verification

Figure 6-3 shows our sample environment after the prerequisites and Content Manager are installed.

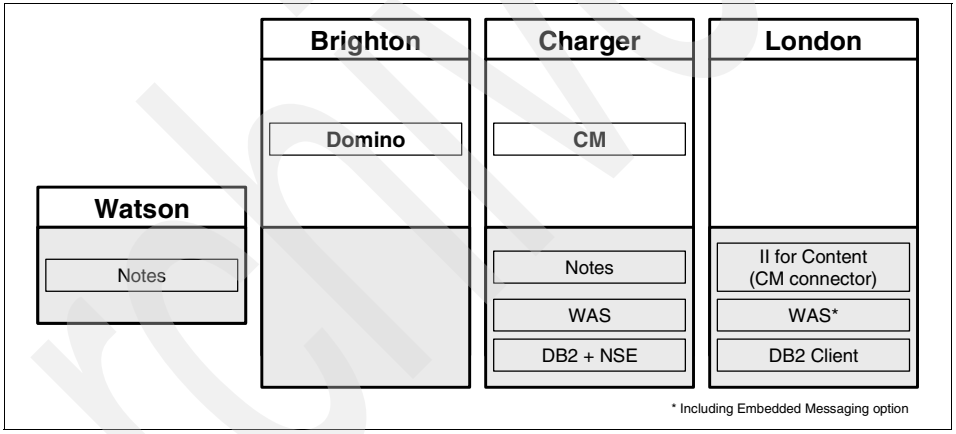

Figure 6-3 Sample environment after Content Manager installation

At the end of the installation, an installation validation utility will run. You should see the following message, which indicates a successful installation:

```
Product validation completed with no detected configuration errors.
```

In addition to this message, perform the following steps to ensure that you installed Content Manager successfully:

1. Launch the system administration client that is automatically installed on the Content Manager server. In our scenario, we launch the system administration client from Charger.

2. Log on to the Content Manager system using the administrative user ID. In our scenario, we use icmadmin.
3. Open the Resource Manager configuration.
4. If the Resource Manager configuration is available, the communication with the Resource Manager is set up properly.

Perform the following steps as the final test of a successful installation:

1. Install a Content Manager Windows client on the Content Manager server. This is the fastest way to configure the client. In our scenario, we install the client on Charger.
2. Launch the Content Manager Windows client.
3. Import a text into the NOINDEX class of type text.
4. Retrieve the document immediately afterward. Make sure that you can retrieve it, open it, and view it on the screen.

6.5.2 Key information to remember

At this point, as part of your overall configuration setup documentation, write down the key configuration information that will be used during the rest of the software installation and configuration and future system operation.

Table 6-15 contains the key configuration input values to remember after the Content Manager installation.

Table 6-15 Key information to remember after Content Manager installation

Configuration data	Sample input value	Description
Content Manager administrator user ID	icmadmin	The administrative user ID for the Content Manager system.
Connection user ID	icmconct	A connection user ID that is used for clients to connect to Library Server database if they do not have a valid database user ID.
Content Manager server host name	charger.redbook.bocaraton.ibm.com	Fully qualified host name of the server that is running the Content Manager Library Server.
Library Server database	icmnlsdb	The name of the Library Server database. Every remote client has to catalog this database before being able to access it.

6.6 CommonStore (CSLD) installation and configuration

In the e-mail archiving and records management solution, CommonStore is used to archive e-mail from mail databases to Content Manager repository. It also provides a user interface to declare e-mail as records if manual records declaration and classification is allowed.

The steps involved include:

1. Install CommonStore for Lotus Domino.
2. Configure Content Manager:
   a. Create the appropriate attributes and item type.
   b. Create the appropriate Content Manager user ID.
3. Configure ArchPro environment:
   a. Set Content Manager connector environment.
   b. Create archint.ini.
4. Start ArchPro:
   a. Submit license.
   b. Save password for Content Manager user ID.
   c. Start ArchPro.
5. Prepare Notes/Domino:
   a. Create Domino user ID.
   b. Copy template files and sign them.
   c. Create and configure configuration database and job database.
6. Configure CommonStore Task environment:
   a. Prepare notes.ini and names.nsf for CSLD Task.
   b. Set environment for CSLD Task.
   c. Save password for Domino user ID.
7. Start CSLD Task.
8. Implement Windows Services.

Features such as full text search and single instance store are not covered. Refer to the *IBM DB2 CommonStore for Lotus Domino: Administrator's and Programmer's Guide* Version 8.3, SH12-6742, to set up those features.

Note: It is not our intention to include all steps of the installation in this section. We recommend using the existing product manual in conjunction with the materials we present here for a successful installation and configuration.

In the sample environment, we install CommonStore in Charger.

Step 1: Installing CommonStore for Lotus Domino

Select a complete installation. This will install all components (ArchPro Server, CSLD Task) on the machine.

In Table 6-16, we provide the input value we used during our installation in our sample environment. Replace our sample input value according to your environment setup.

Table 6-16 CommonStore installation input for the sample environment

Required input field	Sample input value	Description
Installation directory	C:\IBM\CSLD	The installation directory in this sample environment is different than the standard installation directory. The only reason for that is to have an easier way to manage Path and Classpath variable.

Step 2: Configuring Content Manager

Log on to the Content Manager System Administrator using Content Manager Administrator ID. In the sample environment, the Administrator user ID is icmadmin (created during installation of the Content Manager, see *6.5.2, "Key information to remember" on page 152*).

Perform these steps after you are in the system administration client:

1. Create the appropriate attributes and item type (DominoMail).
2. Create the appropriate Content Manager user ID (CSLD).

Creating the appropriate attributes and item type (DominoMail)

Attributes are used to store metadata. In an e-mail environment, the Notes fields (such as Subject or Sender) are mapped through CommonStore to Content Manager attributes.

To create the new attributes:

1. Select **Data Modeling** → **Attributes**.
2. Right-click **New**. Enter the appropriate information, and click **Save**.

   The attributes CSLDOrigUser, CSLDOrigDB, CSLDDocUNID, CSLDDocSeqNum, CSCDISIS, CSCRISIS and BCC are mandatory and must be created.

   Other attributes can be created to map additional Notes fields to Content Manager attributes.

   Table 6-17 on page 155 lists the attributes used in the sample environment.

Table 6-17 Content Manager attributes used in the sample environment

Attribute	Required	Attribute type	Character type	Character length
CSLDOrigDB	YES	Var. char.	Other	1 ... 254
CSLDOrigUser	YES	Var. char	Extended alphanumeric	1 ... 254
CSLDDocUNID	YES	Character	Alphanumeric	32
CSLDDoc SeqNum	YES	Character	Extended alphanumeric	25
CSCDISIS	YES	Character	Alphanumeric	32
CSCRISIS	YES	Character	Alphanumeric	32
BCC	YES	Var. char.	Other	0 ... max. possible
SUBJECT	NO	Var. char.	Other	0 ... 254
SENDER	NO	Var. char.	Other	0 ... 100
TO	NO	Var. char.	Other	0 ... max. possible
CC	NO	Var. char.	Other	0 .... max. possible
POSTEDDATE	NO	Time stamp	N/A	N/A

While CSCDISIS identifies the message archived (Document Identifier) along with the attributes common for all instances, CSCRISIS (Record Identifier) identifies the specific instance of that message. This instance holds the attributes CSORIGINATOR and BCC.

To create the new item type, DominoMail:

1. Select **Data Modeling** → **Item types**.
2. Right-click **New**.
3. Enter `DominoMail` as the name of the item type.

4. Click the **Attributes** tab, select these attributes, and assign them to the item type:

   ```
   CSCDISIS
   SUBJECT
   SENDER
   TO
   CC
   POSTEDDATE
   ```

5. For single instance store, create a child component, `CSLDMailChild`, and add the following attributes to this child component:

   ```
   CSCRISIS
   CSLDOrigDB
   CSLDOrigUser
   BCC
   CSLDDocUNID
   CSLDocSeqNum
   ```

   Table 6-18 shows the input values we used to create the DominoMail item type in the sample environment.

Table 6-18 Item type, DominoMail, created in the sample environment

Configuration tab / field	Sample input value	Description
Definition		
Name	DominoMail	You can enter any name for the item type. The CommonStore configuration file (archint.ini) refers to this name.
Text search	unchecked	Text search is not configured during this installation. To enable full text search, follow the instructions in the CommonStore document *Text Search Configuration for IBM DB2 Content Manager V8*.
Access Control List		
Item type access control list	PublicReadACL	In order to make the integration with Records Manager fully functional, Public Read Access is necessary.
Check ACL at	on Item level	An access control list is applied to every document inserted to this item type.

Configuration tab / field	Sample input value	Description
Attributes		
Attributes	CSCDISIS SUBJECT SENDER TO CC POSTEDDATE	In an single instance store environment, these attributes have the same values for all e-mail. If one attribute is different, this implies that the e-mail has changed and it must be stored as a separate item.
Child component	CSLDMailChild	One Content Manager item can have multiple children. In a single instance store environment, child components are used to store those attributes that might be different for every e-mail, even though the content of the e-mail might be the same. Note, the name of a child component must be unique within a Content Manager system. This name will be used to configure the CommonStore ArchPro archint.ini file.
Attributes for child component	CSCRISIS CSLDOrigDB CSLDOrigUser BCC CSLDocUNID CSLDDocSeqNum	These attributes (or mapped Domino fields) can be different per e-mail, even though the rest of the e-mail is the same. For example, the BCC field will be different for a user who is on the BCC list than it is for a user who is not on the BCC list of the same e-mail.
Document Management		
Document part	ICMBASE	Only this document part is needed as no text search is configured. In case of text search, the ICMBASETEXT would be necessary.

Creating the Content Manager user ID (CSLD)

CommonStore uses a Content Manager ID to communicate with the Content Manager system. This user ID requires access to the e-mail stored in the DominoMail item type.

To create a new Content Manager user ID:

1. Select **Authentication** → **Users**.
2. Right-click **New**.

3. Enter `CSLD` as the user name and appropriate values. Click **Save**.

   Table 6-19 shows the input values we used to create the Content Manager user ID in the sample environment.

Table 6-19 Content Manager user ID, CSLD, created in the sample environment

Configuration tab / field	Sample input value	Description
Define Users		
Name	CSLD	You can enter any name for the user. The CommonStore configuration file (archint.ini) refers to this name.
Password		During startup of the CommonStore Server (ArchPro), this password has to be provided so that the ArchPro can use the ID to log on to the Content Manager system.
Password expiration	Never expires	To ensure that the CommonStore will start up correctly, make sure that the password never expires. If the password can expire, the CommonStore startup will fail if the Content Manager password changes.
Maximum privilege set	AllPrivs	The Content Manager used ID used by CommonStore has to be a Super User, which is necessary for the Records solution. Therefore, it is necessary to assign the AllPrivs privilege set.
Set Defaults		
Default item access control list	PublicReadACL	It is necessary to provide an ACL in this field; otherwise, the user ID cannot be created. However this ACL is only used if, during an item type creation, the field “User’s default ACL” is chosen to be the ACL that defines the ACL to be set for items stored in the item type.

Tip: To make sure that the Content Manager user ID used by CommonStore can access the Content Manager system, install a Content Manager Windows client on the CommonStore Server. Use this client to log on to the Content Manager using the ID (CSLD) created to be used by CommonStore. Try to import a document (for example, a text file) to the newly created item type.

Step 3: Configuring the ArchPro environment

To configure the ArchPro environment:

1. Set the Content Manager connector environment.
2. Create an archint.ini file.

Setting the Content Manager connector environment

To apply the correct environment settings for the Content Manager V8 connector, run a batch program called Agentenv_CM8.bat. This program is delivered with CommonStore and it can be found in the bin directory of the CommonStore installation directory.

Run Agentenv_CM8.bat from a Windows command prompt as follows:

1. Open a Windows command prompt.
2. Change to the `bin` directory.

   In our sample environment, it is C:\IBM\CSLD\bin, where C:\IBM\CSLD is the chosen installation directory (see “Step 1: Installing CommonStore for Lotus Domino” on page 154).
3. Run Agentenv_CM8.bat.

Creating the archint.ini file

The archint.ini file configures the ArchPro Server. It defines whether logging and tracing are activated. Furthermore, logical archives are defined in archint.ini. These logical archives point to a specific Content Manager server including the item type that is used.

To create a server configuration profile, archint.ini, follow these steps:

1. Open the sample profile archint_sample_cm8.ini in an editor. This file resides in the instance directory of the CommonStore installation directory:

   ```
   Instance Directory: C:\IBM\CSLD\Server\instance01
   ```
2. Save the file as `archint.ini` in the same directory.

> **Important:** Make sure the file is saved as `archint.ini` and not as archint.ini.*txt*.

3. Use the search function of your editor to locate the following section:

```
ARCHIVE              CSLDMail
   STORAGETYPE          CM
   LIBSERVER            sampleLibServer
   ITEM_TYPE            sampleItemType
   CMUSER               sampleUser
   ARCHIVETYPE          GENERIC_MULTIPART
```

Following the sample environment, configure the logical archive as listed in Table 6-20.

Table 6-20 Input values for archint.ini file in the sample environment

Parameter	Sample input value	Description
ARCHIVE	EMail	You can enter any name for the logical archive that will be used by CommonStore Task. A Task is not aware of a Library Server or an item type but only of the logical archive name. ArchPro defines logical archives in order to make them transparent to a Task which archive system (Content Manager, Content Manager OnDemand, Tivoli Storage Manager) is used. The Task only refers to the logical archive or archives.
STORAGETYPE	CM	ArchPro supports various archive systems in addition to Content Manager. The value CM used here specifies that the archive is a Content Manager system.
ITEM_TYPE	DominoMail	The item type used to store e-mail. It was created in step 2a of 6.6, “CommonStore (CSLD) installation and configuration” on page 153.
LIBSERVER	icmnlsdb	The name of the Content Manager Library Server in which the Item type is created. If the Content Manager is running on a system other than the CommonStore Server, this is the name under which the remote database is cataloged on the CommonStore system. In the sample environment, the Content Manager is installed on the same system and therefore it is not necessary to catalog the database. The name of the database is configured during the Content Manager installation. See 6.5.2, “Key information to remember” on page 152.

Parameter	Sample input value	Description
CMUSER	CSLD	Content Manager user ID that is used by CommonStore to communicate with the archive system. This ID was created in step 2b of 6.6, “CommonStore (CSLD) installation and configuration” on page 153.
ARCHIVETYPE	GENERIC_MULTIDOC	This specifies how documents are stored in the Content Manager. This is relevant if Component Archiving is selected for e-mail archiving. In this case, the e-mail gets separated from its attachment. Every component (e-mail and every attachment) is stored as a separate document in Content Manager if GENERIC_MULTIDOC is defined.
SISCHILDNAME	CSLDMail Child	To enable single-instance store functionality in your archive, add this line. This value has to be exactly the same as the one defined in “Step 2: Configuring Content Manager” on page 154.

After changing the values, the logical archive section of the archint.ini file should look like this:

```
ARCHIVE              EMail
   STORAGETYPE          CM
   LIBSERVER            icmnlsdb
   ITEM_TYPE            DominoMail
   CMUSER               CSLD
   ARCHIVETYPE          GENERIC_MULTIDOC
```

4. Save the changes.

Step 4: Starting ArchPro

To start ArchPro:

1. Submit a license.
2. Save the password for the Content Manager user ID.
3. Start ArchPro.

Submitting a license

To submit a license, perform the following steps:

1. Open a Windows command prompt.

2. Change to the instance01 subdirectory of the CSLD installation path.

   Instance Directory: `C:\IBM\CSLD\server\instance01`

3. Enroll a CommonStore license by entering the following command:

   ```
   archpro -f license
   ```

4. The location of the license file is requested. Provide the full path including the file name:

   ```
   C:\IBM\CSLD\licensekey\csld8.lic
   ```

   Important: If you skip this step, a Try and Buy licence will be installed and it will expire after 90 days. After the 90-day period, the CommonStore Server will not start.

Saving the password for the Content Manager user ID

To save the password for the Content Manager user ID:

1. Open a Windows command prompt (if not already open).
2. Change to the instance01 subdirectory of the CSLD installation path.

   Instance Directory: `C:\IBM\CSLD\server\instance01`

3. Set the password for the item type (which is created in “Creating the appropriate attributes and item type (DominoMail)” on page 154) by entering the following command:

   ```
   archpro -f serverpasswd
   ```

   The password for the Content Manager user ID used by CommonStore (which is defined in the archint.ini and created in “Creating the Content Manager user ID (CSLD)” on page 157) is requested.

4. Type the password and press **Enter**.

Important: If this command is not issued from the instance directory, it is necessary to point to the archint.ini file to be used. To do so, include the -i parameter as follows:

```
archpro -f serverpasswd -i <path to archint.ini file>
```

If the **archpro** command is issued without the -i parameter, then ArchPro searches in the starting directory for the archint.ini file. If no file can be found, the startup will fail.

Important: If the password of the Content Manager user ID used by CommonStore is changed, it is necessary to run the **archpro -f serverpasswd** command again, to provide the new password to ArchPro.

Starting ArchPro

To start ArchPro, perform the following steps:

1. Open a Windows command prompt (if not already open).
2. Change to the instance01 subdirectory of the CSLD installation path.

   Instance Directory: `C:\IBM\CSLD\server\instance01`
3. Start the CommonStore Server by entering the following command:

   `archpro`

Example 6-1 shows the messages displayed during archpro startup.

Example 6-1 Archpro startup message

```
C:\IBM\CSLD\server\instance01>archpro

          ********************************************************************
          * IBM DB2 CommonStore - Server 8.3.0.0                             *
          * (c) Copyright IBM Corporation, 1997, 2004 All Rights Reserved. *
          * Build 8.3.0.0,    Compiled at Mar  9 2005.                       *
          ********************************************************************

CSS0030I: ArchPro is using INI file 'C:\IBM\CSLD\server\instance01\archint.ini'
CSS0910I: Trying to get a LUM Production License for IBM DB2 CommonStore for
Lotus Domino
CSS0929I:
          ****************************************************************
          *                Got a Production License for                  *
          *          IBM DB2 CommonStore for Lotus Domino                *
          ****************************************************************
CSS0158I: ArchPro 3464 started on UNICODE Port 4336.
CSS0157I: ArchPro 3464 is waiting for external connections on fixed port 8013.
CSS0100I: IBM Content Manager CommonStore HTTP Task 8.3.0.0
CSS0333I: ArchPro is informed that Web dispatcher 'HTTP_TASK_1' has started and
is ready (socket 1840).
CSA0100I: IBM Content Manager CommonStore Agent for Content Manager 8.3.0.0
CSS0010I: HTTP WORKER #0: has been initialized and started
CSS0001I: HTTP_TASK_1: successfully started, ready to process jobs
CSS0330I: ArchPro is informed that CM agent 'CM-AGENT_1' has started (socket
1804).
CSS0010I: HTTP LISTENER : has been initialized and started
CSS0103I: HTTP LISTENER : listens for HTTP requests on port 8085
ADMU0116I: Tool information is being logged in file C:\IBM\CSLD\Search
Server\logs\server1\startServer.log
ADMU3100I: Reading configuration for server: server1
CSA0300I: Connection for repository 'DOMINOMAIL' on server 'charger' with item
type 'DOMINOMAIL' for user 'icmadmin' OK
CSA0001I: CM-AGENT_1: successfully started, ready to process jobs
```

```
ADMU3200I: Server launched. Waiting for initialization status.CSS0325I: ArchPro
is informed that  Agent 3480 is ready to obtain order.
CSS0166I: ArchPro is fully initialized. Queue processing is enabled now.
CSA0010I: CM-AGENT WORKER #0: has been initialized and started
CSA0010I: CM-AGENT WORKER #1: has been initialized and started
ADMU3000I: Server server1 open for e-business; process id is 3848
```

The message prefix (the first three characters of the message) identifies the component generating the message. The following mappings apply:

CSS CommonStore Server (ArchPro)

CSA CM8 agent

The CommonStore Server has completed initialization when the following message is displayed:

```
Archpro is fully initialized. Queue processing is enabled now.
```

Step 5: Preparing Notes/Domino

To prepare for Notes/Domino, you need to:

1. Create a Domino user ID (CSLD Task).
2. Copy template files and sign them.
3. Create and configure the configuration database and job database.

Creating a Domino user ID (CSLD Task)

CommonStore uses a Domino user ID to authenticate with the Domino server. This user ID needs access to all Notes databases that are archived or involved in the records declaration process. This user ID also needs access to the configuration database and to the job database. The ACL of the job database needs to grant the role “CSLD User” to this user ID. This role enables an ID to see all jobs within the job database in addition to the jobs created by the ID.

To create a Domino user ID:

1. Launch the Domino Administrator Client.
2. Log on with a user ID that has the proper rights to register new users.
3. Create the user ID. Use the sample input value in Table 6-21 on page 165 as a reference.

Table 6-21 Input values for user, CSLD Task

Configuration field	Sample input value	Description
User ID	CSLD Task	User ID that will be used by CommonStore Tasks to authenticate with the Domino server. First Name: CSLD Last Name: Task
password		

Copying template files and signing them

CommonStore ships with three Notes template files. After the installation, these templates are located in the data directory in the CommonStore installation directory. A Domino administrator has to copy those files onto the Domino server and sign them to avoid unnecessary security prompts for the e-mail user.

To make the templates available, copy them directly to the Domino server data directory.

To sign a template:

1. Open the Domino Administrator.
2. Select **Files**, and select the files to be signed.
3. Right-click **sign**.
4. Copy the files from the CommonStore installation directory on Charger to the Domino data directory on Brighton. See Table 6-22.

Table 6-22 Original and destination location of the Notes template files

Configuration data	Sample input value	Description
Original location	Charger C:\IBM\CSLD\data	CommonStore install directory
Destination	Brighton C:\IBM\Domino\data	Domino data directory

5. Go to Notes administration client.
6. Select **Files**, and select the three templates.
7. Right-click and select **Sign using the Administrator ID**.
8. Add the user ID that was used to sign the templates to the Domain ECL:
   a. In Notes administration client, select **People&Groups**.
   b. Select **Action** → **Edit Administration ECL**.

Creating and configuring configuration database and job database

A configuration database stores all information needed for a CommonStore Task, such as the task profile. The database must be created using the template that ships with the CommonStore Server. If an older version of an template is used, errors will occur.

To create a new configuration database:

1. Open a Notes client using an ID with proper rights to create a database on the server.
2. Select **File** → **Database** → **New**.
3. Select the Domino server, the template, and enter the appropriate information for the new database using the values in Table 6-23 as reference.

Important: If the Sample Mail Template is used or the CommonStore script libraries are copied to an existing Mail Template, the CreateCSNJobs script library must be updated to point to the newly created job database; therefore, the Jobdatabaseserver and Jobdatabasename entries need to be updated.

Table 6-23 Configuration database setup

Configuration field	Sample input value	Description
Server	Mail/ITSO	Name of the Domino server where the configuration database is created.
Title	CSLDConfig	Title of the configuration database.
Filename	CSLDConfig.nsf	File name of the configuration database created.
Server	Mail/ITSO	Name of the Domino server that contains the configuration database template file.
Template	CSLD Configuration Database 8.3	Name of the template to be selected.

4. Add the Domino user ID used by CommonStore (CSLD Task) to the ACL.

To create a new job database:

1. Start Notes Client.
2. Select **File** → **Database** → **New**.
3. Enter appropriate values for the database using the information listed in Table 6-24 on page 167 as a reference.

Table 6-24 Job database setup

Configuration field	Sample input value	Description
Server	Mail/ITSO	Name of the Domino server where the job database is created.
Title	CSLDJobs	Title of the job database.
Filename	CSLDJobs.nsf	File name of the job database created.
Server	Mail/ITSO	Name of the Domino server that contains the job database template file.
Template	CSLD Job Database 8.3	Name of the template to be selected.

4. Add the Domino user ID used by CommonStore (CSLD Task) to the ACL and assign the role `CSLDUsers` to it.
5. Open the newly created Configuration database and create a profile for the archiving task and a profile for the retrieve task, using values in Table 6-25 and Table 6-26 on page 168 as references.

Table 6-25 Database profile for archive task

Configuration field	Sample input value	Description
Basics		
Name	Archiver	This profile name is used during the startup of a CommonStore Task.
Working DBs		
Working DBs	All	The value All specifies that all jobs in the job database will be processed by this task.
Job DB		
Database name	CSLDJobs.nsf	Name of the job database that this task is assigned to.
Server	Mail/ITSO	Server on which the job database is located.

Configuration field	Sample input value	Description
Security		
Restrict retrieval to point of origin	Yes	CommonStore security is activated, so that only retrievals to the original database are possible. The Content Manager attribute, CSLDOrigDB, has to exist.
Environment		
Task TCP/IP port	9000	
CommonStore TCP/IP port	47111	ArchPro server listens for requests from CSLD Tasks on this port. It is defined in the archint.ini with the variable DOMINOPORT.
CommonStore host name	charger.redbook.bocaraton.ibm.com®	This information is used to create the URL when http links are created.
CommonStore Web port	8095	This information is used to create the URL when http links are created. Together with the CommonStore host name, a URL will look similar to this: http://charger.redbook.bocaraton.ibm.com:8095 This port has to point to the HTTP Dispatcher port, which is defined in the archin.ini with the parameter WEBPORT.
Folder Archive ID		Folder Archiving is not used in this environment, so the value is not set.

Table 6-26 Database profile for retrieve task

Configuration field	Sample input value	Description
Basics		
Name	Retriever	This profile name is used during the startup of a CommonStore Task.
Working DBs		
Working DBs	All	The value All specifies that all jobs in the job database will be processed by this task.

Configuration field	Sample input value	Description
Job DB		
Database name	CSLDJobs.nsf	Name of the job database that this task is assigned to.
server	Mail/ITSO	Server on which the job database is located.
Security		
Restrict retrieval to point of origin	Yes	CommonStore security is activated, so that only retrievals to the original database are possible. The Content Manager attribute, CSLDOrigDB, has to exist.
Environment		
Task TCP/IP port	9001	
CommonStore TCP/IP port	47111	ArchPro server listens for requests from CSLD Tasks using this port. Defined in the archint.ini with the variable DOMINOPORT.
CommonStore host name	charger.redbook.bocaraton. ibm.com	This information is used to create the URL when http links are created.
CommonStore Web port	8095	This information is used to create the URL when http links are created. Together with the CommonStore host name, a URL will look similar to this: http://charger.redbook.bocaraton.ibm.com:8095 This port has to point to the HTTP Dispatcher port, which is defined in the archin.ini with the parameter WEBPORT.
Folder Archive ID		Folder archiving is not used in this sample environment, so the value is not set.

6. Set document mapping and content type mapping using Table 6-27 and Table 6-28 on page 171 as reference.

Table 6-27 Document mapping

Configuration window / field	Sample input value	Description
Form		
Define mapping for	Document form	Specifies that a document mapping based on a form is created.
Notes form name	Memo	Specifies the name of the form that is mapped to a logical archive.
Optional form aliases	Reply, Forward	Specify all other forms that will be mapped to the logical archive. **Note:** Only mapped forms are archived. If only Memo is mapped, and other forms (such as reply memo or forward memo) are not specified in the optional form aliases, they will not be archived.
CommonStore Archive ID	EMail	The name of the logical archive this document mapping is using. The value is defined in the archin.ini as shown in “Step 3: Configuring the ArchPro environment” on page 159. This value is case sensitive.
Configuration		
Notes fields to display in hit lists	Subject, From, PostedDate	A hit list is created when CommonStore finishes a search within the Content Manager. The values define which values are shown in the columns of the hit list.
Form for result documents	Memo	The Notes form used when a document gets selected in a hit list for retrieve.
Attribute		
Notes document field names	Subject From PostedDate	List of Notes/Domino fields that are mapped to the Content Manager attributes.

Configuration window / field	Sample input value	Description
Archive attribute names	CSLDSubject CSLDFrom CSLDPostedDate	List of the Content Manager attributes that are mapped to Notes/Domino fields. These attributes are created in “Step 2: Configuring Content Manager” on page 154. In this sample environment, the Notes field Subject is mapped to the Content Manager attribute CSLDSubject, the Notes field From is mapped to the Content Manager attribute CSLDFrom, and PostedDate to CSLDPostedDate.

Table 6-28 Content type mapping

Configuration field	Sample input value	Description
File extension	csn	Files with that extension are created when e-mail is archived in Notes Native Format.
Content type	csn/Application	

Step 6: Configure the CommonStore Task environment

To configure the CommonStore Task environment, you need to:

1. Prepare notes.ini and names.nsf for the CSLD Task.
2. Set the environment for the CSLD Task.
3. Save the password for Domino user ID.

Preparing notes.ini and names.nsf for CSLD Task

The CommonStore Task uses the Notes API and therefore needs a notes.ini file and a local names and address book database (usually called names.nsf). An easy way to create a proper notes.ini is to log on to a Notes client with the Domino user ID created for CommonStore. If the configuration database and the job database are accessible from the user ID, the notes.ini and names.nsf can be used for the CommonStore Task.

Table 6-29 Notes.ini and names.nsf setup

Configuration data	Sample input value	Description
Location of copied notes.ini and names.nsf	C:\IBM\CSLD\	The startup command of a CommonStore Task includes this directory path to point to the notes.ini.
Directory value in copied notes.ini	C:\IBM\CSLD\	This parameter tells a Notes API where to find the necessary names.nsf.

Tip: The notes.ini and the names.nsf file in the Notes installation directory should be copied to a different directory, where it is only accessed by the CommonStore Task and not changed by any other Notes application.

Important: After both files are copied, make sure the directory entry in the copied notes.ini file is changed to the directory that contains the copied names.nsf.

Also, copy the Domino ID file used by the task to the same directory and make sure the *keyfilename* entry points to the right directory.

Setting environment for CSLD Task

The nnotes.dll must be in the system path because the CommonStore Task needs access to that file.

If you have more than one nnotes.dll files, choose a path of a Lotus Notes client installation.

Table 6-30 nnotes.dll setup

Configuration data	Sample input value	Description
Location of nnote.dll	C:\IBM\notes\	This directory is added to the system PATH variable.

Saving the password for Domino user ID

In order to start a CSLD Task without user interaction, it is necessary to store the password of the Domino user ID. The Domino user ID a task uses is specified by the notes.ini that is used by the task.

Use the following command to be prompted for the password to be stored:

```
csld -f serverpasswd -i <Path to notes.ini>
```

The password will be stored encrypted in a file. Every time the password of the used Domino user ID is changed, this process must be repeated. If the password is changed without repeating this process, a CommonStore Task will fail to connect to the Domino server because it is using an invalid password.

After the password is saved, notes.ini has to be updated to inform the Notes API to make use of this stored password. Add the following line to the notes.ini file:

```
EXTMGR_ADDINS=CSLDExtPwd.dll
```

Important: Do not add this line to a notes.ini file that is used by a Notes client, or else the Notes client will use the DLL and will not start up properly.

Table 6-31 summarizes the setup for saving the Domino user ID password.

Table 6-31 Domino user ID password saving setup

Configuration data	Sample input value	Description
location of used notes.ini file	C:\IBM\CSLD\	
additional line in this notes.ini	EXTMGR_ADDINS=CSLDExtPwd.dll	
store password	csld -f serverpasswd -i <Path to notes.ini>	

Step 7: Start CSLD Task

A CommonStore Task requires an up-and-running ArchPro. The startup command syntax is:

```
csld -s <servername> -n <configdatabasename> -p <profilename> -i <notesinifile>
```

In this syntax:

- `<servername>` is the Domino server of the configuration database.
- `<configdatabasename>` is the name of the configuration database.
- `<profilename>` is the name of the profile to be used.
- `<notesinifile>` is the notes.ini file to be used.

To start the CSLD Tasks:

1. Open two Command prompt windows.
2. Start the Archive task:

   ```
   c:\ibm\csld\bin\csld.exe" -s Mail/ITSO -n CSLDConfig.nsf -p Archiver -i
   c:\ibm\csld\notes.ini
   ```

   Substitute with values in your environment.

3. Start the Retrieve task:

   ```
   c:\ibm\csld\bin\csld.exe" -s Mail/ITSO -n CSLDConfig.nsf -p Retriever -i
   c:\ibm\csld\notes.ini
   ```

   Substitute with values in your environment.

Step 8: Implementing Windows services

CommonStore provides the ability to run all components (ArchPro, tasks) as Windows service. To implement a component as a service, it must be installed as a service using the archservice program that ships with CommonStore.

The archservice program must have an .ini file to install a program as a service. The .ini file contains the startup command for the component to be started and a location for the trace file of the service.

Attention: Before installing the ArchPro and the Task as Windows services make sure previously started components (for example, in Windows command prompts) are closed down.

To implement a Window service:

1. Create a directory C:\IBM\CSLD\WindowsService.
2. Open a text editor and create three .ini files, one per component. Each file contains the following parameters:
   - `SERVICE_ TRACEFILE`: The path (including the file name) of the trace file, that this service creates. The directory must exist already.
   - `PROCESS1`: The startup command of the component.

   Example 6-2, Example 6-3 on page 175, and Example 6-4 on page 175 show the .ini files setup for the sample environment. Use them as references to set up files in your environment.
3. Save the files in directory C:\IBM\CSLD\WindowsService.

Example 6-2 Sample ArchProService.ini file

```
#-------------------------------------------------------------------------------#
# full file name of SERVICE_TRACEFILE
#-------------------------------------------------------------------------------#
SERVICE_TRACEFILE  'C:\ibm\csld\WindowsService\ArchProService.trace'
#-------------------------------------------------------------------------------#
# start sequences for archpro and csld tasks
#-------------------------------------------------------------------------------#
PROCESS1 '"c:\ibm\csld\bin\archpro.exe" -i
"c:\ibm\csld\server\instance01\archint.ini"'
```

Example 6-3 Sample ArchiveTaskService.ini file

```
#-------------------------------------------------------------------------------#
# full file name of SERVICE_TRACEFILE
#-------------------------------------------------------------------------------#
SERVICE_TRACEFILE  'C:\ibm\csld\WindowsService\ArchiveTaskService.trace'
#-------------------------------------------------------------------------------#
# start sequences for archpro and csld tasks
--------------------------------------------------------------------------------#
PROCESS1 '"c:\ibm\csld\bin\csld.exe" -s mail/ITSO -n CSLDConfig.nsf -p Archiver
-i "c:\ibm\csld\notes.ini"'
```

Example 6-4 Sample RetrieveTaskService.ini file

```
#-------------------------------------------------------------------------------#
# full file name of SERVICE_TRACEFILE
#-------------------------------------------------------------------------------#
SERVICE_TRACEFILE  'C:\ibm\csld\WindowsService\RetrieveTaskService.trace'
#-------------------------------------------------------------------------------#
# start sequences for archpro and csld tasks
--------------------------------------------------------------------------------#
PROCESS1 '"c:\ibm\csld\bin\csld.exe" -s mail/ITSO -n CSLDConfig.nsf -p
Retriever -i "c:\ibm\csld\notes.ini"'
```

4. Open a Command prompt window.

   The command syntax used to install the Window services is:

   ```
   archservice install -n <name> -c <config file>
   ```

   In this syntax:

   <name> appears in the Windows services list as part of the service name. A CommonStore Service always starts with CommonStore_ and ends with the specified value. In the examples, the services will appear as:

   ```
   CommonStore_ArchPro
   CommonStore_ArchiveTask
   CommonStore_RetrieveTask
   ```

   <config file> is the path (including the file name) of the configuration file to be used.

5. Execute the following command to install ArchPro service:

   ```
   archservice install -n ArchPro -c
   C:\IBM\CSLD\WindowsService\ArchProService.ini
   ```

6. Execute the following command to install the Archive Task service:

   ```
   archservice install -n ArchiveTask -c
   C:\IBM\CSLD\WindowsService\ArchiveTaskService.ini
   ```

7. Execute the following command to install the Retrieve Task service:

```
archservice install -n RetrieveTask -c
C:\IBM\CSLD\WindowsService\RetrieveTaskService.ini
```

6.6.1 Installation summary and verification

To test the installation, create a test Notes user. Replace its e-mail's database template with the sample mail template that ships with CommonStore. Set the job database name and job database server value in that template.

In the sample environment, these values are:

- Job database name: CSLDJobs.nsf
- Job database server: Mail/ITSO

After the template is applied, make sure that the CSLD Task has access to that database. You can test this by using CSLD Task user ID to open the mail database.

Open the mail database using the regular mail database user ID. Select an e-mail for archiving. Select **CommonStore** → **Archive Selected Documents** and click **OK**.

The archive is successful when the document moves to the "Archived documents" category in the Inbox. To check whether the retrieve is working, open the archived document and click **Fetch**.

If both operations are successful, the installation and basic configuration of CommonStore are completed.

Reviewing the sample environment, Figure 6-4 on page 177 shows the components that are installed and configured after this section is completed.

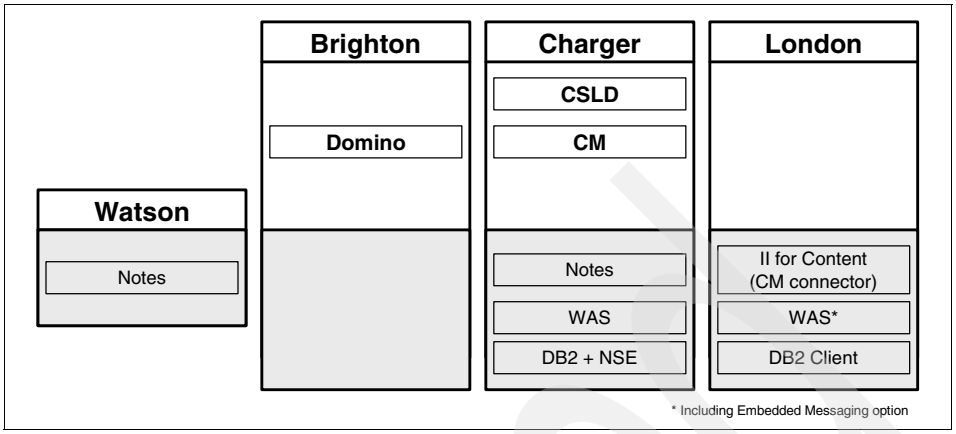

Figure 6-4 Sample environment after CSLD is installed and configured

Note: The various forms, views, and libraries from the CSLD sample template are not meant for production use but to be used as a guide to incorporating CSLD and Records functions into your corporate Notes template.

6.6.2 Key information to remember

At this point, as part of your overall configuration setup documentation, write down the key configuration information that will be used during the rest of the software installation and configuration and future system operation.

Table 6-32 contains the key configuration input values to remember after the CommonStore for Lotus Domino installation and configuration.

Table 6-32 Key information to remember after CSLD installation

Configuration data	Sample input value
Domino user ID used by CommonStore	CSLD Task
Content Manager user ID used by CommonStore	CSLD
Item type used by CommonStore	DominoMail

6.7 Records Manager installation and configuration

In the e-mail archiving and records management solution, Records Manager is used as a records administration application and as an engine for records enabling Content Manager (via Records Enabler for Content Manager), thus

records enabling the entire e-mail archiving solution. In this section, we describe the main steps involved in installing and configuring Records Manager.

These steps are as follows:

1. Install Records Manager engine V4.1.1.
2. Install Records Manager database V4.1.1.
3. Upgrade Records Manager engine V4.1.2.
4. Upgrade Records Manager database V4.1.2.
5. Run engine configuration utility.

Note: It is not our intention to include all detailed steps of the installation in this section. We recommend using the existing product manual in conjunction with the materials we present here for a successful installation and configuration.

In our sample environment, we install Records Manager engine on London and Records Manager database on Charger.

Step 1: Install Records Manager engine V4.1.1

In Table 6-33, we provide the input value we used during our installation on London in the sample environment. Replace the sample input values according to your environment setup.

Table 6-33 Records Manager engine installation input for sample environment

Configuration window / field	Sample input value	Description
Installation Destination		
Directory Name	C:\IBM\IRM	Records Manager installation directory. **TIP**: Use a directory without the version number. For updates, the same directory can be used and no second directory with a different version number is created.
Installation Type		
Setup type	typical	
Deployment and Configuration		
I want the installer to do deployment and configure for me	selected	The setup program deploys all J2EE applications and configures them.

Configuration window / field	Sample input value	Description
WebSphere Application Server Connection Information		
Connector Type	SOAP	Specifies the type of communication interface between the WebSphere Application Server and the installation program.
Connector Port	8880	Specifies the port used by the Connector Type.
Cell	london	Specify the cell name of the WebSphere Application Server installed under 6.4.3, "WebSphere Application Server installation" on page 145. **Tip**: To view the cell name for a server, start the WebSphere Application Server, open the WebSphere Application Console. In the console navigation tree, click **Servers** → **Application servers** → **Server 1** → **Runtime**.
Node	london	Specify the node name of the WebSphere Application Server installed under 6.4.3, "WebSphere Application Server installation" on page 145. **Tip**: To view a node name, go to WebSphere Application Console, click **Servers** → **Application servers** → **Server1** → **Runtime**.
Server	server1	As this book is being written, the only possible working value is server1.
Security Enabled	unchecked	*Required* if WebSphere security is enabled. If security is enabled, this option has to be checked and a valid user name and password must be provided. To verify whether security is currently enabled, ensure that the WebSphere Application Server is started, open the WebSphere Application Console. In the console navigation tree, click **Security** → **Global Security**, then verify the **Enabled** setting in the **Configuration** tab.

Configuration window / field	Sample input value	Description
Connection Factories Authentication		
Connection Factories Authentication User	irmwas	This Windows user ID has to exist. It does not need any special rights and is used just for authentication. The user ID must not be longer than eight characters (which is the current WebSphere limitation). An application component uses a connection factory to access a connection instance, which the component then uses to connect to the underlying enterprise information system (EIS). Examples of connections include database connections, Java Message Service connections, and SAP R/3 connections.
Connection Factories Authentication password		
Mail Session Configuration		
Mail Transport Host	Brighton	Specifies the name of the server to access for the engine to send e-mail.
SMTP User Name		This is required only if SSL is configured on the SMTP Server. Specifies the name of an e-mail user who has access to send e-mail through the specified transport host. Leave this field blank if the transport host does not require authentication.
SMTP User password		Required only if SSL is configured on the SMTP Server. *Required* for the engine to send e-mail. Specifies the password of an e-mail user who has access to send e-mail through the specified transport host. Leave this field blank if the transport host does not require authentication.

Configuration window / field	Sample input value	Description
Administration Client Configuration		
Root	IRMClient	Specifies the context root for your Records Manager Administrator client. This is the name that you use to access the client for Records Manager in your browser (the virtual directory name). For example: `http://london.redbook.bocaraton.ibm.com:9080/IRMClient`
Engine Server Name	London.redbook.bocaraton.ibm.com	Specifies the host name of the computer where the Records Manager engine is installed.
Engine Server ORB Port	8880	Specifies the JNDI service port for the host where you are installing the Records Manager engine. This field specifies the port number on which the application server Object Request Broker (ORB) listens for requests.
Web Services Configuration		
Web Service Configuration Root	IRMWebServices	Specifies the context root for the Records Manager Web server. Name used to access the Web services for Records Manager in a browser (the virtual directory name).
Web Services Node Name	London.redbook. bocaraton. ibm.com	Specifies the host name of the computer where the Records Manager engine is installed.
Web Services HTTP Port	9080	Specifies the number for the port that the WebSphere Application Server uses for message queues.
Import Export Configuration		
Engine Server Name	London.redbook. bocaraton. ibm.com	The host name of the computer where the Records Manager engine is installed.

Configuration window / field	Sample input value	Description
Engine Server ORB Port	2809	Specifies the JNDI service port for the host where the Records Manager engine is installed. This field specifies the port number on which the application server Object Request Broker (ORB) listens for requests.
WebSphere Location		
WebSphere Location	C:\IBM\WS\WAS	The WebSphere Application Server installation directory.

Step 2: Install Records Manager database V4.1.1

Table 6-34 shows the input values we used during our installation on Charger in the sample environment. Replace our sample input values according to your environment setup.

Table 6-34 Records Manager database installation input for sample environment

Configuration window / field	Sample input value	Description
Installation Directory		
Directory Name	C:\IBM\IRM\Database	This directory is not the default installation directory.
Database Type		
Database Type	DB2	
JDBC™ driver class path	C:\IBM\DB2\SQLLIB\java	Path to the *db2java.zip* file. In the sample environment, this file is located in C:\IBM\DB2\SQLLIB\java.
DB2 Database Configuration		
DB2 Node/Instance Name	DB2	The instance name if the Records Manager database is being installed directly onto a DB2 server. If a remote database is being used, then this is the name of a cataloged DB2 instance. **Note:** This name cannot exceed eight characters in length.

Configuration window / field	Sample input value	Description
Database Name	irmdb	The name of the DB2 database that is being created. **Note:** The database name cannot exceed eight characters in length, and it must be unique for each database you create.
Default Disk	C	The default location where the database and dataset files are being created. For example, for Windows, type C for the C:\ drive.
Folder for Database container	C:\IBM\IRM\Database	The default location where the database and dataset files are being created. This location will have the containers of the table spaces for the database to create.
User name	irmadmin	Specifies the name of the DB2 user that will be the owner of the Records Manager schema. This user will have Database administration privileges in the newly created database. **Important:** This user must exist, as it is not created automatically during the installation.
User password		
Territory	default	Specifies a portion of the locale mapped to the country code for the internal processing by the database manager.
Collating System	System	Specifies the sequence in which characters are ordered for the purpose of sorting, merging, comparing, and processing indexed data sequentially.

Configuration window / field	Sample input value	Description
DB Language	English	(optional) Specifies a language identifier. Set this field when the database language is different from the default language on your computer. The language you specify must be available on the computer where you are performing the installation.
System administrator user name	db2admin	Specifies the name of the database user with system administrator privileges for the DB2 database instance. This user is created during the installation of the DB2 database in 6.4.1, “DB2 server installation” on page 143.
System administrator user password		Specifies the password of the database user with system administrator privileges for the DB2 database instance.
Database File Plan Population		
Select a plan to populate database	Sample	A sample file plan is created.

Step 3: Upgrade Records Manager engine V4.1.2

The Records Manager engine upgrade is basically a redeploy of the WebSphere Application Server. When running the upgrade, re-enter the values you provided earlier.

Tip: Use the same installation directory (C:\IBM\IRM\); otherwise, a second directory will be created and the old one will not be deleted.

Step 4: Upgrade Records Manager database V4.1.2

After the Records Manager engine is upgraded to V4.1.2 level, upgrade the Records Manager database to the same level. Table 6-35 on page 185 shows the input values we used while we upgraded the sample environment.

Table 6-35 Records Manager database upgrade input for sample environment

Configuration window / field	Sample input value	Description
Installation Directory		
Directory Name	C:\IBM\IRM\Database	Use the same directory as specified during the original installation of the Records Manager database.
Custom or Automatic upgrade		
Automatic	Selected	
Database Type		
Database Type	DB2	
JDBC driver class path	C:\IBM\DB2\SQLLIB\java	
DB2 Database Configuration		
DB2 Node/Instance Name	DB2	Instance name as specified in "Step 2: Install Records Manager database V4.1.1" on page 182. This is the instance name if the Records Manager database is installed directly on a DB2 server. It is the name of a cataloged DB2 instance, if a remote database is being used.
Database Name	irmdb	The name of the database that is created in step 2.
Folder for Database container	C:\IBM\IRM\Database	The folder for the database container specified in step 2.
User name	irmadmin	Name of the user specified in step 2.
User password		
System administrator user name	db2admin	Specifies the name of the database user with system administrator privileges for the DB2 database instance.

Configuration window / field	Sample input value	Description
System administrator user password		Specifies the password of the database user with system administrator privileges for the DB2 database instance.
Database Back Up		
Selected database was backed up	Selected	**Important**: If this is not selected, the upgrade cannot proceed.

Step 5: Run engine configuration utility

The Records Manager engine needs to access the Records Manager database. Because the database can be on different platforms (DB2,Oracle, SQL Server), a data source must be configured. This data source is used by the Records Manager engine to access the database.

If the Records Manager engine and Records Manager database are running on different machines, the database must be cataloged on the engine machine. In the sample environment, the engine runs on London, and the database is on Charger. We need to catalog the database on London. Use the DB2 Configuration Utility to catalog the database that was created in “Step 2: Install Records Manager database V4.1.1” on page 182. In the sample environment, the remote database on Charger (irmdb) is cataloged as irmdb on London.

The Records Manager engine configuration utility must be run on the same machine that the Records Manager engine is running. To start the utility, select **Start → Program Files → IBM DB2 Records Manager → Engine Configuration Utility**.

In Table 6-36, we provide the input values we used during the start up of the utility for the sample environment. Replace our sample input values according to your environment.

Table 6-36 Engine configuration utility startup

Configuration field	Sample input value	Description
Connector Type	SOAP	
Port Number	8880	
Cell	london	The WebSphere Application Server cell on which the Records Manager Engine is deployed.

Configuration field	Sample input value	Description
Node	london	The WebSphere Application Server node on which the Records Manager Engine is deployed.
Server	server1	The WebSphere Application Server server into which the Records Manager Engine is deployed.

After the engine configuration tool is started, a data source (the Records Manager database created in "Step 2: Install Records Manager database V4.1.1" on page 182) must be created.

To create the data source, select **Action** → **New**.

Table 6-37 shows the input values we used when creating the new data source. Replace these values with the appropriate ones for your environment.

After creating the new data source, select **File** → **Save Changes**.

Table 6-37 Data source input for the sample environment

Configuration field	Sample input value	Description
DB2 Universal JDBC driver location	C:\IBM\DB2\SQLLIB\java	
Data Source Name	irmdb	Name of the database. You can use any name.
Database Name	irmdb	Name of the cataloged database.
User name	irmadmin	Name of the user created in step 2 who has administrative rights for the database.
User password		

Important: After configuring the data source, if the utility is closed without saving, the provided information will be lost and the data source will not be available during the Records Manager Administration client startup.

Before the Records Manager Administration client can be used, the WebSphere Application Server must be restarted.

Use the Windows Services utility to restart the service “IBM WebSphere Application Server V5 - server1,” or go to the WebSphere bin directory (C:\IBM\WS\WAS\bin) and use the following commands to restart the server:

```
stopserver server1
startserver server1
```

6.7.1 Installation summary and verification

At this point, the Records Manager engine and the Records Manager database are installed.

In the sample environment, the recommended scenario of installing them on two different computers is performed. Figure 6-5 shows the sample environment after the successful installation.

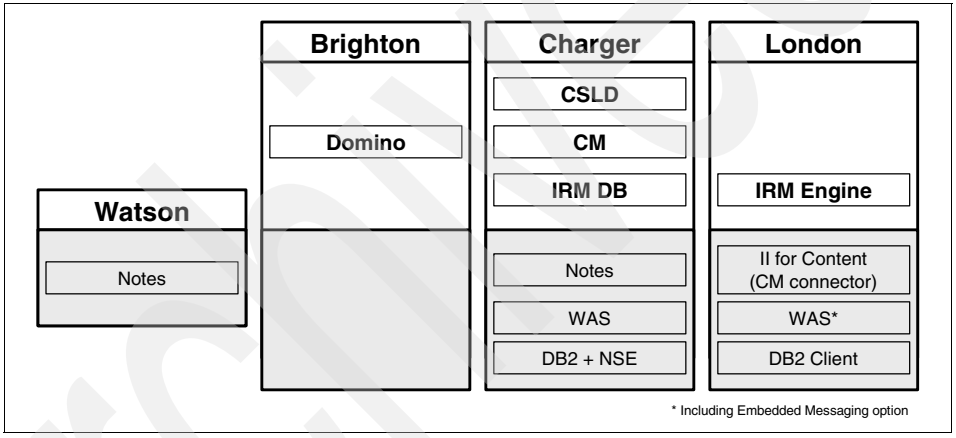

Figure 6-5 Sample environment after Records Manager is installed

To verify the installation, start WebSphere Application Console, go to **Servers** → **Application servers** → **server1** → **Server Components** → **JMS servers**, and make sure that `Started` appears in the Initial State field.

To verify that messaging queue is up and running, go to Windows Task Manager and look for processes with amq* and runq*. If both of them are there, then the message queue has started. If not, perform these steps:

1. Go to **Servers** → **Application servers** → **server1** → **Server Components** → **JMS servers**.
2. Make sure **Started** is selected in the Initial State.
3. Click **Reset**.

4. Go back and check the amq* and runq* processes in the Task Manager. They should be there now.

Attention: The procedure above is very important. If the message queue is not started, you will encounter problems when working with Records Manager.

To make sure Records Manager is installed properly, log on to Records Manager Administration client using `Administrator` as the user ID and `cronos` as the password. Make sure you can log in.

You can also use a set of basic Records Manager activities to ensure that the Records Manager system is up and running. The suggested activities include:

1. Create a file plan.
2. Create a record.
3. Perform record scheduling.
4. Turn the crank.
5. Destroy record via retention rule.

6.7.2 Key information to remember

At this point, as part of your overall configuration setup documentation, write down the key configuration information that will be used during the rest of the software installation and configuration and future system operation.

Configuration data	Sample input value	Description
Records Manager Administrator user ID	Administrator password: cronos	
WebSphere Application Server server name	server1	Name of the server into which the Records Manager engine is deployed.

6.8 CRME installation and configuration

Records Enabler for Content Manager (CMRE) is the bridge between Content Manager and Records Manager. It works with both products to provide the records control capability in the Content Manager system. In this section, we describe the main steps involved in installing and configuring CMRE.

These steps include:

1. Set environment variables and create users.

2. Install Records Enabler (CMRE server, Host Interface server, and Permission Synchronization server).
3. Install Records Manager Extension.
4. Implement Windows Services.

Important: Make sure that the Content Manager V8 connector is installed. It is necessary to use the Information Integrator for Content installation to install the connector. A Content Manager client installation is not sufficient.

Note: We do not include all of the detailed steps of the installation in this section. We recommend using the existing product manual in conjunction with the materials we present here for a successful installation and configuration.

In the sample environment, we install Records Enabler on London.

Step 1: Set environment variables and create users

Set the environment with the following variables:

```
IBMCMROOT = C:\IBM\CM\db2cmv8
WAS_HOME = C:\IBM\WS\WAS
JDBCPATH = C:\IBM\DB2\SQLLIB\java\db2java.zip
```

Substitute the values according to your system installation setup.

Important: The JDBCPATH must include the db2java.zip file name; otherwise, the installation will fail.

Create a local user on the machine that runs the Content Manager Library Server. This user has to be in the DB2 administrator group.

For the sample environment, we create:

- User name: CMREID
- User group: DB2ADMNS

Step 2: Install Content Manager Records Enabler (CMRE)

You need to install CMRE. To help your CMRE installation process, we provide the input values we used during our installation in Table 6-38 on page 191. Replace our sample input values according to your environment setup.

