

Content Manager OnDemand for i
Version 7 Release 2

*Common Server Planning and
Installation Guide*



IBM Confidential IBM Confidential IBM Confidential IBM Confidential

IBM Confidential IBM Confidential IBM Confidential IBM Confidential

Content Manager OnDemand for i
Version 7 Release 2

*Common Server Planning and
Installation Guide*



Notices

Before using this information and the product it supports, read the information in "Notices" on page 121.

Edition notice

This edition applies to version 7, release 2 of IBM Content Manager OnDemand for i (product number 5770-RD1) and to all subsequent releases and modifications until otherwise indicated in new editions.

© Copyright IBM Corporation 2001, 2014.

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

Contents

IBM Navigator for i vii

Accessibility information for Content Manager OnDemand ix

System requirements. xi

Part 1. Introduction 1

About Content Manager OnDemand. 3

Content Manager OnDemand overview 3

System overview 4

Concepts 4

 Application. 4

 Application Group 5

 Folder 5

 Cabinet 6

Indexing methods 7

Documents 7

Server 8

 Request manager 9

 Database manager 10

 Storage manager. 10

 Data indexing and loading 10

 Management programs 11

 Content Manager OnDemand Web Enablement Kit 12

Preparing for Content Manager OnDemand 15

Administrative roles and responsibilities. 16

Application programming interfaces 17

 Client customization 17

 Server APIs 18

 Server logging 18

Archive storage management 21

Archive storage management 21

Content Manager OnDemand storage objects 21

Defining the storage configuration. 21

Part 2. Planning information 23

Reports and other data. 25

Collecting requirements 25

Input data formats 25

Indexing data. 27

 Document indexing. 27

 Page indexing 28

 Indexing data with Content Manager OnDemand programs 29

 Indexing reports using date fields 30

Content Manager OnDemand objects 31

Overview 31

Folders 32

Application groups. 33

Applications 35

Users and groups 36

Content Manager OnDemand user ID relationship to IBM i user profiles 36

Permissions 37

 Overview 37

 Folder permissions 37

 Application group permissions 38

Naming rules. 39

Data types and field types 41

Storage requirements 43

Overview 43

Storage hierarchy 45

Data compression 45

Calculating disk storage requirements 46

 System software 46

 Temporary space for indexing 46

 Disk storage for stored reports 46

 Content Manager OnDemand database storage 47

 Database journal storage 47

 Temporary space for importing index data 48

Calculating archive storage requirements 48

 Report storage space 48

 Migrated index storage space 48

Backup and recovery 51

Backup and recovery overview 51

Server software 51

Server information 51

Content Manager OnDemand database 52

 Database backup 52

 Database journaling 52

 Database recovery 52

Reports. 52

 Disk storage 53

 Archive storage 53

Save your configuration files 55

Content Manager OnDemand files. 55

Windows font files 55

Part 3. Installing software 57

Installation checklist. 59

Installing Content Manager OnDemand server software 61

Obtaining Content Manager OnDemand software 61

Installing Content Manager OnDemand	61
Content Manager OnDemand installation procedures	61
Installing the Content Manager OnDemand end-user client	63
Installing the Content Manager OnDemand administrative client	63

Part 4. Configuring and initializing the system 65

Creating an instance. 67

About instances	67
Configuring an instance	67
Changing an instance configuration	69
Deleting an instance	70
Starting and stopping servers	71
Starting the servers	71
Stopping the servers	72
Server work management	73
Automatically starting instances	74
Connecting to instances	74
Connecting from Content Manager OnDemand clients	74
Working with instances	74

Defining a locale 77

Important consideration regarding locales	77
Steps to take after setting locales	78
Solving problems	78

Configuring the ARS.CFG file 81

ARS_AUTOSTART_INSTANCE parameter	81
ARS_DISABLE_ARSLLOG parameter	81
ARS_LANGUAGE parameter	82
ARS_LDAP_ALLOW_ANONYMOUS parameter.	82
ARS_LDAP_BASE_DN parameter.	82
ARS_LDAP_BIND_ATTRIBUTE parameter	82
ARS_LDAP_BIND_MESSAGES_FILE parameter	82
ARS_LDAP_IGN_USERIDS parameter	82
ARS_LDAP_MAPPED_ATTRIBUTE parameter	83
ARS_LDAP_PORT parameter	83
ARS_LDAP_SERVER parameter	83
ARS_MESSAGE_OF_THE_DAY parameter	83
ARS_MSGS_LANGUAGE parameter	83
ARS_NUM_DBSRVR parameter	83
ARS_PRINT_PATH parameter	83
ARS_SUPPORT_CFSOD parameter	84
ARS_SUPPORT_HOLD parameter.	84
ARS_TMP parameter	84
ARSSOCK_RESOLVE_CLIENT_NAME parameter	84

Configure your system at start up . . . 85

Start the server jobs	85
Start the monitors	85

Scheduling maintenance programs . . 87

Maintaining application group data	87
--	----

Backing up the Content Manager OnDemand database	88
--	----

Part 5. Preparing the system for use 89

Saving Content Manager OnDemand passwords into encrypted files 91

Examples of ARSSTASH commands	91
Syntax of the ARSSTASH command	92

Verifying the installation 93

Defining migration policies and storage sets 95

Configuring the System Log application group 97

Maintaining system log data in archive storage	97
Reducing system log size.	98

Configuring the System Load application group 101

Assigning the System Load application group to a storage set	101
--	-----

Configuring the System Migration application group 103

Assigning the System Migration application group to a storage set.	103
--	-----

Backing up the Content Manager OnDemand data 105

Part 6. Appendixes 107

Appendix A. Database servers 109

Description	109
-----------------------	-----

Appendix B. National language support 111

Overview	111
Conversion between different code pages	111
When does character conversion occur?.	111
Character mapping	112
Configuring your system for NLS	112
Language identifiers, CCSIDs, and locales	113
Mapping AFP fonts	114
Creating application groups	114
Creating applications	115
Data Type	115
Indexing	116
Logical Views	116
Running Content Manager OnDemand programs	116

**Appendix C. System log and the
system log user exit 117**

Notices 121
Trademarks 123

Index 125

IBM Navigator for i

IBM Navigator for i is a powerful graphical interface for managing your IBM® i servers.

IBM Navigator for i functionality includes system navigation, configuration, planning capabilities, and online help to guide you through your tasks. IBM Navigator for i makes operation and administration of the server easier and more productive and is the only user interface to the new, advanced features of the operating system. It also includes Management Central for managing multiple servers from a central system.

Accessibility information for Content Manager OnDemand

For complete information about accessibility features that are supported by this product, see the IBM *Content Manager OnDemand for i: Common Server Administration Guide*.

System requirements

For system requirements, see the following: Hardware and software system requirements

Part 1. Introduction

This section provides an overview of the Content Manager OnDemand system and contains information that can help you better understand how Content Manager OnDemand works.

This section describes how Content Manager OnDemand manages reports and index data, includes important information about how Content Manager OnDemand, the database manager, and the storage manager work to index, load, and retrieve documents, and contains a list of the tasks that Content Manager OnDemand administrators typically perform to manage a Content Manager OnDemand system.

About Content Manager OnDemand

Content Manager OnDemand overview

A Content Manager OnDemand system can support small office environments and large enterprise installations with hundreds of system users.

Content Manager OnDemand can dramatically improve productivity and customer service in many businesses by providing fast access to information stored in the system.

Content Manager OnDemand processes the 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 Content Manager OnDemand, you can archive newly created and frequently accessed data on high speed, disk storage volumes and automatically migrate them to other types of storage volumes as they age.

Content Manager 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 workstation or Content Manager OnDemand server.

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

Content Manager OnDemand provides users the ability to view documents, print, email, and fax copies of documents, and attach electronic notes to documents.

Some of the advantages that Content Manager OnDemand offers include:

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

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

Content Manager 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 data
- 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 Content Manager OnDemand can help you quickly retrieve the specific data that you need to provide fast customer service.

System overview

A Content Manager 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.

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

Content Manager OnDemand provides the capability to do most client functions from almost any operating system, by using a Web browser.

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

Content Manager 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 Content Manager 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, folder, and cabinet represent how Content Manager OnDemand stores, manages, retrieves, views, and prints reports and data, along with the associated index information.

When defining a new report or type of data to Content Manager 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. To help users find folders quickly, administrators can create cabinets.

Application

An application describes the physical characteristics of a report to Content Manager OnDemand. Typically you define an application for each program that produces output that will be stored in Content Manager 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 Content Manager 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 index fields for the data that you load into Content Manager OnDemand. When you load a report into Content Manager OnDemand, you must identify the application group where Content Manager OnDemand will load the index data and store the documents.

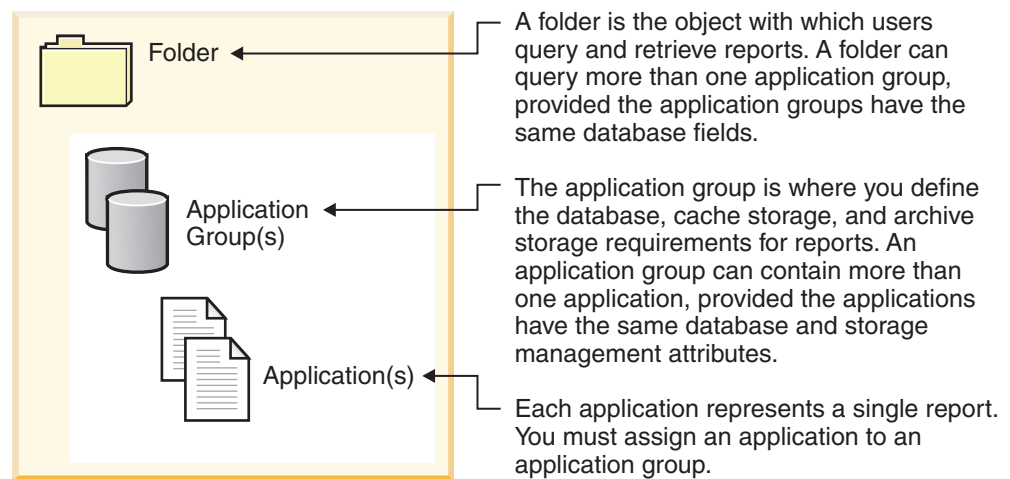
An application group is a collection of one or more Content Manager OnDemand applications with common indexing and storage management attributes. You typically group several different reports in an application group so that users can access the information contained in the reports with a single query.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

Folder

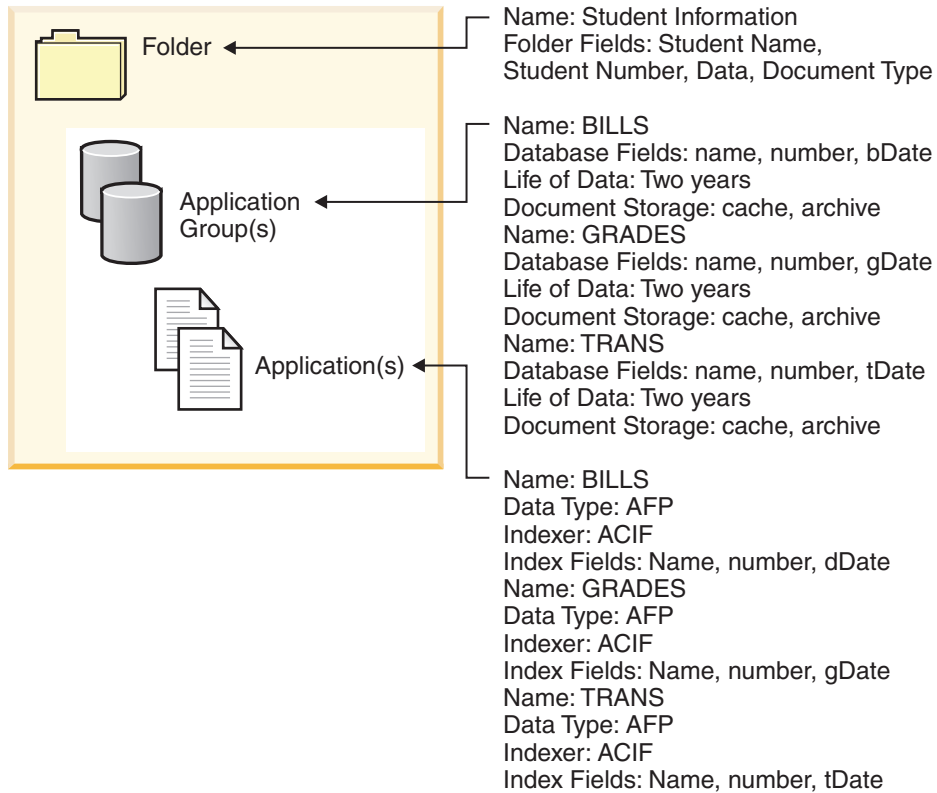
A folder is the user's way to query and retrieve data stored in Content Manager OnDemand. A folder provides users with a convenient way to find related information stored in Content Manager OnDemand, regardless of the source of the information or how the data was prepared.

A folder allows an administrator to set up a common query screen for several application groups that might 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.

The following folders, application groups, and applications (part 1 of 2) graphic illustrates the concepts described in this section.



The following folders, application groups, and applications (part 2 of 2) graphic shows an example.



Cabinet


If users have many folders, they can group their folders into cabinets, which enable users to navigate to folders more easily.

Cabinets follow these rules:

- A cabinet can contain one or more folders.
- A folder can belong to zero or more cabinets.

You can organize the folders that a user needs for generating different types of reports. For example, the user needs to pull together information on fund performance, fund balance, and fund transactions for a monthly report. The "Monthly Report" cabinet contains folders for each type of information that the user needs to collect. The user also needs to generate investments performance reports for clients. Investment performance reports include information on stock performance, bond performance, and fund performance. The "Client Report" cabinet contains folders for stocks, bonds, and funds. Both cabinets contain the folder for fund performance because the user needs information on fund performance to generate both reports.

Optional: You can organize folders in cabinets to enable users to navigate to folders more easily. Each cabinet is a collection of folders.



A folder can belong to more than one cabinet.

Figure 1. Using cabinets to organize folders

Indexing methods

You can index data in two ways; document or report indexing.

