

Content Manager OnDemand for i
Version 7 Release 2

Common Server Administration Guide



Content Manager OnDemand for i
Version 7 Release 2

Common Server Administration Guide



Note

Before using this information and the product it supports, read the information in "Notices" on page 323.

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

© **Copyright IBM Corporation 1991, 2014.**

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

Contents

Prerequisite and related information . . . vii

IBM Navigator for i	vii
Accessibility information for Content Manager OnDemand	vii

System requirements ix

Database concepts 1

System	1
Instance	1
Database.	1
Table	1
Index	2
Journals and journal receivers	2

Database maintenance 3

Before you begin	3
Expiring index data	3
How to expire index data	4
Migrating indexes	5
How to migrate indexes	6

Migrating and importing index data . . . 7

Configuring the system.	7
System Log messages	7
System Log user exit program	7
Archive Storage Manager	8
Storage sets.	8
Application groups	8
What happens when a user queries migrated data.	9
Message to the user	9
Message to the System Log	9
Importing index data	9
Run the STRIMPOND command	9
After you import index data.	10
Expiring imported migrated indexes	10
Configuring index migration	10

Document storage 11

Defining document storage management	11
Application groups	11
Disk storage manager	12
Archive storage manager	12
Migrating documents	13
Processes that should not be run simultaneously	15
Migration processing in the system log	15
Removing documents	15
Removing documents from disk storage.	16
Removing documents from archive storage.	17
Eliminating the need to run Disk Storage Manager (DSM)	17

Using a Network File System (NFS)

directory for document storage 21

Setting up an NFS disk pool.	21
Example scenario	22
On the Storage System (RDR400X).	22
On the Archive System (RDR400Y)	23
Special consideration for UIDs	24

Using Tivoli Storage Manager 25

Using Tivoli Storage Manager as a separate storage manager	25
Setting up Tivoli Storage Manager as a separate storage manager.	25
Use instructions	30
Using Tivoli Storage Manager as an ASM migration policy level	32
Setting up Tivoli Storage Manager as an ASM migration policy level	32
Use instructions	35

Backup and recovery 37

Backup considerations.	37
Recovery considerations	38
Reports.	38

Installing the administrative client . . . 41

System requirements	41
Installing the IBM Navigator for i interface.	41
Installing the OnDemand Administrator client.	42
Running Setup	42
To install on a PC	42
To install on a network file server	43
To use automated install	43
To uninstall	43

About the administrative client 45

Getting started	46
Using online help	46
Adding a server	46
Logging on a server	47
Changing passwords	47
System parameters	48
Maximum Password Age.	48
Minimum Password Length.	48
Inactivity Time Out.	49
System Logging	49
User Exit Logging	49
Login Processing	50
Annotations	51
System Log Comments	51
LDAP Authentication	51
Setting system parameters	52
Setting trace parameters	53
Adding items to a server	53

New command	53
Copy command	53
Export command	54
Drag and drop operation	54
Report Wizard	55
Starting the Report Wizard	55
Using the Report Wizard	56
Concepts.	59
Migration policies	59
Monitor definitions	60
Tape devices	61
Tape volumes.	61
Optical storage groups.	62
Optical volumes	62
Disk pool storage groups	62
Users	62
User types.	63
Authority	64
Groups	65
Printers.	65
Storage sets	66
Application groups	66
Applications	67
Folders	67
Cabinets	68
Holds	68
About application groups, applications, and folders	68
OnDemand permissions	70
Folder permissions	71
Application group permissions	71
Specifying permissions	71
Hints and tips	75
LDAP (Lightweight Directory Access Protocol) authentication support	77
LDAP authentication processes	78
OnDemand logon with the LDAP authentication and anonymous binding	79
OnDemand logon with the LDAP authentication and non-anonymous binding	80
Bypassing the LDAP server authentication	82
Sample LDAP configurations	83
Examples	85
System configuration	85
Creating monitor definitions.	86
Creating disk pools.	87
Creating optical storage groups.	87
Adding optical volumes	88
Adding tape devices	88
Adding tape volumes	89
Creating migration policies	89
Adding users.	90
Adding a user with the administrative client	90
Examples	91
Adding groups	93
Choose a server	93
Two ways to add a group	93
Adding users.	94

Adding the group	94
Examples	94
Adding server printers	95
Choose a server	96
Two ways to add a server printer	96
Adding the server printer	96
Examples	96
Adding storage sets	97
Adding a report	97
An example	98
Adding a field to an existing application group	108
Local server setup for offline administration	111

Loading spooled file data	113
Overview.	113
Preparing to load reports	114
Storage space	114
Defining the application group	114
Defining the application.	114
Loading reports	115
Running the ADDRPTOND command	115
Using an output queue monitor	115
Indexing reports	116
Processing the input data	116
Processing index data	117
Processing reports and resources	117
Loading index data	118
Loading storage objects	118
Disk storage	118
Archive storage.	118
Resources.	119
Verifying processing	119
Backing up databases.	120

Loading image files.	121
Overview.	121
Defining the application group	121
Database Organization	121
Expiration Type	122
Permissions	122
Field Definition.	122
Defining the application.	122
Application Group	123
Data Format.	123
Indexer	123
Data Compression.	123
Defining the folder	123
Application Group	123
Permissions	123
Field Definition.	124
Field Mapping	124
Accessing the image files	124
Creating index data	124
Configuring the ARSLOAD program	125
Processing the input data	126
Processing index data	126
Processing the image files	126
Verifying processing	127

Loading user-defined data 129

Defining the application group	129
Database organization	130
Expiration type	130
Permissions	130
Field definition	130
Defining the application	131
Application Group	131
Data Format	131
File Extension	131
Indexer	131
Data Compression	131
Defining the folder	131
Application Group	132
Permissions	132
Folder fields	132
Field Mapping	132
Accessing the input files	132
Creating the index data	133
Configuring the ARSLOAD program	133
Processing the input data	134
Processing index data	134
Processing the Lotus WordPro files	135
Verifying processing	135

Restarting a load process 137

Deleting a report 139

Managing the server 141

Errors and alerts	141
System logging facility	141
Searching for and viewing messages	141
System Log user exit	142
Monitoring users	143
Generating usage statistics	143
ARSSUPPORT utility	143
Finding or changing the server job and its attributes for a particular instance	145
Controlling the run priority of instance server jobs	146
Using Content Manager OnDemand data areas	147
Restarting journaling	149

Importing and exporting objects through batch administration 151

Overview	151
Installing batch administration	151
Prerequisites	151
Installation verification	152
Common problem during installation verification	153
Importing an XML file into an OnDemand system	153
Preparing an XML file for the import process	154
Creating an XML file	155
Importing the XML file by using the ARSXML command	163
Exporting OnDemand administrative objects to an XML file	163

Objects and data model used in the XML file 165

XML objects in the Content Manager OnDemand XML file	165
User object	168
Group object	171
Application object	176
Application group object	204
Storage set object	215
Folder object	216
Printer object	224
Cabinet object	228
Hold object	233
System Parameters object	237
Data field limitations	240

Command reference 243

Content Manager OnDemand server commands	243
ADDRPTOND	243
CHGPLDOND	243
CRTINSTOND	243
ENDMONOND	244
FNDKEYOND	244
MGRMEDRDAR	244
MRGSPLFOND	245
PRTRPTOND	245
PRTTXTOND	245
RMVRPTOND	245
STRASMOND	246
STRDSMOND	246
STRIMPOND	247
STRMONOND	247

System log messages 249

Overview	249
--------------------	-----

5250 host connection to client viewer 251

Operational and environmental considerations	251
Workstation installation tasks	251
Server configuration	253

Server printing and faxing 255

Using an IBM i printer file to define server print parameters	255
Server fax setup	256

API and user exit reference 259

API reference	259
Using quotes when executing the APIs	259
Using the QSHELL environment	259
Calling QSHELL commands from an IBM i command line	260
How to read a syntax diagram	261
ARSDATE	263
ARSDOC	267
ARSLOAD	294
ARSSLU	303
ARSXML	304

Saving Content Manager OnDemand passwords into encrypted files	314
Syntax of the ARSSTASH command	314
Examples of ARSSTASH commands	315
User exit reference.	315
Output queue or directory monitor user exit program	315
Facsimile user exit program	318
Automating ARSLOAD data loading	319
ARSLOAD	319

Automating ARSLOAD	319
Accessibility features	321
Notices	323
Trademarks	325
Index	327

Prerequisite and related information

Use the IBM® i documentation as your starting point for looking up IBM i technical information.

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

What's new?

You can see what is new for Content Manager OnDemand for i here:

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

Database concepts

System

A database *system* represents a physical machine that contains a copy of the database manager. In IBM Content Manager OnDemand, the *server* is the physical machine referred to as the *system*.

Instance

An IBM Content Manager OnDemand *instance* is a logical server environment consisting of a server and its own separate database and disk space. Each Content Manager OnDemand instance (the server, database and disk):

- Has its own definition of folders, application groups, applications and printers
- Must run in a single CCSID
- Has different security (users, groups, folder and application group permissions)
- Must have its name specified on commands if it is not the default instance
- Has its own System Log

Some reasons you might have multiple instances on a machine are:

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

For Content Manager OnDemand, IBM recommends that you name your primary production, or only, instance QUSROND. This is the default OnDemand instance that is used on all OnDemand commands unless you explicitly specify a different instance name. Using this as your primary OnDemand instance will save you from having to specify an instance name in most cases.

Database

A *database* is a collection of data that is stored in tables. In IBM Content Manager OnDemand, generally speaking, there are two types of tables:

- System tables, which contain information about the objects you define to the system, such as users, groups, application groups, applications, folders, storage sets, and printers. There are also system tables that contain information Content Manager OnDemand uses to control and maintain the system.
 - Application group tables, which contain the index data for the reports that you load on the system
-

Table

A *table* consists of data logically arranged in columns and rows. For example, when you create an application group, the system creates a table definition that contains one column for each field that you define. When you load a report into an application group, the system adds one row to an application group table for each document contained in the report.

Index

In IBM Content Manager OnDemand, an *index* is a key that points to a document. An index allows more efficient access to documents by creating a direct path to a document through pointers.

You define indexes when you create an application group. The indexes should contain information that uniquely identify a document, such as date, account number, and customer name. Indexes are populated by values extracted from a report when you load a report on the system. Each row in an application group table identifies one document.

However, keep in mind that you do not want lots of indexes on a report just to have indexes. You should have a good business reason to have an index. While indexes can help you find documents faster, having too many of them can slow you down when you load reports on the system. Every time you add a new row (document) to a table, you have to add a row to each and every one of the indexes for that table. So the more indexes that you have, the longer it may take when you load a report.