Table 6-38 Records Enabler installation

Configuration window / field	Sample input value	Description
WebSphere deployment information		
DB2 Content Manager Records Enabler Server	Selected	
DB2 Content Manager Records Manager Host Interface	Selected	
DB2 Content Manager Records Enabler Permissions Synchronization	Selected	
WebSphere Application Server cell name	london	Specify the cell name of the WebSphere Application Server installed under 6.4.3, "WebSphere Application Server installation" on page 145. **Tip**: To view the cell name for a server, start the WebSphere Application Server and open the WebSphere Application Console. In the console navigation tree, click **Servers** → **Application servers** → **Server 1** → **Runtime**.
WebSphere Application Server node name	london	Specify the node name of the WebSphere Application Server installed under 6.4.3, "WebSphere Application Server installation" on page 145. **Tip**: To view a node name, go to WebSphere Application Console, click **Servers** → **Application servers** → **Server1** → **Runtime**.
Host name	London.redbook.bocaraton.ibm.com	The fully qualified host name of the machine running the WebSphere Application Server.

Configuration window / field	Sample input value	Description
WebSphere Application Security Enabled	unchecked	Required if WebSphere security is enabled. If security is enabled, this option has to be checked and a valid user name and password has to be provided. To verify whether security is currently enabled, ensure that the WebSphere Application Server is started, go to the WebSphere Application Console, click **Security** → **Global Security**, and verify the **Enabled** setting on the **Configuration** tab.
Records Manager Server configuration		
Records Manager Web services address	Charger.redbook.bocaraton.ibm.com:2809	
Records Manager Administration client URL	http://charger.redbook.bocaraton.ibm.com:9080/IRMClient	This value is checked during the installation. If it is not available, installation will not go further.
Records Manager database	irmdb	The name of the Records Manager database.
Records Manager Administrator	Administrator	Records Manager administrative user ID. The default ID is Administrator.
Password	cronos	The default password for the Records Manager Administrator is `cronos` if it is not changed after the IRM installation.
Content Manager Server configuration		
Server name	icmnlsdb	The name of the Library Server database.
Content Manager authentication	icmadmin	Administrative user ID for the Content Manager System. See 6.5.2, “Key information to remember” on page 152.
Password		
Content Manager Records Enabler Connection ID	cmreid	This user ID is created by the installation program within Content Manager, but it has to exist on the operating system level on the machine running the Content Manager Library Server.

Configuration window / field	Sample input value	Description
password		Password of the user ID on the Windows operating system level.
confirm password		
eClient rendering Content URL	http://....	In the sample environment, eClient is not installed. **Important:** Leave the default value and do not erase that field. With an empty field, the installation will fail. This value can be configured later using the CMRE Administration client.
eClient document list URL	http://...	The eClient is not installed in the sample environment. **Important:** Leave the default value and do not erase that field. With an empty field, the installation will fail. This value can be configured later using the CMRE Administration client.
Database System used for Content Manager	DB2	Specifies the database type of the Content Manager Library Server.
Records Enabler configuration		
CMRE Server	cmresvr	The server will be created during the installation if it does not exist.
Records Manager Host Interface Server	rmecmhost	The server will be created during the installation if it does not exist.
Add Host Configuration record to DB2 Records Manager	checked	Checking this creates a “Host” entry within the Records Manager. The specified Content Manager system will be registered within the Records Manager.
Content Manager Records Enabler Permissions Synchronization	cmrepsproc	The server will be created during the installation if it does not exist.
Permissions Synchronization Scheduler	checked	

Configuration window / field	Sample input value	Description
Permissions Synchronization engine	checked	

Step 3: Installing Records Manager Extension

You must install Records Manager Extension. To help your installation process, we provide the input values we used during our installation in Table 6-39. Replace our sample input values according to your environment setup.

Table 6-39 Records Manager Extension installation input for sample environment

Configuration window / field	Sample input value	Description
WebSphere Deployment information		
WebSphere Application Server cell name	london	Specify the cell name of the WebSphere Application Server installed under 6.4.3, "WebSphere Application Server installation" on page 145. **Tip**: To view the cell name for a server, start the WebSphere Application Server, open the WebSphere Application Console, and click **Servers** → **Application servers** → **Server 1** → **Runtime**.
WebSphere Application Server node name	london	Specify the cell name of the WebSphere Application Server installed under 6.4.3, "WebSphere Application Server installation" on page 145. **Tip**: To view the cell name for a server, start the WebSphere Application Server, open the WebSphere Application Console, and click **Servers** → **Application servers** → **Server 1** → **Runtime**.
Host name	london.redbook.bocaraton.ibm.com	The fully qualified host name of the machine running the WebSphere Application Server.

WebSphere Application Security Enabled	unchecked	Required if WebSphere security is enabled. If security is enabled, this option has to be checked and a valid user name and password have to be provided. To verify whether security is enabled, ensure that WebSphere Application Server is started, open the WebSphere Application Console, and in the console navigation tree, click **Security** → **Global Security**, and verify the **Enabled** setting on the **Configuration** tab.
Records Manager Application server name	server1	At the time this book was written, the only possible working value was server1.

Step 4: Implementing Windows Services

To install a WebSphere Application Server as a Windows Service, use the following command:

```
wasservice -add <Windows Service Name> -serverName <WebSphere Server>
```

To start the Windows services tasks:

1. Open a Command prompt window.
2. Go to the WebSphere Application Server bin directory.
3. Install CMRE server (cmresvr), CMRE Host Interface (rmecmhost), and CMRE Permission Synchronization server (cmrepspro) as Windows services, substituting the appropriate values from your environment:

```
wasservice -add RMEServer -serverName cmresvr
wasservice -add RMEHostInterface -serverName cmresvr
wasservice -add RMEPermSync -serverName cmresvr
```

6.8.1 Installation summary and verification

Figure 6-6 on page 196 shows the components that are installed and configured after this section is completed.

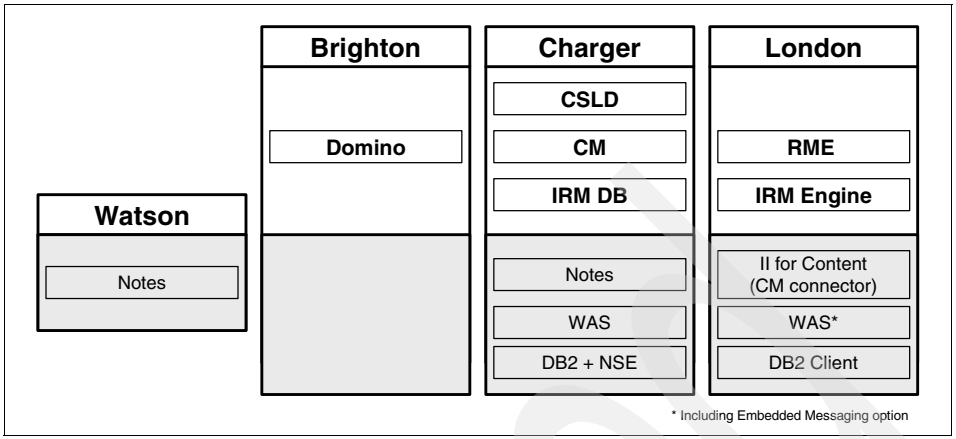

Figure 6-6 Sample environment after Records Enabler is installed

To verify your installation, perform the following steps:

1. Check the Windows services.
2. Import a Content Manager administrator ID into Records Manager.
3. Log on to the CMRE.
4. Records enable an item type.
5. Use Content Manager Windows client to declare a record.

Step 1: Check the Windows services

To verify the installation, make sure all necessary services are running. They are:

- IBM WebSphere Application Server V5 - server1
- IBM WebSphere Application Server V5 - RMEServer
- IBM WebSphere Application Server V5 - RMEHostInterface
- IBM WebSphere Application Server V5 - RMEPermSyncServer
- ICM LS Monitor icmnlsdb
- IBM HTTP Server 1.3.28

Step 2: Import a CM administrator ID into Records Manager

To be able to log on to the CMRE Administration client, a Content Manager user ID that has the administrator rights must be imported to the Records Manager System. The user ID must have administrator rights in Records Manager.

To import a Content Manager user:

1. Start the Records Manager client and log on as an administrator:

   User ID: Administrator

   Password: cronos

2. Go to **Security** → **Users** → **Host Filer** and select the host system that is enabled during the CMRE installation.

   In the sample environment, the host system with the name icmnlsdb is enabled. (See Add Host Configuration record to DB2 Records Manager of “Step 2: Install Content Manager Records Enabler (CMRE)” on page 190.)

3. Click **Import**, select **icmadmin**, and click **Import** again. In the next window, select all permissions by checking **Function Access**. Check the **Is Active** check box. Click **Save** to finish the import.

4. The Content Manager user ID icmadmin is now imported to the Records Manager system and has all necessary rights to act as Administrator with the Records Manager system.

Step 3: Log on to the CMRE

Use icmadmin and its password to log on to the CMRE Administration client.

Important: Do not use the Records Manager administrator user ID, (Administrator) and its password (cronos) to log on to the CMRE Administration client because this ID is not defined within Content Manager.

Also, do not use the Content Manager user ID, cmreid, that is created during the CMRE installation, because this ID is not (and cannot be) imported into Records Manager.

Step 4: Records enable an item type

After successfully logging on to the CMRE Administration, enable an item type:

1. Go to **Content Manager Server Configuration** → **eRecord enable item type**.

   A list of all item types in the records-enabled Content Manager system is displayed.

2. Select the CSLD item type (DominoMail) that was created in step 2a of 6.6, “CommonStore (CSLD) installation and configuration” on page 153.

3. Check the box left of the item type name and select the record type to the right of the item type name.

In the sample environment, the item type `DominoMail` and the record type `email` are created.

Step 5: Using Windows client to declare records

Using Windows client, declare records as follows:

1. Search for documents in the records-enabled item type DominoMail. E-mail archived during the CommonStore verification is located in this item type.
2. Right-click one e-mail and select **Declare Record**. The Records Manager declare/classify window appears.
3. Select a file plan and a unique name for that record. Click **Finish**.

The CMRE communicates with the Content Manager system and the Records Manager system. A Content Manager item type is records enabled and documents within this item type can be declared as a records using the Content Manager Windows client. The records attributes (eRecord, eRecordID, FlPlnCmpntNm, FlPlnCmpntTtl, and RMEAclOri) that are associated with the e-mail are updated.

6.8.2 Key information to remember

At this point, as part of your overall configuration setup documentation, write down the key configuration information that will be used during the rest of the software installation and configuration and future system operation.

Table 6-40 contains the key configuration input values to remember after the Content Manager installation.

Table 6-40 Key information to remember after CMRE installation and configuration

Configuration data	Sample input value
Records enabled item type	DominoMail

6.9 Configuring the CommonStore Server and Notes

In this section, we describe the main steps involved in configuring CommonStore Server and Notes. These steps must be performed to add records declaration and classification capability to the end-to-end integrated e-mail solution.

The steps include:

1. Records enable the Content Manager item type.
2. Prepare CommonStore.
3. Prepare Domino.
4. Records enable the Notes client.

For detailed information about CSLD and Records Enabler integration, refer to Chapter 21, “Using Content Manager Records Enabler in the CSLD environment” in *IBM DB2 CommonStore for Lotus Domino: Administration and User’s Guide, Version 8.3*, SH12-6742.

Step 1: Records enabling the Content Manager item type

This should have been done in the previous section when you validated the installation and configuration of CMRE. If it has not been completed, refer to “Step 4: Records enable an item type” on page 197.

Step 2: Preparing CommonStore

A Lotus Notes user must be mapped to a Content Manager user ID in order for the Notes user to declare a record because the backend repository for records is the Content Manager system. This Content Manager user also must be imported to Records Manager and has to have the appropriate rights to declare records. For the mapping of Lotus Notes user IDs to Content Manager user IDs, an additional component (usermapper.jar) has to be activated on the CommonStore ArchPro server.

To activate the usermapper, follow the steps below:

1. Copy usermapper.jar from the installation bin directory into the instance directory of the ArchPro that will be activated.

   Source directory: C:\IBM\CSLD\bin

   Target directory: C:\IBM\CSLD\server\instance01

2. Update CSExit.properties.

   CSExit.properties is located in the instance directory and contains three parameters that must be configured:

   ```
   DB_DIR
   HASH_MODULO
   PROXY_PORT
   ```

   Set DB_DIR to the directory to store the mapping database. (Note the double backslashes, as a single backslash would indicate an escape sequence, such as \n.) This directory will contain a collection of files that are serialized hash tables containing string keys of the format CM-server:mail-user and CSRepUserDef values. Make sure this folder exists and is empty. The files will be generated automatically, along with a file that indicates the current HASH_MODULO value so that it can be changed.

   Set HASH_MODULO to the maximum number of files to be found in the DB_DIR directory. This is to prevent the entire database from ever needing to be in memory at one time so that a huge number of users can be supported. Bigger values mean smaller memory usage (but more files).

Set PROXY_PORT to the port on which the usermapper proxy is listening.

See Table 6-41 for the input values we used for the sample environment.

Table 6-41 CSExit properties input values

Parameter	Sample input value
DB_DIR	C:\IBM\CSLD\server\instance01\database
HASH_MODULO	1000
PROXY_PORT	8067

3. Add usermapper.jar and Notes.jar into the classpath as well as the directory of CSExit.properties.

Table 6-42 Classpath for usermapper.jar, Notes.jar, and CSExit.properties

File/directory	Sample value to add to classpath
usermapper.jar	C:\IBM\CSLD\server\instance01\ usermapper.jar
Notes.jar	C:\IBM\Notes\Notes.jar
CSExit.properties directory	C:\IBM\CSLD\server\instance01

4. Update archint.ini.

   To activate usermapper, add the values for these parameters in archini.ini:

```
ACCESS_CTL
CM_SECURITY_EXIT
CM_EXIT_LOCATION
```

   The input value for ACCESS_CTL YES specifies whether you want Retrieve operations to be subject to the user's Content Manager permissions.

   The input value for CM_SECURITY_EXIT specifies the name of the security exit class as com.ibm.rme.csexit.CSExit.

   The input value for CM_EXIT_LOCATION specifies the file location of the usermapper.jar file.

   In our sample environment, we set these values as shown in Table 6-43.

Table 6-43 archint.ini file update

Parameter	Sample input value
ACCESS_CTRL	YES
CM_SECURITY_EXIT	com.ibm.rme.csexit.CSExit
CM_EXIT_LOCATION	'C:\IBM\CSLD\server\instance01\usermapper.jar'

Step 3: Preparing Domino

Prepare Domino as follows:

1. Modify the CSLD configuration database:
   a. Open the profile of the archiving task.
   b. Go to **Advanced**.
   c. In the section Write state to, select **Special field**.
   d. Accept the default value and save the profile.
2. Modify the template CSLDStdMail.ntf.

   The sample mail template is records enabled and must be configured:

   a. Open **CSLDStdMail.ntf** with the Notes Designer.
   b. Go to **Shared Code** → **Script Library** → **RMEScriptLibrary** and provide information using the sample input values in Table 6-44 for reference.

Table 6-44 RMEScriptLibrary update

Parameter	Sample input value	Description
RMEServerURL_Default	http://LONDON.redbook.bocaraton.ibm.com:9082/RMEServer/RMEClientServlet	
CMHostName_Default	charger.rebdook.bocaraton.ibm.com	
CMItemTypeName_Default	icmnlsdb	
UserProxyServerName_Default	charger.rebdook.bocaraton.ibm.com	
UserProxyServerPort_Default	8067	Defined in CSExit.properties
CSLDArchiveStatusField_Default		
RefreshInterval_Default		
RefreshTotal_Default		
RMEFolderClassifyTotal		
CSarchAction_Default		
WebServerPort	80	

Parameter	Sample input value	Description
CScreatePlaceholderAs URL_Default		This value is configured as true if you want the system to create hot links for attachments in e-mail after archiving; otherwise, this value is configured as false.

Attention: Changes to those variables will not affect a database that has already had a template applied, because a configuration document is already created and this document will not be refreshed if the values are updated.

Note: The various forms, views, and libraries from the CSLD sample template are not meant for production use, but rather as a guide to incorporating CSLD and Records functions into your corporate Notes template.

3. Set up Domino security:
   a. Log on to the Notes administration client using an Administrator user ID.
   b. Create a new user group (RMEUserGroup):
      i. Go to the **People&Group** tab.
      ii. Click **Groups** → **Add group**.
   c. Add Notes users that need to declare records to this group.
   d. Grant security to the group:
      i. Go to the **Configuration** tab.
      ii. Expand **Server** → **Current Server Document**.
      iii. Click **Security** and grant with the following security permission:

```
Run unrestricted methods and operations
Run restricted Lotus Script/Java agents
Run simple and formula agents
Run restricted Java/Javascript/COM
Run unrestricted Java/Javascript/COM
```

   e. Install the RMEAuth filter in the Domino server:
      i. Find the RMEAuth.dll in the CSLD package and copy it to the Domino data directory.
      ii. Install the RMEAuth filter by specifying the name of the filter in the Domino server record, in the field DSAPI filter file name in the Internet Protocols → HTTP table. You can specify just the name of the filter file

(RMEAuth) if it is located in the Domino program or data directories; otherwise you must specify the fully qualified path name.

iii. Restart the Domino server and the HTTP task.

Step 4: Records enable the Notes client

To records enable the Notes client, you must enable menus and buttons for the records management functions.

Complete the following steps:

1. Incorporate the CSLD functions into the e-mail template as outlined in *CSLD V8.3 Administrator's and Programmer's Guide,* SH12-6742.
2. Update the notes.ini file to include the following line:

   ```
   $RecordsEnabler=yes
   ```

3. Restart the Notes client.

The records management related menus and buttons will be available to use.

6.9.1 Verification

To verify that the CommonStore Server and Notes client are configured properly:

1. In Notes client, create an e-mail message and send it off.
2. Manually declare the e-mail message as a record.
3. Make sure that the e-mail is archived in Content Manager using the Content Manager Windows client.
4. After the e-mail is archived, the Records Manager classification window should come up. Specify the bucket in the file plan; for example, Account Receivable in our sample environment. Enter a unique ID into the e-mail name field, and click **Finish**.
5. Make sure that in Notes client reflects that the e-mail is a record.
6. Using Content Manager Windows client, verify that the e-mail Record attribute is set to yes.
7. Using the Records Manager Administration client, check that the record is there.
8. Use Notes client to view the e-mail message.

7

Installation and configuration in a Microsoft Exchange environment

This chapter describes the installation and configuration of an e-mail archiving and records management solution using CommonStore for Exchange Server, Records Manager, Content Manager, and Content Manager Records Enabler. Using a sample environment, we describe the major steps involved in installing and configuring the various components in a Windows environment. For more detailed information, see the appropriate product documentation.

We cover the following topics in the chapter:

- Overview
- Introduction to the sample environment
- Prerequisites and perquisite software installation
- Content Manager installation and configuration
- CommonStore (CSX) installation and configuration
- Records Manager installation and configuration
- CMRE installation and configuration
- Records enable CommonStore Server and Outlook

7.1 Overview

In this section, we provide an overview for the e-mail archiving and records management integrated solution installation and configuration.

We cover:

- Software used for the integrated solution
- Installation and configuration steps and recommendation

7.1.1 Software used for the integrated solution

Several products are used in this end-to-end integrated solution. We list the software used and the purpose in the solution in Table 7-1.

For clarity, fix pack details have been omitted. These are referenced later under solution installation and configuration.

Table 7-1 Software used in the integrated solution and the purpose in the solution

Product	Purpose
IBM DB2 Content Manager (CM)	Repository used to store the documents and metadata for both the archive and records management systems.
IBM CommonStore for Exchange Server (CSX)	E-mail archive system for Microsoft Exchange mail system.
IBM DB2 Records Manager (IRM)	Engine and administration for Records Manager.
Records Enabler for Content Manager (CMRE)	Records enables the Content Manager repository. Also provides records management functions for Lotus Notes or Outlook users.
IBM DB2 ESE UDB (DB2)	Enterprise-class database used to hold both system configuration and objects metadata.
IBM DB2 Net Search Extender	Extension to DB2 that adds full text search capabilities for both object metadata and documents including attachments.
IBM Web Sphere Application Server with Embedded Messaging	Web application server that hosts the Content Manager Resource Manager, Content Manager Records Enabler servers, the Records Manager applications, and Content Manager eClient.
Microsoft Exchange server	E-mail system.

7.1.2 Installation and configuration steps and recommendation

Four main products involved in the end-to-end solution:

- IBM DB2 Content Manager
- IBM DB2 CommonStore for Exchange Server
- IBM DB2 Records Manager
- IBM DB2 Records Enabler for Content Manager

Plan your system configuration first (see 5.3, “System configuration” on page 120). Decide where you want to install each main product and review the prerequisites for each product (especially if you decide to separate some components onto different machines). Before you start, make sure you know exactly what must be installed on each machine and the sequence of the installation.

Some of the product installation and configuration steps can be done at the same time and others are better done sequentially. We recommend the following sequence of steps for installation and configuration:

1. Install a working Content Manager system on a designated machine.

   This includes the installation of the prerequisites first (DB2, Net Search Extender, WebSphere Application Server), and then the Content Manager product.

   Validate that the Content Manager system is working by importing some documents, retrieving them, and viewing them.

   If you decide to separate the installation of Content Manager Resource Manager from Library Server onto different servers, you may need to install different prerequisites onto the machines. Refer to the product documentation to install the proper prerequisites for each server. You should also validate the system after the installation and configuration as mentioned above.

2. Install a working CommonStore system on a designated machine.

   This includes the installation of the prerequisites (Content Manager V8 Connector from the Information Integration for Content installation, DB2 Runtime Client or DB2 Administration Client, and Outlook client), and then the CommonStore for Exchange Server product.

   Validate that CommonStore is working properly with your mail database and Content Manager by setting up some policies, archiving some e-mail, retrieving the archived e-mail and viewing it. Also look into the Content Manager repository to ensure that the archived e-mail is there as expected.

3. Install a working Records Manager on a designated machine.

   This step can be done in conjunction with step 2. If you have multiple people doing the installation, you can work together in parallel. Otherwise, we recommend performing this task after the CommonStore installation.

   The step includes the installation of the prerequisites and then the Records Manager product.

   We recommend installing the Records Manager engine and its database on separate machines. The prerequisites for the engine include WebSphere Application Server and DB2 Runtime or DB2 Administration client. The prerequisite for the Records Manager database is DB2 server. As discussed in 5.3.1, “Configuration options” on page 122), you can optionally put the Records Manager database where Content Manager is installed.

   Validate that Records Manager is working properly by using the Records Manager administration client to create a default file plan, add a record to the system, and view the added record.

4. Install and configure Records Enabler for Content Manager on a designated machine. Configure Content Manager and Records Manager. Install and configure Records Manager Extension.

   Records Manager Extension must be deployed on the same WebSphere Application Server as Records Manager.

   Validate the system: Import an item of the record enabled item type, declare it as a record, and make sure that the record is marked as declared in Content Manager and the record’s metadata is in Records Manager. After this, archive an e-mail and declare the e-mail as a record. Check both Content Manager and Records Manager to ensure that proper information is stored in each. (Content Manager should have the e-mail information based on your archiving method, and Records Manager should have the record’s related metadata.)

Tip: For the machine that will run Permission Synchronization Server of the Records Enabler for Content Manager product, make sure that WebSphere Application Server with Embedded Messaging is installed on it. This is especially important if you start your installation from an existing system (for example, a working Content Manager system).

If the Embedded Messaging feature was not installed, you must deinstall Content Manager, uninstall WebSphere Application Server, then reinstall WebSphere Application Server, including the Embedded Messaging feature (installed by default in V5.1.1 and later), and then Content Manager again. Do not take shortcuts or you may experience strange results.

If the existing Content Manager is being used and cannot be deinstalled, we recommend installing the Permission Synchronization Server on another machine.

7.2 Introduction to the sample environment

We use a sample environment to show how to install and configure an e-mail archiving and records management solution.

Figure 7-1 on page 210 shows the sample environment before any e-mail archiving and records management software is installed.

This environment has one client machine, Watson, and three servers, Brighton, Charger, and London. An Exchange Server is running on Brighton. An Outlook client (including the CDO component, which is part of the Outlook installation package but not automatically installed) is running on Watson. This should be your starting point.

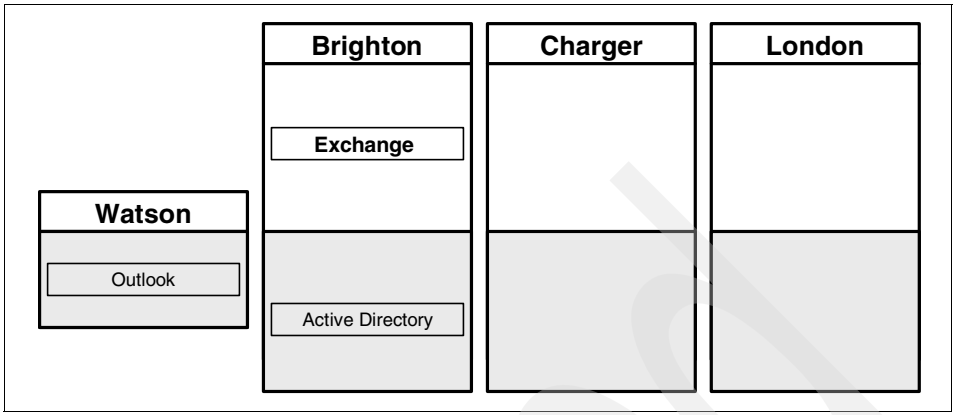

Figure 7-1 Sample environment before any software installation

Figure 7-2 shows the sample environment after all components are installed, including the necessary prerequisites on each server.

On Charger, we install prerequisite software including an Outlook client (Outlook), WebSphere Application Server, and DB2 server with Net Search Extender (DB2 + NSE). In addition, we install CommonStore for Exchange Server (CSX), Content Manager (CM), and IBM Records Manager database (IRM DB).

On London, we install prerequisite software including the Content Manager V8 connector (CM connector) from Information Integrator for Content, WebSphere Application Server with Embedded Messaging (WAS*), and a DB2 client. In addition, we install Content Manager Records Enabler (CMRE) and IBM Records Manager Engine (IRM Engine).

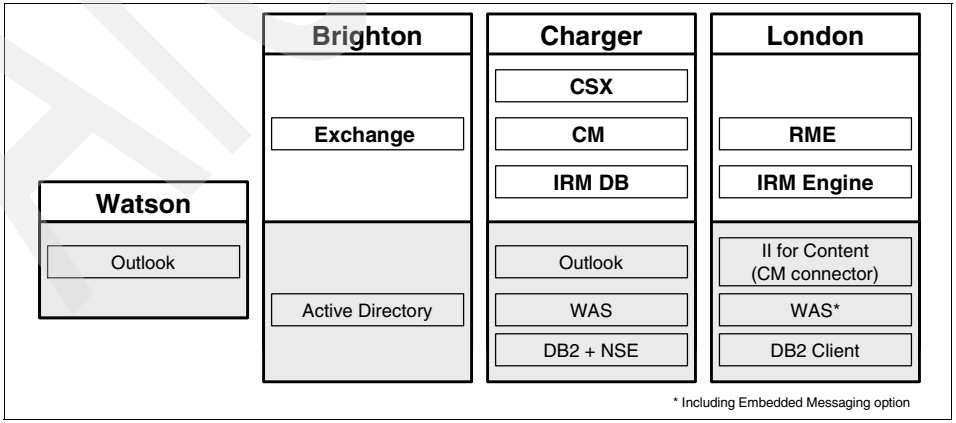

Figure 7-2 Sample environment after all software installation

7.3 Prerequisites

Four core components are involved in the e-mail archiving and records management solution: Content Manager for Multiplatforms, CommonStore for Exchange Server, IBM Records Manager, and Records Enabler for Content Manager.

Each core component has different prerequisites. We list what they are and their version requirements. In addition, we explain why you should install the prerequisite. Understanding this should help you when you build a system that fulfills your business needs.

Prerequisites for Content Manager V8.3

Table 7-2 describes the prerequisites for Content Manager V8.3.

Table 7-2 Prerequisites for Content Manager V8.3

Product	Version	Reason why we need it
WebSphere Application Server	5.1.1.2	Resource Manager is a J2EE application, thus needs WebSphere Application Server.
DB2	8.2	Library Server uses stored procedures and needs a relational database (icmnlsdb). Resource Manager stores metadata in a relational database (rmdb).
DB2 Net Search Extender	8.1	For full text search in Content Manager.

Prerequisites for CommonStore for Exchange Server V8.3

Table 7-3 lists the prerequisites for CommonStore for Exchange Server V8.3.

Table 7-3 Prerequisites for CommonStore for Exchange Server V8.3

Product	Version	Reason why we need it
Information Integrator (II) for Content - CM V8 connector	8.3.0	The CommonStore agent needs the APIs (connector) to communicate with the Content Manager Library Server.
DB2 Runtime Client	8.2	If CommonStore for Exchange Server is not on the same machine as the Content Manager Library Server, the CM V8 connector (which is needed by the CommonStore agent) needs the Library Server database to be cataloged on the machine where CommonStore is installed. DB2 Runtime Client is thus needed. For ease of installation, DB2 Administration Client is recommended.
Outlook client	2000, XP(2002), and 2003	The CommonStore Task communicates with the Exchange Server, so it needs the Microsoft API (MAPI) and the Collaborative Data Objects (CDO).

Prerequisites for IBM Records Manager V4.1.2

Table 7-4 describes the prerequisites for the Records Manager engine for IBM Records Manager V4.1.2.

Table 7-4 Prerequisites for IBM Records Manager Engine V4.1.2

Product	Version	Reason why we need it
WebSphere Application Server	5.1.1.2	Resource Manager is a J2EE application, thus needs WebSphere Application Server.
DB2 Runtime Client	8.2	The Records Manager Engine needs to communicate with the Records Manager database. If the Records Manager database is installed on another machine, then the machine with the Records Manager Engine installed must have the database cataloged on it, so it needs DB2 Runtime Client. For ease of installation, DB2 Administration Client is recommended.

Table 7-5 describes the prerequisites for the Records Manager database for IBM Records Manager V4.1.2.

Table 7-5 Prerequisites for IBM Records Manager (Database) V4.1.2

Product	Version	Reason why we need it
DB2	8.2	The Records Manager database is a relational database. The machine that installs the Records Manager database needs to install DB2 server.

Prerequisites for Content Manager Records Enabler V8.3

Table 7-6 describes the prerequisites for the CMRE server, PermSync server, and Host Interface server.

Table 7-6 Prerequisites for CMRE server, PermSync server, and Host Interface

Product	Version	Reason why we need it
WebSphere Application Server with Embedded Messaging	5.1.1.2	CMRE server, PermSync server, and the Host Interface require WebSphere Application Server. PermSync Server also needs WebSphere with Embedded Messaging.
Information Integrator (II) for Content - CM V8 connector	8.3.0	All three components communicate with Content Manager Library Server and therefore need Content Manager APIs.
DB2 Runtime Client	8.2	The CM V8 connector communicates with the Content Manager Library Server database. If the Content Manager Library Server is installed on another machine, then the machine with these servers installed must have the database cataloged on it; thus it needs DB2 Runtime Client. For ease of installation, DB2 Administration Client is recommended.

Table 7-7 describes the prerequisites for the Records Manager extension.

Table 7-7 Prerequisites for Records Manager Extension V8.3.0

Product	Version	Reason why we need it
WebSphere Application Server	5.1.1.2	The CMRE server, PermSync server, and the Host Interface are WebSphere Application Server applications. They need WebSphere Application Server V5.1 with Fix Pack 1 and cumulative Fix 2.

7.4 Prerequisite software installation

In this section, we describe the main steps for installing the basic software that must go in prior to the installation of the main components of the e-mail archiving and records management solution.

The prerequisite installations include:

- "DB2 server installation" on page 215
- "DB2 Administration Client installation" on page 217
- "WebSphere Application Server installation" on page 217
- "Information Integrator for Content (CM connector) installation" on page 220

Table 7-8 summarizes the prerequisites for each software product from 7.3, "Prerequisites" on page 211. Note the following information:

- DB2 client (runtime or administration) is needed for Records Manager (IRM) engine, and DB2 server is required for Records Manager database if they are installed on separate machines; otherwise, they need only DB2 server.
- WebSphere Application Server with Embedded Messaging (WAS* in Table 7-8) is required for Permission Synchronization server from CMRE.
- DB2 Net Search Extender is required if you need full text search capabilities.

Table 7-8 Prerequisites requirement per software product

CM	CSX	IRM	CMRE
WAS		WAS	WAS* (PermSync svr)
DB2 +NSE	DB2 client	DB2 client (engine) DB2 (database)	
	II for Content - CM V8 connector		II for Content - CM V8 connector
	Outlook client		

For the sample environment that we described in 7.2, "Introduction to the sample environment" on page 209, we install the following prerequisites on two servers:

- London:
  - DB2 Administration Client V8.2
  - Information Integrator for Content - Content Manager V8 connector
  - WebSphere Application Server (including Embedded Messaging) V5.1.1.2
- Charger:
  - DB2 server V8.2

- DB2 Net Search Extender V8.2
- WebSphere Application Server V5.1.1.2 (Embedded Messaging not necessary)
- Outlook

Note: We do not include all detailed steps of the installation in this section. We recommend using the existing product manuals in conjunction with the materials we present here for successful installations and configurations.

7.4.1 DB2 server installation

Content Manager and Records Manager use relational databases to store content (objects and records) metadata and system configuration information.

You must install DB2 server to the machine (or machines) that will store Content Manager databases and Records Manager databases.

The steps involved are summarized as follows:

1. Install DB2 server software.
2. Verify DB2 server installation.
3. Install DB2 Net Search Extender.
4. Install DB2 Fix Pack 8.

In our sample environment, both the Content Manager database and Records Manager database are located on Charger. We choose to install the DB2 server on this machine.

Installing DB2 server

To help your DB2 server installation process, Table 7-9 lists the input values we used during the DB2 server installation for our sample environment. Replace our sample input values according to your environment setup.

Table 7-9 DB2 server installation input summary for the sample environment

Required input field	Sample input value	Description
DB2 Version	8.1.7	Content Manager V8.3 requires at least DB2 V8.1.7.
Installation type	Typical	
Drive	C:\	
Installation directory	C:\IBM\DB2\SQLLIB	

Required input field	Sample input value	Description
DB2 Administration server user ID	db2admin	This Windows user ID is used by the DB2 Administration server (DAS) to log on to the system as a service.
DB2 administration group	DB2ADMNS	This group is created automatically. The name can not be changed. Every user ID that needs to be an DB2 administrator has to be in this group.
DB2 instance	DB2	

Verifying installation

To verify the DB2 server installation, create the sample database from the First Step menu. Make sure that the sample database is created successfully and that you can view the data.

Installing Net Search Extender

This is an optional step. If you need to use the full text search feature of Content Manager, install the Net Search Extender before installing Content Manager. After Content Manager is installed, you can activate the full text search feature.

Use the database administrative user ID (db2admin) to run the Net Search Extender as a service.

In our sample environment, we choose not to configure for full text search.

Installing DB2 Fix Pack 8

After installing the DB2 server and the Net Search Extender, install the DB2 Fix Pack 8. The Records Manager database requires DB2 V8.2 (which is equivalent of V8.1.8) as a prerequisite.

See appropriate documentation for details of the steps to install the fix pack.

Key information to remember

At this point, as part of your overall configuration setup documentation, write down the key configuration information that will be used during the rest of the software installation and configuration and future system operation.

Table 7-10 on page 217 contains the key configuration input values to remember after the DB2 server installation.

Table 7-10 Key information to remember after the DB2 server installation

Configuration data	Sample input value
DB2 Administration server user ID	db2admin
DB2 administration group	DB2ADMNS
DB2 database server host name	charger.redbook.bocaraton.ibm.com

7.4.2 DB2 Administration Client installation

We recommend installing the Records Manager database and Records Manager engine on two separate servers.

The Records Manager engine needs access to the Records Manager database. If they are not installed on the same machine, the Records Manager database has to be cataloged on the Records Manager engine machine. To catalog a database, a DB2 Runtime Client is required.

In addition to installing DB2 client to the machine where the Records Manager engine will run (if it and the Records Manager database are installed on separate machines), you should install the DB2 client on the machine where you will run CommonStore.

Although the DB2 Runtime Client is the minimum requirement for both instances, we recommend installing the DB2 Administration Client because it provides a simple graphical user interface to catalog a database and other powerful DB2 tools to administer a remote database.

The installation process is straightforward and we do not cover it here.

In the sample environment, we install the Records Manager database on Charger and the Records Manager engine on London; therefore, we install the DB2 client on London.

7.4.3 WebSphere Application Server installation

The three core components of the solution (Content Manager, Records Manager, and Records Enabler) rely on WebSphere Application Server, so you must install it on the machines where these products are installed.

Important: Before you continue, make sure that WebSphere Embedded Messaging is installed on the server that runs the Records Enabler; otherwise, you may encounter problems in future.

If the Embedded Messaging feature was not installed, you must deinstall Content Manager, uninstall WebSphere Application Server, then reinstall WebSphere Application Server, including the Embedded Messaging feature (installed by default in the latest version), and then Content Manager again. Do not take shortcuts or you may experience strange results.

The steps involved in WebSphere Application Server installation are summarized as follows:

1. Install WebSphere Application Server software.
2. Verify installation.
3. Install Fix Pack 1 and any accumulative fix packs.

In the sample environment, we have to install WebSphere Application Server on both Charger and London. We also must install Embedded Messaging on London because Records Enabler will be installed on the machine.

Installing WebSphere Application Server software

To help your WebSphere Application Server installation process, Table 7-11 shows the input values we used during our installation on both servers. Replace our sample input values according to your environment setup.

Table 7-11 Installation input summary for the sample environment

Required input field	Sample input value	Description
WebSphere Application Server version	5.1.0	
Installation type	Full	This includes Embedded Messaging. This must be installed on the machine that will run the Records Manager engine. **Note:** Full installation includes all of the sample applications, and that may increase the install time.
Installation directory for WebSphere Application Server	C:\IBM\WS\WAS	

Required input field	Sample input value	Description
Installation directory for IBM HTTP Server	C:\IBM\IHS\	
Installation directory for Embedded Messaging server and client	C:\IBM\WS\WSMQ	
Node name	for Charger: charger for London: london	We install WebSphere Application Server on both Charger and London servers.
Host name	for Charger: charger.redbook.bocaraton.ibm.com for London: london.redbook.bocaraton.ibm.com	
WebSphere Administrator user ID	wsadmin	This is the Windows user ID used to run the WebSphere services.

Verifying installation

To verify successful WebSphere Application Server installation, at the "WebSphere Application Server First Steps window, click **Verify Installation**. The message `Installation Verification is complete` should show up.

Installing Fix Pack 1 and cumulative fixes

In order to fulfill the prerequisite, install Fix Pack 1 and the cumulative Fix 3.

Important: Before the installation, stop all WebSphere services (server1) as well as the IBM HTTP Server.

Key information to remember

At this point, as part of your overall configuration setup documentation, write down the key configuration information that will be used during the rest of the software installation and configuration and future system operation.

Table 7-12 on page 220 contains the key configuration input values to remember after the WebSphere Application Server installation.

Table 7-12 Key information to remember after the installation

Configuration data	Sample input value
WebSphere cell	for Charger: charger for London: london
WebSphere node	for Charger: charger for London: london

7.4.4 Information Integrator for Content (CM connector) installation

Content Manager V8 connector (which comes with the Information Integrator for Content) is a prerequisite for Records Enabler for Content Manager (CMRE) and CommonStore. Install the Content Manager V8 connector on both the CommonStore machine and the Records Enabler machine if these products are not installed on the same system where Content Manager is installed.

Content Manager V8 connector must be installed before CMRE installation.

We recommend installing Content Manager first before you install the connector to any of the machines in question. This is because during the Content Manager installation, many values necessary to configure the connector will be defined.

Although we describe the connector installation in this section, defer the installation until the Content Manager installation is done as described in 7.5, “Content Manager installation and configuration” on page 221.

Installing Content Manager V8 connector

To install Content Manager V8 connector, launch the Information Integrator for Content installation process by selecting **Connector** → **Content Manager V8 Connector** from the appropriate installation windows. Refer to the product manual for specific installation instructions.

Table 7-13 on page 221 shows the input values we used during installation in our sample environment. Replace our sample input values for your environment setup.

Table 7-13 CM V8 connector installation input for sample environment

Required input field	Sample input value	Description
Database server type	DB2 Universal Database	This is the database used by the Content Manager System.
Library Server database	icmnlsdb	Library Server stores metadata in a relational database. This value specifies the name of the database.
Library server schema name	icmadmin	
Authentication type	Server	
Connection user ID	icmconct	A connection user ID that is used for clients to connect to Library Server database if they do not have a valid database user ID.
Connection user ID password		
Local	Checked	

7.5 Content Manager installation and configuration

In the e-mail archiving and records management solution, Content Manager is used as a repository for archived e-mail. When e-mail becomes records, it is still stored in the Content Manager repository.

In this section, we describe the main steps involved in installing and configuring Content Manager, a major component in the integrated solution:

1. Install Content Manager system.
2. Implement Windows service.
3. Verify installation.

Note: We do not include all of the detailed steps of the installation in this section. We recommend using the existing product manual in conjunction with the materials we present here for a successful installation and configuration.

In our sample environment, we install Content Manager on Charger.

Installing Content Manager system

We choose a typical Content Manager system installation here.

Attention: During a typical installation, full text search is *not* configured. If you need the full text search feature, choose a custom installation.

Important: The installation directory that we used in the sample environment, *C:\IBM*, is *not* the default installation directory. We used it to keep the PATH system variable short and avoid potential problems with long PATH value.

Follow the product manual for detailed installation steps.

To help your installation process, Table 7-14 shows the input values we used during our installation. Input fields are grouped by the input window. Replace our sample input values as needed for your environment setup.

Table 7-14 Content Manager installation input values

Input window / field	Sample input value	Description
Installation destination input window		
Installation directory	C:\IBM\CM	Root directory of the Content Manager installation.
Installation type input window		
Installation type	typical	Includes the HTTPs configuration for Content Manager internal communication. If you need to configure full text search, use custom installation type.
System information input window		
host name	charger.redbook.bocaraton.ibm.com	The server's fully qualified network name.
Library Server database input window		
Library Server database name	icmnlsdb	Library Server stores metadata in a relational database. This value specifies the name of the database.

Input window / field	Sample input value	Description
Library Server administration ID	icmadmin	This is the Content Manager administrator ID. If it does not exist as a Windows user ID, it will be created during the installation and will be added to the Windows Administrator group with appropriate system rights.
Resource Manager database input window		
Resource Manager database name	rmdb	Resource Manager stores information about the saved documents (such as location on a disk) in a relational database. This value specifies the name of that database.
Resource Manager database administrator	rmadmin	This is the Resource Manager administrator ID. If it does not exist as a Windows user ID, it will be created during the installation and will be added to the Windows' Administrator group with appropriate system rights.
Resource Manager volume mounting point	C:\	This is the partition on which Resource Manager stores the archived documents. During a typical installation, the directory `staging` is used as a cache area for documents retrieved from a connected Tivoli Storage Manager system.
Resource Manager application input window		
Application server node name	Charger	Resource Manager is a J2EE application. This value specifies on which node the application is deployed and a new WebSphere Application Server is created. During a typical installation, a new server called *icmrm* is installed. The application running in this server is deployed as icmrm.

Implementing Windows service

We recommend setting the Resource Manager application, a WebSphere Application Server, as a Windows service for ease of management.

To install a WebSphere Application Server as a Windows Service, use the following command:

```
wasservice -add <Windows Service Name> -serverName <WebSphere Server>
```

For our sample environment, we use the following command to implement the Windows service:

```
wasservice -add ResourceManager -serverName icmrm
```

7.5.1 Installation summary and verification

Figure 7-3 shows our sample environment after the prerequisites and Content Manager are installed.

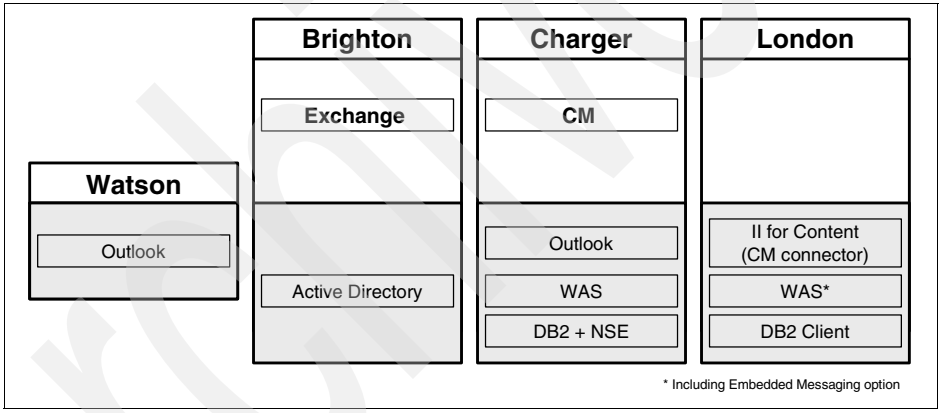

Figure 7-3 Sample environment after Content Manager is installed

At the end of the installation, an installation validation utility runs. You should see the following message, which indicates a successful installation:

```
Product validation completed with no detected configuration errors.
```

In addition to this message, perform the following steps to ensure that you installed Content Manager successfully:

1. Launch the system administration client that is automatically installed on the Content Manager server. In our scenario, we launch system administration client from Charger.
2. Log on to the Content Manager system using the administrative user ID. In our scenario, we use icmadmin.

3. Open the Resource Manager configuration.
4. If the Resource Manager configuration is available, the communication with the Resource Manager is set up properly.

Perform the following steps as the final test of a successful installation:

1. Install a Content Manager Windows client on the Content Manager server. This is the fastest way to configure the client. In our scenario, we install the client on Charger.
2. Launch the Content Manager Windows client.
3. Import a text file into the NOINDEX class of type text.

Retrieve the document immediately afterward. Make sure that you can retrieve it, open it, and view it on the screen.

7.5.2 Key information to remember

At this point, as part of your overall configuration setup documentation, write down the key configuration information that will be used during the rest of the software installation and configuration and future system operation.

Table 7-15 contains the key configuration input values to remember after the Content Manager installation.

Table 7-15 Key information to remember after Content Manager installation

Configuration data	Sample input value	Description
Content Manager administrator user ID	icmadmin	The administrative user ID for the Content Manager system.
Connection user ID	icmconct	A connection user ID that is used for clients to connect to Library Server database if they do not have a valid database user ID.
Content Manager server host name	charger.redbook.bocaraton.ibm.com	Fully qualified host name of the server that is running the Content Manager Library Server.
Library Server database	icmnlsdb	The name of the Library Server database. Every remote client has to catalog this database before being able to access it.

7.6 CommonStore (CSX) installation and configuration

In the e-mail archiving and records management solution, CommonStore is used to archive e-mail from mail databases to the Content Manager repository. It also provides a user interface to declare e-mail as records if manual records declaration and classification is allowed.

The steps involved include:

1. Install CommonStore for Exchange Server.
2. Configure Content Manager:
   a. Create the appropriate attributes and item type.
   b. Create the appropriate Content Manager user ID.
3. Configure the ArchPro environment:
   a. Set Content Manager connector environment.
   b. Create archint.ini.
4. Start ArchPro:
   a. Submit license.
   b. Save password for Content Manager user ID.
   c. Start ArchPro.
5. Prepare the task environment:
   a. Create Windows/Exchange user ID.
   a. Configure Outlook on the CommonStore machine.
   b. Install Active Directory Extension.
   c. Add forms and public folders to the Exchange system.
6. Configure the CommonStore Task:
   a. Set general properties.
   b. Define content type mappings.
   c. Configure search server.
   d. Configure task.
   e. Configure archive.
   f. Assign policy to user.
7. Start CSX Task.
8. Implement Windows Services.
9. Installing the CommonStore Outlook Extension.

Features such as full text search and single instance store are not covered. Refer to *IBM DB2 CommonStore for Exchange Server: Administration and User's Guide Version 8.3*, SH12-6741, to set up those features.

Note: We do not include detailed steps of the installation in this section. We recommend using the existing product manual in conjunction with the materials we present here for a successful installation and configuration.

In the sample environment, we install CommonStore on Charger.

Step 1: Installing CommonStore for Exchange Server

Select a complete installation. This installs all server components (ArchPro Server, CSX Task, Search Server), the CSX System Manager, and the tools required for setup (Directory Extension Tool, Forms and Public Folders) on the machine.

For production, we recommend installing and using the CSX System Manager from a different machine because both the CSX Task and the CSX System Manager use a MAPI session. (Using two MAPI sessions on one machine simultaneously can result in negative effects.) For the first configuration of the system, however, the CSX System Manager can be used on the CommonStore machine. In the sample environment, the CSX System Manager is installed on Charger (CommonStore machine) and also on Watson (client machine) for later configurations.

Table 7-16 shows the input value we used during installation in our sample environment. Replace our sample value according to your environment setup.

Table 7-16 CommonStore installation input for the sample environment

Required input field	Sample input value	Description
Installation directory	C:\IBM\CSX	The installation directory in this sample environment is different from the standard installation directory. The only reason for that is to have an easier way to manage Path and Classpath variable.

Step 2: Configuring Content Manager

Log on to the Content Manager System Administrator using your Content Manager Administrator ID. In the sample environment, the Administrator user ID is icmadmin (created during Content Manager installation; see Table 7-15 on page 225).

Perform the following steps when you are in the system administration client:

1. Create the appropriate attributes and item type (CSXMail).
2. Create the appropriate Content Manager user ID (CSX).

Creating the appropriate attributes and item type (CSXMail)

Attributes are used to store metadata. In an e-mail environment, the Exchange fields (such as Subject or Sender) are mapped through CommonStore to Content Manager attributes.

To create the new attributes:

1. Select **Data Modeling** → **Attributes**.
2. Right-click **New**. Enter the appropriate information, and click **Save**.

   The attributes CSORIGINATOR, CSUNIQUEID, CSCDISIS, CSCRISIS and BCC are mandatory and must be created.

   Other attributes can be created in order to map additional Exchange fields to Content Manager attributes.

   Table 7-17 lists the attributes used in the sample environment.

Table 7-17 Content Manager attributes used in the sample environment

Attribute	Required	Attribute type	Character type	Character length
CSORIGINATOR	YES	Var. char.	Other	1 ... 256
CSUNIQUEID	YES	Character	Alphanumeric	32
CSCDISIS	YES	Character	Alphanumeric	32
CSCRISIS	YES	Character	Alphanumeric	32
BCC	YES	Var. char.	Other	0 ... max. possible
SUBJECT	NO	Var. char.	Other	0 ... 254
SENDER	NO	Var. char.	Other	0 ... 100
TO	NO	Var. char.	Other	0 ... max. possible
CC	NO	Var. char.	Other	0 .... max. possible
DATE_TIME_C	NO	Time stamp	N/A	N/A
DATE_TIME_M	NO	Time stamp	N/A	N/A

Note: The CSORIGINATOR field holds the name of the mailbox or public folder where the message was archived from. For example, the CSORIGINATOR for a mailbox might be:

```
/O=REDBOOK/OU=FIRST ADMINISTRATIVE GROUP/CN=RECIPIENTS/CN=EMILYW
```

Important: The CSORIGINATOR for a public folder might be longer than 256 characters. If you plan to archive content in existing public folders, you might need to create a Content Manager attribute longer than 256 characters for CSORIGINATOR. An example CSORIGINATOR for a public folder could be:

```
PUBLICFOLDER:\FOUR-HUNDRED\NINETY-ONE\NINETY-TWO\NINETY-THREE\NINETY-FOUR
\NINETY-FIVE\NINETY-SIX\NINETY-SEVEN\NINETY-EIGHT\NINETY-NINE\ONE-HUNDRED
```

While CSCDISIS identifies the message archived (Document Identifier) along with the attributes common for all instances, CSCRISIS (Record Identifier) identifies the specific instance of that message. This instance holds the attributes CSORIGINATOR, CSUNIQUEID, and BCC.

To create the new item type, `CSXMail`:

1. Select **Data Modeling** → **Item types**.
2. Right-click **New**.
3. Enter `CSXMail` as the name of the item type.
4. Click the **Attributes** tab, select the following attributes, and assign them to the item type:

   ```
   CSCDISIS
   SUBJECT
   SENDER
   TO
   CC
   DATE_TIME_C
   DATE_TIME_M
   ```

5. For single-instance store, create a child component, `CSXMailChild`, and add the following attributes to this child component:

   ```
   CSCRISIS
   CSORIGINATOR
   CSUNIQUEID
   BCC
   ```

   Table 7-18 on page 230 shows the input values we used to create the CSXMail item type in the sample environment.

Table 7-18 Item type CSXMail, created in the sample environment

Configuration tab / field	Sample input value	Description
Definition		
Name	CSXMail	You can enter any name for the item type. The CommonStore configuration file (archint.ini) refers to this name.
Text search	unchecked	Text search is not configured during this installation. To enable full text search, make sure to follow the instructions in the CommonStore document *Text Search Configuration for IBM DB2 Content Manager V8*, sections "Creating a text-searchable MIME type" and "Creating a text-searchable item type."
Access Control List		
Item Type access control list	PublicReadACL	To make the integration with Records Manager fully functional, Public Read Access is necessary.
Check ACL at	on Item level	An access control list is applied to every document inserted to this item type.
Attributes		
Attributes	CSCDISIS SUBJECT SENDER TO CC DATE_TIME_C DATE_TIME_M	In a single instance store environment, two messages will be regarded as being the same e-mail if the value of the CSCDISIS attribute is identical. This will be true if both messages have the same message ID, the messages were not modified, and they have been archived with the same archiving type.
Child component	CSXMailChild	One Content Manager item can have multiple children. In an single instance store environment, child components are used to store those attributes that might be different for every e-mail, even though the content of the e-mail might be the same. **Note:** The name of a child component must be unique within a Content Manager system. This name will be used to configure the CommonStore ArchPro archint.ini file.

Configuration tab / field	Sample input value	Description
Attributes for child component	CSCRISIS CSORIGINATOR CSUNIQUEID BCC	These attributes (or mapped Exchange fields) can be different per e-mail, even though the rest of the e-mail is the same. The attribute CSCRISIS identifies the record and is unique. The CSORIGINATOR identifies the mailbox the message was in. The BCC field will contain a value only if an e-mail with a BCC recipient was archived from the Sent Items folder of the Sender.
Document Management		
Document part	ICMBASE	Only this document part is needed as no text search is configured. In case of text search, the ICMBASETEXT would be necessary.

Creating the Content Manager user ID (CSX)

CommonStore uses a Content Manager ID to communicate with the Content Manager system. This user ID needs access to the e-mail stored in the CSXMail item type.