Content Manager 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. Content Manager OnDemand supports up to 128 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 Content Manager 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. Content Manager OnDemand stores the report as groups of pages and indexes each group. When reports include a sorted transaction value (for example, invoice number), Content Manager 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

Content Manager OnDemand documents represent indexed groups of pages. Typically a Content Manager OnDemand document is a logical section of a larger report, such as an individual customer statement within a report of thousands of statements.

A Content Manager 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, Content Manager 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.

The diagram below illustrates Content Manager 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 Content Manager 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 Content Manager 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. Content Manager OnDemand retrieves the group that contains the value entered by the user.

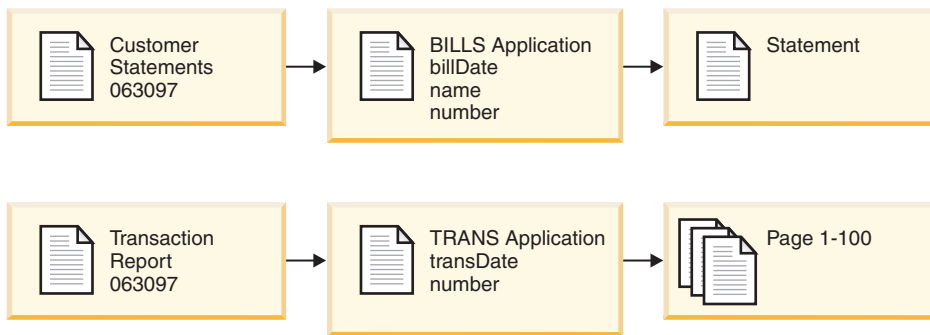


Figure 2. Applications and documents

Server

The Content Manager OnDemand server maintains a central database about the reports stored in Content Manager 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 Content Manager 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). Content Manager OnDemand provides several indexing programs. The OS/400[®] Indexer can be used to index a variety of data types, and is the most common Content Manager OnDemand indexer for spooled files. The Content Manager OnDemand PDF Indexer can be used to create index data for Adobe Acrobat PDF files. The Content Manager OnDemand Generic Indexer can be used to process 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. Content Manager OnDemand requires PSF (a product option of IBM i), if you plan to use the Content Manager OnDemand server print function for AFP documents or if you plan to use server fax for any documents in Content Manager OnDemand.
- Content Manager OnDemand management programs to maintain the Content Manager 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 Content Manager OnDemand.

The following topics provide additional information:

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

Request manager

The request manager processes search requests from Content Manager 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 Content Manager OnDemand client program decompresses and displays the document.

Content Manager 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, Content Manager 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 Content Manager OnDemand programs in the System Log, for example, logon, query, print, and so forth. The messages contain the transaction number, user ID, time stamp, and other information. Administrators can open the System Log folder and view the messages. Content Manager 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

Content Manager OnDemand uses a database manager called DB2 Universal Database™ (DB2®; provided with IBM i), to maintain the index data for the reports that you load into the system.

The database manager also maintains the Content Manager OnDemand system tables that describe the applications, application groups, storage sets, folders, groups, users, holds, cabinets, and printers that you define to the system.

Storage manager

The Content Manager 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.

Content Manager 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 Content Manager 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 Content Manager OnDemand must be indexed. Content Manager 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. Content Manager OnDemand data loading programs read the index data generated by the OS/400 Indexer and load it into the Content Manager OnDemand database. The data loading programs obtain other processing parameters from the Content Manager 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 Content Manager OnDemand administrative client. The administrative client includes a report wizard that lets you create indexing parameters by visually marking up sample report data. Content Manager 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 i: Common Server Indexing Reference* for details about the indexing programs provided with Content Manager OnDemand.

The illustration shows an overview of the data indexing and loading process.

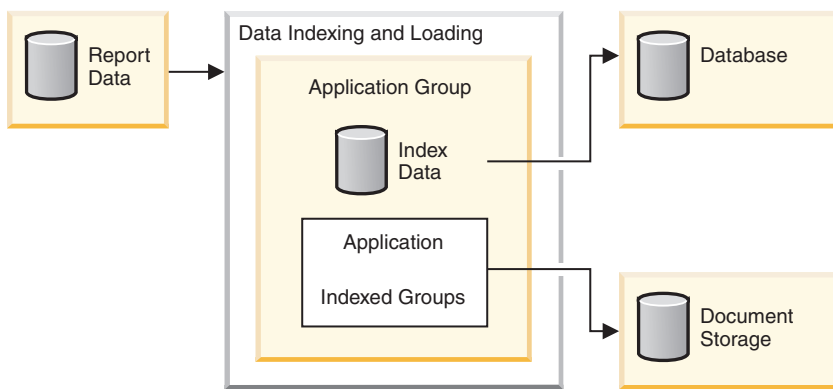


Figure 3. Data preparation, indexing, and loading

The Content Manager 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 Content Manager 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 Content Manager OnDemand database with index data extracted from the report. Depending on the storage management attributes of the application group, the data loading program might 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

Content Manager OnDemand provides a disk storage manager 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 disk storage manager uses 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 Content Manager OnDemand can reclaim the disk space released by expired and migrated data. You should configure your Content Manager 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, if the migration policy contains an Expire level. An administrator defines management information to the archive storage manager to support the Content Manager OnDemand data it manages. The management information includes storage volumes that can contain Content Manager OnDemand data, the number of copies of a report to maintain, and how long to keep data in the archive management system.

Content Manager OnDemand maintains documents on disk and on archive media independently of each other, until the documents expire. Expiration in one storage manager expires data from the other storage manager. Each use their own criteria to determine when to expire and delete data; each use their own commands to remove documents. You need to decide how data will be expired from your system and consistently apply those standards when creating your application groups and migration policies.

Content Manager OnDemand Web Enablement Kit

The Content Manager OnDemand Web Enablement Kit is an optional feature of Content Manager OnDemand that allows people in an organization to use a Web browser to access data stored in a Content Manager 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 Content Manager 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 Content Manager OnDemand interfaces and protocols to access data stored in an OnDemand server. No additional code is needed on the Content Manager 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 IBM i 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 workstation.
- The Image Web Viewer. The Image Web Viewer lets users search, retrieve, view, navigate, and print BMP, GIF, JPEG, PCX, PNG, and TIFF documents from a Web browser. Each person in your organization that plans to use the Image Web Viewer to view documents might need to install it on their workstation. This is required when you want to view the image within the browser window, rather than allowing a default workstation 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 Content Manager OnDemand Advanced Function Presentation Transformations. The IBM Content Manager OnDemand Advanced Function Presentation Transformations 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 Advanced Function Presentation Transformations from IBM and install and configure it on the Web server. See your IBM representative for more information about the IBM Content Manager OnDemand Advanced Function Presentation Transformations.

Preparing for Content Manager OnDemand

This section contains an outline that you might find helpful as you prepare your organization for the Content Manager 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 Content Manager OnDemand will affect their daily work.
- Develop an end-user training course or contract with IBM to provide training for Content Manager 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 Content Manager 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 might need. For example, if you need to convert AFP data to HTML data, you can use the IBM Content Manager OnDemand Advanced Function Presentation Transformations for IBM i Transform, 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 Content Manager OnDemand to support the reports that you plan to store on the system.
- Use the Content Manager 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 Content Manager 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 Content Manager OnDemand system. Update your company's operations and recovery manuals with information required to operate, support, and backup the Content Manager OnDemand system.

Administrative roles and responsibilities

Content Manager OnDemand administrators assume responsibility for and take care of the Content Manager OnDemand system.

The Content Manager OnDemand system includes all sorts of things, including hardware, application and system software, reports, and users.

- Hardware includes the IBM i 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 Content Manager OnDemand applications and decide how Content Manager OnDemand will manage data on the servers.
- Administrators define Content Manager OnDemand groups and users to the system and make sure that the client software is installed and operating properly.

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

Depending on the size of your organization, there might be one person or many people administering the system. If your organization is large, the administrative tasks might be divided among several people. For example, a Content Manager OnDemand system administrator could maintain Content Manager OnDemand storage sets, system printers, holds, groups, and users; a Content Manager 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 a Content Manager OnDemand system. Some of these tasks might be the responsibility of a person other than a Content Manager OnDemand administrator.

- Installing and upgrading equipment
- Installing and maintaining Content Manager 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 Content Manager OnDemand system printers
- Defining reports to the system
- Defining Content Manager 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 Content Manager 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

Content Manager OnDemand provides an administrative client to allow administrators to maintain Content Manager 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 applications, application groups, folders, cabinets, storage sets, holds, system printers, 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.

Content Manager OnDemand provides a set of administrative commands to help administrators maintain the system. For example, Content Manager 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

Content Manager OnDemand provides several kinds of application programming interfaces that you can use to customize Content Manager 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 Monarch¹ 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.

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

Server APIs

Content Manager 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.

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

Server logging

System logging facility

Content Manager OnDemand provides the system logging facility to help an administrator track activity and monitor the system.

Content Manager 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, Content Manager OnDemand records information about your actions in the system log. You can use one of the Content Manager OnDemand client programs to search for and view messages from the system log by time stamp, severity, message number, user ID, and other search criteria.

System log user exit

Content Manager OnDemand provides a user exit that can be used to process the messages that are written to the Content Manager 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 Content Manager 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.

Content Manager OnDemand sends parameters to the system log user exit, such as the name of the Content Manager OnDemand instance, a time stamp, a log record identifier, the user ID that is associated with the action, accounting information for the user ID, 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

Content Manager OnDemand to provide specific information to the system log user exit by setting system and application group parameters with the administrative client.

Archive storage management

Archive storage management

You can configure a Content Manager OnDemand system to maintain copies of reports on disk and in archive storage.

The copies in archive storage are for long-term storage. Content Manager 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 Content Manager OnDemand.
- Server commands that control the server program activities such as data migration and expiration.
- A web-based Content Manager OnDemand component of IBM Navigator for i that you can use to define storage management policies and archive media definitions. The storage management policies determine where data is stored and how long Content Manager OnDemand maintains the data.

Content Manager OnDemand storage objects

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

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

When you load a report into Content Manager 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. Content Manager 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 Content Manager 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.

For more information on migrating documents and recommendations for storage management criteria defined in your application groups, storage sets, and migration policies, see the *IBM Content Manager OnDemand for i: Common Server Administration Guide*.

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.

About this task

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 Content Manager OnDemand. For example, you define an optical library to the IBM i server. Then the Content Manager 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 Content Manager 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. Planning information

This section is a planning source for Content Manager OnDemand administrators.

Other people in an organization interested in this section might 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 Content Manager OnDemand administrators perform to plan for the installation of Content Manager OnDemand and prepare for the operation of Content Manager OnDemand.

Reports and other data

This section contains information that can help you plan for the reports that you will be storing into Content Manager OnDemand.

You can use the information to help determine the hardware configuration that you need to support your Content Manager 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 Content Manager OnDemand, and provide information about indexing reports.

Collecting requirements

Planning for Content Manager 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 Content Manager 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 sorted transaction or ledger data.
 - Document organization: logical groups of information, such as statements or policies.
 - Data that might not have a consistent format, such as reference materials or product literature.
- Will Content Manager 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 Content Manager 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

The types of data listed below, Content Manager OnDemand allows you to store almost any other type of data on the system.

Content Manager 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.
- System Network Architecture (SNA) Character Stream. Input data that is generated from an IBM i printer file with a device type of *SCS.
- The same as SCS, except that the data includes extended print attributes that might be ignored in an SCS data stream. The extended print attributes include variable lines per inch and variable characters per inch.
- Line data with ANSI or machine carriage control characters.
- Workstation files such as text documents, spreadsheets, and presentations, for example.
- Adobe Portable Data Format (PDF) files.
- 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.
 - PNG (Portable Network Graphics). A lossless compression format that can handle multiple colors and large pictures. Offers significant file size savings over GIF, but not as much over JPEG.
 - 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.

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 HTM, then the client could start Internet Explorer to view the document.

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

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. For example, users can retrieve and view customer statements that Content Manager 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 Content Manager OnDemand. The resources include overlays, page segments, and form definitions. 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 Content Manager OnDemand is to index reports.

When you index a report, Content Manager 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, Content Manager OnDemand displays a list of search fields, which represent the database fields. To perform a query, the user enters values in the search fields. Content Manager 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 might produce a large print stream made up of thousands of individual customer bills. With Content Manager OnDemand, you can identify the individual customer bills within the report as smaller, separate information units, or logical items (known as documents in Content Manager 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 Content Manager OnDemand generally fit into one of two categories:

- Document. For reports made up of logical items, such as statements, bills, policies, and invoices.
- Page. For reports that contain sorted values that either increase from the beginning of the report to the end of the report, or decrease from the beginning of the report to the end of the report, 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 document 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, Content Manager OnDemand returns a list of the items that match the user's query and transfers the individual items to the Content Manager OnDemand client program for viewing and printing. Content Manager OnDemand supports up to 128 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. The graphic shows an example of a report file and document indexing.