The SQL optimizer automatically chooses the most efficient way to access data in tables. The optimizer takes indexes into consideration when determining the fastest access path to data.

Journals and journal receivers

Each database includes recovery journals and journal receivers, which are used to recover from application or system errors. In combination with database backups, journals are used to recover the consistency of the database right up to a point in time when an error occurred.

All instances have journals associated with them. These journals keep records of database changes. Journals prevent a failure (system power, application error) from leaving a database in an inconsistent state. They restore the state of a database to the point before the change. Journals can also enable forward recovery to any point in time before the failure.

In planning for disaster recovery, be sure to remember that journals must be stored off site, or at least safely away from the disaster, in order to recover your database beyond the point of the last full, off line backup.

Database maintenance

You need to maintain the IBM Content Manager OnDemand database to keep it performing in an optimal manner. IBM recommends that you run the following database maintenance tasks on a regular basis:

- Expire index data that has reached its life of data and indexes period
- Migrate indexes to archive storage, if your organization needs to keep indexes for some period of time after users no longer need to retrieve the documents to which they point. This is not recommended.

Before you begin

IBM recommends that you do the following:

- Keep the index data for any given version of a report in the database on disk at least until such time that 99 percent of the requests for the report have passed. As a rule of thumb, if there is any chance that someone in your organization will need to retrieve a version of a report, keep the index data in the database – don't allow it to be migrated.
- Expire data periodically, and migrate index data only when absolutely necessary.
- Select Multiple Loads per Database Table when you define the Database Organization within your application group definitions. With this selection, each time that you load a report into an application group, IBM 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 (unless you select Single table for all loads). Content Manager OnDemand maintains a segment table for each instance. 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. This method is the default, and should be used in most cases.

Expiring index data

Indexes *expire* (are eligible for removal) because their life of data period has passed. The indexes, and the documents that they point to, can then be removed from the system. When you remove an index, information about the document to which it points is removed from the database (the document can no longer be retrieved). However, because indexes are eligible to be removed does not mean that they will be deleted from the database. IBM Content Manager OnDemand does not delete expired index data from the database until expiration processing runs.

The application group expiration policy determines when index data is eligible for deletion from the database. You define the expiration policy when you create the application group. The following properties on the Storage Management page comprise the expiration policy:

- Life of data and indexes. The length of time in days to maintain index data and documents on the system. After the index data is on the system for this number

of days, it is eligible to be deleted. The value of Life of Data and Indexes is set in the Storage Management tab of the Content Manager OnDemand application group definition. If you change this value after you have loaded data, the change affects the data that is already in Content Manager OnDemand as well as any new data loaded after the change is made.

Tip: If you specify Never Expire, then expiration processing is disabled for the application group. (That is, index data will not be removed from the database.)

- **Expiration Type.** Determines whether individual indexes or an entire table of index data is deleted at a time. When Content Manager OnDemand deletes index data, it either deletes a row (if the Expiration Type is Document), deletes all rows for the specific load (if the Expiration Type is Load), or deletes a table (if the Expiration Type is Segment). The amount of index data in a table and the number of reports the data represents is determined by the Database Organization. If the Database Organization is Multiple Loads per Database Table, then by default, a table of index data can hold up to 10 million indexes (unless you select Single table for all loads, in which case there is no maximum number of records for the index table). These types of tables usually hold the indexes for many reports. If the Database Organization is Single Load per Database Table, then each table holds the indexes for one and only one load.

A table of index data is not eligible to be deleted until the latest date in any of its rows reaches the Life of Data and Indexes period. For example, suppose that the Life Of Data and Indexes is set to 365 days, the Expiration Type is set to Segment, and the Database Organization is set to Multiple Loads per Database Table. By default, a table will contain approximately 10 million rows. Further, suppose that a report is loaded into the application group once every month and that each report adds one million rows to the database. Each table can hold the index data from approximately ten reports. Using these assumptions, the data that is loaded into the application group in January will not be eligible to be deleted by expiration processing until November of the following year. If you need to remove the index data for a report as soon as it reaches its Life of Data and Indexes period, then set the Database Organization to Single Load per Database Table and set the Expiration Type to Load. (And run expiration processing at least once a month.) You should consider selecting Single table for all loads if you have a small number of documents to be archived over time. You would not want to select Single table for all loads if your Expiration Type is Segment.

Content Manager OnDemand uses the application group's expiration policy to determine when indexes and documents expire and should be removed from the system. The archive storage manager marks documents for removal based on the Expire level specified in the migration policy. However, you should specify the same criteria to the disk storage manager and the archive storage manager. For more information on migrating and expiring documents, and recommendations for storage management criteria defined in your application groups, storage sets, and migration policies, see "Defining document storage management" on page 11.

How to expire index data

IBM Content Manager OnDemand does not delete expired index data from the database until expiration processing runs. The STRDSMOND command is the expiration utility. You can schedule the STRDSMOND command to run automatically or you can run it manually. You should make sure that the STRDSMOND command runs periodically so that Content Manager OnDemand deletes indexes when it is time to do so (so that expired documents can no longer be retrieved).

When the STRDSMOND command removes indexes, it saves a message similar to this in the System Log: 128 App1Grp Segment Expire (App1Grp) (Segment)

One message is saved in the System Log for each table that was deleted during expiration processing.

While not recommended, if you have migrated indexes to archive media, then the STRASMOND command will perform expiration processing on that index data.

Migrating indexes

This section provides an overview of the process of migrating index data from the database to archive storage. See “Migrating and importing index data” on page 7 for information about configuring the system for migration processing.

IBM Content Manager OnDemand provides automatic migration to move indexes from the database to archive storage to maintain seldom used indexes for long periods of time.

Important: If you use migration to move indexes to archive storage, make sure that you migrate them after there is no longer a need to retrieve the documents to which they point.

The STRDSMOND command uses an application group's migration policy to control when migration of indexes for an application group occurs:

- **Migration of Indexes.** If you specify No Migration, then migration of indexes is disabled for the application group. (That is, index data will not be migrated.) If you specify Migrate After *n* Days, then index data is eligible to be migrated after reaching the specified number of days. Indexes will be migrated the next time that the STRDSMOND command runs.
- **Life of Data and Indexes.** The length of time in days to maintain index data on the system. For migration, this value must be greater than the Migrate After *n* Days value. The value of Life of Data and Indexes is set in the Storage Management tab of the Content Manager OnDemand application group definition. If you change this value after you have loaded data, the change affects both the data that already exists in Content Manager OnDemand as well as any new data loaded after the change has been made.

Content Manager OnDemand does not migrate index data from the database to archive media until migration processing runs. The STRDSMOND command is the migration utility. You can control automatic migration processing by scheduling the command to run with the appropriate options. You can also manually start migration processing by running the command from the command line.

After a migrated table is successfully loaded into the System Migration application group, the table is deleted from the database. However, Content Manager OnDemand keeps track of all migrated tables. That way, if index data in a migrated table is needed, then Content Manager OnDemand can alert an administrator to take action (such as manually import the table back into the database).

How to migrate indexes

You can control automatic migration processing by scheduling the STRDSMOND command to run with the appropriate options. You can also manually start migration processing by running the STRDSMOND command from the command line.

When the STRDSMOND command migrates indexes, it saves the following messages in the System Log. A set of three messages should be saved in the System Log for each table that is migrated from the database to archive storage:

166 App1Group Segment Export (App1Grp) (Segment)

14 DB Info Exported (SQL Code)

87 App1Grp Load (System Migration)

The first message identifies a table of application group index data that is to be migrated from the database to archive storage. The second message reports the status of exporting the table from the database to temporary storage. The third message reports the loading of information about the migrated table into the System Migration application group. The System Migration application group must be assigned to a storage set that identifies an archive storage media type (such as optical or tape).

Migrating and importing index data

Index migration is the process by which IBM Content Manager OnDemand moves index data from the database to archive storage. This process optimizes database storage space while allowing you to maintain index data for a very long time. You typically migrate index data after users no longer need to access the information, but for legal or other business requirements, you still need to maintain the data for some number of years. If a user queries index data that has been migrated, an administrator must import a copy of the migrated table into the database. After maintaining the imported table in the database for a specified number of days, Content Manager OnDemand deletes it from the database.

This section provides information about importing index data into the database, including what happens when a user queries for migrated data, how to import the index table or tables required by the query, and what happens after you import a table into the database.

IBM assumes that an experienced Content Manager OnDemand administrator will use the information provided in this section. If you have questions about any of the topics in this section or if you would like help configuring your system to support migrating and importing of index data, contact the IBM support center.

In general, migrating index data is not recommended.

Configuring the system

There are a number of things you should consider regarding system configuration before you make any decisions about index migration.

System Log messages

IBM Content Manager OnDemand provides the System Log for administrators to monitor the system. When you install and configure Content Manager OnDemand, you initialize the System Log tables. The System Log is critical to the operation of the system.

When Content Manager OnDemand processes a query for application group indexes that have been migrated to archive storage, it saves a message in the System Log and sends a message to the System Log user exit program. A message is also sent to the QSYSOPR message queue.

You can configure the system to examine the messages that Content Manager OnDemand sends to the System Log user exit and mail them to an administrator or send them to another program. You can also configure the System Log user exit program to determine what action to take when a user queries for data that has been migrated to archive storage. See “System log messages” on page 249 for more information and a list of the most common System Log messages.

System Log user exit program

When a client queries index data that has been migrated to archive storage, IBM Content Manager OnDemand saves message number 168 ApplGrp Segment Not Available in the System Log. Content Manager OnDemand also sends the message to the System Log user exit program. If you have defined your own System Log

user exit program, then you can determine the action to take when Content Manager OnDemand sends the message to the System Log user exit program. For example, you may want the program to notify an administrator that a request for a table of migrated index data has occurred. See the *IBM Content Manager OnDemand for i: Common Server Planning and Installation Guide* for more information about the System Log user exit program.

Archive Storage Manager

Before IBM Content Manager OnDemand can migrate index data to archive storage, you must configure a migration policy with the information that the archive storage manager uses to maintain the data. The migration policy should maintain the data indefinitely. If you need the system to maintain a backup copy of the index data, then you should specify this in the migration policy.

Storage sets

IBM Content Manager OnDemand uses the System Migration application group to manage all index data that is migrated to archive storage. You must assign the System Migration application group to a storage set that identifies an archive storage media type.

Application groups

When you define an application group, you specify the storage management information that determines how long IBM Content Manager OnDemand maintains data stored in the application group and when Content Manager OnDemand takes certain actions. For example:

- **Life of Data and Indexes:** Determines the length of time that Content Manager OnDemand maintains index data and report data stored in the application group.
- **Migration of Indexes:** Determines the number of days before Content Manager OnDemand moves index data from the database to archive storage.