To create a new Content Manager user ID:

1. Select **Authentication** → **Users**.
2. Right-click **New**.
3. Enter `CSX` as the user name and appropriate values. Click **Save**.

   Table 7-19 shows the input values we used to create the Content Manager user ID in the sample environment.

Table 7-19 Content Manager user ID CSX, created in the sample environment

Configuration field	Sample input value	Description
Define Users		
Name	CSX	You can enter any name for the user. The CommonStore configuration file (archint.ini) refers to this name.
Password		During startup of the CommonStore Server (ArchPro), this password has to be provided so that the ArchPro can use the ID to log on to the Content Manager system.

Password expiration	Never expires	To ensure that the CommonStore will start up correctly, make sure that the password never expires. If the password can expire, CommonStore startup will fail if the Content Manager requests a password change.
Maximum privilege set	AllPrivs	The Content Manager used ID used by CommonStore has to be a Super User, which is necessary for the Records solution. Therefore, it is necessary to assign the AllPrivs privilege set.
Set Defaults		
Default item access control list	PublicReadACL	It is necessary to provide an ACL in this field; otherwise, the user ID cannot be created. However, this ACL is used only if, during an item type creation, the field `User's default ACL` is chosen to be the ACL that defines the ACL to be set for items stored in the item type.

Tip: To make sure that the Content Manager user ID used by CommonStore can access the Content Manager system, install a Content Manager Windows client on the CommonStore Server. Use this client to log on to the Content Manager using the ID (CSX) created to be used by CommonStore. Try to import a document (for example, a text file) to the newly created item type.

Step 3: Configuring ArchPro environment

To configure ArchPro environment:

1. Set the Content Manager connector environment.
2. Create an archint.ini file.

Setting the Content Manager connector environment

To apply the correct environment settings for the Content Manager V8 connector, run a batch program called Agentenv_CM8.bat. This program is delivered with CommonStore and it can be found in the bin directory of the CommonStore installation directory.

Run Agentenv_CM8.bat from a Windows command prompt as follows:

1. Open a Windows command prompt.

2. Change to the bin directory.

   In our sample environment, it is C:\IBM\CSX\bin, where C:\IBM\CSX is the chosen installation directory (see “Step 1: Installing CommonStore for Exchange Server” on page 227).

3. Run Agentenv_CM8.bat.

Creating archint.ini file

The archint.ini file configures the ArchPro Server. It defines whether logging and tracing are activated, and defines logical archives, which point to a specific Content Manager server including the item type that is used.

To create a server configuration profile, archint.ini, follow these steps:

1. Open the sample profile archint_sample_cm8.ini in an editor. This file resides in the instance directory of the CommonStore installation directory:

   ```
   Instance Directory: C:\IBM\CSX\Server\instance01
   ```

2. Save the file as `archint.ini` in the same directory.

Important: Make sure the file is saved as `archint.ini` and not as archint.ini.*txt*.

3. Use the search function of your editor to locate the following section:

   ```
   ARCHIVE             CSXSAMPLEARCHIVE
      STORAGETYPE         CM
      LIBSERVER           sampleLibServer
      ITEM_TYPE           sampleItemType
      CMUSER              sampleUser
      ARCHIVETYPE         GENERIC_MULTIPART
   ```

   Following the sample environment, configure the logical archive as listed in Table 7-20 on page 234.

Table 7-20 Input values for archint.ini file in the sample environment

Parameter	Sample input value	Description
ARCHIVE	CSXMAIL	You can enter any name for the logical archive that will be used by CommonStore Task. A Task is not aware of a Library Server or an item type but only of the logical archive name. ArchPro defines logical archives in order to make it transparent for a Task on which archive system (Content Manager, Content Manager OnDemand, Tivoli Storage Manager) is used. The Task only refers to the logical archive or archives.
STORAGETYPE	CM	ArchPro supports various archive systems in addition to Content Manager. The value CM used here specifies that the archive is a Content Manager system.
LIBSERVER	icmnlsdb	The name of the Content Manager Library Server in which the Item type is created. If the Content Manager is running on a different system than the CommonStore Server, this is the name under which the remote database is cataloged on the CommonStore system. In the sample environment, the Content Manager is installed on the same system and therefore it is not necessary to catalog the database. The name of the database is configured during the Content Manager installation; see 7.5.2, “Key information to remember” on page 225.
ITEM_TYPE	CSXMail	The item type used to archive e-mail; created in step 2b of 7.6, “CommonStore (CSX) installation and configuration” on page 226.
CMUSER	CSX	Content Manager user ID that is used by CommonStore to communicate with the archive system. This ID was created in step 2b of 7.6, “CommonStore (CSX) installation and configuration” on page 226.
ARCHIVETYPE	GENERIC_MULTIDOC	This specifies how documents are stored within the Content Manager. This is relevant if Component Archiving is selected for e-mail archiving. In this case, the e-mail gets separated from its attachment. Every component (e-mail and every attachment) is stored as a separate document within Content Manager if GENERIC_MULTIDOC is defined.

Parameter	Sample input value	Description
SISCHILDNAME	CSXMailChild	To enable single-instance store functionality in your archive, add this line. This value has to be exactly the same as the one defined in “Step 2: Configuring Content Manager” on page 227.

After changing the values, the logical archive section of the archint.ini file should look like this:

```
ARCHIVE              CSXMAIL
   STORAGETYPE          CM
   LIBSERVER            icmnlsdb
   ITEM_TYPE            CSXMail
   CMUSER               CSX
   ARCHIVETYPE          GENERIC_MULTIDOC
   SISCHILDNAME         CSXMailChild
```

4. Save the changes.

Step 4: Starting ArchPro

To start ArchPro, you need to:

1. Submit license.
2. Save password for Content Manager user ID.
3. Start ArchPro.

Submitting license

To submit license, perform the following steps:

1. Open a Windows command prompt.
2. Change to the instance01 subdirectory of the CSX installation path:

   Instance Directory: `C:\IBM\CSX\server\instance01`

3. Enroll a CommonStore license by entering the following command:

   ```
   archpro -f license
   ```

4. The location of the license file is requested. Provide the full path including the file name:

   ```
   C:\IBM\CSX\licensekey\csx8.lic
   ```

Important: If you skip this step, a Try and Buy licence will be installed and it will expire after 90 days. After the 90-day period, the CommonStore Server will not start.

Saving a password for Content Manager user ID

To save a password for the Content Manager user ID:

1. Open a Windows command prompt (if not already open).
2. Change to the instance01 subdirectory of the CSX installation path.

   Instance Directory: `C:\IBM\CSX\server\instance01`
3. Set the password for the CSX item type (CSXMail), created in step 2a of 7.6, “CommonStore (CSX) installation and configuration” on page 226, by entering the following command:

   ```
   archpro -f serverpasswd
   ```

   The password for the Content Manager user ID used by CommonStore (which is defined in the archint.ini and created in “Creating the Content Manager user ID (CSX)” on page 231) is requested.
4. Enter the password and press **Enter**.

Important: If this command is not issued from the instance directory, it is necessary to point to the archint.ini file to be used. To do so, include the -i parameter as follows:

```
archpro -f serverpasswd -i <path to archint.ini file>
```

If the archpro command is issued without the -i parameter, then ArchPro searches in the starting directory for the archint.ini file. If no file can be found, the startup will fail.

Important: If the password of the Content Manager user ID used by CommonStore is changed, it is necessary to run the **`archpro -f serverpasswd`** command again to provide the new password to ArchPro.

Starting ArchPro

To start ArchPro, perform the following steps:

1. Open a Windows command prompt (if not already open).
2. Change to the instance01 subdirectory of the CSX installation path.

   Instance Directory: `C:\IBM\CSX\server\instance01`
3. Start the CommonStore Server by entering the **`archpro`** command.

Example 7-1 on page 237 shows the messages displayed during archpro startup.

Example 7-1 ArchPro startup messages

```
C:\IBM\CSX\server\instance01>archpro
         ****************************************************************
         * IBM DB2 CommonStore - Server 8.3.0.0                         *
         * (c) Copyright IBM Corporation, 1997, 2004 All Rights Reserved. *
         * Build 8.3.0.0,    Compiled at Mar  9 2005.                   *
         ****************************************************************

CSS0030I: ArchPro is using INI file 'C:\IBM\CSX\server\instance01\archint.ini'.
CSS0910I: Trying to get a LUM Production License for IBM DB2 CommonStore for
Exchange Server
CSS0929I:
         ****************************************************************
         *                Got a Production License for                  *
         *           IBM DB2 CommonStore for Exchange Server            *
         ****************************************************************
CSS0158I: ArchPro 3464 started on UNICODE Port 4336.
CSS0157I: ArchPro 3464 is waiting for external connections on fixed port 8013.
CSS0100I: IBM Content Manager CommonStore HTTP Task 8.3.0.0
CSS0333I: ArchPro is informed that Web dispatcher 'HTTP_TASK_1' has started and
is ready (socket 1840).
CSA0100I: IBM Content Manager CommonStore Agent for Content Manager 8.3.0.0
CSS0010I: HTTP WORKER #0: has been initialized and started
CSS0001I: HTTP_TASK_1: successfully started, ready to process jobs
CSS0330I: ArchPro is informed that CM agent 'CM-AGENT_1' has started (socket
1804).
CSS0010I: HTTP LISTENER : has been initialized and started
CSS0103I: HTTP LISTENER : listens for HTTP requests on port 8085
ADMU0116I: Tool information is being logged in file C:\IBM\CSX\Search
Server\logs\server1\startServer.log
ADMU3100I: Reading configuration for server: server1
CSA0300I: Connection for repository 'CSXMAIL' on server 'charger' with item typ
e 'CSXMail' for user 'icmadmin' OK
CSA0001I: CM-AGENT_1: successfully started, ready to process jobs
ADMU3200I: Server launched. Waiting for initialization status.CSS0325I: ArchPro
is informed that  Agent 3480 is ready to obtain order.
CSS0166I: ArchPro is fully initialized. Queue processing is enabled now.
CSA0010I: CM-AGENT WORKER #0: has been initialized and started
CSA0010I: CM-AGENT WORKER #1: has been initialized and started

ADMU3000I: Server server1 open for e-business; process id is 3848
```

The message prefix (the first three characters of the message) identifies the component generating the message. The following mappings apply:

CSS	CommonStore Server (ArchPro)
CSA	CM8 agent
ADM	Search server

The CommonStore Server has completed initialization when the following message is displayed:

```
Archpro is fully initialized. Queue processing is enabled now.
```

The search server is ready for processing when this message is displayed:

```
Server server1 open for e-Business;
```

Step 5: Preparing the task environment

To prepare for the task environment:

1. Create a Windows/Exchange user ID (CSX Admin).
2. Configure Outlook on the CommonStore machine.
3. Install Active Directory Extension.
4. Add forms and public folders to the Exchange system.

Creating a Windows/Exchange user ID (CSX Admin)

We recommend creating a user ID that is specifically used for CSX. In the context of this book, the user *CSX Admin* with ID *csxadmin* is used.

To create a Windows/Exchange user ID:

1. Log on to the Exchange server using an administrative user ID.
2. Select **Start** → **Settings** → **Control Panel** → **Administrative Tools** → **Active Directory Users and Computers**.
3. Right-click **Users** and select **New** → **User** from the context menu.
4. Create the user ID as shown in Table 7-21.

Table 7-21 Input value for user CSX Admin

Configuration field	Sample input value	Description
First Screen		
First Name	CSX	
Last Name	Admin	
User logon name	csxadmin	This user ID will be used by the CommonStore Task to log on to the Exchange server. This ID needs the appropriate access rights in order to access the e-mail stores.

Second Screen		
Password		This password has to be provided during the CommonStore Task startup so that the Task can use the ID to log on to the Exchange system.
Password never expires	checked	To ensure that the CommonStore Task (CSX) will start up correctly, make sure that the password never expires. If the password can expire, the CommonStore Task start up will fail if the password change is requested.
Third Screen		
Create an Exchange mailbox	checked	
Alias	csxadmin	The fields `User logon name` and `Alias` must have the same values. Normally, the user logon name is taken as the alias by default.
Server	Redbook/ First Administrative Group /BRIGHTON	In the Server field, make sure to select an Exchange server that will be serviced by the CSX Task.

Add the user ID csxadmin to the following groups:

- Domain Admins
- Schema Admins

To add csxadmin to a group:

1. Select **Start** → **Settings** → **Control Panel** → **Administrative Tools** → **Active Directory Users and Computers** → **Users**.
2. Right-click **CSX Admin** and select **Properties** → **Member Of** → **Add**.

   We recommend also creating a test user to verify your setup. Create the user `Test User` with ID `testuser` with a mailbox on the same Exchange server.

Configure Outlook on the CommonStore machine

Refer to the Outlook manual for the installation and configuration.

When installing Outlook, be sure to select to have the Collaborative Data Objects (CDO) be installed.

Install the Active Directory Extension

The CommonStore Task stores the configuration data within the Active Directory. Before this can be done, the logical structure to store this information must be defined within the Active Directory. The installation of the CommonStore Active Directory Extension creates this logical structure.

Attention: On Windows 2000 or Windows XP, make sure that schema modifications are allowed on the domain controller:

1. Select **Start** → **Run**, and enter `MMC`.
2. Select **File** → **Add/Remove Snap-in**.
3. Click **Add**.
4. Select **Active Directory-Schema** and click **Add**.
5. Click **Close** and then **OK**.
6. Select **Active Directory Schema** → **Operations Master**.
7. Check **The schema may be modified on this domain controller**.

Although this step can be performed on another workstation, we recommend doing this on the domain controller with the user ID csxadmin. In the sample environment, Brighton is the domain controller, so we performed this installation on Brighton. The CommonStore Directory Extension Tool (included in the CommonStore installation) is installed on Brighton.

To install Active Directory Extension:

1. Start a Windows command prompt.
2. Run the following scripts:
   - CsxImportSchema.vbs

     This extends the Active Directory schema with CSX configuration objects.
   - CsxImportEmptyDirectory.vbs

     This creates a node in the Active Directory to hold CSX configuration data.
   - CsxActiveDirectoryExtensionVerification.vbs

     This verifies that the Active Directory Extension has been configured successfully.

For detailed step-by-step instructions, refer to the CSX book section “Extending the Active Directory schema.”

If the Active Directory Schema Extension installation is successfully finished, the configuration data of the CommonStore Task can now be stored in the Active Directory using the CommonStore System Manager. The CommonStore Task

configuration is done in “Step 6: Configuring the CommonStore Task” on page 244.

Attention: On Windows 2000 or Windows XP, make sure that schema modifications are *not* allowed on the domain controller:

1. Select **Start** → **Run**, and enter `MMC`.
2. Select **File** → **Add/Remove Snap-in**.
3. Click **Add**.
4. Select **Active Directory-Schema** and click **Add**.
5. Click **Close** and then **OK**.
6. Select **Active Directory Schema** → **Operations Master**.
7. Uncheck **The schema may be modified on this domain controller**.

Adding forms and public folders to the Exchange system

In an Exchange environment, forms are needed to display objects. For example, the form IPM.Note is used to display e-mail. The CommonStore Task uses Job documents to communicate with the CommonStore Outlook Extension. The job documents are stored in a public folder and also require a form to be displayed.

To add the necessary forms:

1. Create the Organizational Forms Library.

   CSX form templates are registered in the Organizational Forms Library. This is a public folder. If there is no Organizational Forms Library on the Exchange server, create one by following the instructions in the CSX book section:

   Configuring the Exchange Organization → **Setting up an Organizational Forms Library** → **Creating an Organizational Forms Library for CSX 2000/2003**.

   Note: If the entry EFORMS REGISTRY is not visible, make sure that you have installed the latest Exchange Server fix pack.

2. Add forms and public folders to your Exchange system.

   To add the CSX forms to your Organizational Forms Library and create public folders required by CSX:

   a. Open a command prompt.

   b. Execute the command **`InstallCSXFormsAndFolders.bat`** using the profile of the CSX Task and administration user.

Note: If you see the error message `ActiveX object could not be created` is displayed, make sure that the Collaborative Data Objects (CDO) are installed. This component is shipped with Outlook, but not installed by default.You can call the Outlook setup program from the Add/Remove Programs window to add this feature.

After successful completion, the following folders are available:

- ...\CommonStore\Configuration

  This folder contains configuration information. It is updated by the CSX System Administration program and read by all CSX clients.

- ...\CommonStore\Job Folders

  Individual subfolders of this folder contain CSX client requests for interactive archiving and retrieval. All CSX clients must have write access to these folders.

- `...\CommonStoreTransferFolder`

  This folder is used only for PST file archiving.

3. Grant user access to the public folders.

   To grant access to the public folders created during the installation of the form templates, use the Exchange System Manager. As each CommonStore installation is different, this section can give only some recommendations. The overall aim is to prevent ordinary users from damaging the existing configuration. See the following list:

   - Only the administrator should be the owner of the public folders. Other users should have limited rights.
   - User IDs that are authorized to create and modify Task Administration Data objects in the CSX System Manager must also have owner rights for the public folder CommonStore\Configuration. Other users should have only read access to this folder. We recommend that you create special user IDs just for administration tasks.
   - As the administrator, grant other users only reviewer rights to the public folder CommonStore so that these users cannot create or delete folders and items.
   - Grant read access to the public folder Configuration, which is a subfolder of the CommonStore folder. The CSX Client Components must be able to read the content of the mapping items therein, which is why the users need read access. Not extending this right any further prevents users from changing or deleting items in this folder.

– For the same reason, grant users only read access to the public folder named Job Folders.

– Users must be able to write to the subfolders of Job Folders because this is part of the job creation process. The subfolders are created at a later point in time, when you configure CSX Task instances in the CSX System Manager. The subfolders are actually the job folders for the interactive job documents of each individual CSX Task instance. It is possible that users need the right to delete items in these subfolders. However, you should not allow them to delete items in job folders belonging to other task instances. Restrict the write access of users to the job folder that belongs to the CSX Task instance assigned to their mail boxes.

– Make sure that the user who starts a particular CSX Task instance has owner rights for all subfolders in the Job Folders folder of this instance.

– Users need access to the transfer folder. Change the permissions for the new public folder CommonStoreTransferFolder by following these steps:

  i. Open Outlook with the user ID that you used to create the CommonStore folders. In the folder list, select **Public Folders** → **All Public Folders** → **CommonStoreTransferFolder**.

  ii. Right-click that folder and select **Properties**.

  iii. In the CommonStoreTransferFolder Properties notebook, select the **Permissions** tab.

  iv. Select the name **Default**. The following settings must be applied:

```
Create items - not set
Read items - not set
Create subfolders - set
Folder owner - not set
Folder visible - set
```

  v. Select **Own** in the Edit items section, select **Own** in the Delete items section, select the name **Anonymous**, and select **None** from the Roles drop-down list in the Permissions section.

  vi. Click **OK** to close the CommonStoreTransferFolder properties notebook.

Step 6: Configuring the CommonStore Task

To configure the CommonStore Task, log on to the CSX System Manager using the administrative ID created in "Step 5: Preparing the task environment" on page 238. Go to **Start** → **Programs** → **IBM DB2 CommonStore Exchange Server** → **System Manager**, and complete the following steps:

1. Set general properties.
2. Define content type mapping.
3. Configure search server.
4. Configure task.
5. Configure archive.
6. Define archive policy.
7. Assign policy to users.

Setting general properties

The general properties are used to specify how the CommonStore Task is working if no other policies are assigned to a user or no content type mapping is found for a certain document.

CommonStore uses content type mappings to map extensions of archived attachments to a Content Manager content type (for example, the doc extension is mapped to application/msword). This content type information is stored in the archive. If the content type is incorrect, the attachment may not be displayed correctly. It is therefore very important to define the correct content types.

If no content type is associated to an extension, the default setting for content type is used. We recommend leaving the Content type field empty until you get used to the behavior of CommonStore and you are sure that all file extensions used in your organization are mapped to a content type.

If an attachment is archived and no content type is mapped for the extension, the default setting is used. If no default setting exists for that extension, CSX generates a warning message, but archives the attachment anyway.

To allow the browser to display those files, add the following lines to the file C:\IBM\CSX\Server\instance01\csmimes.properties:

```
=application/unknown
null=application/unknown
```

Note: If you use a default setting for the content type, make sure that this type follows the MIME type syntax (for example, application/UNKNOWN). Using this approach, some browsers may still be able to display the content correctly.

Use the values in Table 7-22 on page 245 as guidance to set the properties for the CommonStore Task.

Table 7-22 CommonStore Task properties setting

Configuration field	Sample input value	Description
General		
Archiving policy	<None>	Specifies which archiving policy is used by default if no archiving policy is assigned to a user. To avoid unexpected side effect sat the beginning, the parameter is set to none.
Mailbox threshold		Specifies the size of a mailbox in megabytes. Only mailboxes bigger than this threshold will be archived.
Content type		If no content type mapping exists for an extension and this parameter is not set, a warning message is generated when archiving.
Tasks refresh every	600	Specifies the time in seconds the Task waits before it reads the configuration from the Active Directory again, after it is started.
Restubbing Schedule		
Cycle	Never	Specifies how often the default restubbing process is used. To avoid unexpected side effects, the value is set to never, so that the general restubbing will never take place.

Defining content type mappings

As mentioned earlier, CommonStore uses content type mappings to map extensions of archived attachments to a Content Manager MIME type (formerly content types), for example, the extension *doc* is mapped to application/msword. This MIME type information is stored in the archive. If the MIME type is incorrect, the attachment may not be displayed correctly. It is important to define correct content types mappings.

If no content type mapping is specified for a document with a certain extension and no default mapping is defined in the General Properties, the archiving of this document will fail.

To add a new content type mapping:

1. Right-click **Content Type Mapping** → **Add new content type mapping**.

2. For every content type mapping, a file extension and a Content Manager content type (MIME type in Content Manager system administration client) has to be provided.

   Create the content type mappings as specified in Table 7-23.

Table 7-23 Content type mapping

File extension	Content type	Description
msg	application/outlook	Entire messages are archived using the MSG-format. A mapping for this extension is mandatory.
doc	application/msword	The doc extension is used as an example for many other possible extensions.

Attention: Make sure the provided content type is spelled *exactly* the same as the MIME type configuration within Content Manager.

To check the correct MIME type:

1. Log on to the Content Manager system administration client.
2. Go to **Data Modeling** → **MIME Types**.
3. Double-click the MIME Type to be checked.

The value MIME Type has to be the same as the one used for the content type value in the CommonStore configuration.

There is no MIME type for the MSG format within the Content Manager system. The appropriate MIME type has to be created.

To create the correct MIME Type in the Content Manager system:

1. Log on to the Content Manager system administration client.
2. Go to **Data Modeling**.
3. Right-click **MIME Types** and select **New** from the context menu.
4. The value for the MIME type can be chosen. It has to follow the syntax of *value/value,* for example: application/outlook.

Configuring search server

To enable Outlook users to search for archived messages from their Outlook clients, you must specify certain configuration parameters. These parameters are stored in search configuration objects in the CSX System Manager so that they can be reused.

On the Privileged User page, special users can be granted the right to search for messages that were archived by other users.

Table 7-24 lists the values we used to configure the search server. Substitute the sample input values with values that are specific to your environment.

Table 7-24 Search server configuration

Configuration field	Sample input value	Description
General		
Configuration name	SearchCM	This is the name for the search configuration profile. You can use any name you want. This name will be selected during the Task Administration Data setup.
Host name	charger	The host name of the server that runs the CommonStore Server.
Port number	7900	Specifies the port the search server is using.
Privileged User		
Users	CSX Admin	List of users with access to all archived documents using the CommonStore search.

Configuring task

Table 7-25 on page 248 shows the values we used to configure the task. Replace our values with the ones suitable for your environment.

Important: The value entered in the Fixed port field must be identical to the ARCHPRO_PORT value in the CommonStore Server configuration file (archint.ini).

While socket communication is used to handle all requests, the file transfer from the CSX Task to the CommonStore Server and vice versa is done using a shared directory. The Transfer path field must address the same directory as the TRANSFERPATH value in the CommonStore Server configuration file (archint.ini).

Table 7-25 Task configuration

Configuration field	Sample input value	Description
General		
Task name	task_brighton	This is the name for the CommonStore Task. You can use any name for this. This name will be used during the Task startup.
Administrator notification	CSX Admin	A list of user IDs that are notified in case of an error.
Parameters		
Worker count	3	Number of Worker threads. This number should be equal to the number of Committer threads. A higher number of threads results in a higher throughput. However, an unlimited increase of this number will not result in better performance, since the ArchPro server will be the bottleneck. Using three threads per instance is a well-known configuration.
Committer count	3	Number of Committer threads. This number should be equal to the number of Worker threads. See description above.
External port	7000	Specifies the port the CSX Task is using for communication.
Trace file name	C:\IBM\CSX\Task\csx_task_brighton.trc	
Error file name	C:\IBM\CSX\Task\csx_task_brighton.err	
Log directory	C:\IBM\CSX\Task\log	
Search configuration	SearchCM	Specifies the search configuration profile defined in “Configuring search server” on page 246.
Records Enabler configuration	<None>	This value will be set after the Records Manager installation.

Configuration field	Sample input value	Description
Trace	Selected	For testing purpose, tracing should be activated. As soon as the system is fully configured and running without any errors, this should be deselected.
CommonStore Server		
Host name	charger	Host name of the CommonStore Server (ArchPro) machine.
Fixed port	8013	Specifies the port on which the CommonStore Server (ArchPro) is listening for tasks. This port is specified in the archint.ini file with the parameter ARCHPRO_PORT.
Transfer path	C:\IBM\CSX\ Server\instanc e01\transfer	Specifies the directory where the Task and the ArchPro exchange their files. The ArchPro reads this information from the archint.ini. The entry in the archint.ini and this entry have to be the same. TRANSFERPATH is the parameter in the archint.ini.
Exchange Server		
Job folder name	task_brighton	A public folder that has to be created (using the New button within the CSX System Manager). This folder is used to store job documents during manual archiving and declaration.
Exchange server (or servers) serviced by this task	BRIGHTON. redbook. bocaraton.ibm .com	Name of the exchange server that will be serviced by this task. One task could handle more than one Exchange server.

Configuring archive

The archive refers to the logical archive defined in the archint.ini. (See “Step 3: Configuring ArchPro environment” on page 232.) This definition is used during the archiving policy definition.

For every archive, an attribute mapping has to be defined. An archive mapping specifies which Outlook/Exchange property is mapped to which Content Manager attribute. None of these mappings are mandatory but can be chosen. However, we recommend creating at least some standard mappings in order to be able to provide an attribute-based search (such as “list all e-mail with ‘CommonStore’ in the subject”) for the e-mail user.

Table 7-26 shows the values we used to configure the archive. Replace sample input values to suit your environment and business needs.

Table 7-26 Archive configuration

Configuration field	Sample input value	Description
General		
Archive id	CSXMAIL	Name of the logical archive defined in the archint.ini. See “Step 3: Configuring ArchPro environment” on page 232. This value is case sensitive.
Archive type	Content Manager 8	Specifies the archiving type. CommonStore supports three archive systems: Content Manager, Content Manager OnDemand, and Tivoli Storage Manager.
Property Mapping		
Mapping list		Contains the property / attribute mapping.

`Subject -> SUBJECT`
`Sender -> SENDER`
`To -> TO`
`Cc -> CC`
`Creation Time -> DATE_TIME_C`
`Last Modification Time -> DATE_TIME_M`
		The left side contains the Outlook/Exchange properties. The right side contains the Content Manager attributes as defined in “Step 2: Configuring Content Manager” on page 227.

Note: The pseudo property _CSX_MESSAGE_LOCATION is not a property defined in Exchange. It is computed by CSX and reflects the name of the folder the message is stored in. This property can also be mapped to an attribute in the archive. It can be used as a condition in an archiving policy.

Defining archive policy

Table 7-27 on page 251 specifies the archive policy we defined. Replace sample input values with ones that are suitable for your environment and business requirements.

Table 7-27 Archive policy sample

Configuration field	Sample input value	Description
General		
Policy Name	Interactive Archiving	You can use any name that is descriptive to the policy.
Restubbing on	checked	Specifies whether the retrieved e-mail is restubbed.
Restubbing time	24	Specifies the time in hours after the retrieve that a restubbing of an e-mail takes place.
Archiving Schedule		
Cycle	Always	The CommonStore Task applies this rule when checking for e-mail.
Automatic Rules		
In the sample environment, no automatic archiving is configured.		
Interactive Rules		
Message Class	IPM.Note	Rules are based on Message classes. Must list every Message class this rule should apply to.
Archiving Type	Entire	The entire e-mail is stored in the archive. Other options are Attachment (only the Attachment is stored in the archive) and Component (attachments and body are stored as separate objects).
Deletion Type	Body	The deletion type specifies what part of an e-mail gets deleted within Exchange. Deleting the body will result in the smallest stub. However, it is possible only to delete the attachments (even though the entire message is archived) to leave more information in the stub document.
Archive	CSXMAIL	Specifies the name of the logical archive that is registered under “Configuring archive” on page 249.

Configuration field	Sample input value	Description
Use archiving type and deletion type provided by the client	checked	Specifies that the client can override the rules. For example, if the rule defines Entire/Body (Archiving type/Deletion type) archiving, the client can archive an e-mail with Entire/Attachment. If this check box is not selected, only the specified archiving method is allowed.

Assigning a policy to a user

Every user has to have a policy assigned to it. If no policy is assigned, the user cannot archive e-mail automatically or interactively. If only a policy with an automatic rule is applied, the user will not be able to archive e-mail interactively.

To assign the Interactive Archiving policy created earlier to the user Test User:

1. Go to **Policies for Mailboxes** → **Users**.
2. Right-click **Test User** and select **Properties** from the context menu.
3. Select **Archiving Policy**, and select **Interactive Archiving**.

Step 7: Start CSX Task

Log on to the CSX Task workstation with user csxadmin and complete the following steps to start the CSX Task:

1. Open a Windows command prompt.
2. Start the CSX Task by executing the command:

   ```
   csx task_brighton
   ```

   The CSX Task reads the configuration data from the Active Directory, where it is stored by the CSX System Manager. It then creates an Outlook profile and logs on to the Exchange Server. The crawler thread is responsible for automatic archiving and declaring according to the automatic rules in the mailbox policies. The crawler thread is also responsible for restubbing. The polling thread handles interactive requests.

Example 7-2 shows sample output from the CSX Task initialization to the console.

Example 7-2 CSX Task initialization sample output

```
C:\IBM\CSX\bin>csx task_brighton
CJS0000I: CSX J/Starter 8.3.0.0
[11:50:24]
```

```
          ****************************************************************
          *  IBM*                                                        *
          *  DB2 CommonStore for Exchange Server - Task                  *
          *  Version 8.3.0.0 build: 20050309                             *
          *                                                              *
          *  Licensed Materials - Property of IBM                        *
          *  5724-B85                                                    *
          *  (c) Copyright IBM Corporation 1997, 2005.                   *
          *  All Rights Reserved.                                        *
          *                                                              *
          * * Trademark of International Business Machines               *
          ****************************************************************
[11:50:24]  task_brighton: No property file found for task: 'task_brighton'
[11:50:24]  CSX0010I: CSX COMMITTER #2: has been initialized and started
[11:50:24]  CSX0010I: CSX COMMITTER #1: has been initialized and started
[11:50:24]  CSX0010I: CSX COMMITTER #0: has been initialized and started
[11:50:26]  CSX0010I: CSX WORKER    #0: has been initialized and started
[11:50:26]  CSX0010I: CSX WORKER    #1: has been initialized and started
[11:50:26]  CSX0010I: CSX WORKER    #2: has been initialized and started
[11:50:26]  CSX0010I: CSX POLLER    : has been initialized and started
[11:50:26]  CSX0001I: task_brighton: successfully started, ready to process jobs
[11:50:26]  CSX0010I: CSX CRAWLER   #brighton.emms.bocaraton.ibm.com: has been
initialized and started
```

The line `task_brighton: successfully started, ready to process jobs` shows that the CommonStore Task is started successfully.

Step 8: Implementing Windows Services

CommonStore provides the possibility to run all components (ArchPro, tasks) as Windows services. To implement a component as a service, it must be installed as a service using the *archservice program* that ships with CommonStore.

The archservice program needs an INI file to install a program as a service. The INI file contains the startup command (PROCESS1) for the component to be started and a location for the trace file (SERVICE_TRACEFILE) of the service.

Attention: Before installing the ArchPro and the Task as Windows services make sure that previously started components (for example, in Windows command prompts) are closed down.

Even though it is possible to have the ArchPro and the Task running as one service, we recommend installing both as separate services. In the sample environment, the INI files are stored in the newly created directory C:\IBM\CSX\WindowsService. Using a text editor, create two INI files, ArchProService.ini and TaskService.ini.

To implement a Window service:

1. Create a directory C:\IBM\CSX\WindowsService.
2. Open a text editor and create two INI files, one per component. Each file contains the following parameters:
   - SERVICE_ TRACEFILE: The path (including the file name) of the trace file that this service creates. The directory must exist already.
   - PROCESS1: The startup command of the component.

   Example 7-3 shows the ArchProService.ini file we set up for the sample environment. Example 7-4 shows the TaskService.ini we set up. Use them as references to set up files in your environment.
3. Save the files in directory C:\IBM\CSX\WindowsService.

Example 7-3 ArchProService.ini for the sample environment

```
SERVICE_TRACEFILE C:\ibm\csx\WindowsService\ArchProService.trace
PROCESS1 C:\ibm\csx\bin\archpro.exe -i C:\ibm\csx\server\instance01\archint.ini
```

Example 7-4 TaskService.ini for the sample environment

```
SERVICE_TRACEFILE C:\ibm\csx\WindowsService\TaskService.trace
PROCESS1 C:\ibm\csx\bin\csx.exe task_brighton -u "CSX Admin"
```

4. Open a Command prompt window.

   The command syntax used to install the Window services is:

   ```
   archservice install -n <name> -c <config file>
   ```

   In this syntax:
   - <name> appears in the Windows services list as part of the service name. A CommonStore Service always starts with CommonStore_ and ends with the specified value. For example, The ArchPro Window service will appear as CommonStore_ArchPro.
   - <config file> is the path (including the file name) of the configuration file to be used.
5. Execute the following command to install ArchPro service:

   ```
   archservice install -n ArchPro -c
   C:\IBM\CSX\WindowsService\ArchProService.ini
   ```
6. Execute the following command to install the Task service:

   ```
   archservice install -n Task -c C:\IBM\CSX\WindowsService\TaskService.ini
   ```

Step 9: Installing the CommonStore Outlook Extension

Install all available features from the client package on a client workstation. To do so, select installation type **Custom** and select all features. The feature “Allow use of client settings” gives you some flexibility when testing different archiving types and deletion types during setup. The feature “Records Enabler” is required to declare and view records and will be needed later on.

In the sample environment, Watson represents a client system. The CommonStore Outlook Extension is installed on that machine.

7.6.1 Installation summary and verification

Figure 7-4 shows the sample environment after CommonStore for Exchange is installed.

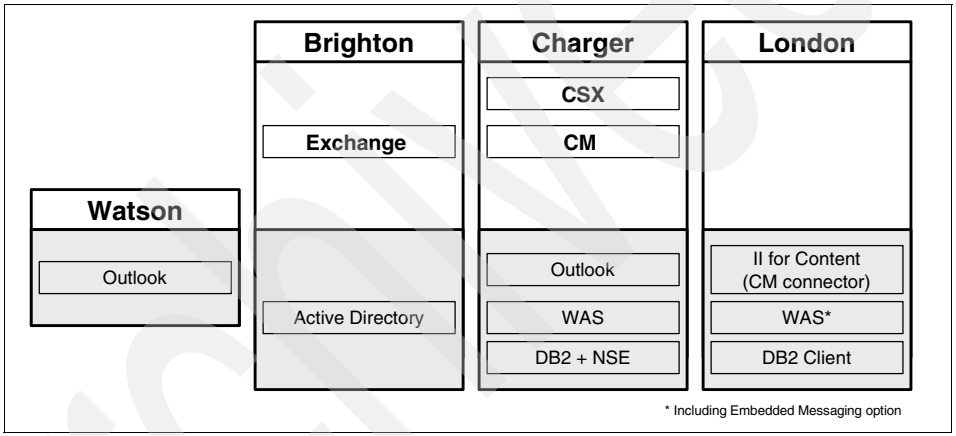

Figure 7-4 Sample environment after CSX is installed and configured

To verify the CSX Task setup, we can archive and retrieve a message and verify that it is stored in the archive correctly. This includes the following steps:

1. Start CommonStore (ArchPro and Task).
2. Start Outlook.
3. Manually archive a message.
4. View an archived message.
5. Retrieve an archived message.
6. SIS setup.
7. Search for an archived message.

Starting CommonStore (ArchPro and Task)

Make sure all CommonStore components are up and running.

> **Tip:** Running the components in Windows command prompt makes it easier to find problems during the test phase.

Either choose to start the Windows Services or run the components from Windows command prompts.

Starting Outlook

Log on to the client workstation with user `testuser` and perform the following steps:

1. Create an Outlook profile for testuser. If you are using Outlook 2003, make sure that the check box Use Cached Exchange Mode is not checked.
2. Start Outlook.

   The CSX tool bar contains the following buttons for CSX operations:

   / Use default settings / Use custom settings
   Archive
   Retrieve
   Search

Manually archiving a message

To perform some interactive archiving:

1. Switch from the default configuration mode to the custom configuration mode.

   To do so, click the toggle button (Use default settings / Use custom settings). A pencil appears in the icon: .
2. Create a test message with the subject `CSX test setup` including an attachment and some body text and send it to testuser. See Figure 7-5 on page 257.

   This creates two copies of the message in the mailbox: one in the Inbox and one in the Sent Items folder.

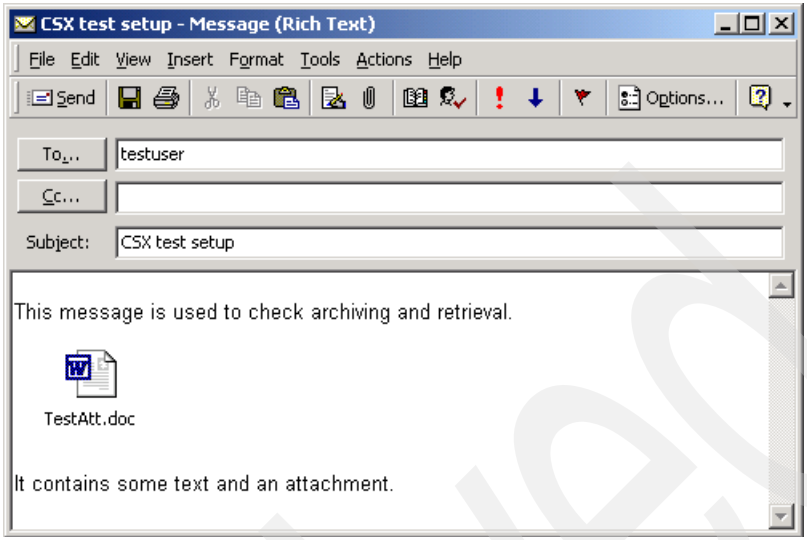

Figure 7-5 Sample test message

3. Create another test message with the subject `CSX test message 2` and also send it to testuser.
4. Select the two test messages in the Inbox and click (Archive).

   The Archiving Settings dialog displays (see Figure 7-6).

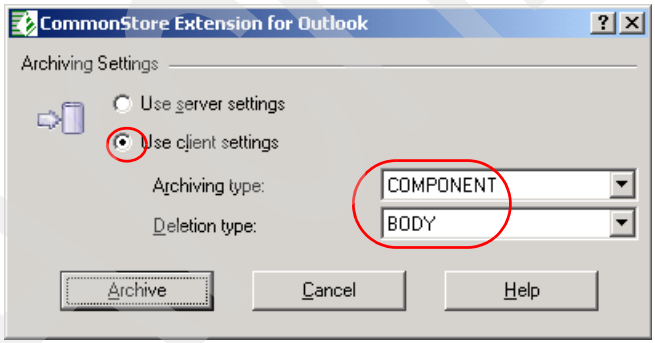

Figure 7-6 Archiving Settings dialog

5. Configure the archiving settings as follows:
   a. Select **Use client settings**.
   b. From the Archiving type pull-down menu, select **COMPONENT**.
   c. From the Deletion type pull-down menu, select **BODY**.

6. Click **Archive**.

   A new job document is created and stored in the public folder. The CSX Task polling thread reads the archive request from the job folder. The CSX Task worker thread reads the message from the mailbox, decomposes it according to the archiving type (COMPONENT), and passes two documents (the message remainder and the attachment) to the CommonStore Server.

   When the documents are archived in the Content Manager repository, the CSX Task committer thread stubs the message. It removes the body text and the attachment according to the deletion type (BODY) and inserts hyperlinks to both the message remainder and the attachment (Figure 7-7).

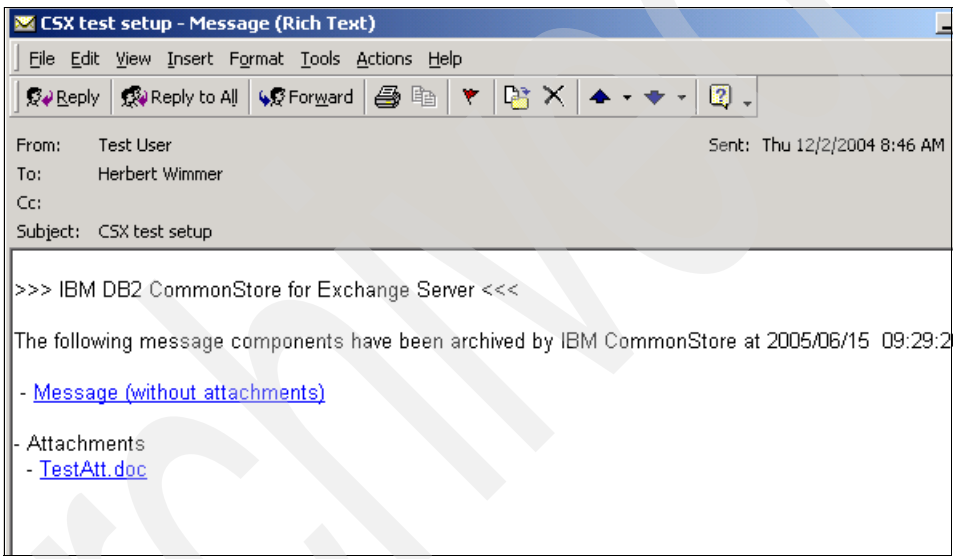

Figure 7-7 Archived sample message

Viewing an archived message

Click the hyperlinks to verify that the data is accessible and that the content of the message and the attachment is displayed correctly.

Note: Outlook is used to display the message remainder. Although you can see the icon for the attachment, the attachment binary data does not exist. Trying to open the attachment causes an error message.

Retrieving an archived message

Select the test message in the Inbox and click (Retrieve).

The CSX Task polling thread reads the retrieve request from the job folder. The CSX Task worker thread forwards the request to the CommonStore Server. When the documents are retrieved from the Content Manager repository, the CSX Task committer replaces the stub with the original message remainder and adds the attachment.

Single instance store (SIS) setup

To verify SIS:

1. Select the test message with subject “CSX test setup” in the Sent Items folder and click (Archive). Since this message was sent from Test User to Test User, there are two messages in Test User’s mailbox that are the same.

   The Archiving Settings dialog displays (see Figure 7-8).

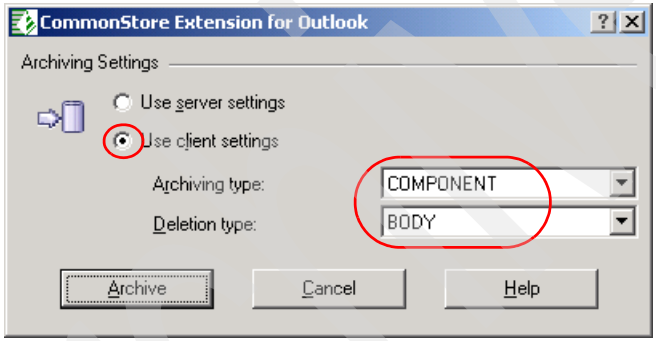

Figure 7-8 Archiving Settings dialog

2. Configure the archiving settings as follows:
   a. Select **Use client settings**.
   b. From the Archiving type pull-down menu, select **COMPONENT**.
   c. From the Deletion type pull-down menu, select **BODY**.
3. Click **Archive**.

The test message, which already is archived from the Inbox, is now archived from a different folder. The CommonStore Server realizes that the message exists in the repository and just adds a child component record to keep additional data.

Use the Content Manager Windows client to verify that only one entry exists for the test message. A search for documents with `SUBJECT LIKE CSX%` results in a hit list as shown in Figure 7-9 on page 260.

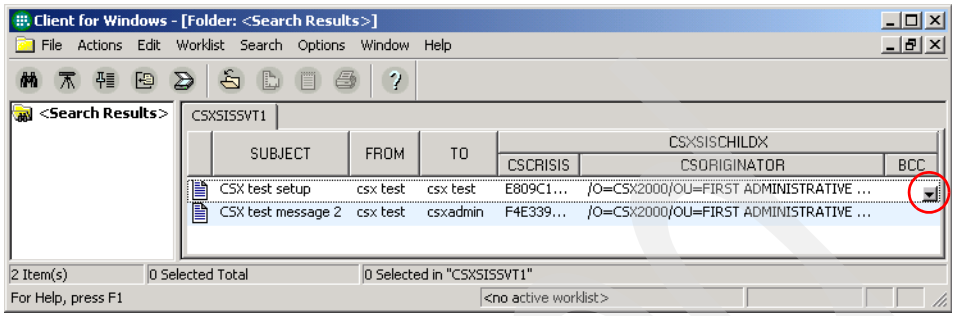

Figure 7-9 Content Manager Windows client's search results

Because the message "CSX test setup" is archived twice, two child entries exist for this message. This is indicated by the pull-down icon near the end of the line. To view the attributes and their values, select **Attributes** on the pop-up menu. This displays Figure 7-10.

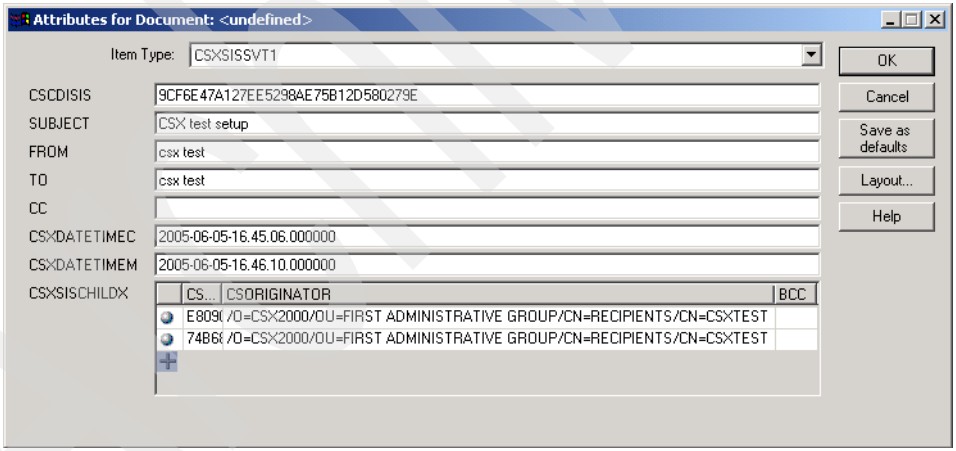

Figure 7-10 Content Manager Windows client's document attributes

Searching for an archived message

To verify the CSX search functionality from the Outlook client, perform the following steps:

1. Click (Search).

   A browser window opens as shown in Figure 7-11.

2. Click the pull-down menu of the Search in field. This displays a list of all attributes defined for the Content Manager item type CSXMail.

3. Select the individual attributes from the Search in pull-down menu. Note that the list of available values in the Search operator pull-down menu changes depending on the attribute selected.

4. Execute a query against the archive by entering the following values:

   Search in SUBJECT

   Search operator LIKE

   Search term CSX%

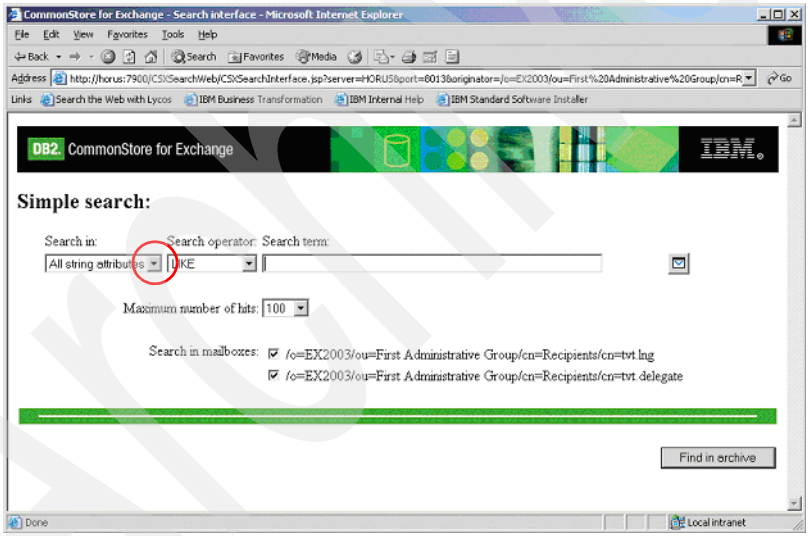

Figure 7-11 Simple search dialog

The resulting list displays as in Figure 7-12 on page 262.

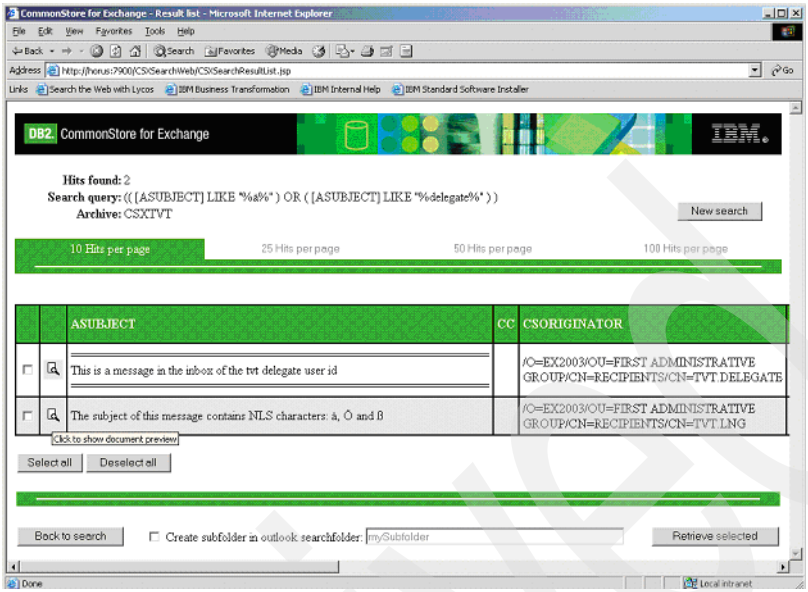

Figure 7-12 Result list window

7.6.2 Key information to remember

At this point, as part of your overall configuration setup documentation, write down the key configuration information that will be used during the rest of the software installation and configuration and future system operation.

Table 7-28 contains the key configuration input values to remember after the CommonStore for Exchange Server installation and configuration.

Table 7-28 Key information to remember after CSX installation and configuration

Configuration data	Sample input value
Windows/Exchange user ID used by CommonStore	CSX Admin
Content Manager user ID used by CommonStore	CSX
Item type used by CommonStore	CSXMail

7.7 Records Manager installation and configuration

In the e-mail archiving and records management solution, Records Manager is used as a records administration application and also as an engine that record enables Content Manager (via Records Enabler for Content Manager), thus records enabling the entire e-mail archiving solution. In this section, we describe the main steps involved in installing and configuring Records.

These steps are as follows:

1. Install Records Manager engine V4.1.1.
2. Install Records Manager database V4.1.1.
3. Upgrade Records Manager engine V4.1.2.
4. Upgrade Records Manager database V4.1.2.
5. Run engine configuration utility.

Note: We do not include all of the detailed steps of the installation in this section. We recommend using the existing product manual in conjunction with the materials we present here for a successful installation and configuration.

In the sample environment, we install Records Manager engine on London and Records Manager database on Charger.

Step 1: Install Records Manager engine V4.1.1

Table 7-29 shows the input value we used during our installation on London in the sample environment. Replace the sample input values according to your environment setup.

Table 7-29 Records Manager Engine installation input for sample environment

Configuration window / field	Sample input value	Description
Installation Destination		
Directory Name	C:\IBM\IRM	Records Manager installation directory. **TIP**: Use a directory without the version number. For updates, the same directory can be used and no second directory with a different version number is created.
Installation Type		
Setup type	typical	

Configuration window / field	Sample input value	Description
Deployment and Configuration		
I want the installer to do deployment and configure for me	selected	The setup program deploys all J2EE applications and configures them.
WebSphere Application Server Connection Information		
Connector Type	SOAP	Specifies the type of communication interface between the WebSphere Application Server and the installation program.
Connector Port	8880	Specifies the port used by the Connector Type.
Cell	london	Specifies the cell name of the WebSphere Application Server installed under 7.4.3, “WebSphere Application Server installation” on page 217. **Tip**: To view the cell name for a server, start the WebSphere Application Server, open the WebSphere Application Console, and in the console navigation tree, click **Servers** → **Application servers** → **Server 1** → **Runtime**.
Node	london	Specify the node name of the WebSphere Application Server installed in 7.4.3, “WebSphere Application Server installation” on page 217. **Tip**: To view a node name, go to WebSphere Application Console and click **Servers** → **Application servers** → **Server1** → **Runtime**.
Server	server1	At the time this book was written, the only possible working value is server1.

Configuration window / field	Sample input value	Description
Security Enabled	unchecked	*Required* if WebSphere security is enabled. If security is enabled, this option has to be checked and a valid user name and password must be provided. To verify whether security is enabled, ensure that the WebSphere Application Server is started, open the WebSphere Application Console, and in the console navigation tree, click **Security** → **Global Security**, and then verify the **Enabled** setting in the **Configuration** tab.
Connection Factories Authentication		
Connection Factories Authentication User	irmwas	This Windows user ID must exist. It does not need any special rights and is used just for authentication. The user ID must not be longer than eight characters (which is the current WebSphere limitation). An application component uses a connection factory to access a connection instance, which the component then uses to connect to the underlying enterprise information system (EIS). Examples of connections include database connections, Java Message Service connections, and SAP R/3 connections.
Connection Factories Authentication password		
Mail Session Configuration		
Mail Transport Host	Brighton	Specifies the name of the server to access for the engine to send e-mail.
SMTP User Name		This is required only if SSL is configured on the SMTP Server. Specifies the name of an e-mail user who has access to send e-mail through the specified transport host. Leave this field blank if the transport host does not require authentication.