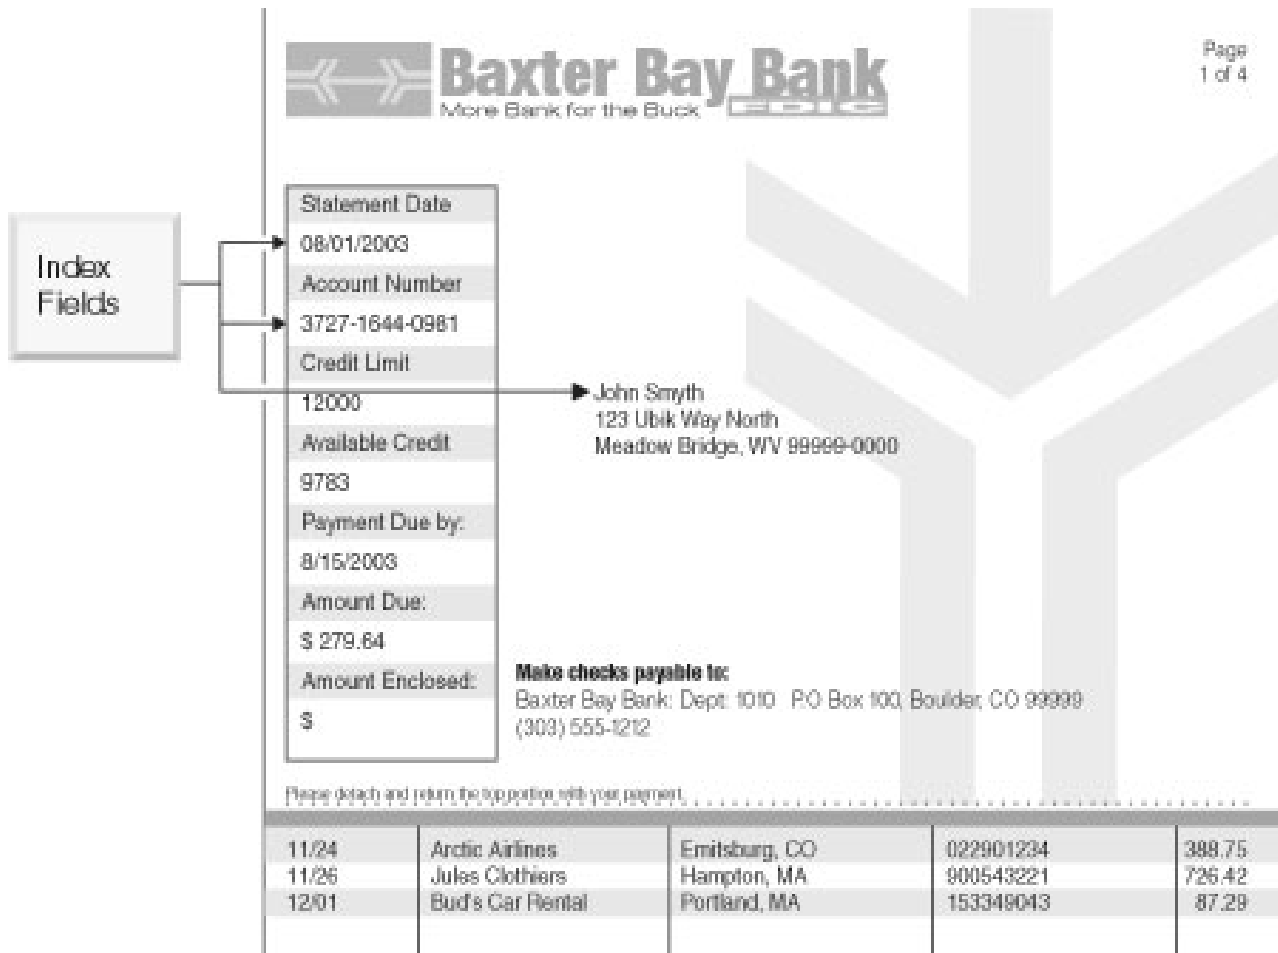


Figure 4. Document indexing method

Page indexing

Page indexing allows users to search sorted report data and retrieve the first occurrence of the value that they specified in the query.

Content Manager 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, Content Manager OnDemand returns a list of the items that match the query. When the user selects an item for viewing, Content Manager OnDemand performs a text search within the item for the value specified by the user. The Content Manager OnDemand client program displays the first page that contains the value specified by the user. Content Manager OnDemand uses a single, unique sorted index value for the retrieval of the report data, for example, an invoice number or a transaction identifier. The graphic shows an example of a report file and page indexing.

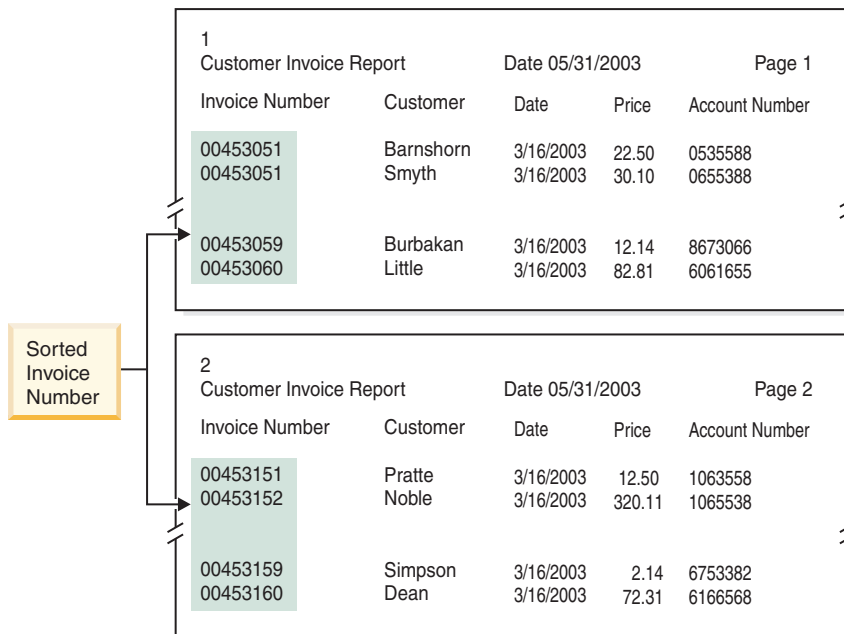


Figure 5. Page indexing method

Indexing data with Content Manager OnDemand programs

Content Manager OnDemand provides a variety of methods to index your print data.

The following topics provide additional information:

- Use the Content Manager 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 Content Manager OnDemand PDF Indexer program to index Adobe PDF files
- Create index data for the Content Manager OnDemand Generic Indexer program

Using the graphical tool

You can use the Content Manager 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 Content Manager OnDemand Administrator client and can be used during your definition of an application.

Generating index data in application programs

You can create index information for AFP data using the application program that generates the report.

The index information is contained in Tag Logical Elements (TLEs) within the data stream. These TLEs can be marked as index values using the graphical indexing tool just like any other data on the print page. Some application programs already provide TLEs in the data stream. If not, you might find it necessary to modify your application programs to add TLEs if you choose to use them.

Indexing PDF input files

The Content Manager 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 Content Manager 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 *IBM Content Manager OnDemand for i: 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 in the Generic Indexer format

Content Manager OnDemand provides the Generic Indexer 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 in the Generic Indexer format 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 Content Manager OnDemand data loading process can find documents within a file. Content Manager 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 i: 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, Content Manager 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 defined as a segment field. In addition, Content Manager 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 Content Manager OnDemand does not contain a date, you can use the date that the report was loaded into the system.

Content Manager 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 Content Manager 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, Content Manager OnDemand adds an application group table structure and other control information to the database.

Content Manager 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 Content Manager OnDemand, such as an application group, Content Manager 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, Content Manager 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 Content Manager OnDemand system open a folder to query and access reports that are stored on the system. A folder is the primary Content Manager OnDemand object that users deal with. A folder provides users the means to search for and retrieve data stored in Content Manager 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 Content Manager OnDemand. The Content Manager 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.

A Content Manager 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 Content Manager 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 Content Manager 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 Content Manager OnDemand, such as application groups, applications, and folders.

Content Manager OnDemand uses properties to describe the appearance, behavior, and internal structure of the objects that make up a Content Manager OnDemand system. For example, Display Format is a property of a folder field that determines

how Content Manager 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. Content Manager 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, Content Manager 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 Content Manager 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.

Content Manager OnDemand maintains information about the name of the folder and its structure in the Content Manager 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 Content Manager OnDemand through properties and values grouped in categories. A category is a set of related properties. Content Manager 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 Content Manager 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, Content Manager 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, Content Manager 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.

Content Manager OnDemand supports up to 128 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. Content Manager 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.

Content Manager OnDemand requires a segment field for each application group that you define. Content Manager 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. Content Manager 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. Content Manager OnDemand can index, store, and retrieve the logical items contained in a report. Each logical item can be indexed on up to 128 values, for example, account number, customer name, and balance. Content Manager OnDemand creates a row in the database for each logical item it finds in the report.

- Other reports might be organized differently, and might not necessarily contain logical items. For example, a report could contain thousands of pages of transaction or general ledger data. Content Manager 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. Content Manager OnDemand divides these types of reports into groups of pages and indexes each group of pages. While these types of reports might 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 Content Manager OnDemand should store the index data for the reports that you load into the application group. Content Manager OnDemand provides a method 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, Content Manager 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. Content Manager 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. Content Manager OnDemand automatically segments the application group data when it grows beyond a certain size. Content Manager 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.

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

Content Manager 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.) Content Manager 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. Content Manager 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, Content Manager 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. Content Manager 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

A Content Manager OnDemand application describes the physical characteristics of a report, processing instructions for the indexing and data loading programs, and information about how Content Manager OnDemand displays and prints pages of a report.

You can specify default settings for viewing and printing pages of a report at the Content Manager 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, Content Manager 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). Content Manager OnDemand uses the information that you specify to properly interpret the data for viewing.

The Content Manager OnDemand application is where you specify information to the indexing and data loading programs, such as the technique that Content Manager OnDemand uses to compress the report file, the parameters used to index the data, and information that Content Manager OnDemand uses to process index data before loading index records into the database. Content Manager 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 Content Manager 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 Content Manager 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 Content Manager OnDemand with a user ID. Content Manager OnDemand authenticates user IDs and determines the usage and administrative authority available to the user based on the log on user ID.

A Content Manager OnDemand user ID does not necessarily have to identify an individual user. However, for accounting and security purposes, you probably want to assign a Content Manager OnDemand user ID to each person that uses the system.

Content Manager OnDemand automatically creates the QONDADM user ID when you initialize the system. The QONDADM user ID 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, cabinets, and printers.

Content Manager OnDemand groups are a means to organize users by function, authorization, holds, or any other purpose you might require. When you define a Content Manager 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.

Content Manager OnDemand groups are not the same as IBM i group profiles, but you might choose to use the same names if that is easier to maintain.

Content Manager OnDemand user ID relationship to IBM i user profiles

When you install Content Manager OnDemand, a default logon security user exit is enabled that forces a relationship between your Content Manager OnDemand user IDs and your IBM i user profiles.

If you use Content Manager OnDemand as installed, then the following is true:

- A Content Manager OnDemand user ID must match an IBM i user profile.
- The password when logging on to Content Manager OnDemand is the same as the password used when logging on to IBM i. If you change the password using Content Manager OnDemand, it actually changes the IBM i password. When you create a Content Manager OnDemand user ID, 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 Content Manager OnDemand might force a user to change their password before it is required by IBM i.
- Minimum Password Length should be set to Allow Blank Password. This prevents Content Manager OnDemand from trying to impose its own rules on the length of a password and allows IBM i to use its own rules.

If the security exit is not enabled, then the Content Manager OnDemand user ID and password have no relation to the IBM i user ID and password and all the Content Manager OnDemand System Parameter settings are honored.

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 the table, 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 i: Common Server Administration Guide* for details.

You can set folder and application group permissions for every user and group defined to Content Manager 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 Content Manager 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 Content Manager 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 Content Manager OnDemand client programs and can modify the folder field information with the administrator interface. Content Manager 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. Content Manager 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 Content Manager 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 Content Manager 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 Content Manager OnDemand displays report file pages. Users can define their own logical views with Content Manager 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 Content Manager OnDemand, you assign names to the various objects.

If you install Content Manager 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 bytes, depending upon whether your Content Manager OnDemand user names are linked to your IBM i user profile names. If you are using Content Manager OnDemand as shipped, your Content Manager OnDemand user names are linked to your IBM i 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 Content Manager OnDemand user names and your IBM i user profile names.
- Cannot include the * (asterisk), % (percentage) + (plus), [(left bracket),] (right bracket), " (double quotation mark)
- Must be unique
- By default, Content Manager OnDemand converts lowercase letters in a user name to uppercase (for example, laguarde is stored as LAGUARDE)

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

When creating a password, the value that you specify:

- Can contain from one to 20 bytes. When creating a password, the value that you specify can be a maximum of 20 bytes. However, the password authentication that is built into Content Manager 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, Content Manager OnDemand converts lowercase letters in a password to uppercase (for example, laguarde is stored as LAGUARDE). 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 bytes
- Cannot include the * (asterisk), ' (apostrophe) % (percentage) + (plus), [(left bracket),] (right bracket), " (double quotation mark)
- Must be unique
- Can be mixed case; however, Content Manager 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 60 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 i: Common Server Administration Guide* for more details about output queue monitors.
- Cannot include the % (percentage), [(left bracket),] (right bracket), or “ (double quotation mark) characters
- Can be mixed case; however, Content Manager 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 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, Content Manager 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 Content Manager OnDemand reserved words:

```

annot      doc_off
comp_len   doc_type
comp_off   prt_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 for i SQL Reference*.)

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

- Can contain from one to 30 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 60 bytes, including embedded blanks
- Can be mixed case; however, Content Manager OnDemand ignores the case (for example, Report Date is the same as report date)
- 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 bytes
- Can be mixed case; however, Content Manager 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 60 bytes
- Can be mixed case; however, Content Manager 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, Content Manager 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, Content Manager 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 Content Manager OnDemand what kind of data can be stored in the field.

When you define a folder to Content Manager 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

See the administrator client help for a list of folder field types.

Storage requirements

Overview

Estimating storage requirements for a Content Manager OnDemand system begins with understanding and documenting end-user requirements for storing and accessing data.

Before you turn requirements into a storage subsystem to support your system, you must also review the various operational and performance issues. For example, Content Manager OnDemand supports up to 128 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 might 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 Content Manager 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 a Content Manager 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 Content Manager OnDemand server software, and any other applications on the IBM i that are required to run your business.
- Storage space for configuration files and control files.
- Storage space for the Content Manager 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 might 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 might 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 might greatly exceed the average report size.
- Temporary storage space for indexing a report on the Content Manager OnDemand server.
- Temporary storage space for loading a report on the Content Manager OnDemand server.

- Storage space for reports stored on disk. This might be zero, for reports that do not require disk storage. However, a very large amount of disk space might be required for reports that must remain on disk for several months or longer.
Content Manager 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. Content Manager 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. For sorted transaction data, the examples and calculations that follow assume that Content Manager OnDemand creates 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 i: Common Server Indexing Reference* provides more information.
- Storage space for the database, which includes Content Manager OnDemand system tables (control information and objects that you define to Content Manager 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, Content Manager OnDemand can divide the report into groups of a fixed number of pages and create one index row for each group of pages. For reports that contain logical items, such as statements, and policies, Content Manager 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 Content Manager OnDemand journal and journal receivers. You should plan for disk space for the Content Manager 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 might decide to remove full storage volumes from a library one year after the last time that Content Manager 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 might 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 Content Manager OnDemand.

For example, Content Manager 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, Content Manager OnDemand can automatically store one copy of the report on disk and another copy of the report on archive storage media. Content Manager 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. Content Manager OnDemand provides a command that you can use to automatically remove expired reports from disk on a regular schedule. After you run expiration processing, Content Manager OnDemand reclaims the space taken by expired documents.

| Content Manager OnDemand maintains documents on disk and on archive media
| independently of each other until the documents expire. Each use their own
| criteria to determine when to expire and delete data; each use their own
| commands to remove documents. You need to decide how data will be expired
| from your system and consistently apply those standards when creating your
| application groups and migration policies.

Data compression

Content Manager OnDemand can compress report data using several different data compression algorithms, before storing the data on disk and archive storage.

The compression ratio that Content Manager OnDemand can achieve has a significant impact on the amount of space required to store reports.

The compression ratios that Content Manager 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 i: Common Server Administration Guide* for more information about the **ADDRPTOND** command.

Calculating disk storage requirements

System software

Content Manager 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 Content Manager OnDemand server software.

Temporary space for indexing

Content Manager OnDemand requires temporary storage space on disk to index reports.

The temporary space required by Content Manager 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:

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

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 varies greatly based on requirements such as the number of reports that you store on the system, the compression ratio that Content Manager 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.

Reports are not automatically deleted unless the program that deletes reports is set up to run on a scheduled basis.

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

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

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:

$$\text{Disk space required} = 2 \text{ GB} * 12 * .33 * 1.1 = 8.71 \text{ GB}$$

Content Manager OnDemand database storage

When you load a report into the system, Content Manager 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, Content Manager OnDemand can create one database row for every item found in the report. For reports that contain sorted transaction data, Content Manager 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 Content Manager 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 Content Manager OnDemand.)

There are four major factors that determine the amount of disk space required for the Content Manager 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 Content Manager OnDemand maintains the index data in the database

See the administrator online help for a list of the types of index fields supported by Content Manager OnDemand and the number of bytes required to store a value in each type of index field.

Database journal storage

The Content Manager 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 create a Content Manager OnDemand instance, a journal called QSQRN is created in your instance library. A journal receiver called QSQRN0001 is automatically created as the first receiver for the journal. QSQRN 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

Content Manager 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.

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.

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:

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

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 Content Manager OnDemand can achieve a compression ratio of 3:1 (0.33), you would require approximately 244 GB of archive storage space:

$$\text{Optical space} = (8 \text{ GB} * 84) * 0.33 * 1.1 = 244 \text{ GB}$$

Backup report storage space

The Content Manager 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 Content Manager OnDemand to maintain a backup copy of your reports, double the archive storage space that you calculated.

Migrated index storage space

Content Manager 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.

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

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 Content Manager OnDemand migrates them to archive storage.

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

$$\text{Archive media database space} = (\text{Database size per month} * \text{compression ratio}) * (\text{life of data} - \text{months before migrating data})$$

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:

$$\text{Archive media database space} = (202 \text{ MB} * .33) * (84 - 24) = 4 \text{ GB}$$

Backup and recovery

Backup and recovery overview

This section of the book describes backup and recovery for Content Manager OnDemand.

Content Manager OnDemand has recommendations about methods and procedures that an administrator can use to make sure that the following critical Content Manager OnDemand components can be recovered when needed:

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

See Chapter 5 of the *Content Manager OnDemand for i: 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 Content Manager 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 Content Manager 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 Content Manager OnDemand, you specified information that customized Content Manager 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 Content Manager OnDemand. In addition, if you periodically make changes to the Content Manager 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.

Content Manager OnDemand database

Database backup

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.

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

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 Content Manager 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 Content Manager OnDemand database can be recovered to any point in time.

Database recovery

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.

There are two types of database recovery. The first type recovers from failures that occur while update transactions are taking place.

The second type of recovery deals with corruption of the Content Manager 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 Content Manager 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 journals. These steps will recreate the Content Manager OnDemand database before the catastrophic failure.

Reports

Content Manager 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 disk, optical, or tape storage volumes, or Tivoli® Storage Manager servers 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.

Content Manager 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 Content Manager OnDemand to use archive storage. You configure Content Manager 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 Content Manager 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 Content Manager OnDemand migration policy identifies where the primary copy of a report is maintained. Content Manager 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.

Save your configuration files

When you install software on a Content Manager 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 Content Manager 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 i: Common Server Administration Guide* for specific libraries and objects to be backed up on a regular basis.

Content Manager OnDemand files

Save a copy of the Content Manager OnDemand configuration files listed in the table.

Replace *instance* with the actual name of the *instance* to which the particular file belongs.

Table 2. Content Manager OnDemand configuration files to save

File	Default Location	Purpose
ars.cache	/QIBM/UserData/OnDemand/ <i>instance</i> /config	Define disk storage details.
ars.cfg	/QIBM/UserData/OnDemand/ <i>instance</i> /config	Content Manager OnDemand server configuration file.
ars.ini	/QIBM/UserData/OnDemand/config	Configure OnDemand instances.
ars.dbfs	/QIBM/UserData/OnDemand/ <i>instance</i> /config	Not used, but file must exist.
arslog	/QIBM/ProdData/OnDemand/bin	The symbolic link to the default System Log user exit program.

Windows font files

If you have modified any font-related files on your Windows client workstations, save a copy of the files listed in the table.

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 3. 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.
ALIAS.FNT	\Program Files\IBM\OnDemand32\Font	The Alias file. Used to map AFP fonts to Truetype fonts.

Part 3. Installing software

This section contains instructions for installing the Content Manager OnDemand server software on the IBM i server.

Installation checklist

Setting up your Content Manager OnDemand system typically requires that you do the following steps.

Procedure

1. Contact the IBM Support Center for the latest PTFs for Content Manager OnDemand, or obtain the list from the Content Manager OnDemand support web page (under the Downloads heading). The list of current PTFs can be found in Information APAR number II14723. You should consider requesting current PTFs for IBM i (DB2 in particular) and IBM i Access, because these two products are an integral part of the Content Manager OnDemand system. After you obtain the current list of PTFs, load and apply them on your system.
2. Obtain a copy of the latest Content Manager OnDemand Read This First document from the web at the support web page (under the Featured links heading). Print and read the entire document before you begin.
3. Check the Content Manager OnDemand prerequisites and verify the required and optional hardware and software products.
4. If you are upgrading to a new version of Content Manager OnDemand, save the configuration files used by the system (see “Save your configuration files” on page 55). Also note any changes you have made to the QRDARS400 user profile. An upgrade will overlay this Content Manager OnDemand user profile and you must make the changes again after your upgrade has completed.
5. Obtain the Content Manager OnDemand software (see “Installing Content Manager OnDemand server software” on page 61).
6. Install the Content Manager OnDemand software on the IBM i server (see “Installing Content Manager OnDemand server software” on page 61).
7. Configure the Content Manager OnDemand software (see Part 4, “Configuring and initializing the system,” on page 65). This step includes the following:
 - a. Redo any changes to the QRDARS400 user profile that you noted in item 4, above
 - b. Create an instance (see “Creating an instance” on page 67)
 - c. Define a locale (see “Defining a locale” on page 77)
 - d. Configure the ARS.CFG file (see “Configuring the ARS.CFG file” on page 81)
 - e. Configure system startup (see “Configure your system at start up” on page 85)
 - f. Define schedules for maintenance programs (see “Scheduling maintenance programs” on page 87)
8. Create and initialize the database on the server (see “Creating an instance” on page 67)

What to do next

Verify the installation of Content Manager OnDemand:

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

Prepare the system for use:

1. Define migration policies and storage sets. 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. 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. 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. After configuring the system, we recommend that you create a full backup image of the Content Manager OnDemand database.

Installing and configuring optional software:

1. If you plan to reprint AFP documents or fax any documents using the Content Manager OnDemand server print function, then you must install PSF for IBM i on the Content Manager OnDemand for i server. Complete the following tasks:
 - a. Install and configure PSF for IBM i (part of IBM i).
 - b. Define a server printer on the Content Manager 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.

Installing Content Manager OnDemand server software

You must install a copy of the Content Manager OnDemand server software on your IBM i server.

Obtaining Content Manager OnDemand software

Content Manager OnDemand software is no longer sent on the stacked media (CDs) you receive with many of your other licensed programs.

About this task

The product is now available from the Passport Advantage web site, from which it can be downloaded. IBM recommends that you print and review the PDF instruction file provided with the information rather than the abbreviated instructions available from the Passport Advantage® site as a TXT file.

Installing Content Manager OnDemand

Existing Spool File Archive implementations must be migrated from the Spool File Archive environment to Common Server before the system on which they are running is upgraded to 6.1, 7.1, or 7.2.

About this task

Content Manager OnDemand releases 5.3 and 5.4 included the Common Server environment, as well as the legacy environments of Spool File Archive, AnyStore, Record Archive, and Object Archive. All of these environments were fully supported through and including 5.4. As stated in IBM Announcement Letter #206-030 dated February 14, 2006, 5.4 was the last release that Spool File Archive, AnyStore, Record Archive and Object Archive would be shipped and supported. Beginning with Content Manager OnDemand 5.3, a Spool File Archive migration utility has been available as part of the Content Manager OnDemand licensed program product, providing capability to migrate report definitions and indexes from the legacy Spool File Archive environment to the Common Server environment. Existing Spool File Archive implementations must be migrated from the Spool File Archive environment to Common Server before the system on which they are running is upgraded to 6.1, 7.1, or 7.2.

Content Manager OnDemand installation procedures

If you are upgrading from a previous release of Content Manager OnDemand, you must be running Content Manager OnDemand server version 8.4.0.3 or higher prior to upgrading to Content Manager OnDemand for i Version 7 Release 2.

About this task

You can determine your current server version using either one of the following methods:

- Log on to either the Content Manager OnDemand Administrator or OnDemand end-user client. Once logged on, the server version is shown on the message line in the lower right portion of the panel.

- Log on to your IBM i server. Start qshell by using the QSH command. Once qshell is started, enter ARSDATE (with no additional parameters) on the qshell command line. The server version is displayed along with other information about the ARSDATE program.

The licensed program number for Content Manager OnDemand for i Version 7 Release 2 is 5770RD1. The individual Content Manager OnDemand features (which install separately) are as follows:

- ***BASE** for the Base Support (required for all others)
*BASE includes Common Server (previously option 10) and Content Manager OnDemand Web Enablement Kit; previously option 11). This option also includes the IBM Navigator for i - Content Manager OnDemand component (previously named IBM Navigator for i - Content Manager OnDemand Administration).
- **12** for PDF Indexer (optional)
- **13** for Advanced Function Presentation Transformations for IBM i (optional)
- **14** for Enhanced Retention Management (optional)

Full text indexing support is also available for use with Content Manager OnDemand for i Version 7 Release 2.

Product options 10 and 11 are now included in *BASE and no longer require installation of a separate option. Important: If product options 10 or 11 exist on your system, they are deleted during the upgrade process.

The System i Navigator - Content Manager OnDemand component is replaced by the web-based IBM Navigator for i - Content Manager OnDemand component that is included in the *BASE product option. Important: The System i Navigator - Content Manager OnDemand component is not supported at 7.2 and is deleted during the upgrade process.

Product options 1 through 5 do not apply to this publication, and are no longer available. They have not been supported since Content Manager OnDemand for iSeries 5.4 (licensed program number 5722RD1). Important: If options 1 through 5 of licensed program number 5722RD1 exist on your system during the upgrade to Content Manager OnDemand for i 7.2, they are deleted during the upgrade process.

The options are:

- 1 for Spool File Archive
- 2 for Object Archive
- 3 for Record Archive
- 4 for AnyStore
- 5 for Server (used with Spool File Archive)

The recommended way to install Content Manager OnDemand is to use the installation licensed programs menu option from the Work with Licensed Programs menu (GO LICPGM). From the installation 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 12, 13, or 14 as needed).

Important: The first time you start the Content Manager OnDemand servers after a server upgrade or licensed program release upgrade, you should start each

| instance separately. After the first server starts, test the storing and viewing of data
| to ensure that no problems exist. Then start a second server, and perform the same
| tests. If you have a test instance and a production instance defined, you might
| consider working with the test instance first.

Installing the Content Manager OnDemand end-user client

For information on installing the Content Manager OnDemand client for Windows, see Content Manager OnDemand: Client Installation Guide.

Installing the Content Manager OnDemand administrative client

The Content Manager OnDemand Administrator client must be at the same version or higher as the Content Manager OnDemand server. For version 7.2 of Content Manager OnDemand for i, the Content Manager OnDemand Administrator client must be at version 9.0.0.3 or higher.

About this task

For information on installing the Content Manager OnDemand administrative client for Windows, see Content Manager OnDemand for i: Common Server Administration Guide.

Part 4. Configuring and initializing the system

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

Creating an instance

About instances

An instance is defined in the ARS.INI file by naming the instance (which identifies the name of the library used by the instance).

A Content Manager OnDemand instance is a logical server environment with its own library containing a unique set of database files. 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 Content Manager 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. Renaming an instance is not supported. If you need to change an instance name, you must create a new instance, export or recreate your definitions such as applications, application groups, folders, users, etc. and reload any data that you need into the new instance.

Configuring 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 might need to change your user profile, then sign off and back on before creating the instance

About this task

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 does 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.

Procedure

To create your Content Manager OnDemand instance:

1. Name the instance or use the default instance name of **QUSROND**. The instance name must be a valid library name for IBM i. The instance name must

start with an alphabetic character or @ followed by any of these characters: 0-9, A-Z, @, #, or underscore (_). Ensure 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**), and must not be named CONFIG or WWW. This instance name is referred to as *[instance]* in the rest of these instructions.

2. Create the instance by using the Create Instance for OnDemand (**CRTINSTOND**) command. At a minimum, you must specify the name of the instance (which then uses system values and defaults for the additional parameters, such as ***DFT** for the PORT parameter, which uses port 1445). You can specify additional parameters in order to customize the instance to meet your requirements. For example, you can specify a three-character language identifier (using the **LANGID** parameter) which must match one of the language identifiers. If you specify the **LOCALE** parameter, the one that you specify should be included in the list of valid locales. If the instance should be located in a user Auxiliary Storage Pool (**ASP**), then the user **ASP** number (2 through 32) must be specified for the **ASP** parameter and ***ASP** must be specified for the **ASPDEV** parameter. If the instance should be located in an Independent Auxiliary Storage Pool (**IASP**), then ***ASPDEV** must be specified for the **ASP** parameter and the **IASP** name (such as **IASP2**) must be specified for the **ASPDEV** parameter.

Important: If you create an instance in an IASP, that instance cannot have a disk pool in the system ASP or a user ASP.

For example, the following command would create an instance called **ONDTEST** with a server language of US English, using TCP/IP port 1445 :

