

IBM Content Manager OnDemand for iSeries Common
Server



Planning and Installation Guide

Version 5 Release 3

IBM Content Manager OnDemand for iSeries Common
Server



Planning and Installation Guide

Version 5 Release 3

Note

Before using this information and the product it supports, read the information in "Notices" on page 137.

Third Edition (May 2004)

This edition applies to IBM Content Manager OnDemand for iSeries Common Server Version 5 Release 3 and to all subsequent releases and modifications until otherwise indicated in new editions. This edition replaces SC27-1158-01.

© Copyright International Business Machines Corporation 2001, 2004. All rights reserved.

US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

About IBM Content Manager OnDemand for iSeries Common Server Planning and Installation Guide (SC27-1158). vii

Who should read this book	vii
How this book is organized	vii
Prerequisite and related information	ix
Other information available on the World Wide Web	ix
iSeries Navigator.	ix
How to send your comments.	x

Summary of changes. xi

Part 1. Introduction 1

Chapter 1. About OnDemand. 3

Overview	3
System overview	4
Concepts	4
Application.	4
Application Group	4
Folder	5
Indexing methods	6
Documents	6
Server	7
Request manager	8
Database manager	9
Storage manager	9
Data indexing and loading.	9
Management programs	10
Web Enablement Kit	11

Chapter 2. Preparing for OnDemand . . . 13

Administrative roles and responsibilities.	14
Application programming interfaces	15
Client customization	15
Server APIs	15
Server logging	16

Chapter 3. Archive storage management 17

Introduction	17
OnDemand storage objects	17
Defining the storage configuration.	17

Part 2. System requirements 19

Chapter 4. Hardware and software . . . 21

iSeries server requirements	21
Tape storage	22
Printing	22

Server printing requirements	23
Windows client	23
Administrative software	23
Installing and running client software on a network	23
Client start up parameters	24
*LOOPBACK interface for TCP/IP.	24
Adobe PDF viewing software	24

Part 3. Planning information 27

Chapter 5. Reports and other data . . . 29

Collecting requirements	29
Input data formats	29
Indexing data.	30
Document indexing.	31
Report indexing	31
Indexing data with OnDemand programs	32
Indexing reports using date fields	33

Chapter 6. OnDemand objects. 35

Overview	35
Folders	36
Application groups	36
Applications	39
Users and groups	39
OnDemand userid relationship to OS/400 user profiles	40
Permissions	40
Overview	40
Folder permissions	41
Application group permissions	41
Naming rules.	42
Data types and field types	44

Chapter 7. Storage requirements . . . 47

Overview	47
Storage hierarchy	48
Data compression	49
Calculating disk storage requirements	50
System software	50
Temporary space for indexing	50
Disk storage for stored reports	50
OnDemand database storage	51
Database journal storage	52
Temporary space for importing index data	52
Calculating archive storage requirements	52
Report storage space	52
Migrated index storage space	53

Chapter 8. Backup and recovery. . . . 55

Overview	55
Server software	55
Server information	55
OnDemand database	55

Database backup	55
Database journaling	56
Database recovery	56
Reports	56
Disk storage	57
Archive storage	57
Chapter 9. Saving configuration files	59
OnDemand files	59
Windows font files	60
Part 4. Installing software	61
Chapter 10. Installation checklist	63
Chapter 11. Installing OnDemand server software	65
Installing OnDemand	65
OnDemand Installation Procedures	65
Installing the OnDemand end-user client	65
Installing the OnDemand administrative client	66
Part 5. Configuring and initializing the system	67
Chapter 12. Creating an instance	69
About instances	69
Configuring an instance	69
Deleting an instance	72
Starting and stopping servers	72
Automatically starting instances	73
Connecting to instances	73
Connecting from OnDemand clients	73
Working with instances	74
Chapter 13. Defining a locale	75
Background information on instances	75
Steps to take for instances that already exist	75
Steps to take for new instances or new installations	76
Solving problems	76
Chapter 14. Configuring the ARS.CFG file	79
ARS_AUTOSTART_INSTANCE parameter	79
ARS_LANGUAGE parameter	79
ARS_MSGS_LANGUAGE parameter	80
ARS_MESSAGE_OF_THE_DAY parameter	80
ARS_NUM_DBSRVR parameter	80
ARS_TMP parameter	80
ARS_PRINT_PATH parameter	80
ARSSOCK_RESOLVE_CLIENT_NAME parameter	80
Chapter 15. Configuring system start up.	81
Start the server jobs	81
Start the output queue monitor(s)	81

Chapter 16. Scheduling maintenance programs	83
Maintaining application group data	83
Backing up the OnDemand database	84

Part 6. Preparing the system for use	85
---	-----------

Chapter 17. Verifying the installation	87
---	-----------

Chapter 18. Defining migration policies and storage sets	89
---	-----------

Chapter 19. Configuring the System Log application group	91
Maintaining system log data in archive storage	91
Maintaining system log data on disk	92

Chapter 20. Configuring the System Migration application group.	95
Assigning the System Migration application group to a storage set	95

Chapter 21. Backing up the OnDemand data	97
---	-----------

Appendix A. Migration from Spool File Archive to Common Server	99
Introduction	99
What can be migrated?	101
Reports that should NOT be migrated	101
Reports that may require attention before they are migrated.	102
Migration prerequisites	103
Backups as part of the migration process	103
Backup steps	104
Description of different phases of migration	104
Phase 0: Coexistence	104
Phase 1: Definition Migration	105
Phase 2: Definition Conversion	106
Phase 3: Index Data Migration.	106
Phase 4: Index Data Conversion	107
Phase 5: Cleanup	107
Background information.	108
Sources of data for the migration of definitions	108
Scheme for creating Common Server applications, application groups, and folders from existing Spool File Archive report definitions	109
Scheme for creating Common Server folders from existing Spool File Archive report group definitions	109
Scheme for creating Common Server users and user groups from existing Spool File Archive authorization lists	110

	Scheme for creating Common Server query restriction from existing Spool File Archive key security definitions	110
	Migration of definitions: Steps to follow	111
	Identifying report definitions to migrate	111
	Creating a TEST instance for definition migration testing	111
	Creating all PRODUCTION instances prior to production definition migration	112
	Creating a backup of your data	112
	Running the definition migration programs	112
	Modifying application group definitions	115
	Testing the migrated definitions	116
	Preparing to capture production reports using the Common Server	116
	Migration of index data: Steps to follow	116
	Important	117
	Creating a backup of your data	117
	Running the index data migration programs	117
	Cleanup following successful migration	120
	Phase 5: Cleanup steps	120
	Appendix B. Including Spool File Archive folders in the Common Server folder selection list	123
	Appendix C. Database servers	125
	Description	125

Appendix D. National language support	127
Overview	127
Conversion between different code pages	127
When does character conversion occur?	127
Character mapping	128
How does OnDemand determine code page values?	128
Configuring your system for NLS	129
Language identifiers	129
Locales	129
Mapping AFP fonts	131
Creating application groups	131
Creating applications	132
Data Type	132
Indexing	133
Logical Views	133
Running OnDemand programs	134

Appendix E. System log and the system log user exit	135
--	------------

Notices	137
Trademarks	139

Index	141
------------------------	------------

About IBM Content Manager OnDemand for iSeries Common Server Planning and Installation Guide (SC27-1158)

This book contains planning and installation information for IBM® Content Manager OnDemand for iSeries™ Common Server Version 5 Release 3 (OnDemand).

Who should read this book

This book is intended primarily for system administrators who need to implement, install, and maintain OnDemand software. It can also be used by other users who require an understanding of the administration and operation of the OnDemand system.

How this book is organized

This book is organized into the following parts:

- Part 1, “Introduction,” on page 1 provides an overview of OnDemand. This part contains the following sections:
 - Chapter 1, “About OnDemand,” on page 3 provides an overview of the system
 - Chapter 2, “Preparing for OnDemand,” on page 13 contains information that can help you prepare your organization for OnDemand, lists the types of administrative tasks required to maintain the system, and describes several ways that you can configure how the server and clients operate
 - Chapter 3, “Archive storage management,” on page 17 describes how OnDemand works with the storage manager to maintain documents in archive storage
- Part 2, “System requirements,” on page 19 lists the hardware, software, network, and printing requirements for and other information about OnDemand servers and clients. This part contains the following sections:
 - Chapter 4, “Hardware and software,” on page 21 lists the hardware, software, and licensing requirements for servers and clients, provides information about the different types of server configurations, contains information about downloading data to the server and printing reports from the server
- Part 3, “Planning information,” on page 27 contains information that can help you define reports¹ to OnDemand, index data, estimate storage requirements, and plan for backup and recovery of data on the system. This part contains the following sections:
 - Chapter 5, “Reports and other data,” on page 29 contains information that can help you plan for the reports that you will be storing on the system
 - Chapter 6, “OnDemand objects,” on page 35 contains information that can help you plan application groups, applications, and folders for your reports
 - Chapter 7, “Storage requirements,” on page 47 provides information that can help you estimate the amount of storage required to maintain reports on the system

1. In OnDemand, the term *report* refers to any type of data that you want to store in the system. A 10,000-page general ledger generated by an OS/390 application and a two-page Lotus WordPro file are both reports to OnDemand.

Chapter 8, "Backup and recovery," on page 55 provides information that you can use to develop a backup and recovery plan for the system

Chapter 9, "Saving configuration files," on page 59 provides information about how to you configure a server after installing OnDemand software to meet the specific requirements of your environment

Part 4, "Installing software," on page 61 provides information for installing the OnDemand server software on the iSeries server. This part contains the following sections:

Chapter 10, "Installation checklist," on page 63 contains information to help prepare your system for software installation.

Chapter 11, "Installing OnDemand server software," on page 65 contains information about installing OnDemand server software, end-user client and administrative software.

Part 5, "Configuring and initializing the system," on page 67 provides information how to configure the OnDemand server to operate and support your organization and create and initialize the database. This part contains the following sections:

Chapter 12, "Creating an instance," on page 69 defines and provides information about configuring OnDemand instances.

Chapter 13, "Defining a locale," on page 75

Chapter 14, "Configuring the ARS.CFG file," on page 79 provides various OnDemand programs and how they obtain information from the ARS.CFG file about a variety of configuration settings.

Chapter 15, "Configuring system start up," on page 81 provides information regarding server functions you may choose to start automatically whenever the system is started.

Chapter 16, "Scheduling maintenance programs," on page 83 provides information on the OnDemand maintenance commands to run automatically on a regular schedule.

- Part 6, "Preparing the system for use," on page 85 how to verify the installation of OnDemand. This part contains the following sections:

Chapter 17, "Verifying the installation," on page 87 provides information on how to logon to a server, update a server and other installation validation steps.

Chapter 18, "Defining migration policies and storage sets," on page 89 define migration policies and storage sets before you can define reports to OnDemand or load data into the system.

Chapter 19, "Configuring the System Log application group," on page 91 provides information about how to initialize the system log. The system log is comprised of the System Log application group, a set of System Log applications, and the System Log folder. When you install and configure OnDemand, you initialize the system log.

Chapter 20, "Configuring the System Migration application group," on page 95 provides information about Index migration, which is the process by which OnDemand moves index data from the database to archive storage.

Chapter 21, "Backing up the OnDemand data," on page 97 provides information on creating a full, offline backup of the OnDemand database on removable media.

- Appendix A, "Migration from Spool File Archive to Common Server," on page 99 describes how to migrate from Spool File Archive (with or without AnyStore or the existing Server feature) to the new Common Server.

- Appendix B, “Including Spool File Archive folders in the Common Server folder selection list,” on page 123 describes how to configure the folder selection list on the OnDemand client to show folders from both the Spool File Archive and Common Server environments.
- Appendix C, “Database servers,” on page 125 provides information on the ARS_NUM_DBSRVN parameter in the ARS.CFG file, which determines the number of processes that OnDemand can start on the server to handle database requests.
- Appendix D, “National language support,” on page 127 provides information about the National Language Support (NLS) provided by OnDemand, including information about the code pages (code sets) supported to provide national language (NL) character support.
- Appendix E, “System log and the system log user exit,” on page 135 provides information about the system messages OnDemand generates.

Prerequisite and related information

Use the IBM iSeries Information Center as your starting point for looking up iSeries technical information.

You can access the Information Center two ways:

- From the following Web site: <http://www.ibm.com/eserver/series/infocenter>
- From CD-ROMs that ship with your Operating System/400® order: *iSeries Information Center*, SK3T-4091-04. This package also includes the PDF versions of iSeries manuals, *iSeries Information Center: Supplemental Manuals*, SK3T-4092-01, which replaces the Softcopy Library CD-ROM.

The Information Center contains advisors and important topics such as Java™, TCP/IP, Web serving, secured networks, logical partitions, clustering, CL commands, and system application programming interfaces (APIs). It also includes links to related IBM Redbooks™ and Internet links to other IBM Web sites such as the IBM home page.

Other information available on the World Wide Web

More iSeries information is available on the World Wide Web. You can access general information from the iSeries home page, which is at the following Web site: <http://www-1.ibm.com/servers/eserver/series/>

To access workshops on advanced iSeries functions, use the Technical Studio, located at: <http://www.iseries.ibm.com/tstudio/>

Worldwide, you can read about, select, order and take delivery of iSeries program temporary fixes (PTF) over the Internet. iSeries Internet PTFs (downloads) and Preventive Service Planning (PSP) information are available at the following Internet location: <http://as400service.ibm.com>

iSeries Navigator

IBM iSeries Navigator is a powerful graphical interface for managing your iSeries servers. iSeries Navigator functionality includes system navigation, configuration, planning capabilities, and online help to guide you through your tasks. iSeries Navigator makes operation and administration of the server easier and more

productive and is the only user interface to the new, advanced features of the OS/400®. It also includes Management Central for managing multiple servers from a central system.

You can find more information on iSeries Navigator in the IBM iSeries Information Center and at the following Web site:
<http://www.ibm.com/eserver/series/navigator/>

How to send your comments

Your feedback is important in helping to provide the most accurate and high-quality information. If you have any comments about this book or any other iSeries documentation, fill out the readers' comment form at the back of this book.

- If you prefer to send comments by FAX, use either of the following numbers: – United States, Canada, and Puerto Rico: 1-800-937-3430 – Other countries: 1-507-253-5192
- If you prefer to send comments electronically, use one of these e-mail addresses: – Comments on books: RCHCLERK@us.ibm.com – Comments on the iSeries Information Center: RCHINFOC@us.ibm.com

Be sure to include the following:

- The name of the book or iSeries Information Center topic.
- The publication number of a book.
- The page number or topic of a book to which your comment applies.

Summary of changes

This edition of *IBM Content Manager OnDemand for iSeries Common Server: Planning and Installation Guide* contains new technical information. There may be some instances where changes were made, but change bars are missing. Significant changes to note are:

- At Version 5 Release 1, Content Manager OnDemand for iSeries (OnDemand) introduced a new server implementation known as the OnDemand Common Server. The Common Server provides enhanced indexing, searching, viewing, security, PDF, and web enablement capabilities for OnDemand users and administrators. Current OnDemand customers who have implemented Spool File Archive (with or without AnyStore or the existing Server feature) can now migrate to the new Common Server using the instructions outlined in Appendix A of the Content Manager OnDemand for iSeries Common Server Planning and Installation Guide. Note that, throughout the documentation, reference to the migration of Spool File Archive data also implies AnyStore data as well, if AnyStore is installed.
- Significant additions have been made to the Content Manager OnDemand for iSeries Common Server Indexing Reference publication regarding functions supported by the OS/400 Indexer. These additions include topics related to defining multi-key indexes, transaction fields, text search fields, SCS spooled files with AFP overlays, and masks for application fields.
- Content Manager OnDemand for iSeries now supports the new iSeries-supported Plasmon optical libraries.
- Two command parameters for the Start Archived Storage Management for OnDemand (STRASMOND) command have been removed to make the use of the command simpler. See Appendix A of the Content Manager OnDemand for iSeries Common Server Administration Guide for details.
- OS/400 has withdrawn the original HTTP server support. In conjunction with this, the Content Manager OnDemand Web Enablement Kit (ODWEK) support for the original HTTP server has also been withdrawn. The HTTP Apache server is now the only supported HTTP server for ODWEK.

Part 1. Introduction

This section provides an overview of the OnDemand system and contains information that can help you better understand how OnDemand works. This section describes how OnDemand manages reports and index data, includes important information about how OnDemand, the database manager, and the storage manager work to index, load, and retrieve documents², and contains a list of the tasks that OnDemand administrators typically perform to manage an OnDemand system.

2. In OnDemand, the term *document* refers to an indexed part of a report, such as a statement, policy, or other logical grouping of pages.

Chapter 1. About OnDemand

Overview

OnDemand supports any organization that can benefit from hard copy or microfiche replacement and instant access to information. An OnDemand system can support small office environments and large enterprise installations with hundreds of system users. OnDemand can dramatically improve productivity and customer service in many businesses by providing fast access to information stored in the system.

OnDemand processes the print output of application programs, extracts index fields from the data, stores the index information in a relational database, and stores one or more copies of the data in the system. With OnDemand, you can archive newly created and frequently accessed reports on high speed, disk storage volumes and automatically migrate them to other types of storage volumes as they age.

OnDemand fully integrates the capabilities of Advanced Function Presentation™ (AFP™), including management of resources, indexes, and annotations, and supports full fidelity reprinting and faxing of documents to devices attached to a PC or OnDemand server.

OnDemand provides administrators with tools to manage OnDemand servers, and authorize users access to OnDemand servers and data stored in the system.

OnDemand provides users the ability to view documents, print, e-mail, and fax copies of documents, and attach electronic notes to documents.

Some of the advantages that OnDemand offers include:

- Easily locate data without specifying the exact report
- Retrieve the pages of the report that you need without processing the entire report
- View selected data from within a report

OnDemand can provide you with an information management tool that can increase your effectiveness when working with customers.

OnDemand does the following:

- Integrates data created by application programs into an online, electronic information archive and retrieval system
- Provides the controlled and reliable access to all of an organization's reports
- Retrieves the data that you need when you need it
- Provides a standard, intuitive client with features such as thumbnails, bookmarks, notes, and shortcuts

These features mean that OnDemand can help you quickly retrieve the specific page of a report that you need to provide fast customer service.

System overview

An OnDemand system consists of client programs and server programs that communicate over a network running the TCP/IP communications protocol, a database manager that maintains index data and server control information, and storage managers that maintain documents on various types of storage devices.

OnDemand client programs run on PCs attached to the network and communicate with OnDemand servers. OnDemand client programs operate on personal computers running Windows® systems³. The client program is the user's way to search for and retrieve reports stored on the system. Using the client program, users can construct queries and search for reports, retrieve documents from OnDemand, view, print, e-mail, and fax copies or pages of documents, and attach electronic notes to pages of a document.

OnDemand servers manage control information and index data, store and retrieve documents and resource group files, and process query requests from OnDemand client programs. The documents can reside on disk, optical, and tape storage volumes. New reports can be loaded into OnDemand every day. That way, OnDemand can retrieve the latest information generated by application programs.

OnDemand client programs and servers communicate over a computer network supported by TCP/IP. When a user submits a query, the client program sends a search request to the OnDemand server. The server returns the list of documents that match the query to the user. When the user selects a document for viewing, the client program retrieves a copy of the document from the server where the document is stored, opens a viewing window, and displays the document.

Concepts

The terms *application*, *application group*, and *folder* represent how OnDemand stores, manages, retrieves, views, and prints reports and index data. When defining a new report or type of data to OnDemand, an administrator must create an application and assign the application to an application group. (If an application group does not exist, the administrator must create one first.) Before users can search for and retrieve documents, an administrator must create or update a folder to use the application group and application.

Application

An application describes the physical characteristics of a report to OnDemand. Typically you define an application for each program that produces output that will be stored in OnDemand. The application includes information about the format of the data, the orientation of data on the page, the paper size, the record length, and the code page of the data. The application also includes parameters that the indexing program uses to locate and extract index data and processing instructions that OnDemand uses to load index data in the database and documents on storage volumes.

Application Group

An application group contains the storage management attributes of and index fields for the data that you load into OnDemand. When you load a report into

3. OnDemand provides the capability to do most client functions from almost any operating system, by using a Web browser. See "Web Enablement Kit" on page 11 for information about accessing data stored in OnDemand servers with a Web browser.

OnDemand, you must identify the application group where OnDemand will load the index data and store the documents. An application group is a collection of one or more OnDemand applications with common indexing and storage management attributes. You typically group several different reports in an application group so that users can access the information contained in the reports with a single query. All of the applications in the application group must be indexed on the same fields, for example, customer name, account number, and date.

Folder

A folder is the user's way to query and retrieve data stored in OnDemand. A folder provides users with a convenient way to find related information stored in OnDemand, regardless of the source of the information or how the data was prepared. A folder allows an administrator to set up a common query screen for several application groups that may use different indexing schemes, so that a user can retrieve the data with a single query. For example, a folder called Student Information might contain transcripts, bills, and grades, which represents information stored in different application groups, defined in different applications, and created by different programs.

Figure 1 illustrates the concepts described in this section.

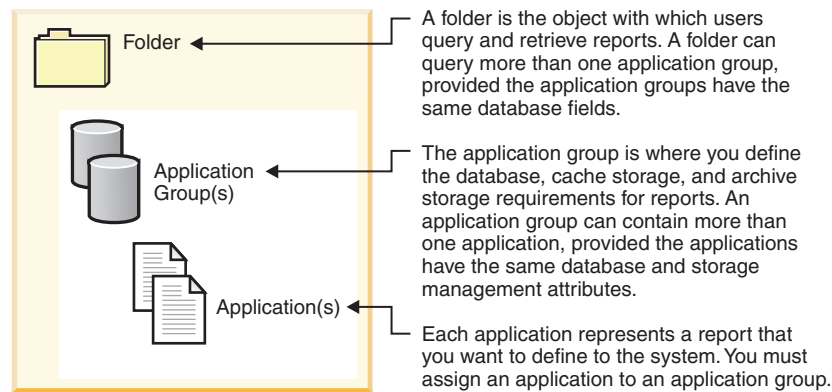


Figure 1. Folders, application groups, and applications (part 1 of 2)

Figure 2 on page 6 shows an example.

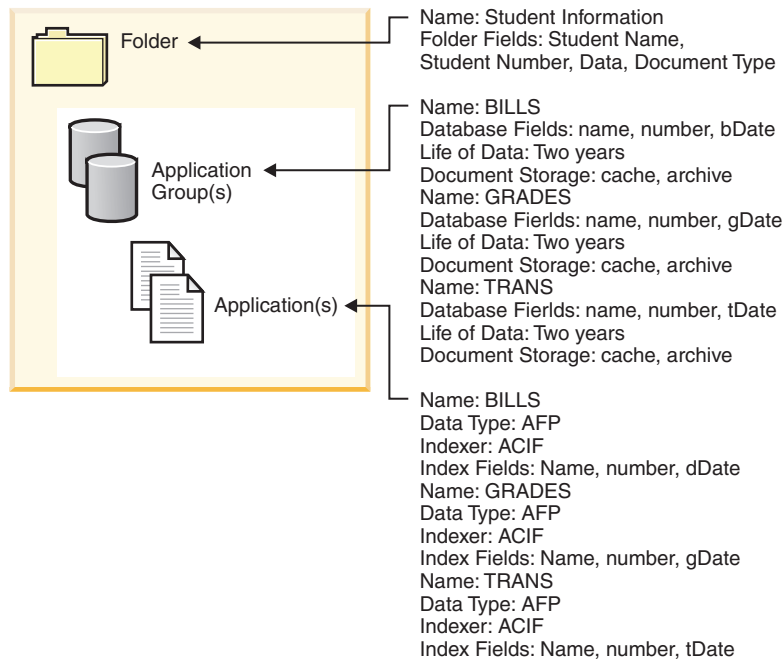


Figure 2. Folders, application groups, and applications (part 2 of 2)

Indexing methods

OnDemand provides two ways to index data:

- Document indexing is used for reports that contain logical items such as policies, and statements. Each of the items in a report can be individually indexed on values such as account number, customer name, and balance. OnDemand supports up to 32 index values per item. With document indexing, the user does not necessarily need to know about reports or report cycles to retrieve a document from OnDemand.
- Report indexing is used for reports that contain many pages of the same kind of data, such as a transaction log. Each line in the report usually identifies a specific transaction, and it would not be cost effective to index each line. OnDemand stores the report as groups of pages and indexes each group. When reports include a sorted transaction value (for example, invoice number), OnDemand can index the data on the transaction value. This is done by extracting the beginning and ending transaction values for each group of pages and storing the values in the database. This type of indexing lets users retrieve a specific transaction value directly.

Documents

OnDemand documents represent indexed groups of pages. Typically an OnDemand document is a logical section of a larger report, such as an individual customer statement within a report of thousands of statements. An OnDemand document can also represent a portion of a larger report. For reports that do not contain logical groups of pages, such as transaction logs, OnDemand can divide the report into groups of pages. The groups of pages are individually indexed and can be retrieved to the client workstation much more efficiently than the entire report. Documents are always identified by date, and usually one or more other ways, such as customer name, customer number, or transaction number.

Figure 3 illustrates OnDemand applications and documents. An administrator could define the BILLS application for a report that contains logical items, such as customer statements. The BILLS application uses the document indexing method to divide the report into documents. Each statement in the report becomes a document in OnDemand. Users can retrieve a statement by specifying the date and any combination of name and number. An administrator could define the TRANS application for a report that contains lines of sorted transaction data. The TRANS application uses the report indexing method to divide the report into documents. Each group of 100 pages in the report becomes a document in OnDemand. Each group is indexed using the first and last sorted transaction values that occur in the group. Users can retrieve the group of pages that contains a specific transaction number by specifying the date and the transaction number. OnDemand retrieves the group that contains the value entered by the user.

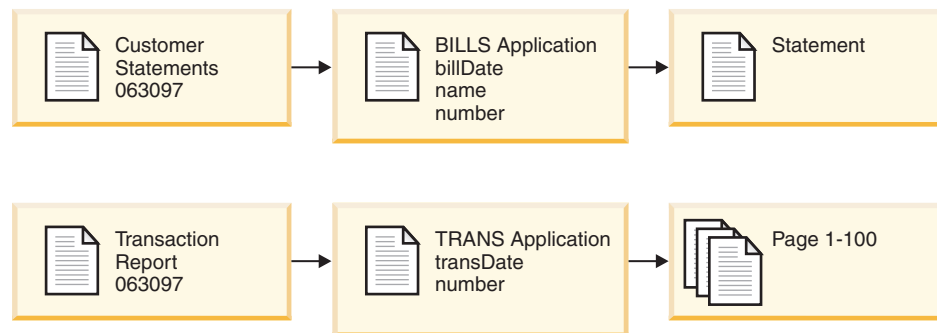


Figure 3. Applications and documents

Server

The OnDemand server maintains a central database about the reports stored in OnDemand. The database also contains information about the objects defined to the system, such as users, groups, printers, application groups, applications, folders, migration policies and storage sets. The server processes client logons, queries, and print requests and updates to the database. The server also maintains documents on disk storage and, optionally, works with the archive storage manager to maintain documents on archive media, such as optical and tape. The server loads data, retrieves documents, and expires documents.

The OnDemand server environment contains several components:

- A *request manager* that provides client, network, and operating system services, security, and accounting.
- A *database manager* that maintains the index data for the reports that you store on the system.
- Database *control information* about the users, groups, application groups, applications, folders, storage sets, and printers that you define on the system. The control information determines who can access the system, the folders that a user can open, and the application group data that a user can query and retrieve.
- A *disk storage manager* that maintains documents on disk. Disk storage is for high-speed access to the most frequently used documents.
- An *archive storage manager*, which is for the long-term storage of one or more copies of documents on archive media, such as optical and tape.
- *Data indexing* programs. These programs create index data and collect required resources (if present). OnDemand provides several indexing programs. The

OS/400 Indexer can be used to index a variety of data types, and is the most common OnDemand indexer for OS/400 spooled files. The OnDemand PDF Indexer can be used to create index data for Adobe Acrobat PDF files. The OnDemand Generic Indexer can be used to create index data for almost any other type of data that you want to store on the system, such as HTML documents, Lotus® WordPro documents, TIFF files, and so forth.

- *Data loading* programs that can be set up to automatically store report data into application groups and update the database.
- Archived reports and resources.
- A *server print* facility that allows users to reprint a large volume of documents at high speed. OnDemand requires PSF/400 Version 5 Release 3 (a product option of OS/400), if you plan to use the OnDemand server print function for AFP documents or if you plan to use server fax for any documents in OnDemand.
- OnDemand *management programs* to maintain the OnDemand database and documents on disk.
- A *system logging facility* that provides administrators with tools to monitor server activity and respond to specific events as they occur. The interface to the system logging facility is through the System Log folder and the System Log user exit within OnDemand.

The following topics provide additional information:

- The OnDemand request manager
- The OnDemand database manager
- The OnDemand storage manager
- Data indexing and loading
- OnDemand management programs

Request manager

The request manager processes search requests from OnDemand client programs. When a user enters a query, the client program sends a request over the network to the request manager. The request manager works with the database manager to compile a list of the items that match the query and returns the list to the client program. When the user selects an item for viewing, the request manager sends a retrieval request to the disk storage manager, if the document resides on disk, or to the archive storage manager, if the document resides in archive storage. The storage manager retrieves the document and, optionally, the resources associated with the item. The OnDemand client program decompresses and displays the document.

OnDemand management programs include utilities that maintain the database and disk storage, including the ability to automatically migrate data from the database and disk storage to archive storage. These programs use the services of the request manager to manage index data, documents, and resource files.

When a user logs on to the system, OnDemand assigns a unique transaction number to that instance of the client program. All activity associated with that instance of the client program contains the same transaction number. The request manager records messages generated by the various OnDemand programs in the System Log, for example, logon, query, print, and so forth. The messages contain the transaction number, userid, time stamp, and other information. Administrators can open the System Log folder and view the messages. OnDemand also provides a System Log user exit so that you can run a user-defined program to process messages. For example, you could design a user-defined program to send an alert

to an administrator when certain messages appear in the System Log. The messages in the System Log can also be used to generate usage and billing reports.

Database manager

OnDemand uses a database manager called DB2 Universal Database™ (DB2®; provided with OS/400), to maintain the index data for the reports that you load into the system. The database manager also maintains the OnDemand system tables that describe the applications, application groups, storage sets, folders, groups, users, and printers that you define to the system.

Storage manager

The OnDemand disk storage manager maintains a copy of documents, usually temporarily, on disk. The disk storage manager then migrates documents from disk to archive media and removes documents that have passed their life of data period.

OnDemand also supports an archive storage manager. The archive storage manager maintains one or more copies of documents on archive media, such as optical or tape. You decide which types of archive media that your OnDemand system must support, configure the storage devices on the system, and define the storage devices to the archive storage manager. To store application group data on archive media, you must assign the application group to a storage set that is managed by the archive storage manager.