Configuration window / field	Sample input value	Description
SMTP User password		Required only if SSL is configured on the SMTP Server. *Required* for the engine to send e-mail. Specifies the password of an e-mail user who has access to send e-mail through the specified transport host. Leave this field blank if the transport host does not require authentication.
Admin Client Configuration		
Root	IRMClient	Specifies the context root for your Records Manager administrator client. This is the name that you use to access the client for Records Manager in your browser (the virtual directory name). For example: `http://london.redbook.bocaraton.ibm.com:9080/IRMClient`
Engine Server Name	London.redbook.bocaraton.ibm.com	Specifies the host name of the computer where the Records Manager engine is installed.
Engine Server ORB Port	8880	Specifies the JNDI service port for the host where you are installing the Records Manager engine. This field specifies the port number on which the application server Object Request Broker (ORB) listens for requests.
Web Services Configuration		
Web Service Configuration Root	IRMWebServices	Specifies the context root for the Records Manager Web server. This is the name that is used to access the Web services for Records Manager in a browser (the virtual directory name).
Web Services Node Name	London.redbook. bocaraton.ibm.com	Specifies the host name of the computer where the Records Manager Engine is installed.
Web Services HTTP Port	9080	Specifies the number for the port that the WebSphere Application Server uses for message queues.

Configuration window / field	Sample input value	Description
Import Export Configuration		
Engine Server Name	London.redbook.bocaraton.ibm.com	The host name of the computer where the Records Manager Engine is installed.
Engine Server ORB Port	2809	Specifies the JNDI service port for the host where the Records Manager engine is installed. This field specifies the port number on which the application server Object Request Broker (ORB) listens for requests.
WebSphere Location		
WebSphere Location	C:\IBM\WS\WAS	The WebSphere Application Server installation directory.

Step 2: Install Records Manager database V4.1.1

In Table 7-30, we provide the input values we used during our installation on Charger in the sample environment. Replace the sample input values according to your environment setup.

Table 7-30 Records Manager database installation input for sample environment

Configuration window / field	Sample input value	Description
Installation Directory		
Directory Name	C:\IBM\IRM\Database	This directory is not the default installation directory.
Database Type		
Database Type	DB2	
JDBC driver class path	C:\IBM\DB2\SQLLIB\java	Path to the *db2java.zip* file. In the sample environment, this file is located in C:\IBM\DB2\SQLLIB\java.

Configuration window / field	Sample input value	Description
DB2 Database Configuration		
DB2 Node/Instance Name	DB2	The instance name if the Records Manager database is being installed directly on a DB2 server. If a remote database is being used, then this is the name of a cataloged DB2 instance. **Note:** This name cannot exceed eight characters in length.
Database Name	irmdb	The name of the DB2 database that is being created. **Note:** The database name cannot exceed eight characters in length, and it must be unique for each database you create.
Default Disk	C	The default location where the database and dataset files are being created. For example, for Windows, type `C` for the C:\ drive.
Folder for Database container	C:\IBM\IRM\Database	The default location where the database and dataset files are being created. This location will have the containers of the table spaces for the database to create.
User name	irmadmin	Specifies the name of the DB2 user that will be the owner of the Records Manager schema. This user will have database administration privileges in the newly created database. **Important:** This user must exist, since it is not created automatically during the installation.
User password		
Territory	default	Specifies a portion of the locale mapped to the country code for the internal processing by the database manager.

Configuration window / field	Sample input value	Description
Collating System	System	Specifies the sequence in which characters are ordered for the purpose of sorting, merging, comparing, and processing indexed data sequentially.
DB Language	English	**Optional:** Specifies a language identifier. Set this field when the database language is different from the default language on your computer. The language you specify must be available on the computer where you are performing the installation.
System administrator user name	db2admin	Specifies the name of the database user with system administrator privileges for the DB2 database instance. This user is created during the installation of the DB2 database in 7.4.1, “DB2 server installation” on page 215.
System administrator user password		Specifies the password of the database user with system administrator privileges for the DB2 database instance.
Database File Plan Population		
Select a plan to populate database	Sample	A sample file plan is created.

Step 3: Upgrade Records Manager engine V4.1.2

The Records Manager engine upgrade is basically a redeployment of the WebSphere Application Server. When running the upgrade, re-enter the values you provided earlier.

Tip: Make sure to use the same installation directory (C:\IBM\IRM\); otherwise, a second directory will be created and the old one will not be deleted.

Step 4: Upgrade Records Manager database V4.1.2

After the Records Manager engine is upgraded to V4.1.2 level, upgrade the Records Manager database to the same level.

Table 7-31 Records Manager database upgrade input for sample environment

Configuration window / field	Sample input value	Description
Installation Directory		
Directory Name	C:\IBM\IRM\Database	Use the same directory as specified during the original installation of the Records Manager database.
Custom or Automatic upgrade		
Automatic	selected	
Database Type		
Database Type	DB2	
JDBC driver class path	C:\IBM\DB2\SQLLIB\java	
DB2 Database Configuration		
DB2 Node/Instance Name	DB2	Instance name as specified in “Step 2: Install Records Manager database V4.1.1” on page 267. This is the instance name if the Records Manager database is installed directly on a DB2 server. It is the name of a cataloged DB2 instance, if a remote database is being used.
Database Name	irmdb	The name of the database that was created in step 2.
Folder for Database container	C:\IBM\IRM\Database	The folder for the database container specified in step 2.
User name	irmadmin	Name of the user that is specified in step 2.
User password		

Configuration window / field	Sample input value	Description
System administrator user name	db2admin	Specifies the name of the database user with system administrator privileges for the DB2 database instance.
System administrator user password		Specifies the password of the database user with system administrator privileges for the DB2 database instance.
Database Back Up		
Selected database was backed up	selected	**Important**: If this is not selected, the upgrade cannot proceed.

Step 5: Run engine configuration utility

The Records Manager engine needs to access the Records Manager database. Because the database can be on different platforms (DB2,Oracle, SQL Server), a data source must be configured. This data source is used by the Records Manager engine to access the database.

If the Records Manager engine and Records Manager database are running on different machines, the database must be cataloged on the Engine machine. In the sample environment, the engine runs on London, and the database is on Charger. We need to catalog the database on London. Use the DB2 Configuration Utility to catalog the database that was created in “Step 2: Install Records Manager database V4.1.1” on page 267. In the sample environment, the remote database on Charger (irmdb) is cataloged as irmdb on London.

The Records Manager engine configuration utility must be run on the same machine that the Records Manager engine is running. To start the utility, select **Start → Program Files → IBM DB2 Records Manager → Engine Configuration Utility**.

In Table 7-32, we provide the input values we used during the startup of the utility for the sample environment. Substitute the sample input values according to your environment.

Table 7-32 Engine configuration utility startup

Configuration field	Sample input value	Description
Connector Type	SOAP	
Port Number	8880	

Configuration field	Sample input value	Description
Cell	london	The WebSphere Application Server cell on which the Records Manager engine is deployed.
Node	london	The WebSphere Application Server node on which the Records Manager engine is deployed.
Server	server1	The WebSphere Application Server server into which the Records Manager engine is deployed.

After the Engine configuration tool is started, a data source (the Records Manager database created in “Step 2: Install Records Manager database V4.1.1” on page 267) must be created.

To create the data source, select **Action** → **New**.

In Table 7-33, we provide the input values we used when creating the new data source. Replace these values with the appropriate ones for your environment.

After creating the new data source, select **File** → **Save Changes**.

Table 7-33 Data source input for the sample environment

Configuration field	Sample input value	Description
DB2 Universal JDBC driver location	C:\IBM\DB2\SQLLIB\java	
Data Source Name	irmdb	Name of the database. You can use any name.
Database Name	irmdb	Name of the cataloged database.
User name	irmadmin	Name of the user created in step 2 who has administrative rights for the database.
User password		

Important: After configuring the data source, if the utility is closed without saving, the provided information will be lost and the data source will not be available during the Records Manager Administration client startup.

Before the Records Manager administration client can be used, the Application Server must be restarted.

Use the Windows Services utility to restart the Service **IBM WebSphere Application Server V5 - server1**, or go to the WebSphere bin directory (C:\IBM\WS\WAS\bin) and use the following commands to restart the server:

```
stopserver server1
startserver server1
```

7.7.1 Installation summary and verification

At this point, the Records Manager engine and the Records Manager database are installed.

In the sample environment, the recommended scenario of installing them on two different computers is performed. Figure 7-13 shows the sample environment after the successful installation.

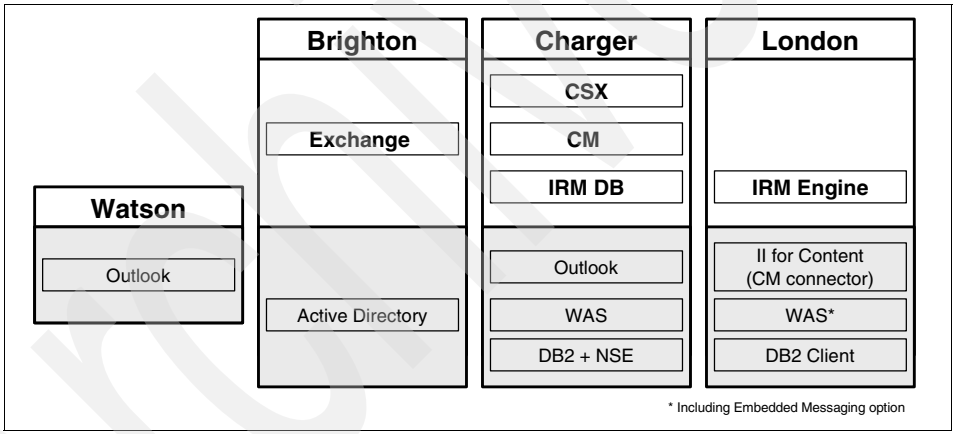

Figure 7-13 Sample environment after Records Manager is installed

To verify the installation, start WebSphere Application Console, go to **Servers** → **Application servers** → **server1** → **Server Components** → **JMS servers**, and make sure in the Initial State field, Started, is there.

To verify that the messaging queue is up and running, go to Windows Task Manager, and look for processes with amq* and runq*. If both of them are there, then the message queue has started. If not, perform these steps:

1. Go to **Servers** → **Application servers** → **server1** → **Server Components** → **JMS servers**.
2. Make sure **Started** is selected in the Initial State.

3. Click **Reset**.
4. Go back and check the Task Manager. The amq* and runq* processes should be there now.

> **Attention:** This procedure is very important. If the message queue is not started, you will encounter problems when working with Records Manager.

To make sure Records Manager is installed properly, log on to Records Manager administration client using `Administrator` as the user ID and `cronos` as the password. Make sure you can log on.

You can also exercise a set of basic Records Manager activities to ensure that the Records Manager system is up and running. Suggested activities include:

1. Create a file plan.
2. Create a record.
3. Perform record scheduling.
4. Turn the crank.
5. Destroy a record via retention rule.

7.7.2 Key information to remember

At this point, as part of your overall configuration setup documentation, write down the key configuration information that will be used during the rest of the software installation and configuration and future system operation.

Configuration data	Sample input value	Description
Records Manager administrator user ID	Administrator password: cronos	
WebSphere Application Server server name	server1	Server into which the Records Manager engine is deployed.

7.8 CMRE installation and configuration

Records Enabler for Content Manager (CMRE) is the bridge between Content Manager and Records Manager. It works with both products to provide the records control capability into the Content Manager system. In this section, we describe the main steps that are involved in installing and configuring CMRE.

These steps include:

1. Set environment variables and create users.

2. Install Records Enabler (CMRE server, Host Interface server, and Permission Synchronization server).
3. Install Records Manager Extension.
4. Implement Windows Services.

Important: Make sure that the Content Manager V8 connector is installed. It is necessary to use the Information Integrator for Content installation to install the connector. A Content Manager client installation is not sufficient.

Note: It is not our intention to include all detailed steps of the installation in this section. We recommend using the existing product manual in conjunction with the materials we present here for a successful installation and configuration.

In the sample environment, we install Records Enabler in London.

Step 1: Set environment variables and create users

Set the environment with the following variables:

```
IBMCMROOT = C:\IBM\CM\db2cmv8
WAS_HOME = C:\IBM\WS\WAS
JDBCPATH = C:\IBM\DB2\SQLLIB\java\db2java.zip
```

Replace these values according to your system installation setup.

Important: The JDBCPATH must include the db2java.zip file name; otherwise, the installation will fail.

Create a local user on the machine that runs the Content Manager Library Server. This user has to be in the DB2 administrator group.

For the sample environment, we create:

User name CMREID
User group DB2ADMNS

Step 2: Install Records Enabler (CMRE)

You need to install CMRE. To help your CMRE installation process, Table 7-34 on page 276 shows the input values we used during our installation. Replace the sample input values as appropriate according to your environment setup.

Table 7-34 Records Enabler installation

Configuration window / field	Sample input value	Description
WebSphere deployment information		
DB2 Content Manager Records Enabler Server	selected	
DB2 Content Manager Records Manager Host Interface	selected	
DB2 Content Manager Records Enabler Permissions Synchronization	selected	
WebSphere Application Server cell name	london	Specify the cell name of the WebSphere Application Server installed in 7.4.3, "WebSphere Application Server installation" on page 217. **Tip**: To view the cell name for a server, start the WebSphere Application Server, open the WebSphere Application Console. In the console navigation tree, click **Servers** → **Application servers** → **Server 1** → **Runtime**.
WebSphere Application Server node name	london	Specify the node name of the WebSphere Application Server installed in 7.4.3, "WebSphere Application Server installation" on page 217. **Tip**: To view a node name, go to WebSphere Application Console, click **Servers** → **Application servers** → **Server1** → **Runtime**.
Host name	London.redbook.bocaraton.ibm.com	The fully qualified host name of the machine running the WebSphere Application Server.

Configuration window / field	Sample input value	Description
WebSphere Application Security Enabled	unchecked	Required if WebSphere security is enabled. If security is enabled, this option has to be checked and a valid user name and password must be provided. To verify whether security is enabled, ensure that the WebSphere Application Server is started, go to the WebSphere Application Console, click **Security** → **Global Security**, and verify the **Enabled** setting on the **Configuration** tab.
Records Manager Server configuration		
Records Manager Web services address	Charger.redbook.bocaraton.ibm.com:2809	
Records Manager Administration Client URL	http://charger.redbook.bocaraton.ibm.com:9080/IRMClient	This value is checked during the installation. If it is not available, installation will not go further.
Records Manager database	irmdb	The name of the Records Manager database.
Records Manager Administrator	Administrator	Records Manager administrative user ID. The default ID is Administrator.
Password	cronos	The default password for the Records Manager Administrator is cronos if it is not changed after the IRM installation.
Content Manager Server configuration		
Server name	icmnlsdb	The name of the Library Server database.
Content Manager authentication	icmadmin	Administrative user ID for the Content Manager System. See 7.5.2, “Key information to remember” on page 225.
Password		
Content Manager Records Enabler Connection ID	cmreid	This user ID is created by the installation program in Content Manager, but it has to exist on the operating system level on the machine running the Content Manager Library Server.

Configuration window / field	Sample input value	Description
password		Password of the user ID on the Windows operating system level.
confirm password		
eClient rendering Content URL	http://....	In the sample environment, eClient is not installed. **Important:** Leave the default value and do not erase that field. With an empty field, the installation will fail. This value can be configured later on using the CMRE Administration client.
eClient document list URL	http://...	The eClient is not installed in the sample environment. **Important:** Leave the default value and do not erase that field. With an empty field, the installation will fail. This value can be configured later on using the CMRE Administration client.
Database System used for Content Manager	DB2	Specifies the database type of the Content Manager Library Server.
Content Manager Records Enabler configuration		
CMRE server	cmresvr	The server will be created during the installation if it does not exist.
Records Manager Host Interface Server	rmecmhost	The server will be created during the installation if it does not exist.
Add Host Configuration record to DB2 Records Manager	checked	Checking this creates a “Host” entry in the Records Manager. The specified Content Manager system will be registered in the Records Manager.
Content Manager Records Enabler Permissions Synchronization	cmrepsproc	The server will be created during the installation if it does not exist.
Permissions Synchronization Scheduler	checked	

Configuration window / field	Sample input value	Description
Permissions Synchronization engine	checked	

Step 3: Installing Records Manager Extension

You must install Records Manager Extension. To help your installation process, Table 7-35 shows the input values we used during our installation. Replace the sample input values according to your environment setup.

Table 7-35 Records Manager Extension installation input for sample environment

Configuration window / field	Sample input value	Description
WebSphere Deployment information		
WebSphere Application Server cell name	london	Specify the cell name of the WebSphere Application Server installed in 7.4.3, “WebSphere Application Server installation” on page 217. **Tip**: To view the cell name for a server, start the WebSphere Application Server, open the WebSphere Application Console, and click **Servers** → **Application servers** → **Server 1** → **Runtime**.
WebSphere Application Server node name	london	Specify the cell name of the WebSphere Application Server installed under 7.4.3, “WebSphere Application Server installation” on page 217. **Tip**: To view the cell name for a server, start the WebSphere Application Server, open the WebSphere Application Console, and click **Servers** → **Application servers** → **Server 1** → **Runtime**.
Host name	London.redbook.bocaraton.ibm.com	The fully qualified host name of the machine running the WebSphere Application Server.

WebSphere Application Security Enabled	unchecked	Required if WebSphere security is enabled. If security is enabled, this option must be checked and a valid user name and password must be provided. To verify that security is enabled, ensure that the WebSphere Application Server is started, open the WebSphere Application Console. In the console navigation tree, click **Security** → **Global Security**, and verify the **Enabled** setting on the Configuration tab.
Records Manager Application server name	server1	At the time this book was written, the only possible working value was server1.

Step 4: Implementing Windows Services

To install a WebSphere Application Server as a Windows Service, use the following command:

```
wasservice -add <Windows Service Name> -serverName <WebSphere Server>
```

To start the Windows services tasks:

1. Open a Command prompt window.
2. Go to WebSphere Application Server's bin directory.
3. Install CMRE server (cmresvr), CMRE Host Interface (rmecmhost), and CMRE Permission Synchronization server (cmrepspro) as Windows services:

```
wasservice -add RMEServer -serverName cmresvr
wasservice -add RMEHostInterface -serverName rmecmhost
wasservice -add RMEPermSync -serverName cmrepspro
```

Substitute values from your environment as appropriate.

7.8.1 Installation summary and verification

Figure 7-14 shows the components that will be installed and configured after this section is completed.

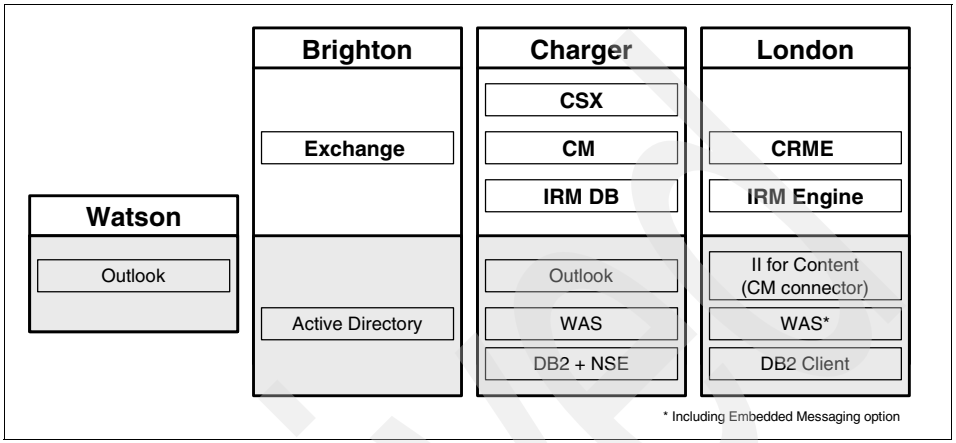

Figure 7-14 Sample environment after CMRE is installed and configured

To verify your installation, perform the following steps:

1. Check the Windows services.
2. Import a Content Manager administrator ID into Records Manager.
3. Log on to the CMRE.
4. Records enable an item type.
5. Use Content Manager Windows client to declare a record.

Step 1: Check the Windows services

To verify the installation, make sure all necessary services are running. They are:

- IBM WebSphere Application Server V5 - server1
- IBM WebSphere Application Server V5 - RMEServer
- IBM WebSphere Application Server V5 - RMEHostInterface
- IBM WebSphere Application Server V5 - RMEPermSyncServer
- ICM LS Monitor icmnlsdb
- IBM HTTP Server 1.3.28

Step 2: Import a CM administrator ID into Records Manager

In order to be able to log on to the CMRE Administration client, a Content Manager user ID with administrator rights must be imported to the Records Manager System, and the user ID needs administrator rights in Records Manager.

To import a Content Manager user:

1. Start the Records Manager client and log on as an administrator:

   User ID Administrator
   Password cronos

2. Go to **Security** → **Users** → **Host Filer** and select the host system that is enabled during the CMRE installation.

   In the sample environment, the host system with the name icmnlsdb is enabled. (See “Add Host Configuration record to DB2 Records Manager” on page 278.)

3. Click **Import**, select **icmadmin**, and click **Import** again. In the next window, select all permissions by checking **Function Access**. Check the **Is Active** check box. Click **Save** to finish the import.

4. The Content Manager user ID icmadmin is now imported to the Records Manager system and has all necessary rights to act as Administrator with the Records Manager system.

Step 3: Log on to the CMRE

Use icmadmin and its password to log on to the CMRE administration client.

Important: Do not use the Records Manager administrator user ID and its password (Administrator/cronos) to log on to the CMRE Administration client because this ID is not defined within Content Manager.

Also, do not use the Content Manager user ID, CMREID, that is created during the CMRE installation, because this ID is not (and cannot be) imported into Records Manager.

Step 4: Records enable an item type

After successfully logging on to the CMRE Administration client, enable an item type as follows:

1. Go to **Content Manager Server Configuration** → **eRecord enable item type**. A list of all item types in the records-enabled Content Manager system is displayed.

2. Select the CSX item type (CSXMail) that was created in step 2a of 7.6, “CommonStore (CSX) installation and configuration” on page 226.

3. Check the box to the left of the item type name and select the record type to the right of the item type name.

In the sample environment, the item type “CSXMail” and the record type “email” are created.

Step 5: Using Windows client to declare records

Using Windows client, declare records as follows:

1. Search for documents in the records-enabled item type CSXMail. The e-mail archived during the CommonStore verification is located in this item type.
2. Right-click one e-mail message and select **Declare Record**. The Records Manager declare/classify window appears.
3. Select a file plan and a unique name for that record. Click **Finish**.

The CMRE communicates with the Content Manager system and the Records Manager system. A Content Manager item type is records enabled, and documents in this item type can be declared as records using the Content Manager Windows client. The records attributes (eRecord, eRecordID, FlPlnCmpntNm, FlPlnCmpntTtl, and RMEAclOri) associated with the e-mail are updated.

7.8.2 Key information to remember

At this point, as part of your overall configuration setup documentation, write down the key configuration information that will be used during the rest of the software installation and configuration and future system operation.

Table 7-36 contains the key configuration input value to remember after the Content Manager installation.

Table 7-36 Key information to remember after CMRE installation and configuration

Configuration data	Sample input value
Records enabled item type	CSXMail

7.9 Records enable CommonStore Server and Outlook

In this section, we describe the main steps involved in configuring CommonStore Server with Records Manager and the Outlook client. These steps must be performed to add records declaration and classification capability to the end-to-end integrated e-mail solution.

The steps involved include:

1. Records enable Content Manager item type.
2. Configure ArchPro.
3. Configure the CSX Task.
4. Records enable Outlook.

For detailed information about CSX and Records Enabler integration, refer to *IBM DB2 CommonStore for Exchange Server: Administration and User's Guide, Version 8.3*, SH12-6741.

Step 1: Records enabling Content Manager item type

This should have been done in the previous section when you validated the installation and configuration of CMRE. If it has not been completed, refer to "Step 4: Records enable an item type" on page 282.

Step 2: Configuring ArchPro

An Exchange user must be mapped to a Content Manager user ID to be able to declare a record, because the solution is records enabled in Content Manager. This Content Manager user also has to be imported to Records Manager and must have the appropriate rights to declare records. For the mapping of Exchange user IDs to Content Manager user IDs, an additional component (usermapper.jar) has to be activated on the CommonStore ArchPro server.

To activate the usermapper, follow the steps below:

1. Copy usermapper.jar from the installation bin directory into the instance directory of the ArchPro that will be activated.

   Source directory: C:\IBM\CSX\bin

   Target directory: C:\IBM\CSX\server\instance01

2. Add usermapper.jar to the classpath and configure the directory of CSExit.properties. In the sample environment, we add the paths according to Table 7-37.

Table 7-37 Classpath update to include usermapper.jar and CSExit.properties file

File/Directory	Value to add to classpath
usermapper.jar	C:\IBM\CSX\server\instance01\ usermapper.jar
CSExit.properties directory	C:\IBM\CSX\server\instance01

3. Generate encryption keys.

   This step involves generating a pair of keys. The services provided by the usermapper.jar on the CSX Server listen on a preconfigured port for requests from Content Manager Records Enabler code in the CSX Task. The e-mail user's Content Manager credentials are included in the information passed to the CSX Task. Because these credentials cannot be passed as is, use a key pair for encryption at the CSX Server side and decryption on the CSX Task Side. After installation of the CSX Task, you must generate the public/private key pair.

To generate the key pair, complete the following steps:

a. Create the key pair by typing `java com.ibm.rme.csexit.KeyGen <name>` where <name> is the name you want to give this client. Two files are created.

b. Copy <name>.prv to the directory C:\IBM\CSX\bin.

c. Place <name>.pub in the CLASSPATH of the CSX Server. For example, place the <name>.pub in C:\IBM\CSX\Server\instance01\.

4. Update CSExit.properties.

CSExit.properties is located in the instance directory and contains three parameters that must be configured:

```
DB_DIR
HASH_MODULO
PROXY_PORT
```

Set DB_DIR to the directory to store the mapping database. (Note the double backslashes; a single backslash indicates an escape sequence, such as \n.) This directory will contain a collection of files that are serialized hash tables containing string keys of the format `CM-server:mail-user` and CSRepUserDef values. Make sure this folder exists and is empty. The files will be generated automatically, along with a file that indicates the current HASH_MODULO value so that it can be changed.

Set HASH_MODULO to the maximum number of files to be found in the DB_DIR directory. This is to prevent the entire database from ever needing to be in memory at one time so that a huge number of users can be supported. Bigger values mean smaller memory usage (but more files).

Set PROXY_PORT to the port on which the usermapper proxy is listening.

In the sample environment, we set these parameters as listed in Table 7-38 on page 285.

Table 7-38 CSExit.properties parameter update

Parameter	Sample input value
DB_DIR	C:\IBM\CSX\server\instance01\database
HASH_MODULO	1000
PROXY_PORT	8067

5. Update archint.ini.

   To activate the usermapper, add the values for the following parameters in the archini.ini file:

   ```
   ACCESS_CTL
   CM_SECURITY_EXIT
   CM_EXIT_LOCATION
   ```

   The input value for ACCESS_CTL YES specifies whether you want Retrieve operations to be subject to the user's Content Manager permissions. The YES setting causes the user mapper to be activated when **archpro** is started. You must activate the user mapper to do declare records.

   The input value for CM_SECURITY_EXIT specifies the name of the security exit class as com.ibm.rme.csexit.CSExit.

   The input value for CM_EXIT_LOCATION specifies the file location of the usermapper.jar file.

   In our sample environment, we set these values as shown in Table 7-39.

Table 7-39 Archint.ini updates

Parameter	Sample input value
ACCESS_CTRL	YES
CM_SECURITY_EXIT	com.ibm.rme.csexit.CSExit
CM_EXIT_LOCATION	'C:\IBM\CSX\server\instance01\usermapper.jar'

Step 3: Configuring the CSX Task

Use the CSX System Manager to add Records Enabler support to your configuration. To records enable the CSX Task:

1. Start the System Manager.
2. Create the Records Enabler configuration:
   a. Go to **Records Enabler** Configuration.
   b. Right-click **Add records enabler configuration**.
   c. Set the configuration as shown in Table 7-40.

Table 7-40 Records enabler configuration

Configuration field	Sample input value	Description
Configuration name		The name of this Records Enabler configuration object.

Configuration field	Sample input value	Description
Host name	london.redbook.bocaraton.ibm.com	The Web Application Server on which the CMRE Client Servlet runs. You can enter a DNS name or an IP address. Do not specify localhost.
Port number	9081	The number of the port for communication between the Web Application Server and the CSX Task.
Application name	RMEServer/RMEClientServlet	The name of the CMRE Client Servlet instance that you want to use.
User mapper proxy port	8067	The value in the User mapper proxy port field must match the PROXY_PORT value in the file CSExit.properties for the archpro instance.
Private key file name	csxkeys	Enter the name you used when generating the encryption key in step 3 on page 284.

3. Update task configuration:
   a. Go to **Task Administration Data**.
   b. Double-click **TaskBrighton**, and select **Parameters**.
   c. In the Records Enabler Configuration field, insert the Records Enabler Configuration that you just created.

Step 4: Records enabling Outlook

In order to get the Outlook Client to support Content Manager Records Enabler, the Content Manager Records Enabler component has to be selected during installation of the CSX Client Components.

During the installation of the Outlook Extension the necessary functions for record declaration are already installed.

7.9.1 Verification

To verify that the CommonStore Server and Outlook client are configured properly, you can perform the following steps:

1. In the Outlook client, create an e-mail message and send it off.
2. Manually declare the e-mail message as a record.

Note: The first time you use an Outlook client to declare a record, you may be prompted for your Content Manager credentials (user ID and password) between step 3 and step 4. After you supply the credentials, the CSX Task authenticates you with the CMRE server. If the credentials can be authenticated, they are then stored in the user mapping table on the archpro server. After this, whenever the user manually declares or views records, these credentials are passed to IRM in order to authenticate the mailbox user prior to launching the IRM classification window. IRM authenticates credentials via Content Manager.

3. Make sure that the e-mail is archived in Content Manager using the Content Manager Windows client.
4. After the e-mail is archived, the Records Manager classification window should appear. Specify the bucket in the file plan (for example, Account Receivable in our sample environment). Enter a unique ID into the e-mail name field, and press **Finish**.
5. Make sure that in the Outlook client, it reflects that the e-mail is a record.

Note: The visual indicator that shows whether a message is a record does not automatically appear in the views of your Outlook client. Outlook allows you to display these additional columns through the Tools → Options on the menu bar. The complete instructions are found in the CSX publication.

6. Using Content Manager Windows client, verify that the e-mail's isRecord attribute is set to yes.
7. Using Records Manager administration client, check that the record is there.
8. Using the Outlook client, view the e-mail record.

8

Installation and configuration in a Lotus Domino and AIX environment

This chapter describes the installation and configuration of an e-mail archiving and records management solution using CommonStore for Lotus Domino (CSLD), Records Manager, Content Manager, and Content Manager Records Enabler. Using a sample environment, we describe the major steps involved in installing and configuring the various components in an AIX environment. For more detailed information, see the appropriate product documentation.

We cover the following topics in this chapter:

- Overview
- Prerequisites and prerequisite software installation
- Content Manager installation and configuration
- CommonStore installation and configuration
- Records Manager installation and configuration
- CMRE installation and configuration
- Configuring the CommonStore Server and Notes

8.1 Overview

In this section, we provide an overview for the e-mail archiving and records management integrated solution installation and configuration.

We cover:

- Software used for the integrated solution
- Installation and configuration steps and recommendation

8.1.1 Software used for the integrated solution

Several products are used in this end-to-end integrated solution. Table 8-1 lists the software used and its purpose in the solution.

For clarity, fix pack details have been omitted. These are referenced later in solution installation and configuration.

Table 8-1 Software used in the integrated solution and its purpose in the solution

Product	Purpose
IBM DB2 Content Manager (CM)	Repository used to store the documents and metadata for both the archive and records management systems.
IBM CommonStore for Lotus Domino (CSLD)	E-mail archive system for Lotus Notes.
IBM DB2 Records Manager (IRM)	Engine and administration for Records Manager.
Records Enabler for Content Manager (CMRE)	Records enables the Content Manager repository. Also provides records management functions for Lotus Notes or Outlook users.
IBM DB2 ESE UDB (DB2)	Enterprise-class database used to hold both system configuration and objects metadata
DB2 Net Search Extender	Extension to DB2 that adds full text search capabilities for both object metadata and documents including attachments.
IBM Web Sphere Application Server with Embedded Messaging	Web application server that hosts the Content Manager Resource Manager, Content Manager Records Enabler servers, the Records Manager applications, and Content Manager eClient.
Lotus Domino server	E-mail system.

8.1.2 Installation and configuration steps and recommendation

Four main products are involved in the end-to-end solution:

- IBM DB2 Content Manager
- IBM DB2 CommonStore for Lotus Domino
- IBM DB2 Records Manager
- IBM DB2 Records Enabler for Content Manager

Plan your system configuration first. (See 5.3, “System configuration” on page 120.) Decide where you want to install each main product and review the prerequisites for each product (especially if you decide to separate some components onto different machines). Before you start, make sure you know exactly what should be installed on each machine, and the sequence of the installation.

Some of the product installation and configuration can be done at the same time and others are better to be done sequentially. We recommend the following sequence of steps for installation and configuration:

1. Install a working Content Manager system on a designated machine.

   This includes the installation of the prerequisites first (DB2, Net Search Extender, WebSphere Application Server), and then the Content Manager product.

   Validate that the Content Manager system is working by importing some documents, retrieving them, and viewing them.

   If you decide to separate the installation of Content Manager Library Server from Resource Manager to a different server, you may need to install different prerequisites onto the machines. Refer to the product documentation to install the proper prerequisites for each server. You also need to validate the system after the installation and configuration as mentioned above.

2. Install a working CommonStore system on a designated machine.

   This includes the installation of the prerequisites (Content Manager V8 connector from the Information Integration for Content installation, DB2 Runtime Client or DB2 Administration client, and Notes client), and then the CommonStore for Lotus Domino product.

   Validate that the CommonStore is working properly with your mail database and Content Manager: Setting up some policies, archiving some e-mail, retrieving the archived e-mail, and viewing them. Also look into the Content Manager repository to make sure that the archived e-mail is there as expected.

3. Install a working Records Manager on a designated machine.

   This step can be done in conjunction with step 2. If you have multiple people doing the installation, you can work together in parallel. Otherwise, we recommend performing this task after the CommonStore installation.

   The step includes the installation of the prerequisites, then the Records Manager product.

   We recommend installing the Records Manager engine and its database on separate servers. The prerequisites for the engine include WebSphere Application Server and DB2 Runtime Client or DB2 Administration Client. The prerequisite for the Records Manager database is DB2 server. As discussed in 5.3.1, “Configuration options” on page 122, you can optionally put the Records Manager database where Content Manager is installed.

   Validate that Records Manager is working properly by using the Records Manager administration client to create a default file plan, add a record to the system, and view the added record.

4. Install and configure Records Enabler for Content Manager on a designated machine. Configure Content Manager and Records Manager. Install and configure Records Manager Extension.

   Records Manager Extension must be deployed on the same WebSphere Application Server as Records Manager.

   Validate the system: Import an item of the record enabled item type, declare it as a record, and make sure that the record is marked as declared in Content Manager and that the record’s metadata is in Records Manager. After this, archive an e-mail and declare the e-mail as a record. Check both Content Manager and Records Manager to ensure that the proper information is stored there. (Content Manager should have the e-mail information, based on your archiving method, and Records Manager should have the record’s related metadata).

Tip: For the machine that will run Permission Synchronization Server of the Records Enabler for Content Manager product, make sure that WebSphere Application Server with Embedded Messaging is installed on it. This is especially important if you start your installation from an existing system, such as a working Content Manager system.

If the Embedded Messaging feature was not installed, you must deinstall Content Manager, uninstall WebSphere Application Server, then reinstall WebSphere Application Server, including the Embedded Messaging feature (installed by default in V5.1.1 and later), and then Content Manager again. Do not take shortcuts or you may experience strange results.

If the Content Manager is in use and cannot be deinstalled, we recommend installing the Permission Synchronization Server on another machine.

8.2 Introduction to the sample environment

We use a sample environment to show you how to install and configure an e-mail archiving and records management solution in an AIX environment.

Figure 8-1 shows the sample environment before any e-mail archiving and records management software are installed.

There are two servers in this environment, Jamaica and Bonnie. A Lotus Domino server is running on Bonnie. This should be your starting point.

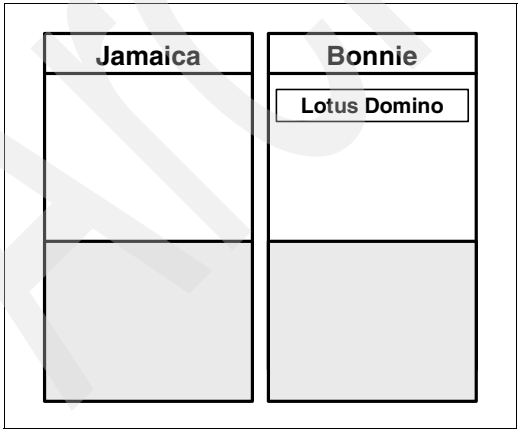

Figure 8-1 Sample environment before any software installation

Figure 8-2 shows the sample environment after all components are installed, including the necessary prerequisites on each server.

On Jamaica, we install prerequisite software including a Notes client (Lotus Domino), a WebSphere Application Server (WAS in Figure 8-2), and DB2 server with Net Search Extender (DB2 + NSE). In addition, we install CommonStore for Lotus Domino (CSLD), Content Manager (CM), and IBM Records Manager database (IRM DB).

On Bonnie, we install prerequisite software including the Content Manager V8 connector (CM connector) from Information Integrator for Content, WebSphere Application Server with Embedded Messaging (WAS*), and a DB2 client (DB2). In addition, we install Content Manager Records Enabler (CMRE), and IBM Records Manager Engine (IRM Engine).

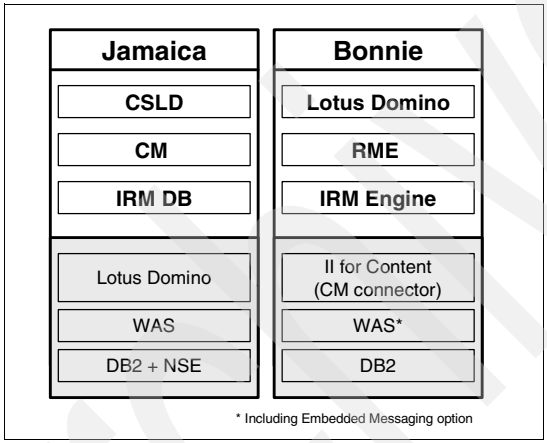

Figure 8-2 Sample environment after all software installation

8.3 Prerequisites

Four core components are involved in the e-mail archiving and records management solution: Content Manager for Multiplatforms, CommonStore for Lotus Domino, IBM Records Manager, and Content Manager Records Enabler.

Each core component has different prerequisites. We list what they are and their version requirements, and explain why you need to install the prerequisite. Understanding this should help you when you build a system that fulfills your business needs.

Prerequisite for Content Manager V8.3

Table 8-2 describes the prerequisites for Content Manager V8.3.

Table 8-2 Prerequisites for Content Manager V8.3

Product	Version	Reason why we need it
WebSphere Application Server	5.1.1.2	Resource Manager is a J2EE application and thus needs WebSphere Application Server.
DB2	8.2	Library Server uses stored procedures and needs a relational database (icmnlsdb). Resource Manager stores metadata in a relational database (rmdb).
DB2 Net Search Extender	8.1	For full-text search within Content Manager.

Prerequisite for CommonStore for Lotus Domino V8.3

Table 8-3 describes the prerequisites for CommonStore for Lotus Domino V8.3.

Table 8-3 Prerequisites for CommonStore for Lotus Domino V8.3

Product	Version	Reason why we need it
Information Integrator for Content - CM V8 connector	8.3.0	The CommonStore agent needs the APIs (connector) to communicate with the Content Manager Library Server.
DB2 Runtime Client	8.2	If CommonStore for Lotus Domino is not on the same machine as the Content Manager Library Server, the Content Manager V8 connector (which is needed by the CommonStore agent) needs the Library Server database to be cataloged on the machine where CommonStore is installed. DB2 Runtime Client is thus needed.
Domino	6.5	The CommonStore Task communicates with the Domino server and requires Notes API. Domino R6 on AIX requires IO Completion Port to be installed and enabled on AIX.
Visual Age for C++	depends	The CommonStore binaries require the C++ runtime libraries. The version depends on the version of your AIX.

Prerequisite for IBM Records Manager V4.1.2

Table 8-4 describes the prerequisites for the Records Manager engine for IBM Records Manager V4.1.2.

Table 8-4 Prerequisites for IBM Records Manager Engine V4.1.2

Product	Version	Reason why we need it
WebSphere Application Server	5.1.1.2	Resource Manager is a J2EE application.
DB2 Runtime Client	8.2	The Records Manager Engine needs to communicate with the Records Manager database. If the Records Manager database is installed on another machine, then the machine with the Records Manager Engine installed needs the database to be cataloged on it and thus it needs DB2 Runtime Client.

Table 8-5 lists the prerequisite for the Records Manager database for IBM Records Manager V4.1.2.

Table 8-5 Prerequisite for IBM Records Manager (Database) V4.1.2

Product	Version	Reason why we need it
DB2	8.2	The Records Manager database is a relational database. The machine that installs the Records Manager database needs to install DB2 server.

Prerequisites for Records Enabler V8.3

Records Enabler for Content Manager (CMRE) is comprised of three components: Records Enabler (CMRE) server, Permission Synchronization (PermSync) server, and Host Interface server.

Table 8-6 lists the prerequisites for the CMRE server, PermSync server, and Host Interface server.

Table 8-6 Prerequisites for CMRE server, PermSync server, and Host Interface

Product	Version	Reason why we need it
WebSphere Application Server with Embedded Messaging	5.1.1.2	CMRE server, PermSync server, and the Host Interface require WebSphere Application Server. PermSync server also needs Embedded Messaging.

Product	Version	Reason why we need it
Information Integrator for Content - CM V8 connector	8.3.0	All three components communicate with Content Manager Library Server and therefore need Content Manager APIs.
DB2 Runtime Client	8.2	The Content Manager V8 connector needs to communicate with the Content Manager Library Server database. If the Content Manager Library Server is installed on another machine, then, the machine with these servers installed must have the database cataloged on it; thus it needs DB2 Runtime Client.

Table 8-7 lists the prerequisite for the Records Manager extension.

Table 8-7 Prerequisite for Records Manager Extension V8.3.0

Product	Version	Reason why we need it
WebSphere Application Server	5.1.1.2	The CMRE server, PermSync server, and the Host Interface are WebSphere Application Server applications. They need Version 5.1 with Fix Pack 1 and cumulative Fix 2.

8.4 Prerequisite software installation

In this section, we describe the main steps involved in installing the basic software that is required before the main components of the e-mail archiving and records management solution are installed:

- "DB2 server installation" on page 298
- "DB2 Administration Client installation" on page 300
- "WebSphere Application Server installation" on page 301
- "Information Integrator for Content (CM connector) installation" on page 304

Table 8-8 on page 298 summarizes the prerequisites for each software product from 8.3, "Prerequisites" on page 294. Note that:

- DB2 client (runtime or administration) is needed for Records Manager (IRM) engine, and DB2 server is required for Records Manager database if they are installed on separate machines; otherwise, they need only DB2 server.
- WebSphere Application Server with Embedded Messaging (WAS* in Table 8-8 on page 298) is required for Permission Synchronization server from CMRE.

- DB2 Net Search Extender is required if you need full text search capabilities.

Table 8-8 Prerequisites per software product

CM	CSLD	IRM	CMRE
WAS		WAS	WAS* (PermSync svr)
DB2 +NSE	DB2 client	DB2 client (engine) DB2 (database)	
	II for Content - CM V8 connector		II for Content - CM V8 connector
	Notes client		

Using the sample environment that we described in 8.2, “Introduction to the sample environment” on page 293, we need to install the following prerequisites on two servers:

- Bonnie:
  - DB2 Runtime Client V8.2
  - Information Integrator for Content - Content Manager V8 connector
  - WebSphere Application Server (including Embedded Messaging) V5.1.1.2
- Jamaica:
  - DB2 server V8.2
  - DB2 Net Search Extender V8.2
  - WebSphere Application Server (Embedded Messaging not needed) V5.1.1.2
  - Lotus Domino

Note: We do not include all detailed steps of the installation in this section. We recommend using the existing product manuals in conjunction with the materials we present here for successful installations and configurations.

8.4.1 DB2 server installation

Content Manager and Records Manager use relational databases to store content (objects and records) metadata and system configuration information.

Install DB2 server to the machine (or machines) that will store Content Manager databases and Records Manager databases.

The steps involved are summarized as follows:

1. Install DB2 server software.
2. Verify DB2 server installation.
3. Install DB2 Net Search Extender.
4. Install DB2 Fix Pack 8.

In our sample environment, both the Content Manager database and Records Manager database are located on *Jamaica*. We choose to install the DB2 server on this machine.

Installing DB2 server

To help your DB2 server installation process, we provide the input values we used during the DB2 server installation in Table 8-9. Substitute the sample input values according to your environment setup.

Table 8-9 DB2 server installation input summary for the sample environment

Required input field	Sample input value	Description
DB2 Version	8.1.7	Content Manager V8.3 requires at least DB2 V8.1.7.
Installation type	Typical	
File system	/usr	
Installation directory	/usr/opt/db2_08_01	
DB2 Administration server user ID	db2inst1	This AIX user ID is the DB2 instance owner. In AIX, the DAS is created under an instance.
DB2 administration group	db2grp1	Every user ID that needs to be a DB2 administrator has to be in this group.
DB2 instance	db2inst1	

Verifying installation

To verify the DB2 server installation, create the sample database from the First Step menu. Make sure the sample database is created successfully and that you can view the data.

Installing DB2 Net Search Extender

This is an optional step. If you need to use the full text search feature of Content Manager, install DB2 Net Search Extender before installing Content Manager. After Content Manager is installed, you can activate the full text search feature.

Use the database administrative user ID (db2admin) to run DB2 Net Search Extender as a service.

In our sample environment, we choose not to configure for full text search.

Installing DB2 Fix Pack 8

After installing the DB2 server and DB2 Net Search Extender, install the DB2 Fix Pack 8. The Records Manager database requires DB2 V8.2 (which is equivalent of V8.1.8) as a prerequisite.

See the appropriate documentation for detailed steps of installing the fix pack.

Key information to remember

At this point, as part of your overall configuration setup documentation, write down the key configuration information that will be used during the rest of the software installation and configuration and future system operation.

Table 8-10 lists key configuration input values to note after DB2 server installation.

Table 8-10 Key information to remember after the DB2 server installation

Configuration data	Sample input value
DB2 Administration server user ID	db2inst1
DB2 administration group	db2grp1
DB2 database server host name	jamaica.almaden.ibm.com

8.4.2 DB2 Administration Client installation

We recommend installing the Records Manager database and Records Manager engine on separate servers.

The Records Manager engine needs access to the Records Manager database. If they are not installed on the same machine, the Records Manager database has to be cataloged on the Records Manager engine machine. To catalog a database, a DB2 Runtime Client is required.

In addition to installing the DB2 client to the machine where the Records Manager engine will run (if it is installed on a separate machine from the Records Manager database), the DB2 client should also be installed on the machine where you will run CommonStore.

Although the DB2 Runtime Client is the minimum requirement for both instances, we recommend installing the DB2 Administration Client because it provides a

simple graphical user interface to catalog a database and other powerful DB2 tools to administer a remote database.

The installation process is straightforward and we do not cover it here.

In the sample environment, we install the Records Manager database on Jamaica and the Records Manager engine on Bonnie. We therefore need to install DB2 client on Bonnie.

8.4.3 WebSphere Application Server installation

The three core components of the solution (Content Manager, Records Manager, and Records Enabler) rely on WebSphere Application Server, which should be installed on the machines where these products are installed.

Important: Before you continue, make sure that WebSphere Application Server with Embedded Messaging is installed on the server that runs the Records Enabler; otherwise, you may encounter problems later.

If the Embedded Messaging feature was not installed, you must deinstall Content Manager, uninstall WebSphere Application Server, then reinstall WebSphere Application Server, including the Embedded Messaging feature (installed by default in the latest version), and then Content Manager again. Do not take shortcuts or you may experience strange results.

The steps involved in WebSphere Application Server installation are summarized as follows:

1. Install WebSphere Application Server software.
2. Verify installation.
3. Install Fix Pack 1 and any accumulative fix packs.

In the sample environment, it is necessary to install WebSphere Application Server on both Jamaica and Bonnie. We also have to install Embedded Messaging on Bonnie because Records Enabler will be installed on the machine.

Installing WebSphere Application Server software

To help your WebSphere Application Server installation process, we provide the input values we used during our installation on both servers in Table 8-11 on page 302. Replace our sample input values according to your environment setup.

Table 8-11 Installation input summary for the sample environment

Required input field	Sample input value	Description
WebSphere Application Server version	5.1.0	
Installation type	Full	This includes Embedded Messaging. Again, this must be installed on the machine that will run the Records Manager engine. **Note:** When choosing full installation, the installation includes all of the sample applications, and that may increase the install time.
Installation directory for WebSphere Application Server	/usr/WebSphere/AppServer	
Installation directory for IBM HTTP Server	/usr/IBMHttpServer	
Installation directory for Embedded Messaging server and client	/mqm	
Node name	for Jamaica: jamaica for Bonnie: bonnie	We install WebSphere Application Server on both servers.
Host name	for Jamaica: jamaica.almaden.ibm.com for Bonnie: bonnie.almaden.ibm.com	
WebSphere Administrator user ID	wsadmin	This is the Windows user ID used to run the WebSphere services.

Verifying installation

To verify successful WebSphere Application Server installation, at the WebSphere Application Server - First Steps window, click **Verify Installation**. The confirmation message `Installation Verification is complete` should appear.

Installing Fix Pack 1 and cumulative fixes

In order to fulfill the prerequisite, run the Fix Pack 1 update wizard, and then run the cumulative Fix 3 update wizard.

Important: Before the installation, stop all WebSphere services (server1) as well as the IBM HTTP Server.

Refer to Table 8-12 for the values we used during the fix pack and cumulative fixes installation. Replace our sample input values for your environment setup.

Table 8-12 Input values for fix pack and cumulative fixes

Required input field	Sample input value	Description
WebSphere Application Server Version	5.1.1.3	
Stop all WebSphere Application Servers	server1	For the fix pack installation, all services must be shut down.
Stop IBM HTTP Server		The fix pack must update the plug-in files, which would be blocked by any running IBM HTTP server.
JAVA_HOME	/usr/WebSphere/AppServer/java	

Key information to remember

At this point, as part of your overall configuration setup documentation, write down the key configuration information that will be used during the rest of the software installation and configuration and future system operation.

Table 8-13 contains the key configuration input values to remember after the WebSphere Application Server installation.

Table 8-13 Key information to remember after WebSphere Application Server installation

Configuration data	Sample input value
WebSphere cell	for Jamaica: jamaica for Bonnie: bonnie
WebSphere node	for Jamaica: jamaica for Bonnie: bonnie

8.4.4 Information Integrator for Content (CM connector) installation

Content Manager V8 connector (which comes with the Information Integrator for Content) is a prerequisite for Records Enabler for Content Manager (CMRE) and CommonStore. Install Content Manager V8 connector on both the CommonStore machine and the Records Enabler machine if these products are not installed on the same system where Content Manager is installed.

Content Manager V8 connector must be installed before CMRE installation.

Attention: We recommend installing Content Manager *before* you install the connector to any of the machines in question. This is because during the Content Manager installation, many values necessary to configure the connector will be defined.

Although we describe the connector installation in this section, defer the installation until the Content Manager installation is done as described in 8.5, “Content Manager installation and configuration” on page 307.

Installing Content Manager V8 connector

To install Content Manager V8 connector, launch the Information Integrator for Content installation process, select **Connector**, then **Content Manager V8 Connector** from the appropriate installation windows. Refer to the product manual for specific installation instructions.

Table 8-14 lists input values we used for installation in our sample environment. Replace our sample input values according to your environment setup.

Table 8-14 CM V8 connector installation input for sample environment

Required input field	Sample input value	Description
Database server type	DB2 Universal Database	This is the database used by the Content Manager system.
Library Server database	icmnlsdb	Library Server stores metadata in a relational database. This value specifies the name of the database.
Library server schema name	icmadmin	This AIX user ID has to be created for the Content Manager installation.
Authentication type	Server	

Required input field	Sample input value	Description
icmcont	icmconct	This AIX user ID is used as a connection user ID for clients to connect to the Library Server database if they do not have a valid database user ID.
icmconct password		
Local		

8.4.5 Domino server installation

The Domino server is a prerequisite for the CommonStore Server because the CommonStore Tasks require the Notes API to use a Notes user ID to connect to the Domino mail servers.

The AIX I/O completion ports are a prerequisite for the Domino server. If your system does not already have them installed, install and configure them.

Installing and configuring AIX I/O completion ports

To install and configure AIX I/O completion ports, complete the following steps:

1. Install bos.iocp.rte from the AIX installation media.
2. Run **cfgmgr** to make the system aware of the new devices.
3. Run **smitty** and select **Devices** → **IO Completion Ports** → **Configure IO Completion Ports**. Set it to be available after system restart.
4. Shut down and restart AIX to make the I/O Completion Ports available.

Installing the Domino server

To install a Domino server, perform the following steps:

1. Create the following AIX accounts as shown in Table 8-15.

Table 8-15 AIX accounts

Configuration field	Sample input value	Description
AIX group	notes	This group is used to set the group permissions on the files in the Domino data directory.

Configuration field	Sample input value	Description
AIX user	notes	The Domino instance owner; used to set owner permissions on the files in the data directory. Set the primary group to notes.
AIX user home path	/home/notes	Make sure this file system has at least 250 MB free space.

2. Install the Domino server with setups as shown in Table 8-16.

Table 8-16 Domino server installation options

Configuration field	Sample input value	Description
Server type	Enterprise Server	
Partitioned server	no	
OS user ID	notes	
OS group ID	notes	
Domino data directory	/home/notes	This is the home path of the AIX user ID created to be the Domino instance owner.
File system	/opt	Make sure that there is at least 550 MB free space.
Installation directory	/opt/lotus	

Note: No configuration of the Domino server is required as we only needed it for the Notes APIs on Jamaica.

Figure 8-3 on page 307 shows the sample environment after all necessary prerequisite software has been installed.