```
CRTINSTOND INSTANCE(ONDTEST) LANGID(ENU) LOCALE('/QSYS.LIB/EN_US.LOCALE')
```

The **CRTINSTOND** command:

- Creates the /CONFIG directory under /QIBM/UserData/OnDemand and the default and model files within this directory (if they do not already exist).
- 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.
- Creates the instance directory: /QIBM/UserData/OnDemand/*[instance]*. If the instance is located in an Independent **ASP**, the instance directory path is preceded by the Independent **ASP** name. For example, if the Independent **ASP** name is **IASP**, the instance directory is created in directory /IASP/QIBM/UserData/OnDemand.
- 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.) If the instance is located in an Independent **ASP**, the instance directory path is preceded by the Independent **ASP** name. For example, if the Independent **ASP** name is **IASP**, the ARS.CFG, ARS.CACHE and ARS.DBFS files are created in directory /IASP/QIBM/UserData/OnDemand/*[instance]*.
- Creates the library and database tables for the instance. If the instance is located in an **IASP**, you must set the **ASP** Group before you can work with files in that library. Use the Set ASP Group (**SETASPGRP**) command to set the **ASP** Group.

- Creates the directories needed for the instance as specified in the ARS.CFG and ARS.CACHE files. And creates a user profile with the same name as the instance, and adds that user to the instance as a Content Manager OnDemand System Administrator.
- Creates an authorization list with the same name as the instance. If the instance is located in an Independent ASP, a record is added to file QARLCASP in library QUSRRDARS.

Changing an instance configuration

You might need to change some of the configuration parameters from the values you specified when you ran the CRTINSTOND command. These parameters exist in either the ARS.INI file or the ARS.CFG file.

To change some of the configuration parameters in the ARS.INI file located in the /QIBM/UserData/OnDemand/config directory, use the Edit File (EDTF) command:
EDTF '/QIBM/UserData/OnDemand/CONFIG/ARS.INI'

Each 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].

Review the following lines:

PORT=0

The port to which the server listens to receive requests from a Content Manager 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. To run multiple instances concurrently, you must specify an unused port on your system. You can run the Work with TCP/IP Network Status (WRKTCPS) OPTION(*CNN) command to see what ports are currently in use on your system.

SRVR_FLAGS_SECURITY_EXIT=1

Specifies that you want to use IBM i user IDs and passwords as the Content Manager OnDemand user IDs 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 Content Manager OnDemand users do not need to have IBM i user IDs then you should specify a value of 0 for this parameter. When you do this, your Content Manager OnDemand passwords have no relationship to IBM i passwords. However, if a Content Manager OnDemand user ID and an IBM i user profile match, some Content Manager OnDemand commands and APIs use the IBM i user profile as the Content Manager OnDemand user ID, even if you have chosen not to relate the two. This situation could permit IBM i users to perform Content Manager OnDemand functions that you did not intend for them to perform. A Content Manager OnDemand user ID should not match an IBM i 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 Content Manager OnDemand System Parameters values (defined using the Content Manager OnDemand administrative client) for the instance that you have changed.

HOST=LOCALHOST

If you are enabling IPv6 on your IBM i system and need some of your Content Manager OnDemand instances to use IPv4 addressing and others to use IPv6 addressing, you might need to change HOST=LOCALHOST to

HOST=IPV6-LOCALHOST within the ARS.INI stanza for each instance that you want to use IPv6 addresses. You might want some instances to run with IPv6 and others with IPv4. This mixed environment is fully supported. Also, during the transition from IPv4 to IPv6, Content Manager OnDemand clients that use IPv4 addresses can connect to the server simultaneously with clients that use IPv6 addresses.

To change some of the configuration parameters in the ARS.CFG file located in the /QIBM/UserData/OnDemand/instancename directory (where *instancename* is the name of the instance you wish to change): Use the Edit File (EDTF) command: EDTF '/QIBM/UserData/OnDemand/MYINSTANCE/ARS.CFG'

You can change these values:

ARS_DISABLE_ARSLLOG=1

Specifies that the ARSLOG user exit is to be disabled. Set this value to 1 to disable the exit; set this value to 0 to enable the exit.

ARS_LANGUAGE=ENU

Specifies the language in which this instance runs. The example shows the use of the English language.

ARS_MSGS_LANGUAGE=ENU

Specifies the language that is used for server messages. The example shows the use of the English language.

ARS_AUTOSTART_INSTANCE=1

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.

Do not modify any of the other values in these instance definition files without first consulting with Content Manager OnDemand Support. You must end and restart the instance server after you make any changes.

Deleting an instance

You can delete an instance. If the server for the instance is already started, you must end the server before deleting the instance.

Before you begin

You can end a specific server by following the instructions to stop a server.

Procedure

To delete an instance:

1. Delete the instance library. Issue the following command: DLTLIB LIB([*instance*]). When you issue this command, you might receive a message indicating that a journal receiver has not been fully saved. You can respond with an I to ignore this message.
2. 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. You can use the **WRKLNK** command to do a recursive delete of all the objects in a directory.

- a. You can do this by issuing the following command: **WRKLNK** '/QIBM/UserData/OnDemand/[instance]' where [instance] is your instance name.
 - b. Specify **option 2** to edit the instance directory you want to delete.
 - c. Then specify **option 9** for every directory listed in that instance directory.
 - d. Select **F12** to go back to the previous screen and then delete the instance directory itself.
3. Delete the instance definition within file **ARS.INI** in directory /QIBM/UserData/OnDemand/config. If the instance has been used with the Content Manager OnDemand Web Enablement Kit (ODWEK), remove any reference to the instance from the **arswww.ini** file in the directory /QIBM/UserData/OnDemand/www. If you are using multilingual CGI support, the directory name will be /QIBM/UserData/OnDemand/www/[ccsid], where *ccsid* is the required CCSID. If you are using WebSphere Application Server V6, the directory name will be /QIBM/UserData/WebSphere.
 4. 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.
 5. If you have created an instance in an Independent ASP, contact IBM software support to discuss steps to delete the instance record from **QUSRRDARS**.
 6. You should delete 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. Use the Work with Objects by Owner (**WRKOBJOWN**) command, specifying the instance name for the user profile (**USRPRF**) parameter to review objects owned by the instance user profile. Either delete each object or change the owner to a different user profile.

Starting and stopping servers

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

Starting the servers

About this task

Servers are started using the **STRTCPSVR *ONDMD** command. The **INSTANCE** parameter of the **STRTCPSVR *ONDMD** command supports the special values of ***DFT**, ***ALL**, and ***AUTOSTART**, as well as the specification of the name of an instance. (An instance is set to autostart if the **ARS.CFG** file for that instance contains **ARS_AUTOSTART_INSTANCE=1**.) The default value for the **INSTANCE** parameter is ***DFT**. You can also create a data area named **STRTCPSVR** to further control the behavior of the **STRTCPSVR** command. See the *Content Manager OnDemand for i: Administration Guide* for more details about the data area.

Without the **STRTCPSVR** data area, the values of ***DFT** and ***AUTOSTART** work identically. All instances that are set to autostart are started. Use of the special value ***ALL** starts all instances configured on the system. You can also specify the name of a single instance to start, for example: **STRTCPSVR SERVER(*ONDMD) INSTANCE(ONDTEST)**.

With the data area, the value of ***DFT** starts only the instance named in the data area. The data area must be named **STRTCPSVR** and located in library **QUSRRDARS**. The data area should be type character with a length of 10. To create the data area, use the command:

```
CRDTAARA DTAARA(QUSRRDARS/STRTCPSVR) TYPE(*CHAR) LEN(10) VALUE(QUSROND) TEXT('Autostart instance name')
```

where QUSROND is the name of the instance to start.

The special values ***ALL** and ***AUTOSTART** work the same with the data area as without the data area.

To determine the instances that are started when **STRTCPSVR SERVER(*ONDMMD) INSTANCE(*AUTOSTART)** is executed, you can look for the ARS_AUTOSTART_INSTANCE=1 in the ARS.CFG file. However, there is an easier way, so that you do not have to check the ARS.CFG file for every instance.

Use the
grep

command in qshell to search the contents of all the ARS.CFG files for the string ARS_AUTOSTART_INSTANCE=1. For example:

```
$
grep -n 'ARS_AUTOSTART_INSTANCE=1' /qibm/userdata/ondemand/*/ars.cfg
/qibm/userdata/ondemand/ONDDEMO/ars.cfg:53:ARS_AUTOSTART_INSTANCE=1
/qibm/userdata/ondemand/ONDDEU/ars.cfg:53:ARS_AUTOSTART_INSTANCE=1
/qibm/userdata/ondemand/ONDENU/ars.cfg:53:ARS_AUTOSTART_INSTANCE=1
/qibm/userdata/ondemand/QUSROND/ars.cfg:53:ARS_AUTOSTART_INSTANCE=1
$
```

From the output of this command, you can determine that instances ONDDDEMO, ONDDDEU, ONDDENU, and QUSROND is started when **STRTCPSVR SERVER(*ONDMMD) INSTANCE(*AUTOSTART)** is run.

The following table summarizes the behavior of the **STRTCPSVR** command with and without the **STRTCPSVR** data area.

Table 4. STRTCPSVR command with and without STRTCPSVR data area

	*DFT	*ALL	*AUTOSTART	Named instance
Without the data area, running STRTCPSVR starts:	All instances set to autostart	All instances configured on the system	All instances set to autostart	The named instance
With the data area, running STRTCPSVR starts:	Only the instance named in the data area	All instances configured on the system	All instances set to autostart	The named instance

Stopping the servers

About this task

Servers are stopped using the **ENDTCPSVR *ONDMMD** command. The instance parameter of the **STRTCPSVR *ONDMMD** command supports the special values of ***DFT** and ***ALL**, as well as the specification of the name of an instance. The default value for the **INSTANCE** parameter is ***DFT**. You also have the option of creating a data area named **STRTCPSVR** to further control the behavior of the **ENDTCPSVR** command. Create the data area as described above. See the Content Manager OnDemand for i: Common Server Administration Guide for more details about the

data area. Note that, even though the data area is named STRTCPSVR, it controls both the STRTCPSVR and ENDTCPSVR commands. This is by design, so that *DFT starts and ends the same instance.

Without the STRTCPSVR data area, the values of *DFT and *ALL work identically. All instances that are active are ended. You can also specify the name of a single instance to end, for example: ENDTCPSVR SERVER(*ONDMD) INSTANCE(ONDTEST).

With the data area, the value of *DFT ends only the instance named in the data area. The data area must be named STRTCPSVR and located in library QUSRRDARS.

The following table summarizes the behavior of the ENDTCPSVR command with and without the STRTCPSVR data area.

Table 5. ENDTCPSVR command with and without STRTCPSVR data area

	*DFT	*ALL	Named instance
Without the data area, ENDTCPSVR ends:	All active instances	All active instances	The named instance
With the data area, ENDTCPSVR ends:	Only the instance named in the data area	All active instances	The named instance

Server work management

Server jobs are started using a job description by the name of the instance (which must be found in the instance library).

If a job description by that name is not found in the instance library, 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 want to change the job queue that instance TEST uses, you create a job description called TEST in the TEST library that specifies the job queue you want to use.

To change the run priority of the server jobs, you must add a routing entry to the subsystem. The server job is always submitted with routing data QRLMSERVER. To change the run priority of all server jobs for all instances to 40, you could add the following routing entry to subsystem QSYSWRK. (You must choose a sequence number (SEQNBR) that is not already in use.)

```
ADDRTGE SBSD(QSYSWRK) SEQNBR(1841) CMPVAL(QRLMSERVER) PGM(QSYS/QCMD)
      CLS(QSYS/QSYSCLS40)
```

After making this change, you must stop and restart all your servers.

Automatically starting instances

To enable an instance to start automatically each time the system restarts, you must add one of the commands to your QSTRUP program. You can also add the commands to a job scheduler.

Connecting to instances

Connecting from Content Manager 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.

About this task

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

Content Manager OnDemand provides commands that you can use to work with an instance.

About this task

For example, Content Manager OnDemand provides commands to add, remove, and print spooled files and other data types and to manage storage. On many of the Content Manager 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 i: Administration Guide* appendix entitled “Command reference” for more information about OnDemand commands.

Content Manager 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 i: Administration Guide* appendix entitled "API and user exit reference" for more information about Content Manager OnDemand APIs.

Defining a locale

About this task

Content Manager 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 Content Manager 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.

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. If the locale you require must be created from source code, you must install option 21 of IBM i (Extended NLS Support; a product option of 5770-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.

Important consideration regarding locales

The user profile named the same as your instance AND all user profiles that perform Content Manager OnDemand data loading functions (such as **ADDRPTOND**, **STRMONOND**, **arsload** or **arsdoc add**) must have their locale set to the locale of the instance.

If the user profiles that perform Content Manager OnDemand data loading functions are also used for other non-Content Manager OnDemand work on your system, changing their LOCALE value might impact other non-Content Manager OnDemand applications. Consider creating specific Content Manager 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 Content Manager 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 Content Manager 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 *DECfmt
          *SRTSEQ) LOCALE('/QSYS.LIB/ES_ES_E.LOCALE')
```

After setting the correct locale on an IBM i user profile, you might choose to copy that user profile to create another similar user profile for use with Content Manager OnDemand. The IBM i copy function 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.

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

Steps to take after setting locales

If you encounter problems while trying to create a Content Manager OnDemand Common Server instance, see the list of problems, causes, and recoveries below.

About this task

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 was not created. The qshell job log also contains OND0588 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.

Problem: Your data will not load into OnDemand. There is no system log message indicating the cause of failure. The job log 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 job log for the qshell job, QP0ZSPWP, for the error message OND0588 You must set a locale.. If the job log contains this message, you must change the user profile running the load job to specify a locale.