You should plan to migrate index data only after users no longer need to access the reports to which it refers. Only in exceptional situations should users need to access index data that has been migrated. If a user needs to access index data that has been migrated to archive media, the process of importing the table back into the database requires manual actions by an administrator, and usually results in a significant delay in completing the query. The import process also requires additional space in the database to hold the imported tables, additional log file storage, and temporary storage on the server to run the import process.

- **Keep Imported Migrated Indexes:** Determines how long that Content Manager OnDemand maintains the imported index data in the database before it is scheduled for deletion.

If you need to maintain index data in archive storage, then you must configure the Migration of Indexes in your application groups. You must specify the number of days to keep the index data on disk in Keep Imported Migrated Indexes. Content Manager OnDemand will schedule imported index data for deletion from the database after it resides in the database for the number of days specified in Keep Imported Migrated Indexes or Life of Data and Indexes, whichever occurs first.

You can use the administrative client to configure your application groups.

What happens when a user queries migrated data

There are several ways to be notified when a user requests queries migrated index data.

Message to the user

When the server determines that the index data required to complete a query has been migrated to archive storage, it sends a message to the client program. The message states that the data required to complete the query is not available and that the user should contact an administrator.

Message to the System Log

When IBM Content Manager OnDemand determines that the index data required to complete a query has been migrated to archive storage, it saves a message in the System Log. An administrator can open the System Log folder to search for and display messages in the System Log.

Content Manager OnDemand also sends a message to the QSYSOPR message queue and the System Log user exit program. You can configure the system to examine the message and send an alert to an administrator or call another program to take some action. See “System log messages” on page 249 for more information and a list of the most common System Log messages.

Importing index data

If index data is to be imported back from archive media, it is important to consider the following:

Verify database storage space

Importing migrated index data from archive storage back into the database requires additional database storage. Before you import the index data, you should verify that sufficient free space is available.

Verify database log file space

Importing migrated index data from archive storage back into the database requires database journal storage. Before you import the index data, you should verify that sufficient free space is available.

Run the STRIMPOND command

IBM Content Manager OnDemand provides the STRIMPOND command to import tables of migrated index data from archive storage back into the database. (The name of the application group and the index table to import can be obtained from the message that Content Manager OnDemand saved in the System Log.)

After the STRIMPOND command completes the import operation, you can open the System Log folder to see the messages that were generated by the import process. The messages will reference the ARSADMIN program name. See “System log messages” on page 249 for more information and a list of the most common System Log messages.

See online help for more information about the STRIMPOND command and its parameters.

After you import index data

Querying data

After you import index data from archive storage back into the database, you should notify the user to retry the query.

Expiring imported migrated indexes

IBM Content Manager OnDemand schedules an imported index table for deletion after it resides in the database for the number of days specified in the Length of Time to Keep Imported Indexes property in application groups. After an imported index data reaches the specified value, the next time that the STRDSMOND command runs, the imported index table is deleted from the database. (However, the table still exists in archive storage.)

You typically configure the STRDSMOND command to run automatically on a regular schedule. You can also run the STRDSMOND command manually.

Configuring index migration

If you find that your users are often querying for index data that has been migrated to archive storage, then IBM recommends that you configure your application groups to increase the length of time that IBM Content Manager OnDemand maintains the index data in the database. This should reduce the number of queries that need migrated index data.

Keeping imported migrated indexes

IBM Content Manager OnDemand schedules imported index data for deletion after the index data resides in the database for the number of days specified in Keep Imported Migrated Indexes or Life of Data and Indexes, whichever occurs first.

Document storage

Defining document storage management

The document storage management definitions determine where and when IBM Content Manager OnDemand stores documents and how it maintains them.

When you load a document into Content Manager OnDemand, you assign it to an application group. The application group is the last document storage management component that you define, because it requires storage set and migration policy definitions, which you must create first. The application group identifies the storage set and determines where documents should be loaded. You assign each application group to a storage set. You can load documents onto disk, onto archive media, or onto both disk storage and archive storage. The disk storage manager maintains documents on disk. The archive storage manager maintains documents on archive media. The archive storage manager uses a migration policy to determine where to store documents and how long to maintain them. After a document ages for the specified number of days, the migration process can move it from disk to archive storage.

This chapter refers to the Content Manager OnDemand Archive Storage Manager (ASM) as the storage manager for your Content Manager OnDemand data. However, Tivoli® Storage Manager can be enabled to be used in addition to, or in place of ASM on your IBM i server. See “Using Tivoli Storage Manager as a separate storage manager” on page 25 for more information on using the Tivoli Storage Manager.

Application groups

The application group is the last component that you must define because it requires storage set and migration policy definitions. The application group provides a way to group related documents. All documents in the application group are loaded on the media that is part of the storage set to which the application group is assigned. All documents in the application group migrate according to the rules that are defined for the application group's migration policy.

Use the administrative client to create the application groups that determine the document storage for your documents. You typically define one application group for each set of your documents that have similar storage requirements. For example, documents that must be retained for a specific length of time, in specific storage locations and stored on specific types of media.

Loading

An application group definition contains the rules for loading documents into an application group. It requires a storage set, which you must create first. The application group determines if documents are loaded onto disk, archive storage, or both. If the application group causes documents to be stored on archive storage, then the migration policy specifies when (or if) documents are copied to archive storage.

See the following properties on the Storage Management tab of the application group:

- **Storage Set Name:** Determines where documents will be loaded. Note that the storage set name will match its associated migration policy name (created using the Content Manager OnDemand component of IBM Navigator for i) except for cache or Tivoli Storage Manager storage nodes, which do not have associated migration policies.

Important: If you specify **Cache Only**, documents can only be loaded onto disk. This value cannot be changed later, so carefully consider possible future requirements before you select **Cache Only**. Unless you are certain that you will never want to migrate the data for this application group from disk (cache), a better choice might be to create your own migration policy or select another storage set from the pulldown list. Your new migration policy/storage set could be defined to use ASP01 (the system Auxiliary Storage Pool on disk on your IBM i system) as the first level of storage, and then later optical or tape could be added.

- **Cache Data:** Determines if documents will be loaded into disk storage. If the storage set is a cache-only storage set, documents must be loaded onto disk. For this reason, you cannot select **No** for **Cache Data** if the storage set is cache-only.
- **Migrate Data from Cache (on the Advanced panel):** If you specify **When Data is Loaded**, then documents will be loaded into archive storage.

Migrating

Migration is the process of copying documents from disk to archive storage as controlled by the rules of the application group's storage management criteria and migration policy. However, because a document is eligible to be migrated does not mean that it will be migrated. Other factors affect migration, such as the frequency with which you run migration processing. (Migration cannot take place until you run migration processing.)

The Storage Management tab of the application group and the application group's migration policy contain the rules for migrating the documents in an application group. They define how long a document stays on disk and, through the storage set and migration policy, where the document will be moved next. The migration policy level identifies the next location.

See the following settings on the Storage Management tab:

- **Storage Set Name:** Determines the next location for documents. If you specify **Cache Only**, then migration is disabled for the application group.
- **Migrate Data From Cache:** Determines when documents are eligible to be migrated. If you specify **When Data is Loaded**, then migration is done at the time the data is loaded. If you specify **No**, then migration is disabled for the application group.

Disk storage manager

The Disk Storage Manager (DSM, which is initiated by using the STRDSMOND command) maintains documents in cache (on disk, not in a disk pool). Documents migrate from disk storage to archive storage based on the migration policy that is defined for the application group. The disk storage manager can delete documents after they exceed the **Cache Data for *n* Days** or **Life of Data**, whichever occurs first. See "Removing documents" on page 15 for more information.

Archive storage manager

The archive storage manager (ASM, initiated by using the STRASMOND command) is the interface to the archive media (typically disk pools, Tivoli Storage

Manager, optical, or tape). The archive storage manager maintains a backup or long-term copy of documents. Before loading documents, you must define storage sets and migration policies, optical volumes, tape devices, and tape volumes. The archive storage manager can delete documents after they exceed the retention period specified in the migration policy levels and reach the expiration level. See “Removing documents” on page 15 for more information.

Migrating documents

IBM Content Manager OnDemand provides automatic migration to copy documents from disk storage to archive storage (for documents that were not copied to archive storage during the load process) and to make documents eligible for deletion to maintain free space on disk. Automatic migration is provided by using the Start Disk Storage Management (STRDSMOND) and Start Archived Storage Management (STRASMOND) commands. Migration helps to ensure that there is sufficient free space on disk, where faster response time can provide the most benefit to your users.

Important:

- You should run migration processing on a regular schedule to make sure that a backup copy of your documents gets created as soon as practically possible. If you defer the migration of documents to archive storage, and disk storage were to become corrupted, then you could be left without a copy of your documents.
- The STRASMOND command must only be run in batch (SBMJOB parameter set to *YES). Running this command interactively (with SBJOB(*NO)) may cause SQL errors.
- By default, the QUSROND default instance is used, and will produce the desired results for most systems. You can use an instance other than QUSROND as your default by defining the QDFTINST data area as described in “Using Content Manager OnDemand data areas” on page 147. You can also specify the instance name directly when you run the commands. If you need to run the STRASMOND command for multiple instances, you must issue the command separately for each instance. Note that if you initiate the archive storage manager by running the STRDSMOND command with RUNASM(*YES), then the instance name is passed from the disk storage manager and no further specifications are needed.
- If you run STRDSMOND for a specific application group (rather than the default of *ALL) and you set the Run ASM (RUNASM) parameter to *YES, be aware that ASM will run for ALL application groups, even though you have named a specific application group for DSM to use. You can, however, name a specific Policy for ASM to process, if desired. Also note that when you specify RUNASM(*YES), Content Manager OnDemand will initiate a separate batch job for ASM.
- If you specify Cache Data for 90 Days on the Storage Management tab within the application group, DSM will keep the data in the IBM i IFS directory for 90 days after the report date (a segment field) before it removes the data from the IFS CACHE directory. Regardless of the settings on the Storage Management tab of the application group definition, DSM will not delete data before it is sent to ASM. To determine when data is sent to ASM, select the **Advanced** button on the Storage Management tab within the application group. Data is passed to ASM based on the criteria specified in the Migrate Data from Cache section on the Advanced panel as shown in the table:

Table 1. Criteria specified in the Migrate Data from Cache section on the Advanced panel of the Storage Management tab

Criteria	Description
No	Data is never passed to the archive media. This option is only allowed if you specify a cache only Storage Set and is not recommended.
When Data is Loaded	Archive objects are passed to ASM when the store process runs from one of the load processes, such as ADDRPTOND, STRMONOND, arslload, arsdoc add.
Next Cache Migration	Archived objects are passed to ASM the next time STRDSMOND is run.
After xx Days in Cache	After reaching xx days, archived objects are passed to ASM the next time that STRDSMOND is run.