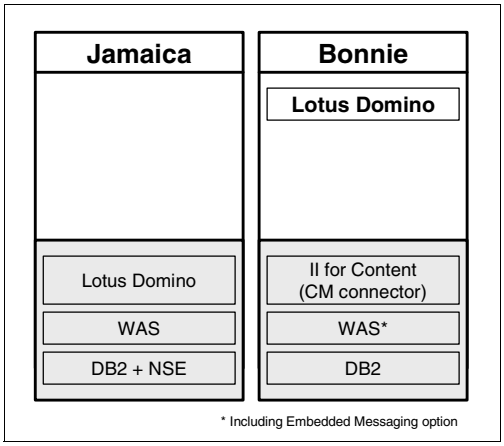

Figure 8-3 Sample environment after prerequisites installed

8.5 Content Manager installation and configuration

In the e-mail archiving and records management solution, Content Manager is used as a repository for archived e-mail. When e-mail becomes records, they are still stored in the Content Manager repository.

In this section, we describe the main steps involved in installing and configuring Content Manager, a major component in the integrated solution.

The main steps involved include:

1. Create the required AIX user IDs.
2. Install Content Manager system.
3. Verify installation.

Note: It is not our intention to include all the detailed steps of the installation in this section. We recommend using the existing product manual in conjunction with the materials we present here for a successful installation and configuration.

In the sample environment, we install Content Manager on Jamaica.

Creating the required AIX user IDs

1. Create an AIX user ID (rmadmin) for Content Manager Resource Manager and an ID (icmadmin), for Content Manager Library Server and general system administration. See Table 8-17 on page 308 for the configuration.

Table 8-17 AIX user IDs for Library Server and Resource Manager

User name	Primary group	Group set	Description
rmadmin	ibmcmgrp	staff, db2grp1	This user ID owns the Resource Manager database (rmdb) schema. It must be added to the group *db2grp1*, and the .profile of this user has to execute the DB2 profile.
icmconct	ibmcmgrp	staff, db2grp1	This user ID is the default ID that is used to connect to the databases if another ID is used and does not have access to the database. It must be added to the group *db2grp1* and the .profile of this user has to execute the DB2 profile.
icmadmin	ibmcmgrp	staff, db2grp1	This user ID owns the Library Server database (icmnlsdb) schema. It is Content Manager's administrator ID. It must be added to the group *db2grp1* and the .profile of this user has to execute the DB2 profile.

2. After the user IDs are created, set their passwords.

Note: All new AIX user accounts require the password to be changed the first time they are used after being created. To do this, Telnet to localhost and log on as the new user. After you enter the password, it will prompt you to change the password. Enter the same password as before and again to confirm it. If this is not done, any applications such as DB2 that require this ID for authentication will fail the authentication.

Installing Content Manager system

We choose a typical Content Manager system installation here.

Attention: During a typical installation, full text search is *not* configured. If you need the full text search feature, choose a custom installation.

Follow the product manual for detailed installation steps.

To help your installation process, we provide the input values we used during our installation in Table 8-18. Input fields are grouped by the input window. Substitute the sample input values according to your environment setup.

Table 8-18 Content Manager installation input values

Input window / field	Sample input value	Description
CM Version	8.3	
Installation directory	/opt/IBM/db2cmv8	Root directory of the Content Manager installation.
Installation type	typical	This includes the HTTPs configuration for the Content Manager internal communication. If you need to configure full text search, use Custom installation type.
Working directory (This option is not there if you choose typical installation type.)	/home/ibmcmadm	The working directory stores log files and configuration files generated when the product is running.
Host name	jamaica.almaden.ibm.com	The fully qualified network name of the server.
Database type	DB2	
Library Server database name	icmnlsdb	Library Server stores metadata in a relational database. This value specifies the name of the database.
Library server schema name	icmadmin	
Library server administration ID	icmadmin	This is the AIX ID you created earlier. It must be part of the db2grp1 group and the .profile of this user has to execute the DB2 profile.
Content Manager Connection ID	icmconct	Users are generally only defined in Content Manager (and not at the operating system level that hosts the Content Manager system). This ID is used to enable users to connect to the Library Server database. It only needs access rights to the database and no other privileges in Content Manager.
Library Server ID	1	

Input window / field	Sample input value	Description
Library Server transaction ID duration	180	
Enable UNICOD	Checked	
Enable text search	Checked	Even though text search is not used at the beginning, it should be checked if it will be used later on.
Enable LDAP	Not checked	
Resource Manager database name	rmdb	
Resource database administrator	rmadmin	This is the ID you created earlier for Resource Manager.
Resource Manager volume mounting point	/rmdb	In AIX it is a good practice to create a separate file system to install the database into.
Resource manager staging directory	/staging	Temporary working space for the Resource Manager (File system should be at lease 200 MB).
Application Server node name	Jamaica	
Resource Manager Web application name	icmrm	
Resource Manager Web application context root	/icmrm	
Token duration	172800	
Starting port for resource manager service	7500	
HTTP Port	80	Provide the following information for the Resource Manager application ports. If you enabled security for your WebSphere Application Server or WebSphere Business Integration Server Foundation, select the check box.

Input window / field	Sample input value	Description
HTTPS Port	443	
Use and configure IBM HTTP Server SSL	Selected	Select the SSL option for securing communication between the Resource Manager application and the system administration client. If you do not use IBM HTTP Server SSL, a default WebSphere SSL configuration will be used.

8.5.1 Installation summary and verification

Figure 8-4 shows the environment after Content Manager has been installed.

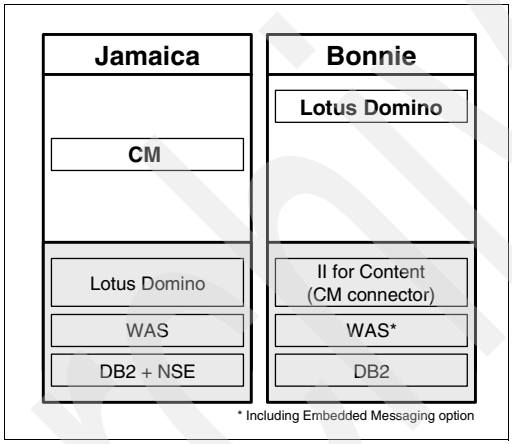

Figure 8-4 Sample environment after Content Manager installation

At the end of the installation, an installation validation utility will run. You should see the following message that indicates a successful installation:

```
Product validation completed with no detected configuration errors.
```

Perform the following steps to ensure that you installed Content Manager successfully:

1. Launch the system administration client that is automatically installed on the Content Manager server. In our scenario, we launch the system administration client from Jamaica.
2. Log on to the Content Manager system using the administrative user ID. In our scenario, we use icmadmin.

3. Open the Resource Manager configuration.
4. If the Resource Manager configuration is available, the communication with the Resource Manager is set up properly.

Perform the following steps as the final test of a successful installation:

1. Install a Content Manager client on a Windows machine.
2. Launch the Content Manager Windows client.
3. Import a text into the NOINDEX class of type text.
4. Retrieve the document immediately afterward. Make sure that you can retrieve it, open it, and view it on the screen.

8.5.2 Key information to remember

At this point, as part of your overall configuration setup documentation, write down the key configuration information that will be used during the rest of the software installation and configuration and future system operation.

Table 8-19 contains the key configuration input values to remember after the Content Manager installation.

Table 8-19 Key information to remember after Content Manager installation

Configuration data	Sample input value	Description
Content Manager administrator user ID	icmadmin	The administrative user ID for the Content Manager system.
Resource Manager Administrator ID	rmadmin	The Resource Manager database schema owner.
Content Manager server host name	jamaica.almaden.ibm.com	Fully qualified host name of the server that is running the Content Manager Library Server.
Library Server database	icmnlsdb	The name of the Library Server database. Every remote client has to catalog this database before being able to access it.

8.6 CommonStore installation and configuration

In the e-mail archiving and records management solution, CommonStore is used to archive e-mail from mail databases to Content Manager repository. It also provides a user interface to declare e-mail as records if manual records declaration and classification is allowed.

The steps involved include:

1. Install CommonStore for Lotus Domino (CSLD).
2. Configure Content Manager:
   a. Create the appropriate attributes and item type.
   b. Create the appropriate Content Manager user ID.
3. Create Domino user ID for CommonStore:
   a. Create a Domino user ID.
   b. Copy template files and sign them.
   c. Create and configure the configuration database and job database.
4. Configure ArchPro environment:
   a. Set the Content Manager connector environment.
   b. Create archint.ini.
5. Start ArchPro:
   a. Submit a license.
   b. Save a password for Content Manager user ID.
   c. Start ArchPro.
6. Configure CommonStore Task environment:
   a. Prepare notes.ini and names.nsf for CSLD Task.
   b. Set the environment for CSLD Task.
   c. Save the password for the Domino user ID.
7. Start CSLD Task.
8. Run CSLD in the background.

Features such as full text search and single instance store are not covered. Refer to *IBM DB2 CommonStore for Lotus Domino: Administrator's and Programmer's Guide* Version 8.3, SH12-6742 to set up those features.

Note: We do not include detailed steps of the installation in this section. We recommend using the existing product manual in conjunction with the materials we present here for a successful installation and configuration.

In the sample environment, we install CommonStore on Jamaica.

Step 1: Installing CommonStore for Lotus Domino

Select a complete installation to install all components (ArchPro Server, CSLD Task) on the machine.

Table 8-20 lists the input value used for installation in our sample environment. Replace the sample input value according to your environment setup.

Table 8-20 CommonStore installation input for the sample environment

Required input field	Sample input value	Description
Installation directory	/usr/lpp/csld	The installation directory in this sample environment is different from the standard installation directory, simply to have an easier way to manage Path and Classpath variables.

Important: If you are installing CommonStore on AIX 5.2 or higher, the Visual Age for C++ runtime libraries that come with CommonStore 8.3 for AIX are missing the libraries for AIX 5.2. In order to make this work, you have to install the correct version of the VACC runtime libraries for your OS level.

A temporary workaround is to install the VACPP that comes with CommonStore and create an aix52 link that points to the aix51 library path:

```
#cd /usr/vacpp/lib
#ln -s aix51 aix52
```

Step 2: Configuring Content Manager

Log on to the Content Manager System Administrator using the Content Manager Administrator ID. In the sample environment, the Administrator user ID is icmadmin (created during installation of the Content Manager; see 8.5.2, “Key information to remember” on page 312).

Perform the following steps when you are in the system administration client:

1. Create the appropriate attributes and item type (CSLDMail).
2. Create the appropriate Content Manager user ID (CSAGENT).

Creating the appropriate attributes and item type (CSLDMail)

Attributes are used to store metadata. In an e-mail environment, Notes fields (such as Subject or Sender) are mapped through CommonStore to Content Manager attributes.

To create the new attributes:

1. Select **Data Modeling** → **Attributes**.

2. Right-click **New**. Enter the appropriate information, and click **Save**.

   Create at least the mandatory attribute CSLDDocUNID. If CommonStore security is used, the attributes CSLDOrigUser and CSLDOrigDB must be created.

   Other attributes can be created in order to map additional Notes fields to Content Manager attributes.

   See Table 8-21 for the attributes added in the sample environment.

Table 8-21 Attributes

Name	Attribute type	Character type	Max length
CSLDDocUNID	Character	Alphanumeric	32
CSLDOrigUser	Variable character	Extended alphanumeric	254
CSLDOrigDB	Character	Extended alphanumeric	17
CSLDSubject	Variable character	Extended alphanumeric	254
CSLDFrom	Variable character	Extended alphanumeric	100
CSLDPostedDate	Time stamp	N/A	N/A

To create the new item type, CSLDMail:

1. Go to **Data Modeling** → **Item types**.
2. Right-click **New**.
3. Enter `CSLDMail` as the name of the item type.
4. Click the **Attributes** tab and assign attributes to the item type CSLDMail as shown in Table 8-22.
5. Document parts **ICMBASE** must be selected. If the item type must be text searchable, add **ICMBASETEXT** and check the **Text Searchable** option.

Table 8-22 Item type CSLDMail

Configuration field	Sample input value	Description
Name	CSLDMail	
ACL	ICMPublic	All users have public read access to that item type. Select **on Item level**.
Text search	unchecked	Text search is not configured.
Document part	ICMBASE	Only this document part is needed because no textuaries are configured.

Creating the appropriate Content Manager user ID (CSAGENT)

CommonStore uses a Content Manager ID to communicate with the Content Manager system. This user ID needs access to the e-mail stored in the DominoMail item type.

To create a new Content Manager user ID:

1. Select **Authentication** → **Users**.
2. Right-click **New**.
3. Enter `CSAGENT` as the user name and appropriate values. Click **Save**.

   Table 8-23 shows the input values we used to create the Content Manager user ID in the sample environment.

Table 8-23 Content Manager user ID

Configuration field	Sample input value	Description
User name	CSAGENT	ID used by CommonStore to archive content in Content Manager.
Password		During startup of the CommonStore Server (ArchPro), this password must be provided so that ArchPro can use the ID to log on to the Content Manager system.
Password expiration	Never expires	To ensure that CommonStore starts up correctly, make sure that the password never expires. If the password can expire, CommonStore startup will fail if Content Manager requests a password change.
Maximum privilege set	AllPrivs	The Content Manager user ID used by CommonStore has to be a Super User, which is necessary for the Records solution. Therefore, it is necessary to assign the AllPrivs privilege set.
Default item access control list	PublicReadACL	You must provide an ACL in this field; otherwise, the user ID cannot be created. However, this ACL is used only if, during an item type creation, the field “User’s default ACL” is chosen to be the ACL that defines the ACL to be set for items stored in the item type.

Step 3: Creating Domino user ID for CommonStore

To create a Domino user ID for CommonStore:

1. Create a Domino user (CSLD Task).
2. Set up the Domino user in the Notes client.

Creating a Domino user (CSLD Task)

Domino uses a Lotus Notes user ID to authenticate with the Domino server. This user ID needs access to all Notes databases that require archiving. This user ID also needs access to the CommonStore configuration database and job database (created later in this chapter).

To create a Domino user ID, start the Domino Administrator Client and log on with a user ID with proper rights to register new users and create a regular notes account. See Table 8-24.

Table 8-24 Input values for user CSLD Task

Configuration field	Sample input value	Description
User ID	CSLD Task	Lotus Notes user that will be used by CommonStore Tasks to authenticate with the Domino server. First Name: CSLD Last Name: Task
Password		Record this password, as it will be required later.

Setting up the Domino user in the Notes client

CommonStore uses a Lotus Notes user ID to authenticate with the Domino server. This user ID needs Editor access to all Notes databases that will use CommonStore archiving. This ID also needs Reader access to the configuration database and Editor access to the job database. The ACL of the job database must grant the role CSLD User to this user ID. This role allows an ID to see all jobs within the job database and not only the jobs created by an ID.

Use a new Notes client to set up the new Notes user ID. After the setup is complete and you can open this user's mailbox, open the personal address book of this user and remove all location documents except the Office location. Rename the Office location to be the name of the Notes user (for example, CSLD Task) and enable only the TCP/IP port. Close and open the Notes client and verify that it can still connect to the user's mailbox. Close the Notes client.

Attention: You will need the ID file, names.nsf, and notes.ini for this Notes user later in this chapter.

Step 4: Configuring ArchPro environment

To configure ArchPro environment:

1. Create an AIX user ID to run ArchPro.

   a. Create the group as shown in Table 8-25.

Table 8-25 AIX group

Configuration field	Sample input value	Description
Group Name	cmstore	This group enables the two CommonStore user IDs to transfer files.

   b. Create the user as shown in Table 8-26.

Table 8-26 AIX user to run Archpro

Configuration field	Sample input value	Description
User Name	ARCHPRO	User ID that will run the ArchPro program.
Primary Group	cmstore	Both the user ID that runs ArchPro and the user ID that runs the Tasks and Crawlers need to have this group as the primary group. This is because they read and write temporary files when archiving and retrieving content.
Group Set	cmstore, staff, ibmcmgrp	

   c. Set the user's password.

> **Note:** All new AIX user accounts require the password to be changed the first time they are used after being created. To do this, Telnet to localhost and log on as the new user. After you enter the password, you are required to change the password. Enter the same password as before and again to confirm it. If this is not done, if applications require this ID for authentication, the authentication will fail.

2. Set the DB2 environment.

   The ARCHPRO user ID has to run the DB2 profile to enable it to communicate with the Content Manager database.

3. Set the CommonStore environment.

   The ARCHPRO user ID has to set environment settings to enable it to access the CommonStore files.

4. Set the Notes environment.

   The ARCHPRO user ID has to set environment settings to enable it to use the Lotus Domino APIs.

5. Set the Content Manager connector environment.

   The ARCHPRO user ID has to set environment settings to enable it to communicate with Content Manager.

6. Create archint.ini.

   The archint.ini configures the ArchPro Server. It defines whether logging and tracing are activated, and logical archives are defined in archint.ini. Those logical archives point to a specific Content Manager server, including the item type that is used.

Example input for our sample environment

We performed the following sequence of steps for our sample environment:

1. Create an AIX user ID to run Archpro:

   a. Create the group:

   ```
   #smitty groups
   Add a Group
   ```

   b. Create the user:

   ```
   #smitty users
   Add a User
   ```

   c. Set user password

   ```
   #passwd ARCHPRO
   ```

2. Set the DB2 environment:

   a. Switch to the ARCHPRO user:

   ```
   su - ARCHPRO
   ```

   b. Add the DB2 profile to the end of this user's .profile:

   ```
   $vi .profile
   ```

   Example 8-1 shows the .profile we use for the sample environment, with the newly added line in bold text.

Example 8-1 .profile

```
PATH=/usr/bin:/etc:/usr/sbin:/usr/ucb:$HOME/bin:/usr/bin/X11:/sbin:.

export PATH
```

```
if [ -s "$MAIL" ]           # This is at Shell startup.  In normal
then echo "$MAILMSG"        # operation, the Shell checks
fi                          # periodically.

. /home/db2inst1/sqllib/db2profile
```

3. Set the CommonStore environment:
   a. As the ARCHPRO user, copy the environment script:

      ```
      $cp /usr/lpp/csld/bin/csenv.sh to ./
      ```

   b. Change the permissions so that the owner can execute it:

      ```
      $chmod 744 csenv.sh
      ```

   c. Add the environment script to the end of this user's .profile:

      ```
      vi .profile
      ```

      Example 8-2 shows the .profile we use for the sample environment, with the newly added line in bold text.

Example 8-2 .profile

```
PATH=/usr/bin:/etc:/usr/sbin:/usr/ucb:$HOME/bin:/usr/bin/X11:/sbin:.

export PATH

if [ -s "$MAIL" ]           # This is at Shell startup.  In normal
then echo "$MAILMSG"        # operation, the Shell checks
fi                          # periodically.

. /home/d1cmius1/sqllib/db2profile
. $HOME/csenv.sh
```

4. Set the Notes environment:
   a. As the ARCHPRO user, copy the environment script:

      ```
      $cp /usr/lpp/csld/bin/notesenv.sh to ./
      ```

   b. Change the permissions so that the owner can execute it:

      ```
      $chmod 744 notesenv.sh
      ```

   c. Add the last line to the end of this user's .profile:

      ```
      vi .profile
      ```

      Example 8-3 shows the .profile we use for the sample environment, with the newly added line in bold text.

Example 8-3 .profile

```
PATH=/usr/bin:/etc:/usr/sbin:/usr/ucb:$HOME/bin:/usr/bin/X11:/sbin:.

export PATH
if [ -s "$MAIL" ]           # This is at Shell startup.  In normal
then echo "$MAILMSG"        # operation, the Shell checks
fi                          # periodically.

. /home/d1cmius1/sqllib/db2profile
. $HOME/csenv.sh
. $HOME/notesenv.sh
```

d. Make a Notes data directory to contain the Notes user files:

```
$mkdir notesdata
```

e. From the data directory of the Notes user created in step 3b above, FTP the three files to this Notes data directory for ARCHPRO.

Example 8-4 shows the action and the results.

Example 8-4 FTP files to notes data directory for ARCHPRO

```
c:\Lotus\Notes\Data> ftp jamaica.almaden.ibm.com

Connected to jamaica.almaden.ibm.com.
220 jamaica.almaden.ibm.com FTP server (Version 4.1 Tue Jul 6 21:20:07 CDT
2004) ready.
User (jamaica.almaden.ibm.com:(none)): ARCHPRO
331 Password required for ARCHPRO.
Password:
230-Last login: Fri Jun 17 15:31:08 CDT 2005 on ftp from ::ffff:9.22.98.74
230 User ARCHPRO logged in.
ftp> cd notesdata
250 CWD command successful.
ftp> ascii
200 Type set to A; form set to N.
ftp> put notes.ini
200 PORT command successful.
150 Opening data connection for notes.ini.
226 Transfer complete.
ftp: 5087 bytes sent.
ftp> bin
ftp> put names.nsf
200 PORT command successful.
150 Opening data connection for names.nsf.
226 Transfer complete.
ftp: 4194304 bytes sent.
ftp> put ctask.id
200 PORT command successful.
```

```
150 Opening data connection for ctask.id.
226 Transfer complete.
ftp: 3424bytes sent.
ftp> bye
```

f. As the ARCHPRO user, modify the Windows directory references in the notes.ini file to be the AIX equivalents, as Example 8-5 shows in bold text.

Example 8-5 Updated notes.ini

```
[Notes]
Directory=/home/ARCHPRO/notesdata
KitType=2
SetupDB=setupweb.nsf
NotesProgram=/opt/lotus/notes/latest/ibmpow
TCPIP=TCP, 0, 15, 0
Ports=TCPIP
KeyFilename=ctask.id
NAMES=names.nsf
Timezone=5
DST=1
MailType=0
$$HasLANPort=1
DSTLAW=4,1,1,10,-1,1
PhoneLog=2
Log=log.nsf, 1, 0, 7, 40000
SHARED_MAIL=0
Location=CSLD,9AA,CN=CSLD TASK/O=ITSO
MailServer=CN=Bonnie/O=ITSO
MailFile=mail\ctask.nsf
```

5. Set the Content Manager connector environment:

   As the ARCHPRO user, add the Content Manager environment script and **umask** setting to the end of this user's .profile:

   ```
   vi .profile
   ```

   Example 8-6 shows the updated .profile, with the newly added line in bold.

Example 8-6 .profile

```
PATH=/usr/bin:/etc:/usr/sbin:/usr/ucb:$HOME/bin:/usr/bin/X11:/sbin:.

export PATH

if [ -s "$MAIL" ]           # This is at Shell startup.  In normal
then echo "$MAILMSG"        # operation, the Shell checks
fi                          # periodically.

. /home/d1cmius1/sqllib/db2profile
```

```
. $HOME/csenv.sh
. $HOME/notesenv.sh
. /opt/IBM/db2cmv8/bin/cmbenv81.sh
umask ug=rwx,o=r
```

6. Create archint.ini:
   a. As the ARCHPRO user, create the CommonStore instance directory:

      ```
      $mkdir inst000
      ```

   b. Copy the sample archint.ini from the install directory to the instance directory.

      ```
      $cd inst000
      $cp /usr/lpp/csld/bin/archint_sample_cm8.ini ./archint.ini
      ```

   c. Change the settings according to Table 8-27.

Table 8-27 Archint.ini setting changes

Parameter	Sample input value	Description
INSTANCEPATH	/home/ARCHPRO/inst000	
TEMPPATH	/tmp/csldxfer	Path where archived content is temporarily stored when being passed between tasks and archpro.
ARCHIVE	Email	Name of the logical archive that will be addressed by CommonStore Task. A Task is not aware of a Library Server or an item type, but only of the logical archive name. ArchPro defines logical archives in order to make them transparent for a Task in which archive system (Content Manager, Content Manager OnDemand, Tivoli Storage Manager) is used. The Task only refers to the logical archive or archives.
STORAGETYPE	CM	ArchPro supports more archive systems than just Content Manager. The value CM specifies that the archive is a Content Manager System.
ITEM_TYPE	CSLDMail	The item type used to archive e-mail. It was created in step 2a of 8.6, “CommonStore installation and configuration” on page 313.

Parameter	Sample input value	Description
LIBSERVER	icmnlsdb	The name of the Content Manager Library Server (DB2 database) in which the item type is created.
CMUSER	CSAGENT	Content Manager user ID used by CommonStore to communicate with the archive system. This ID was created in step 2b of 8.6, “CommonStore installation and configuration” on page 313.
ARCHIVETYPE	GENERIC_MULTIDOC	Default archive type for CommonStore

d. As the root user, create the TEMPPATH directory and set ownership and permissions:

```
#cd /tmp
#mkdir csldxfer
#chown ARCHPRO:cmstore
#chmod 774 csldxfer
```

Step 5: Starting ArchPro

When the configuration is completed, start the ArchPro server as follows:

1. Submit license.

   To submit the CommonStore license use this command:

   ```
   archpro -f license
   ```

   The ArchPro prompts for the path of the license file. This path must include the filename.

2. Save the password for CM user ID.

   In order to start ArchPro without user interaction, you must store the password of the CM user ID. Use **`archpro -f serverpasswd`** to be prompted for the password to be stored. The password will be stored encrypted in a file. Every time the password of the used CM user ID is changed, this process must be repeated. If the password is changed without repeating this process, ArchPro will fail to connect to the Content Manager because it is using an invalid password.

3. Start ArchPro.

   To start the ArchPro, use **`archpro`**. If any error occurs, the ArchPro will shut down immediately.

Example input for our sample environment

We performed the following sequence of steps for our sample environment:

1. Submit license:
   a. As the ARCHPRO user, change to the instance directory:

      ```
      $cd $HOME/inst000
      ```

   b. Register the CommonStore license:

      ```
      $archpro -f license
      ```

2. Save the password for the Content Manager user ID:

   ```
   archpro -f serverpasswd
   ```

   You will be prompted to specify the password for Content Manager ID CSAGENT (as configured in the logical archive definition in “Creating the appropriate Content Manager user ID (CSAGENT)” on page 316).

3. Start ArchPro:

   ```
   archpro
   ```

Attention: A successfully started ArchPro is necessary before the other steps can be accomplished. The ArchPro is the connection to the archive system, so it must be running before any archiving can be take place.

Step 6: Configuring the CommonStore Task environment

To configure the CommonStore environment, complete the following steps:

1. Create an AIX user ID to run tasks and crawlers:
   a. Create the user as specified in Table 8-28.
   b. Set its password.

Table 8-28 Create user

Configuration field	Sample input value	Description
User Name	CSLDTASK	User ID that will run the Task and Crawler programs.
Primary Group	cmstore	Both the user ID that runs Archpro and the user ID that runs the Tasks and Crawlers must have this group as the primary group. This is because they read and write temporary files when archiving and retrieving content.
Group Set	cmstore, staff, ibmcmgrp	

2. Set the CommonStore environment.

   The CSLDTASK user ID must set environment settings to enable it to access the CommonStore files.

3. Set Notes environment.

   The CSLDTASK user ID must set environment settings to enable it to use the Lotus Domino APIs.

4. Create archint.ini.

   The archint.ini configures the ArchPro Server. It defines whether logging and tracing are activated, and logical archives are defined in the archint.ini. Those logical archives point to a specific Content Manager server including the item type that is used.

5. Copy template files and sign them.

   The templates are used to create the job database and configuration database. The sample mail template is used to replace the design of a user's mail database to enable it to support archiving and retrieval of mail.

6. Create and configure the configuration database and job database.

   The job and configuration databases must exist on a Domino server. The configuration database holds all task and crawler configurations.

Example input for our sample environment

For our sample environment, we perform the following steps:

1. Create an AIX user ID to run tasks and crawlers:

```
#smitty users
Add a User
#passwd CSLDTASK
```

2. Set the CommonStore environment (use the CSLDTask user to copy the environment script):

```
$cd $HOME
$cp /usr/lpp/csld/bin/csenv.sh to ./
$chmod 744 csenv.sh
vi .profile
```

   Example 8-7 shows the updated .profile; add a new line as shown in bold text.

Example 8-7 .profile

```
PATH=/usr/bin:/etc:/usr/sbin:/usr/ucb:$HOME/bin:/usr/bin/X11:/sbin:.

export PATH

if [ -s "$MAIL" ]           # This is at Shell startup.  In normal
```

```
then echo "$MAILMSG"        # operation, the Shell checks
fi                          # periodically.

. $HOME/csenv.sh
```

3. Set the Notes environment (use CSLDTASK user to copy the script):

   a. As user CSLDTASK, copy the environment script and change permission of the script file:

      ```
      $cp /usr/lpp/csld/bin/notesenv.sh to ./
      $chmod 744 notesenv.sh
      ```

   b. Update .profile with the notesenv.sh information:

      ```
      vi .profile
      ```

      See the newly added lines in bold text in Example 8-8.

Example 8-8 .profile

```
PATH=/usr/bin:/etc:/usr/sbin:/usr/ucb:$HOME/bin:/usr/bin/X11:/sbin:.

export PATH

if [ -s "$MAIL" ]           # This is at Shell startup.  In normal
then echo "$MAILMSG"        # operation, the Shell checks
fi                          # periodically.

. $HOME/csenv.sh
. $HOME/notesenv.sh
umask ug=rwx,o=r
```

   c. Make a notesdata directory to contain the Notes user files:

      ```
      $mkdir notesdata
      ```

   d. Copy the three Notes user files from the notesdata directory of the ARCHPRO user to the notesdata directory of the CSLDTASK user:

      ```
      $cd $HOME/notesdata
      $cp /home/ARCHPRO/notesdata/notes.ini ./
      $cp /home/ARCHPRO/notesdata/names.nsf ./
      $cp /home/ARCHPRO/notesdata/ctask.id ./
      ```

   e. Change the ARCHPRO directory references in the notes.ini file to be the CSLDTASK equivalents (Example 8-9).

Example 8-9 Updated notes.ini file

```
[Notes]
Directory=/home/CSLDTASK/notesdata
KitType=2
SetupDB=setupweb.nsf
```

```
NotesProgram=/opt/lotus/notes/latest/ibmpow
TCPIP=TCP, 0, 15, 0
Ports=TCPIP
KeyFilename=ctask.id
NAMES=names.nsf
Timezone=5
DST=1
MailType=0
$$HasLANPort=1
DSTLAW=4,1,1,10,-1,1
PhoneLog=2
Log=log.nsf, 1, 0, 7, 40000
SHARED_MAIL=0
Location=CSLD,9AA,CN=CSLD TASK/O=ITSO
MailServer=CN=Bonnie/O=ITSO
MailFile=mail\ctask.nsf
```

4. Create archint.ini:

   a. As the CSLDTASK user, create the CommonStore instance directory:

      ```
      $cd $HOME
      mkdir inst001
      ```

   b. Copy the archint.ini file from the instance directory of the ARCHPRO user to the instance directory of the CSLDTASK user:

      ```
      $cd inst001
      $cp /home/ARCHPRO/inst000/archint.ini ./
      ```

   c. Update archint.ini as shown in Table 8-29.

Table 8-29 archint.ini update

Parameter	Sample input value	Description
INSTANCEPATH	/home/CSLDTASK/inst001	

5. Copy template files and sign them.

   CommonStore ships with three Notes template files. After the installation, those templates are located under the CommonStore installation directory.

   ```
   /usr/lpp/csld/data/CSLDConfig.ntf
   /usr/lpp/csld/data/CSLDJobs.ntf
   /usr/lpp/csld/data/CSLDStdMail.ntf
   ```

   A Domino administrator has to copy these files onto the Domino server and sign them to avoid unnecessary security prompts for the e-mail user.

   To make the templates available, copy them directly to the Domino server data directory.

Specify the Job database and Domino server in the CreateCSNJobs script library. Example 8-10 shows the modified lines in bold text.

Example 8-10 CreateCSNJobs.script

```
Public Function JobDatabaseName As String
   Dim session As New NotesSession
   Dim db As NotesDatabase
   Dim doc As NotesDocument

   Dim jobDBName As Variant

   Set db=session.CurrentDatabase

   ' The "CSLD Profile Document" is created with the "CSLD Wizard" database
   Set doc = db.GetProfileDocument("CSLD Profile Document")
   ' If the "CSLD Profile Document" exists and contains "jobDatabaseName"
   If (Not doc Is Nothing And doc.HasItem("jobDatabaseName")=True) Then
      jobDBName=doc.GetItemValue("jobDatabaseName")
      JobDatabaseName =jobDBName(0)
   Else
      ' Job database name. Users of this script library must modify this:
      JobDatabaseName="CSLDJobs.nsf"
   End If
End Function


Public Function JobDatabaseServer As String
   Dim session As New NotesSession
   Dim db As NotesDatabase
   Dim doc As NotesDocument

   Dim jobDBServer As Variant

   Set db=session.CurrentDatabase

   ' The "CSLD Profile Document" is created with the "CSLD Wizard" database
   Set doc = db.GetProfileDocument("CSLD Profile Document")
   ' If the "CSLD Profile Document" exists and contains "jobDatabaseServer"
   If (Not doc Is Nothing And doc.HasItem("jobDatabaseServer")=True) Then
      jobDBServer=doc.GetItemValue("jobDatabaseServer")
      JobDatabaseServer = jobDBServer(0)
   Else
       ' Job database server. Users of this script library must modify this:
      JobDatabaseServer="Bonnie/ITSO"
   End If
End Function
```

Sign the templates: Open the Domino Administrator client, select the **Files** tab, select the templates to be signed, right-click the template, and click **sign**.

Attention: The ID used to sign the templates must be listed in the Domino domain Execution Control List and be granted the appropriate access.

6. Create and configure the configuration database and job database.

   A configuration database stores all necessary information for a CommonStore Task, such as the task profile. The database must be created using the template that ships with the CommonStore Server. If an older version of a template is used, errors will occur.

   To create the configuration database, open a Notes client using an ID with proper rights to create a database on a server. Go to: **File** → **Database** → **New**. Select **Domino server**, define a name, and select the template.

   a. Set up a job database as shown in Table 8-30.

Table 8-30 Job database setup

Configuration field	Sample input value	Description
Server	Bonnie/ITSO	Name of the Domino server where the job database is created.
Title	CSLDJobs	Title of the job database.
Filename	CSLDJobs.nsf	File name of the job database created.
Server	Bonnie/ITSO	Name of the Domino server that contains the job database template file.
Template	CSLD Job Database 8.3	Name of the template to be selected.

   b. Add the Domino user ID used by CommonStore (`CSLD Task`) to the ACL and assign the role `CSLDUsers` to it.

   c. Set up the configuration database as shown in Table 8-31.

Table 8-31 Configuration database setup

Configuration field	Sample input value	Description
Server	Bonnie/ITSO	Name of the Domino server where the configuration database is created.
Title	CSLDConfig	Title of the configuration database.

Configuration field	Sample input value	Description
Filename	CSLDConfig.nsf	File name of the configuration database.
Server	Bonnie/ITSO	Name of the Domino server that contains the configuration database template file.
Template	CSLD Configuration Database 8.3	Name of the template to be selected.

d. Add the Domino user ID used by CommonStore (CSLD Task) to the ACL.

e. Set up the database profile for the archive task. See Table 8-32 for input values.

Table 8-32 Database profile for archive task

Configuration field	Sample input value	Description
Name	Archive	The profile name used to during the startup of a CommonStore Task.
Working DBs	All	`All` specifies that all jobs in the job database will be processed by this task.
Job database name	CSLDJobs.nsf	Name of the job database that this task is assigned to.
Job database server	Bonnie/ITSO	Server on which the job database is located.
Restrict retrieval to point of origin	No	CommonStore security is activated, so that only retrievals to the original database are possible. The Content Manager attribute CSLDOrigDB has to exist.
Task TCP/IP port	9000	
CommonStore TCP/IP port	47111	The port ArchPro server listens on for requests from CSLD Tasks. Defined in the archint.ini with the *variable DOMINOPORT.*
CommonStore host name	jamaica.almaden.ibm.com	This information is used to create the URL when HTTP links are created.

Configuration field	Sample input value	Description
CommonStore Web port	8095	This information is used to create the URL when HTTP links are created. Together with the CommonStore host name a URL will look like: http://jamaica.almaden.ibm.com:8095 This port has to point to the HTTP Dispatcher port, which is defined in the archin.ini with the parameter WEBPORT.
Folder Archive ID		Folder Archiving is not used in this environment, so the value is not set.

f. Set up the database profile for the retrieve task. Table 8-33 shows input values.

Table 8-33 Database profile for retrieve task

Configuration field	Sample input value	Description
Name	Retrieve	The profile name used during the startup of a CommonStore Task.
Working DBs	All	`All` specifies that all retrieve jobs in the job database will be processed by this task.
Job database name	CSLDJobs.nsf	Name of the job database that this task is assigned to.
Job database server	Bonnie/ITSO	Server on which the job database is located.
Restrict retrieval to point of origin	No	CommonStore security is activated, so that only retrievals to the original database are possible. The Content Manager attribute CSLDOrigDB has to exist.
Task TCP/IP port	9001	
CommonStore TCP/IP port	47111	The port ArchPro server listens on for requests from CSLD Tasks. Defined in the archint.ini with the variable DOMINOPORT.
CommonStore host name	jamaica.almaden.ibm.com	Used to create the URL when HTTP links are created.

Configuration field	Sample input value	Description
CommonStore Web port	8095	This information is used to create the URL when HTTP links are created. Together with the CommonStore host name, a URL will look like: http://jamaica.almaden.ibm.com:8095 This port has to point to the HTTP Dispatcher port, which is defined in the archin.ini with the parameter WEBPORT.
Folder Archive ID		Folder Archiving is not used in this environment, so the value is not set.

g. Set up document mapping. See Table 8-34 for input values.

Table 8-34 Document mapping

Configuration field	Sample input value	Description
Define mapping for	Document form	Specifies that a document mapping based on a form is created.
Notes form name	Memo	Specifies the name of form that is mapped to an logical archive.
Optional form aliases	Reply, Forward	Only mapped forms are archived. If only Memo is mapped, all other forms are not archived, so specify all other forms that are also mapped to the logical archive.
CommonStore Archive ID	Email	The name of the logical archive this document mapping is using. The value is defined in the archin.ini as shown in “Step 4: Configuring ArchPro environment” on page 318. The value is case sensitive.
Notes fields to display in hit lists	Subject, From, PostedDate	A hit list is created when CommonStore finishes a search within the Content Manager. The values define which values are shown in the columns of the hit list.
Form for result documents	Memo	The Notes form used when a document is selected in a hit list for retrieval.

Configuration field	Sample input value	Description
Notes document field names	Subject From PostedDate	List of Notes/Domino fields that are mapped to Content Manager attributes.
Archive attribute names	CSLDSubject CSLDFrom CSLDPostedDate	List of Content Manager attributes that are mapped to Notes/Domino fields. These attributes were created in “Step 2: Configuring Content Manager” on page 314. In this example, the Notes field Subject is mapped to the Content Manager attribute CSLDSubject.

h. Set up content type mapping. See Table 8-35 for input values.

Table 8-35 Content type mapping

Configuration field	Sample input value	Description
File extension	csn	Files with that extension are created when e-mail is archived in Notes Native Format.
Content type	csn/Application	

i. Configure the Lotus Notes password to bypass for CommonStore Task.

As the CSLDTASK user, store the password for the Lotus Notes user ID:

```
$cd $HOME/inst001
$csld -f serverpasswd -i $HOME/notesdata/notes.ini
```

Important: Every time the password of the Lotus Notes user ID is changed, this process must be repeated. If the password is changed without repeating this process, a CommonStore Task will fail to connect to the Domino server because it is using an invalid password.

j. As the root user, copy the extension manager add-in library file from the CSLD media to the CSLD bin directory:

```
#cd /usr/lpp/csld/bin
#cp <libextpwd.a from CDLD media> ./
#chmod 755 libextpwd.a
```

k. As the CSLDTASK user, add a line to the notes.ini file to make the extension available to the Notes API:

```
$cd $HOME/notesdata
$vi notes.ini
```

Table 8-11 shows the updated notes.ini file.

Example 8-11 Notes.ini file

```
[Notes]
Directory=/home/CSLDTASK/notesdata
EXTMGR_ADDINS=libextpwd.a
KitType=2
SetupDB=setupweb.nsf
NotesProgram=/opt/lotus/notes/latest/ibmpow
TCPIP=TCP, 0, 15, 0
Ports=TCPIP
KeyFilename=ctask.id
NAMES=names.nsf
Timezone=5
DST=1
MailType=0
$$HasLANPort=1
DSTLAW=4,1,1,10,-1,1
PhoneLog=2
Log=log.nsf, 1, 0, 7, 40000
SHARED_MAIL=0
Location=CSLD,9AA,CN=CSLD TASK/O=ITSO
MailServer=CN=Bonnie/O=ITSO
MailFile=mail\ctask.nsf
```

7. Start the tasks.

A CommonStore Task needs an up-and-running ArchPro. The startup command includes the following information: Domino server of the configuration database, name of the configuration database, name of the profile to be used, and the notes.ini file to be used.

Example input for the sample environment

For the sample environment, we perform the following steps:

1. Start the task using the following command:

```
csld -s <servername> -n <configdatabasename> -p <profilename> -i
<notesinifile>
```

See Table 8-36 on page 336 for input values.

Table 8-36 Parameters for starting CSLD Task

Parameter	Sample input value	Description
servername	Bonnie/ITSO	Name of the Domino server on which the configuration database is located.
configdatabase name	CSLDConfig.nsf	Name of the configuration database that contains the profile of the task. This database is created in step 6 on page 326.
profilename	Archive / Retrieve	Name of the profile in the configuration database as configured in step 6 on page 326.
notesinifile	$HOME/notesdata/notes.ini	Path including the filename to the notes.ini file that was copied and modified as described under step of the CommonStore installation and configuration chapter.

2. Start the Archive task.

   As the CSLDTASK user, start the Archive task:

```
$cd $HOME/inst001
$csld -s Bonnie/ITSO -n CSLDConfig.nsf -p Archive -i
$HOME/notesdata/notes.ini
```

3. Start the Retrieve task.

   As the CSLDTASK user, start the Retrieve task:

```
$cd $HOME/inst001
$csld -s Bonnie/ITSO -n CSLDConfig.nsf -p Retrieve -i
$HOME/notesdata/notes.ini
```

Step 8: Running CSLD in the background

To run CSLD in the background:

1. Create scripts using **nohup** to run the programs in the background.
2. Example 8-12 on page 337 shows a script that can start the archpro program either in the current session (default), in the background, or in debug mode (captures the console logging of stdout and stderr to files). Log on as user ARCHPRO and run this:

```
$cd $HOME
$./start_Archpro.sh
or
$./start_Archpro.sh -background
or
$./start_Archpro.sh -debug
```

Example 8-12 Content of /home/ARCHPRO/start_Archpro.sh

```
#!/usr/bin/ksh
cd $HOME/inst000
case $1 in
        -background)
                nohup archpro 1> /dev/null 2> /dev/null &;;
        -debug)
                nohup archpro 1> $HOME/Archpro.stdout 2> $HOME/Archpro.stderr
&;;
        *)
                archpro;;
esac
cd $HOME
```

l. Example 8-13 shows a script that can start the Archive task either in the current session (default), in the background, or in debug mode (captures the console logging of stdout and stderr to files). Log on as user CSLDTASK and run this:

```
$cd $HOME
$./start_ArchiveTask.sh
or
$./start_ArchiveTask.sh -background
or
$./start_ArchiveTask.sh -debug
```

Example 8-13 [contents of /home/CSLDTASK/start_ArchiveTask.sh]

```
#!/usr/bin/ksh
cd $HOME/inst001
case $1 in
        -background)
                nohup csld -s Bonnie/ITSO -n CSLDConf.nsf -i
$HOME/notesdata/notes.ini -p Archive 1> /dev/null 2> /dev/null &;;
        -debug)
                nohup csld -s Bonnie/ITSO -n CSLDConf.nsf -i
$HOME/notesdata/notes.ini -p Archive 1> $HOME/Archive.stdout 2>
$HOME/Archive.stderr &;;
        *)
                csld -s Bonnie/ITSO -n CSLDConf.nsf -i
$HOME/notesdata/notes.ini -p Archive;;
esac
cd $HOME
```

m. Example 8-14 on page 338 shows a script that can start the Retrieve task either in the current session (default), in the background, or in debug mode

(captures the console logging of stdout and stderr to files). Log on as user CSLDTASK and run this:

```
$cd $HOME
$./start_RetrieveTask.sh
or
$./start_RetrieveTask.sh -background
or
$./start_RetrieveTask.sh -debug
```

Example 8-14 [contents of /home/CSLDTASK/start_RetrieveTask.sh]

```
#!/usr/bin/ksh
cd $HOME/inst001
case $1 in
        -background)
                nohup csld -s Bonnie/ITSO -n CSLDConf.nsf -i
$HOME/notesdata/notes.ini -p Retrieve 1> /dev/null 2> /dev/null &;;
        -debug)
                nohup csld -s Bonnie/ITSO -n CSLDConf.nsf -i
$HOME/notesdata/notes.ini -p Retrieve 1> $HOME/Retrieve.stdout 2>
$HOME/Retrieve.stderr &;;
        *)
                csld -s Bonnie/ITSO -n CSLDConf.nsf -i
$HOME/notesdata/notes.ini -p Retrieve;;
esac
cd $HOME
```

8.6.1 Installation summary and verification

To test the installation, create a test Notes user. Replace its e-mail database template with the sample mail template that ships with CommonStore. Set the job database name and job database server value in that template.

In the sample environment, these values are:

Job database name CSLDJobs.nsf
Job database server Bonnie/ITSO

After the template is applied, make sure that the CSLD task has access to that database. You can test this by using user ID CSLD Task to open the mail database.

Open the mail database using the regular mail database user ID. Select an e-mail for archiving. Select **CommonStore** → **Archive Selected Documents** and click **OK**.

The archive is successful when the document moves to the Archived documents category within the inbox. To confirm that the retrieve is working, open the archived document and click **Fetch**.

If both operations are successful, the installation and basic configuration of CommonStore are completed. Figure 8-5 shows the components that are installed after this section is completed.

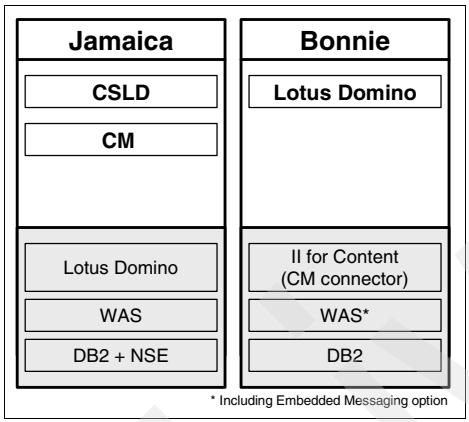

Figure 8-5 Sample environment after CommonStore is installed

Note: The various forms, views, and libraries from the CSLD sample template are not meant for production use but to be used as a guide to incorporating CSLD and Records functions into your corporate Notes template.

8.6.2 Key information to remember

At this point, as part of your overall configuration setup documentation, write down the key configuration information that will be used during the rest of the software installation and configuration and future system operation.

Table 8-37 contains the key configuration input values to remember after the CommonStore for Lotus Domino installation and configuration.

Table 8-37 Key information to remember after CSLD installation

Configuration data	Sample input value
Lotus Notes user ID used by CommonStore	CSLD Task
Content Manager user ID used by CommonStore	CSAGENT
Item type used by CommonStore	CSLDMail

8.7 Records Manager installation and configuration

In the e-mail archiving and records management solution, Records Manager is used as a records administration application and as an engine that records enables Content Manager (via Records Enabler for Content Manager) and thus records enables the entire e-mail archiving solution. In this section, we describe the main steps involved in installing and configuring Records.

These steps are as follows:

1. Install Records Manager engine V4.1.1.
2. Install Records Manager database V4.1.1.
3. Upgrade Records Manager engine V4.1.2.
4. Upgrade Records Manager database V4.1.2.
5. Run engine configuration utility.

Note: We do not include details of all of the installation steps in this section. We recommend using the existing product manual in conjunction with the materials we present here for a successful installation and configuration.

In the sample environment, we install Records Manager engine on Bonnie and Records Manager database on Jamaica.

Step 1: Install Records Manager engine V4.1.1

Make sure user irmwas is created using information in Table 8-38 as reference.

Table 8-38 Create user

Configuration field	Sample input value	Description
User Name	irmwas	This user ID is used for WebSphere connection factory authentication.
Primary Group	staff	
Group Set	staff	

Table 8-39 on page 341 lists the input values we used during our installation on Bonnie in the sample environment. Replace our sample input values as needed for your environment setup.

Table 8-39 Records Manager engine installation input for sample environment

Configuration window / field	Sample input value	Description
Installation Destination		
Directory Name	/opt/IBM/DB2RecordsManager	Records Manager installation directory. **Tip:** Use a directory without the version number. For updates, the same directory can be used so no second directory with a different version number will be created.
Installation Type		
Setup type	Custom	
Deployment and Configuration		
I want the installer to do deployment and configure for me	selected	The setup program deploys all J2EE applications and configures them.
WebSphere Application Server Connection Information		
Connector Type	SOAP	The type of communication interface between the WebSphere Application Server and the installation program.
Connector Port	8880	The port used by the Connector Type.
Cell	bonnie	The cell name of the WebSphere Application Server.
Node	bonnie	The node name of the WebSphere Application Server installed in 8.4.3, “WebSphere Application Server installation” on page 301. **Tip:** To view a node name, go to WebSphere Application Console; click **Servers** → **Application servers** → **Server1** → **Runtime**.
Server	server1	When this book was written, the only possible working value was server1.

Configuration window / field	Sample input value	Description
Security Enabled	unchecked	*Required* if WebSphere security is enabled. If security is enabled, this option has to be checked and a valid user name and password must be provided. To verify that security is enabled, ensure that WebSphere Application Server is started, open the WebSphere Application Console, and in the console navigation tree, click **Security** → **Global Security** to verify the **Enabled** setting in the **Configuration** tab.
Connection Factories Authentication		
Connection Factories Authentication User	irmwas	This AIX user ID has to exist. It does not need any special rights and is used just for authentication. The user ID must not be longer than eight characters (which is the current WebSphere limitation). An application component uses a connection factory to access a connection instance, which the component then uses to connect to the underlying enterprise information system (EIS). Connection examples include database connections, Java Message Service connections, and SAP R/3 connections.
Connection Factories Authentication password		
Mail Session Configuration		
Mail Transport Host	bonnie.redbook.bocaraton.ibm.com	The name of the server to access for the engine to send e-mail.

Configuration window / field	Sample input value	Description
SMTP User Name		This is required only if SSL is configured on the SMTP Server. Specifies the name of an e-mail user who has access to send e-mail through the specified transport host. Leave this field blank if the transport host does not require authentication.
SMTP User password		Only required if SSL is configured on the SMTP Server. *Required* for the engine to send e-mail. Specifies the password of an e-mail user who has access to send e-mail through the specified transport host. Leave this field blank if the transport host does not require authentication.
Administration Client Configuration		
Root	IRMClient	The context root for your Records Manager Administrator client. This is the name that you use to access the client for Records Manager in your browser (the virtual directory name). For example: `http://bonnie.redbook.bocaraton.ibm.com:9080/IRMClient`
Engine Server Name	bonnie.redbook.bocaraton.ibm.com	The host name of the computer where the Records Manager engine is installed.
Engine Server ORB Port	2809	Specifies the JNDI service port for the host where you are installing the Records Manager engine. This field specifies the port number on which the application server Object Request Broker (ORB) listens for requests.
Web Services Configuration		
Web Service Configuration Root	IRMWebServices	Specifies the context root for the Records Manager Web server. This is the name that is used to access the Web services for Records Manager in a browser (the virtual directory name).

Configuration window / field	Sample input value	Description
Web Services Node Name	bonnie.redbook.bocaraton.ibm.com	Specifies the host name of the computer where the Records Manager engine is installed.
Web Services HTTP Port	9080	Specifies the number for the port that the WebSphere Application Server uses for message queues.
Import Export Configuration		
Engine Server Name	bonnie.redbook.bocaraton.ibm.com	The host name of the computer where the Records Manager engine is installed.
Engine Server ORB Port	2809	Specifies the JNDI service port for the host where the Records Manager engine is installed. This field specifies the port number on which the application server Object Request Broker (ORB) listens for requests.
WebSphere Location		
WebSphere Location	/usr/WebSphere/AppServer	The WebSphere Application Server installation directory.

Step 2: Install Records Manager database V4.1.1

Make sure irmadmin is created using Table 8-40 as reference.

Table 8-40 Create user

Configuration field	Sample input value	Description
User Name	irmadmin	This user ID is used to connect to the Records Manager database.
Primary Group	ibmcmgrp	
Group Set	staff, db2grp1	

In Table 8-41 on page 345, we provide the input values we used during our installation on Jamaica in the sample environment. Replace our sample input values according to your environment setup.

Table 8-41 Records Manager database installation input for sample environment

Configuration window / field	Sample input value	Description
Installation Directory		
Directory Name	/opt/IBM/DB2RecordsManager/Database	This directory is not the default installation directory.
Database Type		
Database Type	DB2	
JDBC driver class path	/usr/opt/db2_08_01/java	Path to the *db2java.zip* file located under the DB2 installation path.
DB2 Database Configuration		
DB2 Node/Instance Name	db2inst1	The instance name if the Records Manager database is being installed directly on a DB2 server. If a remote database is being used, then this is the name of a cataloged DB2 instance. **Note:** This name cannot exceed eight characters in length.
Database Name	irmdb	The name of the DB2 database that is being created. **Note:** The database name cannot exceed eight characters in length, and it must be unique for each database you create.
Default Disk	/	The default location where the database and dataset files are being created.
Folder for Database container	irmdb	The default location where the database and dataset files are being created. This location will have the containers of the table spaces for the database to create.

Configuration window / field	Sample input value	Description
User name	irmadmin	Specifies the name of the DB2 user that will be the owner of the Records Manager schema. This user will have database administration privileges in the newly created database. **Important:** This user must exist, as it is not created automatically during the installation.
User password		
Territory	default	Specifies a portion of the locale mapped to the country code for the internal processing by the database manager.
Collating System	System	Specifies the sequence in which characters are ordered for the purpose of sorting, merging, comparing and processing indexed data sequentially.
DB Language	English	**Optional:** Specifies a language identifier. Set this field when the database language is different from the default language on your computer. The language you specify must be available on the computer where you are performing the installation.
System Administration User Name	db2inst1	Specifies the name of the database user with system administrator privileges for the DB2 database instance.
System Administration User Password		Specifies the password of the database user with system administrator privileges for the DB2 database instance.
Database File Plan Population		
Select a plan to populate database	Sample	A sample file plan is created.

Step 3: Upgrade Records Manager engine V4.1.2

The Records Manager engine upgrade is basically a redeploy of the WebSphere Application Server. When running the upgrade, re-enter the values you provided earlier.

Tip: Make sure to use the same installation directory; otherwise, a second directory will be created and the old one will not be deleted.

Step 4: Upgrade Records Manager database V4.1.2

After the Records Manager engine is upgraded to V4.1.2, upgrade the Records Manager database to the same level.