Problem: Your spooled file fails to archive. The system log contains message 88, including text similar to the following Row 1: The string " 1/12/13" 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 might, 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 is compatible with the instance into which the data is being archived.

Configuring the ARS.CFG file

When you run the various Content Manager OnDemand programs, they obtain information from the ARS.CFG file about a variety of configuration settings.

About this task

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 Content Manager 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 might 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 Content Manager 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 Content Manager 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 does not automatically start by issuing STRTCPSVR SERVER(*ONDM) INSTANCE(*AUTOSTART) 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.

ARS_DISABLE_ARSLOG parameter

The national language environment used to create the database. Content Manager OnDemand uses a three-character language code to derive the locale and code set for the server operating system.

You must provide the correct language code for your operating environment before you create the Content Manager OnDemand instance database.

This parameter specifies that the ARSLOG exit is to be disabled completely. Without specifying this parameter, even though all user exit logging is disabled in the System Parameters window of the Content Manager OnDemand Administrator client, an attempt is still made to call the ARSLOG exit for certain messages. By specifying ARS_DISABLE_ARSLOG=1 in the ARS.CFG file, the ARSLOG exit is not invoked at all. If you do not plan on using the ARSLOG exit, specify ARS_DISABLE_ARSLOG=1 to minimize the overhead of attempting to call the ARSLOG exit.

ARS_LANGUAGE parameter

The national language environment used to create the database. Content Manager OnDemand uses a three-character language code to derive the locale and code set for the server operating system.

You must provide the correct language code for your operating environment before you create the Content Manager OnDemand instance database.

ARS_LDAP_ALLOW_ANONYMOUS parameter

Specifies whether or not anonymous bind connections are allowed on this LDAP server. Valid values are TRUE and FALSE. If FALSE, you must also specify an LDAP user ID and password in the stash file.

ARS_LDAP_BASE_DN parameter

Specifies the base distinguished name to use.

This parameter is required for LDAP authentication.

Example 1:

```
ARS_LDAP_BASE_DN=ou=mycity,o=xyzcompany
```

Example 2:

```
ARS_LDAP_BASE_DN=dc=ondemand,dc=xyzcompany
```

ARS_LDAP_BIND_ATTRIBUTE parameter

Specifies the attribute being bound and is the attribute name to be searched on the LDAP server.

This parameter is required for LDAP authentication.

Example:

```
ARS_LDAP_BIND_ATTRIBUTE=mail
```

ARS_LDAP_BIND_MESSAGES_FILE parameter

Specifies the location of a file containing the LDAP message strings the Content Manager OnDemand server looks for during login.

This is used for issuing messages when the user's password is about to expire, or their LDAP account is locked. **ARS_LDAP_BIND_MESSAGES_FILE** is used in conjunction with the ARSLDAP.INI file to implement this functionality.

ARS_LDAP_IGN_USERIDS parameter

This parameter specifies the user IDs that Content Manager OnDemand ignores when you enable LDAP for authentication. If the parameter does not exist or you do not specify a value, Content Manager OnDemand defaults to ADMIN.

You can specify up to 10 user IDs, delimited by a comma. If you specify a list of user IDs and you want to include ADMIN, you must specify it on the list.

ARS_LDAP_MAPPED_ATTRIBUTE parameter

Specifies the attribute being returned to Content Manager OnDemand as a user ID.

This is the attribute name to be returned from the LDAP server once the bind attribute name is found. It can be the same as the bind attribute or different. This parameter is required for LDAP authentication.

Example:

```
ARS_LDAP_MAPPED_ATTRIBUTE=sAMAccountName
```

ARS_LDAP_PORT parameter

Specifies the port on which LDAP is listening. The default value is 389. This parameter is optional.

ARS_LDAP_SERVER parameter

Specifies the IP address or the fully-qualified hostname of the LDAP server. This parameter is required for LDAP authentication.

ARS_MESSAGE_OF_THE_DAY parameter

Use to specify the message of the day for the Content Manager 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_MSGS_LANGUAGE parameter

If you want something other than English for server messages, change this parameter in 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 IBM i server after you install Content Manager OnDemand.

ARS_NUM_DBSRVR parameter

Determines the number of processes that Content Manager OnDemand starts on the server to handle connections to the database.

The default value is 5.

ARS_PRINT_PATH parameter

The location where the Content Manager 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 Content Manager OnDemand instance).

ARS_SUPPORT_CFSOD parameter

To use Content Federation Services for Content Manager OnDemand, you must set this parameter to 1.

See the Web Index document for additional details about the set up and use of Content Federation Services for Content Manager OnDemand.

ARS_SUPPORT_HOLD parameter

Enables the Enhanced Retention Management feature, a separately priced feature.

For more information about this feature, see *Enhanced Retention Management Guide*.

ARS_TMP parameter

The location where Content Manager 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 Content Manager 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 might correct the problem: ARSSOCK_RESOLVE_CLIENT_NAME=0

Configure your system at start up

There are server functions you might 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 additions to the **QSTRUP** program:

- Start the Content Manager OnDemand server jobs
- Start the Content Manager OnDemand monitor

Start the server jobs

The following example shows the command that starts the Content Manager OnDemand server jobs: `STRTCPSVR SERVER(*ONDMD)`

Start the monitors

Use the Start Monitor for OnDemand (**STRMONOND**) command to start Content Manager OnDemand output queue monitors for output queues or directory monitors to receive files destined for Content Manager OnDemand.

Before you begin

Review the online help text to determine the specific parameters you wish to specify for each monitor.

About this task

You can also create monitor definitions using the Content Manager OnDemand component of IBM Navigator for i. After starting IBM Navigator for i, select Content Manager OnDemand, then click Directory Monitors or Output Queue Monitors as needed.

See Loading spooled file data in the *IBM Content Manager OnDemand for i: Common Server Administration Guide* for more information on 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 the IBM Navigator for i as described above.

Scheduling maintenance programs

This section describes the Content Manager OnDemand maintenance commands to run automatically on a regular schedule.

Before you begin

The maintenance programs that you might want to run on a regular schedule include those that maintain application group data in the database and on disk, and maintain Content Manager OnDemand system data in the database.

About this task

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**, **arsload**, **arsdoc add**) at the same time that you run the disk storage maintenance command (**STRDSMOND**).

However, the **STRASMOND** command can be run while archiving data into Content Manager OnDemand. Backups should not be run while **STRDSMOND**, **STRASMOND**, or loading of data are running. 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 might 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

About this task

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 might 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 i: Common Server Administration Guide* for more information about the STRDSMOND and STRASMOND commands.

Backing up the Content Manager OnDemand database

After you finish loading data into Content Manager OnDemand and maintaining the database, we recommend that you create a backup image of the Content Manager OnDemand database.

About this task

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 might want to backup the database once a day, after the last load process completes.

See the section on backing up data and *IBM Content Manager OnDemand for i: Common Server Administration Guide* for more information about backing up your Content Manager OnDemand data.

Part 5. Preparing the system for use

This section of the book describes how to verify the installation of Content Manager 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 might 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 disk pool, optical, or tape) 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 Content Manager OnDemand, IBM recommends that you create a backup copy of the Content Manager OnDemand database.

Saving Content Manager OnDemand passwords into encrypted files

You can store user IDs and passwords in encrypted files (also called stash files).

About this task

Storing passwords in stash file can help you improve security because you do not need to specify the password on the command line, where the password might be visible to others. You can store the user ID and password for the following situations in one stash file:

- Each Content Manager OnDemand instance
- Each Content Manager OnDemand program that runs as a daemon or service (for example, arslload)

You store the stash file in a directory and specify that directory in the **SRVR_OD_STASH** parameter of the **ARS.INI** file. Content Manager OnDemand programs locate the stash file in that directory. If you need to override the user ID and password stored in the stash file, create a stash file and store it in a directory where you run a Content Manager OnDemand program. For security reasons, limit access to the file through file permissions or delete it when you no longer need it.

To store the user IDs and passwords into a stash file, do the following steps:

Procedure

1. Create a stash file by running the **arsstash** command. The command prompts you for the password. For a description of the syntax of the **ARSSTASH** command, see “Syntax of the **ARSSTASH** command” on page 92. For examples of **ARSSTASH** commands, see “Examples of **ARSSTASH** commands.”
2. Save the stash file in a directory and limit access to that file through file permissions.

Examples of ARSSTASH commands

To store the user ID **USER2** and password for the LDAP service in an existing stash file called **ars.stash**, enter the following command: `arsstash -a 7 -s ars.stash -u USER2`


To store the user ID **USER3** and password for the **ARSLOAD** daemon in an existing stash file called **ars.stash**, enter the following command: `arsstash -a 3 -s ars.stash -u USER3`

To store the user ID **DEMOUSER** and password for use by any OnDemand instance, daemon, or service, and store that ID and password in an existing stash file called **ars.stash**, enter the following command: `arsstash -a 1 -s ars.stash -u DEMOUSER`

To view the user ID currently assigned to **ARSLOAD**, enter the following command: `arsstash -a 3 -s ars.stash`

Syntax of the ARSSTASH command

The following diagram describes the syntax of the ARSSTASH command:



The following list describes the parameters of the ARSSTASH command:

action

Specifies whether to store or delete a user ID and password from the stash file. If you do not specify the **-a** parameter, the command runs action 1. If you do not specify the **-u** parameter with actions, then the ARSSTASH command displays the user ID currently assigned to that instance, daemon, or service. You can specify one of the following values for the *action* parameter:

- 1 Store the Content Manager OnDemand user ID and password. If you do not specify an action, this action is the default action.
- 2 Delete the Content Manager OnDemand user ID and password.
- 3 Store the Content Manager OnDemand user ID and password for the ARSLOAD program.
- 6 Store the Content Manager OnDemand user ID and password for Production Data Distribution.
- 7 Store the LDAP user ID and password.

-c Create a stash file.

-s *stash_file_name*
Name of the stash file.

-u *userid*
User ID for the Content Manager OnDemand instance, daemon, or service specified by the *action* variable.

Verifying the installation

You can verify the Content Manager OnDemand installation by completing the following tasks.

Procedure

To verify the installation:

1. Complete all of the required steps to install server software.
2. Complete all of the required steps to configure the server.
3. If you have not already done so, install the Content Manager OnDemand client program on a workstation.
4. Start the Content Manager OnDemand client program. Content Manager OnDemand displays the Logon to a Server dialog box.
5. Click **Update Servers**. Content Manager OnDemand displays the Update Servers dialog box.
6. Enter the name of the Content Manager OnDemand server. Click Help for information about the fields and options.
7. Click **Add > Close** to return to the Logon to a 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 a Content Manager OnDemand user ID and password in the fields provided. The first time that you log on to the system, you must specify the built-in Content Manager OnDemand user ID, QONDADM. Initially, the password is QONDADM1. However, you will be prompted to enter and verify a new password.
10. Open and search the **System Log** folder.

Results

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

Defining migration policies and storage sets

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

You can set up Content Manager 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 media such as disk pool, optical, or tape. The Content Manager OnDemand online help provides details about defining storage characteristics. Also see the *IBM Content Manager OnDemand for i: 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.

Configuring the System Log application group

When you install and configure Content Manager 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 Content Manager 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 migrate a copy of the system log data to archive media.

IBM recommends that, before you begin defining reports to Content Manager OnDemand, loading data into the system, or allowing users to access the system, 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 disk pool, optical, or tape, and then by updating the System Log application group and assigning it to the storage set.

No matter where you decide to keep your System Log data, it is important to note that searching the System Log is one of a number of ways to locate the Load ID for archived data. Every time data is loaded into Content Manager 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.

Maintaining system log data in archive storage

IBM recommends 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.

About this task

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 the archive storage you have specified.

Procedure

After you create the migration policy, complete the following steps to assign the System Log application group to a storage set:

1. Start the Content Manager OnDemand Administrator client.
2. Log on to the server with a user ID that has system administrator authority. (The default user named QONDADM has system administrator authority.)
3. Click **Application Groups**.

4. Right-click on the **System Log application group**.
5. From the pop-up menu, select **Update** to open the Update an Application Group window.
6. Click the **Storage Management** tab.
7. In the Storage Set Name list, select the name of the storage set. The storage set name selected should match the migration policy definition created for your System Log data.
8. Under Cache Data, specify at least 1 day for the Cache Document Data for x Days value. The System Log application group is created with a default value of 10000 days in cache. You can reduce this number when setting up the System Log application group to use a Storage Set. You must specify at least 1 day in cache so that the system log entries are appended to previous entries each day, for more efficient storage management, prior to them being sent to the storage manager.
9. Click **Advanced** to open the Advanced Storage Management dialog box.
10. Select **Next Cache Migration** under Migrate Data from Cache. Content Manager OnDemand copies the system log data to archive storage the next time the maintenance commands execute. You can also select **After x days in cache** if you want your system log data to stay in cache (on disk) for some period of time before it is copied to your archive storage. You can select **Next Cache Migration** or **After x Days in Cache** instead of **When Data is Loaded** so that the system log entries are appended to previous entries each day for more efficient storage management, prior to them being sent to the storage manager.
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.

Reducing system log size

Another way to manage the space required for your system log data is to remove or limit messages in the log that are not important to you.

Perhaps you do not feel the need to keep every log entry for every document retrieval, every annotation addition, and every logoff, for example. However, you might choose to keep system log entries for message number 87, which indicates a successful load and contains the Load ID for the particular data being loaded (which is required for **PRTRPTOND** and **RMVRPTOND** as described earlier in this chapter).

Existing system log entries can be selectively deleted, using the ARSDOC DELETE API, run in the QSHHELL environment on your IBM i server. An example is shown below, and assumes you are already in QSHHELL at the time the API is executed. (See the *IBM Content Manager OnDemand for i: Common Server Administration Guide* for more details regarding the use of the ARSDOC DELETE API.) System log entries can also be limited such that only those you feel are important are ever logged at all. See below the ARSDOC DELETE example for more information on that option as well.

Be sure to save any system log files before you perform any cleanup activities. The first system log that OnDemand uses is a physical file named SL2, found in the instance library. When the SL2 file becomes full, the next system log file is automatically created and named SL3, and continues to be created and incremented by one as needed.