Data indexing and loading

The reports that you store in OnDemand must be indexed. OnDemand supports several types of index data and indexing programs. For example, you can use the OS/400 Indexer to extract index data from the reports that you want to store on the system. An administrator defines the index fields and other processing parameters that the OS/400 Indexer uses to locate and extract index information from reports. OnDemand data loading programs read the index data generated by the OS/400 Indexer and load it into the OnDemand database. The data loading programs obtain other processing parameters from the OnDemand database, such as parameters used to segment, compress, and store report data on disk and on archive media. You can define your indexing parameters with the OnDemand administrative client. The administrative client includes a *report wizard* that lets you create indexing parameters by visually marking up sample report data. OnDemand also provides indexing programs that can be used to generate index data for Adobe PDF files and other types of source data, such as TIFF files. See the *IBM Content Manager OnDemand for iSeries Common Server: Indexing Reference* for details about the indexing programs provided with OnDemand.

Figure 4 on page 10 shows an overview of the data indexing and loading process.

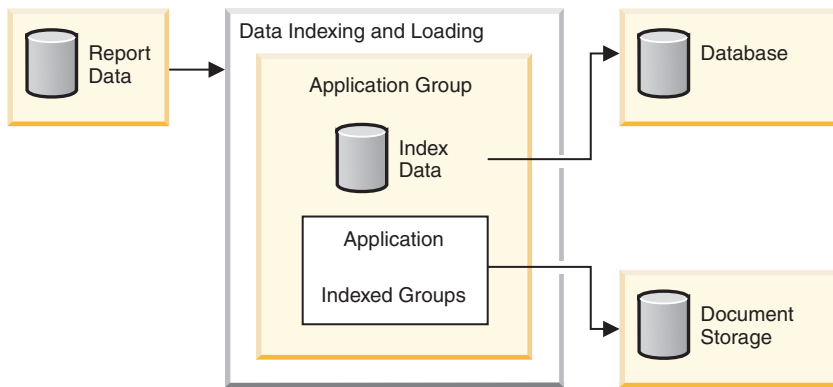


Figure 4. Data preparation, indexing, and loading

The OnDemand data loading program first determines whether the report needs to be indexed. If the report needs to be indexed, the data loading program calls the appropriate indexing program. The indexing program uses the indexing parameters from the OnDemand application to process the report data. The indexing program can extract and generate index data, divide the report into indexed groups, and collect the resources required to view and reprint the report. After indexing the report, the data loading program processes the index data, the indexed groups, and the resources using other parameters from the application and application group. The data loading program works with the database manager to update the OnDemand database with index data extracted from the report. Depending on the storage management attributes of the application group, the data loading program may work with the disk storage manager to segment, compress, and copy report data to disk and the archive storage manager to copy report data to archive storage.

Management programs

OnDemand provides programs to maintain and optimize the database and maintain documents on disk. An administrator usually determines the processing parameters for these programs, including the frequency with which the programs should run. When someone in your organization creates an application group, they specify other parameters that these programs use to maintain the report data stored in the application group. For example, when creating an application group, the administrator specifies how long documents should be maintained on the system and whether index data should be migrated from the database to archive media. The programs use the information to migrate documents from disk to archive media, delete documents from disk, migrate index data from the database to archive media, and delete index data from the database. These functions are useful because OnDemand can reclaim the disk space released by expired and migrated data. You should configure your OnDemand system to automatically start these management programs on a regular schedule, usually once every night or week.

The archive storage manager deletes data from archive media when it reaches its storage expiration date. An administrator defines management information to the archive storage manager to support the OnDemand data it manages. The management information includes storage volumes that can contain OnDemand data, the number of copies of a report to maintain, and how long to keep data in the archive management system.

The disk storage manager and the archive storage manager delete data independently of each other. Each uses its own criteria to determine when to remove documents. Each uses its own utilities and schedules to remove documents. However, for final removal of documents from the system, you should specify the same criteria to the disk storage manager and the archive storage manager.

Web Enablement Kit

The OnDemand Web Enablement Kit (ODWEK) is an optional feature of OnDemand that allows people in an organization to use a Web browser to access data stored in an OnDemand system. For example, you can provide some people with the Uniform Resource Locator (URL) of a Web page that permits them to log on to an OnDemand server and choose a folder to open; you can provide other people with the URL of a Web page that permits them to search a specific folder. ODWEK verifies that the user has permission to use the OnDemand server, has permission to use the folder, and has permission to view the data stored in an application group. After the user submits a search, ODWEK displays a Web page that contains a list of the documents that match the query. The user selects a document to view and ODWEK sends the document to the browser.

ODWEK contains several components:

- The Web server program. The server program uses standard OnDemand interfaces and protocols to access data stored in an OnDemand server. No additional code is needed on the OnDemand server to support ODWEK. You can use one of the following Web server programs to control ODWEK:
 - CGI program. The CGI program runs on an iSeries system that is running the IBM HTTP Server.
 - Java servlet. The servlet runs on a Java-enabled HTTP server with a Java application server, such as the IBM WebSphere[®] Application Server.
- The AFP Web Viewer. The AFP Web Viewer lets users search, retrieve, view, navigate, and print AFP documents from a Web browser. Each person in your organization that plans to use the AFP Web Viewer to view AFP documents must install it on their PC.
- The Image Web Viewer. The Image Web Viewer lets users search, retrieve, view, navigate, and print BMP, GIF, JPEG, PCX, and TIFF documents from a Web browser. Each person in your organization that plans to use the Image Web Viewer to view documents may need to install it on their PC. This is required when you want to view the image within the browser window, rather than allowing a default PC application to launch for viewing.
- The Line Data Java applet. The Line Data applet lets users view line data documents from a Web browser.
- The AFP2HTML Java applet. The AFP2HTML applet lets users view the output generated by the IBM AFP2WEB Transform service offering. The AFP2WEB Transform converts AFP documents and resources into HTML files that can be displayed with the AFP2HTML applet. If you plan to use the AFP2HTML applet, then you must obtain the AFP2WEB Transform from IBM and install and configure it on the Web server. See your IBM representative for more information about the AFP2WEB Transform.

To view other types of documents stored in OnDemand, you must obtain and install the appropriate viewer. For example, to view Adobe Portable Data Format (PDF) documents, get the Adobe Acrobat viewer for the browsers that are used in your organization. If you plan to convert AFP documents that are stored in OnDemand into PDF documents that you can view with the Adobe Acrobat

| viewer, then you must obtain the AFP2PDF Transform from IBM. See your IBM
| representative for more information about the AFP2PDF Transform. To use the
| viewers and applets, the browser must be Netscape Navigator Version 4.7 or later
| or Internet Explorer Version 5.5 or later.

Note: ODWEK is an enablement kit and may require additional programming to provide the desired Web environment. See the *IBM Content Manager OnDemand for iSeries Common Server: Web Enablement Kit Installation* for more information.

Chapter 2. Preparing for OnDemand

This section contains an outline that you may find helpful as you prepare your organization for the OnDemand environment and perform a pilot roll out of the system.

- Work with a single department or group of end-users. Send a memo to the users to explain how OnDemand will affect their daily work.
- Develop an end-user training course or contract with IBM to provide training for OnDemand.
- Establish a support plan for the users. The plan should include the names and phone numbers of persons to contact for assistance and a list of troubleshooting tips.
- Choose a report or set of reports for an initial migration to OnDemand. Obtain hardcopy of the reports.
- Review the reports and determine the type of indexing required. Then select the fields from the reports for index, search, and display fields.
- Review the selections with the users. Verify that the index, search, and display fields allow the users to retrieve the data that they need.
- Determine the viewing requirements of your users.
- Obtain, install, and test any data transforms that you may need. For example, if you need to convert AFP data to HTML data, you can use the AFP2WEB Transform, a service offering available from IBM.
- Identify the type of data contained in the report and determine how you will create index data.
- If you plan to archive AFP data to index the report, identify the resources used by the report. Resources are reusable objects found on pages of a report, such as overlays and page segments. Overlays contain constant data that is merged with variable report data during printing and viewing. Page segments are graphics and images that appear on pages of a report file, such as a company logo.
- Configure archive storage (optical and tape storage devices) on the server. Define archive media devices and storage management policies to OnDemand to support the reports that you plan to store on the system.
- Use the OnDemand administrative client to create the application groups and applications required to support your reports.
- Use the administrative client to define the folders that users open to access data stored on the system.
- Use the administrative client to define users and groups to OnDemand.
- Index the reports.
- Load the report, resources, and index data into the application group.
- Begin end-user testing. Survey the users about initial testing and index, search, and display fields.
- Update your company's vital records list to include the hardware and software required by the OnDemand system. Update your company's operations and recovery manuals with information required to operate, support, and backup the OnDemand system.

Administrative roles and responsibilities

OnDemand administrators assume responsibility for and take care of the OnDemand system. The OnDemand system includes all sorts of things, including hardware, application and system software, reports, and users.

- Hardware includes the iSeries server, backup devices, archive media devices, client PCs, terminals, printers, and the networking equipment.
- Software includes the base operating system, prerequisite software, and client and server programs and configuration files.
- Administrators define OnDemand applications and decide how OnDemand will manage data on the servers.
- Administrators define OnDemand groups and users to the system and make sure that the client software is installed and operating properly.

While OnDemand administrators are responsible for this collective environment from the viewpoint of OnDemand users, it is likely the OnDemand administrators are not the only people in an organization working on all these components.

Depending on the size of your organization, there may be one person or many people administering the system. If your organization is large, the administrative tasks may be divided among several people. For example, an OnDemand system administrator could maintain OnDemand storage sets, system printers, groups, and users; an OnDemand application administrator could maintain application groups, applications, and folders; an operating system administrator could apply base operating system upgrades and perform problem determination; and a service administrator could maintain records of system and network hardware and software and make equipment changes.

The following list of items is typical of the tasks required to administer and maintain an OnDemand system. Some of these tasks may be the responsibility of a person other than an OnDemand administrator.

- Installing and upgrading equipment
- Installing and maintaining OnDemand programs and other software
- Defining and labeling storage volumes
- Monitoring the space used by the database and the space available on the system
- Monitoring the space used for disk storage and the space available on the system
- Monitoring the space used for archive storage and the space available on the system
- Scheduling jobs to maintain the database, disk storage, and archive storage
- Working with users to determine report indexing and retrieval requirements
- Defining migration policies and storage sets
- Defining OnDemand system printers
- Defining reports to the system
- Defining OnDemand groups and users
- Loading reports on the system, or scheduling the jobs that load them
- Managing the backup and recovery process for the database and other areas that contain data critical to the operation of the system
- Monitoring server activity and tuning system parameters
- Solving server, network, and application problems

- Answering end-user questions
- Establishing security and audit policies, for example: set and maintain passwords and permissions; use OnDemand's audit facilities to monitor application group and user activity; develop, document, and maintain change control procedures to prevent unauthorized changes to the system

OnDemand provides an administrative client to allow administrators to maintain OnDemand objects through an easy-to-use, graphical user interface. The administrative client is a 32-bit Windows application. The administrative client allows administrators to define and maintain application groups, storage sets, folders, system printers, applications, groups, and users. The administrative client includes features that allow administrators to process sample report data and create indexing parameters and logical views by visually marking up a sample of a report.

OnDemand provides a set of administrative commands to help administrators maintain the system. For example, OnDemand provides commands for loading and unloading reports, maintaining the database, disk, and archive storage, and querying and retrieving documents. Many of the administrative commands can be configured to run automatically, on a regular schedule.

Application programming interfaces

OnDemand provides several kinds of application programming interfaces that you can use to customize OnDemand clients and work with objects on the server.

Client customization

OnDemand provides information about the Object Linking and Embedding (OLE) control and how to customize the Windows client by specifying command line parameters, by invoking and manipulating OnDemand from other Windows 32-bit applications with the Dynamic Data Exchange (DDE) interface, or by creating a Product Information File (PIF).

The *IBM Content Manager OnDemand: Windows Client Customization Guide* provides information about customizing the Windows clients.

For example, you can integrate Monarch⁴ Version 5 with the Windows client so that users can load OnDemand documents into Monarch. The user can then do complex data manipulation in Monarch, such as creating derived columns and generating charts and reports. See the *IBM Content Manager OnDemand Windows: Client Customization Guide* for more information.

Server APIs

OnDemand provides programs that you can use to work with objects on the system. For example:

- The ARSDOC program is a multi-purpose document processing program. You can use the ARSDOC program to query the server and generate a list of items that match a query; retrieve documents from the system; add, delete, and update documents; and send documents to the server print facility. You can run the ARSDOC program from the command line or invoke it from a user-defined program.

4. Monarch is a software program that is available from Datawatch Corporation.

The *IBM Content Manager OnDemand for iSeries Common Server: Administration Guide* provides details about these and other server API programs.

Server logging

System logging facility

OnDemand provides the system logging facility to help an administrator track activity and monitor the system. OnDemand can record the messages that are generated by the various client and server programs. For example, you can configure the system to record a message in the system log every time a user logs on the system; you can configure the system to record a message in the system log every time an unsuccessful log on attempt occurs; and so forth. When you use the administrative client to add objects to the system and update the database, OnDemand records information about your actions in the system log. You can use one of the OnDemand client programs to search for and view messages from the system log by time stamp, severity, message number, userid, and other search criteria.

System log user exit

OnDemand provides a user exit that can be used to process the messages that are written to the OnDemand system log. A common use of the system log user exit is to watch for error conditions or certain messages and take the appropriate action, such as notifying an administrator or operator or running some other program.

The system log user exit runs the ARSLOG program after writing a record to the system log. However, the ARSLOG program that is provided with OnDemand does not perform any functions. You must replace the one that is provided by IBM with your own program that performs the functions that you require. For example, you could create a program to check the message number and severity of each message written to the system log and, when appropriate, send an alert to the system console.

OnDemand sends parameters to the system log user exit, such as the name of the OnDemand instance, a time stamp, a log record identifier, the userid that is associated with the action, accounting information for the userid, a message severity, a message number, and the text of the message. The information that appears in the accounting information part of the message can be specified for each user defined to the system by using the administrative client. You can customize the text of the messages by selecting the application group fields (and values) to include in the message. You can further configure OnDemand to provide specific information to the system log user exit by setting system and application group parameters with the administrative client.

See Appendix E, "System log and the system log user exit," on page 135 for more information about the system log user exit.

Chapter 3. Archive storage management

Introduction

You can configure an OnDemand system to maintain copies of reports on disk and in archive storage. The copies in archive storage are for long-term storage. OnDemand supports a variety of optical and tape storage devices. Archive Storage Management includes the following components:

- A server program that maintains a database of information about the devices and data that it manages. The server program also controls the storage media and devices that you define to OnDemand.
- Server commands that control the server program activities such as data migration and expiration.
- An iSeries Navigator client program that you can use to define storage management policies. The storage management policies determine where data is stored and how long OnDemand maintains the data.

OnDemand storage objects

The storage management criteria that you specify determines where and when OnDemand stores reports and how it maintains them.

An administrator creates an OnDemand application for each report that is to be stored on the system. Applications with similar storage characteristics can be placed into a collection called an application group.

When you load a report into OnDemand, you assign it to an application group. The application group identifies a storage set which points to an associated migration policy. A migration policy identifies location(s) on which data is stored. OnDemand will automatically store a copy of the report on disk, unless you specify otherwise. If a migration policy also identifies archive storage locations, then OnDemand automatically stores a copy of the report in archive storage.

One or more application groups can specify the same storage set. All of the data that is associated with a particular storage set (and corresponding migration policy) will be maintained using the same type of media, devices, length of time to maintain data on the system, and so forth.

You should specify the same storage management criteria in both the application group and the migration policy definition. In other words, the sum of the Duration values for all migration policy storage levels plus value for the Days in Cache should equal the Life of Data and Indexes value in the application group.

Defining the storage configuration

Before you begin loading reports on the system, you need to determine the amount of storage required to hold the report data. You should also determine how long the system should maintain a version of a report, how many copies of a report the system should maintain, on what type of media a report should be stored, and any other business, legal or operational requirements for storing and maintaining data.

After collecting the storage requirements, you configure storage devices on the system and define devices to OnDemand. For example, you define an optical library to the iSeries server. Then the OnDemand administrator defines storage management policies, using the information that you collected about the reports that you plan to maintain on the system. For example, the policy information includes the length of time that OnDemand should keep the data that it manages.

When you load a report into the system, you identify an application group. The application group identifies a storage set. The storage set information determines how many copies of a report are maintained and where the copies are maintained.

Part 2. System requirements

This section contains information about the hardware, software, networking, and printing requirements for OnDemand servers and clients.

Chapter 4. Hardware and software

The exact hardware and software configuration that you need for OnDemand to support your organization depends on the volume of data that you plan to maintain on the system, the number of concurrent users that the system must support, the backup and recovery requirements of your organization, and the performance levels that the system must meet.

iSeries server requirements

The minimum configuration for an OnDemand for iSeries server requires the following:

- OS/400 Version 5 Release 3
- For OnDemand systems being upgraded from Version 5 Release 2, OnDemand Version 5 Release 2 HIPER PTF SI06587 (along with corequisite PTFs SI06287 and SI06581 and prerequisite PTF SI07038), provide an important set of index data correction programs that are mandatory for upgrading to OnDemand Version 5 Release 3. Full instructions are provided at www.ibm.com/software/data/ondemand/pubs/qrldmtafix_instructions.pdf. Limited instructions are also included in the Special Instructions section of the SI06587 PTF cover letter, but should be used only if you are unable to obtain the full instructions from the internet. It is important to understand these instructions fully before running the correction programs. It is also important to understand that your CM OnDemand Common Server system will be UNAVAILABLE as soon as you apply these PTFs. The correction process will need to be completed before you can successfully restart the server for your Common Server environment AND must be performed prior to upgrading to OnDemand Version 5 Release 3.
- For OnDemand systems being upgraded from Version 5 Release 1, OnDemand Version 5 Release 1 HIPER PTF SI06182 (along with corequisite PTFs SI06166 and SI05778 and prerequisite PTF SI06509) provide an important set of index data correction programs that are mandatory prerequisites for upgrading to OnDemand Version 5 Release 3. Full instructions are provided at www.ibm.com/software/data/ondemand/pubs/qrldmtafix_instructions.pdf. Limited instructions are also included in the Special Instructions section of the SI06182 PTF cover letter, but should be used only if you are unable to obtain the full instructions from the internet. It is important to understand these instructions fully before running the correction programs, AND to understand that your OnDemand Common Server system will be UNAVAILABLE as soon as you apply these PTFs. The correction process will need to be completed before you can successfully restart the server(s) for your Common Server environment AND must be performed prior to upgrading to OnDemand Version 5 Release 3.
- iSeries Access Express Version 5 Release 3 (for administrative client functions)
- QShell Interpreter Version 5 Release 3 (product option 30 of OS/400)
- International Components for Unicode (product option 39 of OS/400), even if you are not using Unicode. You must retain this OS/400 product option on your system as you run OnDemand. Now that OnDemand supports locale specifications, this option of OS/400 is mandatory.
- Extended NLS Support Version 5 Release 3 (product option 21 of OS/400), if the locale you require must be created from source code. Extended NLS Support

provides all the IBM-supplied locale source files. See Chapter 13, "Defining a locale," on page 75 for more information about locales.

- Portable Application Solutions Environment (PASE) Version 5 Release 3 (product option 33 of OS/400), if you plan to use the new text search function for AFPDS documents. It is also possible that, in the future, other new functions of OnDemand may require PASE.
- PSF/400 Version 5 Release 3 (product option 36, 37, or 38 of OS/400, depending upon the number and speed of your printers), if you plan to use the OnDemand server print function for AFP documents or if you plan to use server fax for any documents in OnDemand.
- TCP/IP
- Magnetic disk space for the database, report data, and temporary work space. A minimum of 24.6MB of space is required, in addition to sufficient disk space for report storage. (A minimum of 42.7MB is required if you install the server **and** the PDF Indexer.)
- An iSeries supported tape drive, automated tape library, or optical library for data archive. An iSeries supported automated tape library or optical library is recommended.

Tape storage

Backup Recovery and Media Services for iSeries is the recommended software to support the use of tape media for your OnDemand archives.

Printing

The OnDemand client program allows users to print directly from the document that they are viewing or print selected items from the document list. The system supports these methods of printing:

- Print to a local printer. The OnDemand client can reprint nearly all types of documents (AFP, line data, and many different types of workstation files) under Windows. This method is most commonly used to print to PostScript and PCL printers that are already defined and used by other programs on the user's PC. When printing AFP documents to local printers, the fidelity of the output can vary. The best fidelity is achieved when printing AFP documents that have been formatted with the IBM Core Interchange Fonts. If an exact reprint of the document is required, you may need to use the server print function, and may require PSF/400.
- Print through the OnDemand server print function. This is the highest performance print option, because the documents are not retrieved to the user's PC before printing. The server print function is also designed to allow many documents to be selected for reprint from the document list. However, the types of documents that are supported with server print are more limited than those that can be printed locally on a workstation printer. When the OnDemand server print option is used, the client sends a print request to the OnDemand server. The server sorts the documents by storage volume before retrieving them. After the documents are retrieved, a print job is submitted. When a user prints an AFP document, the resource group that was archived at the time that the document was loaded into the system is put into the print data stream, to make sure that the document prints with the correct resources.

An OnDemand server printer is an interface between the user and a server print device. OnDemand supports two types of server print devices: a fax machine and a physical printer.

To configure the system to support the server print function, you must define a server printer to OnDemand with the administrative client; the server printer identifies an output queue on the server.

Server printing requirements

If you plan to use the OnDemand server print function for AFP documents, or if you plan to use server fax function for any documents, then you must install PSF/400.

Windows client

The OnDemand Windows client runs under Windows 2000, Windows XP, and Windows Server 2003.

The client requires the following hardware and software:

- Physical connection to the network, such as a Token Ring or Ethernet network adapter
- 256 MB of memory
- A Pentium® or Pentium compatible 800 MHz or faster processor
- A super-VGA display and adapter with a minimum resolution of 800x600
- Sufficient temporary space for documents being viewed and printed
- The Windows TCP/IP networking protocol
- To view PDF documents from the Windows client, IBM recommends that you use Adobe Acrobat. See “Adobe PDF viewing software” on page 24 for more information.

Administrative software

OnDemand provides the following software to help you administer the system:

- Administrative commands. The server commands used to load reports and manage disk and archive storage. These commands run on the iSeries server.
- Administrative Client. The iSeries Navigator client used to maintain users, groups, application groups, applications, folders, migration policies, storage sets, and printers. The administrative client runs under Windows 2000, Windows XP, and Windows Server 2003.

The administrative client software requires the following:

- A minimum of 256 MB of memory
- A Pentium or Pentium compatible 800 MHz or faster processor
- A super-VGA display and adapter with a minimum resolution of 800x600
- A minimum of 10 MB of free hard disk space
- The Windows TCP/IP networking protocol
- iSeries Access Express Version 5 Release 3

Installing and running client software on a network

A personal computer attached to the network can share a single copy of the OnDemand client software. In general, you install client software on a network file server to save hard disk space on PCs and make it easier to upgrade software later. You typically install OnDemand client software on an individual PC for better performance and to reduce network traffic.

OnDemand provides a network installation procedure to copy OnDemand control files from a network file server to a user's PC (no program files are copied to the user's PC). As part of a network installation, the setup program builds a directory structure on the user's PC and creates folders and icons on the desktop. The OnDemand program icons point to a copy of the software installed on the network file server.

When the user runs an OnDemand client program from a network file server, the operating system loads the OnDemand programs from the network file server into memory on the user's PC. When the user selects items for viewing, OnDemand allocates temporary work space on the user's PC for documents and resources.

Before you install a client program using the network installation option, you must install the client software on a network file server. See the *IBM Content Manager OnDemand: Windows Client Customization Guide* for information about the installation options and an overview of the network installation procedure.

Client start up parameters

OnDemand provides command line parameters that you can use to customize the operation of the client program. For example, you can automate the logon process so that the user does not have to specify a server, userid, or password. See the *IBM Content Manager OnDemand: Windows Client Customization Guide* for information about the parameters that you can specify for the Windows client.

*LOOPBACK interface for TCP/IP

Every iSeries system using TCP/IP has a *LOOPBACK interface for TCP/IP. If you use the command `WRKTCPPSTS *IFC` you will see an address of 127.0.0.1 with a line description of *LOOPBACK. This interface is required for OnDemand to operate properly. If you have deleted the *LOOPBACK interface, you may experience problems such as:

- Unable to start the OnDemand servers
- Unable to archive documents
- Poor performance logging on to the OnDemand Client

If you have deleted the *LOOPBACK interface, you should recreate it using the command:

```
ADDTCPIFC INTNETADR('127.0.0.1') LIND(*LOOPBACK) SUBNETMASK('127.0.0.0')
```

Adobe PDF viewing software

IBM recommends that you use Adobe Acrobat to view PDF documents from the Windows client or work with PDF input files from the administrative client. Adobe Acrobat is not provided by IBM; you should purchase Adobe Acrobat from Adobe or some other software vendor.

The following Adobe PDF viewing software can be used with OnDemand:

- Adobe Acrobat. Licensed software to view and create PDF files. You can use Adobe Acrobat to do the following:
 - Integrate Adobe Acrobat with the Windows client. If the Data Type of the OnDemand application is PDF, when a user retrieves a PDF document from the system, the client opens the PDF document in the client viewing window.

| IBM recommends that you purchase Adobe Acrobat for users that need to
| view PDF documents from the Windows client.

- | – Integrate Adobe Acrobat with the administrative client. If you plan to use the
| report wizard or the graphical indexer to process PDF input files, you must
| first install Adobe Acrobat on the PC from which you plan to run the
| administrative client. IBM recommends that you purchase Adobe Acrobat for
| users that need to work with PDF input files from the administrative client.
- | • Adobe Reader. Free of charge software (available for download from the Adobe
| Web site) to view PDF documents. The Windows client can start Adobe Reader
| in a separate window provided that all of the following conditions are met:
 - | – The Data Type of the OnDemand application is User Defined
 - | – The File Type of the OnDemand application is PDF
 - | – Adobe Reader is associated with the PDF file type on the client operating
| system

| You can find out more about Adobe software from Adobe on the Web at
| www.adobe.com.

Part 3. Planning information

This section is a planning source for OnDemand administrators. Other people in an organization interested in this section may include technical and service support personnel, database administrators, network administrators, application administrators, and anyone else who has responsibility for making decisions about business systems, such as people responsible for physical site planning, operations, and backup and recovery.

This section describes activities that IBM recommends OnDemand administrators perform to plan for the installation of OnDemand and prepare for the operation of OnDemand.

Chapter 5. Reports and other data

This section contains information that can help you plan for the reports that you will be storing into OnDemand. You can use the information to help determine the hardware configuration that you need to support your OnDemand system. We list questions that you might ask users of the reports, provide information about the types of data that you can store in OnDemand, and provide information about indexing reports.

Collecting requirements

Planning for OnDemand requires that you understand how the system will be deployed, who will use the system and how they will use it, and other end-user requirements. Answers to these questions provide information that allows you to properly configure your OnDemand system, including the storage and network configuration, to support your applications and users:

- What types of print data streams will the system support? Are transforms required to convert input data to other display formats (such as AFP to HTML)?
- What is the logical organization of the print data streams?
 - *Page* organization: a consistent stream of pages of transaction or ledger data.
 - Logical *groups* of information, such as statements or policies.
 - Data that may not have a consistent format, such as reference materials or product literature.
- Will OnDemand support short-term report management, long-term archival storage, or both?
- What is the volume of input to process? How large are your reports (in pages and bytes); how many reports; how many versions of reports?
- What index values do the users of a report need to retrieve a specific version of a report (or a document)?
- How much time is available to load reports into OnDemand? Daily? Weekly?
- How long do you plan to maintain report data on the system?
- How many concurrent, logged-on users do you anticipate on average; at peak times?
- How many active users do you anticipate?
- What is the transaction rate of the active users?

Input data formats

OnDemand supports several types of input data:

- AFP print data streams, including line data mixed with AFP structured fields and line data formatted with a page definition.
- Line data with ANSI or machine carriage control characters.
- Unformatted ASCII data that is typically generated in the workstation environment. This data type is not to be confused with USERASCII OS/400 spooled files.
- Adobe Portable Data Format (PDF) files. OnDemand provides Adobe PDF viewing software so that users can view PDF documents stored in OnDemand.
- Image files in the following formats:

- BMP (Bitmap). A file that contains a bit-mapped graphic.
- GIF (Graphic Interchange Format). A bit-mapped color graphics file format for IBM-compatible computers. GIF uses an efficient compression technique for high resolution graphics.
- JFIF (JPEG Format Image File). A file that contains image data compressed using the JPEG (Joint Photographic Experts Group) standard.
- PCX (Picture Exchange Format). A file that contains a graphic in the PCX file format, widely used by PC applications, such as the PC Paintbrush program. Compressed using PackBytes compression.
- TIFF (Tagged Image File Format). A bit-mapped graphics image format for scanned images with resolutions up to 300 DPI. TIFF simulates gray-scale shading. OnDemand supports single and multipage TIFF files that are uncompressed or are compressed using JPEG, CCITT Group 3, CCITT Group 3 / 2D, and CCITT Group 4 compression.

In addition to the types of data listed above, OnDemand allows you to store almost any other type of data on the system. For example, you can define an application for HTML documents. When you define the application, you must identify the file type of the data. The file type determines the program that the client starts when the user retrieves a document. For example, if the file type is HTML, then the client could start Netscape Navigator to view the document.

AFP supports graphics, presentation text, image, and bar code objects. Storing AFP data on the system allows full-fidelity viewing of presentation text and image objects.⁵ For example, users can retrieve and view customer statements that OnDemand presents using an electronic form, fonts, and images. The user views a copy of the statement that appears the same as the statement the customer received in the mail. AFP also supports navigation within a report file, using a table of contents.

When you store reports that contain AFP data, you must also store the resources into OnDemand. The resources include overlays, page segments, form definitions, and fonts. The resources must be resident on the processor where the data is to be indexed and loaded.

Indexing data

One of the main operations that you do with OnDemand is to index reports. When you index a report, OnDemand extracts index values from the report and stores them in the database. The database fields that you define for your application groups hold the index values. When a user opens a folder, OnDemand displays a list of search fields, which represent the database fields. To perform a query, the user enters values in the search fields. OnDemand compares the values from the search values with the values in the database fields and retrieves the items that match the query.

When you index a report, you can divide a large report into smaller, uniquely identifiable units of information. For example, when an application program generates customer bills, it may produce a large print stream made up of thousands of individual customer bills. With OnDemand, you can identify the individual customer bills within the report as smaller, separate information units,

⁵ Viewing of bar code objects is not currently supported. However, you can view bar codes that have been rendered using fonts.

or logical items (known as documents in OnDemand). Your users can search for and retrieve the logical items using identifiers such as account number, customer name, and date.

The reports that you process with OnDemand generally fit into one of two categories:

- Document. For reports made up of logical items, such as statements, bills, policies, and invoices.
- Report. For reports that (typically) contain line data, with sorted values on each page, such as a transaction log or general ledger.