Table 8-42 Records Manager database upgrade input for sample environment

Configuration window / field	Sample input value	Description
Installation Directory		
Directory Name	/opt/IBM/DB2RecordsManager/Database	Use the same directory as specified during the original installation of the Records Manager database.
Custom or Automatic upgrade		
Automatic	Selected	
Database Type		
Database Type	DB2	
JDBC driver class path	/usr/opt/db2_08_01/java	
DB2 Database Configuration		
DB2 Node/Instance Name	db2inst1	DB2 Instance user name.
Database Name	irmdb	The name of the database that is created in step 2 on page 353.
Folder for Database container	/irmdb	The folder for the database container that is specified in step 2 on page 353.
User name	irmadmin	Name of the user that is specified in step 2 on page 353.
User password		

Configuration window / field	Sample input value	Description
System Administration User Name	db2admin	The name of the database user with system admin privileges for the DB2 database instance.
System Administration User Password		The password of the database user with system administrator privileges for the DB2 database instance.
Database Back Up		
Selected database was backed up	Selected	**Important**: If this is not selected, the upgrade cannot proceed. Ensure that the database is backed up before attempting to upgrade the database.

Step 5: Run engine configuration utility

The Records Manager engine must be able to access the Records Manager database. Because the database can be on different platforms (DB2,Oracle, SQL Server), a data source must be configured. This data source is used by the Records Manager engine to access the database.

If the Records Manager engine is running on a different machine than the Records Manager database, the database must be cataloged on the engine machine. In the sample environment, the engine runs on Bonnie, and the database is on Jamaica. We need to catalog the database on Bonnie. Use the DB2 Configuration Utility to catalog the database created in “Step 2: Install Records Manager database V4.1.1” on page 344. In the sample environment, the remote database on Jamaica (irmdb) is cataloged as irmdb on Bonnie.

The Records Manager engine configuration utility must be run on the machine where the Records Manager engine is running. To start the utility, export your display and run the script /opt/IBM/DB2RecordsManager/EngineConfiguration.sh

Table 8-43 lists input values used during the startup of the utility for the sample environment. Replace our sample input values for your environment.

Table 8-43 Engine configuration utility startup

Configuration field	Sample input value	Description
Connector Type	SOAP	
Port Number	8880	

Configuration field	Sample input value	Description
Cell	bonnie	The WebSphere Application Server cell on which the Records Manager Engine is deployed.
Node	bonnie	The WebSphere Application Server node on which the Records Manager Engine is deployed.
Server	server1	The WebSphere Application Server server into which the Records Manager Engine is deployed.

After the engine configuration tool is started, a data source (the Records Manager database created in "Step 2: Install Records Manager database V4.1.1" on page 344) must be created.

To create the data source, select **Action** → **New**.

In Table 8-44, we provide the input values we used when creating the new data source. Replace these values with the appropriate ones for your environment.

After creating the new data source, select **File** → **Save Changes**.

Table 8-44 Data source input for the sample environment

Configuration field	Sample input value	Description
Data Source Name	irmdb	Name of the database. You can use any name.
Database Name	irmdb	Name of the cataloged database.
User name	irmadmin	Name of the user created in step 2 on page 353 that has administrative rights for the database.
User password		

Important: After configuring the data source, if the utility is closed without saving, the provided information will be lost and the data source will not be available during the Records Manager Administration client startup.

Before the Records Manager Administration client can be used, the Application Server must be restarted.

As the root user, execute the WebSphere stopServer.sh and startServer.sh scripts:

```
# /usr/WebSphere/AppServer/bin/stopServer.sh server1
# /usr/WebSphere/AppServer/bin/startServer.sh server1
```

8.7.1 Installation summary and verification

At this point, the Records Manager engine and the Records Manager database are installed.

In the sample environment, the recommended scenario of installing them on two different computers is performed. Figure 8-6 shows the sample environment after the successful installation.

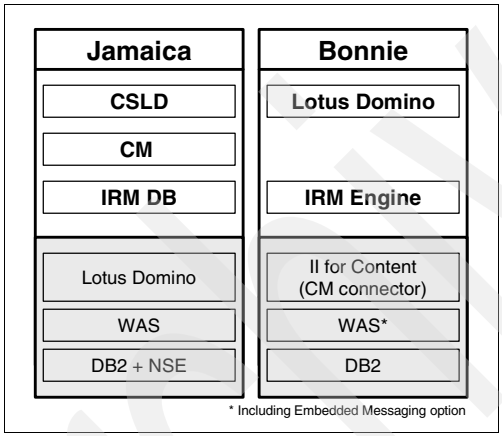

Figure 8-6 Sample environment after Records Manager is installed

To verify the installation, start WebSphere Application Console and go to **Servers** → **Application Servers** → **Server1** → **Message Listener Service** → **Listener Ports.** The following services should be listed and in a started state.

```
IRMCacheListenerPort
IRMTaskListenerPort
```

If either are not started, start them and verify that they are set to automatically start by clicking on them and verifying that Initial State is Started.

Important: The procedure above is very important. If the message queue is not started, you will encounter problems when working with Records Manager.

To make sure that Records Manager is installed properly, log on to Records Manager administration client using `Administrator` as the user ID and `cronos` as the password. Make sure you can log on.

You can also exercise a set of basic Records Manager activities to ensure that the Records Manager system is up and running. Suggested activities include:

1. Create a file plan.
2. Create a record.
3. Perform record scheduling.
4. Turn the crank.
5. Destroy record via retention rule.

8.7.2 Key information to remember

At this point, as part of your overall configuration setup documentation, write down the key configuration information that will be used during the rest of the software installation and configuration and future system operation.

Configuration data	Sample input value	Description
Records Manager administrator user ID	Administrator password: cronos	
WebSphere Application Server server name	server1	Server into which the Records Manager engine is deployed.

8.8 CMRE installation and configuration

Records Enabler for Content Manager (CMRE) is the bridge between Content Manager and Records Manager. It works with both products to introduce the records control capability into the Content Manager system. In this section, we describe the main steps involved in installing and configuring CMRE.

These steps include:

1. Set environment variables and create users.
2. Install Records Enabler (CMRE server, Host Interface server, and Permission Synchronization server).
3. Install Records Manager Extension.

Important: Make sure that the Content Manager V8 connector is installed. It is necessary to use the Information Integrator for Content installation to install the connector. A Content Manager client installation is not sufficient.

Note: It is not our intention to include all detailed steps of the installation in this section. We recommend using the existing product manual in conjunction with the materials we present here for a successful installation and configuration.

In the sample environment, we install Records Enabler on Bonnie.

Step 1: Set environment variables and create users

Set the environment with the following variables:

```
IBMCMROOT = /opt/IBM/db2cmv8
WAS_HOME = /usr/WebSphere/AppServer
JDBCPATH = /usr/opt/db2_08_01/java/db2java.zip
```

Substitute the values according to your system installation setup.

Important: The JDBCPATH must include the db2java.zip file name; otherwise, the installation will fail.

Create a local user on the machine that runs the Content Manager Library Server. This user must be in the DB2 administrator group. For the sample environment, create user cmreid using Table 8-45 as a reference.

Table 8-45 Create user

Configuration field	Sample input value	Description
User Name	cmreid	This user ID is used to connect to the Records Manager database. **Note:** All AIX user IDs used to authenticate with DB2 must be lowercase.
Primary Group	system	This user ID must belong to the system group (basically gives it root privileges).
Group Set	db2grp1	

Modify the user's .profile to include the following data:

```
IBMCMROOT=/opt/IBM/db2cmv8
export IBMCMROOT
PATH=$PATH:.:$IBMCMROOT/java/jre/bin:/usr/bin:$IBMCMROOT/bin
export PATH
. /home/db2inst1/sqllib/db2profile
```

Step 2: Install Records Enabler (CMRE)

You need to install CMRE. To help your CMRE installation process, we provide the input values we used during our installation in Table 8-46. Substitute the sample input values according to your environment setup.

Table 8-46 Records Enabler installation

Configuration field / window	Sample input value	Description
WebSphere deployment information		
DB2 Content Manager Records Enabler Server	Selected	
DB2 Content Manager Records Manager Host Interface	Selected	
DB2 Content Manager Records Enabler Permissions Synchronization	Selected	
WebSphere Application Server cell name	bonnie	Specify the cell name of the WebSphere Application Server installed in 8.4.3, “WebSphere Application Server installation” on page 301. **Tip**: To view the cell name for a server, start the WebSphere Application Server, open the WebSphere Application Console, and in the console navigation tree, click **Servers** → **Application servers** → **Server 1** → **Runtime**.
WebSphere Application Server node name	bonnie	Specify the node name of the WebSphere Application Server installed in 8.4.3, “WebSphere Application Server installation” on page 301. **Tip**: To view a node name, go to WebSphere Application Console and click **Servers** → **Application servers** → **Server1** → **Runtime**.

Configuration field / window	Sample input value	Description
Host name	bonnie.redbook.bocaraton.ibm.com	The fully qualified host name of the machine running the WebSphere Application Server.
WebSphere Application Security Enabled	unchecked	Required if WebSphere security is enabled. If security is enabled, this option must be checked and a valid user name and password have to be provided. To verify whether security is currently enabled, ensure that the WebSphere Application Server is started. Go to the WebSphere Application Console, click **Security** → **Global Security**, and verify the **Enabled** setting on the **Configuration** tab.
Records Manager Server configuration		
Records Manager Web services address	bonnie.redbook.bocaraton.ibm.com:2809	
Records Manager Administration Client URL	http://bonnie.redbook.bocaraton.ibm.com:9080/IRMClient	This value is checked during the installation. If it is not available, installation will not go further.
Records Manager database	irmdb	The name of the Records Manager database.
Records Manager Administrator	Administrator	Records Manager administrative user ID. The default ID is `Administrator`.
Password	cronos	The default password for the Records Manager Administrator is `cronos` if it is not changed after the IRM installation.
Content Manager Server configuration		
Server name	icmnlsdb	The name of the Library Server database.
Content Manager authentication	icmadmin	Administrative user ID for the Content Manager System. See 8.5.2, “Key information to remember” on page 312.
Password		

Configuration field / window	Sample input value	Description
Content Manager Records Enabler Connection ID	cmreid	This AIX user ID has to exist on the machine running the Content Manager Library Server.
password		
confirm password		
eClient rendering Content URL	http://....	In the sample environment, eClient is not installed. **Important:** Leave the default value and do not erase that field. With an empty field, the installation will fail. This value can be configured later on using the CMRE Administration client.
eClient document list URL	http://...	The eClient is not installed in the sample environment. **Important:** Leave the default value and do not erase that field. With an empty field, the installation will fail. This value can be configured later on using the CMRE Administration client.
Database System used for Content Manager	DB2	Specifies the database type of the Content Manager Library Server.
Content Manager Records Enabler configuration		
CMRE server	cmresvr	The server will be created during the installation if it does not exist.
Records Manager Host Interface server	rmecmhost	The server will be created during the installation if it does not exist.
Add Host Configuration record to DB2 Records Manager	checked	Checking this creates a “Host” entry within the Records Manager. The specified Content Manager system will be registered within the Records Manager.
Content Manager Records Enabler Permissions Synchronization	cmrepsproc	The server will be created during the installation if it does not exist.

Configuration field / window	Sample input value	Description
Permissions Synchronization Scheduler	checked	
Permissions Synchronization engine	checked	

Step 3: Installing Records Manager Extension

You need to install Records Manager Extension. To help your installation process, Table 8-47 lists the input values we used during our installation. Replace our sample input values as appropriate for your environment setup.

Table 8-47 Records Manager Extension installation input for sample environment

Configuration window / field	Sample input value	Description
WebSphere Deployment information		
WebSphere Application Server cell name	bonnie	Specify the cell name of the WebSphere Application Server installed in 8.4.3, "WebSphere Application Server installation" on page 301. **Tip**: To view the cell name for a server, start the WebSphere Application Server, open the WebSphere Application Console, and click **Servers** → **Application servers** → **Server 1** → **Runtime**.
WebSphere Application Server node name	bonnie	Specify the cell name of the WebSphere Application Server installed in 8.4.3, "WebSphere Application Server installation" on page 301. **Tip**: To view the cell name for a server, start the WebSphere Application Server, open the WebSphere Application Console, and click **Servers** → **Application servers** → **Server 1** → **Runtime**.
Host name	bonnie.redbook.bocaraton.ibm.com	The fully qualified host name of the machine running the WebSphere Application Server.

Configuration window / field	Sample input value	Description
WebSphere Application Security Enabled	unchecked	Required if WebSphere security is enabled. If security is enabled, this option has to be checked and a valid user name and password have to be provided. To verify whether security is currently enabled, ensure that the WebSphere Application Server is started, open the WebSphere Application Console, and in the console navigation tree, click **Security** → **Global Security**, and verify the **Enabled** setting on the **Configuration** tab.
Records Manager Application server name	server1	When this book was written, the only possible working value was server1.

8.8.1 Installation summary and verification

Figure 8-7 shows the components that will be installed and configured after this section is completed.

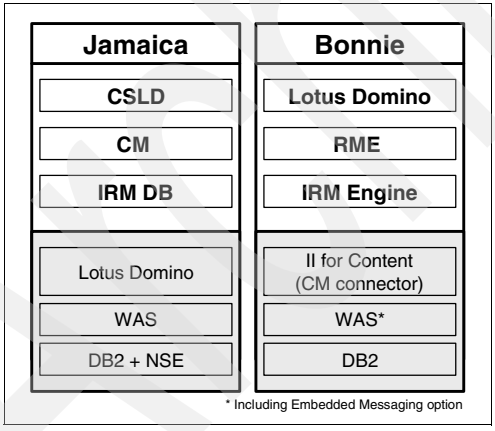

Figure 8-7 Sample environment after Records Enabler is installed

To verify your installation, perform the following steps:

1. Check that applications are running.
2. Import a Content Manager administrator ID into Records Manager.
3. Log on to the CMRE.
4. Records enable an item type.
5. Using Content Manager Windows client to declare a record.

Step 1: Check that the applications are running

To verify the installation, make sure all necessary services are running. They are:

- IBM WebSphere Application Server - server1
- IBM WebSphere Application Server - RMEServer (cmresvr)
- IBM WebSphere Application Server - RMEHostInterface (rmecmhost)
- IBM WebSphere Application Server - RMEPermSyncServer (cmrepsproc)
- IBM HTTP Server

Step 2: Import a CM administrator ID into Records Manager

To be able to log on to the CMRE Administration client, a Content Manager user ID that has the administrator rights must be imported to the Records Manager System, and the user ID needs administrator rights in Records Manager.

To import a Content Manager user:

1. Start the Records Manager client and log on as an administrator:

   User ID Administrator
   Password cronos

2. Go to **Security** → **Users** → **Host Filer** and select the host system that is enabled during the CMRE installation.

   In the sample environment, the host system with the name icmnlsdb is enabled. (See Add Host Configuration record to DB2 Records Manager of “Step 2: Install Records Enabler (CMRE)” on page 353.)

3. Click **Import**, select **icmadmin**, and click **Import** again. In the next window, select all permissions by checking **Function Access**. Check the **Is Active** check box. Click **Save** to finish the import.

4. The Content Manager user ID icmadmin is now imported to the Records Manager system and has all necessary rights to act as Administrator with the Records Manager system.

Step 3: Log on to the CMRE

Use icmadmin and its password to log on to the CMRE Administration client.

Important: Do not use the Records Manager administrator user ID (Administrator) and its password (cronos) to log on to the CMRE Administration client because this ID is not defined within Content Manager.

Also, do not use the Content Manager user ID, cmreid, that is created during the CMRE installation, because this ID is not (and cannot be) imported into Records Manager.

Step 4: Records enable an item type

After successfully logging on to CMRE Administration client, enable an item type:

1. Go to **Content Manager Server Configuration → eRecord enable item type**. This displays a list of all item types in the records enabled Content Manager system.
2. Select the CSLD item type (CSLDMail) that was created as shown in Table 8-22 on page 315.
3. Check the box to the left of the item type name, and select the record type to the right of the item type name.

In the sample environment, the item type `DominoMail` and the record type `email` are created.

Step 5: Using the Windows client to declare records

Using the Windows client, declare records as follows:

1. Search for documents in the records enabled item type DominoMail. The e-mail archived during the CommonStore verification is located in this item type.
2. Right-click one e-mail and select **Declare Record**. The Records Manager declare/classify window comes up.
3. Select a file plan and a unique name for that record. Click **Finish**.

The CMRE communicates with the Content Manager system and the Records Manager system. A Content Manager item type is records enabled, and documents within this item type can be declared as a records using the Content Manager Windows client. The records attributes (eRecord, eRecordID, FlPlnCmpntNm, FlPlnCmpntTtl, and RMEAclOri) associated with the e-mail are updated.

8.8.2 Key information to remember

At this point, as part of your overall configuration setup documentation, write down the key configuration information that will be used during the rest of the software installation and configuration and future system operation.

Table 8-48 contains the key configuration input value to remember after the Content Manager installation.

Table 8-48 Key information to remember after CMRE installation and configuration

Configuration data	Value
Records enabled item type	CSLDMail

8.9 Configuring the CommonStore Server and Notes

In this section, we describe the main steps involved in configuring CommonStore Server and Notes. These steps must be performed to add records declaration and classification capability to the end-to-end integrated e-mail solution.

The steps involved include:

1. Records enable Content Manager item type.
2. Prepare CommonStore.
3. Prepare Domino.
4. Records enable Notes client.

For detailed information about CSLD and Records Enabler integration, refer to Chapter 21, “Using Content Manager Records Enabler in the CSLD environment” in *IBM DB2 CommonStore for Lotus Domino: Administration and User’s Guide, Version 8.3*, SH12-6742.

Step 1: Records enabling Content Manager item type

This should be done in the previous section when you validate the installation and configuration of the CMRE. If it has not been completed, refer to “Step 4: Records enable an item type” on page 359.

Step 2: Preparing CommonStore

A Lotus Notes user must be mapped to a Content Manager user ID in order for the Notes user to declare a record because the back-end repository for records is the Content Manager system. This Content Manager user also has to be imported to Records Manager and must have the appropriate rights to declare records. For the mapping of Lotus Notes user IDs to Content Manager user IDs, an additional component (usermapper.jar) has to be activated on the CommonStore ArchPro server.

To activate the usermapper, follow these steps:

1. Copy usermapper.jar from the installation bin directory into the instance directory of the user that runs Archpro.

   Source directory /usr/lpp/csld/bin
   Target directory home/ARCHPRO/inst000

2. Update or create the CSExit.properties file.

   CSExit.properties must exist in the instance directory. It contains three parameters that must be configured: DB_DIR, HASH_MODULO, and PROXY_PORT.

   Set DB_DIR to the directory where you want to store the mapping database. This directory will contain a collection of files that are serialized hash tables

containing string keys of the format CM-server:mail-user and CSRepUserDef values. Make sure that this folder exists and is empty. The files will be generated automatically, along with a file that indicates the current HASH_MODULO value so that it can be changed.

Set HASH_MODULO to the maximum number of files to be found in the DB_DIR directory. This is to prevent the entire database from ever needing to be in memory at one time so that a huge number of users can be supported. Bigger values mean smaller memory usage (but more files).

Set PROXY_PORT to the port on which the usermapper proxy is listening.

Example 8-15 shows the values of these parameters in CSExit.properties.

Example 8-15 CSExit.properties

```
DB_DIR=/home/ARCHPRO/inst000/mapper_db
HASH_MODULO=100
PROXY_PORT=12345
```

3. Add usermapper.jar, csrepexit.jar, Notes.jar, and NCSO.jar into the classpath as well as the *directory* that contains CSExit.properties. To do this, edit the ARCHPRO user's notesenv.sh script.

   See Example 8-16 and Example 8-17 on page 362 as references.

Example 8-16 ARCHPRO users csenv.sh script

```
# set CommonStore environment
PATH=$PATH:/usr/lpp/csld/bin
LIBPATH=$LIBPATH:/usr/lpp/csld/bin
NLSPATH=$NLSPATH:/usr/lpp/csld/nls/%L/%N
CSNBASE=/usr/lpp/csld/nls/$LANG
CSNINSTANCEPATH=$HOME
CSNTAFDATAPATH=/usr/lpp/csld

CLASSPATH=$CLASSPATH:/home/ARCHPRO/inst000/usermapper.jar:/home/ARCHPRO/inst000
:/usr/lpp/csld/bin/csrepexit.jar

export PATH
export LIBPATH
export NLSPATH
export CSNBASE
export CSNINSTANCEPATH
export CSNTAFDATAPATH
export CLASSPATH
```

Example 8-17 ARCHPRO users notesenv.sh script

```
#!/bin/sh
#
# This script helps setting up the notes environment
# needed to run CommonStore for Notes Domino
#

# Set the PATH and LIBPATH environment for the user
# Adjust this for the local installation
PATH=$PATH:/opt/lotus/notes/latest/ibmpow:/opt/lotus/notes/latest/ibmpow/res/C
LIBPATH=$LIBPATH:/opt/lotus/notes/latest/ibmpow

# Add the users notes data directory to the path
PATH=$HOME/notesdata:$PATH

# Settings for the notes client
# Adjust this for the local installation
Notes_ExecDirectory=/opt/lotus/notes/latest/ibmpow
LOTUS=/opt/lotus

# Users notes data directory
NOTES_DATA_DIR=$HOME/notesdata

CLASSPATH=$CLASSPATH:/opt/lotus/notes/latest/ibmpow/Notes.jar:$HOME/inst000/NCS
O.jar

# Export all settings
export PATH LIBPATH Notes_ExecDirectory LOTUS NOTES_DATA_DIR CLASSPATH
```

Copy the NCSO.jar file from the data directory of the Domino server you installed for the APIs to the Archpro user's inst000 directory.

4. Update archint.ini (see Table 8-49).

   To activate the usermapper, add the values for the following parameters in the archini.ini file:

   ```
   ACCESS_CTL
   CM_SECURITY_EXIT
   CM_EXIT_LOCATION
   ```

   The input value for ACCESS_CTL YES specifies whether you want Retrieve operations to be subject to the user's Content Manager permissions.

   The input value for CM_SECURITY_EXIT specifies the name of the security exit class as com.ibm.rme.csexit.CSExit.

   The input value for CM_EXIT_LOCATION specifies the file location of the usermapper.jar file.

Table 8-49 archint.ini file update

Parameter	Value
ACCESS_CTRL	YES
CM_SECURITY_EXIT	com.ibm.rme.csexit.CSExit
CM_EXIT_LOCATION	'/usr/lpp/csld/bin/usermapper.jar'

Step 3: Preparing Domino

Prepare Domino as follows:

1. Modify the CSLD configuration database:
   a. Open the profile of the archiving task.
   b. Go to **Advanced**.
   c. In the section Write state to, select **Special field**.
   d. Accept the default value and save the profile.
2. Modify the template CSLDStdMail.ntf.

   The sample mail template is records enabled and must be configured:
   a. Open the CSLDStdMail.ntf with the Notes Designer.
   b. Go to **Shared Code** → **Script Library** → **RMEScriptLibrary** and provide the information as outlined in Table 8-50 on page 363.

Table 8-50 RMEScriptLibrary update

Parameter	Sample input value	Description
RMEServerURL_Default	http://bonnie.redbook.bocaraton.ibm.com:9082/RMEServer/RMEClientServlet	
CMHostName_Default	icmnlsdb	
CMItemTypeName_Default	CSLDMail	
UserProxyServerName_Default	jamaica.rebdook.bocaraton.ibm.com	
UserProxyServerPort_Default	12345	Defined in CSExit.properties
CSLDArchiveStatusField_Default	CSNDStatus	

Parameter	Sample input value	Description
RefreshInterval_Default	2	
RefreshTotal_Default	30	
RMEFolderClassifyTotal		This integer corresponds to the number of File Plan buckets that you want Notes users to be able to use.
WebServerPort	80	The port number for the HTTP server on the Domino mail server.

Attention: Changes to these variables in the template will not effect a database that has already had the template applied. This is because a configuration document is created the first time a user uses the database, and these values are not refreshed automatically. In order for a user to update these values, they need to clear the current values in their mail database by accessing the Records Configuration windows via their inbox. The next time they open the Records Configuration window, it will contain the updated values from their mailbox design.

Note: The various forms, views, and libraries from the CSLD sample template are not meant for production use but to be used as a guide to incorporating CSLD and Records functions into your corporate Notes template.

3. Set Domino security.
   a. Log on to the Notes administration client using an Administrator user ID.
   b. Create a new user group (RMEUserGroup):
      i. Go to the **People&Group** tab.
      ii. Click **Groups** and then **Add group**.
   c. Add Notes users who need to declare records to this group.
   d. Grant security to the group:
      i. Go to the **Configuration** tab.
      ii. Expand **Server** → **Current Server Document**.
      iii. Click **Security** and grant the following security permission:

```
Run unrestricted methods and operations
Run restricted Lotus Script/Java agents
Run simple and formula agents
Run restricted Java/Javascript/COM
Run unrestricted Java/Javascript/COM
```

e. Install the RMEAuth filter in the Domino server:

   i. Find the librmeauth_r.a file in the CSLD package, copy it to the mail server's Domino binary install path, and specify the owner and group to be the same as the other existing files.

```
/usr/lotus/notes/latest/ibmpow/librmeauth_r.a

owner = root
group = bin
permissions = -r-xr-xr-x
```

   ii. Install the RMEAuth filter by specifying the name of the filter in the Domino server record, in the DSAPI filter file name field in the Internet Protocols → HTTP table. You can specify just the name of the filter file (librmeauth_r.a) if it is located in the Domino program or data directories; otherwise you must specify the fully qualified path name.

   iii. Restart the Domino server.

Step 4: Records enabling Notes client

To records enable a Notes client, enable menus and buttons for the records management functions.

Complete the following steps:

1. Incorporate the CSLD functions into the user's e-mail template as outlined in the CSLD documentation.
2. Update notes.ini file to include the following line:

```
$RecordsEnabler=yes
```

3. Restart the Notes client.

The records management related menus and buttons will be available to use.

8.9.1 Verification

To verify that the CommonStore Server and Notes client are configured properly, you can follow these steps:

1. In Notes client, create an e-mail message and send it to yourself.
2. Manually declare the e-mail message as a record.
3. Make sure that the e-mail is archived in Content Manager using the Content Manager Windows client.
4. After the e-mail is archived, the Records Manager classification window should come up via a Web browser. Specify the bucket in the file plan, for

example Account Receivable in our sample environment. Enter a unique ID into the e-mail name field, and click **Finish**.

5. Make sure that the Notes client reflects that the e-mail is a record. (It can take anywhere from two minutes to an hour for the padlock icon to show because the RMEDeclareProgAgent mailbox agent must run.)
6. Using Content Manager Windows client, verify that the e-mail's isRecord attribute is set to yes.
7. Using Records Manager administration client, check that the record is there.
8. Use Notes client to view the e-mail message.

Part 3

Advanced topics

This part introduces some advanced topics to address when working with the integrated e-mail archiving and records management solution.

We cover the topics in the following chapters:

- Chapter 9, “Deployment considerations” on page 369
- Chapter 10, “Records Manager configuration and administration” on page 391
- In this chapter, we address records hold and disposition.
- Chapter 11, “Discovery” on page 409

9

Deployment considerations

This chapter examines the deployment of an integrated e-mail archiving and records management solution using CommonStore, Records Manager, and Content Manager from a planning perspective.

We review:

- Establishing a test system
- Piloting the system
- Configuration management in a production environment
- Post-system implementation

Information contained in this chapter should be used to supplement the processes your organization has for deploying software solutions.

9.1 Establishing a test system

One often-overlooked aspect of deploying a system to build a test system that can be used to test components of the production system without the risks associated with using a live system.

Too often, organizations do not set up a test system, and continue to make modifications to a production system that have not been thoroughly tested first. At best, these modifications may succeed, but more often than not, users or functionality will be disrupted and in the worst cases, systems may have to be restored or rebuilt.

We discuss the following aspects of the installation of a test system:

- Uses for the test system
- Mimicking the production system
- Change management

9.1.1 Uses for a test system

The use of a test system is important in a number of ways (some of the following ways assume that the production system is up and running in either its pilot or complete form):

- Configuration modifications that will affect any of the functionality or the user interface can be checked prior to rolling out to users.

  Some of the configuration changes that can be performed will be relatively minor (such as archive policy modifications where date changes are required or minor permission modifications in Records Manager) and may not necessarily require testing first. However, you might have a change management procedure (see 9.1.3, “Change management” on page 372) that dictates that *any* modification, however minor, must be tested on a test system first.

- Testing of code changes can be performed without risking the live system.

  IBM Content Manager, CommonStore, and Records Manager all have mechanisms to support extensive customization of the products. Usually through their APIs, this customization enables developers to write custom routines, user interfaces, or other code for these products. Use of the test system to trial customization such as this (assumes the code has already been developed and tested off the test system usually on a development platform) is critical to a full system test of the code prior to production implementation.

- New software versions can be implemented in the test system and thoroughly checked for interdependencies before being rolled into production.
  - IBM will issue version-specific fix packs from time to time and, as with all our software products, release notes will identify the reason for the fix and what steps should be taken to apply the fix.
  - With the IBM Content Management portfolio of products, IBM has made identification of product interdependencies clear in the product release notes. If one of the products in this portfolio has a new version, use of the test system will enable you to confirm correct operation of this new version in your own environment.
- Additional software may be added to the current implementation that must have interoperability testing carried out before production implementation.

  An example here is if you extend the current e-mail archive and records management system to include Document Management. This major modification to the system *must* be thoroughly tested in your environment before you deploy it in the production environment.

In the following sections, we include some of the issues to consider in the construction of a test system for the integrated e-mail archiving and records management system.

9.1.2 Mimicking the production system

As much as possible, the test system should mimic the production system.

Hardware

In most cases, this presents challenges to an organization. In almost every case, it will not be possible, cost-effective, or practical to set up your test system using the same scale of hardware (server hardware) as your production system. But unless performance testing (which should have been completed as part of an earlier implementation component) is something your organization does on an ongoing basis (which is unlikely), then equivalently specified hardware may not be important in your test system.

Multi-site environment

If you have operations in multiple sites, it may be a useful requirement for your test system to establish multiple servers across several geographically dispersed sites. In a Content Manager environment, multiple Resource Managers can assist in spreading the content load and placing it closer to regional user groups.

If this is the case, a test system that encompasses multiple servers and sites will test a number of components:

- Correct setup and behavior of the storage of objects onto correct Resource Manager.
- Correct access through organization WAN firewalls.
- Correct replication of objects to remote Resource Managers.

On a slightly different aspect, this multi-site test environment will also facilitate the testing of remote access and management of resources located off-site.

A similar concept could be used to test remote access to the system from client PCs.

Multi-machine environment

It is too tempting (from a resources perspective) to think that the test system is a system assembled from old pieces of hardware lying around the server room. After all, the cost of building a dedicated, multi-machine environment just to do testing may seem hard to justify. But many issues can be tested when your test environment is built using the same number of machines as your production environment.

One of the key benefits of this multi-machine environment is that communication between the machines can be fully tested using real-world protocols and IP addressing rather than relying on, perhaps, localhost as the machine IP addresses if all servers are co-located on a single machine.

Software

This aspect of the test system is extremely important, both for ensuring that the same applications (and versions) are installed on your test system as exist on the production system, and that the configurations between and within each software package mimic the production environment.

Much of the configuration used in the integrated e-mail archive and records management system can be exported.

9.1.3 Change management

Many organizations have effective IS change management procedures in place to ensure that system modifications are planned and implemented in a controlled fashion. Too often, modifications thought of as minor are implemented on a system without correct procedures being followed. These changes could introduce problems into the system. It is extremely important to ensure that these change control policies exist and that they are implemented effectively and every time.

Change management examples

Examples of good change management practices include:

- Production server shutdowns limited to out-of-business hours.

  Be aware that out-of-business hours on your site may still be business hours elsewhere, particularly where a specific server may be centrally located.
- All required configurations ranked according to priority, dependence (upon other systems), or criticality.
- All configurations require justification and sign-off.

  There is nothing like having to justify a change to your upper management to make you think of its consequences.
- All configuration changes require testing on a test system to confirm effects.
- All configuration changes require a roll-back plan in case of any unforeseen problems.
- All configuration changes of a significant nature must pass scrutiny by a testing board or a user group.

Extend procedures to a test system

Change management procedures must be extended to include the test system, both to encourage effective management techniques to the test system and, more important, to ensure that the test system is kept up to date and in sync with the production system.

If the test system becomes out of sync with the production system, its value as a mechanism to ensure that modifications made to the production should not cause any problems will be undermined severely, and confidence that changes tested on the test system will work on the production system will be lost.

There will be occasions where changes required on the production system cannot be fully tested on the test system because, for example, the hardware is different. But the benefits of at least testing those changes that do apply will help to ensure that the two systems remain synchronized.

Tiered testing systems

Some organizations that run critical or large systems apply a layered approach to their test systems. Each test system may differ slightly from the others. Most of the differences will be on the hardware side so that when configuration modifications are required, tests performed on different test systems may identify results that differ between different hardware configurations.

Many other organizations employ perhaps two systems in addition to the main system. The additional system is a development system on which initial

configuration or customization changes can be tested in an environment that is similar to (but not necessarily exactly the same as) the production system.

9.2 Piloting the system

With such a relatively complex system, piloting the system in some form will be crucial to the success of the deployment. Some organizations do not use the term “pilot.” Instead, terms such as *phased rollout* are used. Whichever terminology you adopt, one thing is clear: Deploying this system in a limited form will greatly assist a later, more extensive deployment. Some of the issues to consider are as follows:

- What goals must be set for the pilot system?
- What user group (or groups) will be used during the pilot?
- What configuration (archiving and records management) should be used during the pilot?

9.2.1 Goals of the pilot

Every pilot (or limited deployment) has to have a goal. Otherwise, how can you measure whether the pilot has been successful? This section lists some examples of pilot goals for the integrated e-mail archiving and records management solution.

System performance

Did the system performance under pilot loads meet or exceed expectations?

The purpose of a pilot is to test the system performance under a limited load. Performance can be perceived to be good or poor depending on your role. For an end user, the performance of the system has a major impact on the success of the project. If users experience poor system performance, then, regardless of the importance of the system, the project may fail. A pilot is a good opportunity to review performance with real users, but what are you measuring? With an integrated system such as e-mail archiving and records management, many aspects could affect the overall, perceived performance of the system; for example:

- If manual archiving is being used, how fast can a user archive an e-mail?

  When can the user get control of the machine?
- What interaction from the user is required?

  The more interaction, the slower the system is perceived to be.

- Is the functionality offered by the system too much of an overhead for the user or the tasks they need to accomplish?
- If an e-mail has not been archived and the user needs to declare it as a record, how long before the user gets the response from Records Manager asking the user to declare and classify?

  The declare and classify options do not appear until the system has received the e-mail or document's PID back from Content Manager; of course, this depends on this action's success and performance of the underlying Content Manager repository.
- If a user declares a record, what metadata is the user required to complete?
  - Requiring too much information to be entered may increase the user's frustration and affect the user's perception of overall system performance (even though the system may be idle, waiting for user's input).
  - Does the user need to drill down through a long or complex file plan hierarchy before reaching the record container to place the record into? (The use of IBM Records Manager functionality such as custom Views would assist here.)
- Are there any LAN or WAN performance issues?
  - Use of remote Content Manager Resource Managers to speed up local access to archived documents (bear in mind here that successfully implementing local or remote Resource Managers requires CommonStore to use multiple administrative IDs unique to each location so that the home location of an object can be identified and the relevant Resource Manager is used to store the content.
  - Is the WAN bandwidth shaped in any way that is likely to reduce performance of the new system? Often, organizations employ bandwidth shaping across their WAN to allow effective use of the bandwidth for existing systems. (An example is a transactional system: An order management system is given a higher priority across the WAN to ensure that a terminal user gets an adequate response). When a new system is introduced, these packet-shaping mechanisms (often based on open ports) must be reviewed to reflect the use of the new system.
  - Are system firewalls having any affect on remote traffic accessing the new system?

A records administrator, on the other hand, would probably be willing to accept certain overheads on the system (such as accurate completion of any metadata requirements) before users became frustrated with overall performance.

Because there are many components involved in the overall performance system, a good plan is to perform subsystem performance tests prior to rolling

the system out even to pilot users. Examples of IBM product resources to assist organizations with subsystem performance testing are:

- DB2:
  - *DB2 II: Performance Monitoring, Tuning and Capacity Planning Guide*, SG24-7073
  - *DB2 UDB/WebSphere Performance Tuning Guide*, SG24-6417
- WebSphere Application Server:
  - *Maximum Performance with WebSphere Application Server V5.1 on iSeries*, SG24-63833
- Content Manager:
  - *IBM DB2 Content Manager V8 Implementation on DB2 Universal Database: A Primer*
  - *Performance Tuning for Content Manager*, SG24-6949

Important: Be sure that the references above apply to your organization's environment. In particular, product versions used in these references may not match that of your organization. However, many of the principles may still be applicable.

Successful subsystem data upload tests

Were e-mail and documents successfully archived and declared as records?

Data inputs or uploads for each of the subsystems in this combined archive and records management system should be tested individually. This process ensures as much as possible that there are no issues when adding documents, e-mail, and records to each of the components. Even though your organization may be using a more automated approach to archiving (policy-based archiving) and records declaration (auto-classify and declare), manual testing of each of these components is important.

Prior to commencing any of the tests outlined next, you must design and document a useful set of use cases. These can be applied either during the subsystem tests or after end-to-end testing has begun. Without the thought and planning that this process requires, you may not have tested all components and systems effectively.

Examples of the subsystem tests that should be performed are as follows:

- Can e-mail be archived through CommonStore and stored in Content Manager?
  - Manually archive an e-mail and test that it can be viewed and retrieved back into the e-mail file (assuming that your configuration supports this functionality).
  - As an added test, use the Content Manager Windows client to view the archived e-mail directly from Content Manager.
- Can documents and attachments be archived and retrieved successfully?
- Can an e-mail be declared as a record and successfully classified?
  - After the e-mail has been declared and classified, security of the object will change once under the control of Records Manager so viewing of the object from within the user's e-mail file by the user that declared the record may not be possible if their access has been removed. If this is the case, an authorized user (one with at least View permissions within Records Manager) must use the Records administration client (RAC) to view that the record has been successfully classified.
- Can records be moved through their life cycle successfully?
  - This test should be performed on documents that use suitably short disposition schedules to allow successful disposition during the pilot period. One configuration that assists here is to set up a test disposition schedule that uses either an immediate disposition or a one-day disposition schedule.
- Does auto-classification work?
  - This test depends on how you have set up auto-classification. Auto-classification usually is initiated by the user (when they declare an e-mail) as they select Auto-Classify. Assuming the rules setup in Records Manager will function correctly with the limited amount of metadata available, the e-mail should be auto-classified.
  - To initially make this subsystem test a little easier, base the auto-classify rule on a simple text string that can be captured from the Subject attribute of e-mail.
- Can Records administrators manage the records management component of the system to effectively administer file plan and security components, to manage records disposition, and to classify records not classified elsewhere?
  - This test applies to a much smaller section of users but unless it is performed, your records administration staff may not be able to sign off on their aspect of the project.

– Multiple tests can be performed here and, as with all testing, some degree of training should have taken place prior to commencement of testing.

System usability

Were users able to perform their archiving and records management tasks effectively?

One of the main goals for any software system pilot is to ensure that users can use the system to perform their archive and records management tasks effectively.

Clearly, the system will have been configured (from both technical and business perspectives) before it is used, and this configuration will have had input from key business users. These users may have had an initial influence on the look and feel or usability of the system prior to its rollout to pilot users. This form of user input is important to ensure that basic usability (at least, the usability that can be configured) meets the requirements of the business.

Related to usability is the aspect of pilot user training. Even before the bulk of users are trained, pilot users must receive some training before they commence the pilot. The project manager should ensure that adequate training or education is carried out prior to pilot users gaining access to the system. This will streamline the pilot and enable users (both end users and administrative users) to concentrate more on testing the system without having to struggle with how they perform a particular action.

A final, related aspect of usability is support: How will pilot users obtain necessary support during the pilot? This aspect of the pilot is covered in a later section.

Post-pilot response

Although not specifically a pilot goal, an organization must plan for the end of the pilot. In other words, when the predefined pilot period has finished, what happens next? Whenever a limited deployment is planned, always have a very definite plan for how the pilot is to proceed after it is completed. Three scenarios must be considered here:

- The pilot succeeded and a larger implementation will commence.
- The pilot did not achieve its goals and will be extended.
- The pilot did not meet its goals and a larger implementation of the system under pilot will not take place.

Scenario 1: The pilot succeeded

In this instance and sometime before the end of the pilot, organizations must begin planning for a wider implementation of the system. Here are some of the issues that should be considered in this instance:

- What lessons can be learned from the pilot that will have an impact on the larger implementation?
  - Is any customization required in mail databases to respond to requests from users to make more effective use of the system?
  - Have any performance issues been addressed that will have an impact on usability?
  - Was the training adequate for most of the users?
  - Does the file plan meet the requirements of the users?
  - Are any automated, policy-based archive rules too harsh where a user's e-mail is archived too quickly?
  - Was the pilot user population suitably broad so as to effectively test most of the new system functionality?
- During the pilot, was any parallel running taking place where documents or e-mail were required to be accessed by both pilot and non-pilot users?

  This may influence any bulk data import planned for existing data.
- How much of the user population will require some or all of the functionality of the new system?

  This may influence:
  - Security permissions set in Records Manager.
  - Whether manual or policy-based archiving is set up in CommonStore.
  - How to segment any customization of e-mail databases.
- Within what time frame will the system be implemented to a larger group of users? Many factors could influence your decision:
  - Is any legislation having an impact on any of your user groups? If so this group (and its size) could affect implementation times.
  - It may be that any relevant legislation is driving a particular group's compliance, in which case this group may require access to the system earlier than other groups.
- Post-pilot consideration:
  - Not necessarily linked to a pilot but one aspect of implementing a new system that is often overlooked is that of return on investment (ROI). Even during a pilot, this important consideration should be considered. The project or business sponsor and project manager must review what key

measurements, indicators, or metrics should be examined to determine whether a pilot or implementation of a new archive or records management system has returned an investment for the organization. Examples of these indicators are:

- The system will offer an increased level of protection in key compliance areas such that the cost of the system, its implementation, and continued use are less than the cost of non-compliance. In this case, non-compliance could be the cost of legislative penalties imposed for not accurately managing corporate or organizational documents.
- The system will reduce the cost of any document discovery litigation that may be brought against the organization.
- The system will form a key or initial element in an enterprise content management system.
- The system will reduce the cost of physical or electronic storage of documents. (This last cost reduction should not be blown out of proportion; often reduction in storage costs could be quite minor. The media costs are relatively small when compared to the system implementation costs.)

– With all of these areas, the key is to find a *measurable* metric, one that can be quantitatively used as a benchmark. Often these metrics are time saved in person-days or dollar savings.

Scenario 2: The pilot did not achieve its goals and will be extended

In this instance, the pilot may have to be extended to all relevant pilot users to fully test all functionality and usability of the system. This may be because additional functionality was added to the initial build or configuration or that some customizations have been performed that need longer to test.

Either way, extending a pilot still has to be managed effectively. As mentioned earlier, there may be a requirement to continue with any parallel running that may be in progress; arrangements here may affect both pilot and non-pilot users.

Pilot user expectations should be re-aligned to the new pilot period and ensure that the business as a whole (project and business sponsors included) is kept informed of progress.

Scenario 3: The pilot did not meet its goals

Occasionally, a pilot will not succeed. If this case arises and the pilot does not succeed, plan how to roll back or what to change to make it a success. Data will be have been placed in the system that will have to be exported. If data has been placed in Content Manager, it is possible to export files to a file system outside Content Manager's control. Documents that have been archived through CommonStore to Content Manager and have subsequently been declared as

records will have their access secured. This must be considered when exporting the data.

An option that can be adopted during a pilot phase is to run in parallel with a manual records management system. This will minimize the disruption if a pilot phase is not extended into production, but it will increase the amount of work placed on pilot users. This method will succeed only where organizations are already running a paper-based recordskeeping system.

9.2.2 Pilot users selection

In the earlier chapter about concepts and solution planning, we made reference to the drivers that govern an organization's desire to implement a records management system. This will have an impact on which users to select in the piloting of this new solution.

A key factor influencing selection of users, beyond whether a particular group will require the functionality offered by records management or archiving, is whether the particular users are positive for change and flexible enough to participate in a pilot. It is only common sense to select your pilot users on the basis that they will address the goals of the pilot, have the time to participate, and would be willing to offer constructive input to your assessment of the pilot.

Uses affected by key legislation

If, for example, the legislation governing your organization's application of records management relates to SEC 17a-4, then the users affected will be anyone who is involved with share trading, dealing, or brokerage activities and is required to maintain records of these activities.

For the full text of the United States' Securities Exchange Act of 1934:

http://www.sec.gov/about/laws/sea34.pdf

Records management and archive users

If one of the purposes of the pilot is to examine both records management and e-mail archiving, then membership of the pilot user group will have to be selected carefully as there should be mechanisms in place to ensure that functions a particular user cannot access are not exposed or are suitably managed through user training. If, on the other hand, a staged approach is to be taken where one of the components (Records Manager for records or CommonStore for archiving) is implemented first, then users may have to be selected on a different basis.

Help desk team

Another very key group of users that must be included in any pilot are the staff supporting the help desk functions. These users must be involved in early discussions so that suitable responses can be issued to user questions during the pilot. Some organizations choose to establish a custom second-line or third-line support channel whereby any user support issues raised during the pilot will be fielded initially by first-line support. Upon discovering that the issue relates to the ongoing records management - archive pilot, they immediately refer the issue to a dedicated second-line or third-line support group or individual. Either way, the help desk team will be involved and must be suitably prepared.

Another group that should be included relatively early in the pilot is any staff involved in the initial training of pilot users. This group must have access to some key implementation decisions and may be involved in system configuration exercises. They must also be given time to prepare material or other training media that will be made available to pilot groups.

9.2.3 Configuration used during the pilot

As mentioned earlier, the specific configuration of the system for the pilot will depend on a number of factors:

- Will the pilot user subset require only part of the file plan to be available?
  - If only a subset of the file plan is required for the pilot, then this may suggest that configuration would be easier and may not necessarily require any user-specific or group-specific Records Manager Views to be constructed. However, it is important that the relative look-and-feel of the system be as similar to a full implementation as possible. If that means that a full implementation would normally include a fairly extensive file plan that would use Views to restrict the file plan to a manageable traverse for the user, then this configuration should be used during the pilot.
- Group and user permissions will be easier to set up because only those users accessing the pilot have be set up.
- With CommonStore for Lotus Domino, limiting the system to pilot users probably will require a second mail template to be used only for those pilot users (unless only Policy-based archiving is required). This task should be reflected in the project plan and the necessary change management procedures applied.
- With CommonStore for Exchange, if public folders are to be used, relevant, limited security must be applied to limit their access to pilot users.
- If the server hardware used during the pilot is likely to be used post-pilot, its configuration should be substantial enough to cater for the full system

implementation. Be aware though that this specification of hardware may well be totally overspecified for the limited pilot users.

This should be reflected in better performance for the pilot. Take this into account when assessing performance, as full load testing may not be possible until after the pilot.

9.3 Configuration management in a production environment

Although it is possible to treat some of the components of a integrated e-mail archiving and records management system as a black box where the software just gets on with its job, configuration of DB2, CommonStore, Content Manager, WebSphere Application Server, and Records Manager can be made relatively straightforward by implementing an effective configuration management system in addition to the change management system we covered previously.

Each of the products mentioned above has its own configurations, and certain aspects of these configurations are independent of other components. However, much of the configuration for each of these components should be treated holistically with the entire system. In other words, if modification is made to one component (for example, CommonStore field mapping), this could affect the other components (for example, affect attribute availability in Content Manager, which in turn could have an impact on attribute mapping in Records Manager). Security is another aspect of configuration that is particularly interdependent. Modification of key configuration components of the system can have a substantial affect on the functioning of the whole system. This is particularly important when the system is being used in production.

Effective management of configuration can therefore be seen as a critical component of system management. Although staff with the appropriate level of access can quite easily make changes to the system, it is hoped that change management procedures will limit the effect of any informal or unplanned configuration changes. However, what will greatly assist the whole process of making changes to the system will be a carefully planned configuration management system or process. Three areas that will assist you in developing an effective configuration management system are:

- Document the configuration early.
- Standardize configuration where possible.
- Apply good document control standards.

Document the configuration early

While either installing the software components or making initial configuration changes, it is too easy to not document what has been configured. For example, an early step with CommonStore is to enroll the CommonStore license. This requires both access to a file and the entering of a command, both of which should be documented. Revisiting a system without a completely documented configuration will be fraught with problems. An unknown configuration will add time to any solution modifications and potentially cause problems.

Some of the early installation and configuration of DB2 requires passwords to be entered. The planning of these must have happened already, and preinstallation planning worksheets are a very useful start to the configuration collection.

Below are some of the components or entries that should be documented as soon as configurations are made:

- Installation passwords (including those stored with any LDAP directories)

  These should be documented and kept secure, releasing them only on approval of any change management requests to approved staff. Changing compromised passwords is an overhead to be avoided.

- Installation directories

  Documenting these will aid in planning your backup and recovery strategy, and it is good practice is to place all software binaries in one directory.

- Machine architecture:
  - A visual representation of the current software and server placement is important. It can assist, when planning modifications, to know which machine hosts what component.
  - Annotate the architecture diagram to include IP addresses of the relevant machines and any logons if using remote or terminal access.
  - Include machine specifications to assist in any future hardware upgrades.

- Key users

  Not quite as important as other configuration areas, but will enable IS administrative staff to identify who to contact (possibly business owners of certain systems) in case of modifications or disruption to key systems.

Standardize the configuration

This may sound like a particularly obvious approach to take but many organizations treat the implementation of a new system as an opportunity to re-invent the wheel! Often guided by contractors or vendors, IS conventions that have already been established are put to one side or ignored when a new system is implemented.

A good example of this is the assignment of administrative user names. For example, your organization may require that all database administrator (DBA) names are unique to the organization and not the default standard in software applications. DB2 uses the user db2admin as the default DBA. Although this can be changed at a later date, it is simpler to apply your organization's convention as the software is installed.

Directories are another example where standardization on an organizational convention will make system maintenance easier. If administrators know that software can always be found in the \APPS\<software> directory, location of logs in troubleshooting will be easier. As a recommendation on Windows systems, avoid the \Program Files directory: It contains more characters than is useful, it is the default directory for many applications and it will soon become a confusing directory listing to navigate, lengthy directory names add length to environment variables, and often the maximum length can be reached, causing problems.

Standardization across development, test, and production environments is another aid to ease system maintenance.

When standard conventions are established, develop organization-wide processes and work instructions to ensure that these standards are maintained. These documents can then be referred to or referenced in change management requests for system modification.

Apply good document control standards

This practice goes without saying, but a set of configuration documents (no matter how detailed) will quickly become less useful if they are not properly maintained. Incorporating your document set into a set of work processes that become part of an overall change management process will go at least partway to ensuring that the document set is better controlled.

It will not be practical for just configuration documents on their own, but a good document management system such as IBM Document Manager (which can use IBM Content Manager as its repository and IBM Records Manager to control its records) would enable your organization to apply solid control mechanisms over the documents. Here are some of the features or standards that you should adopt in controlling your IS configuration document sets:

- Version control

  If multiple staff have access to change documents, keeping track of the current version of a document is extremely hard. Of course, access to the wrong version of a document can cause lost time spent on system changes.

- Security

  Unless only the right staff have access to configuration documents, elements of your configuration can become compromised; in addition to unauthorized modification to the document set, inexperienced or unauthorized users may be able to gain access to systems by referring to documents that normally should not be available to particular groups of staff.

- Review mechanisms:
  - If effective, these will ensure that any suggested changes are correctly reviewed, ideally by both business and skilled technical staff.
  - Business reviews should consider the impact the proposed changes will have on the day-to-day operations of the affected part of the business. If a proposed change (such as a version upgrade to a core product) is likely to mean that a system is down for a prolonged length of time, this may reduce the impact of any advantages the proposed change will bring.
  - A technical review performed by an appropriately skilled subject matter expert will allow consideration to be given to the broader impact of the proposed change. For example, a modification to the attributes of an item type may require that appropriate modifications are also made to full text index components.

Note the following points:

- If you are not going to install administration clients on servers, designate one machine as the master client machine and install them on that.
- Standardize the configuration.
- Establish a configuration document (spreadsheet) early on with absolutely *all* parameters used in the install.
- Duplicate Test and Production configuration.
- Establish a configuration management process early and enforce it.
- Establish a document control process (standards, versioning) and enforce it.
- Get the customer trained as early as possible.
- Assume that it will take longer than originally anticipated.
- It will be more complex.
- Know the import and export limitations of Records Manager.
- Allow for records declaration in the sizing of Content Manager.
- Applications will behave badly; prepare for this during the pilot.
- Knowledge of DB2 will assist greatly in understanding IBM Content Manager.

9.4 Post-system implementation

After the system has been implemented, there are a number of ongoing tasks that should assist in ensuring that the system continues to provide an effective archiving and records management system to all users. Many of these tasks could be considered as maintenance tasks and will include a number of soft tasks such as meeting with system users on a regular basis.

Regular meetings with business owners will enable the IS department to anticipate many of the issues that may occasionally occur, such as performance issues or changes to configuration, security, or personnel.

Other, perhaps more traditional, IS tasks under the maintenance heading will include performance monitoring, storage availability, and possible server outage recovery. Some of the tasks that should be reviewed are discussed in this section:

- Tuning the system
- Upgrades
- Disaster recovery

9.4.1 Tuning the system

Specific information about performance tuning was covered in an earlier chapter, but there are some performance tasks that can be applied to specific aspects of the system.

Much of the data in the system (much configuration data as well as all of the document metadata) is stored in a relational database system, in the sample environment, which is DB2. When you receive reports of poor performance, ask the following questions to quickly determine the best place to start looking for a potential cause:

- When did the problem first occur?

  If the problem has been occurring for some time, and if a database monitor schedule has been implemented, you can use historical data to find differences. This will enable you to focus on changes in system behavior and then focus on why these changes were introduced. On Windows-based systems, Performance Monitor can be used to begin to identify areas of concern.

- Is the performance issue constant or intermittent?

  If the poor performance is continual, check to see whether the system has started to handle more load or if a shared database resource has become a bottleneck. Other potential causes of performance degradation include increased user activity, multiple large applications, and removal of hardware

devices. If performance is poor only for brief periods, begin by looking for common applications or utilities running at these times, such as DB2 backup.

- Does the problem appear to be system-wide or isolated to DB2 and its applications?

  System-wide performance problems suggest an issue outside of DB2. It is possible that something at the operating system level should be addressed. If the poor performance is confined to DB2applications, we will focus on what DB2 is doing. If isolated to one application, is there a particular query that appears to be problematic?