For a basic approach to the expiration of data, the Life of Data and Indexes should be the total of Cache Data days from the application group and the sum of the Duration at this level for all levels of the migration policy that is used for this application group. For example: The value for Cache Data days is 90 days, the migration policy for this application group has two levels, 100 days at the disk pool level and 7 years at the optical level so the Life of Data and Indexes value should be set to 2745 days.

You can, instead, take a more advanced approach to managing the expiration of your data. If you want to continue to use DSM to manage the expiration of your data, by using Life of Data and Indexes to control expiration, you should consider setting the duration of the last level of the migration policy to Never Expire. This allows controlled movement to a new level if one should be added in the future. If you want to manage the expiration of your data using ASM, using an expire level as the last level of the migration policy, you should consider setting the Life of Data and Indexes to Never Expire. This ensures that DSM will never expire the data. See “Eliminating the need to run Disk Storage Manager (DSM)” on page 17 for more information regarding Archived Storage Manager-based expiration.

After the data is sent to ASM and has entered a level as specified in the migration policy, the number of days at that level can only be changed using the Change Policy Level Date (CHGPLDOND) command for that particular data. If you change any of these values in the migration policy (instead of using the CHGPLDOND command), only newly archived documents are affected. All previously archived documents are unaffected.

You control automatic migration processing by scheduling the STRDSMOND and STRASMOND commands to run with the appropriate options. See your operating system information for details about how to schedule tasks. You can also start migration processing by running the command manually.

The STRDSMOND command uses an application group's storage management information to control when movement of data for an application group occurs:

- If you use Next Cache Migration to control when migration for an application group occurs, then the disk storage manager migrates data from cache each time that you start the STRDSMOND command with the appropriate options.

- If you use After xx Days in Cache to control when migration for an application group occurs, then a document must be stored in cache for at least the specified number of days before it is eligible to be migrated.

The disk space that migrated documents occupy can be reused after expiration processing completes. When you run migration processing, you should also run expiration processing so that the disk storage manager can reclaim the disk storage space occupied by migrated documents.

Processes that should not be run simultaneously

The following list identifies Common Server processes that should not be run at the same time:

- Do not run multiple Start Archived Storage Management (STRASMOND) commands against the same migration policy.
- Do not run multiple Start Disk Storage Management (STRDSMOND) commands.
- Do not run STRDSMOND while loading data using the following commands:
 - Add Report to OnDemand (ADDRPTOND)
 - Start Monitor for OnDemand (STRMONOND)
 - ARSLOAD API
 - ARSDOC add API
- Do not run STRDSMOND and STRASMOND at the same time. (ASM can start after DSM completes.)
- Do not perform system backups while any of the following processes are running:
 - ADDRPTOND, STRMONOND, ARSLOAD, ARSDOC add
 - STRDSMOND
 - STRASMOND
 - The instance server is running (for example, the QUSROND server job)

Migration processing in the system log

When you run the STRDSMOND command, it saves messages about its activities in the system log. The types of messages saved in the system log depend on the options that you specify when you run the STRDSMOND command and can be found by searching for the user ARSMAINT for the time that STRDSMOND was running. The number of messages saved in the system log during a migration process depends on the options that you specify for the STRDSMOND command and the number of application groups and segments of data processed. The viewable detail of message number 197: Cache Migration contains a list of each document that is migrated during the STRDSMOND process. Although you can run multiple STRDSMOND commands for different application groups within the same instance or different instances, it is not recommended due to possible locking issues.

Removing documents

Documents *expire* (are eligible for removal) because their disk expiration date or archive retention period has passed. Expired documents can then be removed by the storage managers. The disk storage manager identifies documents for removal by using the application group's expiration information. The archive storage manager marks documents for removal based on the criteria defined in the migration policy.

Documents expire from disk when they reach their disk expiration date. If a document's disk expiration date is less than its Life of Data period, then the document is simply removed from disk storage. Subsequent requests for the document are satisfied by the archive storage manager. When the document reaches its Life of Data period, information about it is removed from the IBM Content Manager OnDemand database (the document can no longer be retrieved). When the document's archive retention period has passed, information about it is removed from the archive storage manager database.

Because a document is eligible to be removed does not mean that it will be deleted from storage. The disk storage manager does not delete expired documents from storage until expiration processing runs. During expiration processing, the archive storage manager deletes information about expired documents from its database. However, the actual documents may remain on archive media until such time that the media on which they reside is reinitialized.

Important: The disk storage manager and the archive storage manager delete documents independently of each other. Each uses its own criteria to determine when documents expire and should be removed from the system. Each uses its own utilities to remove documents. However, for final removal of documents from the system, you should specify the same criteria to the disk storage manager and the archive storage manager. For more information on recommendations for storage management criteria defined in your application groups, storage sets, and migration policies, see “Migrating documents” on page 13.

Removing documents from disk storage

The options specified on the Storage Management tab when the application group is created determines when documents are eligible for deletion from disk. The options on the Storage Management tab are as follows:

- **Cache Data for n Days.** The length of time in days to keep documents on disk. After a document reaches this value, it is eligible to be deleted from disk.
- **Life of Data.** The length of time in days to maintain documents on the system. If you specify **Never Expire**, then expiration processing is disabled for the application group.
- **Expiration Type.** Determines whether one or more documents are eligible to be deleted at a time. For example, by default, the **Load** expiration type means that the set of documents that were loaded together will also expire together.

Alternative: **Segment** (in contrast to **Load**) is another possible choice for expiration type. This is the first time that a *segment* has been mentioned. Up to now, the term *documents* has been used, which is the data object that most people associate with the IBM Content Manager OnDemand system. However, administrators who maintain the system may also work with *segments*, which represent many documents, and *storage objects*, which are containers of compressed documents that are maintained by the storage managers.

The disk storage manager does not delete expired documents from disk until expiration processing runs. The STRDSMOND command is the expiration utility. You can schedule the STRDSMOND command to run automatically or you can run it manually. You should make sure that the STRDSMOND command runs periodically so that the disk storage manager can reclaim the space that is occupied by expired documents.

Expiration processing in the system log

When you run the STRDSMOND command, it saves messages about its activities in the system log. These messages can be found by searching on user ARSMAINT for the time period that STRDSMOND was running. The types of messages saved in the system log depend on the options that you specify when you run the STRDSMOND command. The number of messages saved in the system log each time that expiration processing runs depends on the options that you specify for the STRDSMOND command and the number of application groups and segments of data processed. The viewable detail of message number 196: Cache Expiration contains a list of each document that is expired during the STRDSMOND process.

Important: In addition to the messages in the system log, you should monitor your system every day for messages that indicate that your disk space is becoming full.

Removing documents from archive storage

Important: Removing a document from archive storage means that the backup or long-term copy of the document will be deleted from the system. You typically remove documents from archive storage when you no longer have a business or legal requirement to keep them.

A migration policy specifies the criteria that makes a document eligible for deletion. Documents become eligible for deletion under the following conditions:

- Administrators delete documents from archive media (using the Remove Report from OnDemand (RMVRPTOND) command)
- An archived document exceeds the time criteria defined in the Expire level of the migration policy (processed by the Start Archived Storage Mgmt (STRASMOND) command) or the Life of Data and Indexes (processed by the Start Disk Storage Management (STRDSMOND) command).

The storage manager does not delete information about expired documents from its database until expiration processing runs. You can run expiration processing either automatically or manually using the STRASMOND command (or you can use the STRDSMOND command with the Run ASM parameter set to *YES). You should make sure that expiration processing runs periodically to allow the archive storage manager to reuse storage space that is occupied by expired documents. When expiration processing runs, the archive storage manager deletes documents from its database. The storage space that these documents occupy then may become reusable.

Eliminating the need to run Disk Storage Manager (DSM)

OnDemand provides both Disk Storage Manager (DSM) and Archive Storage Manager (ASM) functions to facilitate a broad range of storage management options for your archived data. DSM manages the archived data stored in "cache" within OnDemand, which is typically used for short term storage. Most customers configure the system to maintain the most recent and frequently used versions of reports in cache storage. ASM manages the archived data that no longer resides in cache, which may include disk storage such as Auxiliary Storage Pools (ASPs) or Independent Auxiliary Storage Pools (IASPs) on your IBM i system, or Network File System (NFS) drives, as well as optical and tape media.

With Archive Storage Manager (ASM)-based expiration, it is possible to eliminate the need to run DSM. There are some advantages to doing this. For example, DSM

should not be run while loading data into OnDemand, so it may be difficult to find a large enough window of time for both loading data and running DSM. Another advantage is eliminating the need to run two processes to manage the migration and expiration of the data archived into OnDemand, namely DSM and ASM.

When choosing whether or not to implement ASM-based expiration processing, you should first understand the differences between how DSM expires data and how ASM expires data. DSM and ASM each use a different date to determine when the data should be expired. DSM uses the last segment date from the document. For example, if you have defined your segment date to be a statement date in the document and you are loading statements that are one year old, the segment date used to determine the life of the data would be a year prior to the date the data is being loaded. If you have set the Life of Data and Indexes to five years in your OnDemand application group definition, the data will actually expire four years after it is loaded into OnDemand, which is five years from the segment date. The same document, when using ASM-based expiration, would be expired five years after being loaded into OnDemand since the data is passed to ASM at the same time it is loaded and ASM uses the date it receives the data to determine the length of time until it is expired.

You should also be aware of the impact of changes made to the Life of Data and Indexes in an application group which is used by DSM to expire data and an expiration level in a migration policy which is used by ASM to expire data. A change to the Life of Data and Indexes will only impact the application group in which you are making the change. A change to the expiration level in the migration policy will affect all application groups that are using that particular migration policy.

To configure OnDemand to eliminate the need for running DSM for new application groups, you must first define your application groups to:

- not cache data
- migrate data at load time

You do this using the OnDemand Administrator client, on the Storage Management and Advanced Storage Management tabs of the application group definition. On the Storage Management tab, you must specify 'No' for the Cache Data option, and on the Advanced Storage Management tab, you must specify 'When Data is Loaded' for the Migrate Data from Cache option. Note that when you specify 'No' for the Cache Data option, the Migrate Data from Cache option is automatically changed to 'When Data is Loaded.' For any OnDemand migration policy specified in an application group definition that does not use cache, you should also define an expiration level as the last level of the policy if you wish to have your data expire after a designated length of time.