Reports that do not contain logical items or sorted line data can usually be indexed with the report indexing method.

Document indexing

Document indexing can be used to index reports that are made up of logical items or to index reports that contain unique values such as an account number or a customer name. When searching and retrieving these types of reports, OnDemand returns a list of the items that match the user's query and transfers the individual items to the OnDemand client program for viewing and printing. OnDemand supports up to 32 fields as indexes or filters for document-type data. The fields do not have to be sorted and can contain numeric or text information. The fields are stored in the database as indexes or filters. Figure 5 shows an example of a report file and document indexing.

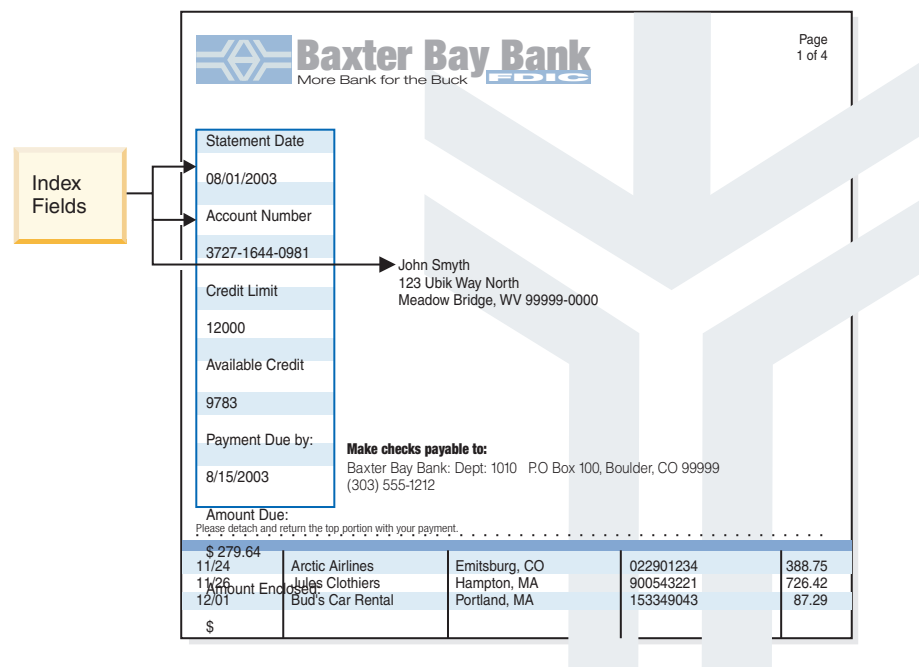


Figure 5. Document indexing method

Report indexing

Report indexing allows users to search sorted report data and retrieve the first occurrence of the value that they specified in the query. OnDemand divides the report data into groups of pages and stores the first and last index values contained in each group of pages in the database. When the user enters a query, OnDemand returns a list of the items that match the query. When the user selects

an item for viewing, OnDemand performs a text search within the item for the value specified by the user. The OnDemand client program displays the first page that contains the value specified by the user. OnDemand uses a single, unique sorted index value for the retrieval of the report data, for example, an invoice number or a transaction identifier. Figure 6 shows an example of a report file and report indexing.

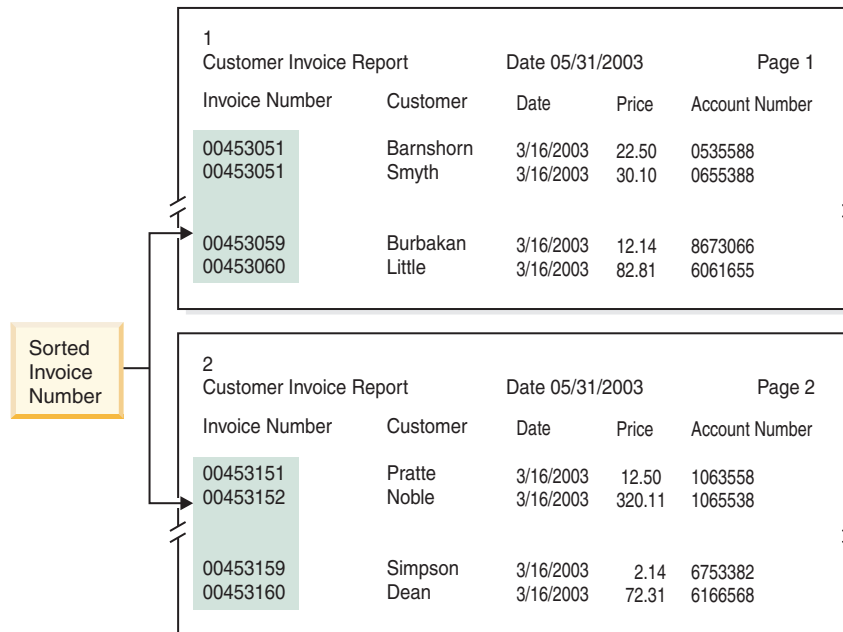


Figure 6. Report indexing method

Indexing data with OnDemand programs

OnDemand provides a variety of methods to index your print data. The following topics provide additional information:

- Use the OnDemand graphical tool to visually locate and mark index values from the print page
- Generate the index data in the application program that generates the report
- Use the OnDemand PDF Indexer program to index Adobe PDF files
- Create index data for the OnDemand Generic Indexer program

Using the graphical tool

You can use the OnDemand graphical indexing tool to visually locate and mark index values from the print data you plan to archive. This tool is part of the iSeries Navigator administrative client and can be used during your definition of an application.

Generating index data in application programs

You can create index information in the application program that generates the report. Some application programs already provide support to add indexing information. However, you may find it necessary to modify your application program to add indexing functions.

Indexing PDF input files

The OnDemand PDF Indexer is a utility that you can use to extract index data from or generate index data about Adobe PDF files. The index data can enhance your ability to store, retrieve, and view PDF documents with OnDemand. The PDF

Indexer processes PDF input files. A PDF file is a distilled version of a PostScript file, adding structure and efficiency. A PDF file can be created by Acrobat Distiller or a special printer driver program called a PDFWriter. The PDF Indexer supports PDF Version 1.3 input and output files. The *IBM Content Manager OnDemand for iSeries Common Server Indexing Reference* provides details about the PDF Indexer program and shows examples on how to use it to process PDF input files.

Generating index data using the Generic Indexer

OnDemand provides the Generic Indexer program so that you can create index data for files that cannot be indexed using other methods, such as the OS/400 Indexer. For example, you can create an index file for the Generic Indexer program that describes a set of input files that contain data in the TIFF image format. The index file contains the index fields and values for each file that you want to process. The index file also describes where the Generic Indexer program can find documents within a file. OnDemand will create a row in the database for each index record contained in the index file. Users can search the database using any combination of the index fields that were defined in the index file. The *IBM Content Manager OnDemand for iSeries Common Server Indexing Reference* describes the generic index file format.

Indexing reports using date fields

| To store data in the system, each report should be indexed with a date field. When
| querying the database, OnDemand uses the date in a report to distinguish one
| report's data from another. Retrieval performance can also be better when using a
| date field that has been also defined as a segment field. In addition, OnDemand
| also uses the report date to determine when to remove reports from disk and how
| long to maintain report data (index data and documents) on the system.

You can use the date that appears in the report, such as the run date, a transaction date, or the statement date. If the data that you want to store in OnDemand does not contain a date, you can use the date that the report was loaded into the system.

OnDemand supports date values in the range of January 1, 1970 to December 31, 2069. OnDemand also supports a date/time field. A date/time field can contain date values from January 1, 1970 to December 31, 2038.

Chapter 6. OnDemand objects

This chapter contains information that can help you plan application groups, applications, and folders for your reports.

Overview

When you install and configure the OnDemand software, you create and initialize a set of database tables that form the internal framework of the system. When you define a report to the system, OnDemand adds an application group table structure and other control information to the database.

OnDemand uses a set of objects to describe the database tables, fields, and data that make up the system. When you define an object to OnDemand, such as an application group, OnDemand stores the choices that you make and the information that you enter about the application group into the database. Every time that you load a report into an application group, OnDemand updates the database with control information, inserts rows of index data into an application group table, and stores report data and resource files on storage volumes.

Users of the OnDemand system open a folder to query and access reports that are stored on the system. A folder is the primary OnDemand object that users deal with. A folder provides users the means to search for and retrieve data stored in OnDemand. Users open a folder to construct queries and retrieve the reports that are stored in the application groups referenced by the folder. A folder can reference one or more application groups.

An application group represents the index and report data that you load into OnDemand. The OnDemand database contains tables of application group data. Records in an application group table contain index values extracted from reports and pointers to report data (documents). An application group can contain one or more applications that have the same storage characteristics and index fields.

An OnDemand application includes a description of the physical characteristics of a report, such as the type of data contained in the report and the record format of the input file, instructions to the indexing and loading programs that process the report, and information that OnDemand uses to display and print pages of the report. Typically, you define an application for each type of report that you plan to store in OnDemand. You can group applications that have the same storage characteristics and index fields into an application group.

You assign a unique name to each object that you define to OnDemand, such as application groups, applications, and folders.

OnDemand uses properties to describe the appearance, behavior, and internal structure of the objects that make up an OnDemand system. For example, Display Format is a property of a folder field that determines how OnDemand client programs display the values of the field in the document list. The properties are grouped in categories. For example, the General category under folders contains properties that describe general information about a folder, such as the name and description of the folder and the application groups contained in the folder.

Folders

A folder provides users the means to search for and retrieve related reports stored on the system. Users open folders, construct queries, and retrieve reports from application groups. (However, it is not necessary that users know about or understand application groups.) When you create a folder, you define the search and display fields that appear when the user opens the folder. You map the folder fields to database fields in the application groups referenced by the folder. The database fields contain index values extracted from the reports that are loaded into the application groups. For example, the folder search field *Customer Account Number* could be mapped to the *acct#* application group database field. OnDemand creates database records that include the index values for the *acct#* field when you load a report into the application group. When the user enters a query, OnDemand retrieves records from the database if the values of the *acct#* database field match the value that the user typed in the *Customer Account Number* search field.

When you define a folder to OnDemand, you add one or more application groups to the folder, select index fields from the application groups to appear as search and display fields when the user opens the folder, and specify the properties of the search and display fields. For example, you can determine the layout of the search fields on the screen and specify values that will automatically appear in the search fields when the user opens the folder.

OnDemand maintains information about the name of the folder and its structure in the OnDemand database. For example, the database contains information that describes the search and display fields the you defined and the database fields that you selected from application groups referenced by the folder.

You define a folder to OnDemand through properties and values grouped in categories. A category is a set of related properties. OnDemand provides folder categories for general information, permissions, field definitions, field information, and field mapping. The general category is where you specify general properties about the folder, such as the name of the folder and the application groups contained in the folder. The permissions category is where you determine the groups and users that can open the folder. You can assign other types of folder authorities in the permissions category, such as specifying someone to administer the folder. The field definitions category is where you define the search and display fields for the folder. The field information category is where you specify the attributes of the search and display fields. For example, you can specify the search operators available for each field and determine the order that the search fields appear on the screen. The field mapping category is where you map the folder search and display fields to database fields in application groups referenced by the folder.

Application groups

An application group is a collection of one or more applications that have the same index fields and storage characteristics. The application group is the object that OnDemand uses to maintain the reports that you load into the system. The application group holds index data for reports, documents, management information, permissions for the groups and users authorized to access application group, and so forth.

When you define an application group, you specify the name and type of the database fields that will hold the index data extracted from the reports that are loaded into the application group. You specify whether a database field is used to

index or filter data, and specify other characteristics of the fields. When you define an application group, OnDemand creates an application group table structure in the database, with a column for each database field that you defined. When you load a report into the application group, OnDemand inserts rows into an application group table for each indexed item found in the report. An indexed item can be a logical item, such as a policy or statement, or a group of pages, depending on how the report is organized and how you decide to index the report. Users search for reports using one or more of the fields that you defined for the application group.

OnDemand supports up to 32 *index* and *filter* fields for each application group:

- Index fields allow fast access to a specific record using a key, but generally require a large amount of disk storage to implement and require longer to load data into the application group. OnDemand uses index fields to locate the records in the database that meet the search criteria entered by the user. The index record contains the physical location of an item on a storage volume.
- Filter fields are used to refine queries, retrieving only a subset of the records found with an index field. Filter fields are generally used with an index field to identify a specific item in the application group. Filter fields can also be used to display additional information in the document list, for example, an address.

OnDemand requires a *segment field* for each application group that you define. OnDemand uses the segment field to organize and maintain application group data and to locate items that match a query. A segment field also allows a user to search one specific index file for an application using the date (which is the segment field). Without a segment field defined, all the segments (index files) within an application group are searched. The segment field must be one of the following date field or a date/time fields:

- Report Date. The date that the application program created the report file. This is typically the date found on pages of the report.
- Load Date. The date that you loaded the report into the application group. Use the load date if the report does not contain a date.

Storage requirements and index fields are the primary considerations when you define an application group and identify the applications that you can place in an application group. A third factor is the organization of the information contained in the report. OnDemand can index, store, and retrieve data contained in a report based on the structure of the data that it contains.

- Some reports are made up of logical groups of information, such as statements, invoices, and policies. These groups, or logical items, can contain one or more pages of information. OnDemand can index, store, and retrieve the logical items contained in a report. Each logical item can be indexed on up to 32 values, for example, account number, customer name, and balance. OnDemand creates a row in the database for each logical item it finds in the report.
- Other reports may be organized differently, and may not necessarily contain logical items. For example, a report could contain thousands of pages of transaction or general ledger data. OnDemand can index, store, and retrieve information from these types of reports using index values such as date, page number, and a sorted value, such as transaction number. OnDemand divides these types of reports into groups of pages and indexes each group of pages. While these types of reports may contain logical items, it probably would not be cost effective to index every item in the report. That is, indexing every item in

these types of reports would probably result in thousands of index records being added to the database each time that a report is loaded into the application group.

When you create an application group, you specify how OnDemand should store the index data for the reports that you load into the application group. OnDemand provides two methods that you can use to determine how index records are loaded into the database and how users can query the application group:

- Multiple Loads per Database Table

With this method, each time that you load a report into the application group, OnDemand inserts the index records into an existing database table. Index records for every report loaded into the application group are stored in the same database table. OnDemand maintains the application group data so that, as far as a user querying the application group knows, they appear to reside in one database table. OnDemand automatically segments the application group data when it grows beyond a certain size. OnDemand maintains a segment table for each application group. The segment table provides faster query performance by limiting searches to a specific table of application group data, using a date value to construct the query. We recommend that you use this method to organize your database when the users that search for data stored in the application group do not necessarily know or care what particular occurrence of a report generated the information that they need. This method is the default, and is used in most cases.

- Single Load per Database Table

With this method, each time that you load a report into the application group, OnDemand stores the index records into a new database table. You can define a *report field* for the application group so that users can easily search for and retrieve a specific version of a report. We recommend that you use this method to organize your database when the users that search for data stored in the application group need to query a occurrence of a report.

When you create an application group, you specify the storage characteristics of the report, such as the length of time that OnDemand maintains data stored in the application group and the data migration values. The storage characteristics also determine whether OnDemand stores a copy of the report on archive media, whether OnDemand should create a backup copy of the report, and when OnDemand removes report data when it is no longer needed.

OnDemand can perform three types of processing on application group data:

- Database expiration processing

Index data *expires* (is eligible for removal from the system) when it reaches its Life of Data and Indexes period. (You specify the Life of Data and Indexes period when you create an application group.) OnDemand provides a command that you can use to remove index data. You typically set up the command to run automatically on a regular schedule. Database expiration processing reclaims the disk space taken by deleted index data.

- Disk migration processing

Disk migration is the process of copying reports from disk to archive storage. You specify when a report should be copied from disk to archive storage when you create an application group. OnDemand provides a command that you can use to copy reports to archive storage. You typically set up the command to run automatically on a regular schedule. Disk migration optimizes the use of disk storage, while providing excellent performance for short-term retrievals of reports. As a report ages, and in all likelihood accesses becomes less frequent,

OnDemand can automatically copy the report to long-term (archive) storage. You can also use disk migration to defer the loading of reports to archive storage to a time when there is little or no other system activity.

- Disk expiration processing

Disk expiration is the process of deleting reports from disk storage. You specify how long a report should remain on disk when you create an application group. OnDemand provides a command that you can use to delete reports from disk. You typically set up the command to run automatically on a regular schedule. Disk expiration reclaims disk space taken by expired reports.

Applications

An OnDemand application describes the physical characteristics of a report, processing instructions for the indexing and data loading programs, and information about how OnDemand displays and prints pages of a report. You can specify default settings for viewing and printing pages of a report at the OnDemand application level. For example, if you select a default printer for the application, when a user prints a document that is associated with the application, OnDemand sends the document to the printer that you specified. Typically you define an application for each different report that you plan to load into the system.

When you create an application, you specify properties of the input data (such as whether the data contains carriage control characters, and the record format of the input data). OnDemand uses the information that you specify to properly interpret the data for viewing.

The OnDemand application is where you specify information to the indexing and data loading programs, such as the technique that OnDemand uses to compress the report file, the parameters used to index the data, and information that OnDemand uses to process index data before loading index records into the database. OnDemand uses the indexing parameters, options, and data values that you specify to locate index data in and extract index data from the report.

You can set up one or more *logical views* of a report. A logical view determines how OnDemand displays line data reports and governs other viewing characteristics. For example, you can set up a logical view so that when a user selects a document for viewing, the OnDemand client program automatically locks the heading of the report in place when the user scrolls up or down on a page.

Users and groups

Each user logs on to OnDemand with a userid. OnDemand authenticates userids and determines the usage and administrative authority available to the user based on the log on userid. An OnDemand userid does not necessarily have to identify an individual user. However, for accounting and security purposes, you probably want to assign an OnDemand userid to each person that uses the system.

OnDemand automatically creates the QONDADM userid when you initialize the system. The QONDADM userid has system administrator authority. A system administrator can perform the basic user functions, such as logging on the system and opening folders, and administrative functions, such as defining users and groups and creating, updating, and deleting application groups, applications, folders, storage sets, and printers.

OnDemand groups are a means to organize users by function, authorization, or any other purpose you might require. When you define an OnDemand group, you can organize users by department or function and set folder and application group permissions that are common to all of the users assigned to the group. The permissions determine the types of actions that users assigned to the group can perform. You do not have to assign a user to a group, but doing so can simplify administration of users with similar requirements and capabilities. **Note:** OnDemand groups are not the same as OS/400 group profiles, but you may choose to use the same names if that is easier to maintain.

OnDemand userid relationship to OS/400 user profiles

When you install OnDemand, a default logon security user exit is enabled that forces a relationship between your OnDemand userids and your OS/400 user profiles. If you use OnDemand as installed, then the following is true:

- An OnDemand userid must match an OS/400 user profile.
- The password when logging on to OnDemand is the same as the password used when logging on to OS/400. If you change the password using OnDemand, it actually changes OS/400 password. When you create an OnDemand userid, you actually leave the password field blank.
- Maximum Password Age should be set to Password Never Expires. If you specify a value for Maximum Password Age, then OnDemand may force a user to change their password before it is required by OS/400.
- Minimum Password Length should be set to Allow Blank Password. This prevents OnDemand from trying to impose its own rules on the length of a password and allows OS/400 to use its own rules.

If the security exit is not enabled, then the OnDemand userid and password have no relation to the OS/400 userid and password and all the OnDemand System Parameter settings are honored.

Note: Enabling or disabling of this exit can be done at an individual instance level. See Chapter 12, “Creating an instance,” on page 69 for more information on instances.

Permissions

Overview

As both a convenience and security measure, you can assign a user to a group. When you assign a user to a group, the user obtains the permissions of the group. For example, suppose you create a group and authorize the group to open the Student Information folder. Any user that you assign to the group automatically obtains permission to open the Student Information folder.

If you assign a user to more than one group, the user normally obtains the permissions of all of the groups. For example, using the group settings listed in Table 1, a user assigned to both groups can access the Student Bills and Student Transcripts folders.

Table 1. Group permissions

Group	Folders
Accounting	Student Bills
Admissions	Student Transcripts

However, there are exceptions to this rule. See information about permissions in the *IBM Content Manager OnDemand for iSeries Common Server Administration Guide* for details.

You can set folder and application group permissions for every user and group defined to OnDemand. If you set permissions for a specific group, the group permissions take precedence over the permissions set at the folder level or the application group level. If you set permissions for a specific user, the user permissions take precedence, regardless of any group that includes the user or the permissions set at the folder level or the application group level.

You can set folder and application group permissions when you add or update a folder or application group. You can also set folder and application group permissions when you add or update a user or a group.

Folder permissions

You can set folder permissions at the folder, group, and user levels. Setting permissions at the folder level provides all OnDemand users and groups that are not otherwise given permissions with the permissions that you define. Setting permissions at the group level provides all of the users that you assign to the group with the permissions that you define. Group level permissions override folder level permissions. Setting permissions at the user level provides a specific user with the permissions that you define. User level permissions override group level permissions and folder level permissions.

By default, only the user that created the folder, users with administrator permission for the folder, application group/folder administrators, and system administrators can access the folder.

You can set the following types of folder permissions:

- **Access.** Users can open the folder with OnDemand client programs and search for and retrieve data from the application groups referenced in the folder. To search for and retrieve items, users must have access permission for the folder, and access permission to one or more of the application groups referenced in the folder.
- **Fields.** Users can open the folder with OnDemand client programs and can modify the folder field information with the administrator interface. OnDemand maintains a set of folder fields for each user given fields permission for the folder.
- **Named Queries.** A named query is a set of search criteria, saved by name, that can be selected and restored into folder search fields. OnDemand supports two types of named queries: public, that is, a named query that is available to all users that can open the folder, and private, that is, a named query available only to the user that created the named query. Users can be given authority to view, create, modify, and delete named queries.
- **Administrator.** A folder administrator can modify and delete the folder. A folder administrator can change user and group permissions, add and remove users and groups from the folder, and make changes to the folder field information.

Application group permissions

You can set application group permissions at the application group, group, and user levels. Setting permissions at the application group level provides all

OnDemand users and groups that are not otherwise given permissions with the permissions that you define. Setting permissions at the group level provides all of the users that you add to the group with the permissions that you define. Group level permissions override application group level permissions. Setting permissions at the user level provides a specific user with the permissions that you define. User level permissions override group level permissions and application group level permissions.

By default, only the user that created the application group, users with administrator permission for the application group, application group/folder administrators, and system administrators can access the application group.

You can set the following types of application group permissions:

- Access. Users can search for and retrieve data stored in the application group using OnDemand client programs.
- Document. Determines the types of document functions users can perform. The default document permissions are view, print, fax, and copy.
- Annotation. Determines the types of annotation functions users can perform. The default annotation permissions are view and add.
- Logical Views. Logical views determine how OnDemand displays report file pages. Users can define their own logical views with OnDemand client programs.
- Administrator. An application group administrator can modify and delete the application group. An application group administrator can change user and group permissions, add and remove users and groups from the application group, change message logging options, update the storage management settings for the application group, and make changes to the application group field information.
- Query restriction. Limits access to application group data. You typically set up a query restriction to limit the data that a specific user or group of users can access.

Naming rules

When you create objects in OnDemand, you assign names to the various objects.

Note: If you install OnDemand with a language that requires multiple bytes per character (for example, Kanji), the number of characters permitted for a name is less than the number listed in the sections that follow.

When naming a user, the name that you specify:

- Can contain from one to 128 characters (bytes), depending upon whether your OnDemand user names are linked to your OS/400 user profile names. If you are using OnDemand as shipped, your OnDemand user names are linked to your OS/400 user profile names. This is the default setup and desirable for most customers. However, you can change this if you like, so that there is no relationship between your OnDemand user names and your OS/400 user profile names. See “OnDemand userid relationship to OS/400 user profiles” on page 40 for more details.
- Cannot include the ' (apostrophe), * (asterisk), % (percentage) + (plus), _ (underscore), [(left bracket),] (right bracket), “ (double quote), or blank characters
- Must be unique

- By default, OnDemand converts lowercase letters in a user name to uppercase (for example, laguarde is stored as LAGUARDE)

Note: If you define your own logon user exit, then you can determine the characteristics of userids on your system.

When creating a password, the value that you specify:

- Can contain from one to twenty characters (bytes)

Note: When creating a password, the value that you specify can be a maximum of 20 characters. However, the password authentication that is built into OnDemand verifies only the first eight characters that are entered by the user. The additional characters are provided for customers who choose to implement their own password security by using the logon user exit. Contact the IBM support center for more information about the logon user exit.

- By default, OnDemand converts lowercase letters in a password to uppercase (for example, laguarde is stored as LAGUARDE)

Note: If you define a logon user exit, then you can determine the characteristics of passwords on your system.

When naming a group, the name that you specify:

- Can contain from one to 128 characters (bytes)
- Cannot include the ' (apostrophe), * (asterisk), % (percentage) + (plus), _ (underscore), [(left bracket),] (right bracket), " (double quote), or blank characters
- Must be unique
- Can be mixed case; however, OnDemand ignores the case (for example, LaGuarde is the same as laguarde)

When naming an application group, application, or folder, the name that you specify:

- Can contain from one to sixty characters (bytes), including embedded blanks. If you choose to use the output queue monitor function of OnDemand to automatically load your reports, then you should plan to name your applications and application groups such that the first ten characters match the spooled file name, userdata, formtype, or one of the other supported spooled file attributes. See the *IBM Content Manager OnDemand for iSeries Common Server Administration Guide* for more details about output queue monitors.
- Cannot include the ' (apostrophe), % (percentage), _ (underscore), [(left bracket),] (right bracket), or " (double quote) characters
- Can be mixed case; however, OnDemand ignores the case (for example, LaGuarde is the same as laguarde)
- An application name must be unique to the application group where you assign the application
- An application group or folder name must be unique to the server

When naming a database field, the name that you specify:

- Can contain from one to eighteen characters (bytes)
- Must begin with the letter A through Z

- Can include the letters A through Z, the numbers 0 through 9, and the @ (at sign), \$ (dollar), _ (underscore), and # (number sign)
- Can be mixed case; however, OnDemand doesn't create a unique name (for example, rDate is the same as rdate)
- Must be unique to the application group
- Cannot be any of the OnDemand reserved words:

annot	doc_off
comp_len	doc_type
comp_off	pvt_nid
comp_type	resource
doc_len	res_comp_type
doc_name	sec_nid
- Cannot be any of the words reserved by the database manager. (For a list of reserved words, see Appendix D of the *DB2 Universal Database for iSeries SQL Reference*.)

When naming a logical view, the name that you specify:

- Can contain from one to thirty characters (bytes)
- Can be mixed case
- A public view must be unique to the application
- A private view must be unique to the user

When naming a folder field, the name that you specify:

- Can contain from one to sixty characters (bytes), including embedded blanks
- Cannot include the ' (apostrophe), % (percentage), _ (underscore) [(left bracket),] (right bracket), or " (double quote) characters
- Can be mixed case
- Must be unique to the folder

When naming a storage set or migration policy, the name that you specify:

- Can contain from one to 60 characters (bytes)
- Can be mixed case; however, OnDemand ignores the case (for example, LaGuarde is the same as laguarde)
- Must be unique

When naming a server printer, the name that you specify:

- Can contain from one to sixty characters (bytes)
- Can be mixed case; however, OnDemand ignores the case (for example, LaGuarde is the same as laguarde)
- Must be unique to the server

When naming a server printer queue, the name that you specify:

- Must be a valid output queue name on the server, entered in the format of *libraryname/outq*

Data types and field types

When you define an application group, OnDemand creates a structure for a database table with the index and filter fields that you define. When you store a report in the application group, OnDemand extracts index data from the report, places the index data into the database fields, and inserts rows into the application group table. The database fields that you define for the application group can

contain different types of data. When you define the database fields, you select a data type for each field. The data type tells OnDemand what kind of data can be stored in the field.

When you define a folder to OnDemand, the fields that you define can be used in two ways:

- For search fields, in which users enter values to construct queries
- For display fields, to identify the items in the document list

Table 2 lists the types of application group and folder fields supported by OnDemand.

Table 2. Application group and folder field types

Field Type	Description
Small Integer	Contains whole numbers between $-32,767$ and $32,767$
Integer	Contains whole numbers between -2147483648 and 2147483647
Big Integer	Contains whole numbers between -922337036854775807 and 922337036854775807 . Big integer fields hold a 64-bit integer representation of a number or a character string in the form of an integer constant.
Decimal	Contains numbers between -10^{307} and 10^{308} with up to 15 significant digits. You can store currency values in a decimal field, and use the precision attribute to format the decimal places.
String (Fixed)	Contains letters, numbers, special symbols, such as the % and #, and any other printable character. A fixed length string field requires one byte per character declared; unused characters do consume storage.
String (Variable)	Contains letters, numbers, special symbols, such as the % and #, and any other printable character. A variable length string field requires one byte per character plus four bytes for overhead; unused characters do not consume storage.
Date	Contains a valid date from January 1, 1970 to December 31, 2069. If you need to index reports that contain dates that occur before January 1, 1970, you must define the date as a String field. OnDemand checks a date value to make sure it is valid.
Time	Contains times of day, stored in three-second increments, since midnight, and limited to 24 hours
Date/Time	Contains both a date and time value. The date can be from January 1, 1970 to December 31, 2038. The time is stored in one-second increments.
Date/Time (TZ)	Contains both a date and time value. The date can be from January 1, 1970 to December 31, 2038. A Date/Time (TZ) field is exactly like a Date/Time field, but uses the time zone set on the client PC.

Table 3 lists additional types of fields that are supported in folders.

Table 3. Additional folder field types

Field Type	Description
Annotation Color Search	Use to search annotations to a document by specifying a color. A match occurs and an item is added to the document list if the color of the text in one or more of the annotations to a document is the same as the color that is specified in the search field. A folder can have one annotation text search field.

Table 3. Additional folder field types (continued)

Field Type	Description
Annotation Text Search	Use to search annotations to a document for the specified string. A match occurs and an item is added to the document list if one or more of the annotations to a document contain the text that is specified in the search field. A folder can have one annotation text search field.
Application Group	For a search field, contains a list of the application groups that can be searched from the folder. When you create a folder that contains more than one application group, you can define an application group field. If enabled for queries, users can select the name of the application group that OnDemand searches, rather than searching all of the application groups contained in the folder (the default). For a display field, lists the name of the application group in which the document was found. A folder can have one application group field.
Segment	Contains a list of the tables of index data that are stored in the application groups that can be searched from the folder. Each item in the list represents a segment of application group data. OnDemand segments application group data by date. If enabled for queries, users can select a specific segment of application group data to search. A folder can have one segment field.
Text Search	<p>Used to find documents that contain a non-indexed word or phrase. A match occurs and an item is added to the document list when one or more lines in a document contain the word or phrase exactly as specified. The search string can contain letters, numbers, special symbols, such as the % and #, and any other printable character. A folder can have one text search field.</p> <p>Note: The (sequential) text search takes place on the server. A text search will delay the generation of the document list. Only documents that meet all of the criteria specified in the other folder fields will be searched for the specified word or phrase.</p> <p>A typical use of a text search field is to provide users an additional search field without incurring database overhead. For example, assume that a report is indexed on date and transaction number. A text search field would allow users to optionally enter a customer's name, phone number, or any other information contained in the document or documents that the user needs to retrieve (the information is not contained in the database). However, a text search field has a direct impact on the time it takes to generate the document list and the performance of the server. A large number of users performing text searches at the same time can usually drain the resources of even the most powerful server.</p>

Chapter 7. Storage requirements

Overview

Estimating storage requirements for an OnDemand system begins with understanding and documenting end-user requirements for storing and accessing data. Chapter 5, “Reports and other data,” on page 29 provides information that can help you gather end-user requirements.