  - Within CommonStore, *Single Instance Store* (SIS) is the mechanism used to ensure that for multiple documents e-mailed to a group of users who each decide to archive their copy, only one copy is physically archived. For this function to work efficiently, CommonStore checks a set of attributes (via their indexes) in the DB2 database to see whether the document already exists. This check has to be performed quickly and in real time. Non-indexed attributes or index problems can cause a delay.
  - If one application is problematic, then you can further examine whether users are reporting one or more queries that are experiencing a slowdown (such as the Single Instance Store–related issue). You might be able to isolate the issue to one application and a potential group of queries.

- Is there any commonality to the poor performance or does it appear to be random?

  You should determine whether any common database tables, table space, indexes, and so forth are involved. If so, this suggests that these objects are a point of contention. Other areas to potentially focus on include referential integrity constraints, foreign key cascades, and locking issues.

9.4.2 Upgrades

Upgrades to software systems should not be considered lightly. They should only be considered where security, major functionality, or incompatibility with other systems issues are a concern. When planning for an upgrade to any of the software packages, consider the following points:

- Can I justify this upgrade? Before performing any upgrades, they must be justified on a number of issues:
  - Does the business case warrant the expense of the upgrade? Upgrades are not a no-cost option even though the software may be part of a maintenance arrangement with your vendor. The cost of planning, testing, any business disruption, or performing the upgrades to the production system will all be considered as cost factors and must be outweighed by the benefits.

- Do the technical advantages offered by the upgrade outweigh the cost?
- What extra functionality will be offered by the new system and does this warrant the cost of upgrading?

- Is this upgrade compatible with other versions of your organization's software?

  Will the upgrade resolve any major security or compatibility issues? Your organization may be considering an upgrade to your operating system, such as a move to Linux®, where all of your key systems need to be upgraded.

- Plan the upgrade well in advance. Any upgrade of a major software system takes planning. Assuming that the upgrade is justified, consider these options:
  - Has the new version been thoroughly tested on the test system against the current release of each of the associated software systems?
  - Have alternative working options for users been planned to mitigate the disruption to the business during the upgrade?
  - Does all technical support staff have the necessary skills to manage the new version?
  - Can I roll back in case anything goes wrong during the upgrade?
  - Will the upgrade be backward compatible with any older versions of my software?

9.4.3 Disaster recovery

Although this book may not do justice to the large topic of disaster recovery (DR), the topic should be highlighted as one aspect of the overall solution. After all, the e-mail archive and records management systems could be considered as a critical system and one that should be available all if not most of the time.

Aside from a hardware-based DR solution, a number of components in our solution can be utilized to assist when considering a DR solution. For example, Content Manager's Resource Manager can be utilized, through its replication abilities, in an overall DR plan. DB2 offers a number of methods to incorporate into a High Availability and Disaster Recovery (HADR) plan. The following references should assist:

- *Disaster Recovery with DB2 UDB for z/OS®*, G24-6370
- http://www.ibm.com/software/data/db2/udb/hadr.html
- http://www.ibm.com/developerworks/ibm/library/i-hiavail1/

10

Records Manager configuration and administration

This chapter describes two key areas of records operations: holds and disposition. We introduce the key aspects of holds and the impact and results of applying holds and performing dispositions.

We cover the following topics:

- Records Manager configuration sequence
- Holds
- Records disposition

10.1 Records Manager configuration sequence

This section summarizes the sequence of steps needed to configure a basic Records Manager system. After these tasks have been completed, you will have the basis of a working system that will enable you to archive, declare, and classify e-mail. Refer to each product's planning and installation guide for a detailed review of all of the configuration steps.

Records Manager configuration includes the following steps:

1. Create a view.
2. Create components.
3. Establish relationships between components.
4. Add any required custom attributes.
5. Add instances of the file plan.
6. Import host users.
7. Assign permissions.
8. Set up life cycle information.

We use the following assumptions:

- All relevant software has been installed (including DB2 and Content Manager).
- The CommonStore system (either CSX or CSLD) has been installed and configured (including license enrolled, password for system users set).
- CommonStore can archive e-mail from either Exchange or Domino into the Content Manager repository.

After the Records Manager server and the Records Enabler for Content Manager are installed and validated, configuration of Records Manager (through the Records administration client) can be performed in the following sequence. (References to access indicate options to select from within the Records Administrator Client.)

Step 1: Create a view

A file plan view is a collection of relationships between components that comprise the file plan. In the same way that a view in a relational database is a collection of joined tables that comprise a schema, file plan views give each component in the file plan a context. No file plan components can exist outside a view. Every file plan component must be in at least one view.

Within Records Manager there can exist more than one view, and components can exist in multiple views. The view in this case is just a method of associating or grouping a collection of components together.

To set up, select **File Plan Design** → **Views** → **Add**.

Our sample environment view is ITSO.

Step 2: Create components

After a view (usually the primary view) has been created, components can be created and associated with the view. Each of the created components becomes a placeholder for an instance of one of the components. As an example, one of the components we created for our sample environment is Department. This particular component was specified when creating an instance of the component, such as Finance. The components will form the hierarchy of the file plan. Each component can be one of two types:

Container component This is one of the containers for records.

Record component This is the component type used for instances of a record. They differ from container components in that they can have content.

To set up, select **File Plan Design** → **Components** → **Add**.

Our sample environment components:

Department	Container component
Region	Container component
Division	Container component
email	Record component
Document	Record component

Step 3: Establish relationships between components

File plan relationship definitions govern how file plan components interact with each other. Two file plan components in a file plan relate to each other through relationship definitions. Defining a relationship definition specifies the source file plan component definition, the target file plan component definition, the view to which the relationship belongs, and the relationship capacity.

In other words, defining file plan component relationships enables the records administrator to build a hierarchy from the previously defined components. For example, in our example, the component Region is a child component to the component Department.

Be aware when developing relationships that a deep hierarchy (depending on how security and permissions are established) may frustrate users as they have to navigate through many levels of the hierarchy.

Figure 10-1 on page 394 shows the file plan hierarchy we set up in our scenario.

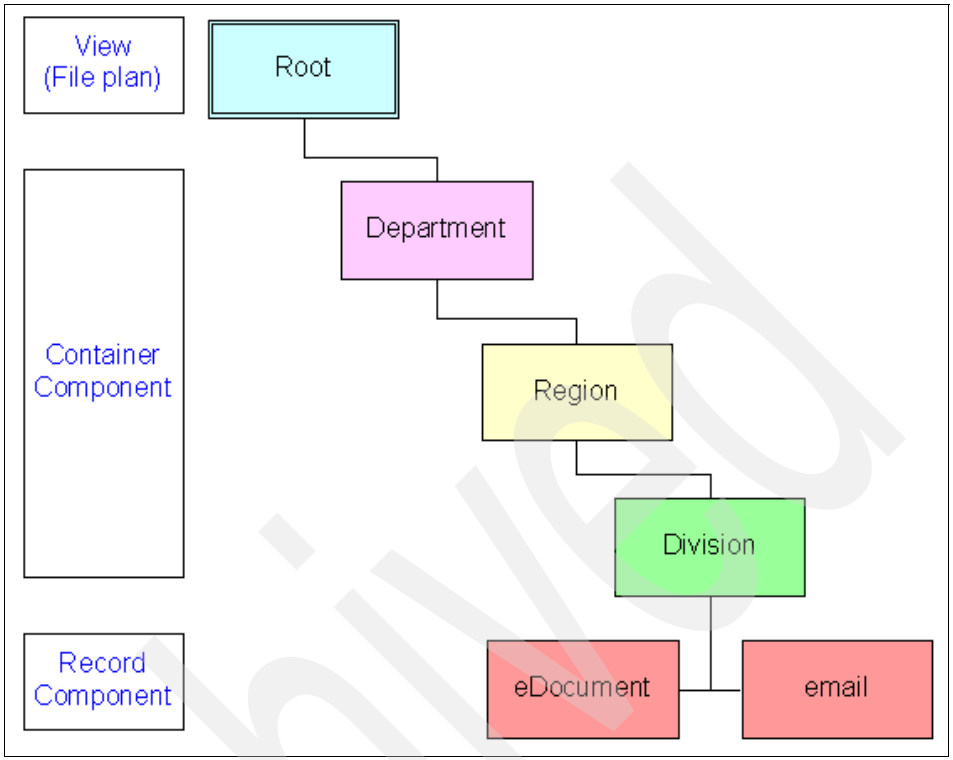

Figure 10-1 File plan used in the redbook

To set up within the Records Administrator Client, perform the following steps:

1. Select **File Plan Design** → **Views**.
2. Locate required View and select **View/Relationship Properties**.
3. For each component in the hierarchy, select **Add Relationship**.

Tip: It is good practice to start the name of each relationship with an integer to enable a sequential sorting of the list of relationships.

Step 4: Add any required custom attributes

As you design a file plan, specify the attribute definitions belonging to each file plan component definition. This includes specifying the name, the data type, the size, and the caption of the attribute definition. Different types of file plan components (known as file plan component definitions) have different sets of attributes that provide additional information about the file plan components. Note that adding custom attributes is not mandatory for a functioning system, but it is good practice to establish an effective and compliant records environment.

After you add an attribute definition to a file plan component definition, for every new file plan component of that type that is added to the file plan, users can add data into an attribute corresponding to the attribute definition.

Custom attributes can be added to both container components and record components. Custom attributes are added to an existing list of system attributes common against all file plan components. It is possible to restrict what attributes a user sees when classifying a record by using *Profiles*.

Custom attributes can be used to model your organization's metadata requirements for records management. An example of how they can be used is where you may be generating documents as a consultant for a particular client. You would like to specify the customer name against each document you classify (assumes manual classification) as a record so that at the close of the engagement, all documents can be processed together through the life cycle management functions within Records Manager.

To set up:

1. In the Records administrator client, select **File Plan Design** → **Components**.
2. Locate the component that you wish to add attributes to and click the **View/Edit Properties** action.
3. Within the selected component, click **Add**.

   No custom attributes were added for our scenario.

Step 5: Add instances of the file plan

After the file plan design has been completed and the relationships that form the hierarchy have been built, instances of each of the file plan components can be added. This task is performed using the records administrator client. In our sample scenarios, we created a file plan component called Department. We now need to create instances of this component.

Note: To create an instance of a file plan component, select the level above this component definition.

For example, using our sample file plan scenario of **ITSO** → **Department** → **Region** → **Division**, if you want to create different departments, you must select **ITSO** and then click the **Add** action. If you want to create different regions under a particular department (such as the Finance department), select **ITSO** → **Finance department** and then click the **Add** action.

To set up:

1. In the Records Administrator Client, select **File Plan Administration**.

2. Next to the primary view name (in our scenario this was called ITSO), click the **Add** action.

   Our sample environment file plan component instances:

View	ITSO
Departments	Finance, Sales
Regions	Europe, Asia, America
Divisions	Accounts Receivables, Accounts Payable

Step 6: Import host users

Users do not have to exist in the Content Manager repository for e-mail to be archived, as CommonStore uses a single administration ID to access and store archived documents. However, when a user declares an e-mail as a record (whether or not the e-mail has yet been archived), the user must exist in Records Manager or the host Content Manager system. Permissions are checked to ensure that record declaration is allowed.

Before Records Manager can be used, host users (users from Content Manager) must be imported into or registered with Records Manager. Additionally, the **IsActive** flag must be selected.

To set up:

1. In the Records Administrator Client, select **Security** → **Users**.
2. Select the CMRE host from the pull-down list.
3. Click **Import**.

Important: For Records Manager to see the host Content Manager users, Records Enabler for Content Manager must be set up correctly, and any users requiring access to declare records must be imported into Records Manager. You can add local users but not host users in Records Manager.

Note: More than one mailbox user can be mapped to a single Content Manager user ID.

Step 7: Assign permissions

After a new file plan has been created and populated and the users are available in Records Manager, assign component access permissions for the file plan components to the users or groups. You can assign the access permissions at either the system level or at the component level. If both permission levels exist, the component level permissions will have precedence over the system level permissions.

When assigning permissions, you must select to which system's users you are granting permissions. For additional clarity, we provide appropriate images to accompany the instructions.

To set up (permissions at system level):

1. In the Records Administrator Client, select **Security** → **System Permissions**.
2. Select one or more components to apply control then click **Permissions** (Figure 10-2).

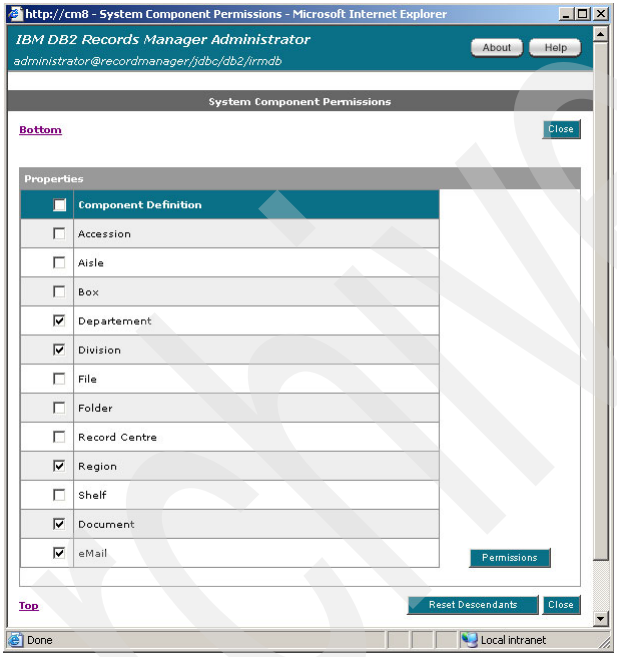

Figure 10-2 System permissions

Note: If individual components had their permissions set individually before this step, their permissions will be overwritten if multiple components are selected in this step.

Attention: Users must be able to traverse the file plan hierarchy to a point where a record component can be declared. Ensure that each level in the file plan hierarchy has been selected from the system permissions as shown in Figure 10-2 on page 397.

3. From the Host Filter list, select the relevant system as a source of users (in our case ICMNLSDB) and ensure that one or more permissions are selected from the list, then click **Add Users** or **Add Groups** (Figure 10-3).
4. Click **OK**.

> **Attention:** The *minimum* permission required is View. To enable the user to declare and retrieve records, the user must have Add, View, and HostRetrieve permissions on the file plans the user will declare into and retrieve records from. In addition, the user must be given File Plan Administration function access when importing the user.

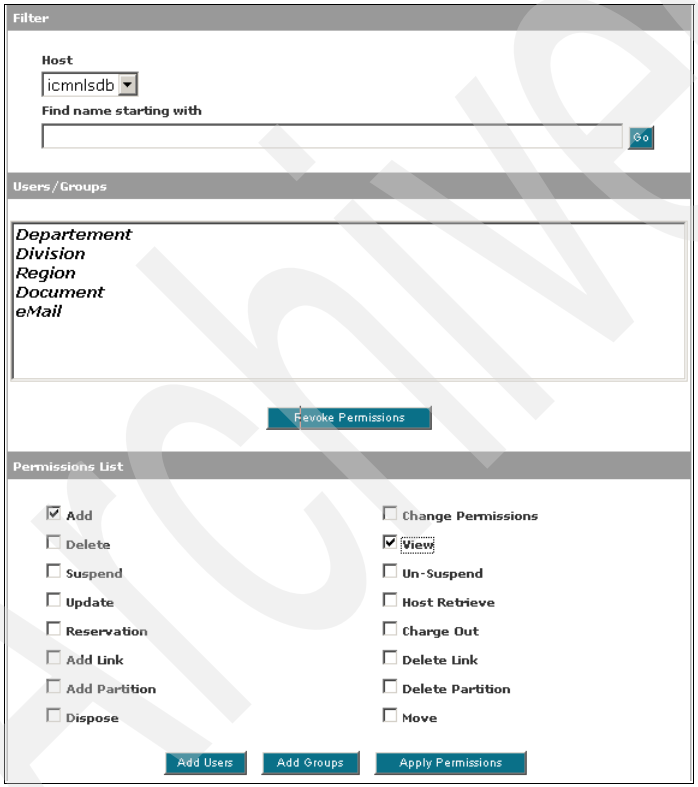

Figure 10-3 User/Group and Permission List showing selection

> **Note:** The previously selected components from the file plan hierarchy are shown (see Figure 10-3). At this stage they appear in the Users/Groups list but without any users identified. In a later step, these file plan components will indicate which users or groups have been given what access.

5. Highlight one or more users or groups then click **OK** (Figure 10-4).

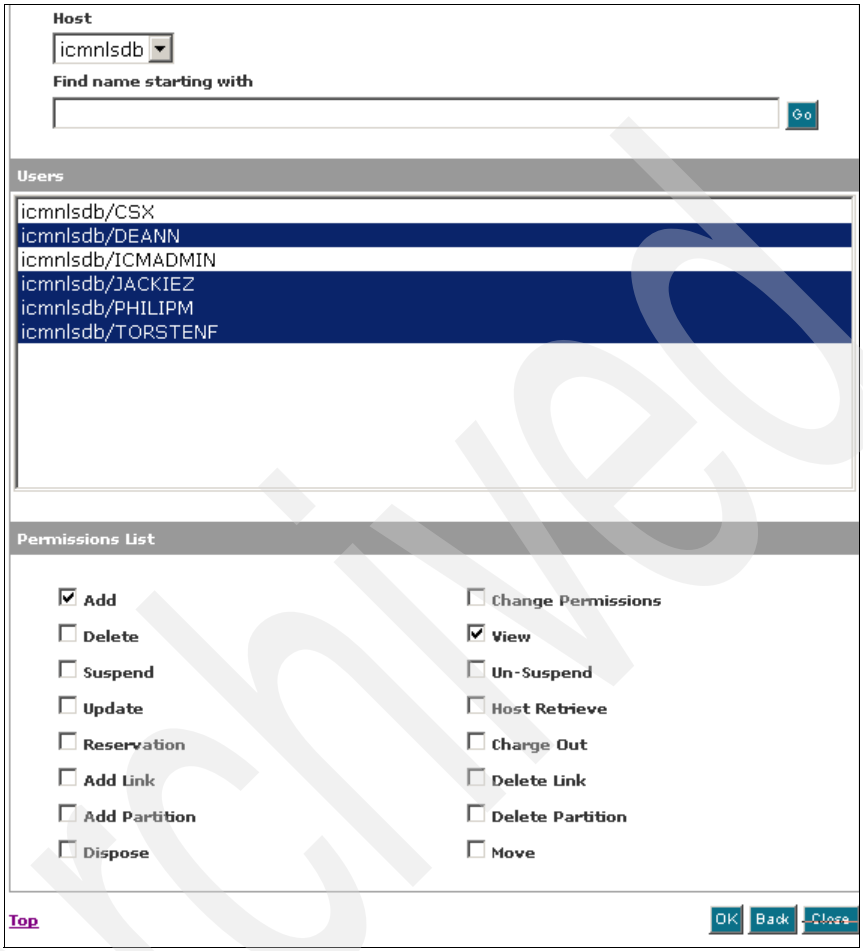

Figure 10-4 Selected users from the host system showing permissions given

6. The users and groups that were selected, plus their permissions and the file plan components where this combination of users and permissions have access, is shown (Figure 10-5).

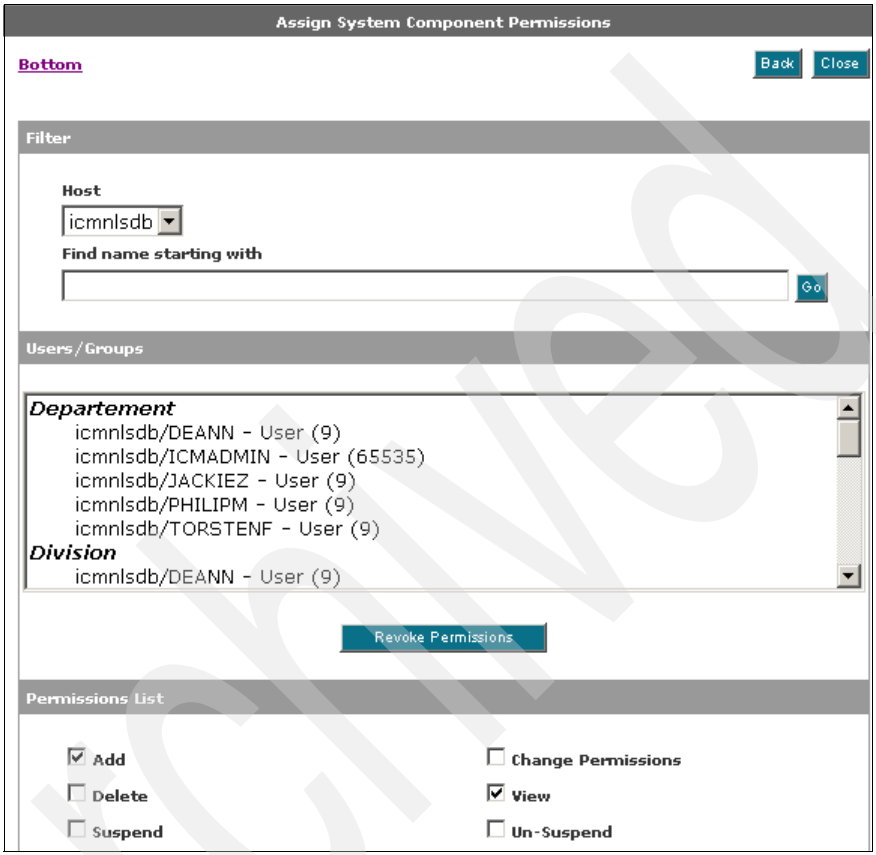

Figure 10-5 Completed permissions showing users and file plan components

Step 8: Set up life cycle information

A *life cycle* is a collection of phases through which any file plan component must transit before it is disposed.

There are two steps in the life cycle management process: designing a life cycle and performing life cycle operations. *Life cycle design* refers to the activities surrounding the configuration of an organization's life cycle and the rules that govern the life cycle operations. *Life cycle operations* refers to the process of executing the rules that govern the life cycle of components in your file plan.

Note: You must have the Life Cycle Management Design function access right to create life cycle phases and codes.

Designing a life cycle is also a two-step process: set up at least one life cycle phase (most organizations use two phases: active and dormant), then create a life cycle code that uses a combination of previously created life cycle phases.

These are all life cycle components:

- Life cycle phase
- Life cycle code
- Life cycle date

Life cycle code is set up in **Life Cycle Administration** → **Life Cycle Codes** and specified within each file plan component.

To set up life cycle component, use the following steps.

Note: The following sequence assists you in understanding the relatively complex components of life cycle management.

1. Create a life cycle phase (or phases) and specify a sequence in which these phases transition. For example, the first phase is `Active`, and the second phase is `Dormant`.
2. Create life cycle code. This really only creates a name for the code. For example, life cycle code can be `Active 5 - Dormant 10`.
3. Edit previously created life cycle code and specify how long a component will exist in any (or all) of the life cycle phases. For example, `Active 5 years - Dormant 10 years`.
4. Continue to edit life cycle code to select the trigger date of the selected component that will cause the component to begin time in any of the life cycle phases.
5. Select which code to apply to a particular file plan component.

For more information about Records Manager administration, refer to the product manual.

10.2 Holds

In records and legal circles, the terms tax, legal, litigation, and records holds are common. We use the generic term of hold.

A *hold* is an action taken on records collections to ensure that they are not dispositioned (deleted or archived) as part of their normal retention schedule life and are kept possibly beyond their scheduled date of destruction. For example, if a retention rule specifies that all e-mail records must be kept for three years before deletion, then any such record on hold (having a hold applied to it) that reaches that three-year period will not be (and cannot be) deleted. A hold may also have other direct impacts such as the freezing of information sharing such as educational records and transcripts triggered by a hold applied due to lack of payment.

Reasons behind holds

A lawsuit is served on a company and may or may not explicitly refer to holds. As part of a lawsuit, the court now commonly instructs the company to freeze its normal records destruction processes to ensure that discovery (searches) and evidence can be found. Regulatory auditors may order a company to apply holds as part of an investigation. Internally, due to discrimination or another issue or allegation, a company may determine that it must apply a hold related to a certain business process, customer, supplier or employee.

In recent years, the media worldwide has portrayed, and continues to do so, graphic terminal examples of companies and employees who either failed to apply holds and stop records destruction, who actively ignored hold orders and actively destroyed records (with commercials even portraying imaginary executives at night on the top floor shredding documents endlessly), or in more recent times failed (according to courts) to adequately forecast or expect to be placed under hold orders. Rapid erosion of market and customer confidence, multi-million dollar fines, jail terms, and company closings have resulted.

Note: In IBM DB2 Records Manager, the term and function *suspension* is used to represent and action all types of holds. In other words, suspension = hold.

Applying holds

Usually a company's legal department receives or determines independently to define and apply holds. The legal department does not usually apply holds onto records themselves but rather passes or directs the holds internally to records owners (business process owners) and, most commonly, to the records department and records managers. It is then the responsibility of the records staff to find and determine the records to apply the holds to, and ensure that such records are preserved and not destroyed if they reach the end of their life.

Organizations must be able to demonstrate that, as well as having a strong central records-keeping solution, they have consistent processes for receiving, applying, and removing holds.

Holds may have a duration specified or a date for next review. This enables the records staff to review along with the legal staff the future review status of the records and holds and any change in status that should occur.

Holds may be applied immediately. As in all things records-related, one size and practice does not fit all. Some records staff with mature records processes and systems can search and apply any number of holds across their whole records collections, physical and electronic records, to ensure that they are identified as on-hold and will never be dispositioned. Other records staff can choose to add the hold review as part of the end-of-life disposition process. In this way they can review any candidate records that are due for destruction to see whether they should be held and not destroyed, although they can end up doing more work later on in the process.

Hold target

Records staff usually intimately know their retention schedule and file plans. There are two main ways to apply holds. One way is to browse to a file plan location or container; for example, for a particular matter or case or customer. Another is to search across all records for related terms, keywords, or metadata that make them candidates for the hold; for example, hold all financial trade records.

Hold length

Holds do not last forever. Most records are not kept forever. At some point, days, weeks, months or years later, the hold can be removed, when the legal or other issue is closed or settled or an internal investigation is completed. Whatever the reason, usually the legal department can determine and then instruct the records department to remove the hold and continue the standard retention schedule on those records.

Sometimes the same retention will remain and in the future the records will expire as planned. If the hold had been applied and removed after the original retention date, then those records could immediately become candidates for destruction. Occasionally the legal, records, and business process may decide to apply a new retention on the records that were previously on hold, in effect extending the retention. The key is to follow documented record processes and practices to achieve this. To be able to demonstrate to courts, consistency is key.

One of the most important aspects of a company's records hold process is, from a legal perspective, to demonstrate a high level of documented consistency in applying holds and, later, disposition.

Figure 10-6 shows the stages that lead to a hold.

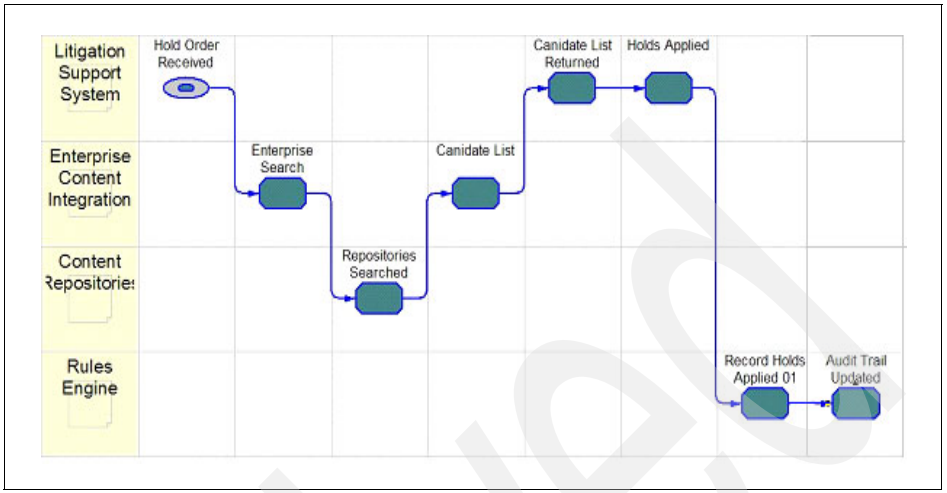

Figure 10-6 Hold stages

IBM DB2 Records Manager provides comprehensive economic hold functionality. Records Manager can:

- Be the central point to define, track, apply, and remove all corporate holds.
- Provide immediate browse and search options to apply holds.
- Support an unlimited number of holds that can be applied across any record types in the record collection.
- Support tracking of multiple holds on the same record.
- Provide reporting of which holds are applied on which records.
- Audit consistent application and removal of holds.
- Provide economic hold support by facilitating the hold application up front, which enables Records Manager to filter any held records from any disposition lists.
- Provide secure hold functionality, so that only designated staff can review, apply, and remove holds (not all holds should be seen by all staff).

10.3 Records disposition

Disposition defines what happens to the records at the end of their retention, a process commonly referred to as *records scheduling*. Most records are destroyed at the end of their life, but not all. Records Managers may use various phases or disposition codes or references to convey the handling required.

10.3.1 Disposition options

Records Manager provides the full range of disposition options as follows:

- Destroy
- Accession
- Review

Destroy

Destroy is the confidential deletion of the records content and metadata. The records term for this is *expunging* (irrevocably deleting the record so that not even document forensics can recover any aspect of it). Electronic records are usually overwritten at the disk bit level (and independent validation can be demonstrated such as provided by gaining US DoD 5015.2 STD certification). For physical records these were usually burnt or shredded but nowadays due to environmental concerns they may be destroyed in acid baths.

Accession

Accession is sometimes called archiving. The records are no longer tracked or kept in your records system but their provenance (history) is maintained as they are passed to some other records holding authority, maybe to the corporate archives. Note that the electronic copies records should be destroyed from your records system after the transfer.

Review

Review highlights the records at the end of their current defined life and enables the records staff and organization to review and possibly change the retention or period of re-reviewing the disposition of the records (for example, in a year's time).

10.3.2 Records scheduling

How often does this structured records review and disposition occur? Generally, records scheduling is performed only once a quarter or once a year. This was the common schedule for dealing with physical (paper) records. This resulted in a sizable collection of records to be reviewed and processed just for physical paper

and objects in boxes. Now with e-mail counted in the millions per day, a more scalable records scheduling solution is critical.

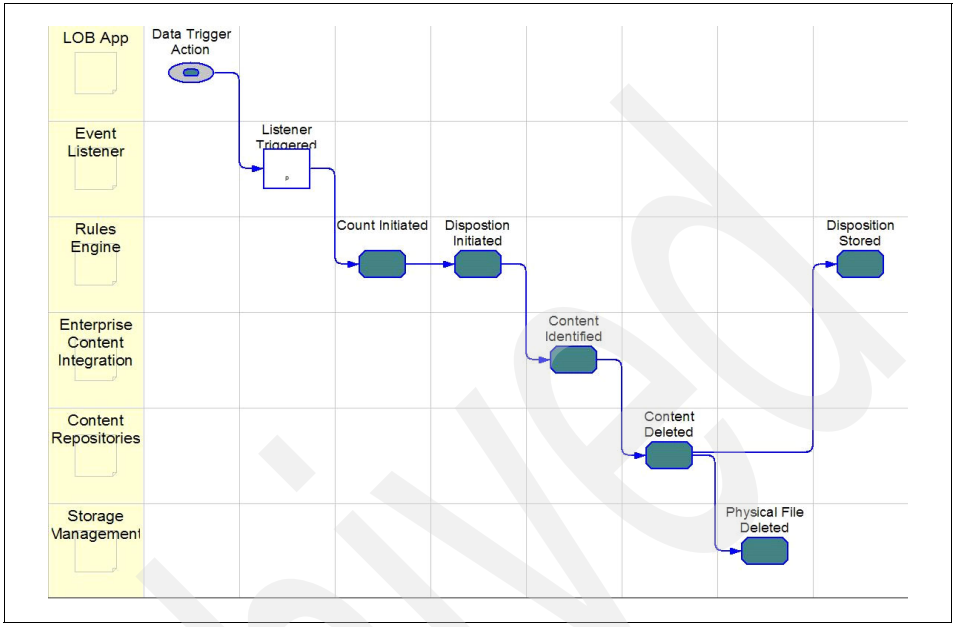

Figure 10-7 Records scheduling

Records Manager provides scheduled processes for the production of candidate record lists for disposition whenever the record staff require, scheduled for the future or on demand. Records staff can then look back in time, or ahead, to determine the number, type, and classification of the candidates for disposition. This provides expedited means for them to also seek disposal authorities from the original records owner departments. When ready, the records staff can then quickly progress to press the button, turn the crank, and begin records disposition (destruction or archiving).

With a central system for records management, companies with Records Manager can achieve more regular, timely scheduling of its large collections of e-mail records, appropriately controlled and audited. Some companies are considering moving to almost daily records scheduling for the large collections of e-mail that will expire per the applied retention schedules. The volume of e-mail they have daily and the cost of storage and management encourages performing scheduling more often. Records Manager can help in making this a repeatable structured process that can be set to run at any time or date and auctioned in the knowledge that any holds applied will ensure that only unheld records will be destroyed.

Another way to view scheduling is in the risk of not doing it. In the past, companies felt that it was best to simply keep everything. Storage keeps getting cheaper (but the management and finding of items stored does not). In the past few years, where records and compliance have risen in importance and have now been mandated, organizations realize that they cannot just keep everything. Records retention schedules help organizations identify the types of e-mail they have and how long to keep it for. Records scheduling provides structured consistent means to destroy the e-mail on a timely basis. If an organization only performs its scheduling of e-mail records twice a year, they would be best to analyze their risk: the risk of having maybe 50% of the e-mail records not deleted when they could have been appropriately and thus a 50% possible increase in discovery risk if, or rather when, a court-ordered discovery is imposed, as well as a possible 50% increase in discovery costs.

Organizations are also well served to revisit their current records scheduling processes where they may not choose to capture and store e-mail as records. A common e-mail management practice is to simply delete all e-mail older than, for example, 90 days. This can open up the organization to be viewed as having little if any corporate e-mail records-keeping in place and could be viewed to be willfully destroying current (or future) evidence. Is that e-mail related to a current employee and should it really be deleted? Is that e-mail related to a customer or supplier contract or order or business transaction or decision, for a customer relationship still ongoing or contracted? E-mail generally is a great source of company decisions. Make them corporate records for everyone's sake.

Disposition results

With e-mail being such an important type of record to companies and courts (it is now one of the first types of records ordered to be put on hold and discovered against by courts), it is critical to have auditable, admissible histories of what actions on the record systems occurred. With so many e-mail messages as candidates for disposition daily, it is best to ensure that the record system can audit the individual destructions, although that volume of audit log and metadata can quickly be unmanageable at that level. More practical is to perform scheduling more often and have the audit results be able to convey the scheduling that occurred (for example, who triggered it, across what type of records, what are the retention rules and date ranges).

Records Manager provides a multi-step process for scheduling, each step of which is audited and can be reported on in summary or down to the individual e-mail record destruction level. Scheduling logs are also kept separate and can even be exported in XML for reuse and reprocessing in the future.

Freedom of Information Act

Freedom of Information Act (FOIA) searches or discovery requests are ever more common. When mandated to do so, organizations must search across their

unclassified information, including e-mail, and provide copies of the original records. They must also perform the search and provide the copies or a statement of when in a reasonable period they will be produced, usually within a set deadline as specified by law (for example, the USA Federal FOIA results or statement are due in 20 business days).

What if an organization has appropriately destroyed or otherwise dispositioned its e-mail records according to its approved retention schedule and records processes? Simply stating “we no longer have e-mail related to that topic or issue” may not hold up in the courts or the court of public opinion. This is where Records Manager can again add value in the use of its “preserve metadata” (but not preserve the e-mail content) and scheduling audit to show that one did in the past have e-mail related to the topic but according to, and with proof of action, the standard records processes and schedules, such e-mail was destroyed.

Consistency in processes and actions related to records is a necessary position of strength to always demonstrate.

11

Discovery

This chapter describes the needs, drivers, and benefits of having the e-mail archiving and records management solution able to perform discovery on e-mail records. This is not meant to be a definitive definition or guide on the greater aspects and practices of discovery, nor replace the need to have your legal council involved.

We cover the following topics:

- Yours to discover
- Poor or inadequate discovery
- Security
- Sample discovery process

11.1 Yours to discover

Discovery is the process of searching across all e-mail records and identifying those that match the discovery and hold order, usually when ordered by courts (and related to hold orders). After being filtered to remove anything under any attorney-client privilege or other classified restriction, and following negotiation with both sides of the legal issue on the terms and issues to search on, the final e-mail record results must be made records (if they are not already), placed on a hold (so they remain until the end of the legal issue), and exported to opposing council (maybe to a CD archive of some size in some agreed format).

Note: The key to practical electronic discovery for e-mail starts with strong processes, manual practices, and knowledge of the e-mail storage and record system infrastructure.

The time and costs of resources and equipment to do discovery is huge. Imagine an organization that has kept years of backup tapes of their e-mail servers. Microsoft was one such organization that was ordered by the courts to restore all their e-mail backups and perform e-mail searches related to the Netscape case. One can only guess that if e-mail had been made records and destroyed consistently and correctly, that the alleged embarrassing e-mail from Bill Gates referring to actions against Netscape would not have been discoverable.

Above all, any organization deploying and using an e-mail record system must engage both their records staff and legal department. Applying consistent e-mail records processes with a records system can reduce the risk and cost of discovery and information management overall. Remember that the party responding to the discovery order usually has to bear the asymmetrical cost of performing that discovery.

11.2 Poor or inadequate discovery

The courts and federal rules of civil procedure mandate timely discovery search and record production. The media continues to portray stories of companies large and small who fail to perform timely or accurate discovery and suffer large business impacts of fines or other restrictions.

Spoliation involves deliberately deleting documents or changing the content of documents. This should be avoided. Organizations performing discovery and e-mail record production must ensure that they do not alter the original format or context of the e-mail records, and can deliver them in some original or representative form. A harder but now very real challenge is to predict or expect

some discovery orders even when such an order has not been received. Companies need to act and predict and ensure that their normal record and non-record e-mail archiving and destruction processes, especially when automated, are held or stopped to ensure future spoliation. Organizations have lost cases or suffered large financial consequences not by what was found in discovery but by being allegedly found to have destroyed or altered records.

11.3 Security

The e-mail record system must support secure discovery. All record access must be logged, as do all actions to package, filter, review, re-search (search refining), hold, and export e-mail records. The ability to control who has discovery search and export options is also critical. You do not want every e-mail user in a company to be able to search all e-mail, for example. Usually such discovery is performed by paralegals or record staff and only via specific user accounts with security permissions to do wide-ranging searches and discovery actions.

11.4 Sample discovery process

We now step through a customized client that provides discovery search and actions using IBM DB2 Records Manager.

1. Figure 11-1 summarizes the flow of information and actions to perform discovery in the e-mail archiving and records management solution using CommonStore, Records Manager, Content Manager, and the custom client, eDiscovery.

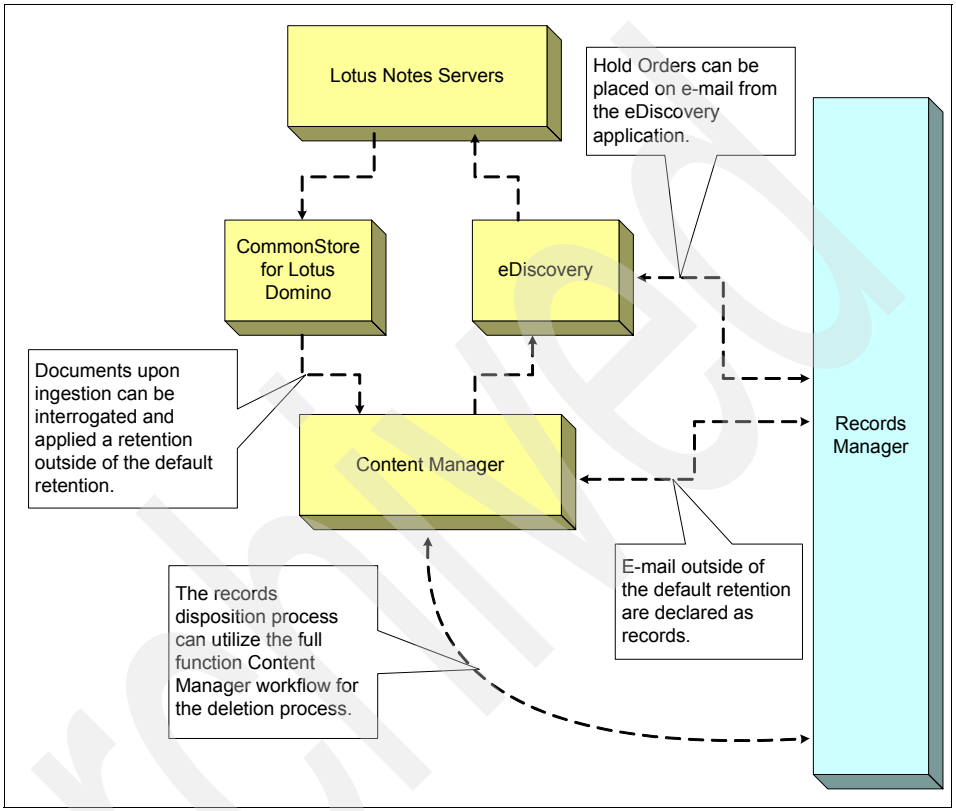

Figure 11-1 Discovery information flow for an e-mail records enabled solution

2. The eDiscovery client can be used to perform searches across all e-mail that the solution stores.

   You can enter a search criteria such as the sender, the recipients, or from which time period, and click **Search**. See Figure 11-2.

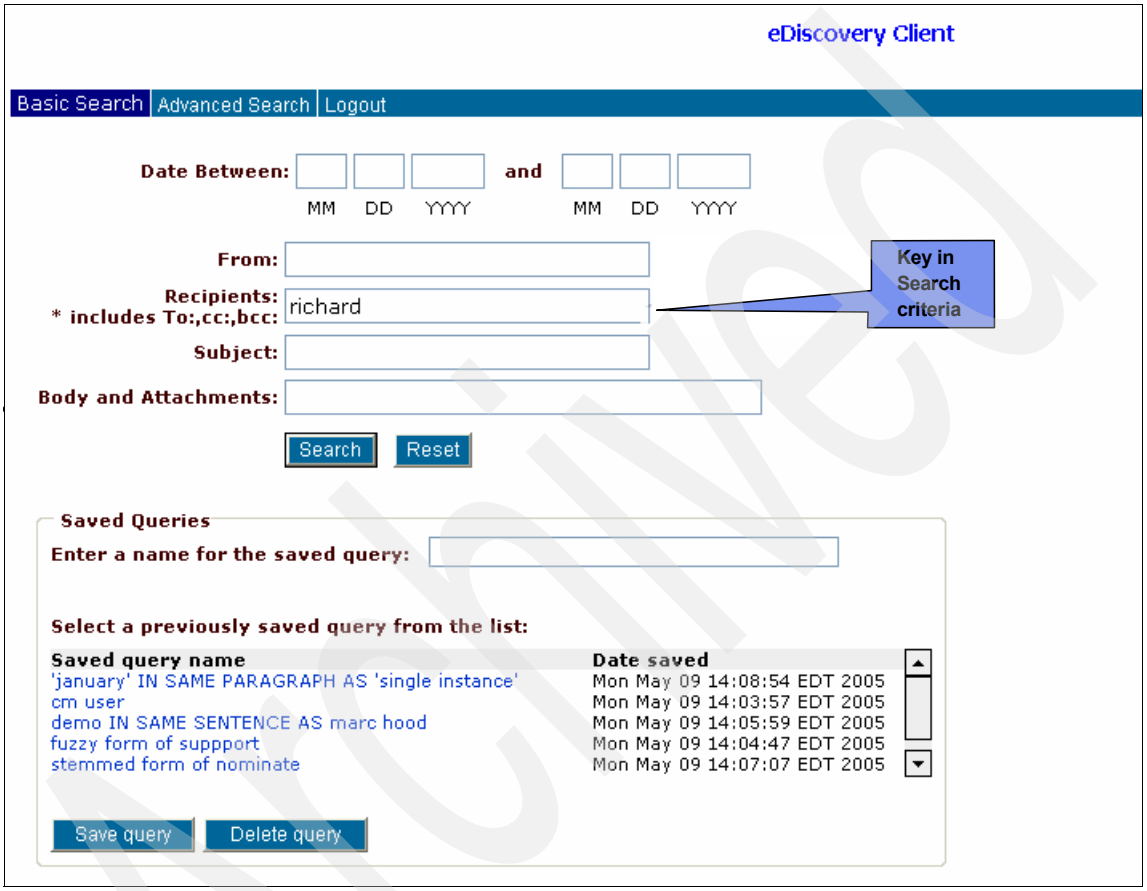

Figure 11-2 eDiscovery search window

Note: Not every e-mail is a record and does not have to be. It may be that only 5% of an organization's e-mail will be candidate records under more long-term retention rules. By default, the solution can provide a more generic retention of the non-record e-mail (for example, keep for 120 days).

The key e-mail record discovery option demonstrated here is the ability to search across any of the core metadata that makes up e-mail, the e-mail body itself, and any attachments.

Search (also known as queries) may be saved and re-used. The advanced search option on the upper-left panel enables searching of e-mail using parametric options (such as searching in the same paragraph and sentence).

3. Figure 11-3 shows the results of the previous discovery search. The small example here shows just two e-mail messages returned.

   One important point here is that the result also displays if the items are already part of another hold, although security may not let the discover know what those holds are.

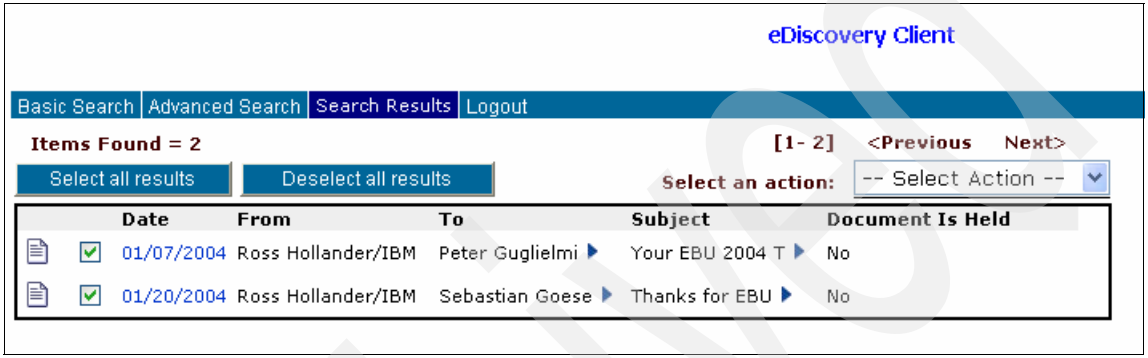

Figure 11-3 Result of the previous search

4. Figure 11-4 shows that both pieces of e-mail (which may not be records yet) have been selected for discovery action. The action menu shows the next step discovery actions available (export, retrieve, hold). In this case, we will apply a hold on the candidate e-mail, which will then be made records as a side action of applying a hold.

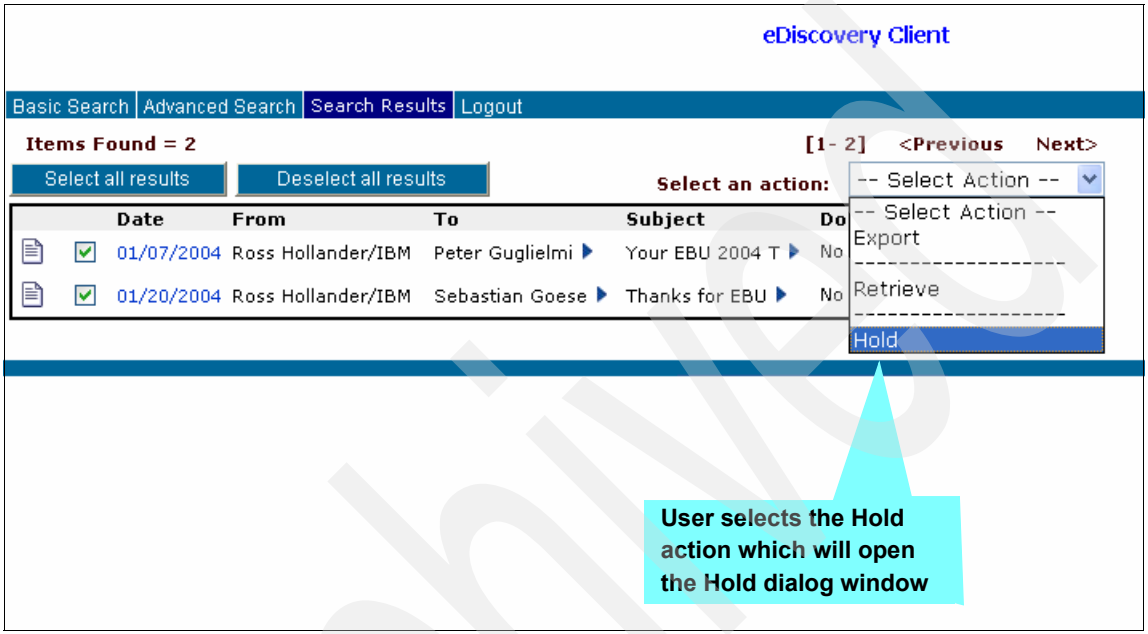

Figure 11-4 Actions you can perform on the discovered e-mail

5. In Figure 11-5, we show the ability to select a hold to apply from the master list of available holds. Remember here that the key is to get the items on hold; you do not need to review them all at this time.

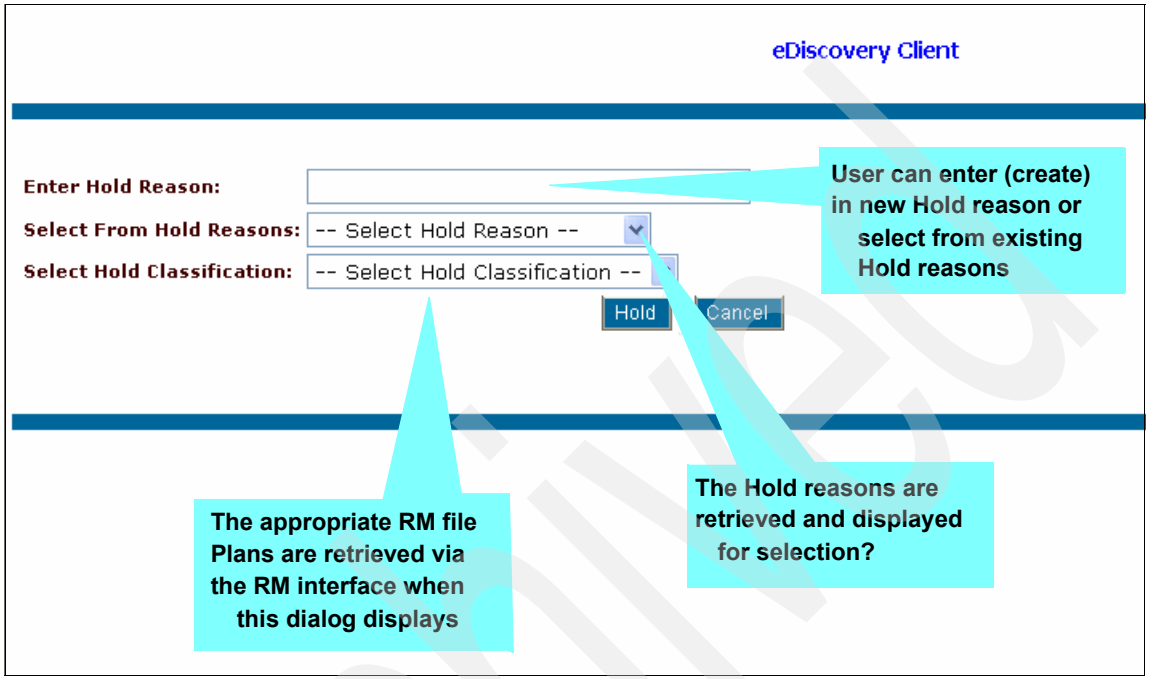

Figure 11-5 Hold a discovered e-mail

6. When a hold reason is selected, the appropriate records classification (bucket in the file plan into which to declare these candidate e-mail as records, which also is how they will then inherit the right retention rule) can be selected.

   See Figure 11-6. E-mail that is already records will be preserved and will get the hold applied.

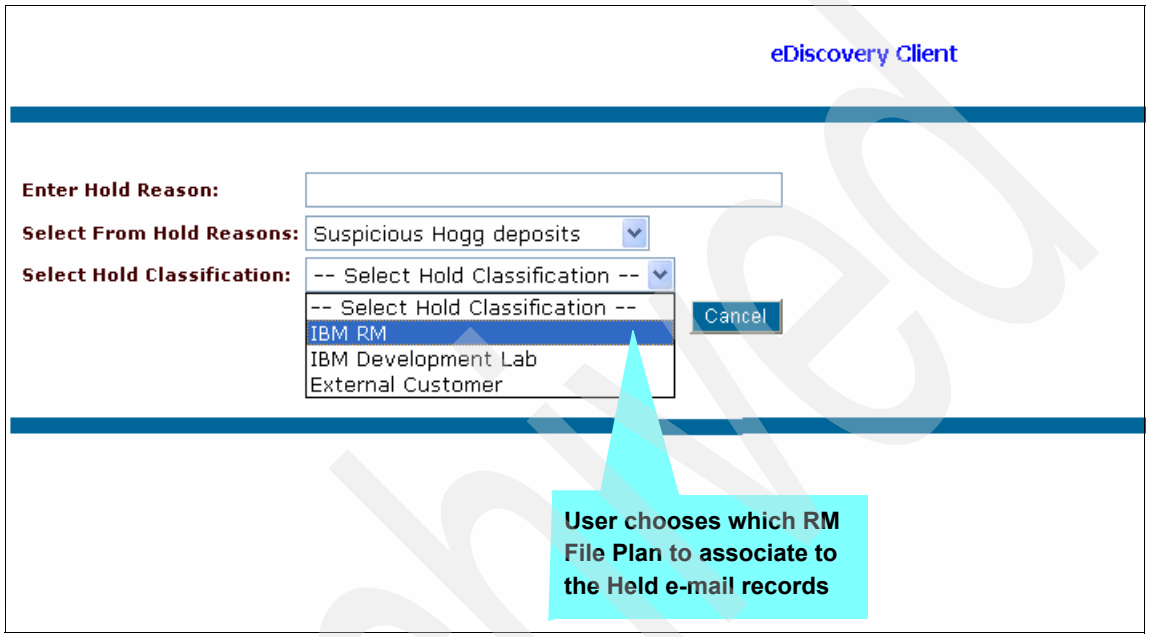

Figure 11-6 Select hold/record classification

7. Figure 11-7 shows the remaining discovery action to be performed. Click **Hold** to apply the hold (and invisibly make the e-mail records where not already).