Important: It is not valid to have a migration policy with only one level, if the only level in the migration policy is an expiration level.

To configure OnDemand to eliminate the need for running DSM for *existing* application groups which already have data archived to them, you must first change your application groups to no longer cache the data. You must also specify to migrate data from cache at load time. You do this using the OnDemand Administrator client, on the Storage Management and Advanced Storage Management tabs of the application group definition. On the Storage Management tab, specify 'No' for the Cache Data option, and on the Advanced Storage Management tab, specify 'When Data is Loaded' for the Migrate Data from Cache

option. Note that when you specify 'No' for the Cache Data option, the Migrate Data from Cache option is automatically changed to 'When Data is Loaded.' Note also that the change to the Cache Data option will be retroactive for data that is already stored in cache. For example, if you originally specified 90 days for the Cache Data option and then you change this option to 'No,' all the data still in the CACHE directory for this application group will be removed the next time DSM is run. Note, if you archive AFP data, you may notice AFP resource objects remaining in the CACHE directory. This behavior is to be expected.

ASM-based expiration also requires that the IBM i user profile that owns the OnDemand instance library be added to the instance as a System Administrator. (The user profile of the instance owner has the same name as the instance, for example QUSROND.) You can add this user to the OnDemand instance using the OnDemand Administrator client. Note that new instances created after installing or upgrading to OnDemand 7.1 or higher will have this user profile added to the instance automatically when the instance is created. Instances that were created prior to OnDemand 7.1 will require that this user profile be added to the OnDemand instance manually.

If you make the changes described, you are no longer required to run DSM. As data is loaded into OnDemand, a copy of the data is placed in the ASMREQUEST directory. When ASM is run, it takes the data from the ASMREQUEST directory and processes it based on the information in the migration policy defined in the application group. ASM will manage the data for the life of the data specified by the levels defined in the migration policy. When the expiration level is reached, assuming one is specified in the migration policy, ASM will remove the data from the storage media and make a request to the server to also remove the index data. If there is no expiration level defined in the migration policy for the data, then the data will never be removed.

Important: Archive Storage Manager-based expiration is not compatible with the Hold and Release capabilities of OnDemand (known as enhanced retention management). This incompatibility is due to the fact that ASM does not check for holds before expiring documents.

Using a Network File System (NFS) directory for document storage

Network File System (NFS) is a distributed file system implementation providing remote, transparent access to files and directories. OnDemand can use an NFS exported directory as a disk pool for document storage.

A few key points regarding NFS include the following:

- You make a directory or an object available to NFS clients by exporting it. Therefore, you have very specific control over which parts of your system you will make available to NFS clients in your network.
- When you export, you can specify which clients have access to the objects. You identify a client by system name or IP address.
- You make a directory or an object available on your NFS client by mounting it.
- When you mount remote server file systems over local client directories, you allow IBM i servers to work with file systems that have been exported from a remote server. The mounted file systems will act and perform as if they exist on the local server.
- The NFS does not provide password protection. It is designed and intended for data sharing within a trusted community of systems. When a user requests access, the server receives the user's User ID number (UID). The UID is used to determine the permissions of the user.
- The use of synchronous writes is recommended. The use of asynchronous writes can result in data loss in the event of a system failure on either the archive or storage system.
- Each NFS mount (across all instances and logical partitions (LPARs)) must use a unique NFS directory.
- When using NFS with OnDemand, note that directories located in an Independent ASP (IASP) cannot be mounted over an NFS exported directory. This limitation affects only installations that locate a complete instance, including the instance library, within an IASP.

More information on NFS can be found within the IBM i information center at <http://www.ibm.com/systems/i/infocenter/>

Setting up an NFS disk pool

About this task

To setup an NFS disk pool to use for OnDemand document storage, follow these steps:

Procedure

1. Access IBM Navigator for i.
2. Click Content Manager OnDemand and verify the current instance.
3. Click **Network File System (NFS) Disk Pools**, and then click **Add**.
4. Enter the pool number (2 through 32), NFS host name and path name, and pool type (primary or backup) for which the NFS disk pool storage group should be defined. For example, you might enter 7 for the pool number, specify host

name MYSYSTEM.MYCOMPANY.COM, /NFSSTG/YQUSROND/PRIMARY as the path name, and select **Primary** for the pool type. Note that if you use an existing pool number, its contents will not be accessible after the file system is mounted. You should use a pool number that is not already in use by a physical ASP on your system.

5. Click **OK** to create the NFS disk pool definition.

Example scenario

In our example NFS scenario, we will use two IBM i servers. One server is called the Archive System, where OnDemand is running and where the NFS directory is mounted. The other server is called the Storage System, where the NFS directory is exported, and where OnDemand stores objects.

Archive System – RDR400Y - OnDemand is running here. Instance name is QUSROND. NFS directory is mounted.

Storage System – RDR400X - NFS directory is exported. OnDemand objects are stored here.

On the Storage System (RDR400X)

In this example, we need to create a user profile on the Storage System with the same UID as the instance profile on the Archive System. Objects created on the Storage System will be owned by this user profile. On the Archive System, QUSROND has a UID of 588. You can use the Create User Profile (CRTUSRPRF) command as shown to create the companion user profile on the Storage System. Note that we specify the same language parameters as used on the Archive System.

```
CRTUSRPRF USRPRF(YQUSROND) USRCLS(*PGMR) INLPGM(*NONE) INLMNU(*SIGNOFF)
LMTCPB(*NO) TEXT('ONDEMAND NFS MOUNT PROFILE') SPCAUT(*IOSYSCFG
*JOBCTL *SAVSYS) GRPPRF(QRDARS400) LANGID(ENU) CNTRYID(US) CCSID(37)
CHRIDCTL(*JOBCCSID) SETJOBATR(*CCSID *DATFMT *DATSEP *DECfmt *SRTSEQ *TIMSEP)
LOCALE('/QSYS.LIB/EN_US.LOCALE') UID(588)
```

If the needed UID is already in use, see “Special consideration for UIDs” on page 24 for information on changing the UID of an existing user profile.

We then create the following directories and subdirectories in IFS on the Storage System for use with NFS:

```
/NFSSTG/YQUSROND/PRIMARY
/NFSSTG/YQUSROND/BACKUP
```

To automatically export the file systems when the NFS server is started, update the /etc/EXPORTS file in IFS on the Storage System (which is ABC in our example).

The /etc/EXPORTS file for our example would look like the following:

```
/NFSSTG/YQUSROND/PRIMARY root=ABC.EXAMPLE.COM, \
access=ABC.EXAMPLE.COM, \
rw=ABC.EXAMPLE.COM \
#HOSTOPT HostName=ABC.EXAMPLE.COM, \
NoWaitForWrites
/NFSSTG/YQUSROND/BACKUP root=ABC.EXAMPLE.COM, \
access=ABC.EXAMPLE.COM, \
rw=ABC.EXAMPLE.COM \
#HOSTOPT HostName=ABC.EXAMPLE.COM, \
NoWaitForWrites
```

Start the NFS servers using the Start NFS Server command:

```
STRNFSSVR SERVER(*ALL)
```

The job log contains these messages:

```
Start of NFS server daemon or daemons of type *RPC was successful.
Start of NFS server daemon or daemons of type *BIO was successful.
2 entries exported, 0 entries not exported.
Start of NFS server daemon or daemons of type *SVR was successful.
Start of NFS server daemon or daemons of type *MNT was successful.
Start of NFS server daemon or daemons of type *NSM was successful.
Start of NFS server daemon or daemons of type *NLM was successful.
Start NFS server command completed successfully.
```

Notes[®]:

- The user starting NFS servers must have input/output (I/O) system configuration (*IOSYSCFG) special authority to use this command.
- The user starting NFS servers must be enrolled in the system distribution directory. Use the Add Directory Entry (ADDDIRE) command to enroll the user.

To determine if an NFS server is running, use the Work with Active Jobs (WRKACTJOB) command and look in the subsystem QSYSWRK for existence of the following jobs:

- QNFSRPCD - the RPCBind daemon
- QNFSBIOD - the block I/O (BIO) daemon
- QNFSNFSD - the NFS server (SVR) daemon
- QNFSMNTD - the mount (MNT) daemon
- QNFSNSMD - the network status monitor (NSM) daemon
- QNFSNLMD - the network lock manager (NLM) daemon

If necessary, you can manually export or unexport the directories. To export all entries in /etc/EXPORTS:

```
EXPORTFS OPTIONS('-A')
```

To unexport all entries in /etc/EXPORTS:

```
EXPORTFS OPTIONS('-A -U')
```

On the Archive System (RDR400Y)

To setup a primary NFS disk pool to use for OnDemand on the Archive System, follow the steps as described in “Setting up an NFS disk pool” on page 21 using the values as shown in the example, such as QUSROND for instance name, 7 for ASP, and so on.

Then, create a backup disk pool using the same values as the primary disk pool, with the exception of selecting Backup instead of Primary for disk pool type.

Next, create or update an OnDemand migration policy to use the NFS disk pools. Migration policies are also created by using IBM Navigator for i. More information on migration policies can be found in “Migration policies” on page 59.

Typically, the NFS mounts are done automatically, when required by the OnDemand Archive Storage Manager (ASM) to migrate or retrieve data. Message queue QSYSOPR will contain message CPCA1B0 when the NFS mount is performed by ASM. For example:

File system /NFSSTG/YQUSROND/PRIMARY was successfully mounted over directory /QIBM/USERDATA/ONDEMAND/QUSROND/ASMASP07/PRIMARY.

If you need to manually mount the NFS, use the MOUNT command. The mount must be done exactly as it was defined when you created your NFS disk pool. For example:

```
MOUNT TYPE(*NFS) MFS('rdr400x.rtp.raleigh.ibm.com:/NFSSTG/YQUSROND/PRIMARY')
MNTOVRDIR('/QIBM/UserData/OnDemand/QUSROND/ASMASP07/PRIMARY')
```

If you need to manually unmount the NFS, use the UNMOUNT command, for example:

```
UNMOUNT TYPE(*NFS)
MNTOVRDIR('/QIBM/UserData/OnDemand/QUSROND/ASMASP07/PRIMARY')
```

To verify that the NFS is mounted, use either of the STATFS or DSPMFSINF commands.

```
STATFS OBJ('/QIBM/UserData/OnDemand/QUSROND/ASMASP07/PRIMARY')
```

Output from the STATFS and DSPMFSINF commands will be similar to the following:

```
Display Mounted FS Information
Object . . . . . : /QIBM/UserData/OnDemand/QUSROND/ASMASP07/PRIMARY
File system type . . . . . : Network File System (NFS)
(If File system type . . . . . : "root" (/) then the NFS is not mounted)
...
Path of mounted file system . . : rdr400x.rtp.raleigh.ibm.com:/NFSSTG/YQUSROND/PRIMARY
Path mounted over . . . . . : /QIBM/UserData/OnDemand/QUSROND/ASMASP07/PRIMARY
Protection . . . . . : Read-write
...
```

Note:

- The NFS servers do not need to be running on the Archive System.