Before you turn requirements into a storage subsystem to support your system, you must also review the various operational and performance issues. For example, OnDemand supports up to 32 index fields for each report. However, users should not need a lot of indexes to locate a specific version of a report or a document within a report. The number of index fields that you define has a direct impact on the amount of disk space that you will need for your database. In addition, the more indexes that you define for a report, the longer it will take to load the report into the system. It is important to work with users and understand their data retrieval requirements. Define only the number of index fields that they need. You may have to balance end-user requirements with disk space, the amount of time required to load a report, and other performance issues.

Maintaining a copy of reports on disk can have a significant impact on the amount of disk storage that you need on your system. Most customers store the latest versions or most frequently accessed reports on disk. You should review how users search for and retrieve information from the reports that you plan to store in OnDemand. For example, if most retrievals occur in the first 90 days after a report is generated, then you probably want to store the report on disk for at least that length of time. You should choose a time frame for each report which meets the requirements of your users and also makes the best use of available disk space.

There are several components that you need to measure to determine the amount of disk, optical, and tape storage required to support an OnDemand system. For example, the following components of the system require disk storage:

- Storage space for application programs and system software, including the base operating system, the OnDemand server software, and any other applications on the iSeries that are required to run your business.
- Storage space for configuration files and control files.
- Storage space for the OnDemand system logging facility.
- Temporary storage space for reports received from other systems. In general, you should plan for enough disk space to hold either the largest single report that you will be loading on the system or the total of several reports that may be staged for loading at the same time, whichever requires the most storage space. In many organizations, most versions of a report are similar in size. However, there may be times when a report is much larger than average. For example, a report generated at the end of the month or the end of the quarter may greatly exceed the average report size.
- Temporary storage space for indexing a report on the OnDemand server.
- Temporary storage space for loading a report on the OnDemand server.
- Storage space for reports stored on disk. This may be zero, for reports that do not require disk storage. However, a very large amount of disk space may be required for reports that must remain on disk for several months or longer.

OnDemand compresses report data before storing it. The compression ratio can have a significant impact on the amount of disk space that you need to store a report. OnDemand can achieve up to 30:1 compression on line data reports. However, for reports that contain AFP data or image data that is already compressed, the compression achieved will be much lower.

- Storage space for the database, which includes OnDemand system tables (control information and objects that you define to OnDemand) and application group tables (index data extracted from reports). The amount of database space that you should plan for a report is a factor of the number of documents or sections contained in the report, the number of index fields that you define for the report, the number of versions of a report (the frequency with which you load a report on the system), and how long you need to maintain a report on the system.

For reports that contain sorted transaction data, OnDemand can divide the report into groups of a fixed number of pages and create one index row for each group of pages.⁶ For reports that contain logical items, such as statements, and policies, OnDemand can create one index row for each logical item in the report. Typically the database space required for indexing sorted transaction data is much less than the database space required for indexing reports that contain logical items. Also, index fields provide fast lookup, but require a significant amount of database space.

- Storage space for the OnDemand journal and journal receivers. You should plan for disk space for the OnDemand journal and receivers used for commitment control.
- Temporary storage space for server print and fax.
- Temporary storage space for importing migrated indexes from archive media to the database.

The following components of the system require archive storage (optical and tape storage):

- Reports that you plan to store on archive media.
- Backup copies of reports stored on archive media. (For critical applications, some customers require that the system maintain two or more copies of a report on archive media.)

You can replace full optical storage volumes with new empty ones in the optical library as needed, if the availability requirements of your system allow you to do so. For example, you may decide to remove full storage volumes from a library one year after the last time that OnDemand wrote report data to the storage volume. You could replace the full storage volumes with newly initialized storage volumes to hold the latest reports stored on the system. That way, the latest versions of a report are always available in the library. However, if you need to keep many years of report data online in the library or you store massive amounts of data in your application groups, then you may need to plan on having several optical libraries for your system.

Storage hierarchy

There are several different storage management strategies that you can use with OnDemand.

6. For sorted transaction data, the examples and calculations that follow assume that OnDemand will create one indexed item for each group of 100 pages in a report. The number of pages in a group is a parameter that you can configure when you index a report. The *IBM Content Manager OnDemand for iSeries Common Server: Indexing Reference* provides more information.

For example, OnDemand allows data to be migrated from one storage media type to another using criteria defined by an administrator. Disk storage should be used for high-speed access to reports. When you load a report on the system, OnDemand can automatically store one copy of the report on disk and another copy of the report on archive storage media. OnDemand also supports the option of storing reports on disk and then later migrating them to archive storage. However, we recommend that you always plan to copy reports to disk and archive media at the same time (when you load the report). Doing so usually eliminates the need for you to periodically backup disk storage, because a backup copy of your reports already exists on archive media. Copying reports to cache storage and archive storage at the same time also eliminates the need for you to migrate reports to archive media later.

Reports *expire* (are eligible to be removed) from disk when they reach their disk expiration date. You specify the disk expiration date for a report when you create an application group. For example, you can specify that a report should expire from disk after it has been stored there for ninety days. OnDemand provides a command that you can use to automatically remove expired reports from disk on a regular schedule. After you run expiration processing, OnDemand reclaims the space taken by expired documents.

OnDemand maintains documents on disk and on archive media independently of each other. Each use their own criteria to determine when data expires and should be removed from the system; each use their own commands to remove documents. However, for removal of documents from the system, you should specify the same criteria for both. For example, the sum of the Duration values for all migration policy storage levels plus the value for the Days in Cache should equal the Life of Data and Indexes value specified in the application group definition.

Data compression

OnDemand can compress report data using several different data compression algorithms, before storing the data on disk and archive storage. The compression ratio that OnDemand can achieve has a significant impact on the amount of space required to store reports.

The compression ratios that OnDemand can achieve vary widely depending on the type of data and the format of the data. You cannot always accurately estimate the compression ratio by simply examining the data. On average, you can expect to achieve between 2:1 and 15:1 compression for AFP documents and up to 30:1 compression for line data reports. Compression for AFP documents is based on the output data file produced by the indexer, and not the input file, which could have been line data. When the indexer formats line data with a page definition, it may increase the size of the data by adding AFP controls for positioning text.

To properly estimate the amount of storage space required by a report, we recommend that you measure the compression ratio achieved by actually storing a sample of the report. You can store a report by using the ADDRPTOND command. The ADDRPTOND command will place entries in the system log, which will provide before and after size values that can be used to calculate the actual compression ratio achieved.

See the *IBM Content Manager OnDemand for iSeries Common Server Administration Guide* for more information about the ADDRPTOND command.

Calculating disk storage requirements

System software

OnDemand servers require disk space for a wide range of software products to support the business applications running on them. This includes the operating system software, business application software, temporary work space, and the OnDemand server software.

Temporary space for indexing

OnDemand requires temporary storage space on disk to index reports. The temporary space required by OnDemand is a factor of the largest occurrence of a report and the number of reports that you plan to index at the same time.

Use the following calculation to determine the amount of temporary space required to index reports:

$$\text{Temporary space} = \text{Largest report file size} * 1.5$$

Figure 7. Calculating temporary space for indexing

where Largest report file size is the size in bytes of the largest version of a report to be indexed or the total size of all of the reports that the server must index at the same time (if you index more than one report at a time).

For example, if the largest report is 400 MB, then the temporary space required to index the report is:

$$\text{Temporary space} = 400 \text{ MB} * 1.5 = 600 \text{ MB}$$

Disk storage for stored reports

The amount of disk space that you should dedicate to stored report data will vary greatly based on requirements such as the number of reports that you store on the system, the compression ratio that OnDemand can achieve, and the amount of time that you need to keep a report on disk. Most customers store reports on disk for a short period of time, to provide the fastest retrieval for the most frequently used reports. As reports age, and retrieval requests for them are much less frequent, the reports can be retrieved from archive media. Another reason to keep reports on disk is if many users access them at the same time. Because the archive storage manager may require from six and sixty seconds to mount an optical or tape storage volume and retrieve a report, it is usually not possible to support a high transaction rate for reports stored on archive media.

Another use of disk storage is for reports that have a short life, such as one week or one month. You can store these types of reports on disk and OnDemand will automatically delete them when they reach their expiration date. Disk storage can also be used to hold reports for which you do not need a backup copy.

Use the following calculation to determine the amount of disk space required for stored reports:

$$\begin{aligned} \text{Disk space required} &= \text{Size of Data per week} \\ &\quad * \text{Number of Weeks on disk} \\ &\quad * \text{Data Compression ratio} \\ &\quad * 1.1 \end{aligned}$$

Figure 8. Calculating cache storage

For example, if you plan to load 2 GB of report data on the system each week, the reports must be maintained on disk for 12 weeks, and the compression ratio is 3:1 (0.33), then the disk space required for cache storage space can be calculated as follows:

$$\begin{aligned} \text{Disk Space} &= 2 \text{ GB} * 12 * .33 * 1.1 = 8.71 \text{ GB} \\ \text{Required} & \end{aligned}$$

OnDemand database storage

When you load a report into the system, OnDemand extracts index data from the report and stores it in an application group table in the database. For reports that contain logical items, such as statements and policies, OnDemand can create one database row for every item found in the report. For reports that contain sorted transaction data, OnDemand can create one database row for every indexed group of pages (by default, 100 pages in a group).

A database row contains a fixed amount of information that OnDemand uses to maintain reports (approximately 40 bytes) and any additional index and filter fields that you define for the application group. Index fields, which allow users to locate documents quickly, require significantly more disk storage space than filter fields. (Index fields also require more time to load into OnDemand.)

There are four major factors that determine the amount of disk space required for the OnDemand database:

- The number of index and filter fields
- The size of the index and filter fields
- The number of indexed items per month
- The number of months that OnDemand maintains the index data in the database

Table 4 lists the types of index fields supported by OnDemand and the number of bytes required to store a value in each type of index field.

Table 4. Index field types and sizes

Field Type	Field Size
Small Integer	2 bytes
Integer	4 bytes
Big Integer	8 bytes
Decimal (floating point)	8 bytes
String (Fixed)	1 - 254; 1 byte per character declared, even if partially used
String (Variable)	25 - 254; 1 byte per character plus 2 bytes overhead; unused characters do not consume storage
Date	4 bytes
Time	4 bytes

Table 4. Index field types and sizes (continued)

Field Type	Field Size
Date/Time	4 bytes
Date/Time (TZ)	4 bytes

Database journal storage

The OnDemand database includes recovery journal and journal receivers which are used to recover from application or system errors. In combination with database backups, they are used to recover the consistency of the database right up to the point in time when an error occurs. When you install OnDemand, a journal called QSQJRN in library QUSROND (the default instance library name; a journal will be created in your particular instance library if you create additional instances). A journal receiver called QSQJRN0001 is automatically created as the first receiver for the journal. QSQJRN is a system-managed journal, which means that the operating system manages the changing of journal receivers as needed. Messages related to this journal are sent to the QSYSOPR message queue.

Temporary space for importing index data

Note: If you do not plan to migrate index data from the database to archive media, then you do not need to allocate temporary storage space for importing the migrated index data. It is important to understand that migration of index data is **not recommended**. Therefore, few customers should need to read this section.

OnDemand requires temporary work space to import migrated index data from archive media into the database. You must allocate enough disk space to support the maximum number of concurrent import requests that the server must manage. The amount of space that you allocate is based on the size of your application group tables and the number of tables that you must import to satisfy a query for migrated data.

Calculating archive storage requirements

Report storage space

When you estimate the amount of space required to store a report in archive storage, you must consider the size of the report, the compression ratio achieved, and the length of time that the archive storage manager maintains the report. Archive media can be optical storage or magnetic tape. Use the following calculation to estimate that amount of space required:

$$\text{ArchiveStorageSpace} = (\text{Data per month} * \text{life of data in months}) * \text{compression ratio} * 1.1$$

Figure 9. Calculating archive storage space

For example, if you plan to store 8 GB of report data per month, the archive storage manager must maintain the data for seven years, and OnDemand can achieve a compression ratio of 3:1 (0.33), you would require approximately 244 GB of archive storage space:

$$\begin{aligned} \text{OpticalSpace} &= (8 \text{ GB} * 84) \\ &\quad * 0.33 \\ &\quad * 1.1 = 244 \text{ GB} \end{aligned}$$

Backup report storage space

The OnDemand system can maintain a backup (second) copy of reports that you store on archive media. You typically maintain multiple copies of reports that are critical to the operation of your company or difficult or impossible to recreate.

If you need OnDemand to maintain a backup copy of your reports, double the archive storage space that you calculated in “Report storage space” on page 52.

Migrated index storage space

Note: It is important to understand that migration of index data is **not recommended**. Therefore, few customers should need to read this section.

OnDemand supports automatic migration of indexes from the database to archive storage so that you can maintain seldom used indexes for long periods of time. However, migration of indexes should be done only after there is no longer a need to retrieve the reports to which they point. For example, suppose that all of the queries for a report occur in the first 24 months after the report is loaded into the system. After that time, there are almost no queries for the report. The indexes could be eligible to be migrated from the database to archive storage. Migration of index data is optional; you can choose to migrate indexes for all, some, or none of the application groups on your system. In addition, you determine the length of time that indexes stay in the database before OnDemand migrates them to archive storage.

You can use the following calculation to determine the archive storage space required to hold migrated indexes:

$$\begin{aligned} \text{ArchiveMediaDBSpace} &= (\text{Database size per month} * \text{compression ratio}) \\ &\quad * (\text{life of data} - \text{months before migrating data}) \end{aligned}$$

Figure 10. Calculating migrated index storage space

For example, if the database size is 202 MB per month, you need to maintain the indexes for 84 months, and the indexes remain in the database for 24 months before being migrated, then the archive storage required to hold the migrated indexes is:

$$\begin{aligned} \text{ArchiveMediaDBSpace} &= (202 \text{ MB} * .33) \\ &\quad * (84 - 24) = 4 \text{ GB} \end{aligned}$$

Chapter 8. Backup and recovery

Overview

This section of the book describes backup and recovery for OnDemand and provides recommendations about methods and procedures that an administrator can use to make sure that the following critical OnDemand components can be recovered when needed:

- OnDemand software
- OnDemand server information, created or modified during installation, configuration, and ongoing operation of OnDemand
- The OnDemand database
- Archived reports

See Chapter 5 of the *IBM Content Manager OnDemand for iSeries Common Server: Administration Guide* for specific libraries and objects to be backed up on a regular basis.

Server software

If a media failure or some other unforeseen event occurs, you may be required to restore the OnDemand software programs, and other application and user-defined software that you use on the system.

It is important that you store the backup media for these software products in a safe location. We recommend that you register OnDemand as part of your business recovery plan and store the product media in the same place that you store the other programs and files that are vital to the operation of your systems.

Server information

When you installed and configured OnDemand, you specified information that customized OnDemand to operate in your environment. This information is stored in control files located in various IFS directories on the server. We recommend that you backup the control files immediately after you have verified the installation of OnDemand. In addition, if you periodically make changes to the OnDemand server information, we recommend that you backup the control files on a regular basis.

See the operating system and device publications for your server for details about backup and restore concepts and commands.

OnDemand database

Database backup

The operating system provides commands so that you can create backup copies of the OnDemand database.

If your production schedule allows, we strongly encourage you to create offline backups on a regular schedule, perhaps once a night or once a week. Regularly

scheduled offline backups can reduce the time required to rebuild the database, if you need to do so. Keep backup media in a safe place, until the next time that you create an offline backup of the database.

If your schedule does not provide time to take offline backups (that is, your system must always remain available to users), you should take online backups on a regular schedule. See the Backup and Recovery Guide for details about backing up a database.

Database journaling

The database manager uses journaling to record information about changes to the OnDemand database. The information in the journal is used to recover from corruption of data in the database. Journaling ensures that no data is lost. By combining the information in the journal with a backup copy of the database, the OnDemand database can be recovered to any point in time.

Database recovery

There are two types of database recovery. The first type recovers from failures that occur while update transactions are taking place. The journal helps correct this type of failure by allowing the transactions received before the failure to either be reapplied to the database or to be *rolled-out*. Rolling-out transactions is a way to return the database to the state it was in before the transaction that caused the failure.

The second type of recovery deals with corruption of the OnDemand database and is usually caused by media failure. The combination of journals and a backup copy of the database can be used to recreate the OnDemand database at a particular point in time.

If a catastrophic failure occurs, the system administrator will need to intervene to recover the database. Recovery from catastrophic failure starts with restoration of the latest full backup copy of the database. Next, the system administrator reapplies the transactions recorded in the journal(s). These steps will recreate the OnDemand database before the catastrophic failure.

Reports

OnDemand can store copies of reports on disk and archive storage:

- The primary purpose of disk storage is short-term, high-speed storage and retrieval of reports.
- The primary purpose of archive storage is long-term storage and retrieval of reports. Reports in archive storage can also be used as backup copies, in the event that disk storage becomes corrupted or unavailable. Archive storage consists of optical or tape storage volumes managed by the archive storage manager.

Most customers configure the system to copy reports to disk and archive storage at the same time, when they load a report into the system.

OnDemand can retrieve a copy of a report from archive storage after the report has been removed from disk or if the copy on disk is unavailable. However, you must configure the system to support multiple copies of reports. You must define devices to the archive storage manager, and configure OnDemand to use archive storage. You configure OnDemand to use archive storage by defining migration

policies and storage sets, assigning application groups to the storage sets, and configuring data migration in application groups.

Disk storage

Disk storage is the primary, short-term storage location for reports.

If you do not copy reports to archive storage when you store them in OnDemand, then you need to consider how you can recover the reports in the event that you need to do so.

Disk storage can be protected by maintaining it on RAID storage subsystems. RAID storage can provide excellent availability, allowing users to access reports even if a disk or controller fails. However, RAID storage is not fail-safe. There may be situations when, because of multiple disk or controller failures, users cannot access reports. We encourage you to always maintain a backup copy of reports in archive storage.

Archive storage

The OnDemand migration policy identifies where the primary copy of a report is maintained. OnDemand retrieves the primary copy of the report from archive storage after the report has been removed from disk. Customers with special business, legal, or performance reasons may want the system to maintain a backup copy of their reports in archive storage. The backup copy can be used if the primary copy becomes corrupted or unavailable.

Chapter 9. Saving configuration files

When you install software on an OnDemand server, the installation programs copy program files, configuration files, and other types of files from the distribution media to directories on the server. When you configure a server to meet the specific requirements of your environment, you make changes to configuration files and you may customize other files, such as user-defined files and font initialization files.

After you make changes to any of your configuration files or before you upgrade to a new version of OnDemand, we recommend that you save a copy of the files listed in this section. See Chapter 5 of the *IBM Content Manager OnDemand for iSeries Common Server: Administration Guide* for specific libraries and objects to be backed up on a regular basis.

OnDemand files

Save a copy of the OnDemand configuration files listed in Table 5. **Note:** Replace <instance> with the actual name of the <instance> to which the particular file belongs.

Table 5. OnDemand configuration files to save

File	Default Location	Purpose
ars.cache	/QIBM/UserData/OnDemand/<instance>/config	Define disk storage details.
ars.cfg	/QIBM/UserData/OnDemand/<instance>/config	OnDemand server configuration file. Changes described in Chapter 14, "Configuring the ARS.CFG file," on page 79.
ars.ini	/QIBM/UserData/OnDemand/config	Configure OnDemand instances. Changes described in Chapter 12, "Creating an instance," on page 69.
ars.dbfs	/QIBM/UserData/OnDemand/<instance>/config	Not used, but file must exist.
arsload.cfg	/QIBM/UserData/OnDemand/<instance>/config	Define a default OnDemand system administrator userid and password for the ARSLOAD program. Changes described in Appendix B of the <i>IBM Content Manager OnDemand for iSeries Common Server Administration Guide</i> (SC27-1161)

Table 5. OnDemand configuration files to save (continued)

File	Default Location	Purpose
arslog	/QIBM/ProdData/OnDemand/bin	The symbolic link to the default System Log user exit program. Described in Appendix E, "System log and the system log user exit," on page 135.

Windows font files

If you have modified any font-related files on your Windows client workstations, save a copy of the files listed in Table 6. When you install a new level of the client software, these files are overwritten and any modifications you have made will be lost. For this reason, you may want to make a backup copy of any files you modify.

The location shows the default installation directory for the Windows client.

Table 6. User-defined code page files (Windows client)

File	Location	Purpose
ICODED.FNT	\Program Files\IBM\OnDemand32\Font	Coded Font file. Contains the list of coded font files used by the Windows client program.
CSDEF.FNT	\Program Files\IBM\OnDemand32\Font	Character Set Definition file. Contains the list of character sets used by the Windows client program.
CPDEF.FNT	\Program Files\IBM\OnDemand32\Font	Code Page Definition file. Contains the list of code page files used by the Windows client program.
*.CP	\Program Files\IBM\OnDemand32\Font\Maps	Code Page Map file. Any user-defined Windows code page map file added to OnDemand.
ALIAS.FNT	\Program Files\IBM\OnDemand32\Font	The Alias file. Used to map AFP fonts to Type 1 or Truetype fonts.

Part 4. Installing software

This section contains instructions for installing the OnDemand server software on the iSeries server.

Chapter 10. Installation checklist

Setting up your OnDemand system typically requires that you do the following:

- ___ 1. Contact the IBM Support Center for the latest PTFs for OnDemand. The list of current PTFs can be found in Informational APAR number III13680. You may also want to consider requesting current PTFs for OS/400 (DB2 in particular) and iSeries Access Express, because these two products are an integral part of the OnDemand system.
- ___ 2. Obtain a copy of the latest OnDemand Read This First document from the web at <http://www.ibm.com/software/data/ondemand/400/library.html>. Print and read the entire file before you begin.
- ___ 3. Check the OnDemand prerequisites (see Chapter 4, "Hardware and software," on page 21) and verify the required and optional hardware and software products.
- ___ 4. If you are upgrading to a new version of OnDemand, save the configuration files used by the system (see Chapter 9, "Saving configuration files," on page 59).
- ___ 5. Install the OnDemand software on the iSeries server (see Chapter 11, "Installing OnDemand server software," on page 65).
- ___ 6. Configure the OnDemand software (see Part 5, "Configuring and initializing the system," on page 67). This step includes the following:
 - ___ a. Create an instance (see Chapter 12, "Creating an instance," on page 69)
 - ___ b. Define a locale (see Chapter 13, "Defining a locale," on page 75)
 - ___ c. Configure the ARS.CFG file (see Chapter 14, "Configuring the ARS.CFG file," on page 79)
 - ___ d. Configure system startup (see Chapter 15, "Configuring system start up," on page 81)
 - ___ e. Define schedules for maintenance programs (see Chapter 16, "Scheduling maintenance programs," on page 83)
- ___ 7. Create and initialize the database on the server (see Chapter 12, "Creating an instance," on page 69)

Verify the installation of OnDemand (see Chapter 17, "Verifying the installation," on page 87):

- ___ 1. To access the system, you must install the OnDemand client program on a PC. See the *IBM Content Manager OnDemand: User's Guide* for details about installing OnDemand client software.
- ___ 2. Log on to the server with an OnDemand client program.

Prepare the system for use:

- ___ 1. Define migration policies and storage sets (see Chapter 18, "Defining migration policies and storage sets," on page 89). Before you add application groups or load data into the system, you must define migration policies and storage sets.
- ___ 2. Configure the System Log application group (see Chapter 19, "Configuring the System Log application group," on page 91). Before you define reports to the system, load data, or let users access the system, you should configure the System Log application group.

- ___ 3. Configure the System Migration application group (see Chapter 20, “Configuring the System Migration application group,” on page 95). If you plan to migrate index data to archive storage, then you must configure the System Migration application group.

Note: Migrating index data to archive storage is not recommended.

- ___ 4. Backup the databases (see Chapter 21, “Backing up the OnDemand data,” on page 97). After configuring the system, we recommend that you create a full backup image of the OnDemand database.

Installing and configuring optional software:

1. If you plan to reprint AFP documents or fax any documents using the OnDemand server print function, then you must install PSF/400 on the OnDemand iSeries server. Complete the following tasks:
 - ___ a. Install and configure PSF/400 (part of OS/400).
 - ___ b. Define a server printer on the OnDemand server with the administrative client.
2. If you plan to configure the messages that are sent to the system log or define your own system log user exit program, see Appendix E, “System log and the system log user exit,” on page 135.

Chapter 11. Installing OnDemand server software

You must install a copy of the OnDemand server software on your iSeries server.

Installing OnDemand

This chapter gives you installation instructions for all of the OnDemand Common Server features.

OnDemand Installation Procedures

To install OnDemand, follow the instructions in the book named *Software Installation* (SC41-5120).

The licensed program number for OnDemand Version 5 Release 3 is **5722RD1**. The individual OnDemand features (which install separately) are as follows:

- ***BASE** for the Base Support (required for all others)
- **10** for Common Server
- **11** for Web Enablement Kit (ODWEK)
- **12** for PDF Indexer

Options 1 through 5 are also available features of 5722-RD1, but do not apply to this publication. See *IBM Content Manager OnDemand for iSeries: Installation Guide* (SC41-5325) for installation details for options 1 through 5.

The recommended way to install OnDemand is to use the Install licensed programs menu option from the Work with Licensed Programs menu (go licpgm). From the Install licensed programs screen, enter a 1 to Add an option, and then enter each of the features listed above that you have ordered (such as *BASE and then options 10, 11, or 12 as needed).

If you install OnDemand with any other method, errors can occur when you attempt to use it.

IBM recommends that you order, load, and apply all PTFs available for OnDemand after successful installation of the licensed program. Refer to Informational APAR II13680 for a complete list of OnDemand Version 5 Release 3 PTFs. The informational APAR can be ordered electronically using the SNDPTFORD command, specifying II13680 for the PTF number. Be sure to read the PTF cover letters and follow any special instructions.

Installing the OnDemand end-user client

For information on installing the OnDemand client for Windows, see *IBM Content Manager OnDemand: User's Guide* (SC27-0836).

Note: The internet FTP web site for the OnDemand client software is <ftp://service.software.ibm.com/software/ondemand/fixes/>.

Installing the OnDemand administrative client

For information on installing the OnDemand administrative client for Windows, see *IBM Content Manager OnDemand for iSeries Common Server: Administration Guide* (SC27-1161).

Important: If the CD-ROM that you receive contains Client software at a level below 7.1.0.4, you must update the Administrator Client software from the Internet to the 7.1.0.4 level or higher. To download the OnDemand Clients, go to the internet FTP web site at <ftp://service.software.ibm.com/software/ondemand/fixes/>.

Part 5. Configuring and initializing the system

This part of the book describes how to configure the OnDemand server to operate and support your organization and create and initialize the database. This part contains the following sections:

- Chapter 12, "Creating an instance," on page 69
- Chapter 13, "Defining a locale," on page 75
- Chapter 14, "Configuring the ARS.CFG file," on page 79
- Chapter 15, "Configuring system start up," on page 81
- Chapter 16, "Scheduling maintenance programs," on page 83

Chapter 12. Creating an instance

About instances

An OnDemand instance is a logical server environment with its own library containing a unique set of database files. An instance is defined in the ARS.INI file by naming the instance (which identifies the name of the library used by the instance). All of the database files that belong to an instance run in one and only one CCSID.

You can run multiple instances on the same server, with each instance configured differently:

- To have separate test and production environments
- To have databases using different CCSIDs

When you work with more than one instance, you must identify the instance name when you run OnDemand commands (such as ADDRPTOND and STRMONOND).

Each instance has different security from other instances on the same machine. You must define users and groups to each instance and set application group and folder permissions for users of each instance. Each instance has its own system log.

Each additional instance requires additional system resources, such as virtual storage and disk space, and more administration.

Configuring an instance

When you install OnDemand, the configuration files for a default instance called QUSROND are created (if they do not already exist). The QUSROND instance itself is not created during the licensed program installation. The language for the default instance is determined from the OnDemand language version you install. If that language version is not one of the supported server languages, then the server language defaults to English. If you are installing multiple language versions for OnDemand option 10 (the OnDemand Common Server feature), you should verify that the server language specified in the ARS.CFG file for the QUSROND instance is what you want. If it is not, modify the ars.cfg file to set the correct server language before you create the QUSROND instance.

Perform the following steps to validate the QUSROND configuration settings (or to create configuration settings for additional instances if required) and create the instance(s):

1. When you are creating an instance, your user profile must have its locale set to the locale of the instance you wish to create. Because the locale is set in the user profile, you may need to change your user profile, then sign off and back on before creating the instance. Use the Change User Profile (CHGUSRPRF) to change (if necessary) your user profile. You should also make sure that other language-related parameters in your user profile are set correctly. (The Change User Profile (CHGUSRPRF) command will not show the current locale setting; it shows *SAME. Use Display User Profile (DSPUSRPRF) to check the locale setting.) The Locale Job Attributes (SETJOBATR) parameter in your user profile is used to determine which values are obtained from the locale. For

OnDemand, at a minimum, you must use SETJOBATR(*CCSID). For example, if you are located in the USA and are using the English language, you would enter the command:

```
CHGUSRPRF USRPRF(user_profile_name) LANGID(ENU) CNTRYID(US) CCSID(37)
  SETJOBATR(*CCSID *DATFMT *TIMSEP *DATSEP *DECfmt *SRTSEQ)
  LOCALE('/QSYS.LIB/EN_US.LOCALE')
```

If you are located in Spain and are using the Spanish language with Euro support, you would enter the command:

```
CHGUSRPRF USRPRF(user_profile_name) LANGID(ESP) CNTRYID(ES) CCSID(1145)
  SETJOBATR(*CCSID *DATFMT *TIMSEP *DATSEP *DECfmt *SRTSEQ)
  LOCALE('/QSYS.LIB/ES_ES_E.LOCALE')
```

See Table 9 on page 129 for a list of locales.

2. Choose a name for the instance, or use the default instance name of QUSROND. The instance name must be a valid library name for OS/400. Be sure that no library, user profile, or authorization list by that name already exists. Further, the instance name must not start with the letter Q (except for QUSROND, which is the default instance), and must not be named CONFIG or WWW. This instance name will be referred to as [instance] in the rest of these instructions.
3. Create the instance by calling the program QRDARS/QRLMINST and passing three parameters. The first parameter is the name of the instance, the second is the three-character language identifier for the instance, and the third is the locale for the instance. The three-character language identifier should match one of the language identifiers listed in Table 8 on page 129. A list of locales can be found in Table 9 on page 129. For example:

```
CALL PGM(QRDARS/QRLMINST) PARM(ONDTEST ENU '/QSYS.LIB/EN_US.LOCALE')
```

would create an instance called ONDTEST with a server language of US English.