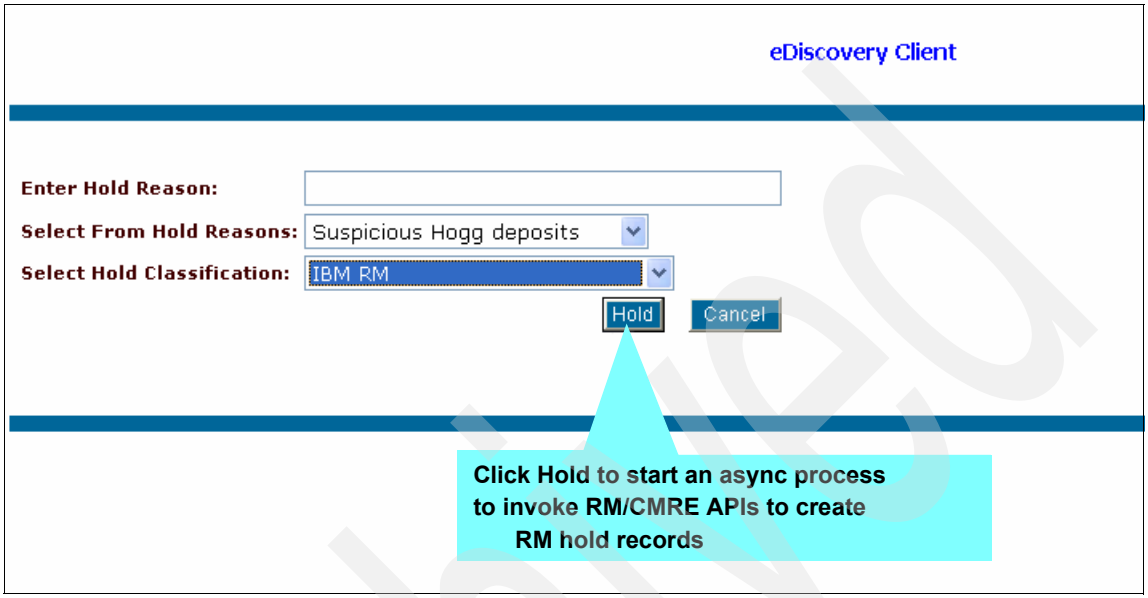

Figure 11-7 Finalize a hold on a discovered e-mail

8. As the volume of e-mail records for discovery processing can be large, these batch actions of searches and hold and export discovery actions can execute for some time. Figure 11-8 shows the processing status of such long-running processes.

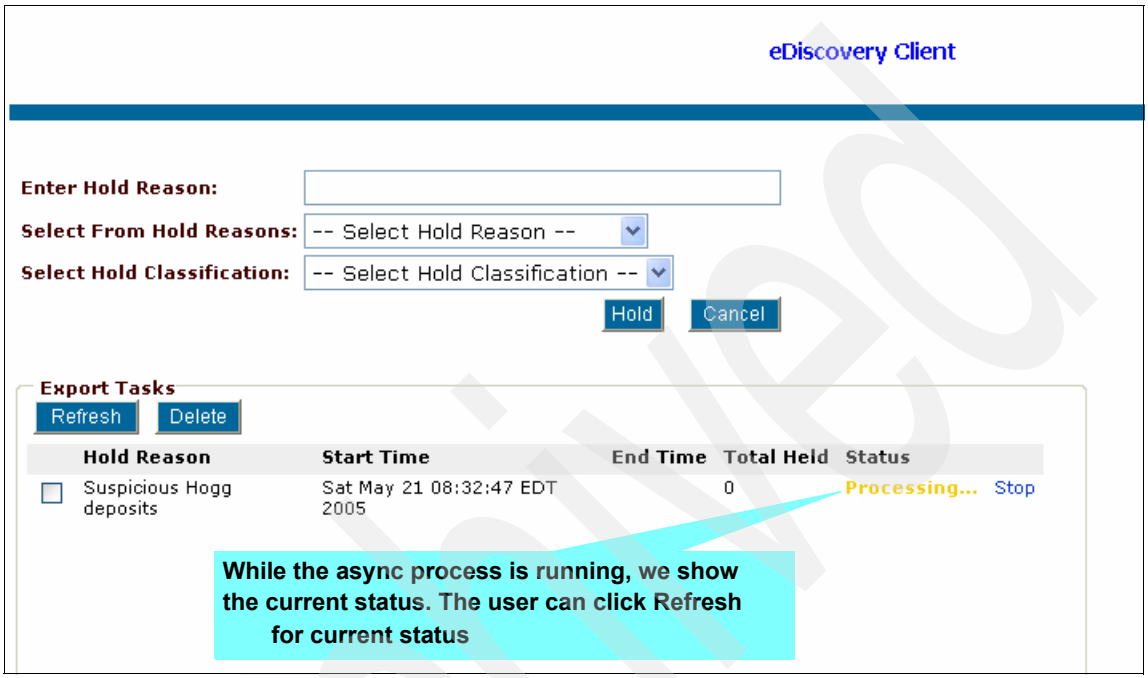

Figure 11-8 Hold process status

After the hold (and records declaration) has occurred, the export option (see Figure 11-4 on page 415) can be selected to render the e-mail into a working directory from where they can then be burnt to CD-R or DVD-R.

This solution also ensures that the exported record e-mail is rendered out in an original representative form that preserves all original metadata and content (again to avoid spoliation issues).

eDiscovery custom client is an IBM services offering that works with Content Manager and Records Manager. For more information, contact your IBM representative.

Part 4

Appendixes

A

File plan used during this Redbook

This chapter describes the file plan we used to test our scenarios during the production of this book. The scenarios we used are applied to both CommonStore for Lotus Domino and CommonStore for Exchange. Additionally, the testing we performed is done across both the Windows and AIX systems.

Purpose of the file plan

Our file plan is created to serve a number of purposes:

- Test correct installation and configuration of our system.
- Test correct user functionality for both e-mail archiving and records declaration and classification.
- Test security of the file plan components.

Our file plan is a subset of the file plan used in the *Records Manager Concepts Guide*, SC18-9182. Further details about the file plan as well as additional Records Manager configurations can be found in this document.

File plan design

The file plan we used has been reduced for clarity. However, it should be clear, from the design we used, how to extend the file plan to include a broader range of business units.

Figure A-1 on page 425 shows the file plan design we used. The instance of this file plan is shown in Figure A-2 on page 426.

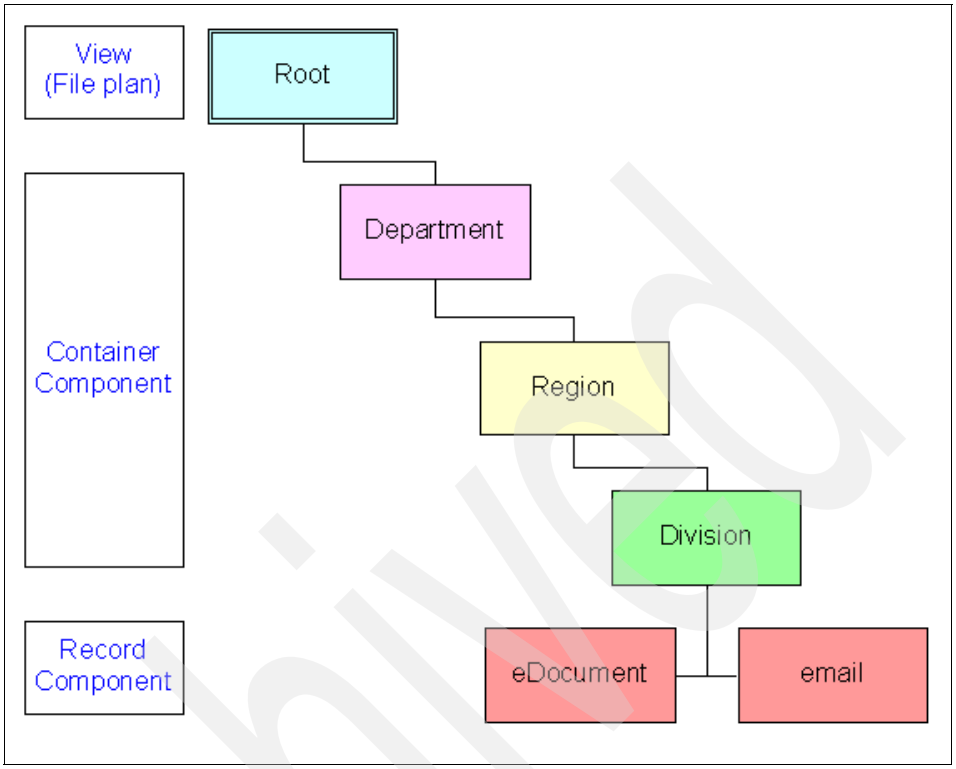

Figure A-1 File plan design

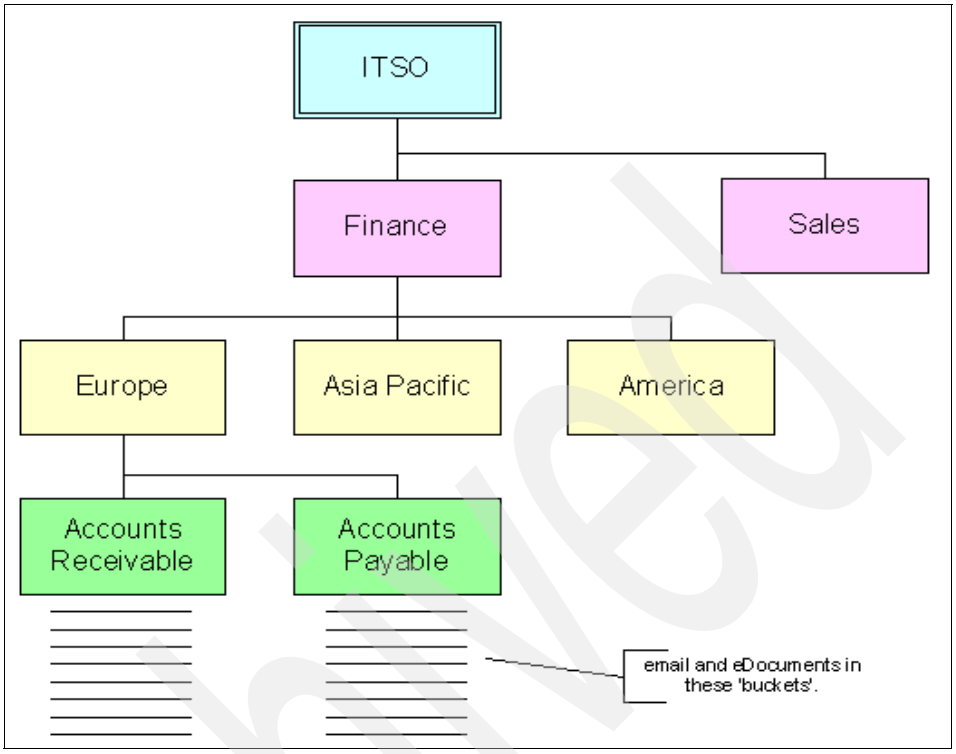

Figure A-2 Instance of the file plan design

In Figure A-2, the documents (eDocuments) and e-mail (email) classified as records are stored in either the Accounts Receivable or Accounts Payable buckets.

This is a simplified file plan. This file plan is just an example and is not proposed, intended, designed, or supported for any production use. We advise that you leverage your company's legal team, business owners, and records staff to determine and sign off on a production file plan that suits your needs. IBM can also advise on external consultants who specialize in the specialist task of file plan analysis and design.

B

Important log files for troubleshooting

The integrated e-mail archiving and records management solution is comprised of multiple products, and each has its log and trace files for troubleshooting purposes.

This appendix provides some of the important log files for each product that is involved in the integrated solution. Use these log files as references when troubleshooting system problems.

We cover the following types of log files:

- CommonStore log files
- CommonStore trace files
- Content Manager log files
- Records Manager log files
- Records Enabler log files
- WebSphere log files
- Lotus Domino log files

CommonStore log files

This section contains the log files related to CommonStore.

We break the group of log files as follows:

- Content Manager agent related log files
- HTTP task related log files
- ArchPro related log files
- Task related log files
- Crawler related log files

Content Manager agent-related log files

These log files are used to report and monitor the operation of the Content Manager API, which is used between the CommonStore Archpro program and Content Manager application. These log files produce limited details as they show program startup and shutdown events and errors. Log files will have a file extension of .log.

There are two types of Content Manager agent related log files:

- CM8Agent program (dklog.log)
- CM8Agent program (cm8errors.log)

CM8Agent program (dklog.log)

This log produces event logging of the Content Manager API calls from the Archpro program to Content Manager. It is useful for identifying problems between Archpro and Content Manager.

This log file is not produced by default. To obtain this log file, you have to configure the logging properties of the Content Manager APIs. By default, the Content Manager logging configuration file is C:\CM\IBM\db2cmv8\cmgmt\connectors\cmblogconfig.properties. Change the DKLogPriority property according to your logging needs.

On Windows the file is named *dklog.log*. It is located in the instance directory of the CommonStore Server.

- For CommonStore for Exchange Server:

  ```
  c:\Program Files\IBM\CSX\server\instance01\dklog.log
  ```

- For CommonStore for Lotus Domino on the Windows platform:

  ```
  c:\Program Files:\IBM\CSLD\server\instance01\dklog.log
  ```

On AIX, the log file has a file name starting with the CommonStore instance user ID followed by ".dklog.log". It is located under the Content Manager working directory:

- For CommonStore for Lotus Domino on the AIX platform:

  ```
  /home/ibmcmadm/log/connectors/CSLD.dklog.log
  ```

CM8Agent program (cm8errors.log)

This log file produces error logging of the CM8Agent program. It is useful for identifying problems with the CM8Agent program.

The log file is located in the CommonStore instance path.

- For CommonStore for Exchange Server:

  ```
  c:\Program Files\IBM\CSX\server\instance01\cm8errors.log
  ```

- For CommonStore for Lotus Domino on the Windows platform:

  ```
  c:\Program Files\IBM\CSLD\server\instance01\cm8errors.log
  ```

- For CommonStore for Lotus Domino on the AIX platform:

  ```
  /home/CSLD/inst001/cm8errors.log
  ```

HTTP task–related log files

These log files are used to report and monitor the operation of the CommonStore HTTP task. These log files produce error details for all requests issued to the CommonStore HTTP task. Log files have a file extension of .log.

There is one type of HTTP task related log file:

- CSHttpTask program (httperror.log)

CSHttpTask program (httperror.log)

This log file produces error logging of the CommonStore Http Listener task. It is useful for identifying problems with the viewing attachments from a URL request.

The log file is located in the CommonStore instance path.

- For CommonStore for Exchange Server:

  ```
  c:\Program Files\IBM\CSX\server\instance01\httperror.log
  ```

- For CommonStore for Lotus Domino on the Windows platform:

  ```
  c:\Program Files\IBM\CSLD\server\instance01\httperror.log
  ```

- For CommonStore for Lotus Domino on the AIX platform:

  ```
  /home/CSLD/inst001/httperror.log
  ```

The errors reported in this log are standardized HTTP return codes as listed in Table B-1.

Table B-1 HTTP return codes

HTTP error code	Short description	Long description
2xx represent success codes		
200	OK	Request has been processed successfully.
201	Created	Request has been processed successfully and a new resource (document) has been created.
4xx represent error codes		
400	Bad request	Request contained syntactical errors.
401	Unauthorized	Resource needs authorization (for example, SAP request is not signed, HTTPS is needed).
404	Not found	Resource not found (document has been deleted?).
408	Time out	Server has not answered in time.
5xx represent server error codes		
500	Internal server error	An internal server error occurred (=> Check csserror.log and traces)

Archpro-related log files

These log files are use to report and monitor the operation of the CommonStore Archpro program. These log files produce details of the CommonStore Archpro task as they show program startup and shutdown events, archive and retrieve events, and errors. Log files have a file extension of .log.

There are two types of Archpro-related log files:

- Archpro program (csserror.log)
- Archpro program (aiYYYYMMDD.log)

Archpro program (csserror.log)

This log file produces error logging of the CommonStore Server (Archpro). It is useful for identifying problems with the Archpro program.

The log file is located in the CommonStore instance path:

- For CommonStore for Exchange Server:

  `c:\Program Files\IBM\CSX\server\instance01\csserror.log`

- For CommonStore for Lotus Domino on the Windows platform:

  `c:\Program Files\IBM\CSLD\server\instance01\csserror.log`

- For CommonStore for Lotus Domino on the AIX platform:

  `/home/CSLD/inst001/csserror.log`

Archpro program (aiYYYYMMDD.log)

This log file produces logging of application startup and shutdown events, archive and retrieve events, and errors. System administrators should monitor this log to identify errors in storing or retrieving information with the backend repository.

Logging can be enabled or disabled by setting the LOG entry in archint.ini (values are ON or OFF). Log output has fields separated by semicolons and thus can be opened with Microsoft Excel® as a delimited file or monitored by third-party monitoring applications.

The log file is located in the CommonStore instance path.

- For CommonStore for Exchange Server:

  `c:\Program Files\IBM\CSX\server\instance01\ai20050602.log`

- For CommonStore for Lotus Domino on the Windows platform:

  `c:\Program Files\IBM\CSLD\server\instance01\ai20050602.log`

- For CommonStore for Lotus Domino on the AIX platform:

  `/home/CSLD/inst001/ai20050602.log`

Task-related log files

These log files are use to report and monitor the operation of the CommonStore Tasks. Log files have a file extension of .log.

There are two type of task related log files, one for CSX Task and one for CSLD Task:

- CSLD Task programs (TaskProfileName.log)
- CSX Task programs

CSLD Task programs (TaskProfileName.log)

These log files produce logging of the CommonStore for Lotus Domino tasks startup and shutdown events and errors.

The CommonStore for Lotus Domino archive and restore tasks are configured in the Lotus Notes CSLD configuration database in the database profile documents. The log files have a file name with the corresponding task's profile name with an extension of .log. The log files' maximum file sizes and their locations are specified in the tasks' profile documents.

If using default parameters, the log files are located in the following directories:

- For CommonStore for Lotus Domino on the Windows platform:

  ```
  c:\Program Files\IBM\CSLD\server\instance01\Archive.log
  ```

- For CommonStore for Lotus Domino on the AIX platform:

  ```
  /home/CSLD/inst001/logs/Archive.log
  ```

CSLD Tasks can be single-threaded or multi-threaded. This is configured in the task's profile document on the Working Databases section. If the All databases option is selected, the task will be single-threaded. If either of the other two options is selected, the task will be multi-threaded, thus using one thread per Domino server or one thread per group of mail databases:

A *single-threaded log file* has a file name the same as the corresponding profile name in the Lotus Notes CSLD configuration database in the database profile documents.

For *multi-threaded log files*, each log has at least two threads, and each produces its own trace file. The log files have a file name the same as the corresponding profile name in the Lotus Notes CSLD configuration database in the database profile documents, and they have a two-digit number appended to the file name identifying the thread that produced it.

By default, the log files are located in the following directories:

- For CommonStore for Lotus Domino on the Windows platform:

  ```
  c:\Program Files\IBM\CSLD\server\instance01\Archive00.log (startup
  information only)
  c:\Program Files\IBM\CSLD\server\instance01\Archive01.log
  ```

- For CommonStore for Lotus Domino on the AIX platform:

  ```
  /home/CSLD/inst001/logs/Archive00.log (startup information only)
  /home/CSLD/inst001/logs/Archive01.log
  ```

At the startup, the CommonStore for Lotus Domino tasks have not read the CSLD configuration database on the Notes server and therefore are unaware of the desired location to write the log files. Until the Notes database is read, the log files will be written to the path specified by the CSNINSTANCEPATH environment variable.

CSX Task programs

These log files produce logging of CommonStore for Exchange task events and errors.

The CommonStore for Microsoft Exchange archive and restore tasks are configured in Active Directory via the CSX System Manager. The log files have a file name with the corresponding date and an extension of .log.

By default, the log files are located in the task's log path.

- For CommonStore for Exchange Server:

  ```
  C:\Program Files\IBM\CSX\Task\log\cs<YYYYMMDD>.log
  ```

  For example, in our sample environment, this is:

  ```
  C:\IBM\CSX\Task\log\cs20050602.log
  ```

Crawler-related log files

These log files are use to report and monitor the operation of the CommonStore crawler. Log files have a file extension of .log.

There are two types of task-related log files, one for CSX and one for CSLD Task:

- CSLD crawler programs (CrawlerPolicyName.log)
- CSX crawler programs (Crawler_FQExchangeServerName.log)

CSLD crawler programs (CrawlerPolicyName.log)

These log files produce logging of startup and shutdown events, job creation events, and errors.

The CommonStore for Lotus Domino crawlers are configured in the Lotus Notes CSLD configuration database in the scheduled task documents. The log files have a file name with the corresponding crawler's scheduled task name with an extension of .log. The log files' maximum size is specified in the crawler's scheduled task document. The trace file is located in the path specified by the CSNINSTANCEPATH environment variable.

By default, the log files are located in the following directories.

- For CommonStore for Lotus Domino on the Windows platform:

  ```
  c:\Program Files\IBM\CSLD\server\instance01\CrawlerArchive.log
  ```

- For CommonStore for Lotus Domino on the AIX platform:

  ```
  /home/CSLD/CrawlerArchive.log
  ```

CSX crawler (Crawler_FQExchangeServerName.log)

These log files produce logging of startup and shutdown events, job creation events, and errors.

The CommonStore for Microsoft Exchange crawlers are configured in Active Directory via the CSX System Manager. The log files have a file name that starts with “crawler_” and then contains the Exchange server’s fully qualified host name and ends with an extension of log.

The log files are located in the CommonStore Task path.

- For CommonStore for Exchange Server:

  ```
  C:\Program Files\IBM\CSX\Task\crawler_<Exchange server's fully qualified
  host name>.log
  ```

  For example, in the sample environment, this would be:

  ```
  c:\IBM\CSX\server\instance01\crawler_charger.redbook.bocaraton.ibm.com.log
  ```

CommonStore trace files

Trace files are not enabled in a normal everyday CommonStore environment. Trace files should be enabled when IBM support requests detailed information about what the CommonStore processes are doing. Trace files produce more details than the log files as they show what area of code is being used at each step of program execution. Trace files have a file extension of .trace or .trc.

We list the group of trace files as follows:

- Content Manager agent related trace files
- HTTP task related trace files
- ArchPro related trace files
- Task related trace files
- CSX Active Directory related trace files
- CSX System Manager traces
- CSX Outlook Extension

In addition, we introduce some best practices when working with a trace file for troubleshooting your system problems.

Content Manager agent related trace files

There is one trace file related to Content Manager agent: cmtrace.trc.

CM8Agent program (cmtrace.trc)

CM8 Agent program produces detailed logging of its operation, in the cmtrace.trc file. This log file is located in the CommonStore instance path.

- For CommonStore for Exchange Server:

  `c:\Program Files\IBM\CSX\server\instance01\cmtrace.trc`

- For CommonStore for Lotus Domino on the Windows platform:

  `c:\Program Files\IBM\CLD\server\instance01\cmtrace.trc`

- For CommonStore for Lotus Domino on the AIX platform:

  `/home/CSLD/inst001/cmtrace.trc`

HTTP task related trace files

There is one trace file related to HTTP task: httpstartup.trc.

CSHttpTask program (httpstartup.trc)

CSHttpTask program produces detailed logging of the CommonStore HTTP listener task, in the httpstartup.trc file. This log file is located in the CommonStore instance path.

- For CommonStore for Exchange Server:

  `c:\Program Files\IBM\CSX\server\instance01\httpstartup.trc`

- For CommonStore for Lotus Domino on the Windows platform:

  `c:\Program Files\IBM\CSLD\server\instance01\httpstartup.trc`

- For CommonStore for Lotus Domino on the AIX platform:

  `/home/CSLD/inst001/httpstartup.trc`

ArchPro related trace files

There are three trace files related to the ArchPro program:

- Archpro program (archint.trace)
- Archpro program (archint_startup.trace)
- Archpro program (archservice.trace)

ArchPro program (archint.trace)

The ArchPro program produces detailed logging of its operation in the archint.trace file. The tracing can be enabled or disabled by setting the TRACE entry in the archint.ini (values are ON and OFF).

The trace file is located in the CommonStore instance path.

- For CommonStore for Exchange Server:

  ```
  c:\Program Files\IBM\CSX\server\instance01\archint.trace
  ```

- For CommonStore for Lotus Domino on the Windows platform:

  ```
  c:\Program Files\IBM\CSLD\server\instance01\archint.trace
  ```

- For CommonStore for Lotus Domino on the AIX platform:

  ```
  /home/CSLD/inst001/archint.trace
  ```

Archpro program (archint_startup.trace)

The ArchPro program produces detailed logging of the startup and shutdown errors of itself in archint_startup.trace file. The tracing can be enabled or disabled by setting the TRACE entry in the archint.ini file (values are ON and OFF).

The trace file is located in the CommonStore instance path.

- For CommonStore for Exchange Server:

  ```
  c:\Program Files\IBM\CSX\server\instance01\archint_startup.trace
  ```

- For CommonStore for Lotus Domino on the Windows platform:

  ```
  c:\Program Files\IBM\CSLD\server\instance01\archint_startup.trace
  ```

- For CommonStore for Lotus Domino on the AIX platform:

  ```
  /home/CSLD/inst001/archint_startup.trace
  ```

ArchPro program (archservice.trace)

The ArchPro program also produces detailed logging of service errors if the program is configured to run as a service, in the archservice.trace file. The tracing can be enabled or disabled by setting the TRACE entry in the archint.ini (values are ON and OFF).

The trace file is located in the CommonStore instance path.

- For CommonStore for Exchange Server:

  ```
  c:\Program Files\IBM\CSX\server\instance01\archservice.trace
  ```

- For CommonStore for Lotus Domino on the Windows platform:

  ```
  c:\Program Files\IBM\CSLD\server\instance01\archservice.trace
  ```

- For CommonStore for Lotus Domino on the AIX platform:

  `/home/CSLD/inst001/archservice.trace`

Task related trace files

There are two types of task-related trace files, one for CSLD Task and one for CSX Task.

The CSLD Task produces four trace files:

- CSLD Task program (startup00.trace)
- CSLD Task program (TaskProfileName.trace)
- CSLD Crawler program (csc_startup.trace)
- CSLD Crawler program (CrawlerPolicyName.trace)

The CSX Task produces two trace files:

- CSX Task program (csx_<task name>.trc)
- CSX Crawler program (crawler_<fully qualified exchange server name>.trc)

CSLD Task program (startup00.trace)

The CSLD Task program produces detailed logging of its startup errors, if any.

Note that at the startup, the CommonStore for Lotus Domino task has not read the CSLD configuration database on the Notes server and therefore it is unaware of the desired location to write the trace file.

The trace file is located in the path specified by the CSNINSTANCEPATH environment variable.

- For CommonStore for Lotus Domino on the Windows platform:

  `c:\Program Files\IBM\CSLD\server\instance01\startup00.trace`

- For CommonStore for Lotus Domino on the AIX platform:

  `/home/CSLD/startup00.trace`

CSLD Task program (TaskProfileName.trace)

The CSLD Task program produces detailed logging of its operation.

The CommonStore for Lotus Domino archive and restore tasks are configured in the Lotus Notes CSLD configuration database in the database profile documents.

The trace files have a file name with the corresponding task's profile name with an extension of trace. The trace files' maximum size is specified in the task's profile document.

The trace files' location is also specified in the task's profile document.

- For CommonStore for Lotus Domino on the Windows platform:

  ```
  c:\Program Files\IBM\CLD\server\instance01\Archive.trace
  ```

- For CommonStore for Lotus Domino on the AIX platform:

  ```
  /home/CSLD/inst001/trace/Archive.trace
  ```

CSLD Tasks can be single-threaded or multi-threaded. This is configured in the task's profile document in the Working Databases section. If the All databases option is selected, the task will be single-threaded. If either of the other two options is selected, the task will be multi-threaded, thus using one thread per Domino server or a thread per group of mail databases.

The *single-threaded trace file* name is the same as the corresponding profile name in the Lotus Notes CSLD configuration database in the database profile documents.

For the *multi-threaded trace files*, each task has at least two threads, and each produces its own trace file.

The trace files have a file name that corresponds to the profile name in the Lotus Notes CSLD configuration database in the database profile documents, and they have a two-digit number appended to the file name identifying the thread that produced it.

- For CommonStore for Lotus Domino on the Windows platform:

  ```
  c:\Program Files\IBM\CSLD\server\instance01\Archive00.trace (startup
  information only)
  c:\Program Files\IBM\CSLD\server\instance01\Archive01.trace
  ```

- For CommonStore for Lotus Domino on the AIX platform:

  ```
  /home/CSLD/inst001/trace/Archive00.trace (startup information only)
  /home/CSLD/inst001/trace/Archive01.trace
  ```

At startup, the CommonStore for Lotus Domino tasks have not read the CSLD configuration database on the Notes server and therefore are unaware of the desired location to write the trace files. Until the Notes database is read, the trace files will be written to the path specified by the CSNINSTANCEPATH environment variable.

CSLD crawler program (csc_startup.trace)

The crawler program produces detailed logging of the startup and shutdown events, job creation events, and errors, in the csc_startup.trace file.

This trace file is located in the path specified by the CSNINSTANCEPATH environment variable.

- For CommonStore for Lotus Domino on the Windows platform:

  ```
  c:\Program Files\IBM\CSLD\server\instance01\csc_startup.trace
  ```

- For CommonStore for Lotus Domino on the AIX platform:

  ```
  /home/CSLD/csc_startup.trace
  ```

Crawler program (CrawlerPolicyName.trace)

The crawler program also produces detailed logging of its operation.

The CommonStore for Lotus Domino crawlers are configured in the Lotus Notes CSLD configuration database in the scheduled task documents.

The trace files have a file name with the corresponding crawler's scheduled task name with an extension of .trace. The trace files's maximum file size is specified in the crawler's scheduled task.

The trace files are located in the CommonStore Instance path.

- For CommonStore for Lotus Domino on the Windows platform:

  ```
  c:\Program Files\IBM\CSLD\server\instance01\Crawler.trace
  ```

- For CommonStore for Lotus Domino on the AIX platform:

  ```
  /home/CSLD/Crawler.trace
  ```

CSX Task program (csx_TaskName.trc)

The CSX Task program produces detailed logging of its operation.

The CommonStore for Exchange Server archive and restore tasks are configured in Active Directory via the CSX System Manager.

The trace files have a file name with the corresponding task's name with an extension of .trc.

The trace files' location is specified in the Task Administration Data object via the CSX System Manager.

- For CommonStore for Exchange

  ```
  c:\Program Files\IBM\CSX\Task\csx_<task name>.trc
  ```

CSX crawler program (crawler_FQExchangeServerName.trc)

The CSX crawler program produces detailed tracing of its operation.

The CommonStore for Exchange Server crawler is configured in Active Directory via the CSX System Manager.

The trace files have a file name with the fully qualified Exchange Server name with an extension of .trc.

The trace files are located in the CSX Task path.

- For CommonStore for Exchange Server:

```
c:\Program Files\IBM\CSX\Task\crawler_<fully qualified Exchange Server
Name>.trc
```

CSX Active Directory related trace files

CsxADData_task_Archive.trc is a trace file related to CSX Active Directory.

Task program (CsxADData_TaskName.trc)

If available, the task program produces additional Active Directory logging information.

The trace file is located under the CommonStore Instance path.

- For CommonStore for Exchange Server:

```
c:\Program Files\IBM\CSX\Task\log\CsxADData_<<Task name>.trc
```

CSX System Manager traces

The CSX System Manager produces detailed logging of its operation. The traces can be configured on the machine that has CSX System Manager installed by setting the environment variable CSXMMCTRACEON to 1.

The trace files are located under the logged-on user's Documents and Settings path.

- For CommonStore for Exchange Server:

```
c:\Documents and Settings\<user>\CommonStore for Exchange System
Manager\Csx.trc
```

CSX Outlook Extension

CSX Outlook Extension produces detailed logging. The logging can be configured on the machine that has Outlook installed via the Outlook menu: **Tools** → **Options** → **CommonStore** - → **Advanced**.

The trace files are located in the logged-on user's Documents and Settings path.

- For CommonStore for Exchange Server:

```
c:\Documents and Settings\<user>\CommonStore for Exchange for
Outlook\Commonstore Extension.trc
```

Best practices

For troubleshooting problems using traces and maintaining a smooth system operation, we recommend the following best practices:

- Do not enable tracing unless requested by support or if an administrator is troubleshooting an issue.
- Tracing causes extra overhead and may produce noticeable performance degradation on an already busy systems.
- With a CSLD system, monitor the CSLD jobs database to detect archive and retrieve errors early.
- Monitor the aiYYYYMMDD.log files daily.
- Use software such as IBM Tivoli Monitor to automate the process.
- Implement a scheduled method to delete the old aiYYYYMMDD.log files, as they are not automatically deleted.
- On AIX, use **crontab** to schedule a shell script.
- On Windows, use the AT command with a third-party log deletion application.

Sample AIX shell script to remove aiYYYYMMDD.log files:

```
#!/bin/sh
# Clean up CommonStore Daily Task log files older than 7 days
find /home/CSLD/inst001/ -name "ai*.log" -atime +7 -exec rm -f {} \;
```

Content Manager log files

In this section, we introduce important log files produced by the Content Manager system. These log files are used to report and monitor the operation of the Content Manager environment. They produce limited details as they show program startup and shutdown events, operational events, and errors.

We cover the following four log files:

- Content Manager install log (cminstall.log)
- Resource Manager application (SystemOut.log / SystemErr.log)
- Library Server (db2diag.log)
- Library Server (icmserver.log)

Content Manager install log (cminstall.log)

The Content Manager installer produces logging to the cminstall.log file.

The log file is located under the Content Manager working directory.

- For Content Manager on the Windows platform:

  `c:\Program Files\db2cmv8\log\cminstall.log`

- For Content Manager on the AIX platform:

  `/opt/IBM/db2cmv8/log/cminstall.log`

Resource Manager application (SystemOut.log / SystemErr.log)

Resource Manager produces logging to the WebSphere Application Server SystemOut.log file for level informational, warning, and error messages.

By default, the Content Manager application is installed in the icmrm WebSphere Application Server server:

- For Content Manager on the Windows platform:

  `c:\Program Files\WebSphere\AppServer\logs\icmrm\SystemOut.log`

- For Content Manager on the AIX platform:

  `/usr/WebSphere/AppServer/logs/icmrm/SystemOut.log`

The Resource Manager produces logging to the WebSphere Application Server SystemErr.log file only for messages of level error.

- For Content Manager on the Windows platform:

  `c:\Program Files\WebSphere\AppServer\logs\icmrm\SystemErr.log`

- For Content Manager on the AIX platform:

  `/usr/WebSphere/AppServer/logs/icmrm/SystemErr.log`

Library Server (db2diag.log)

The Content Manager Library Server produces logging to the DB2 instance owner's log.

This log file is located under the DB2 instance path on the server where the database resides.

- For DB2 on the Windows platform:

  `c:\Program Files\DB2\SQLLIB\DB2\db2diag.log`

- For DB2 on the AIX platform:

  `/home/db2inst1/sqllib/db2dump/db2diag.log`

Library Server (icmserver.log)

The Content Manager Library Server can produce detailed tracing to a DB2 trace file.

This can be enabled following the process shown in “Configuring the Library Server trace file settings and location” from *Planning and Installing your Content Manager System*, SG27-1332.

Records Manager log files

In this section, we introduce important log files produced by Records Manager. These log files are used to report and monitor the operation of the Records Manager environment. They produce limited details as they show program startup and shutdown events, operational events, and errors.

We cover the following six log files:

- Records Manager engine (SystemOut.log)
- Records Manager engine (record_manager.log)
- Records Manager engine (record_manager_err.log)
- Records Manager engine (record_manager_extensions.log)
- Records Manager engine (record_manager_host.log)
- Records Manager Client (irm_records_manager_client.log)

Records Manager engine (SystemOut.log)

The Records Manager engine produces general logging to the WebSphere Application Server SystemOut.log file for messages of level informational, warning, and error.

By default, the Records Manager engine is installed in the WebSphere Application Server called Server1.

To find the log file:

- For Records Manager engine on the Windows platform:

  ```
  c:\Program Files\WebSphere\AppServer\logs\server1\SystemOut.log
  ```

- For Records Manager engine on the AIX platform:

  ```
  /usr/WebSphere/AppServer/logs/server1/SystemOut.log
  ```

Records Manager engine (record_manager.log)

The Records Manager engine also produces logging to the record_manager.log file. The logging level can be configured by editing the log4j.properties file located under the deployed WebSphere Application Server application called IBM DB2 Records Manager.ear in the IRMJAVA.jar file.

To find the log4j.properties file:

- For Records Manager engine on the Windows platform:

  ```
  c:\Program Files\WebSphere\AppServer\installedApps\london\IBM DB2 Records
  Manager.ear\IRMJAVA.jar\log4j.properties
  ```

- For Records Manager engine on the AIX platform:

  ```
  /usr/WebSphere/AppServer/installedApps/bonnie/IBM DB2 Records
  Manager.ear/IRMJAVA.jar/log4j.properties
  ```

This log file is located under the WebSphere install path.

- For Records Manager engine on the Windows platform:

  ```
  c:\Program Files\WebSphere\AppServer\record_manager.log
  ```

- For Records Manager engine on the AIX platform:

  ```
  /usr/WebSphere/AppServer/record_manager.log
  ```

Records Manager engine (record_manager_err.log)

The Records Manager engine also produces logging to the record_manager_err.log file. The logging level can be configured by editing the log4j.properties file located under the deployed WebSphere Application Server application called IBM DB2 Records Manager.ear in the IRMJAVA.jar file.

To find the log4j.properties file:

- For Records Manager engine on the Windows platform:

  ```
  c:\Program Files\WebSphere\AppServer\installedApps\london\IBM DB2 Records
  Manager.ear\IRMJAVA.jar\log4j.properties
  ```

- For Records Manager engine on the AIX platform:

  ```
  /usr/WebSphere/AppServer/installedApps/bonnie/IBM DB2 Records
  Manager.ear/IRMJAVA.jar/log4j.properties
  ```

The log file is located under the WebSphere install path.

- For Records Manager engine on the Windows platform:

  ```
  c:\Program Files\WebSphere\AppServer\record_manager_err.log
  ```

- For Records Manager engine on the AIX platform:

  ```
  /usr/WebSphere/AppServer/record_manager_err.log
  ```

Records Manager engine (record_manager_extensions.log)

The Records Manager engine also produces logging to the record_manager_extensions.log file. The logging level can be configured by editing the log4j.properties file located under the deployed WebSphere application called IBM DB2 Records Manager.ear in the IRMJAVA.jar file.

To find the log4j.properties file:

- For Records Manager engine on the Windows platform:

  ```
  c:\Program Files\WebSphere\AppServer\installedApps\london\IBM DB2 Records
  Manager.ear\IRMJAVA.jar\log4j.properties
  ```

- For Records Manager engine on the AIX platform:

  ```
  /usr/WebSphere/AppServer/installedApps/bonnie/IBM DB2 Records
  Manager.ear/IRMJAVA.jar/log4j.properties
  ```

The log file is located under the WebSphere install path.

- For Records Manager engine on the Windows platform:

  ```
  c:\Program Files\WebSphere\AppServer\record_manager_extensions.log
  ```

- For Records Manager engine on the AIX platform:

  ```
  /usr/WebSphere/AppServer/record_manager_extensions.log
  ```

Records Manager engine (record_manager_host.log)

The Records Manager engine also produces logging to the record_manager_host.log file. The logging level can be configured by editing the log4j.properties file located under the deployed WebSphere application called IBM DB2 Records Manager.ear in the IRMJAVA.jar file.

To find the log4j.properties file:

- For Records Manager engine on the Windows platform:

  ```
  c:\Program Files\WebSphere\AppServer\installedApps\london\IBM DB2 Records
  Manager.ear\IRMJAVA.jar\log4j.properties
  ```

- For Records Manager engine on the AIX platform:

  ```
  /usr/WebSphere/AppServer/installedApps/bonnie/IBM DB2 Records
  Manager.ear/IRMJAVA.jar/log4j.properties
  ```

The log file is located under the WebSphere install path.

- For Records Manager engine on the Windows platform:

  ```
  c:\Program Files\WebSphere\AppServer\record_manager_host.log
  ```

- For Records Manager engine on the AIX platform:

  ```
  /usr/WebSphere/AppServer/record_manager_host.log
  ```

Records Manager Client (irm_records_manager_client.log)

The Records Manager Engine also produces logging to the irm_records_manager_client.log file. The logging level can be configured by editing the log4j.properties file located under the deployed WebSphere application called IBM DB2 Records Manager.ear in the IRMJAVA.jar file.

To find the log4j.properties file:

- For Records Manager engine on the Windows platform:

  ```
  c:\Program Files\WebSphere\AppServer\installedApps\london\IBM DB2 Records
  Manager.ear\IRMJAVA.jar\log4j.properties
  ```

- For Records Manager engine on the AIX platform:

  ```
  /usr/WebSphere/AppServer/installedApps/bonnie/IBM DB2 Records
  Manager.ear/IRMJAVA.jar/log4j.properties
  ```

The log file is located under the WebSphere install path.

- For Records Manager engine on the Windows platform:

  ```
  c:\Program Files\WebSphere\AppServer\irm_records_manager_client.log
  ```

- For Records Manager engine on the AIX platform:

  ```
  /usr/WebSphere/AppServer/irm_records_manager_client.log
  ```

Records Enabler log files

In this section, we introduce important log files produced by Content Manager Records Enabler. These log files are used to report and monitor the operation of the Records Enabler environment. They produce limited details as they show program startup and shutdown events, operational events, and errors.

We cover the following eight log files:

- Records Enabler engine (SystemOut.log)
- Records Enabler engine (rmeserver.log)
  These logs exclude manual declare or view record activities.
- Records Enabler engine (rmeserver.log): manual declare and view
  These logs include only manual declare and view record activities.
- Records Enabler Host Interface (SystemOut.log)
- Records Enabler Host Interface (rmehostlog4j.log)
- Records Enabler Permission Synchronization (SystemOut.log)
- Records Enabler Permission Synchronization (rmepslog4j.log)
- Records Manager Extensions (rmepslog4j.log)

Records Enabler engine (SystemOut.log)

The Content Manager Records Enabler engine produces logging to the WebSphere Application Server SystemOut.log file for messages related to application startup and shutdown events.

By default, the Records Enabler engine is installed in the cmresrv WebSphere Application Server server.

- For Records Enabler engine on the Windows platform:

  `c:\Program Files\WebSphere\AppServer\logs\cmresrv\SystemOut.log`

- For Records Enabler engine on the AIX platform:

  `/usr/WebSphere/AppServer/logs/cmresrv/SystemOut.log`

Records Enabler engine (rmeserver.log)

The Content Manager Records Enabler engine also produces logging to the rmeserver.log file.

It logs all Records Enabler server functions *except* manual declare and view records activities. The logging level can be configured by editing the cmblogconfig.properties file located under the Content Manager working directory path.

The cmblogconfig.properties file is located in the following directory:

- For Records Enabler on the Windows platform:

  `c:\Program Files\IBM\db2cmv8\cmmgt\connectors\cmblogconfig.properties`

- For Records Enabler on the AIX platform:

  `/home/ibmcmadm/cmgmt/connectors/cmblogconfig.properties`

The rmeserver.log file is located under the Content Manager working directory path.

- For Records Enabler on the Windows platform:

  `c:\Program Files\IBM\db2cmv8\log\rme\rmeserver.log`

- For Records Enabler on the AIX platform:

  `/home/ibmcmadm/log/rme/rmeserver.log`

Records Enabler engine (rmeserver.log): manual declare and view

This logs the manual declare and view records activities related to Content Manager Records Enabler server.

By default, all Content Manager Records Enabler server activities are actually logged into one file, rmeserver.log. However, you can set a different log file to track the manual declare and view records activities for troubleshooting purposes.

The logging level can be configured by editing the rmeguilog.properties file located under the Content Manager working directory path.

The rmeguilog.properties file can be found in the following directories:

- For Records Enabler on the Windows platform:

  `c:\Program Files\IBM\db2cmv8\config\rme\rmeguilog.properties`

- For Records Enabler on the AIX platform:

  `/home/ibmcmadm/config/rme/rmeguilog.properties`

The default log file is located under the Content Manager working directory path.

- For Records Enabler on the Windows platform:

  `c:\Program Files\IBM\db2cmv8\log\rme\rmeserver.log`

- For Records Enabler on the AIX platform:

  `/home/ibmcmadm/log/rme/rmeserver.log`

Records Enabler Host Interface (SystemOut.log)

The Content Manager Records Enabler Host Interface produces logging to the WebSphere Application Server SystemOut.log file for messages related to application startup and shutdown events.

By default, the Records Enabler Host Interface is installed in the rmecmhost WebSphere Application Server server.

The SystemOut.log file can be found under the following directories:

- For Records Enabler on the Windows platform:

  `c:\Program Files\WebSphere\AppServer\logs\rmecmhost\SystemOut.log`

- For Records Enabler on the AIX platform:

  `/usr/WebSphere/AppServer/logs/rmecmhost/SystemOut.log`

Records Enabler Host Interface (rmehostlog4j.log)

The Content Manager Records Enabler Host Interface also produces logging to the rmehostlog4j.log file.

The logging level can be configured by editing the rmehostlog.properties file located in the Information Integrator for Content (II4C) working directory path.

The rmehostlog.properties file can be found under the following directories:

- For Records Enabler on the Windows platform:

  `c:\Program Files\IBM\db2cmv8\config\rme\rmehostlog.properties`

- For Records Enabler on the AIX platform:

  `/home/ibmcmadm/config/rme/rmehostlog.properties`

The log file is located under the Content Manager working directory path.

- For Records Enabler on the Windows platform:

  `c:\Program Files\IBM\db2cmv8\log\rme\rmehostlog4j.log`

- For Records Enabler on the AIX platform:

  `/home/ibmcmadm/log/rme/rmehostlog4j.log`

Records Enabler Permission Synchronization (SystemOut.log)

Content Manager Records Enabler Permission Synchronization produces logging to the WebSphere Application Server SystemOut.log file for messages related to application startup and shutdown events.

By default, the Records Enabler Permission Synchronization is installed in the cmrepsproc WebSphere Application Server server.

The SystemOut.log can be found under the following directories:

- For Records Enabler on the Windows platform:

  `c:\Program Files\WebSphere\AppServer\logs\cmrepsproc\SystemOut.log`

- For Records Enabler on the AIX platform:

  `/usr/WebSphere/AppServer/logs/cmrepsproc/SystemOut.log`

Records Enabler Permission Synchronization (rmepslog4j.log)

Records Enabler Permissions Synchronization also produces logging to the rmepslog4j.log file.

The logging level can be configured by editing the rmehostlog.properties file located under the Information Integrator for Content working directory path.

The log file is located under the Content Manager log directory.

- For Records Enabler on the Windows platform:

  `c:\Program Files\IBM\db2cmv8\log\rme\rmepslog4j.log`

- For Records Enabler on the AIX platform:

  `/home/ibmcmadm/log/rme/rmepslog4j.log`

Records Manager Extensions (rmepslog4j.log)

Records Enabler Records Manager Extensions also produces logging to the rmepslog4j.log file.

The logging level can be configured by editing the rmehostlog.properties file located under the Information Integrator for Content working directory path.

The log file can be found under the Content Manager log directory.

- For Content Manager Records Enabler on the Windows platform:

  ```
  c:\Program Files\IBM\db2cmv8\log\rme\rmepslog4j.log
  ```

- For Content Manager Records Enabler on the AIX platform:

  ```
  /home/ibmcmadm/log/rme/rmepslog4j.log
  ```

WebSphere log files

The WebSphere Application Server log files are used to report and monitor the operations of the WebSphere Application Server environment. The log files produce detailed information such as program startup and shutdown events, operational events, and errors.

The log files include SystemOut.log and SystemErr.log.

WebSphere Application Servers (SystemOut.log/SystemErr.log)

The WebSphere Application Servers produce logging to the WebSphere servers' SystemOut.log file for informational, warning, and error-level messages related to WebSphere Application Server and the applications that run under it.

The WebSphere Application Servers produce logging to the SystemErr.log file for error messages related to WebSphere Application Server and the applications that run under it.

By default, the WebSphere Application Server logs are installed under the WebSphere log path:

- For WebSphere Application Server on the Windows platform:

  ```
  c:\Program Files\WebSphere\AppServer\logs\server1\SystemOut.log
  c:\Program Files\WebSphere\AppServer\logs\server1\SystemErr.log
  ```

- For WebSphere Application Server on the AIX platform:

  ```
  /usr/WebSphere/AppServer/logs/server1/SystemOut.log
  /usr/WebSphere/AppServer/logs/server1/SystemErr.log
  ```

Lotus Domino log files

The Domino server log files are used to report and monitor the operation of the Domino environment. Log files produce limited details as they show program startup and shutdown events, operational events, and errors.

Domino servers (log.nsf)

Domino servers produce logging to the Domino servers' console and log.nsf file.

By default, the Domino server log is installed in the Domino data directory and is called Notes Log.

C

Additional material

This book refers to additional material that can be downloaded from the Internet as described below.

Locating the Web material

The Web material associated with this book is available in softcopy on the Internet from the IBM Redbooks Web server. Point your Web browser to:

ftp://www.redbooks.ibm.com/redbooks/SG246795

Alternatively, you can go to the IBM Redbooks Web site at:

ibm.com/redbooks

Select the **Additional materials** and open the directory that corresponds with the redbook form number, SG246795.

Using the Web material

The additional Web material that accompanies this redbook includes the following file:

File name	*Description*
sg246795.zip	User mapper sample code

System requirements for downloading the Web material

The following system configuration is recommended:

Hard disk space: 200 MB
Operating System: Windows
Processor: Pentium® IV or higher
Memory: 512 MB

How to use the Web material

Create a subdirectory (folder) on your workstation and unzip the contents of the Web material zip file into this folder.

Related publications

The publications listed in this section are considered particularly suitable for a more detailed discussion of the topics covered in this redbook.

IBM Redbooks

For information on ordering these publications, see “How to get IBM Redbooks” on page 456. Note that some of the documents referenced here may be available in softcopy only.

- *Content Manager Implementation and Migration Cookbook*, SG24-7051
- *DB2 II: Performance Monitoring, Tuning and Capacity Planning Guide*, SG24-7073
- *DB2 UDB/WebSphere Performance Tuning Guide*, SG24-6417
- *Disaster Recovery with DB2 UDB for z/OS*, SG24-6370
- *Maximum Performance with WebSphere Application Server V5.1 on iSeries*, SG24-6383
- *Performance Tuning for Content Manager*, SG24-6949

Other publications

These publications are also relevant as further information sources:

- *IBM DB2 CommonStore for Lotus Domino: Administrator’s and Programmer’s Guide* Version 8.3, SH12-6742
- *IBM DB2 CommonStore for Exchange Server: Administration and User’s Guide Version 8.3*, SH12-6741
- *IBM DB2 Records Manager: Concepts Guide*, SC18-9182
- *IBM DB2 Records Manager: Installation Guide*, SC18-9185
- *IBM DB2 Records Manager: Administrator’s Guide,* SC18-9180
- *IBM DB2 Records Manager: Technical Reference Guide*, SC18-9181
- *IBM DB2 Content Manager V8.3: Planning and Installing Your Content Management System*, GC27-1332
- *IBM DB2 Content Manager V8.3: System Administration Guide*, SC27-1335

- *IBM DB2 Content Manager Records Enabler:Installing and Configuring*, GC18-7570
- *IBM DB2 Content Manager Records Enabler:User's Guide*, SC18-7571

Online resources

These Web sites and URLs are also relevant as further information sources:

- *IBM DB2 Content Manager V8 Implementation on DB2 Universal Database: A Primer*

  http://www.ibm.com/developerworks/db2/library/techarticle/0305chen/0305chen.html

- DB2 Content Management library

  http://www.ibm.com/software/data/cm/library.html

- UDB disaster recovery information

  http://www.ibm.com/software/data/db2/udb/hadr.html

- Disaster recovery related information

  http://www.ibm.com/developerworks/ibm/library/i-hiavail1/

How to get IBM Redbooks

You can search for, view, or download Redbooks, Redpapers, Hints and Tips, draft publications, and Additional materials, as well as order hard-copy Redbooks or CD-ROMs, at this Web site:

ibm.com/redbooks

Help from IBM

IBM Support and downloads

ibm.com/support

IBM Global Services

ibm.com/services

Index

B

C

D

E

F

G

H

I

J

L

M

N

O

P

Q

R

S

T

U

V

W

E-mail Archiving and Records Management Integration Solution Guide

Using DB2 CommonStore and DB2 Records Manager

(1.0” spine)
0.875”<->1.498”
460 <-> 788 pages

E-mail Archiving and Records Management Integration Solution Guide

Using DB2 CommonStore and DB2 Records Manager

Redbooks

E-mail archiving type, storage model, and archiving policies

E-mail record declaration, classification, and disposition

End-to-end solution installation and integration

As more companies need to manage their e-mail to support regulatory compliance, litigation, and corporate policy and to improve system performance and productivity, the e-mail archiving and records management solution presented in this IBM Redbook is here to rescue.

In this book, we provide a general solution guide to address e-mail archiving and records management issues using the following IBM products:

- IBM DB2 CommonStore for Lotus Domino V8.3 or IBM DB2 CommonStore for Exchange Server V8.3
- IBM DB2 Records Manager V4.1.2
- IBM DB2 Content Manager V8.3
- IBM DB2 Content Manager Records Enabler V8.3

We describe the products' roles in the solution and introduce the basic concepts behind e-mail archiving and e-mail records enabling. We cover features and functions of CommonStore, Records Manager, and Records Enabler, and address key areas to understand and consider when planning and designing each piece of the solution and the overall integrated solution. In addition, we discuss different system configurations, implementation paths, security, the end-to-end solution installation and configuration, and some advanced topics such as records disposition and discovery.

This book is intended for IT architects and specialists who will be responsible in planning, designing, and implementing an e-mail archiving and records management solution.

INTERNATIONAL TECHNICAL SUPPORT ORGANIZATION

BUILDING TECHNICAL INFORMATION BASED ON PRACTICAL EXPERIENCE

IBM Redbooks are developed by the IBM International Technical Support Organization. Experts from IBM, Customers and Partners from around the world create timely technical information based on realistic scenarios. Specific recommendations are provided to help you implement IT solutions more effectively in your environment.

For more information:
ibm.com/redbooks

SG24-6795-00 ISBN 0738494127