Special consideration for UIDs

You might discover that you need to change UIDs, even for IBM-supplied user profiles, to have compatibility with other systems in your network. When you change the UID for a user profile, you also need to change the UID for all the objects that the profile owns in either the root directory or the QOpenSrv directory.

An API is available to make it simpler to change the UID for a user profile. The QSYCHGID API automatically changes the UID in both the user profile and all the owned objects. Source code examples for a sample program and sample command using the QSYCHGID API are provided in the QSAMPLES2 source file in the QUSRRDARS library. Source member CHGUID with a Type of CLP contains a sample CL program that calls the QSYCHGID API. Source member CHGUID with a Type of CMD contains sample command source that runs the CHGUID CL program. An example of how to create the CL program or command is included in the comments section at the top of each sample source member.

Using Tivoli Storage Manager

By default, Archive Storage Manager (ASM) is the standard (and only) storage manager for OnDemand on an IBM i server. You can also use IBM Tivoli Storage Manager (TSM) in addition to ASM or in place of ASM, or simply as another media choice for use in ASM migration policies. Use the instructions found in “Using Tivoli Storage Manager as a separate storage manager” or “Using Tivoli Storage Manager as an ASM migration policy level” on page 32 to set up and use IBM Tivoli Storage Manager in your OnDemand environment.

Using Tivoli Storage Manager as a separate storage manager

You can use IBM Tivoli Storage Manager (TSM) in addition to ASM or in place of ASM using the instructions in this section. Using this configuration, data that is already stored into ASM cannot be migrated to Tivoli Storage Manager, nor can data in Tivoli Storage Manager be migrated to ASM. A second option is to simply use IBM Tivoli Storage Manager as another media choice within a migration policy for use with ASM. For information on that configuration, see “Using Tivoli Storage Manager as an ASM migration policy level” on page 32.

Setting up Tivoli Storage Manager as a separate storage manager

About this task

To install and configure IBM Tivoli Storage Manager as a separate storage manager, follow these general steps:

Procedure

1. Install the Tivoli Storage Manager APIs on the IBM i system.
2. Update your Content Manager OnDemand `ars.cfg` file.
3. Create or update your `dsm.opt` file. IBM recommends that you create a `dsm.opt` file for each instance. This method provides maximum flexibility in configuring TSM support.
4. Create or update your `dsm.sys` file.
5. Configure Content Manager OnDemand to use Tivoli Storage Manager by doing one of the following:
 - Create a storage set and storage node for Tivoli Storage Manager.
 - Update an existing storage set with a new storage node (if you plan to change an existing application group from using ASM to Tivoli Storage Manager).
6. Verify your Tivoli Storage Manager storage set and storage node.

Install the Tivoli Storage Manager APIs on your IBM i system

About this task

To install the IBM Tivoli Storage Manager APIs:

Procedure

1. Go to `ftp://ftp.software.ibm.com/storage/tivoli-storage-management/maintenance/client/v6r1/OS400`.

2. Open the folder of the highest release level listed (v615 at the time this document was created).
3. Copy the file 6.1.5.0-TIV-TSMBAC-OS400.exe to your workstation.
4. Read and follow the installation instructions in the file named 6.1.5.0-TIV-TSMBAC-OS400.README_api_enu.htm on the FTP site.

Update your OnDemand ars.cfg file

Add the ARS_STORAGE_MANAGER=TSM entry:

About this task

IBM Tivoli Storage Manager support is enabled by adding specific lines to the ars.cfg file. The ars.cfg file is located in the IFS on your IBM i server in the /QIBM/UserData/OnDemand/*instancename* directory, where *instancename* is the name of the Content Manager OnDemand instance for which you want to enable Tivoli Storage Manager support.

Procedure

1. Modify the ars.cfg file by adding these lines:

```
#
# TIVOLI STORAGE MANAGER (TSM) ENABLEMENT
ARS_STORAGE_MANAGER=TSM
```

2. Save the file.

Results

When the Content Manager OnDemand server detects this entry, it enables both Tivoli Storage Manager and ASM as storage managers for this particular Content Manager OnDemand instance.

Add the DSMI_DIR=, DSMI_CONFIG=, and DSMI_LOG= entries:

About this task

When IBM Tivoli Storage Manager is enabled, the server looks for particular configuration files.

To tell Tivoli Storage Manager where these configuration files reside:

Procedure

1. Open the ars.cfg file located in /QIBM/UserData/OnDemand/*instance_name*, where *instance_name* is the name of the OnDemand instance for which you are modifying the ars.cfg file.
2. Add the following lines:

```
#
# TIVOLI STORAGE MANAGER (TSM) CONFIGURATION FILE LOCATIONS
DSMI_DIR=/QIBM/UserData/Tivoli/TSM/Client/API/bin
DSMI_CONFIG=/QIBM/UserData/OnDemand/QUSROND/dsm.opt
DSMI_LOG=/QIBM/UserData/OnDemand/QUSROND/TMP
```

where *QUSROND* might need to be replaced by your instance name, if you are not using the default QUSROND instance. The DSMI_DIR entry specifies the directory that contains the TSM API files. The directory must be set to /QIBM/UserData/Tivoli/TSM/Client/API/bin. The DSMI_CONFIG entry specifies the full path name of the TSM API options file. The DSMI_LOG entry

specifies the directory in which TSM stores the TSM API error log. You should copy the above lines verbatim, with no modifications except for the instance name.

What to do next

After you have updated the `ars.cfg` file, stop and restart your instance for the changes to take effect.

Create or update your `dsm.opt` file

To enable IBM Tivoli Storage Manager support, a file named `dsm.opt` must exist, and it must contain the name of a server that is listed in your `dsm.sys` file. The `dsm.opt` file must exist in IFS on your IBM i server in the `/usr/tivoli/tsm/client/api/bin` directory. This path is specified in the `ars.cfg` file, as the value of the `DSMI_CONFIG` entry.

IBM recommends creating a separate `dsm.opt` file for each instance configured for use with TSM. In our example, for instance `QUSROND`, the `dsm.opt` file is located at `/QIBM/UserData/OnDemand/QUSROND`. This path is the path specified in the `DSMI_CONFIG=` entry you just added to the `ars.cfg` file in the previous step.

The following example shows a `dsm.opt` file, where the Tivoli Storage Manager server name is `TMSERVER`. Note that this example can also be found on your IBM i server in the `/QIBM/UserData/Tivoli/TSM/Client/API/bin` directory in the file `dsm.opt.smp`. You can copy the `dsm.opt.smp` file to a file named `dsm.opt` in the same directory to help ensure the `dsm.opt` file is correct.

```
*****
* Tivoli Storage Manager                                     *
*                                                         *
* Client User Options file for OS/400 (dsm.opt)           *
*****
* This file contains an option you can use to specify the TSM
* server to contact if more than one is defined in your client
* system options file (dsm.sys).
*****

SErvername      TMSERVER
```

Create or update your `dsm.sys` file

To enable IBM Tivoli Storage Manager support, a file named `dsm.sys` must also exist, and it must contain a section for each Tivoli Storage Manager server that Content Manager OnDemand will communicate with. The server name specified in the first line of the section must match the name of a server that is listed in your `dsm.opt` file. The `dsm.sys` file must exist in the `/QIBM/UserData/Tivoli/TSM/Client/API/bin` directory on your IBM i server. The directory path was specified in the `DSMI_DIR=` entry that was added to the `ars.cfg` file in a previous step.

The following example shows a `dsm.sys` file, in which the Tivoli Storage Manager server name is `TMSERVER`. Note that this example can also be found on your IBM i server in the `/QIBM/UserData/Tivoli/TSM/Client/API/bin` directory in the file `dsm.sys.smp`. You can copy the sample `dsm.sys.smp` file to a file named `dsm.sys` (in the same directory) to help ensure the `dsm.sys` file you create is correct.

```
*****
* Tivoli Storage Manager                                     *
*                                                         *
* Client System Options file for OS/400 (dsm.sys)         *
*****
* This file contains the minimum options required to get started
```

* using TSM. Copy dsm.sys.smp to dsm.sys. In the dsm.sys file,
 * enter the appropriate values for each option listed below.
 * If your client node communicates with multiple TSM servers, be
 * sure to add a stanza, beginning with the SERVERNAME option, for
 * each additional server.

```
Servername          TMSERVER
COMMethod          TCPip
TCPPort            1500
TCPServeraddress   tsmserver.company.com
COMPRESSION        OFF
ENABLEARCHIVERETENTIONPROTECTION  YES
```

Enter these lines as shown, with no modifications, except for the following:

TMSERVER

Tivoli Storage Manager Server alias. Use this value or replace it with a name of your choice. The name does not need to match any particular part of your Content Manager OnDemand or Tivoli Storage Manager configuration, but it must match the value used in the dsm.opt file.

tsmserver.company.com

Replace this with the host name or TCP/IP address of your Tivoli Storage Manager server.

Configure OnDemand to use Tivoli Storage Manager

Controlling which storage manager an application group uses is determined by the storage set (and the corresponding storage nodes) associated with the application group. By design, you cannot use the IBM Content Manager OnDemand Administrator Client to create or maintain storage sets and storage nodes. This use of the administrator client is disabled for IBM i servers, because the storage sets and storage nodes needed for ASM are automatically added by the IBM Navigator for i Content Manager OnDemand component when you create a migration policy. However, storage sets and storage nodes for use with IBM Tivoli Storage Manager must be added or modified by using the Content Manager OnDemand ARSXML API running within the QSHHELL environment on your i server.

Depending on your requirements, you can choose to create a new storage set and corresponding storage node for Tivoli Storage Manager, or you can choose to update an existing storage set with a new storage node for Tivoli Storage Manager. Only one of the two sets of instructions needs to be performed, although you can choose to do both if you have requirements to use both new and existing storage sets for Tivoli Storage Manager.

Create the OnDemand storage set and storage node:

Use this procedure if you plan to create a new storage set and storage node to use with a new application group.