The QRLMINST program does the following:

- a. Creates the /CONFIG directory under /QIBM/UserData/OnDemand and the default and model files within this directory (if they do not already exist).
- b. Appends the model ARS.INI file (in directory /QIBM/ProdData/OnDemand/config) to the actual ARS.INI file (in directory /QIBM/UserData/OnDemand/CONFIG) and uses the name of the instance wherever it finds [instance] in the model file.
- c. Creates the instance directory (a directory with the same name as the instance) in directory /QIBM/UserData/OnDemand.
- d. Creates the ARS.CFG, ARS.CACHE and ARS.DBFS files in directory /QIBM/UserData/OnDemand/[instance] and uses the name of the instance wherever it finds [instance] and the language identifier wherever it finds [language] in the model file. (The model files for these three new files are located in directory /QIBM/ProdData/OnDemand/config.)
- e. Creates the library and database tables for the instance.
- f. Creates the directories needed for the instance as specified in the ARS.CFG and ARS.CACHE files.
- g. Creates a user profile with the same name as the instance.
- h. Creates an authorization list with the same name as the instance.

- |
- | 4. You may need to change some of the ARS.INI configuration parameters from
- | the default values before you use this instance the first time. You need to edit
- | the ARS.INI file located in the /QIBM/UserData/OnDemand/config directory.
- | The values listed below may be changed in the definition for this instance. The
- | instance definition starts with the line [@SRV@_ [instance]] where [instance] is
- | the name of the instance. For example, the instance ONDTEST would start with
- | the line [@SRV@_ONDTEST].

| The following lines need to be reviewed:

| **PORT=0.** This specifies the port to which the server listens to receive

| requests from an OnDemand client. The value of 0 means to use the default

| port of 1445. Only one server can be listening to a particular port at any

| given time. If you want to be able to run multiple instances concurrently,

| you will need to specify a port that is not being used on your system for

| some other purpose. If you are running OnDemand Spool File Archive as

| well as Common Server, you are most likely using port 1445 for Spool File

| Archive, so you must make sure you change PORT=0 to some other value

| so that it does not default to 1445. You can use the Work with TCP/IP

| Network Sts (WRKTCPSTS) OPTION(*CNN) command to see what ports

| are currently in use on your system.

| **SRVR_FLAGS_SECURITY_EXIT=1.** This specifies that you want to use

| OS/400 userids and passwords as the OnDemand userids and passwords.

| This is the default value and makes it simpler for your end users because

| they do not have to maintain multiple passwords. If your OnDemand users

| do not need to have OS/400 userids, then you should specify a value of 0

| for this parameter. When you do this, your OnDemand passwords have no

| relationship to OS/400 passwords. However, if there is a match between an

| OnDemand userid and an OS/400 user profile, some OnDemand commands

| and APIs will use the OS/400 user profile as the OnDemand userid, even if

| you have chosen not to relate the two. This situation could permit OS/400

| users to perform OnDemand functions that you did not intend for them to

| perform. Therefore, an OnDemand userid should not match an OS/400 user

| profile name unless the two ids are used by the same individual. If you

| make a change to the SRVR_FLAGS_SECURITY_EXIT value, you should

| review the OnDemand System Parameters values (defined using the

| OnDemand administrative client) for the instance that you have changed.

| See "OnDemand userid relationship to OS/400 user profiles" on page 40 for

| details.

- |
- | 5. You may need to change some of the ARS.CFG configuration parameters from
- | the default values before you use this instance the first time. You need to edit
- | the ARS.CFG file located in the /QIBM/UserData/OnDemand/instancename
- | directory (where instancename is the name of the instance you wish to review.
- | The values listed below may be changed in the definition for this instance.

| **ARS_LANGUAGE=ENU.** This specifies the language in which this instance

| runs. This example shows the use of the English language. Valid languages

| are listed in "Locales" on page 129.

| **ARS_MSGS_LANGUAGE=ENU.** This specifies the language that is used

| for server messages. This example shows the use of the English language.

| Valid languages are listed in "Locales" on page 129.

| **ARS_AUTOSTART_INSTANCE=1.** This specifies whether or not to

| automatically start the server for this instance when using the Start TCP/IP

| Server (STRTCPSVR) command. Set this value to 1 to automatically start

| this instance's server; set this value to 0 if you do not want to automatically

| start this instance's server.

6. Do not modify any of the other values in these instance definition files without first consulting with OnDemand Support.

Deleting an instance

To completely delete an instance, you need to do the following:

1. If the server for the instance is already started, you must end the server before deleting the instance. You can end a specific server by following the instructions in “Starting and stopping servers.”
2. Delete the instance library. You can do this by issuing the following command: `DLTLIB LIB([instance])` When you issue this command, you may receive a message indicating that a journal receiver has not been fully saved. You can respond with an I to ignore this message.
3. Delete the instance directory which is under `/QIBM/UserData/OnDemand`. If you use the `WRKLNK` command each directory must be empty before it can be deleted. If files exist in a subdirectory you will have to go into each subdirectory under the instance to delete the files and directories in it before removing the directory.

Note: You can use the `WRKLNK` command to do a recursive delete of all the objects in a directory. You can do this by issuing the following command: `WRKLNK '/QIBM/UserData/OnDemand/[instance]'` where `[instance]` is your instance name. Specify option 2 to edit the instance directory you want to delete. Then specify option 9 for every directory listed in that instance directory. Finally press F12 to go back to the previous screen and then delete the instance directory itself.

4. Delete the instance definition within file `ars.ini` in directory `/QIBM/UserData/OnDemand/config`.
5. Delete the instance authorization list. You can do this by issuing the following command: `DLTAUTL AUTL[instance]` where `[instance]` is the name of the instance you are deleting.
6. Delete definitions related to storage management that you created using the iSeries Navigator administrative functions, if they are no longer needed.
7. You can consider deleting the instance user profile (which has a name that matches the instance name), but be careful to verify that the user profile does not own objects that will still be required after the instance is deleted.

Starting and stopping servers

You must start a server for an instance before clients can connect to the instance.

To start all instances that have `ARS_AUTOSTART_INSTANCE=1` set in the `ars.cfg` configuration file:

1. Start the instances:
`STRTCPSVR SERVER(*ONDMD)`
2. Use the `WRKACTJOB` command to verify that an `<instance>` started:
`WRKACTJOB JOB(<instance>)`

where `<instance>` is the name of the instance.

To stop all instances:

```
ENDTCPSVR SERVER(*ONDMD)
```

To start a specific instance:

```
CALL PGM(QRDARS/QRLMCTL) PARM(*STRTCPSVRQUSROND)
```

where QUSROND is the name of the instance you wish to start. Note that there is no space between *STRTCPSVR and the name of the instance.

To stop a specific instance:

```
CALL PGM(QRDARS/QRLMCTL) PARM(*ENDTCPSVRQUSROND)
```

where QUSROND is the name of the instance you wish to stop. Note that there is no space between *ENDTCPSVR and the name of the instance.

Server jobs are started using a job description by the name of the instance (which must be found in the QUSRRDARS library). If a job description by that name is not found in QUSRRDARS, then job description QOND400 in library QRDARS is used (and can be changed if necessary).

The job description controls the following attributes of the server job:

- JOBQ
- JOBPTY
- OUTPTY
- PRTDEV
- OUTQ
- INLLIBL
- LOG
- LOGCLPGM
- INQMSGRPY
- HOLD
- DATE
- SWS
- JOBMSGQMX
- JOBMSGQFL

| For example, if you wanted to change the job queue that instance TEST used, you
| would create a job description called TEST in library QUSRRDARS that specified
| the job queue you wanted to use. This would enable you to use the TEST job
| description to send your server jobs to a different subsystem than the default.

Automatically starting instances

To enable an instance to start automatically each time the system restarts, you must add one of the commands described in “Starting and stopping servers” on page 72 to your QSTRUP program. You can also add the commands to a job scheduler.

Connecting to instances

Connecting from OnDemand clients

To connect to a particular instance, the client must log on to the server using the correct port. Use the Update Servers dialog box to add servers to the client. When you add a server, you identify the name of the server and the port number and communications protocol used to communicate with the server. The port number

that you specify in the Update Servers dialog box must be the same as the value that you specified for the instance in the ARS.INI file.

Working with instances

OnDemand provides commands that you can use to work with an instance. For example, OnDemand provides commands to add, remove, and print spooled files and other data types and to manage storage. On many of the OnDemand commands, one of the parameters that you can specify to a command is the name of the instance. By default, if you do not specify otherwise, the name of the instance is QUSROND. If you name your only instance something other than QUSROND or you work with more than one instance, then you must always specify the name of the instance when you run a command. That way you can make sure you are working with the correct instance.

See the *IBM Content Manager OnDemand for iSeries Common Server: Administration Guide* appendix entitled “Command reference” for more information about OnDemand commands.

OnDemand also provides APIs that you can use to work with an instance.

The ARSLOAD API uses the `-I` instance parameter. OnDemand retrieves the name of the server that is associated with the instance from the ARS.INI file. For example (using QSHELL):

```
arsload -I TEST -d /test/file1 -c /test/file2
```

where TEST is the name of the instance. In the example, the PORT parameter in the `[@SRV@_TEST]` section of the ARS.INI file identifies the server.

The ARSDOC API provides the `-h` flag to specify the name of the instance.

The following example shows how to run the programs in QSHELL by specifying the name of the instance:

```
/arsdoc get -h TEST -F parmfile
```

where TEST is the name of the instance. In the example, OnDemand uses the PORT parameter in the `[@SRV@_TEST]` section of the ARS.INI file to identify the server.

See the *IBM Content Manager OnDemand for iSeries Common Server: Administration Guide* appendix entitled “API and user exit reference” for more information about OnDemand APIs.

Chapter 13. Defining a locale

OnDemand requires the specification of a locale for each Common Server instance, in order to ensure that your data has the correct language identifier associated with it during the archive process. The steps outlined below must be performed for each OnDemand instance before the server for that instance can be started. Even if you are running a U.S. English system, you must still perform these steps. Some of these steps have also been discussed in Chapter 12, "Creating an instance," on page 69 but are included again here for completeness.

IBM ships a number of locale objects (object type *LOCALE) in library QSYS. Others are provided as source code that must be created using the CRTLOCALE command. (Locale names followed by an asterisk (*) in the table of locales values in "Locales" on page 129) must be created using the CRTLOCALE command.) If the locale you require must be created from source code, you must install option 21 of OS/400 (Extended NLS Support; a no-charge product option of 5722-SS1), which contains all the IBM-supplied locale source files. If you need a locale that is not supplied by IBM (no source or object), you will have to copy and modify an existing locale to meet your needs. Assistance in doing this can be obtained from the IBM support organization in your geography.

Background information on instances

When an OnDemand Common Server instance is created, the following objects are created on the iSeries system:

- A user profile, named the same as the instance
- An IFS directory, located at /QIBM/UserData/OnDemand/instance_name (where instance_name is the name of the instance)
- An authorization list, named the same as the instance
- A library, named the same as the instance
- A section for the instance added to the ARS.INI file located at /QIBM/UserData/OnDemand/CONFIG/ARS.INI

There are two different sets of instructions included here: The first set describes the steps to take for instances that already exist on your iSeries system (if you have not already done so at V5R1 or V5R2). The second set describes the steps to take for new instances being added to your current OnDemand system OR for new installations.

Steps to take for instances that already exist

These instructions describe the steps required to successfully modify the user profiles for an existing OnDemand Common Server instance in order for the correct locale to be used for the instance. **The user profiles for all your instances must be modified.**

1. The user profile named the same as your instance AND all user profiles that perform OnDemand data loading functions (such as ADDRPTOND, STRMONOND, arslod or arsdod add) must have their locale set to the locale of the instance. **WARNING:** If the user profiles that perform OnDemand data loading functions are also used for other non-OnDemand work on your system, changing their LOCALE value may impact other non-OnDemand applications.

Consider creating specific OnDemand user profiles for data loading instead, so as not to impact your other applications, and then only modify the LOCALE value for the specific OnDemand user profiles. If necessary, change the locale of the user profiles using the Change User Profile (CHGUSRPRF) command. You should also make sure that other language related parameters in the user profiles are correctly set. The SETJOBATR parameter is used to determine which values are obtained from the locale. For OnDemand, at a minimum, you must use SETJOBATR(*CCSID). For example, if you are located in Spain and are using the Spanish language with Euro support you would enter the command:

```
CHGUSRPRF USRPRF(user_profile_name) LANGID(ESP) CNTRYID(ES)
          CCSID(1145) SETJOBATR(*CCSID *DATFMT *TIMSEP *DATSEP *DECFTM
          *SRTSEQ) LOCALE('/QSYS.LIB/ES_ES_E.LOCALE')
```

Note that copying one user profile to create another similar user profile does not copy the LOCALE parameter value. After the copy is complete, you must issue the CHGUSRPRF command as shown above to set the LOCALE. You can use the Display User Profile (DSPUSRPRF) command to display or verify the locale of an existing user profile.

2. After you have completed step 1 for all your OnDemand instances, you can start the OnDemand servers for all instances that are set to “autostart,” using the command:

```
STRTCPSVR *ONDM
```

To start only the server for a particular instance, use the command (noting no space between *STRTCPSVR and instance_name):

```
CALL QRDARS/QLMCTL *STRTCPSVRinstance_name
```

3. You can now log on to the instance using the OnDemand administrative client and the OnDemand end-user client.
4. When loading data into OnDemand, you must ensure that the job loading the data is running with the correct locale and CCSID. Follow the instruction in step 1 to correct the user profile used to load data before starting any load jobs (if you haven’t already made the necessary changes). Load jobs are those that use any of the following commands: ADDRPTOND, STRMONOND, arslod, or arsdod add.

If you do not have a locale set in the user profile used to load data, the load job will fail. The joblog for the qshell job, QP0ZSPWP, will contain the error message YOU MUST SET A LOCALE.

Steps to take for new instances or new installations

The instructions outlined in Chapter 12, “Creating an instance,” on page 69 describe the steps required to successfully create, start, and access a new instance of the OnDemand Common Server.

If you encounter problems while trying to create an OnDemand Common Server instance, see the list of problems, causes, and recoveries below. If you need additional assistance contact the software support organization in your geography.

Solving problems

The following list includes some of the locale-related problems you might encounter when creating an instance or archiving data. The cause of each possible problem is described, and the steps to recover from the problem are explained.

| **Problem:** You receive the message: CPFA0A9 Object not found. LOCALE
| /QSYS.LIB/EN_GB_E.LOCALE NOT FOUND

| **Cause:** You specified a locale that does not exist on the system.

| **Recovery:** Either specify a locale that does exist on the system, or create the locale
| that you need using the CRTLOCALE command.

| **Problem:** You receive the message: OND1005 OnDemand instance MYINSTANCE failed
| to instantiate. (V5R1) or OND1005 OnDemand instance MYINSTANCE was not
| created. (V5R2). The joblog also contains RDR2798 YOU MUST SET A LOCALE.

| **Cause:** You did not specify a locale in the user profile creating the instance, or you
| did not signoff and signon again after changing the user profile to specify a locale.

| **Recovery:** Delete the IFS directories created for the instance. Delete the instance
| user profile. Delete the instance authorization list. Return to step 2 of these
| instructions.

| **Problem:** Your data will not load into OnDemand. There is no system log message
| indicating the cause of failure. The joblog for the load job does not contain a
| specific error message.

| **Cause:** You do not have a locale specified in the user profile running the load job.

| **Recovery:** Check the joblog for the qshell job, QP0ZSPWP, for the error message
| YOU MUST SET A LOCALE. If the joblog contains this message, you must change the
| user profile running the load job to specify a locale. See these instructions for
| details.

| **Problem:** Your spooled file fails to archive. The system log contains message 88,
| including text similar to the following Row 1: The string " 0/12/11" could not
| be converted to a date from the format of %m/%d/%y. You check the Load
| Information tab in the application definition and see that the date format is
| specified as %y/%m/%d and that leading spaces are to be removed.

| **Probable Cause:** The user profile archiving the data has a locale specified, but that
| locale is not compatible with the instance into which the data is being loaded. For
| example, specifying a Czech locale but trying to archive US English data will cause
| this type of failure. Note that you may, in fact, have an actual date definition error,
| but it is likely related to locales if you have already verified the date information
| as being correct.

| **Recovery:** Specify a locale in the user profile archiving the data that compatible
| with the instance into which the data is being archived.

Chapter 14. Configuring the ARS.CFG file

When you run the various OnDemand programs, they obtain information from the ARS.CFG file about a variety of configuration settings. The information includes the language setting for the database and the location of configuration and temporary work space directories, for example. Before you begin using OnDemand, you should review the parameters in the ARS.CFG file. Most of the values that IBM provides are sufficient for many customers. However, you may need to change some of the values for your environment. If you make changes to the ARS.CFG file while the system is running, you must stop and restart the server before OnDemand can use any values that you modify.

To change the ARS.CFG file, log on as QSECOFR. Issue the WRKLNK command and change to the /QIBM/UserData/OnDemand/<instance> directory, where <instance> is the name of your OnDemand instance. The default instance name is QUSROND. Make a backup copy of the file provided by IBM. Edit the ARS.CFG file by using Option 2 (Edit) on the WRKLNK screen.

ARS_AUTOSTART_INSTANCE parameter

Defaults to 1 (which means YES, autostart this instance's server) for instance QUSROND. Defaults to 0 (which means NO, do not autostart this instance's server) for all other instances. When you create a new instance (other than QUSROND), the server jobs for that instance will not automatically start by issuing STRTCPSVR *ONDMD unless you change the value of this parameter from 0 to 1.

If you choose not to autostart an instance's server, you can start the server manually by doing the following:

```
CALL QRDARS/QRLMCTL *STRTCPSVRinstance
```

where instance is the name of the instance. There should not be any spaces between the *STRTCPSVR and the instance name.

You can end any instance's server manually by doing the following:

```
CALL QRDARS/QRLMCTL *ENDTCPSVRinstance
```

where instance is the name of the instance you want to stop. Again, there should not be any spaces between the *STRTCPSVR and the instance name.

Note: ENDTCPSVR *ONDMD ends all running OnDemand servers regardless of how they were started.

ARS_LANGUAGE parameter

The national language environment used to create the database. OnDemand uses a three-character *language code* to derive the locale and code set for the server operating system. See Appendix D, "National language support," on page 127 for a list of the language codes and information about configuring the system for national language character support.

Important: You must provide the correct language code for your operating environment **before** you create the OnDemand instance database.

ARS_MSGS_LANGUAGE parameter

If you want something other than English for server messages, add this parameter to the ars.cfg file for the instance and specify the three-character language id that you want to use. To see a list of valid language id values, see the directory names within /QIBM/ProdData/OnDemand/locale on your iSeries server after you install OnDemand.

ARS_MESSAGE_OF_THE_DAY parameter

Use to specify the message of the day for the OnDemand client. Set to the full path name of a file that contains the message that you want the client to show. For example:

```
ARS_MESSAGE_OF_THE_DAY=/QIBM/ProdData/OnDemand/tmp/message.txt
```

The contents of the message file can contain a maximum of 1024 characters of text. The administrative client and the end-user client show the message after the user logs on to the server. To close the message box and continue, the user must click OK. If you do not specify a message file, then the normal client processing occurs.

ARS_NUM_DBSRVR parameter

Determines the number of processes that OnDemand starts on the server to handle connections to the database. The default value is 5 (five). See Appendix C, "Database servers," on page 125 for more information about this parameter.

ARS_TMP parameter

The location where OnDemand programs temporarily store data. You must allocate sufficient free space in the specified file system to support tasks such as migrating and importing index data. The default value is:

```
/QIBM/UserData/OnDemand/<instance>/tmp (where <instance> is the name of the OnDemand instance).
```

ARS_PRINT_PATH parameter

The location where the OnDemand server print function temporarily stores print data. You must make sure that there is enough space in the specified location to hold the print files for the maximum number of concurrent print requests that the server will handle. The default value is: /QIBM/UserData/OnDemand/<instance>/tmp (where <instance> is the name of the OnDemand instance).

ARSSOCK_RESOLVE_CLIENT_NAME parameter

If you are experiencing longer than normal logon response times in a DHCP environment, adding the following line to the ARS.CFG file may correct the problem:

```
ARSSOCK_RESOLVE_CLIENT_NAME=0
```

Chapter 15. Configuring system start up

This section provides information regarding server functions you may choose to start automatically whenever the system is started.

You can modify the operating system's QSTRUP program to automatically issue server commands each time the server is restarted.

Consider the following for possible addition to the QSTRUP program:

- Start the OnDemand server jobs
- Start the OnDemand output queue monitor(s)

Start the server jobs

The following example shows the command that starts the OnDemand server jobs. For more details and options for starting and stopping servers, see "Starting and stopping servers" on page 72.

```
STRTCPSVR SERVER(*ONDMD)
```

Start the output queue monitor(s)

| Use the Start Monitor for OnDemand (STRMONOND) command to start
| OnDemand output queue monitors for output queues that will receive spooled
| files destined for OnDemand. Review the online help text to determine the specific
| parameters you wish to specify for each monitor.

| You can also create monitor definitions using the OnDemand Archive plug-in of
| iSeries Navigator. After starting iSeries Navigator, select OnDemand Archive, then
| use Monitor Defintions under Common Server Administration. These monitor
| definitions allow you to predefine the monitor startup parameters and specify one
| of a number of ways to automate the starting of the monitor(s).

| See "Loading spooled file data" in *IBM Content Manager OnDemand for iSeries
| Common Server: Administration Guide* for more information on output queue
| monitors.

| You can later end the monitor if desired, using the End Monitor for OnDemand
| (ENDMONOND) command or setting conditions for ending when you issue the
| STRMONOND command or set up the monitor definition using iSeries Navigator
| as described above.

Chapter 16. Scheduling maintenance programs

This section describes the OnDemand maintenance commands to run automatically on a regular schedule. The maintenance programs that you may want to run on a regular schedule include those that maintain application group data in the database and on disk, and maintain OnDemand system data in the database.

Note: Most maintenance programs should run when no other programs are updating the database or need exclusive access to the database. For example, you should not load data (ADDRPTOND, STRMONOND, arslod, arsdod add) at the same time that you run the maintenance commands (STRDSMOND, STRASMOND).

The time of day and frequency with which you run the maintenance commands and the processing options that you specify will vary, based on your environment. For example, if you do not load data every day or you add few rows to the database when you do load data, then you should be able to schedule maintenance less often, perhaps once a week. However, if you load a high volume of data every day, then you should maintain the database after the load processing completes. If your system has plenty of disk space available to store index data and reports, then you may be able to schedule the maintenance programs less often.

The examples that follow illustrate one method of daily maintenance. The examples use the typical options to maintain the database and disk storage and take full backup images of the databases. The maintenance programs are scheduled to run early in the morning, before users begin accessing the system. Such a schedule would assume that all data loading completed before the maintenance programs begin. You should use the examples as a guideline.

Maintaining application group data

Note: The OnDemand server must be running, otherwise the maintenance commands will fail.

You can use the STRDSMOND and STRASMOND commands to maintain application group data in the database and on disk and archive storage. These maintenance commands determine the actions to take from parameters that you specify and by retrieving storage management information from application groups. The commands can do the following for your application groups:

- Copies report data from disk storage to archive storage
- Expires report data from disk storage and archive storage
- Migrates index data from the database to archive storage (if required, although not recommended)
- Expires index data from the database

You typically schedule the maintenance commands to do these things on a regular schedule. If your organization loads a high volume of data into the system every day, then we recommend that you schedule the commands to run after all load processing completes. For low to average volumes, you may choose to run the

maintenance commands weekly. You should schedule the commands to run during a period of little or no other activity on the system.

See the *IBM Content Manager OnDemand for iSeries Common Server: Administration Guide* for more information about the STRDSMOND and STRASMOND commands.

Backing up the OnDemand database

After you finish loading data into OnDemand and maintaining the database, we recommend that you create a backup image of the OnDemand database. The backup image can be used to recover the database, if a hardware failure occurs or application data becomes corrupted. If you process several loads each day, then you may want to backup the database once a day, after the last load process completes.

| See Chapter 21, "Backing up the OnDemand data," on page 97 and *IBM Content Manager OnDemand for iSeries Common Server: Administration Guide* for more
| information about backing up your OnDemand data.
|

Part 6. Preparing the system for use

This section of the book describes how to verify the installation of OnDemand. Other important tasks described in this section include:

- Defining migration policies and storage sets. You must add migration policies and storage sets to the system before you can create application groups or assign the system-defined application groups to a storage set. Depending on the storage management characteristics of the reports that you plan to store on the system, you may need to add more than one migration policy and storage set.
- Configuring the system log application group. IBM recommends that you assign the System Log application group to a storage set that specifies archive media (such as optical or tape rather than disk) so that the system can maintain a permanent copy of the system log data.
- Configuring the system migration application group. If you plan to migrate index data from the database to archive storage, then you must create a storage set that specifies archive media. After you add the storage set to the system, you can assign the System Migration application group to the storage set. (Migrating index data from the database to archive storage is not recommended.)
- Creating a backup copy of the database. After installing and configuring OnDemand, IBM recommends that you create a backup copy of the OnDemand database.

Chapter 17. Verifying the installation

1. Complete all of the required steps in Chapter 11, "Installing OnDemand server software," on page 65.
2. Complete all of the required steps in Part 5, "Configuring and initializing the system," on page 67.
3. If you have not already done so, install the OnDemand client program on a PC. See the *IBM Content Manager OnDemand: User's Guide* for details.
4. Start the OnDemand client program. OnDemand displays the Logon to Server dialog box.
5. Click Update Servers. OnDemand displays the Update Servers dialog box.
6. Enter the name of the OnDemand server. Click Help for information about the fields and options.
7. Click Add then Close to return to the Logon to Server dialog box.
8. Select the name of the server that you added in the Update Servers dialog box, if it is not already selected.
9. Type an OnDemand userid and password in the fields provided. (The first time that you log on to the system, you must specify the built-in OnDemand userid, QONDADM. Initially, the password is QONDADM1. However, you will be prompted to enter and verify a new password.)
10. Press Enter.
11. Open and search the System Log folder.

If you were able to view messages stored in the system log, then you can consider the installation of OnDemand successful.

If the client program does not start, check the drive, path name, and program name values used to start the program. Then try the command again.

If the client program issues a message indicating a problem, follow the instructions in the message window. If the problem persists, contact the IBM support center for help with resolving the problem.

Chapter 18. Defining migration policies and storage sets

You must define migration policies and storage sets before you can define reports to OnDemand or load data into the system.

You can set up OnDemand to copy data to disk storage or archive storage (or both). The storage management attributes of the application groups that you add to the system will determine the types of media that you need and how you configure storage sets on the system.

A storage set must contain at least one primary storage location. A primary storage location can use disk storage (the default) or specify archive storage such as optical or tape (or both). The administrative client online help provides details about defining storage characteristics. Also see the *IBM Content Manager OnDemand for iSeries Common Server: Administration Guide* for more information about migration policies and storage sets.

Although not recommended, if you plan to migrate index data to archive storage, then you must assign the System Migration application group to a storage set that specifies archive storage. Also, we recommend that you assign the System Log application group to a storage set that specifies archive storage so that the system can maintain a permanent copy of the data that is written to the system log. The following topics provide additional details:

- Chapter 19, “Configuring the System Log application group,” on page 91
- Chapter 20, “Configuring the System Migration application group,” on page 95

Chapter 19. Configuring the System Log application group

When you install and configure OnDemand, you initialize the system log. The system log is comprised of the System Log application group, a set of System Log applications, and the System Log folder. The System Log application group contains the storage management information that OnDemand uses to maintain the data written to the system log. When you initialize the system, the application group is not assigned to a migration policy and storage set. Because the application group is not assigned to a storage set, the system does not maintain a permanent copy of the system log data.

Before you begin defining reports to OnDemand, loading data on the system, or allowing users to access the system, we recommend that you configure the System Log application group to maintain a permanent copy of the data that is written to the system log. You can do this by first defining a migration policy (and its corresponding storage set) that specifies archive storage media such as optical or tape and then by updating the System Log application group and assigning it to the storage set. See Chapter 18, “Defining migration policies and storage sets,” on page 89 and “Maintaining system log data in archive storage” for more information.

If your system will not use archive storage media, then you should assign the System Log application group to a disk-only migration policy and storage set and change the length of time that OnDemand maintains the system log data to the maximum permitted value. See “Maintaining system log data on disk” on page 92 for more information.

No matter where you decide to keep your System Log data, it is important to note that the System Log is the only place where the Load ID for archived data is maintained. Every time data is loaded into OnDemand, message number 87 is placed in the System Log and the Load ID is recorded as part of that message. The Load ID (also called Report ID) is a required parameter for the Remove Report (RMVRPTOND) and Print Report (PRTRPTOND) commands. Without this Load ID information, you will be unable to use these two commands.

Maintaining system log data in archive storage

We recommend that you create a migration policy that specifies archive storage media. You must add at least one primary storage location to the migration policy. The primary storage location must identify archive storage media to maintain data indefinitely.

After you create the migration policy (which automatically creates a storage set with the same name), you must update the System Log application group and assign it to the storage set. After assigning the application group to the storage set and restarting the server, the system automatically maintains a copy of the system log data in archive storage.

Complete the following steps to assign the System Log application group to a storage set:

1. Start iSeries Navigator. Select the OnDemand server, and sign on if prompted.

2. Start the administrative client by right-clicking on Common Server Administration under the OnDemand Archive section of iSeries Navigator. Then select Client Administrative Functions. (See *IBM Content Manager OnDemand for iSeries Common Server: Administration Guide* for information on installing the administrative client functions if these items do not appear under iSeries Navigator on your workstation.)
3. Log on to the server with a userid that has system administrator authority. (The built-in userid QONDADM has system administrator authority.)
4. Click Application Groups.
5. Point to the System Log application group and click the right mouse button.
6. From the pop-up menu, select Update to open the Update an Application Group window.
7. Click the Storage Management tab.
8. In the Storage Set Name list, select the name of the storage set. The storage set name that you select should match a migration policy definition that maintains data indefinitely.
9. Click Advanced to open the Advanced Storage Management dialog box.
10. Select Next Cache Migration under Migrate Data from Cache. This causes OnDemand to copy the system log data to archive storage the next time that the maintenance commands run.
11. Click OK to close the Advanced Storage Management dialog box.
12. Click OK to save your changes and close the Update an Application Group window.

The administrative client online help provides information about the options on the Storage Management page.