In the following example, use the ARSDOC DELETE function to delete all messages from the system log for the specified date range, except for system log message 87:

```
$
> arsdoc delete -h QUSROND -v -i "WHERE NOT(MSG_NUM = 87)" -G 'System Log'
-f 'System Log' -S 01.05.2007,10.05.2007,%d.%m%Y
06/06/07 11:30:39: Starting arsdoc. Version: 7.2
06/06/07 11:30:39: QRDARS/ARSDOC delete -h QUSROND -v -i WHERE NOT(MSG_NUM = 87)
-G System Log -f System Log -S 01.05.2007,10.05.2007,%d.%m%Y
06/06/07 11:30:39: Searching for folder 'System Log' ...
06/06/07 11:30:39: Search successful
06/06/07 11:30:39: Searching for documents in 'System Log' ...
06/06/07 11:30:39: Querying database with SQL string 'WHERE NOT(MSG_NUM = 87)'
06/06/07 11:30:41: Search successful
06/06/07 11:30:41: Deleting 6161 documents
06/06/07 11:30:45: Delete successful
06/06/07 11:30:41: arsdoc completed.
```

After you run ARSDOC DELETE to clean up your unneeded system log entries, Content Manager OnDemand reuses the space that had been used by the records you just deleted. Alternatively, you might choose to reorganize the system log files using the IBM i Reorganize Physical File Mbr (RGZPFM) command. For example, you could run RGZPFM instance_name/SL2 where instance_name is the name of your OnDemand instance. At a minimum, you might want to reorganize all but the most current system log file.

As a final step, you might want to limit what is sent to the system log in the future. The OnDemand Administrator Client allows you to select which OnDemand activities create entries in the system log. To select which activities create system log entries for an OnDemand server:

1. Log on to the server with the Content Manager OnDemand Administrator Client.
2. Right-click the server.
3. From the pop-up menu, select System Parameters, to invoke the System Parameters dialog box.
4. Select the OnDemand activities you want to log within the System Logging and User Exit Logging boxes.

Configuring the System Load application group

When you install and configure Content Manager OnDemand, you also initialize the system load logging facility.

About this task

The system load logging facility is comprised of the System Load application group, a set of system load applications, and the System Load folder. The System Load application group contains the storage management information that Content Manager OnDemand uses to maintain the data written to the system load logging facility. 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 migrate a copy of the system load data to archive media.

Before you begin defining reports to Content Manager OnDemand, loading data on the system, or allowing users to access the system, you should configure the System Load application group to maintain a permanent copy of the data that is written to the system load logging facility. You can do this by first defining a migration policy (and its corresponding storage set) that specifies archive storage media such as a disk pool, optical, or tape, and then by updating the System Load application group and assigning it to the storage set.

No matter where you decide to keep your System Load data, it is important to note that searching the System Load folder is the most direct way to locate the Load ID for archived data. Every time data is loaded into Content Manager OnDemand, a record is written to the System Load logging facility and the Load ID is recorded as part of that record. The Load ID (also called "Report ID") is a required parameter for the Remove Report (RMVRPTOND) and Print Report (PRTRPTOND) commands.

Assigning the System Load application group to a storage set

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

Before you begin

The Content Manager OnDemand online help provides information about the options on the Storage Management page of application group definitions.

About this task

After you create the migration policy (which automatically creates a storage set with the same name) you must update the System Load 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 load data in the archive storage you have specified. Complete the following steps to assign the System Load application group to a storage set:

Procedure

1. Start the Content Manager OnDemand Administrator client.
2. Log on to the server with a user ID that has system administrator authority. (The default user named QONDADM has system administrator authority.)
3. Click **Application Groups**.
4. Right-click on **System Load application group**.
5. From the pop-up menu, select **Update** to open the Update an Application Group window.
6. Click the **Storage Management** tab.
7. 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.
8. Click **OK**.

Configuring the System Migration application group

Index migration is the process by which Content Manager OnDemand moves index data from the database to archive storage.

Index migration is not recommended.

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. Content Manager 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 Content Manager 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 Content Manager OnDemand uses to maintain index data migrated to archive storage. Until you assign the application group to a storage set that specifies archive media, Content Manager OnDemand cannot migrate index data from the database to archive storage.

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.

Before you begin

The Content Manager OnDemand online help provides information about the options on the Storage Management page.

About this task

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:

Procedure

1. Start the Content Manager OnDemand Administrator client.
2. Log on to the server with a user ID that has system administrator authority. (The default user named QONDADM has system administrator authority.)
3. Click **Application Groups**.
4. Right-click on **System Migration application group**.
5. From the pop-up menu, select **Update** to open the Update an Application Group window.
6. Click the **Storage Management** tab.

- |
 - |
 - |
 - |
7. 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.
 8. Click **OK**.

Backing up the Content Manager OnDemand data

To complete the installation and configuration process, we recommend that you create a full, offline backup of the Content Manager 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 i: Common Server Administration Guide* for more information about backing up your Content Manager 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.

Part 6. Appendixes

Appendix A. Database servers

The **ARS_NUM_DBSRVR** parameter in the **ARS.CFG** file determines the number of processes that Content Manager 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.

Each access to the database requires a database agent. Content Manager 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 Content Manager OnDemand handles the database load. For example, you can define **ARS_NUM_DBSRVR** so that Content Manager OnDemand starts a fixed number of database agents, regardless of the number of concurrent database requests. While this might 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.

Description

Each access to the database requires a database agent. Content Manager 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 Content Manager OnDemand handles the database load. For example, you can define **ARS_NUM_DBSRVR** so that Content Manager OnDemand starts a fixed number of database agents, regardless of the number of concurrent database requests. While this might 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 B. National language support

Overview

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

Conversion between different code pages

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

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, Content Manager OnDemand clients and applications should always use the same code page as the Content Manager OnDemand instance. However, this is not always practical or possible. Content Manager OnDemand provides support for character conversion that allows clients, applications, and instances to use different code pages. While a Content Manager 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 Content Manager OnDemand can contain characters encoded in any code page.

However, when you use different code pages, Content Manager OnDemand might 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 various situations.

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 Content Manager 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 a Content Manager 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 Content Manager OnDemand. The resulting Unicode data is then converted to the code page of the instance. For example, the user enters a user ID, password, and server name to logon to a server. Content Manager 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 Content Manager OnDemand. The resulting Unicode data is then converted to the code page of the client. For example, after authenticating the user ID and password, the

server builds a list of folder names that the user is authorized to open. Content Manager 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 Content Manager OnDemand index data is in a code page that is different than the code page of the Content Manager OnDemand instance.

Character data conversion takes place on the server using Unicode code page mapping tables. Content Manager 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 Content Manager OnDemand.

When you store documents in Content Manager 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 Content Manager OnDemand.

When a user retrieves a document from Content Manager 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 might 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 double-byte character set (DBCS) AFP data and DBCS and single-byte character set (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 ICU 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 ICU converters automatically map the user-defined area of a code page to the standard user-defined area of the corresponding ICU table. If the code page contains UDC mappings outside the standard user-defined area, you can create and use your own ICU converter. (Otherwise, the viewing program cannot display the characters correctly.) Use the Character Data Representation Architecture (CDRA) utility to create your own ICU converter. For AFP data, change the **IconvLocalePath** parameter in the Preferences section of the `FLDPORT2.INI` file to use your own ICU converter. For information concerning character data conversion, font mapping, and related matters, see the *IBM Content Manager OnDemand Windows: Client Customization Guide*.

The AFP Viewer operates internally in Unicode and uses ICU to convert data.

Configuring your system for NLS

When you create a Content Manager OnDemand instance, you specify the language identifier and locale for the instance. These values determine language settings for your instance.

Language identifiers, CCSIDs, and locales

Language identifiers are used to tag difference languages that Content Manager OnDemand supports.

The information can help you determine the language identifier and locale to use when creating a Content Manager OnDemand server instance. An instance can be created in any language, even if that language is not installed for IBM i. Starting with V5R4, the base locale for each language was changed to include support for the Euro currency. New locales were created without Euro support. For example, locale ES_ES now includes Euro support, as does the existing ES_ES_E locale. If you need a Spanish locale without Euro support, you must use ES_ES_PE. The table does not provide a complete list of all language identifiers, CCSIDs, or locales.

Table 6. Supported language identifiers, CCSIDs, and locales

Language Identifier	Country or region	CCSID	Locale
CHS	China	1388	ZH_CN
HRV	Croatia	870	HR_HR
CZE	Czech Republic	870	CS_CZ
DAN	Denmark	1142	DA_DK
ARA	Egypt	420	AR_AA
FIN	Finland	1143	FI_FI
FRA	France	1147	FR_FR
DEU	Germany	1141	DE_DE
ELL	Greece	4971	EL_GR
HUN	Hungary	870	HU_HU
ITA	Italy	1144	IT_IT
JPN	Japan	5035	JA_5035
KOR	Korea	933	KO_KR
NLD	Netherlands	1140	NL_NL
NOR	Norway	277	NO_NO
PLK	Poland	870	PL_PL
PTB	Brazil	37	PT_BR
RUS	Russia	1025	RU_RU
SKY	Slovakia	870	SK_SK
SLO	Slovenia	870	SL_SI
ESP	Spain	1145	ES_ES
SVE	Sweden	1143	SV_SE
CHT	Taiwan	937	ZH_TW
ENU	United States	37	EN_US

Note:

- For a Japanese language instance, the instance must be created with CCSID 5035 and locale JA_5035. Using a CCSID and locale combination that only supports uppercase characters causes unreadable Content Manager OnDemand messages in the System Log.

- The Czech language code in IBM i is CSY; the Czech language code in Content Manager OnDemand is CZE.

Mapping AFP fonts

You might 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 Content Manager 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 might not be able to display the characters properly. (If your documents use fonts that are not defined to Content Manager 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 Content Manager OnDemand was installed (by default, \Program Files\IBM\OnDemand32).

Content Manager 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 might 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 Content Manager 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 Information 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 might 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 i: Indexing Reference* for more information.

Logical Views

If you use the sample data window to create logical views for a report that contains double-byte character data, you might 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 might 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 Content Manager OnDemand programs

Certain Content Manager 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 7 lists the DBCS countries or regions that are supported by Content Manager 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 7. DBCS code pages

Country or Region	Host Code Page
China	935
Japan	290 or 1027
Korea	933
Taiwan	937

Appendix C. System log and the system log user exit

Content Manager OnDemand generates messages about the various actions that occur on the system. For example, when a user logs on the system, Content Manager OnDemand generates a message that contains the date and time, the type of action, the user ID, and other information.

You can configure Content Manager 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 Content Manager OnDemand to generate system messages and specify the types of messages generated by selecting the appropriate options in the System Parameters dialog box.
- Enable Content Manager OnDemand to generate application group messages by selecting the appropriate option in the System Parameters dialog box.
- Specify the types of application group messages generated by selecting options on the Message Logging page in application groups.
- Enable Content Manager 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, Content Manager OnDemand automatically saves the messages in the system log and sends the messages to the user-defined program.

Content Manager 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 Content Manager OnDemand sends to the **ARSLOG** program contain the parameters listed in the table.

Table 8. ARSLOG program parameters (system log user exit program)

Parameter	Purpose	Size	Example
\$1	Content Manager OnDemand instance	10 char	QUSROND
\$2	Time stamp	20 char	08/13/95 14:24:31
\$3	Log record identifier	10 char	57049
\$4	Content Manager OnDemand user ID	128 char	QONDADM
\$5	User's accounting information	60 char	Z76-001J/999999

Table 8. **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	256 char	This parameter is used only if the system log message being passed to the user exit has additional information not included in the message text. The file is deleted immediately after the exit program returns control to Content Manager OnDemand. ¹

Note:

- For example, when viewing system log messages 87 or 88 from the Content Manager OnDemand Client, Yes is indicated under the View column heading of the document list. View=Yes indicates that additional information is available for those messages. The document file parameter contains the complete path name of the file that contains the additional text for the user exit to process. The file is stored in the directory that is specified by the **ARS_TMP** parameter in the **ARS.CFG** file. An example of the contents of this parameter as it is passed to the user exit might be: /QIBM/USERDATA/ONDEMAND/QUSROND/TMP/ARS.155.0000000000000013.SYSLOG.

If you create your own **ARSLOG** program, remember that the **ARSLOG** program and any programs that it might call run as user QRDARS400. QRDARS400 must have authority to any objects that your **ARSLOG** program uses.

Content Manager 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 QUSR RDARS. Any program source code that you modify should not be placed in QRDARS because that library is replaced during software upgrades. However, QUSR RDARS library is not replaced and can be used for your modified source.

Once you have created your program, you should compile the program and place it in your own library (not **QRDARS**). Then, you must change the symbolic link at /QIBM/PRODDATA/ONDEMAND/BIN/ARSLOG to point to your program. For example, if you called your program **MYARSLOG** and placed it in library **QUSR RDARS**, you

would first delete the current symbolic link using this command: **RMVLNK OBJLNK('/QIBM/PRODDATA/ONDEMAND/BIN/ARSLOG')**. Then, you would add the link to your program using this command: **ADDLNK OBJ('/QSYS.LIB/QUSRRDARS.LIB/MYARSLOG.PGM') NEWLNK('/QIBM/PRODDATA/ONDEMAND/BIN/ARSLOG')**

The last step is to ensure that the ARSLOG program is enabled. Check the ARS.CFG file to confirm that the ARS_DISABLE_ARSLOG entry is not set to a 1 (to disable) for the instance for which you can enable logging. ARS_DISABLE_ARSLOG=1 disables the ARSLOG exit program.

The online Help for the administrative client provides information about enabling Content Manager 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 Content Manager OnDemand generates and sends to the **ARSLOG** program.

Notices

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

IBM might 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 might be used. Any functionally equivalent product, program, or service that does not infringe on any IBM intellectual property right might 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:

Intellectual Property Licensing
Legal and Intellectual Property Law
IBM Japan Ltd.
1623-14, Shimotsuruma, Yamato-shi
Kanagawa 242-8502 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. The sample programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

If you are viewing this information softcopy, the photographs and color illustrations may not appear.

Trademarks

IBM, the IBM logo, and [ibm.com](http://www.ibm.com)[®] are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. If these and other IBM trademarked terms are marked on their first occurrence in this information with a trademark symbol ([®] or [™]), these symbols indicate U.S. registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at “Copyright and trademark information” at www.ibm.com/legal/copytrade.shtml

Adobe, Acrobat, Portable Document Format (PDF), and PostScript are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, other countries, or both.