This is an example of an XML file (that might be named addtsmstorageset.xml as shown in the example ARSXML API call) that could be used as input to the ARSXML API to create the new storage set and storage node:

```
*****Beginning of data*****
<?xml version="1.0" encoding="UTF-8" ?>
<onDemand xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xsi:noNamespaceSchemaLocation="/QIBM/ProdData/OnDemand/bin/xml/ondemand.xsd">
  <storageSet name="TSMSET" description="Tivoli Storage Manager" >
  <node name="TSMNODE" description="TSM Node" server="*ONDEMAND"
    logon="tsmnode" password="password" loadData="true" accessMethod="TSM" />
```

```

    </storageSet>
  </onDemand>
  *****End of Data*****

```

This is an example of the ARSXML API call (run within QSHELL on your IBM i server) to add a storage set and storage node for IBM Tivoli Storage Manager:

```
arsxml add -h QUSROND -i /home/dbryant/addtmsstorageset.xml -v
```

This is an example of the messages resulting from the ARSXML API call:

```

ondemandTitleStr=[QRDARS/ARSXML add -h QUSROND -i
/home/dbryant/addtmsstorageset.xml -v ]
command=[java -Djava.library.path=/qsys.lib/qrdars.lib -cp
"/QIBM/PRODDATA/ONDEMAND/BIN/xml/ODAdmin.jar:/usr/xerces-2_9_1/xercesI
mpl.jar:/usr/xerces-2_9_1/xml-apis.jar" com.ibm.cm.od.ArsXMLbat add
-h QUSROND -i /home/dbryant/addtmsstorageset.xml -v ]
Starting arxml. Version: 8.4.0.3
Command Line: arxml add -h QUSROND -i /home/dbryant/addtmsstorageset.xml -v
10/07/09 14:21:16: Attempting login for userid '' on server 'QUSROND' ...
10/07/09 14:21:16: Login successful
10/07/09 14:21:19: Adding storageSet, TSMSET
10/07/09 14:21:19: Add of storageSet, TSMSET was successful.
10/07/09 14:21:19: Adding storageSet-node, TSMSET-TSMNODE
10/07/09 14:21:19: Add of storageSet-node, TSMSET-TSMNODE was successful.
10/07/09 14:21:19: Finished processing file /home/dbryant/addtmsstorageset.xml.

```

For more information about the ARSXML API and invoking Content Manager OnDemand APIs within QSHELL, see "API and user exit reference" on page 259.

Update an existing storage set with a new storage node:

Use this procedure if you plan to change an existing application group from using ASM to using IBM Tivoli Storage Manager.

If you have an existing application group that is currently using ASM, and you want to change it to use Tivoli Storage Manager, you can update the existing storage group (that is named in the application group definition) to add a Tivoli Storage Manager storage node, and then change the storage node designated for loading data to point to the new Tivoli Storage Manager storage node.

This is an example of an XML file (that might be named addtmsnode.xml as shown in the example ARSXML API call) that could be used as input to the ARSXML API to create the storage node for an existing storage set named LONGTERM, and to designate the new node to be used for loading data, instead of the node also named LONGTERM:

```

*****Beginning of data*****
<?xml version="1.0" encoding="UTF-8" ?>
<onDemand xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xsi:noNamespaceSchemaLocation="/QIBM/ProdData/OnDemand/bin/xml/ondemand.xsd">
  <storageSet name="LONGTERM" >
    <node task="update" name="LONGTERM" loadData="false" logon=" " password=" " /
    <node task="add" name="TSMNODE" description="TSM Node" server="*ONDEMAND"
      logon="tsmnode" password="password" loadData="true" accessMethod="TSM" />
  </storageSet>
</onDemand>
*****End of Data*****

```

This is an example of the ARSXML API call (run within QSHELL on your IBM i server) to create the storage node for IBM Tivoli Storage Manager:

```
arsxml update -h QUSROND -i /home/dbryant/addtmsnode.xml -v
```

This is an example of the messages resulting from the ARSXML API call:

```

ondemandTitleStr=[QRDARS/ARXML update -h QUSROND -e c -i
/home/dbryant/addtsmnode.xml -v ]
command=[java -Djava.library.path=/qsys.lib/qrdars.lib -cp
"/QIBM/PRODDATA/ONDEMAND/BIN/xml/ODAdmin.jar:/usr/xerces-2_9_1/xercesI
mp1.jar:/usr/xerces-2_9_1/xml-apis.jar" com.ibm.cm.od.ArsXMLbat update
-h QUSROND -e c -i /home/dbryant/addtsmnode.xml -v ]
Starting arxml. Version: 8.4.1.2
Command Line: arxml update -h QUSROND -e c -i /home/dbryant/addtsmnode.xml -v
10/07/09 14:59:37: Attempting login for userid '' on server 'QUSROND' ...
10/07/09 14:59:37: Login successful
10/07/09 14:59:39: Updating storageSet, LONGTERM
10/07/09 14:59:39: Updating storageSet-node, LONGTERM-LONGTERM
10/07/09 14:59:40: Update of storageSet-node, LONGTERM-LONGTERM was successful.
10/07/09 14:59:39: Adding storageSet-node, LONGTERM-TSMNODE
10/07/09 14:59:40: Add of storageSet-node, LONGTERM-TSMNODE was successful.
10/07/09 14:59:40: Update of storageSet, LONGTERM was successful.
10/07/09 14:59:40: Finished processing file /home/dbryant/addtsmnode.xml.

```

For more information about the ARXML API and invoking Content Manager OnDemand APIs within QSHELL, see “API and user exit reference” on page 259.

Verify your Tivoli Storage Manager storage set and storage node

After you have successfully run the ARXML API, you can use the OnDemand Administrator client to confirm the correct setup by viewing the storage set definition you are using for IBM Tivoli Storage Manager. In the Storage Nodes list box, nodes that are used to load data are marked with an asterisk (*).

Use instructions

IBM Tivoli Storage Manager processing is invoked in one of two ways:

- When data is loaded into Content Manager OnDemand (if Migrate Data From Cache is set to **When Data is Loaded** on the Advanced panel of the Storage Management tab within the application group definition). This process is invoked either by issuing the Add Report to OnDemand (ADDRPTOND) command, or by issuing the Start Monitor for OnDemand (STRMONOND) command.
- As part of the OnDemand Disk Storage Management (DSM) process (if Migrate Data From Cache is set to **After x Days in Cache** on the Advanced panel of the Storage Management tab within the application group definition). This process is invoked by issuing the Start Disk Storage Management (STRDSMOND) command.

No special parameters for data loading or the STRDSMOND command are required to trigger the Tivoli Storage Manager step. By configuring Tivoli Storage Manager support as described earlier, data loading and DSM will automatically invoke Tivoli Storage Manager as needed during processing. The Archive Storage Management process (ASM, started by the Start Archived Storage Mgmt (STRASMOND) command or the RUNASM(*YES) parameter of the STRDSMOND command) is not required unless other Content Manager OnDemand data is managed by ASM in addition to the Content Manager OnDemand data managed by Tivoli Storage Manager.

Checking objects in Tivoli Storage Manager

After you start storing data, you can use the following commands to check the status of objects archived to IBM Tivoli Storage Manager. For more information about administering a Tivoli Storage Manager server, see the Tivoli Storage Manager documentation at <http://publib.boulder.ibm.com/infocenter/tivihelp/v1r1/index.jsp>.

Integrated Solutions Console:

The IBM Integrated Solutions Console (ISC) with the Tivoli Storage Manager Administration Center can be used to check statistics, such as capacity and percent utilized. For information on installing and using the Administration Center, see the Tivoli Storage Manager documentation at <http://publib.boulder.ibm.com/infocenter/tivihelp/v1r1/index.jsp>.

Command Line: About this task

You can also use the IBM Tivoli Storage Manager Administrative Client Command Line (**dsmadm**) to check statistics. To install the Administrative Client Command Line files, you must select **Custom Setup** when installing the Tivoli Storage Manager clients. From the Windows **Start** menu, select **All Programs > Tivoli Storage Manager > Administrative Command Line**. Log onto the Tivoli Storage Manager server and enter your commands.

Object naming on the Tivoli Storage Manager server

Objects sent to IBM Tivoli Storage Manager use the following naming convention. (Note that objects sent to IBM Tivoli Storage Manager by ASM using a IBM Tivoli Storage Manager migration policy level have a different naming convention.)

Item	Example	Comments
FILESPACE_NAME:	/ONDGY/ CGA	/instance name/AGID_Name (unless instance is QUSROND, then the instance name is omitted)
HL_NAME:	/DOC	always /DOC
LL_NAME:	/399FAA1	object name (which is the name created during the load process)

Additional information

- If Migrate Data from Cache is set to **When Data is Loaded** on the Advanced panel of the Storage Management tab in the Application group definition, the load will fail if IBM Tivoli Storage Manager is not running, or the communication connection cannot be made with the Tivoli Storage Manager server. The system log message 88 will contain entries similar to this:

```
Connection cannot be established for the >tmsserver.company.com< server
Unable to store the object >10FAAA<. Object size 40034
Connection cannot be established for the >tmsserver.company.com< server
Unable to unload data from OnDemand - LoadId(5048-14-0-10FAAA-11303-11303)
  Rows Deleted (0)
05/30/07 13:25:43 -- Unloading of data failed
```
- If Migrate Data from Cache is set to **Next Cache Migration** on the Advanced panel of the Storage Management tab in the Application group definition, you must run DSM to migrate the data to Tivoli Storage Manager.
- If the password of the Tivoli Storage Manager node expires, attempting to migrate data to Tivoli Storage Manager generates system log message 20 with entries similar to this:

```
SM Error: ANS1352E (RS52) The session is rejected. Your password has expired.,
RC=52, Reason=0, File=arssms.C, Line=531 Svr- >ondemand.server.company.com
10.44.122.55<-
```
- If the Tivoli Storage Manager node name is spelled incorrectly when the storage set is created, attempting to migrate data to Tivoli Storage Manager generates system log message 20 with entries similar to this:

```
SM Error: ANS1352E (RC53) Session rejected: Unknown or incorrect ID entered,
RC=53, Reason=0, File=arssmsms.C, Line=531, Srvr- >rdr400x.rtp.raleigh.ibm.com
9.42.125.55<-
```

- After modifying the ars.cfg file to add Tivoli Storage Manager support, if the dsm.sys or dsm.opt files can not be found when the instance is started, the instance job will end with an OND1002 message in the job log. Message data will be similar to this:

```
SM Error: ANS1087E (RC106) Access to the specified file or directory
is denied , RC=106, Reason=0, File=arssmsms.C, Line=1042, Srvr-
>rdr400x.rtp.raleigh.ibm.com
```

- If the server specified in the dsm.opt file on the SErvername entry is not found in the dsm.sys file when the instance is started, the instance job will end with an OND1002 message in the job log. Message data will be similar to this:

```
SM Error: ANS1217E (RC409) Server name not found in System Options
File , RC=409, Reason=0, File=arssmsms.C, Line=1042, Srvr-
>rdr400x.rtp.raleigh.ibm.com
```

Using Tivoli Storage Manager as an ASM migration policy level

Follow these instructions to use IBM Tivoli Storage Manager as another media choice within a migration policy for use with ASM. Using this configuration, IBM Tivoli Storage Manager can be used in new or existing migration policies. Existing data can be migrated to IBM Tivoli Storage Manager. Aggregation can also be used for IBM Tivoli Storage Manager data. Dual write to two different IBM Tivoli Storage Manager storage groups is also supported. The two IBM Tivoli Storage Manager storage groups can be located on two physically separated IBM Tivoli Storage Manager servers.