Maintaining system log data on disk

If your system will not use archive storage, then you should assign the System Log application group to a disk-only storage set and change the length of time that OnDemand maintains the data to the maximum permitted value. Doing so ensures that OnDemand does not delete the data from disk for a very long time. Complete the following steps to configure the System Log application group:

1. Start iSeries Navigator. Select the OnDemand server, and sign on if prompted.
2. Start the administrative client by right-clicking on Common Server Administration under the OnDemand Archive section of iSeries Navigator. Then select Client Administrative Functions. (See *IBM Content Manager OnDemand for iSeries Common Server: Administration Guide* for information on installing the administrative client functions if these items do not appear under iSeries Navigator on your workstation.)
3. Log on to the server with a userid that has system administrator authority. (The built-in userid QONDADM has system administrator authority.)
4. Click Application Groups.
5. Point to the System Log application group and click the right mouse button.
6. From the pop-up menu, select Update to open the Update an Application Group window.
7. Click the Storage Management tab.
8. In the Storage Set Name list, select the name of the storage set. The storage set named Cache Only - Library Server is a disk-only storage set created on the server when you initialized the system.

9. Replace the contents of the Cache Data for ___ Days field with 99999. This value causes OnDemand to maintain data for approximately 273 years.
10. Click OK.

The administrative client online help provides information about the options on the Storage Management page.

Chapter 20. Configuring the System Migration application group

Note: Index migration is not recommended.

Index migration is the process by which OnDemand moves index data from the database to archive storage. This process optimizes database storage space while allowing you to maintain index data for a very long time. You typically migrate index data after users no longer need to access the reports, but for legal or other requirements, you still need to maintain the data for some number of years or months. OnDemand uses the storage management settings in application groups to determine whether or not to migrate index data to archive storage. All migrated data is managed through the System Migration application group.

When you install and configure OnDemand, you initialize the system migration function. The system migration function is comprised of the System Migration application group, a set of System Migration applications, and the System Migration folder. The System Migration application group contains the storage management information that OnDemand uses to maintain index data migrated to archive storage. Until you assign the application group to a storage set that specifies archive media, OnDemand cannot migrate index data from the database to archive storage. See Chapter 18, “Defining migration policies and storage sets,” on page 89 and “Assigning the System Migration application group to a storage set” for more information.

Assigning the System Migration application group to a storage set

If you need the system to maintain index data in archive storage, then you must assign the System Migration application group to a storage set that maintains data indefinitely.

After assigning the application group to the storage set and restarting the server, the system automatically migrates index data to archive storage, whenever the maintenance commands run. Complete the following steps to assign the System Migration application group to a storage set:

1. Start iSeries Navigator. Select the OnDemand server, and sign on if prompted.
2. Start the administrative client by right-clicking on Common Server Administration under the OnDemand Archive section of iSeries Navigator. Then select Client Administrative Functions. (See *IBM Content Manager OnDemand for iSeries Common Server: Administration Guide* for information on installing the administrative client functions if these items do not appear under iSeries Navigator on your workstation.)
3. Log on to the server with a userid that has administrator authority. (The built-in userid QONDADM has system administrator authority.)
4. Click Application Groups.
5. Point to the System Migration application group and click the right mouse button.
6. From the pop-up menu, select Update to open the Update an Application Group window.
7. Click the Storage Management tab.

8. In the Storage Set Name list, select the name of the storage set. The storage set that you select should identify an archive media that maintains data indefinitely.
9. Click OK.

The administrative client online help provides information about the options on the Storage Management page.

Chapter 21. Backing up the OnDemand data

To complete the installation and configuration process, we recommend that you create a full, offline backup of the OnDemand database on removable media. A full backup of the database is required to rebuild the database, in the event that you need to do so. You cannot rebuild the database unless you have a full database backup (and any journal receivers generated since the last full database backup). See the *IBM Content Manager OnDemand for iSeries Common Server: Administration Guide* for more information about backing up your OnDemand data.

Make a record of the information about the database backup, including the date and time that the backup was taken and the label of the storage volume. Keep the backup copy in a safe location, preferably offsite. Save the backup copy at least until the next time that you create another full backup of the database.

Appendix A. Migration from Spool File Archive to Common Server

Introduction

At Version 5 Release 1, OnDemand introduced a new server implementation known as the OnDemand Common Server. The Common Server provides enhanced indexing, searching, viewing, security, PDF, and web enablement capabilities for OnDemand users and administrators.

Indexing enhancements

- up to 32 index keys
- up to 254 positions for index key length
- default key values
- string, decimal, integer, and date data types

Searching tools

- full text search
- default search operators
- required search fields
- search result volume limits prevent performance degradation from inadequately defined searches

Viewing

- unique logical views for each user, or group of users
- reports viewed from multiple folders
- creation of graphical annotations

Security enhancements

- permit/deny view, print, fax
- allow/prohibit creation of logical views
- permit/deny adding/deleting annotations

Optional software features

- OnDemand Web Enablement Kit (ODWEK) provides browser access
- PDF Indexer can capture and index PDF output from IFS directories (adding to existing support for line data, SCS, SCS-extended, and AFPDS)

Existing OnDemand customers who have implemented Spool File Archive (with or without AnyStore or the existing Server feature) can now migrate to the new Common Server using these instructions. Note that, throughout these instructions, reference to the migration of Spool File Archive data will also imply AnyStore data as well, if AnyStore is installed.

For those customers who have had experience upgrading previous releases of R/DARS or OnDemand, each migration was incremental. Often, new fields were introduced to existing database files, new files were created, and the program functionality was, of course, enhanced. Since Common Server is a complete replacement for Spool File Archive, migrating from Spool File Archive to the new

Common Server will be quite a different experience. The files are all new, the terminology has changed, and the program architecture is client/server for most functions. A number of OS/400 commands are provided to perform operational tasks, but the majority of administrative and end-user functions are performed using the workstation client software. Visit the Content Manager OnDemand for iSeries support page on the web at <http://www.ibm.com/software/data/ondemand/400/support.html> and search for *comparison* to find a document that provides a comparison between Spool File Archive and Common Server functions, commands, and APIs.

There are two main areas of focus during the migration from Spool File Archive to Common Server:

- **Migrating the definitions.** This migration includes report definitions, report group definitions, public and private logical view definitions, OnDemand users and user groups, migration policies, and optical storage groups. Named queries are not migrated. A warning message prints on the Definition Migration report if a named query exists for a report definition that is being migrated. Users can recreate their named queries after the definition has been migrated.
- **Migrating the index data.** This migration allows the data already stored in Spool File Archive to be accessed after migrating to the Common Server, and includes index records, annotations, and AFP resources.

Migration is done through a series of program calls which first migrate the definitions, and then migrate the index information for your current Spool File Archive data. The actual archived data does not need to be migrated. The Common Server has been designed to access data stored in both Spool File Archive and Common Server formats.

Notes:

1. The migration steps described in these instructions do not apply to OnDemand Object Archive or Record Archive. Object Archive and Record Archive will continue to function as they have in previous versions of OnDemand.
2. You **must read these instructions in their entirety** before beginning your migration activities. Migration steps are very interdependent, and a thorough understanding of the entire process is mandatory for a successful migration. You may also want to consider attending a Common Server class to fully understand Common Server concepts in order to take full advantage of this new OnDemand environment.
3. After reading these instructions thoroughly, **you must read the remainder of this Content Manager OnDemand Common Server Planning and Installation Guide.** In addition to installing the prerequisite Common Server iSeries software and related workstation software, you will need to have done some configuration work described in other chapters of this publication. Terminology will also be used here that is introduced earlier in this publication.
4. For planning purposes, you may wish to review various benchmarks that have been established related to migration. Visit the Content Manager OnDemand for iSeries support page on the web at <http://www.ibm.com/software/data/ondemand/400/support.html> and search for *benchmarks* to find a document that provides more information on this topic.
5. There are some characteristics of Spool File Archive report definitions that may cause problems with migration. See “Reports that should NOT be migrated” on page 101 and “Reports that may require attention before they are migrated” on page 102 for a list of areas that need to be reviewed prior to migration.

6. When using the OnDemand end-user client to access data stored in both Spool File Archive and Common Server systems, the user is prompted for a single user ID at logon. The system logs the user onto the Common Server. The user ID and password are also used by Spool File Archive, so if you have disabled the default configuration setting that links OnDemand Common Server user ids to OS/400 user profiles, you must ensure that users' OnDemand user IDs and passwords match their OS/400 user profiles and passwords. See "OnDemand userid relationship to OS/400 user profiles" on page 40 if you are unsure about whether you have disabled this password linkage.
7. The Common Server does not provide a 5250 ("greenscreen") end-user interface. All administrator and end-user interfaces are performed on the workstation. Some OS/400 operational functions are initiated using 5250 commands, but these are not end-user functions.
8. Changes made to a Spool File Archive report definition **after** it has been migrated in Phase 1 to Common Server will not be reflected in the Common Server. If you require these changes to be reflected in Common Server, you must make the required changes in Common Server to match the changes made in Spool File Archive.

What can be migrated?

The following two sections list categories of reports that either should not be migrated at the current time, or may need some changes to prepare them for migration. Because migration can be done in pieces, specifying individual report names, report group names, or generic names, you are able to select or exclude particular reports if desired or necessary. Note that the "report definition analysis program" in Phase 1 (described further in this appendix) provides a report to identify any of your report definitions that fall into either of these categories. In addition, the actual report definition migration program will list any attempts to migrate report definitions in either of these categories.

Reports that should NOT be migrated

- Content Manager (CM; previously known as VisualInfo™ or VI) Integration. There is currently no integration between Common Server and Content Manager/VisualInfo. Integration of new reports will cease once report definition conversion has completed (Phase 2 described below).
 - If you no longer require integration between OnDemand and CM, you can simply migrate the report definition. A warning message will be issued on the report definition migration report to indicate that you have just migrated a CM-integrated report. The warning will also state that integration of new reports will cease after report definition conversion (Phase 2) occurs for this report definition. Until Phase 2 occurs, the Spool File Archive to CM integration will continue, even though you have begun the migration process to Common Server. During the Cleanup Phase (Phase 5) for the report definition, all CM integration records that had previously been written to the CM control files will be deleted. Until Cleanup is run, however, old data that had been integrated into CM from Spool File Archive will continue to be retrievable from the CM client.
 - If you require continued integration between OnDemand and CM, you can choose either not to migrate the report definition at this time; or to consider other options that offer a consolidated hit list across both OnDemand Common Server and CM. These options include:
 - The Content Manager Web Enablement Kit, an optional feature of DB2 Content Manager for iSeries,

- DB2 Information Integrator for Content, or
- Products or services available from IBM Business Partners.
- Unbundled reports (reports with report type of UBND). You must now separate the bundled spooled file into individual spooled files prior to processing in Common Server. There is no Common Server equivalent to Spool File Archive UBND definitions. The UBND definitions will not be migrated. However, the individual report definitions WITHIN the bundle can be migrated.
- AnyStore reports with report type of ANYS where report data type is *SPLF, *FILE or *BUFFER. You may be able to store this data using the OnDemand generic indexer. See the *IBM Content Manager OnDemand for iSeries Common Server: Indexing Reference* for more information about the generic indexer.
- AnyStore reports with report type of ANYS where object class has been set to a numeric value other than 1 - 7 or 255. Note that numeric values of 1 - 7 are predefined as JPEG, BMP, TIFF, and so on and 255 is predefined as USER-DEFINED. The only known situation where your object class would cause a failure would be if you manually entered a value other than one of these standard, supported values. You may be able to store this data using the OnDemand generic indexer. See the *IBM Content Manager OnDemand for iSeries Common Server: Indexing Reference* for more information on the generic indexer.
- *AFPDSLNE. Reports with report data type of *AFPDSLNE. You will need to consider changing the way this spooled file data is created, perhaps changing it to *APFDS instead of *AFPDSLNE.

Reports that may require attention before they are migrated

- Reports with index exits. Before considering the need to write new Common Server post-processor exit programs (see examples in QSAMPLES2 source file in QRDARS library), review the Common Server capabilities available that may eliminate the requirement for an index exit. See the *IBM Content Manager OnDemand for iSeries Common Server: Administration Guide* for more information about these capabilities. Some examples include:
 - Application options such as the removal of leading, trailing, or imbedded characters before index values are stored.
 - Indexing options such as concatenating two fields into one index value.
 - Folder options such as the ability to use a common date format when searching, regardless of the date format in the spooled file.
- Reports with input exits. These reports require that you write your own program to pre-process the data before OnDemand receives it because the migrated definition will not trigger your current Spool File Archive input exit program. If no Common Server preprocessor program is written, there is a high probability that the report will not store correctly after migration of the definition.
- Reports with viewer exit. The viewer exit does not exist in Common Server; data will be presented in the standard OnDemand Client viewer. If you are using a viewer exit, you should review your need for that exit and see if your requirements can be met by the OnDemand Client viewer. If you truly need an alternate viewer, you might consider the ARSDOC API to retrieve your archived data to be presented in whatever viewer your program chooses.
- Text report overlays (option 3 - Work with Report Overlays - from the OnDemand Report Administration Menu). These text-based overlays are not supported in Common Server. The report definition will migrate, but no text (screen) overlays will appear when viewing. These text-based overlays were primarily used for 5250 (greenscreen) users who did not have the ability to view

documents with AFP overlays on a graphical display. This item is **not** to be confused with AFP overlays, which **are** supported.

- Reports with duplicate key names, such as “Part Number” for key 1 and key 2. If you have reports with duplicate key names, you should rename one of the keys in the Spool File Archive report definition prior to migration to Common Server.
- Reports with blank key names. If you have reports with blank key names, you should specify a key name in the Spool File Archive report definition prior to migration to Common Server.

Migration prerequisites

- Before you can begin migrating your data from Spool File Archive to Common Server, you must already have Common Server installed **on the same iSeries server** on which Spool File Archive is installed. (Common Server is shipped as product option 10 of 5722-RD1.) The OnDemand administrative client and end-user client must also be installed on a Windows workstation so you can validate the migrated definitions and index data. **Note:** See the earlier chapters of this publication for detailed instructions for installing the Common Server. See the Content Manager OnDemand Common Server Administration Guide for detailed instructions for installing the OnDemand Client software.
- If you are using the Kofax Ascent Capture Release Script PRPQ or services offering to archive images using Spool File Archive and AnyStore, you will need to acquire the Common Server Release Script to replace the Release Script you are using. Contact your software provider for more information about acquiring the Common Server Release Script.
- If you are using the AnyStore or Advanced Spool File Archive APIs in your own custom-written application programs, you will need to replace them with the corresponding Common Server APIs. Documentation on the Common Server APIs is available in the *IBM Content Manager OnDemand for iSeries Common Server: Administration Guide*. Visit the Content Manager OnDemand for iSeries support page on the web at <http://www.ibm.com/software/data/ondemand/400/support.html> and search for *comparison* to find a document that provides a comparison between Spool File Archive and Common Server functions, commands, and APIs.

Backups as part of the migration process

It is very important that you take backups of your data at various points during the migration process. Ideally, you would take backups as described below **before and after** each step so that if you have a catastrophic error, you can recover. However, realistically you may not be able to do that due to time constraints. You need to determine the proper balance between the amount of risk and the amount of time you are willing to accept. You also need to consider that restoring from a backup will cause you to write over (and therefore lose) any new data that has been loaded since the backup was taken. The more current the backup, the less data will be overwritten.

At a very minimum, it is recommended that you perform backups at the following times:

- Before you begin the migration process
- Before you begin Phase 3: Index migration
- Before you run Phase 5: Cleanup

Backup steps

1. QUSRRDARS library (which is the library that contains the Spool File Archive index data and control information that points to your archived Spool File Archive data, as well as the Common Server Archive Storage Management (ASM) information for all instances).

```
SAVLIB SAVLIB(QUSRRDARS) DEV(your device name)
```

2. Spool File Archive IFS path (which contains the Spool File Archive data that still resides on disk, along with AFP resources, and so forth, if they exist).

```
SAV DEV(your device name) OBJ('/QIBM/UserData/RDARS/SpoolFile/*')
```

3. QUSROND and/or other instance libraries (which are the Common Server instance libraries that contain index data and control information that points to your archived Common Server data.) You need only backup instance libraries for instances that you are migrating to during migration processing.

If you use the default instance named QUSROND:

```
SAVLIB SAVLIB(QUSROND) DEV(your device name)
```

If you have additional Common Server instances on your system that are being migration to during migration processing:

```
SAVLIB SAVLIB(your instance lib) DEV(your device name)
```

where your instance lib is the name of the instance.

4. Common Server instance IFS path (which contains the Common Server data that still resides on disk, along with configuration files, and so forth, for each of your instances that are being migrated to during migration processing).

```
SAV DEV(your device name) OBJ('/QIBM/UserData/OnDemand/your instance name/*')
```

where your instance name is the name of any additional Common Server instances you have on your system that are being migrated to during migration processing.

Description of different phases of migration

Phase 0	Phase 1	Phase 2	Phase 3	Phase 4	Phase 5
Coexistence	Definition Migration	Definition Conversion	Index Migration	Index Conversion	Clean Up

Phase 0: Coexistence

At this phase, no migration has been done.

When you logon to the end-user client, you can select either your existing Spool File Archive server or a new Common Server instance as the server. If you select a Spool File Archive server, then the folder list that appears looks the same as it did prior to installing the Common Server software. Nothing has changed.

If you select a Common Server instance as the server, you may wish to first modify the instance configuration file named ars.cfg by adding an ARS_MIGR_SERVERS= specification to identify the Spool File Archive server to which this new Common Server instance relates. Then, when you logon to the Common Server instance that contains the ARS_MIGR_SERVERS= specification, the list of folders will contain a combination of user-authorized Spool File Archive folders that contain archived data as well as all of the user-authorized Common Server instance folders.

Included in the folder list may also be the System Log and/or the Migration Log if you are authorized to them. As you can see in Figure 11, the Spool File Archive server name is shown in parentheses to distinguish Spool File Archive folders from the Common Server folders.

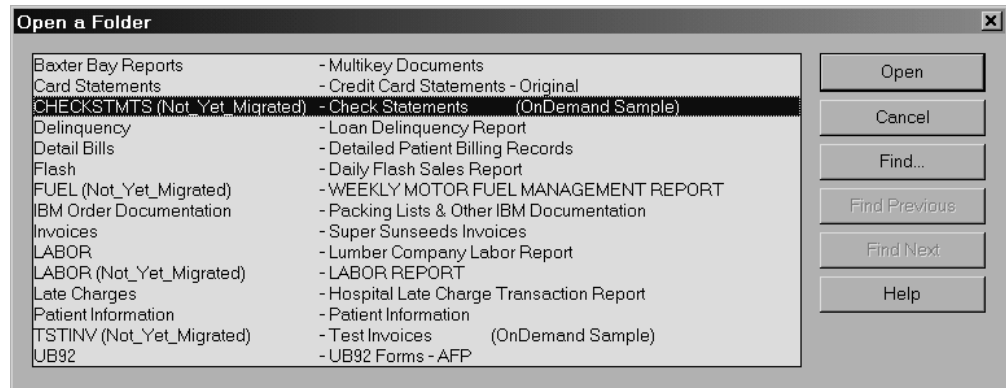


Figure 11. Example of Spool File Archive Folders and Common Server Folders

See Appendix B, “Including Spool File Archive folders in the Common Server folder selection list,” on page 123 for details about setting this configuration specification.

You would still use STRCDSRDAR to store reports for Spool File Archive. New reports could be defined and stored in Common Server.

Phase 1: Definition Migration

During this phase the following actions will occur:

- Definitions are migrated to Common Server
- STRCDSRDAR and QrlsOpenAnyStoreObject give warning message that definition has been migrated
- WRKADMRDAR, WRKSECRDAR, and Operations Navigator give warning messages that definition has been migrated

IBM recommends that definition migration be done in small groups of reports so it can be done in stages.

When this phase is complete, the OnDemand Client communicates to both Spool File Archive and Common Server simultaneously and presents a composite view of both environments just as in the description of Phase 0. As in Phase 0, for Spool File Archive folders, until you have actually stored some reports in Common Server (using the migrated definitions), the only Spool File Archive folders that will appear for the Common Server instance will be those with archived data. The Common Server system log and migration log folders will also appear on the folder list (if you are authorized to them) because they were created when Common Server was installed. You still use STRCDSRDAR to store reports for Spool File Archive, since you have only migrated the **definitions** at this point. New reports could be defined and stored in Common Server.

At this point, it is imperative that you test the migrated definitions by actually storing data in Common Server using the newly-migrated definitions. Pay special attention to the index (key) values that are extracted for the stored data, because converting the indexing parameters (that is, “where the keys and segmentation

criteria are”) is the most critical part of the definition migration. You may wish to store these test reports into a Common Server test instance and then delete that instance when you have completed your testing. Then, you can reissue the definition migration program calls, pointing to your production instance(s) when you are sure they are ready. Be careful not to run the Definition Conversion step (see below) or you will not be able to migration to multiple instances.

Phase 2: Definition Conversion

IMPORTANT: After you run this step, you cannot start over, because you risk the creation of duplicate data in your Common Server archive. For this reason, you must be absolutely certain that you are satisfied with the results of your parallel testing described in Phase 1.

During this phase the following actions will occur:

- Convert to storing the data using
 - ADDRPTOND instead of STRCDSRDAR for storing individual reports
 - STRMONOND TYPE(*OUTQ) instead of STRMONRDAR for output queue monitoring
 - arslod API instead of calling the QrlsOpenAnyStoreSegment, QrlsStoreAnyStoreSegment, and QrlsCloseAnyStoreSegment APIs for AnyStore
 - STRMONOND TYPE(*DIR) or arslod API instead of STRMONANYS command or calling the QRLSSTART program for Kofax integration
- STRCDSRDAR and QrlsOpenAnyStoreObject give escape message that data archiving for this report definition is no longer allowed.
- STRMONRDAR will fail to store the report, and the spooled file will be placed in HLD (Held) status on the Error output queue specified on the monitor command.
- WRKADMRDAR and WRKSECRDAR prohibit changes to the report definition, since it has already been migrated. Operations Navigator will prohibit update or adding of a new version of the migrated report definitions.

When this phase is complete, the data and indexes for migrated report definitions continue to reside in Spool File Archive. The report definitions that were migrated now use ADDRPTOND, STRMONOND, or ARSLOAD (see bullets above) to store reports. This conversion phase protects you from loading the same spooled file data in both Spool File Archive **and** Common Server. If you were able to load the same report into both environments, you would end up with duplicate data after Index Migration (Phase 3).

Phase 3: Index Data Migration

During this phase the following actions will occur:

- Migrate index data, resources, and annotations to Common Server. Note that a date range parameter is provided for Phases 3 and 4 to allow for gradual migration of index data.
- Each stored report occurrence in Spool File Archive will be flagged when it has been migrated.
- DLTRPTRDAR is not allowed; an escape message is sent, indicating that the report cannot be deleted from Spool File Archive since it has been migrated to Common Server.

When this phase is complete, the index records, annotations, and AFP resources for the migrated Spool File Archive reports can be retrieved by Common Server using the migrated index records. The report objects themselves will not move to the Common Server; they will remain in their current Spool File Archive location and will be accessed from that location. No Spool File Archive functions, other than storage management, are used (including the running of the Report Management Cycle (RMC)). If a date range has been specified in order to perform a gradual migration over time, end-users should continue to use Spool File Archive for production retrievals until all indexes for all dates have been migrated. If end-users use Common Server while the individual date ranges of index records are being migrated, they will only be working with a subset of the entire archive until all indexes have been migrated.

To verify the migrated index data, a variety of retrievals should be performed in parallel between Spool File Archive and Common Server. For example, a Spool File Archive retrieval of '*ALL' in the first key could be compared to a Common Server retrieval using LIKE as the operator and '%' (percent) as the value. Compare the counts of items on the list, and ALSO compare the actual values for the keys. If mismatches are found, contact your OnDemand support provider.

Phase 4: Index Data Conversion

During this phase the following actions will occur:

- You manually verify migrated index data, annotations, and AFP resources.
- FNDKEYRDAR, FNDRPTRDAR, QrlsRetrieveAnyStoreSegment, QrlsRetrieveAnyStoreList, QrlrRetrieveReportSegment, and QrlrRetrieveReportKeyList all give escape messages indicating retrieve not allowed.
- PRTRPTRDAR will no longer work; RCLRPTRDAR will still work as long as the archived data is still in the expected location.

When this phase is complete, migration is complete for the specified report, or for the specified date range, if a gradual approach is used. Spool File Archive storage management is still required (including the running of the Report Management Cycle (RMC) for as long as you have data on your system that was originally archived with Spool File Archive). No other Spool File Archive functions are required.

Phase 5: Cleanup

During this phase the following actions will occur:

- Removes all index data, all resources, all annotations for the specified report definition from Spool File Archive.
- Removes all Content Manager/VisualInfo (CM/VI) integration records from the CM/VI control files for the specified report definition, if applicable.

When this phase is complete, you have removed the Spool File Archive index, annotation, and AFP resource data for the specified report(s) that has been migrated to Common Server, along with any CM/VI control data if applicable.

This step does not actually reclaim the disk space from the data that the cleanup process deletes. OS/400 requires that you run the Reorganize Physical File (RGZPFM) command in order to actually reclaim the disk space. However, access paths will shrink immediately upon record deletion. Instructions for the Cleanup phase include a list the files that you may wish to reorganize.

IMPORTANT

If you neglect to run Phase 4 and Phase 5, Spool File Archive's Report Management Cycle (RMC) will never expire any report data or delete any Stored Report Table (QARLSRT) records, nor will the index records for those reports ever be deleted from Common Server. For this reason, it is very important that you run these two final phases of migration.

Background information

Sources of data for the migration of definitions

- Information for Users and User Groups is gathered from user profiles and authorization lists associated with Spool File Archive.
- Information for Application Groups is gathered from the report definitions file (QARLRACT), authorization lists, and security files (QARLRSECn).
- Information for Applications is gathered from the report definition file (QARLRACT). Each version of a report definition is converted into a separate application definition with an application name reflecting the report name and version. For example, CHECKS version 01 would be Common Server application name CHECKS-01. In addition, an extra application definition with just the name of the report is created that can be used to store new information. Public and private logical views are also migrated (QARLRLV).
- Information for Folders is determined from the report definition file (QARLRACT), authorization lists and from the folder field records file (QARLRFFU).
- Information for Common Server migration policies is gathered from Spool File Archive migration policies.
- Information for Printers is not migrated.

Notes:

1. Because the system cannot determine passwords from the OS/400 user profiles, the internal OnDemand passwords for migrated users will be set to ONDEMAND. This should not be a problem because most customers will be using the Common Server default which links OS/400 passwords to OnDemand passwords rather than using a separate, internal OnDemand password. In the default case, using OnDemand as shipped, these OnDemand-specific passwords will be ignored. See "OnDemand userid relationship to OS/400 user profiles" on page 40 if you are unsure about whether you have disabled the default linkage between OS/400 passwords and OnDemand internal passwords.
2. If authorization list QRDARS400 has *PUBLIC set to anything other than *EXCLUDE and at least one report authorization list has *PUBLIC set to anything other than *EXCLUDE, then all non-IBM user profiles will be migrated as OnDemand users. This could cause a large number of OnDemand users to be created that you do not need. If this is the case on your system, you may want to take the time to analyze the OnDemand authorization lists and make changes if you want to limit the number of OnDemand users that are created by the migration.

Scheme for creating Common Server applications, application groups, and folders from existing Spool File Archive report definitions

One Common Server **application** will be created for each Spool File Archive **version** of a report definition. A new application will be created only for Spool File Archive report definition versions for which reports are actually stored, based on the existence of a record in the QARLSRST file. In addition, an extra application definition with just the name of the report is created that can be used to store new information.

One Common Server **application group** and **folder** will be created for each Spool File Archive **report name**. Just as with applications described previously, a new application group and folder will be created only for Spool File Archive report names which actually have data stored, based on existence of a record in the QARLSRST file.

If there is a mismatch between versions regarding the number of key fields or the size of the key fields, the resulting **application group** and **folder** will be created as a “composite” of all versions. Each key field will be defined with the largest of all versions’ key sizes. (For example, if Invoice Number is 8 positions long in version 01 and 10 positions long in version 02, the migrated Invoice Number key would be 10 positions long.) The number of keys will equal the greatest number of keys specified in any of the versions. (For example, if version 01 has 3 keys and version 02 has 4 keys, the migrated application group and folder would have 4 keys defined.)

A report listing all of the mapping between Spool File Archive report definitions and Common Server applications, application groups, and folders is generated during the definition migration step. The report will list key mapping between versions as well.

Example

Based on the approach described above, a Spool File Archive system with 3 versions of a PAYROLL report would result in a Common Server setup that looks like this:

Applications: PAYROLL-01, PAYROLL-02, PAYROLL-03, and PAYROLL

Note here that PAYROLL would be identical to PAYROLL-03. PAYROLL is created for the purpose of storing of new reports directly into Common Server. PAYROLL-03 is reserved for the most recent archives of PAYROLL in Spool File Archive.

Application Group and Folder: PAYROLL, which matches the PAYROLL application name above

Scheme for creating Common Server folders from existing Spool File Archive report group definitions

One Common Server **folder** will be created for each Spool File Archive **report group**. No application or application group by that name will be created. Individual reports within the group will (as described previously) map to application groups within the folder. All applications and application groups within the report group that have not already been migrated will be migrated. If any failures occur in the migration of the individual report definitions within the report group, the folder will not be created.

Scheme for creating Common Server users and user groups from existing Spool File Archive authorization lists

An OS/400 group profile would become an OnDemand Common Server user group if one of the following is true. ("Authorized" means the authority is not *EXCLUDE, and "group" is any group profile not beginning with 'Q', but includes QRDARS400 and QRDARSADM.)

- OS/400 group profile is QRDARS400 or QRDARSADM
- OS/400 group profile authorized on QRDARS400 authorization list
- OS/400 group profile authorized on any report authorization list

Important Note: When an OnDemand group gets created, a group id number is given to the group which determines the precedence that the group has over other groups within OnDemand with respect to authority. The lower the group id number, the higher the precedence. For example, user JBROWN is part of an OnDemand group called PAYROLL and also part of a group called MANAGERS, The PAYROLL group id number is 80020. The MANAGERS group id number is 80001 (a lower number). When JBROWN wants to look at payroll data, the authority that he has to the payroll data is determined by the authority that the MANAGERS group has, since it has a lower group id number. It is important to understand that your Spool File Archive group profiles will be migrated to Common Server in alphabetical order, resulting in the group id numbers being lower for the group profiles with names that begin with letters near the beginning of early in the alphabet. You may need to review the migrated groups to make sure they will provide the authorities that you require, and modify which users are on which groups if necessary, since group id numbers cannot be changed once the group is created. You may also choose to delete/recreate or copy migrated groups to make new groups with different group id values in order to provide the security checking sequence that you require. (Also note that this OnDemand group id field is not related to the GID field found in the OS/400 user profile.)

An OS/400 user profile would become an OnDemand Common Server user if one of the following is true. ("Authorized" means the authority is not *EXCLUDE, and "user" is any profile not beginning with 'Q', but includes QRDARS400 and QRDARSADM.)