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.

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

Lotus is a trademark of Lotus Development Corporation in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

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

Special characters

ARS_NUM_DBSRVR parameter 109

A

administrative
 roles and responsibilities 16
administrative client
 installing 63
AFP data
 indexing 29
 input 25
AFP fonts 114
annotation color search field 41
API 117
APIs 17
application
 about 35
 naming 39
 overview 4, 35
application group
 about 33
 database field types 41
 field types 41
 indexing 33
 naming 39
 overview 5, 33
 permissions 38
application groups 4
 maintaining 87
 migrating data to archive storage 87
 scheduling maintenance 87
 system load 101
 system log 97
 system migration 103
application programming interface
 system log 117
 user exit programs 117
application programming interfaces 17
application programs
 indexing with 29
applications 4
archive storage 21, 53
 about 21
 maintaining system log data 97
 migrated indexes 48
 overview 21
 reports on archive media 48
 requirements 48
ARS_AUTOSTART_INSTANCE
 parameter 81
ARS_DISABLE_ARSLOG parameter 81
ARS_LANGUAGE parameter 82
ARS_MESSAGE_OF_THE_DAY
 parameter 83
ARS_MSGS_LANGUAGE parameter 83
ARS_NUM_DBSRVR parameter 83, 109
ARS_PRINT_PATH parameter 84
ARS_SUPPORT_CFSOD parameter 84
ARS_SUPPORT_HOLD parameter 84

ARS_TMP parameter 84
ARS.CACHE file 67
ARS.CFG file
 ARS_AUTOSTART_INSTANCE
 parameter 81
 ARS_DISABLE_ARSLOG
 parameter 81
 ARS_LANGUAGE parameter 82
 ARS_MESSAGE_OF_THE_DAY
 parameter 83
 ARS_MSGS_LANGUAGE
 parameter 83
 ARS_NUM_DBSRVR parameter 83,
 109
 ARS_PRINT_PATH parameter 84
 ARS_SUPPORT_CFSOD
 parameter 84
 ARS_SUPPORT_HOLD parameter 84
 ARS_TMP parameter 84
 ARSSOCK_RESOLVE_CLIENT_NAME
 parameter 85
autostarting instances 81
 configuring 81
 creating 67
database connections, specifying 83
database manager parameters 109
language parameter 81, 82
message language parameter 83
message of the day parameter 83
printing parameter 84
server parameters 83
server print parameter 84
starting instances 81
temporary space parameter 84
ARS.DBFS file 67
ARS.INI file 67
ARSLOG 18
ARSLOG program 117
ARSSOCK_RESOLVE_CLIENT_NAME
 parameter 85
autostarting instances 81

B

back up media 51
backing up the database 88, 105
backup and recovery 51
 database 52
 documents 52
 overview 51
 reports 52
backup report storage space 48

C

cabinet
 overview 6
cabinets 4
cache expiration processing 33
cache migration processing 33

cache storage file systems
 maintaining 87
 scheduling maintenance 87
CCSIDs 67
CFSOD 84
character conversions 111
character mapping 112
checklist 59
client
 APIs 17
 customization 17
client/server system overview 4
code pages 111, 115
 conversions 111
code sets 111
collecting requirements 25
commands 17
compression 45
concepts
 application 4
 application group 5
 archive storage manager 10
 cabinet 6
 cache storage manager 10
 client/server system overview 4
 data loading 10
 document 8
 expiration 11
 folder 5
 indexing 10
 indexing methods 7
 library server 8
 loading data 10
 migration 11
 object server 8
 server 8
 storage manager 10
 system overview 4
configuration files
 saving 55
configuring the system 65
connecting to the database 83, 109
Content Manager OnDemand overview
 annotations 3
 full fidelity reprinting 3
 indexes 3
 resource management 3
conventions
 names 39
CRTINSTOND command 69, 70

D

data
 compression 45
 formats 25
 indexing 27
database
 backup 88, 105
 backup and recovery 52
 connections 83, 109

- database (*continued*)
 - creating 67
 - expiration processing 33
 - field names 39
 - field types 41
 - importing migrated index data 48
 - logs 47
 - maintaining 87
 - migration 103
 - recovery 52
 - recovery log storage requirements 47
 - scheduling maintenance 87
 - storage requirements 43
 - storage space requirements 47
 - transaction log storage requirements 47
- database agents 109
- database backup 52
- database manager 9
 - expiration 11
 - migration 11
 - overview 10
- database manager parameters
 - ARS_NUM_DBSRVR 109
- database organization
 - multiple loads per database table 33
- database recovery 52
- database rows 47
- database storage 47
- databases 109
 - backing up 52
 - journaling 52
 - recovering 52
- date fields
 - database field type 41
 - folder field type 41
 - indexing 30
- DB2
 - ARS_NUM_DBSRVR parameter 109
 - connections 83, 109
- DBCS
 - application group fields 114
 - code pages 115
 - database fields 114
 - fonts 114
 - generic indexer 116
 - graphical indexer 116
 - indexing data 116
 - languages, support for 111
 - logical views 116
 - programs 116
 - report wizard 116
 - Shift Out Shift In structured fields 115
 - SOSI structured fields 115
- defining a report 31
- defining a storage node 95
- defining a storage set 95
- defining the storage configuration 22
- DHCP
 - ARSSOCK_RESOLVE_CLIENT_NAME 83
- disk storage 53
 - requirements 46
- document
 - about 8
 - defined 1
 - expiration 11
- document (*continued*)
 - indexing 10, 27
 - migration 11
- documents
 - backup and recovery 52
- E**
 - end-user client
 - installing 63
 - environment variables
 - time zone (TZ) 41
 - TZ (time zone) 41
 - expiration
 - overview 11
 - expiration processing 33
- F**
 - field names 39
 - field types 41
 - file formats supported by
 - SCS data 25
 - SCS-Extended data 25
 - file systems
 - cache storage 87
 - files 55
 - filter fields
 - fields, filter 33
 - folder
 - about 32
 - annotation color search field 41
 - field names 39
 - field types 41
 - naming 39
 - overview 5, 32
 - permissions 37
 - query field types 41
 - search field types 41
 - folders 4
 - font files 55
 - fonts 114
- G**
 - generic data
 - indexing 29
 - generic indexer 116
 - generic indexer program 30
 - graphical indexer 116
 - graphical tool 29
 - group permissions overview 37
 - groups
 - about 36
 - names 39
 - overview 36
- H**
 - HOLD 84
- I**
 - i Navigator
 - configuration vii
 - i Navigator (*continued*)
 - Management Central vii
 - system navigation vii
 - identifiers
 - language 113
 - IFS directories 51
 - importing migrated index data
 - storage space requirements 48
 - index data
 - migrating 103
 - index fields
 - fields, index 33
 - indexes
 - archive storage space requirements 48
 - indexing
 - AFP data 29
 - alternatives 29
 - application groups 33
 - application programs 29
 - date fields 30
 - DBCS data 116
 - documents 27
 - generic data 29
 - generic indexer program 30
 - line data 27
 - methods 7
 - other data 29
 - other indexing programs 29
 - overview 10
 - PDF data 29
 - PDF indexer program 30
 - programs 7, 29
 - reports 28
 - requirements 27
 - space requirements 46
 - temporary space 46
 - initialization files 85, 87
 - initializing the system 65
 - input data formats 25
 - installation checklist 59
 - installation procedures 61
 - installing 61
 - Content Manager OnDemand software 61
 - software 59
 - instances 67
 - automatically starting 74
 - autostarting 81
 - changing an instance 69
 - configuring 67, 69
 - configuring an instance 67
 - connecting 74
 - creating 74
 - deleting 70
 - deleting an instance 70
 - starting 74, 81
 - working 74
- J**
 - jobs
 - servers 85

L

- language identifiers 113
- language parameter 81, 82
- LDAP
 - AIX 82, 83
 - anonymous bind connections 82
 - attribute 83
 - attribute binding 82
 - base distinguished name 82
 - hostname of LDAP server 83
 - LDAP port 83
 - login message strings 82
 - login user ID strings 82
- line data
 - indexing 27
 - input 25
- load, system
 - application group, configuring 101
- loading data
 - overview 10
- locale 81, 82
- locales
 - commands 77
 - defining 77
 - setting 78
- log files 17
- log sizes 98
- log, system
 - API definition 117
 - user exit programs 117
- logical server environments 67
- logical views 116
- logon
 - ARSSOCK_RESOLVE_CLIENT_NAME 85
- logon user exit 17
- logs
 - storage space requirements 47

M

- maintaining software 55
- maintenance programs
 - scheduling 87
- management programs
 - overview 11
- Map Coded Font Format 2 structured fields 116
- mapping
 - characters 112
- mapping AFP fonts 114
- MCF2 structured fields 116
- MCF2REF parameter 116
- message language parameter 83
- message of the day parameter 83
- messages
 - API definition 117
 - system log 117
 - user exit programs 117
- migrated indexes
 - storage space requirements 48
- migrating index data 103
- migration
 - application group data 87
 - application group, configuring 103
 - cache storage to archive storage 87
 - overview 11

- migration policies 95
- Monarch
 - integrating with the client 17
- multiple loads per database table 33

N

- names 39
- national language support 81, 82, 83, 111
- National language support 113
- NLS 111, 113
- notices 121

O

- objects 31
- objects for reports 31
- ODWEK 12
- OS/400 Indexer
 - overview 10
- output queue monitors
 - starting 85
- overview 31

P

- page
 - indexing 28
- parameters
 - ARS_NUM_DBSRVR 109
 - ARS_AUTOSTART_INSTANCE 81
 - ARS_DISABLE_ARSLOG 81
 - ARS_LANGUAGE 82
 - ARS_MESSAGE_OF_THE_DAY 83
 - ARS_MSGS_LANGUAGE 83
 - ARS_NUM_DBSRVR 83, 109
 - ARS_PRINT_PATH 84
 - ARS_SUPPORT_CFSOD 84
 - ARS_SUPPORT_HOLD 84
 - ARS_TMP 84
 - ARSSOCK_RESOLVE_CLIENT_NAME 85
 - MCF2REF 116
- Passport Advantage 61
- password length 39
- passwords 36, 39
- PDF (Portable Data Format) 25
- PDF data
 - indexing 29
- PDF indexer program 30
- permissions
 - about 37
 - application group 38
 - folder 37
 - overview 37
- planning 15
 - installation 23
- policies, migration 95
- Portable Data Format (PDF) 25
- preparing 15
- preparing the system for use 89
- prerequisite 89
- printing software
 - installing 59
 - temporary space 84
- problem solving 78

Q

- QSTRUP program 85
- query field names 39
- query field types 41

R

- recording messages 18
- recovery 51
 - database 52
- recovery log
 - storage space requirements 47
- reinstalling software 55
- report
 - expiration 11
 - formats 25
 - indexing 10, 27
 - migration 11
 - storage on archive media 48
 - storage requirements 43
- report storage space 48
- report wizard 116
- reports 25
 - backup and recovery 52
 - storing 53
- reports on disk 21
- request manager 9
- requirements
 - archive storage 48
 - database 48
 - database storage 47
 - disk storage 46
 - importing migrated index data storage space 48
 - indexing reports 46
 - migrated index storage space 48
 - recovery log storage requirements 47
 - reports on archive media 48
 - software storage 46
 - storage 43
 - temporary space 46
 - transaction log storage requirements 47
- requirements overview 43
- responsibilities
 - administrative 16
- roles
 - administrative 16

S

- saving configuration files 55
- scheduling maintenance programs 87
- SCS data
 - input 25
- SCS-Extended data
 - input 25
- search field names 39
- search field types 41
- server
 - overview 8
- server commands 17
- server jobs
 - starting 85
- server logging 17, 18

- server print software
 - installing 59
 - temporary space 84
- server software 51
- server work management 73
- servers
 - ARS.CFG file 81
 - IFS directories 51
 - managing work 73
 - starting 71
 - stopping 71
- setting objects 31
- setting the time zone environment
 - variable 41
- Shift Out Shift In structured fields 115
- short-term storage 53
- software
 - configuration files 55
 - installing 59
 - maintaining 55
 - obtaining 61
 - printing software
 - installing 59
 - reinstalling 55
 - saving configuration files 55
 - server print software
 - installing 59
 - upgrading 55
- software requirements
 - disk storage 46
 - storage 46
- SOSI structured fields 115
- starting instances 81
- storage
 - archive 53
 - archive storage requirements 48
 - database 43, 47, 48
 - disk space 46
 - disk storage requirements 46
 - estimating 43
 - importing migrated index data 48
 - indexing reports 46
 - migrated indexes 48
 - planning for 43
 - recovery log storage requirements 47
 - report 43
 - reports on archive media 48
 - requirements 43
 - temporary space 46
 - transaction log storage
 - requirements 47
- storage configuration
 - defining 22
- storage hierarchy 45
- storage management 45
- storage manager
 - expiration 11
 - migration 11
 - overview 10
- storage nodes 95
- storage objects 21
- storage sets 95, 101, 103
- system initialization 85, 87
- system load
 - application group, configuring 101
 - system load application groups 101

- system log
 - API definition 117
 - user exit programs 117
- system log application groups
 - configuring 97
- system log data 97
- system log file 17
- system log size
 - reducing 98
- system log user exit 18
- system logging facility 18
- system migration
 - application group, configuring 103
- system migration application
 - groups 103
- system overview 4

T

- tasks
 - administrative 16
- temporary space 84
 - storage requirements 46
- time zone environment variable 41
- transaction log
 - storage space requirements 47
- transforming data 25
- types of data 25
- TZ environment variable 41

U

- upgrading software 55
- user exit programs
 - ARSLOG 117
 - system log 117
- user exits
 - logon 17
 - system log 117
- user ID relationships 36
- user profiles 36
- users
 - about 36
 - names 39
 - overview 36
 - passwords 39

V

- verifying the installation 93

W

- Web enablement kit 12
- Windows client
 - APIs 17
 - customization 17
- Windows font files 55

IBM Confidential IBM Confidential IBM Confidential IBM Confidential

IBM Confidential IBM Confidential IBM Confidential IBM Confidential



Printed in USA

IBM Confidential IBM Confidential IBM Confidential IBM Confidential

SC19-2790-01



IBM Confidential IBM Confidential IBM Confidential IBM Confidential