- all users, if *PUBLIC is authorized on QRDARS400 authorization list and at least 1 report authorization list
- user has *ALLOBJ authority
- user authorized on QRDARS400 authorization list user authorized on any report authorization list and authorized through QRDARS400
- user is a member of a group which has become an OnDemand group, as selected above

Scheme for creating Common Server query restriction from existing Spool File Archive key security definitions

The contents of the Spool File Archive key security files will be "transcribed" into an SQL statement to be placed into the corresponding Common Server application group's Query Restriction field.

Important Note: If you use group profiles in your Spool File Archive key security definitions and you have a user profile that belongs to more than one group or you have the user profile explicitly listed as well as part of a group, the "key security" results in Common Server may be different than the results in Spool File Archive. Carefully check the resulting query restrictions in your application group

definition(s) and actually test some restricted retrievals to confirm that the migrated "key security" authorizations provide the results that you require. See the OnDemand administrator client online help topic *Query Restriction* for more information and examples.

Migration of definitions: Steps to follow

These migration instructions assume most customers will migrate their users and user groups, then policies and optical storage groups, and then begin with a few report definitions (as opposed to trying to move all report definitions at once, and put them all into production in the Common Server environment at one time.) This phased approach to report definition migration allows for ample time to be spent verifying the migrated definitions, which is the most important part of the definition migration process.

Identifying report definitions to migrate

In Spool File Archive, you capture reports and later view segments (also called documents) of the report by searching on indexes that have been defined for the report or report group. In Common Server, you capture applications and later view documents of the applications by searching on indexes that have been defined for the folder. The concept is the same, but the terminology is different.

The migration of a report definition (with one or more versions) results in a set of applications, an application group, and a folder. As stated in the introduction, it is expected that customers will migrate a few report definitions at a time, validate the migrated definitions, then begin capturing the reports using the Common Server. The process would then be repeated for another set of report definitions until all reports are being captured as applications in the Common Server environment. The actual definition migration is the quick and easy part; validating the migrated definitions by storing the same spooled file in Spool File Archive and in Common Server and comparing the results can take a considerable amount of time. However, it is time well spent, because careful verification at that time will ensure the integrity and usability of the data once it is in production in the Common Server environment.

The migration software provides an analysis report that can be run prior to the actual definition migration to help identify any reports that may need some attention prior to migration. The report can also help you to identify the report definitions that appear ready to migrate. See "Phase 1: Definition migration steps" on page 113 for more information about the *ANZDFN run type used to generate the analysis report.

Creating a TEST instance for definition migration testing

IBM highly recommends that you create a test instance (such as ONDTEST) for use in Phase 1: Definition Migration. This will allow for parallel testing and significant user verification prior to running your "production" Definition Migration. (To do significant user verification for Phase 1, you will have to actually store data using the migrated definitions before moving to Phase 2.) This instance can then be deleted when you have finished making any changes and are ready for Definition Migration into your production instance(s). If you actually load data during this parallel testing without using (and then deleting) a test instance (in other words, if you load data directly into a production instance), then you would end up with duplicate data after Index Migration (Phase 3) unless you manually delete all the loaded data with the Remove Report (RMVRPTOND) command. Just remember

not to run the Definition Conversion (phase 2) step until you have migrated to your **production** instance(s) so that OnDemand will permit multiple runs of the definition migration program.

In most cases, your production instance will be QUSROND, which is the default Common Server instance. QUSROND is the default value for the instance parameter on the OnDemand iSeries commands, so using this instance name as your production instance saves you from having to specify an instance name when you issue a command. Refer to Chapter 12, “Creating an instance,” on page 69 for instructions about creating an instance.

Creating all PRODUCTION instances prior to production definition migration

All the migration steps require specification of which Common Server *instance* to migrate to. This means that the one or more required instance(s) must be created prior to running migration “in production.” Refer to Chapter 12, “Creating an instance,” on page 69 for instructions about creating an instance, and for an explanation of why you might want more than one.

Creating a backup of your data

Refer to “Backups as part of the migration process” on page 103 for specific information about the data to be backed up.

Running the definition migration programs

The steps to migrate the Spool File Archive definitions to the new Common Server are outlined below. These migration steps are initiated using program calls (rather than commands). As an added control feature, none of the parameters have defaults; all parameters must be specified, or the program call will fail.

Requirements for running definition migration steps

- You must run the migration programs using QSECOFR or a user profile with *ALLOBJ authority.
- The target instance must have an administrator user defined called QONDADM. (This is automatically created when you create an instance.)
- The user profile that issues the migration program calls must have its locale set correctly, matching the locale associated with the instance to which the definitions or index data is being migrated. Because the user profile must be QSECOFR or another profile with *ALLOBJ special authority, you will have to confirm or change the locale for QSECOFR or the user profile with *ALLOBJ special authority. **(For this reason, you may prefer to use a profile other than QSECOFR. Also, do not forget to change it back to its original value after migration, if necessary.)** Further, if you are migrating to more than one instance and the instances require different locales, you will need to change the locale of the user profile, signoff and sign back on again in between migration program calls to the different instances. See Chapter 13, “Defining a locale,” on page 75 for more information on locales.
- The user running the migration process must have Maximum storage set to *NOMAX in their user profile.
- The MAXTMPSTG parameter of their job class must be set to *NOMAX.
- The utilization of the system auxiliary storage pool must not exceed 90% when a migration process is started.

Phase 1: Definition migration steps

1. Run users/user groups migration program. Two parameters are required, which are the run type (*MGRUSR) and the name of the instance to which the users and user groups are to be migrated. All users and groups are migrated when this program is called. The migration program produces a report that lists all OnDemand users and user groups that were created for Common Server from the Spool File Archive authorization lists.

For example:

```
CALL PGM(QRLRMIG) PARM('*MGRUSR' 'QUSROND')
```

where *MGRUSR identifies the run type and QUSROND is the name of the instance to which the information will be migrated.

2. Run policy migration program (which also migrates optical storage groups). Two parameters are required, which are the run type (*MGRPCY) and the name of the instance to which the policies and optical storage groups are to be migrated. All policies and optical storage groups are migrated when this program is called. The migration program produces a report that lists all Spool File Archive report migration policies and optical storage groups that were used to create Common Server storage nodes, storage sets, migration policies, and optical storage groups.

For example:

```
CALL PGM(QRLRMIG) PARM('*MGRPCY' 'QUSROND')
```

where *MGRPCY identifies the run type and QUSROND is the name of the instance to which the information will be migrated.

3. Run report definition analysis program. The analysis program can be used to identify report definitions that may require attention prior to migrating. The analysis report lists:
 - Applications, Application Groups, and Folders that will be created for Common Server from the Spool File Archive report and report group definitions.
 - Report definitions that have key definitions that don't match from one version to the next. The mismatch could be due to differences in the number of keys defined, or in the sizes of the key fields as one version is compared to the next.
 - Report definitions for report versions that indicate any of the conditions listed above under "Reports that Should NOT be Migrated" or "Reports that May Require Attention Before They Are Migrated".

Three parameters are required to run the analysis program:

- a. Run type (*ANZDFN)
- b. Instance name
- c. Report name (specific name, generic, or *ALL). All versions of a report will be analyzed at one time. Both report definitions and report group definitions can be processed in this one program run. If report group name is specified, it will pull in all report definitions in the group and analyze them as well. Note that a maximum of 200,000 report or report group definitions can be analyzed in a single run of this program. A message will be issued if your request exceeds this maximum. If this occurs, you can rerun the program and provide a more specific report name for this parameter.

For example:

```
CALL PGM(QRLRMIG) PARM('*ANZDFN' 'QUSROND' '*ALL')
```

where *ANZDFN identifies the run type, QUSROND is the name of the instance for which the information will be analyzed., and *ALL defines the report definition name(s) to be analyzed. The program produces a report titled “Migrate Report Definitions” with a spooled file name of QPRLRMGRD.

4. Run report definition migration program. (This step will migrate public and private logical views as well.) The migration program produces a report that lists:
 - Applications, Application Groups, and Folders that were created for Common Server from the Spool File Archive report and report group definitions.
 - Report definitions that have key definitions that don’t match from one version to the next. The mismatch could be due to differences in the number of keys defined, or in the sizes of the key fields as one version is compared to the next.
 - Report definitions for report versions that indicate any of the conditions listed above under “Reports that Should NOT be Migrated” or “Reports that May Require Attention Before They Are Migrated”.

Three parameters are required to run the migration program:

- a. Run type (*MGRDFN)
- b. Instance name
- c. Report name (specific name, generic, or *ALL). (All versions of a report will be migrated at one time.) Both report definitions and report group definitions can be processed in this one program run. If report group name is specified, it will pull in all report definitions in the group and migrate them as well. If *ALL or a generic name is specified, the program will monitor for duplicates and insure that none is created. Note that a maximum of 200,000 report or report group definitions can be migrated in a single run of this program. A message will be issued if your request exceeds this maximum. If this occurs, you can rerun the program and provide a more specific report name for this parameter.

For example:

```
CALL PGM(QRLRMIG) PARM('*MGRDFN' 'QUSROND' 'PAY*')
```

where *MGRDFN identifies the run type, QUSROND is the name of the instance to which the information will be migrated, and PAY* represents the report definition name(s) to be migrated. **Note:** There may be some report definitions for which indexer parameters cannot automatically be generated. The definitions will be migrated, and a warning message will be listed on the resulting report to describe the issue.

Phase 2: Definition conversion step

IMPORTANT: After you run this step, there is no way to go back and start over. After you run this step, you cannot re-migrate any definitions (due to either changes you have chosen to make to the definition or due to the need to re-migrate in order to populate a production instance after first testing with a test instance). After you run this step, OnDemand will not allow storage of reports into Spool File Archive. For these reasons, be very careful before choosing to run this step.

1. Run definition conversion program. A report is produced that will list all report definitions that have been converted.

Three parameters are required:

- a. Run type (*CVTDFN)

- b. Instance name
- c. Report name (specific name, generic, or *ALL). (All versions of a report will be converted at one time.) Both report definitions and report group definitions will be processed in this one program run. If a report definition is specified, it must have been migrated successfully during Phase 1. If a report group definition is specified, the reports which are members of the group will also be converted, if they were migrated in Phase 1. If all reports in the group were not migrated successfully during Phase 1, the report group definition will not be converted. If a generic name is specified, the program will process only report definitions (not report group definitions) that satisfy the generic specification. The reason for this is that, if report group definitions were processed when using a generic specification, then there is a high probability that report definitions within the group would be converted even though they did not match the original generic specification. In all cases, the program will monitor for duplicates and insure that none is processed.

For example:

```
CALL PGM(QRLRMIG) PARM('*CVTDFN' 'QUSROND' 'PAY*')
```

where *CVTDFN identifies the run type, QUSROND is the name of the instance for which the information will be converted, and PAY* represents the report definition name(s) to be converted.

Modifying application group definitions

Because Spool File Archive report definitions do not distinguish between string (character), integer (numeric), and other special data types for key fields, all keys are stored as string fields. For this reason, the definition migration program migrates the key fields as string fields in the resulting application group definitions in Common Server. When searching for documents with the OnDemand Client, string fields allow for the use of Like searches (for example, Customer Name Like Smit%), but integer works better for numeric type operations such as greater than and between, for example. There are also query restriction (key security) considerations. As an example, if Account Number is defined as a string, specifying a query restriction of "Account Number > 10000" may not work as you might expect. But if Account Number is an integer, you would be more likely to get the result you expect.

If there are a few critical reports with key fields that your users need to perform numeric operations on, for example, you might want to consider creating your own application, application group, and folder, and defining the key fields as desired. It is important to understand, however, that Common Server folders cannot contain application groups with dissimilar key field attributes. In other words, you cannot map fields from different application groups to one folder field where one of the application groups has customer number defined as a string and the other application group has customer number defined as a numeric. Therefore, if you want to take advantage of the new data types for your key values, you will need to have two separate folders; one with the Spool File Archive data (all keys defined as string fields) and one that you create with your new Common Server data (keys defined with data types of your choosing).

If you have a critical report that must take advantage of the new data types for keys, you may also choose to use the Spool File Archive's Print Report for

OnDemand (PRTRPTRDAR) command to respool the data, then load into Common Server using a new application, application group, and folder that have the new data types defined for the keys.

In either case, you should rename the migrated application, application group, and folder and then use the copy function if you choose to make your own definitions using the new data types. That approach will preserve the migrated definitions from Spool File Archive for your reference.

Testing the migrated definitions

You should validate that the migrated definitions came through correctly by viewing the new applications, application groups, and folders using the OnDemand Administrative client. Pay particular attention to the Indexer panel of the application definition, as there may be some manual intervention required, such as changes to the trigger fields or range values. There may be some situations where this migration may not produce the results you desire. Review these parameters, store some test reports for each application, and adjust the indexing parameters as necessary.

IMPORTANT: Spool File Archive reports with report type of PAGE will migrate successfully but must be modified prior to use in Common Server. The Common Server application definitions that result from PAGE report migration will contain a *mask* that identifies which data on the print page should be used as index data. For example, a mask of #####.## would cause the indexer to select the field only if the data in the field (from left to right) contains four numeric characters, followed by a decimal point, followed by two numeric characters.) Because the migration program will have no way of knowing the nature or format of the data that is to be used for indexing, it will create a very generic mask that allows ANY data to be considered valid index data. You will need to update the resulting application definition for each migrated Spool File Archive PAGE report to specify a mask that more appropriately describes the data, or else the header lines of the report will likely be picked up as index values, making the index values meaningless and making retrieval of the data unpredictable.

You may also want to pay attention to the Load Information tab in the application definition. The specifications to strip leading, imbedded, and trailing characters are defined on this tab, and can be a significant benefit, especially if you had written index exit programs in Spool File Archive to accomplish the same thing.

Preparing to capture production reports using the Common Server

When you are satisfied that the application, application group, and folder definitions meet your needs, then you can prepare for capturing production reports in the Common Server. If you have the output queue monitor (STRMONRDAR) or report capture (STRCDSRDAR) command set up in a job scheduler or in a production CL program, this will need to be modified to use the new Common Server command (STRMONOND and ADDRPTOND, respectively).

Migration of index data: Steps to follow

Migration of index data (which includes annotations and AFP resources as well) is the next important step in this migration process. It is not mandatory that you run index migration immediately after definition migration; in fact, it is more likely that you will **not** run index migration immediately after definition migration.

The index migration program will check to make sure the actual migration of the report definition has occurred, prior to the attempt to migrate indexes. In fact, however, you will likely discover these omissions yourself as you do testing during the definition migration phase.

After indexes are migrated, users cannot view the data from Spool File Archive interfaces. All storage and retrieval functions are performed using the appropriate Common Server function.

Important

DURING the index migration process for a particular report, the total space required for index migration could equal THREE times the size of the report's index data as it currently resides in Spool File Archive. AFTER the index migration process completes for the report, as much as TWO times the space may still be required until you run Cleanup (Phase 5) to delete the index data from Spool File Archive. Further, if you have archived any of your Spool File Archive **indexes** to optical or tape, they will have to be recalled to disk during the migration process, which will take additional disk space. (The recalled Spool File Archive indexes will be deleted at the end of a successful index migration to Common Server, so they will only require disk space for a limited period of time.) However, the migrated indexes will then be back on disk in the Common Server environment, which will consume additional disk space that the Spool File Archive indexes on optical or tape did not consume.

A date range parameter is provided for index migration to address two important subjects:

- The disk space requirement described above. The date range allows you to run the index migration process in manageable increments, being careful not to overload your disk space.
- The index migration elapsed time requirement. The index migration process could also be very time-consuming. If you have limited time in which to run the migration process, you should plan to perform the migration using the date range parameters. Note that if you have migrated indexes to optical or tape, index migration will take extra time due to the slower nature of those media types.

Creating a backup of your data

Refer to "Backups as part of the migration process" on page 103 for specific information about the data to be backed up.

Running the index data migration programs

The steps to migrate the Spool File Archive index data to the new Common Server are outlined below. These migration steps are initiated using program calls (rather than commands). As an added control feature, none of the parameters have defaults; all parameters must be specified, or the program call will fail.

Requirements for running index migration steps

- Must run the migration programs using QSECOFR or a user profile with *ALLOBJ authority.
- The target instance must have an administrator user defined called QONDADM. (This is automatically created when you create an instance.)
- The user profile that issues the migration program calls must have its locale set correctly, matching the locale associated with the instance to which the

definitions or index data is being migrated. Because the user profile must be QSECOFR or another profile with *ALLOBJ special authority, you will have to confirm or change the locale for QSECOFR or the user profile with *ALLOBJ special authority. **(For this reason, you may prefer to use a profile other than QSECOFR. Also, do not forget to change it back to its original value after migration, if necessary.)** Further, if you are migrating to more than one instance and the instances require different locales, you will need to change the locale of the user profile, signoff and sign back on again in between migration program calls to the different instances. See Chapter 13, “Defining a locale,” on page 75 for more information on locales.

- The user running the migration process must have Maximum storage set to *NOMAX in their user profile.
- The MAXTMPSTG parameter of their job class must be set to *NOMAX.
- The utilization of the system auxiliary storage pool must not exceed 90% when a migration process is started.

Phase 3: Index migration steps

1. Run index analysis program. A report is produced that lists the number of index records that will require migration (as well as the number of index records that must be recalled from optical or tape, if you have chosen to archive indexes in Spool File Archive).

Five parameters are required:

- a. Run type (*ANZIDX)
- b. Instance name
- c. Report name (specific name, generic, or *ALL). Index analysis should not be run “massively” but in small, controlled pieces that coincide with your plan for the actual index migration. **Generic specifications and *ALL should be used with caution, as they could initiate a very long-running job.**
- d. Date range - starting date (in YYYYMMDD format)
- e. Date range - ending date (in YYYYMMDD format)

For example:

```
CALL PGM(QRLRMIG) PARM('*ANZIDX' 'QUSROND' 'PAYCHECKS' '19900101' '19920101')
```

where *ANZIDX identifies the run type, QUSROND is the name of the instance to which the information will be analyzed, PAYCHECKS represents the report definition name(s), and 19900101 and 19920101 represent the from- and to- date range of the index records to be analyzed (in YYYYMMDD format). The program produces a report titled “Migrate Indexes” with a spooled file name of QPRLRMIDX.

2. Run prefetch program to recall (back to disk) any Spool File Archive index records that have been archived to optical or tape. THIS STEP IS NOT REQUIRED if you did not choose to archive **index records** in your Spool File Archive policy definition. Also, it is **not** mandatory that you run the prefetch program before migrating indexes. If you prefer, you can allow the index migration program to recall the indexes as part of its processing.

If you prefer to prefetch your indexes, run the prefetch program at this time.

Five parameters are required:

- a. Run type (*RCLIDX)
- b. Instance (*NONE)

- c. Report name (specific name, generic, or *ALL). **Generic specifications and *ALL should be used with extreme caution, as they could initiate an extremely long-running job with significant disk space consumption.**
- d. Date range - starting date (in YYYYMMDD format)
- e. Date range - ending date (in YYYYMMDD format)

For example:

```
CALL PGM(QRLRMIG) PARM('*RCLIDX' '*NONE' 'PAYCHECKS' '19900101' '19920101')
```

where *RCLIDX identifies the run type, *NONE is the name of the instance (which does not apply in this step, PAYCHECKS represents the report definition name(s), and 19900101 and 19920101 represent the from- and to- date range of the index records to be recalled to disk (in YYYYMMDD format).

3. Run index migration program. (This step will migrate annotations and AFP resources as well.)

Five parameters are required:

- a. Run type (*MGRIDX)
- b. Instance
- c. Report name (specific name, generic, or *ALL). **Generic specifications and *ALL should be used with extreme caution, as they could initiate an extremely long-running job with significant disk space consumption.** Index migration should not be run “massively” but in small, controlled pieces since the process can be very time-consuming and very disk space-consuming if large amounts of index records are involved.
- d. Date range - starting date (in YYYYMMDD format)
- e. Date range - ending date (in YYYYMMDD format)

For example:

```
CALL PGM(QRLRMIG) PARM('*MGRIDX' 'QUSROND' 'PAYCHECKS' '19900101' '19920101')
```

where *MGRIDX identifies the run type, QUSROND is the name of the instance to which the information will be migrated, PAYCHECKS represents the report definition name(s), and 19900101 and 19920101 represent the from- and to- date range of the index records to be migrated (in YYYYMMDD format).

Phase 4: Index conversion step

IMPORTANT: After you run this step, there is no way to go back and start over. After you run this step, you cannot re-migrate any report indexes. For these reasons, be very careful before choosing to run this step.

1. Run index conversion program.

Five parameters are required:

- a. Run type (*CVTIDX)
- b. Instance name
- c. Report name (specific name, generic, or *ALL) **Generic specifications and *ALL should be used with caution, as they could initiate an extremely long-running job.** Index conversion should be run specifically for reports ready for conversion; not “massively” with the risk of converting before you are really ready.
- d. Date range - starting date (in YYYYMMDD format)
- e. Date range - ending date (in YYYYMMDD format)

For example:

```
CALL PGM(QRLRMIG) PARM('*CVTIDX' 'QUSROND' 'PAYCHECKS' '19900101' '19920101')
```

where *CVTIDX identifies the run type, QUSROND is the name of the instance to which the information will be converted, PAYCHECKS represents the report definition name(s), and 19900101 and 19920101 represent the from- and to- date range of the index records to be converted (in YYYYMMDD format).

Cleanup following successful migration

Cleanup of migrated Spool File Archive definition records and index records will not automatically occur after migration, even if the migration appears to have run successfully. Cleanup is a separate step that can and should be deferred until all parallel testing and verification is complete to the satisfaction of all interested parties. It is important that YOU (and not OnDemand code) own the responsibility to determine when the time is appropriate to remove the old Spool File Archive data that has been migrated to Common Server. Additionally, if there is some kind of error or oversight in the migration process that is discovered late in the process, you have not lost the input data that would allow you (along with software support) to fix and rerun.

The steps to run cleanup are outlined below. The cleanup program is initiated using a program call (rather than a command). As an added control feature, none of the parameters have defaults; all parameters must be specified, or the program call will fail.

Phase 5: Cleanup steps

1. Backup your data prior to running the cleanup program. Refer to “Backups as part of the migration process” on page 103 for specific information about the data to be backed up.
2. Run cleanup program.
Five parameters are required:
 - a. Run type (*RMVIDX)
 - b. Instance name
 - c. Report name (specific name, generic, or *ALL) **Generic specifications and *ALL should be used with caution, as they could initiate a very long-running job.** Cleanup should not be run “massively” but in small, controlled pieces since the process can be time-consuming if large amounts of index records are involved.
 - d. Date range - starting date (in YYYYMMDD format)
 - e. Date range - ending date (in YYYYMMDD format)

For example:

```
CALL PGM(QRLRMIG) PARM('*RMVIDX' '*NONE' 'PAYCHECKS' '19900101' '19920101')
```

where *RMVIDX identifies the run type, *NONE is the name of the instance (which does not apply for this step), PAYCHECKS represents the report definition name(s), and 19900101 and 19920101 represent the from- and to- date range of the data to be cleaned up (in YYYYMMDD format).

3. The following is a list of files in library QUSRRDARS that you should reorganize after ALL reports have been migrated successfully:
 - QARLRRC
 - QARLRANN

- QARLR000PF
- QARLRxxxPF, where *xxx* is the three-character report group abbreviation. (These files exist only if you are using report groups within Spool File Archive. You will have one of these files for every report group.)

You have now completed your Content Manager OnDemand Spool File Archive to Common Server migration. Please remember that you will be required to keep Spool File Archive (OnDemand product option 1) installed on your system until all data stored by Spool File Archive has expired.

Appendix B. Including Spool File Archive folders in the Common Server folder selection list

As you begin your work on migration, you may find it valuable for your end-users to see a folder selection list on the OnDemand Client that contains folders from both Spool File Archive and Common Server environments. (Keep in mind that the term *folder* in Common Server is similar to the term *report name* or *report definition* in Spool File Archive.) Setting the configuration parameter for this requires only that you update the `ars.cfg` file for the Common Server instance name for which you choose to include Spool File Archive folders. (Note that, for OnDemand Web Enablement Kit (ODWEK) users, the Spool File Archive folders will appear in the ODWEK folder list and can be searched. However, attempting to retrieve a document will fail.)

The `ars.cfg` file is found on your iSeries server in an IFS directory. Using the OS/400 command to Work with Object Links (WRKLNK) command, position yourself to the following directory:

```
/QIBM/UserData/OnDemand/<instance name>
```

where `<instance name>` is the name of the Common Server instance for which you choose to include Spool File Archive folders.

Locate the `ars.cfg` file in your instance's directory. Edit it and add a line to the bottom of the file that looks similar to the following:

```
ARS_MIGR_SERVERS=NotYetMigrated my400.mynetwork.mycompany.com 2 1445 0 0
```

which shows six positional values following the `ARS_MIGR_SERVERS=` parameter name. **Note:** You will change the first, second, and possibly the fourth parameter; the third, fifth, and sixth parameters must be set as shown in Table 7.

Table 7. *ARS_MIGR_SERVERS* Configuration Parameters

Parameter	Meaning
NotYetMigrated	Name that will appear in parentheses after all Spool File Archive folder names in the OnDemand Client folder selection list. (Common Server folders on the folder selection list will have nothing following their names.) This will help to distinguish Spool File Archive folders from Common Server folders.
my400.mynetwork.mycompany.com	Host name or TCP/IP address of the iSeries server. This is the iSeries server name where both Spool File Archive and Common Server reside.
2	THIS VALUE MUST ALWAYS BE SET™ TO THE NUMERIC VALUE OF 2

Table 7. ARS_MIGR_SERVERS Configuration Parameters (continued)

Parameter	Meaning
1445	Port number of the Spool File Archive environment. THIS VALUES WILL USUALLY BE SET TO 1445, which is the default that Spool File Archive uses for communications with the server. (You may notice that the server definition on your OnDemand Client actually says Port 0. This is normal. Specifying Port 0 indicates to the client that you wish to use the default port of 1445.) If you have changed this port number for your Spool File Archive environment for any reason, you would specify the new port number here.
0	THIS VALUE MUST ALWAYS BE SET TO THE NUMERIC VALUE OF 0.
0	THIS VALUE MUST ALWAYS BE SET TO THE NUMERIC VALUE OF 0.

Appendix C. Database servers

The ARS_NUM_DBSRVR parameter in the ARS.CFG file determines the number of processes that OnDemand can start on the server to handle database requests. You can use the ARS_NUM_DBSRVR parameter to balance performance, by distributing the database workload over several processes, and use of system resources.

Description

Each access to the database requires a database agent. OnDemand can start a database agent for each access. However, this consumes a lot of system resources because a new job is started each time a database agent is started. You can use the ARS_NUM_DBSRVR parameter to optimize the way that OnDemand handles the database load. For example, you can define ARS_NUM_DBSRVR so that OnDemand starts a fixed number of database agents, regardless of the number of concurrent database requests. While this may appear restrictive, database requests typically process very quickly.

Factors to consider when setting the ARS_NUM_DBSRVR parameter include:

- Number of concurrent users
- Number of data loading jobs active
- Number of database maintenance processes active

We recommend that you set the ARS_NUM_DBSRVR parameter to support the peak number of concurrent database accesses that you expect the server to handle. A low value limits access to the database during periods of high database activity. A high value requires more system resources during periods of high database activity. The value that you choose also depends on the characteristics of the queries. For example, general queries typically use a database agent longer than a more specific query.

Appendix D. National language support

Overview

This chapter contains information about the National Language Support (NLS) provided by OnDemand, including information about the code pages (code sets) supported to provide national language (NL) character support.

Conversion between different code pages

A code page maps each character from a character set, such as the Latin alphabet, to a numeric representation. Each code page is identified by a numeric identifier. For example, code page 850 represents the character A as hexadecimal 41.

Ideally, for optimal performance, OnDemand clients and applications should always use the same code page as the OnDemand instance⁷. However, this is not always practical or possible. OnDemand provides support for character conversion that allows clients, applications, and instances to use different code pages. This means that, while an OnDemand instance must run in a single code page, clients that access the instance can operate in any code page and reports that you store in OnDemand can contain characters encoded in any code page.

However, when you use different code pages, OnDemand may need to convert characters from one code page to a different code page in order to maintain the meaning of the data.

When does character conversion occur?

Character conversion can occur in the following situations:

- When a client is operating in a code page that is different from the code page of the OnDemand instance.

Character data conversion takes place on the server using Unicode code page mapping tables. Unicode code page mapping tables exist for all single- and double-byte languages. For example, a Windows client operating in the Latin 1 code page 1252 can access an OnDemand instance that has character data encoded in the Latin 1 code page 819 (code set ISO 8859-1).

Any data that the user enters (or default values) is converted to Unicode by OnDemand. The resulting Unicode data is then converted to the code page of the instance. For example, the user enters a userid, password, and server name to logon to a server. OnDemand converts the characters from the code page of the client to Unicode and then from Unicode to the code page of the instance.

Any data sent to the client is converted to Unicode by OnDemand. The resulting Unicode data is then converted to the code page of the client. For example, after authenticating the userid and password, the server builds a list of folder names that the user is authorized to open. OnDemand converts the characters from the code page of the instance to Unicode and then from Unicode to the code page of the client.

- When OnDemand index data is in a code page that is different than the code page of the OnDemand instance.

7. An OnDemand instance is a logical OnDemand server environment, with one library (and set of database files). Each OnDemand instance must run in a single code page.

Character data conversion takes place on the server using Unicode code page mapping tables. OnDemand converts the characters from the code page of the index data to Unicode and then from Unicode to the code page of the instance. For example, index data with code page 1252 can be stored in an instance that has character data encoded in code page 37.

Character conversion will **not** occur for:

- Documents stored in OnDemand.

When you store documents in OnDemand, they are stored on the server as a byte stream and no character conversion occurs. For example, if the characters in the document are encoded in code page 500, the characters remain encoded in code page 500 when stored in OnDemand.

When a user retrieves a document from OnDemand, the server sends the document to the client without converting the characters from one code page to the other. For example, a document is stored in OnDemand with characters encoded in code page 500. When the user retrieves the document, it remains encoded in code page 500, although the client may be running in a code page that is different than the instance, such as 1252. However, the client viewing program maps characters in a document from the code page of the server to the code page of the client.

Character mapping

For DBCS AFP data and DBCS and SBCS line data, the OnDemand client automatically converts characters in a document from the code page of the server to the code page of the client using ICONV converters. This method of character mapping works with the Windows client and supports DBCS (for AFP and line data) and SBCS (for line data) languages, including most DBCS User Defined Character (UDC) mappings. The ICONV converters automatically map the user-defined area of a code page to the standard user-defined area of the corresponding ICONV table.

For SBCS AFP data, the OnDemand client automatically converts characters in a document from the code page of the server to the code page of the client using character mapping files supplied with OnDemand.

How does OnDemand determine code page values?

The client code page is determined from the operating environment when the connection to the instance is made. For example, in Windows, the code page is derived from the locale as specified in the Regional Settings under Control Panel.

The instance code page is derived from the value specified at the time the instance is created. The instance is in one and only one code page.

The code page of index data generated by an indexer is determined by the code page specified by the input. When index data is stored in an instance, it is converted from the specified code page to the code page of the instance. (When index data is retrieved from an instance, it is converted from the code page of the instance to the code page of the client.)

The code page of a line data document is derived from the application (View Information page). For all other types of documents, the code page is derived from the data. The server never performs character conversion on documents.

Configuring your system for NLS

When you create an OnDemand instance, you specify the language identifier and locale for the instance. These values determine language settings for your instance. See Chapter 12, “Creating an instance,” on page 69 for more information about creating instances and specifying the language identifier and locale.

Language identifiers

Table 8 lists the language identifiers that OnDemand supports.

Table 8. Supported language identifiers

Language Identifier	Territory
PTB	Brazil
FRC	Canada
CHS	China
DAN	Denmark
ARA	Egypt
FIN	Finland
FRA	France
DEU	Germany
ITA	Italy
JPN	Japan
KOR	Korea
NLD	Netherlands
NOR	Norway
ESP	Spain
SVE	Sweden
CHT	Taiwan
ENU	U.S.A. (English)

Locales

The information in Table 9 can help you determine the language abbreviation and locale to use when creating an OnDemand server instance. An instance can be created in any language, even if that language is not installed for OS/400. The DBCS languages can only be used when you are running on a DBCS-enabled system.

The table does not provide a complete list of all language codes, CCSIDs, countries, or locales. Additional reference information can be obtained from IBM iSeries Globalization on the Web at www.ibm.com/servers/eserver/series/software/globalization/.

Table 9. Language Codes, CCSIDS, and Locales

Language Code	In OD ¹	Language	SBCS CCSID	SBCS CCSID (Euro)	DBCS CCSID	Territory	Locale ²	Locale (Euro) ²
ARA	Y	Arabic	420			Arabic Speaking Egypt Qatar	AR_AA AR_EG* AR_QA*	

Table 9. Language Codes, CCSIDS, and Locales (continued)

Language Code	In OD ¹	Language	SBCS CCSID	SBCS CCSID (Euro)	DBCS CCSID	Territory	Locale ²	Locale (Euro) ²
CHS	Y	Simplified Chinese	836		935 1388	China	ZH_CH_GBK* ZH_CN	
CHT	Y	Traditional Chinese	37		937	Taiwan	ZH_TW	
CSY		Czech	870	1154		Czech Republic	CS_CZ	
DAN	Y	Danish	277	1142		Denmark	DA_DK	DA_DK_E*
DES		Swiss German	500	1148		Switzerland	DE_CH	
DEU	Y	German	273	1141		Austria Germany	DE_AT* DE_DE	DE_AT_E* DE_DE_E
ELL		Greek	875	4971		Greece	EL_GR	EL_GR_E*
ENA		English Australian	37	1140		Australia	EN_AU*	
ENB		Belgium English	500	1148		Belgium	EN_BE	EN_BE_E*
ENG		UK English	285	1146		United Kingdom	EN_GB	EN_GB_E*
ENP		English Upper Case	37	1140			EN_UPPER	
ENU	Y	US English	37	1140		United States	EN_US	
ESP	Y	Spanish	284	1145		Spain Mexico	ES_ES ES_MX*	ES_ES_E
FIN	Y	Finnish	278	1143		Finland	FI_FI	FI_FI_E
FRA	Y	French	297	1147		France	FR_FR	FR_FR_E
FRB		Belgian French	500	1148		Belgium	FR_BE	FR_BE_E
FRC	Y	Canadian French	500	1148		Canada	FR_CA	
FRS		Swiss French	500	1148		Switzerland	FR_CH	
HEB		Hebrew	424			Israel	HE_IL	
HUN		Hungarian	870	1153		Hungary	HU_HU	
ITA	Y	Italian	280	1144		Italy	IT_IT	IT_IT_E
ITS		Swiss Italian	500	1148		Switzerland	IT_CH*	
JPN	Y	Japanese Katakana with lowercase	290	1027	5035 ³	Japan	JA_5035	
KOR	Y	Korean	833		933	South Korea	KO_KR	
NLB		Belgian Dutch	500	1148		Belgium	NL_BE	NL_BE_E*
NLD	Y	Dutch	37	1140		Netherlands	NL_NL	NL_NL_E
NON		Norwegian - Nynorsk	277	1142		Norway	NO_NO	
NOR	Y	Norwegian - Bokmal	277	1142		Norway	NB_NO*	
PLK		Polish	870	1153		Poland	PL_PL	
PTB	Y	Brazilian Portuguese	37	1140		Brazil	PT_BR	
PTG		Portuguese	37	1140		Portugal	PT_PT	PT_PT_E
RUS		Russian	1025	1154		Russia	RU_RU	
SVE	Y	Swedish	278	1143		Sweden	SV_SE	SV_SE_E*
THA		Thai	838	1160		Thailand	TH_TH	
TRK		Turkish	1026	1155		Turkey	TR_TR	

Table 9. Language Codes, CCSIDS, and Locales (continued)

Language Code	In OD ¹	Language	SBCS CCSID	SBCS CCSID (Euro)	DBCS CCSID	Territory	Locale ²	Locale (Euro) ²
Notes:								
1. A Y in this column indicates that OnDemand will issue system log messages in this language. For all other languages OnDemand will issue system log messages in English.								
2. Locale names followed by an asterisk (*) must be created using the CRTLOCALE command. All other locale names listed are shipped in QSYS library. Note that the asterisk is not part of the actual locale name. It is simply used in this table for identification purposes. The following command shows an example of how to create a locale:								
<pre>CRTLOCALE LOCALE('/QSYS.LIB/SV_SE_E.LOCALE') SRCFILE('/QSYS.LIB/QSYSLOCALE.LIB/QLOCALESRC.FILE/SV_SE_E.MBR') CCSID(1143) TEXT('SWEDEN/SWEDISH LOCALE WITH EURO')</pre>								
3. For a Japanese language instance, the instance must be created with CCSID 5035 and locale JA_5035. Using a CCSID/locale combination that only supports uppercase will cause unreadable OnDemand messages in the system log.								

Mapping AFP fonts

You may need to map the fonts your documents use to fonts that can be displayed at the client.

You need to map AFP fonts if your documents contain fonts that are not defined to OnDemand, if your documents contain modified IBM AFP fonts, or if your documents contain AFP fonts that you or others in your organization created. Otherwise, the viewing program may not be able to display the characters properly. (If your documents use fonts that are not defined to OnDemand, then users will see a message when viewing them.) The *IBM Content Manager OnDemand: Windows Client Customization Guide* provides details about mapping AFP fonts.

You map AFP fonts by using the font definition files supplied with OnDemand. The files are stored in the FONT subdirectory under the directory in which OnDemand was installed (by default, \Program Files\IBM\OnDemand32).

OnDemand provides font definition files for the following multi-byte languages:

- Japanese (ICODED.JPN, CSDEF.JPN, CPDEF.JPN)
- Korean (ICODED.KOR, CSDEF.KOR, CPDEF.KOR)
- Simplified Chinese (ICODED.CHS, CSDEF.CHS, CPDEF.CHS)
- Traditional Chinese (ICODED.CHT, CSDEF.CHT, CPDEF.CHT)

For example, for Japanese, if a font is listed in the CSDEF.JPN file, it will be supported in the AFP viewer. The actual font mapping from AFP font to Windows font happens in the [FGID] section of the CSDEF.JPN file (the [FGID] section is at the end of the file). If a font is not listed in the CSDEF.JPN file, you can add the font by editing the CSDEF.JPN file with a standard text editor.

Creating application groups

An application group is a container that holds report data. You store reports and the index data used to retrieve and maintain them in an application group. You define database fields for each application group. The database fields represent categories of information in a report. When you load a report into an application group, you store index information about the report in the database.

When you define database fields, you specify attributes of the fields. Attributes include the field name, type, and length. For character data, the field length must specify the number of bytes required to hold the field data in the database. For

double-byte languages, character string conversion between code pages may result in either an increase or decrease in the length of the string when data is loaded into the database. For example, the client does not use shift-out/shift-in (SOSI) characters to separate SBCS and DBCS, while the server does use these characters. You need to allow for this possible expansion when specifying the length of a string field. The maximum length of a string field in OnDemand is 254 bytes. Verify the length of each database (string) field you define:

- If you use the Report Wizard to generate application groups, the Report Wizard converts strings you select to the code page of the database and displays the number of characters required to hold the string in the database. You can accept the value generated by the Report Wizard or replace it with another value.
- If you use the Add an Application Group command to add application groups, you must calculate the number of bytes required to hold the field in the database and enter the value on the Field Information page.

Creating applications

You typically create an OnDemand application for each type of report or source of data that you plan to store in OnDemand. When you create an application, you specify attributes of the application. The attributes include:

- The data type of the report as it is stored in OnDemand (for example, AFP). The data type determines the viewing program used to display pages of the report.
- The program used to index the report. If you use one of the indexing programs provided with OnDemand, the OnDemand application typically includes the parameters that the indexing program uses to process the report and generate the index data.
- Logical views of report data. Logical views provide different ways to view pages of a line data report.

You can create an application by using the Report Wizard or by using the Add an Application command. You can create indexing information by entering parameters and values directly into the application, specifying the name of a parameter file that contains the information, or using the Graphical Indexer to generate indexing information. You can create logical views by entering values directly into the application or using the sample data window to generate the logical view information.

Data Type

The Data Type of the application identifies the format of the data as it is stored in OnDemand and the viewer that the client calls to display documents stored in the application. If you plan to store line data in OnDemand or create indexing parameters with the graphical indexer (using a line data source file), you must set the Data Type of the application to Line. (**Note:** References to a data type of Line in this discussion also include the data type of SCS, because SCS is processed as Line data. SCS-extended data type is processed differently, and therefore cannot be used.) When you set the Data Type to Line:

- Verify the code page of the data. The code page of the data is typically the code page of the operating system where the data was created. In OnDemand, the default code page for line data is 500 (ISO EBCDIC).
- If the line data contains shift-in and shift-out (SOSI) codes, indicate how OnDemand handles them. Shift-in and shift-out codes indicate when the code points in a record change from single byte to double byte and double byte to single byte. Select from SOSI1, SOSI2, and SOSI3. (These values are specified

| using the OnDemand Administrator Client, using the PR Mode field on the
| View Info tab within an application definition.)

For all other types of data, the code page is encapsulated in the data. For AFP data, it is possible that characters are encoded in more than one code page. The AFP viewer uses mapping files to display single- and double-byte data in the proper code page. You may need to map the AFP fonts a document uses to outline fonts on the PC so that the client properly displays all characters. The *IBM Content Manager OnDemand: Windows Client Customization Guide* provides details about mapping AFP fonts.

Indexing

If you use the Generic Indexer to index the input data, the default code page is 500. If you need to generate index data in some other code page, then you must specify the CODEPAGE: parameter in the parameter file used by the Generic Indexer. See the *IBM Content Manager OnDemand for iSeries Common Server: Indexing Reference* for more information.

If you use the graphical indexer to generate indexing parameters for a report that contains double-byte character data, you may need to change the starting column, ending column, and column width values generated by the graphical indexer. When you open a report with the graphical indexer, the client converts the characters from the host code page to the PC code page. If the report contains double-byte character data, the SOSI codes in the data are not converted. As a result, when you select a string that contains double-byte character data, the starting column of the string may be displayed incorrectly by the graphical indexer. In addition, if the string in the original report (generated on the host) contained SOSI codes, the ending column or column width value will be displayed incorrectly by the graphical indexer. To generate correct indexing parameters for double-byte character data:

- Create the indexing parameters using the graphical indexer
- Using a copy of the original report (generated on the host), determine the actual starting column, ending column, and column width values for the strings you selected in the report (triggers and fields)
- Change the indexing parameters by using the OnDemand administrator client, clicking the Modify button after selecting the Keyboard option within the Parameters Source section of the Indexer Information tab.

Important: After you modify the indexing parameters, you must not reopen the report with the graphical indexer. If you do so, your modifications will be lost.

Logical Views

If you use the sample data window to create logical views for a report that contains double-byte character data, you may need to change the starting column, ending column, and column width values generated by OnDemand. When you open a report in the sample data window, the client converts the characters from the host code page to the PC code page. If the report contains double-byte character data, the SOSI codes in the data are not converted. As a result, when you select a string that contains double-byte character data, the starting column of the string may be displayed incorrectly by the graphical indexer. In addition, if the string in the original report (generated on the host) contained SOSI codes, the ending column or column width value will be displayed incorrectly by the graphical indexer. To generate correct logical views for double-byte character data:

- Create the logical views using the sample data window
- Using a copy of the original report (generated on the host), determine the actual starting column, ending column, and column width values for the strings you selected in the report (fields and the validation string)
- Change the logical view information by replacing the values on the Logical View Fields page

Important: After you modify the logical view information, you must not reopen the report in the sample data window. If you do so, your modifications will be lost.

Running OnDemand programs

Certain OnDemand programs accept input data (parameters and values) from a parameter file. When you work with a double-byte language database, the data in the parameter file must be encoded in the code page of the database. Table 10 lists the DBCS countries or regions that are supported by OnDemand and the code page that you must run under when you create the parameter file. When you need to create a parameter file, we recommend that you work from a 5250 emulator and use option 2=Edit from the WRKLNK screens to edit a stream file. Do not create the parameter file on a PC and then use the FTP program to send it to a server.

Table 10. DBCS code pages

Country or Region	Host Code Page
China	935
Japan	290 or 1027
Korea	933
Taiwan	937

Appendix E. System log and the system log user exit

OnDemand generates messages about the various actions that occur on the system. For example, when a user logs on the system, OnDemand generates a message that contains the date and time, the type of action, the userid, and other information. You can configure OnDemand to save the messages to a system log. In addition, you can configure the system to send the messages to a user-defined program. The user-defined program can process the messages in any way that you want. For example, it could send alerts to administrators, compile statistics, or generate accounting information.

You must do the following to configure the system to send the messages to a user-defined program:

- Enable OnDemand to generate system messages and specify the types of messages generated by selecting the appropriate options in the System Parameters dialog box.
- Enable OnDemand to generate application group messages by selecting the appropriate option in the System Parameters dialog box.
- Specifying the types of application group messages generated by selecting options on the Message Logging page in application groups.
- Enable OnDemand to send messages to a user-defined program by selecting the appropriate options in the System Parameters dialog box.

After you have completed these steps, OnDemand automatically saves the messages in the system log and sends the messages to the user-defined program.

OnDemand sends the messages to a program called ARSLOG. The ARSLOG program supplied by IBM does not perform any functions. However, you can replace the ARSLOG program supplied by IBM with a user-defined program that does specific functions, such as checking the message number and issuing alerts to administrators.

The messages that OnDemand sends to the ARSLOG program contain the parameters listed in Table 11.

Table 11. ARSLOG program parameters (system log user exit program)

Parameter	Purpose	Size	Example
\$1	OnDemand instance	10 char	QUSROND
\$2	Time stamp	20 char	08/13/95 14:24:31
\$3	Log record identifier	10 char	57049
\$4	OnDemand userid	128 char	QONDADM
\$5	User's accounting information	60 char	Z76-001J/999999

Table 11. ARSLOG program parameters (system log user exit program) (continued)

Parameter	Purpose	Size	Example
\$6	Severity: 1 Alert 2 Error 3 Warning 4 Information 5 Debug	1 char	3
\$7	Message number	5 char	31
\$8	Message text	255 char	Failed Login: od400lib1 7.52.365.12
\$9	Document file	variable	The file is stored in the directory that is specified by the ARS_TMP parameter in the ARS.CFG file (see "ARS_TMP parameter" on page 80). The file is deleted immediately after the exit program returns control to OnDemand.

If you create your own ARSLOG program, remember that the ARSLOG program and any programs that it may call run as user QRDARS400. QRDARS400 must have authority to any objects that your ARSLOG program uses.

OnDemand programs are coded in the C language. However, the ARSLOG program can be written in any language. It is your responsibility to develop the ARSLOG program. You must validate the quality and performance of the ARSLOG program and any other programs that it calls.

A sample ARSLOG CL program can be found in source file QSAMPLES2 in libraries QRDARS and QUSRRDARS. Any program source code that you modify should not be placed in QRDARS because that library is replaced during software upgrades. However, QUSRRDARS library is not replaced and can be used for your modified source.

The online Help for the administrative client provides information about enabling OnDemand to generate messages and send them to the ARSLOG program. The online Help also provides information about how to select which application group messages that OnDemand generates and sends to the ARSLOG program.

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 the IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe on 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.

For license inquiries regarding double-byte (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

IBM World Trade Asia Corporation
Licensing
2-31 Roppongi 3-chome, Minato-ku
Tokyo 106, Japan

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.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Corporation
Software Interoperability Coordinator
3605 Highway 52 N
Rochester, MN 55901-7829
U.S.A.

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this information and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement, or any equivalent agreement between us.

Any performance data contained herein was determined in a controlled environment. Therefore, results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

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.

If you are viewing this information softcopy, the photographs and color illustrations may not appear.

Trademarks

Advanced Function Presentation, AFP, DB2, DB2 Universal Database, IBM, iSeries, Operating System/400, OS/390[®], OS/400, Redbooks, and WebSphere are trademarks of International Business Machines Corporation in the United States, other countries, or both.

Adobe, the Adobe logo, Acrobat, and the Acrobat logo are trademarks of Adobe Systems Incorporated, which may be registered in certain jurisdictions.

Intel and Pentium are trademarks of Intel Corporation in the United States, other countries, or both.

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

Lotus is a trademark of Lotus Development Corporation in the United States, other countries, or both.

Windows and Windows NT[®] are trademarks of Microsoft[®] Corporation in the United States, other countries, or both.

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

Index

A

- Acrobat 24
- administrative
 - roles and responsibilities 14
- administrative client
 - hardware requirements 23
 - software requirements 23
- Adobe Acrobat 24
- Adobe PDF viewing software 24
- AFP data
 - indexing 32
 - input 29
- AFP fonts 131
- annotation color search field 45
- annotation text search field 46
- API
 - See* application programming interface
- APIs 15
- application
 - about 39
 - naming 42
 - overview 4, 39
- application group
 - about 36
 - database field types 44
 - field types 44
 - indexing 37
 - naming 42
 - overview 4, 36
 - permissions 41
- application group field 46
- application groups
 - maintaining 83
 - migrating data to archive storage 83
 - scheduling maintenance 83
- application programming interface
 - system log 135
 - user exit programs 135
- application programming interfaces 15
- application programs
 - indexing with 32
- archive storage
 - about 17
 - migrated indexes 53
 - overview 17
 - reports on archive media 52
 - requirements 52
- ARS_AUTOSTART_INSTANCE parameter 79
- ARS_LANGUAGE parameter 79
- ARS_MESSAGE_OF_THE_DAY parameter 80
- ARS_MSGS_LANGUAGE parameter 80
- ARS_NUM_DBSRVR parameter 80, 125
- ARS_PRINT_PATH parameter 80
- ARS_TMP parameter 80
- ARS.CACHE file 70
- ARS.CFG file
 - ARS_AUTOSTART_INSTANCE parameter 79
 - ARS_LANGUAGE parameter 79
 - ARS_MESSAGE_OF_THE_DAY parameter 80
 - ARS_MSGS_LANGUAGE parameter 80
 - ARS_NUM_DBSRVR parameter 80, 125

- ARS.CFG file (*continued*)
 - ARS_PRINT_PATH parameter 80
 - ARS_TMP parameter 80
 - ARSSOCK_RESOLVE_CLIENT_NAME parameter 80
- autostarting instances 79
 - configuring 79
 - creating 70
 - database connections, specifying 80
 - database manager parameters 125
 - language parameter 79
 - message language parameter 80
 - message of the day parameter 80
 - printing parameter 80
 - server parameters 80
 - server print parameter 80
 - spool file archive folders 123
 - starting instances 79
 - temporary space parameter 80
- ARS.DBFS file 70
- ARS.INI file 70
- ARSLOG program 135
- ARSSOCK_RESOLVE_CLIENT_NAME parameter 80
- autostarting instances 79

B

- backing up the database 84, 97
- backup and recovery 55
 - database 55
 - documents 56
 - reports 56
- Bitmap (BMP) 30
- BMP (Bitmap) 30

C

- cache expiration processing 39
- cache migration processing 38
- cache storage file systems
 - maintaining 83
 - scheduling maintenance 83
- checklist 63
- client
 - APIs 15
 - customization 15
 - hardware requirements 23
 - software requirements 23
- client/server system overview 4
- code pages 132
- commands 15
- compression 49
- concepts
 - application 4
 - application group 4
 - archive storage manager 9
 - cache storage manager 9
 - client/server system overview 4
 - data loading 9
 - document 6
 - expiration 10

- concepts (*continued*)
 - folder 5
 - indexing 9
 - indexing methods 6
 - library server 7
 - loading data 9
 - migration 10
 - object server 7
 - server 7
 - storage manager 9
 - system overview 4
- configuration files
 - saving 59
- connecting to the database 80, 125
- conventions
 - names 42

D

- data
 - compression 49
 - formats 29
 - indexing 30
- database
 - backup 84, 97
 - backup and recovery 55
 - connections 80, 125
 - creating 69
 - expiration processing 38
 - field names 42
 - field types 44
 - importing migrated index data 52
 - logs 52
 - maintaining 83
 - migration 95
 - recovery 55
 - recovery log storage requirements 52
 - scheduling maintenance 83
 - storage requirements 47
 - storage space requirements 52
 - transaction log storage requirements 52
- database manager
 - expiration 10
 - migration 10
 - overview 9
- database manager parameters
 - ARS_NUM_DBSRVR 125
- database organization
 - multiple loads per database table 38
 - single load per database table 38
- date fields
 - database field type 45
 - folder field type 45
 - indexing 33
- DB2
 - ARS_NUM_DBSRVR parameter 125
 - connections 80, 125
- DBCS
 - application group fields 131
 - code pages 132
 - database fields 131
 - fonts 131
 - generic indexer 133
 - graphical indexer 133
 - indexing data 133
 - languages, support for 127
 - logical views 133

- DBCS (*continued*)
 - programs 134
 - report wizard 133
 - Shift Out Shift In structured fields 132
 - SOSI structured fields 132
- defining a storage node 89
- defining a storage set 89
- DHCP
 - ARSSOCK_RESOLVE_CLIENT_NAME 80
- disk storage
 - requirements 50
- document
 - about 6
 - defined 1
 - expiration 10
 - indexing 9, 31
 - migration 10
- documents
 - backup and recovery 56

E

- environment variables
 - time zone (TZ) 45
 - TZ (time zone) 45
- expiration
 - overview 10
- expiration processing 38

F

- field names 42
- field types 44
- file formats supported by OnDemand
 - AFP data 29
 - Bitmap (BMP) 30
 - BMP (Bitmap) 30
 - GIF (Graphic Interchange Format) 30
 - Graphic Interchange Format (GIF) 30
 - JFIF (JPEG Format Image File) 30
 - JPEG Format Image File (JFIF) 30
 - line data 29
 - PCX (Picture Exchange Format) 30
 - PDF (Portable Data Format) 29
 - Picture Exchange Format (PCX) 30
 - Portable Data Format (PDF) 29
 - Tagged Image File Format (TIFF) 30
 - TIFF (Tagged Image File Format) 30
 - user-defined data 30
- file systems
 - cache storage 83
- filter fields
 - fields, filter 36
- folder
 - about 36
 - annotation color search field 45
 - annotation text search field 46
 - application group field 46
 - field names 42
 - field types 44
 - naming 42
 - overview 5, 36
 - permissions 41
 - query field types 44
 - search field types 44
 - segment field 46

- folder (*continued*)
 - text search field 46
- folders
 - common server 123
 - migration 123
 - spool file archive 123
- fonts 131

G

- generic data
 - indexing 32
- generic indexer 133
- generic indexer program 33
- GIF (Graphic Interchange Format) 30
- Graphic Interchange Format (GIF) 30
- graphical indexer 133
- groups
 - about 39
 - names 42
 - overview 39

H

- hardware requirements
 - administrative client 23
 - client 23

I

- importing migrated index data
 - storage space requirements 52
- index data
 - migrating 95
 - migration 100
- index fields
 - fields, index 36
- indexes
 - archive storage space requirements 53
- indexing
 - AFP data 32
 - alternatives 32
 - application groups 37
 - application programs 32
 - date fields 33
 - DBCS data 133
 - documents 31
 - generic data 32
 - generic indexer program 33
 - line data 30
 - methods 6
 - other data 32
 - other indexing programs 32
 - overview 9
 - PDF data 32
 - PDF indexer program 32
 - programs 6, 32
 - reports 31
 - requirements 30
 - space requirements 50
 - temporary space 50
- initialization files 81, 83
- input data formats 29
- installation checklist 63
- installation procedures 65

- installing
 - OnDemand software 65
 - software 61
- instances
 - autostarting 79
 - starting 79

J

- JFIF (JPEG Format Image File) 30
- JPEG Format Image File (JFIF) 30

L

- language parameter 79
- line data
 - indexing 30
 - input 29
- loading data
 - overview 9
- locale 79
- log files 15
- log, system
 - API definition 135
 - application group, configuring 91
 - user exit programs 135
- logical views 133
- logon
 - ARSSOCK_RESOLVE_CLIENT_NAME 80
- logon user exit 15
- logs
 - storage space requirements 52

M

- maintaining software 59
- maintenance programs
 - scheduling 83
- management programs
 - overview 10
- Map Coded Font Format 2 structured fields 133
- mapping AFP fonts 131
- MCF2 structured fields 133
- MCF2REF parameter 133
- memory requirements 21
- message language parameter 80
- message of the day parameter 80
- messages
 - API definition 135
 - system log 135
 - user exit programs 135
- migrated indexes
 - storage space requirements 53
- migrating index data 95
- migration
 - application group data 83
 - application group, configuring 95
 - backups 103
 - cache storage to archive storage 83
 - cleanup 120
 - creating Common Server folders from spool file archive
 - report group definitions 109
 - creating Common Server objects from spool file archive
 - definitions 109
 - creating Common Server query restrictions from spool file archive
 - key security definitions 110

- migration (*continued*)
 - creating Common Server users and user groups from spool
 - file archive authorization lists 110
 - definitions 100
 - index data 100
 - limitations 101, 102
 - overview 10
 - phases 104, 111, 116
 - prerequisites 103
 - reports that should not be migration 101, 102
 - spool file archive 99
 - spool file archive folders 123
 - steps 104, 111, 116
 - tasks 104, 111, 116
- migration policies 89
- Monarch
 - integrating with the client 15
- multiple loads per database table 38

N

- names 42
- national language support 79, 80, 127
- NLS 127

O

- OS/400 Indexer
 - overview 9

P

- parameters
 - ARS_AUTOSTART_INSTANCE 79
 - ARS_LANGUAGE 79
 - ARS_MESSAGE_OF_THE_DAY 80
 - ARS_MSGS_LANGUAGE 80
 - ARS_NUM_DBSRV 80, 125
 - ARS_PRINT_PATH 80
 - ARS_TMP 80
 - ARSSOCK_RESOLVE_CLIENT_NAME 80
 - MCF2REF 133
- password length 42
- passwords 42
- PCX (Picture Exchange Format) 30
- PDF (Portable Data Format) 29
- PDF data
 - indexing 32
- PDF indexer program 32
- PDF viewing software 24
- permissions
 - about 40
 - application group 41
 - folder 41
 - overview 40
- Picture Exchange Format (PCX) 30
- policies, migration 89
- Portable Data Format (PDF) 29
- printing software
 - installing 61
 - requirements 22
 - temporary space 80
- processor requirements 21
- PTFs 65

Q

- query field names 42
- query field types 44

R

- recovery 55
 - database 55
- recovery log
 - storage space requirements 52
- reinstalling software 59
- report
 - defined vii
 - expiration 10
 - formats 29
 - indexing 9, 30, 31
 - migration 10, 99
 - migration limitations 101, 102
 - storage on archive media 52
 - storage requirements 47
- report definition
 - common server folder 123
 - migration 99
 - migration of spool file archive 123
- report name
 - common server folder 123
 - migration of spool file archive 123
- report wizard 133
- reports
 - backup and recovery 56
- requirements
 - archive storage 52
 - database 52
 - database storage 52
 - disk storage 50
 - hardware 21
 - importing migrated index data storage space 52
 - indexing reports 50
 - memory 21
 - migrated index storage space 53
 - processor 21
 - recovery log storage requirements 52
 - reports on archive media 52
 - software 21
 - software storage 50
 - storage 47
 - temporary space 50
 - transaction log storage requirements 52
- responsibilities
 - administrative 14
- roles
 - administrative 14

S

- saving configuration files 59
- scheduling maintenance programs 83
- search field names 42
- search field types 44
- segment field 46
- server
 - overview 7
- server commands 15
- server logging 15
- server print software
 - installing 61

- server print software (*continued*)
 - requirements 22
 - temporary space 80
- servers
 - ARS.CFG file 79
- setting the time zone environment variable 45
- Shift Out Shift In structured fields 132
- single load per database table 38
- software
 - configuration files 59
 - installing 61
 - maintaining 59
 - printing software
 - installing 61
 - reinstalling 59
 - saving configuration files 59
 - server print software
 - installing 61
 - upgrading 59
- software requirements
 - administrative client 23
 - client 23
 - disk storage 50
 - printing 22
 - server print 22
 - storage 50
- SOSI structured fields 132
- spool file archive
 - authorization lists, creating Common Server users and groups from 110
 - creating Common Server folders from 109
 - creating Common Server query restrictions from 110
 - folders 123
 - key security definitions, creating Common Server query restrictions from 110
 - migration 99, 123
 - migration phases 104, 111, 116
 - report definitions, creating Common Server objects from 109
 - report group definitions, creating Common Server folders from 109
- starting instances 79
- storage
 - archive storage requirements 52
 - database 47, 52
 - disk space 50
 - disk storage requirements 50
 - estimating 47
 - importing migrated index data 52
 - indexing reports 50
 - migrated indexes 53
 - planning for 47
 - recovery log storage requirements 52
 - report 47
 - reports on archive media 52
 - requirements 47
 - temporary space 50
 - transaction log storage requirements 52
- storage manager
 - expiration 10
 - migration 10
 - overview 9
- storage nodes 89
- storage sets 89
- system initialization 81, 83
- system log
 - API definition 135

- system log (*continued*)
 - application group, configuring 91
 - user exit programs 135
- system log file 15
- system migration
 - application group, configuring 95
- system overview 4

T

- Tagged Image File Format (TIFF) 30
- tasks
 - administrative 14
- temporary space 80
 - storage requirements 50
- text search field 46
- TIFF (Tagged Image File Format) 30
- time zone environment variable 45
- transaction log
 - storage space requirements 52
- transforming data 29
- TZ environment variable 45

U

- upgrading software 59
- user exit programs
 - ARSLOG 135
 - system log 135
- user exits
 - logon 15
 - system log 135
- user-defined data 30
- users
 - about 39
 - names 42
 - overview 39
 - passwords 42

V

- verifying the installation 87

W

- Windows client
 - APIs 15
 - customization 15



Program Number: 5722-RD1

SC27-1158-02