Another alternative is to use IBM Tivoli Storage Manager (TSM) as a separate storage manager, in addition to ASM or in place of ASM on your IBM i server. For information on that configuration, see “Using Tivoli Storage Manager as a separate storage manager” on page 25.

Setting up Tivoli Storage Manager as an ASM migration policy level

About this task

To install and configure IBM Tivoli Storage Manager as an ASM migration policy level, follow these general steps:

Procedure

1. Install the Tivoli Storage Manager APIs on the IBM i system.
2. Create or update your dsm.opt file. IBM recommends that you create a dsm.opt file for each instance. This method provides maximum flexibility in configuring TSM support.
3. Create or update your dsm.sys file.
4. Create a Tivoli Storage Manager storage group using IBM Navigator for i.
5. Create a new, or update an existing, ASM migration policy to use the new Tivoli Storage Manager storage group.

Install the Tivoli Storage Manager APIs on your IBM i system

About this task

To install the IBM Tivoli Storage Manager APIs:


```

* Client System Options file for OS/400 (dsm.sys)
*****
* This file contains the minimum options required to get started
* using TSM. Copy dsm.sys.smp to dsm.sys. In the dsm.sys file,
* enter the appropriate values for each option listed below.
* If your client node communicates with multiple TSM servers, be
* sure to add a stanza, beginning with the SERVERNAME option, for
* each additional server.
*****

```

```

Servername          TSMSERVER
COMMethod           TCPip
TCPPort             1500
TCPServeraddress    tsmserver.company.com
COMPRESSION         OFF
ENABLEARCHIVERETENTIONPROTECTION YES

```

It is recommended that you enter these lines as shown, with no modifications, except for the following:

TSMSERVER

Tivoli Storage Manager Server alias. Use this value or replace it with a name of your choice. The name does not need to match any particular part of your Content Manager OnDemand or Tivoli Storage Manager configuration, but it must match the value used in the dsm.opt file.

tsmserver.company.com

Replace this with the host name or TCP/IP address of your Tivoli Storage Manager server.

Create a Tivoli Storage Manager storage node

Use the Tivoli Storage Manager (TSM) Storage Nodes option of the Content Manager OnDemand component of IBM Navigator for i to create a Tivoli Storage Manager storage node. Click Tivoli Storage Manager Storage Nodes and then click Add to display the Tivoli Storage Manager (TSM) Storage Node Definition panel. On this panel, you specify the name of the Client node, which is the Tivoli Storage Manager node name created for use with Content Manager OnDemand. You also specify the Password, which is the Tivoli Storage Manager node password used by the Tivoli Storage Manager APIs to sign on to Tivoli Storage Manager. The Directory code location must be the path to the Tivoli Storage Manager APIs. The /QIBM/USERDATA/TIVOLI/TSM/CLIENT/API/BIN default value is the correct path and must not be changed. The Options and Log file locations are filled in based on the instance name but can be changed if, for example, you want to use one dsm.opt file for all instances (which is not recommended).

Create a new or update an existing migration policy

After you have created a Content Manager OnDemand Tivoli Storage Manager storage node, you need to create or update a Content Manager OnDemand migration policy to use it. To do this, use the Content Manager OnDemand component of IBM Navigator for i. Whether you are creating a new migration policy or updating an existing one, click to select an existing level if appropriate, and then click the **Add before** or **Add after** button in the bottom section of the first migration policy panel. Select **Tivoli Storage Manager** for the **Primary media type**. Then select your new Tivoli Storage Manager storage node from the pull-down for **Primary storage group**.

If you add a Tivoli Storage Manager level to an existing migration policy, you can use the Change Policy Level Date (CHGPLDOND) command to change the next level date for existing archives. Then run ASM to move existing data to Tivoli Storage Manager.

Use instructions

IBM Tivoli Storage Manager processing is invoked in one of two ways:

- When data is loaded into Content Manager OnDemand (if Migrate Data From Cache is set to **When Data is Loaded** on the Advanced panel of the Storage Management tab within the application group definition). This process is invoked either by issuing the Add Report to OnDemand (ADDRPTOND) command, or by issuing the Start Monitor for OnDemand (STRMONOND) command.
- As part of the OnDemand Disk Storage Management (DSM) process (if Migrate Data From Cache is set to **After x Days in Cache** on the Advanced panel of the Storage Management tab within the application group definition). This process is invoked by issuing the Start Disk Storage Management (STRDSMOND) command.

No special parameters for data loading or the STRDSMOND command are required to trigger the Tivoli Storage Manager step. By configuring Tivoli Storage Manager support as described earlier, data loading and DSM automatically invokes Tivoli Storage Manager as needed during processing. The Archive Storage Management process (ASM, started by the Start Archived Storage Mgmt (STRASMOND) command or the RUNASM(*YES) parameter of the STRDSMOND command) is not required unless other Content Manager OnDemand data is managed by ASM in addition to the Content Manager OnDemand data managed by Tivoli Storage Manager.

Testing your Tivoli Storage Manager setup

To test your Tivoli Storage Manager setup, follow these steps:

1. Archive data.
2. Run Disk Storage Manager (DSM) if Migrate data is set to "Next Cache Migration."
3. Run ASM. You may need to run ASM more than once if using aggregation.
4. Check the objects created, using the information provided in "Object naming on the Tivoli Storage Manager server" to identify the objects.

Object naming on the Tivoli Storage Manager server

Objects sent to Tivoli Storage Manager when using Tivoli Storage Manager as an ASM migration policy level use the following naming convention. Note that objects sent to Tivoli Storage Manager when using Tivoli Storage Manager as a separate storage manager have a different naming convention.

Item	Example	Comments
FILESPACE_NAME:	/QUSROND/ TSMSTGGRP	/instance name/Tivoli Storage Manager storage group name
HL_NAME:	/Primary/20080222	/Tivoli Storage Manager storage group type/date object was moved to Tivoli Storage Manager
LL_NAME:	/117	object number in ASM

Backup and recovery

It is important to implement a backup plan to protect your data archives in the event your organization experiences a system failure. You need to have a plan and prepare so that you can recover.

Backup considerations

Items you need to consider when planning your regular backups include:

- How to recover the optical or tape media itself
You should consider having a copy of the media stored at an offsite location.
- How to recover the data on disk
You should have procedures in place that backup all of the IBM Content Manager OnDemand data that normally resides on disk, as well as any stored data that has not yet migrated to optical or tape.

Remember: Disk backups will contain database files that reflect the actual archive data location and status at the time of backup. DO NOT restore from a backup that is one week old. Archived data might have migrated to optical or tape, but database control files identify incorrectly the archived data location as disk. In order to minimize this situation, you should save your database libraries and save your Content Manager OnDemand Integrated File System (IFS) directories on the same schedule to keep them synchronized. Perform the backups at least as frequently as you run the disk storage management (STRDSMOND) or archive storage management (STRASMOND) commands. (Possibly even more frequently if you often manually delete reports that are stored in Content Manager OnDemand.)

The following objects need to be saved by the appropriate IBM i commands:

- The Content Manager OnDemand licensed program (5770-RD1)
- User profiles QRDARS400, QRDARS4001, QRDARS4002, QRDARS4003, QRDARS4004, QRDARS4005, QONDADM, and QRDARSADM, and one for each instance that exists (which has the same name as the instance)
- Authorization lists QRDARS400, QONDADM, QRDARSADM, and one for each instance that exists (which has the same name as the instance)
- All objects in QUSRRDARS library
- All objects in any other libraries that contain data for your OnDemand instances. For example, the QUSROND library if you are using the default QUSROND instance, or any other library by the name of any other instance on your system.
- Output queues that are being monitored and any output queues that are being used as processed queues or error queues for your monitors.
These output queues can contain processed spooled files.
- Integrated File System directories

Each instance that is created in Content Manager OnDemand has an Integrated File System (IFS) directory that is named the same as the instance. To back up only the Content Manager OnDemand data on disk, you could back up all items found in the Integrated File System directory `/QIBM/UserData/OnDemand/instance`, where *instance* is the name of your OnDemand instance. Remember that the name of the default instance is QUSROND.

Content Manager OnDemand Common Server migration policies allow you to specify a disk pool as one of the possible storage levels (as an alternative to optical or tape media, for example). If you use disk pools in any of your migration policy storage levels, the simplest approach to make sure that the data in the disk pool is saved and restored properly for an instance is to do the following:

To perform the save:

1. End the instance server `ENDTCPSVR SERVER(*ONDM) INSTANCE(QUSROND)`
2. Unmount the file system `CALL QRDARS/QRLCASMUFS PARM('QUSROND')`
3. Save the directory structure `/dev/QASP01/ONDEMAND_QUSROND*`
4. Save the directory structure `/QIBM/UserData/OnDemand/QUSROND`
5. Restart the instance server `STRTCPSVR SERVER(*ONDM) INSTANCE(QUSROND)`

To perform the restore, if necessary:

1. End the instance server `ENDTCPSVR SERVER(*ONDM) INSTANCE(QUSROND)`
2. Unmount the file system `CALL QRDARS/QRLCASMUFS PARM('QUSROND')`
3. Restore the directory structure `/dev/QASP01/ONDEMAND_QUSROND*`
4. Restore the directory structure `/QIBM/UserData/OnDemand/QUSROND`
5. Restart the instance server `STRTCPSVR SERVER(*ONDM) INSTANCE(QUSROND)`

Note that the instance name (such as the default QUSROND instance in the example) and Auxiliary Storage Pool name (such as ASP01 in the example) must be changed to match your actual environment. Also note that you must name a specific instance when you call the QRLCASMUFS program. The QRLCASMUFS program does not support a value of *ALL for the instance name.

- Integrated File System directories for Content Manager OnDemand Web Enablement Kit (ODWEK)

If you have implemented ODWEK you should also back up all items that are found in the Integrated File System directory `/QIBM/UserData/OnDemand/WWW`.

- Symbolic links in IFS directory `/usr/bin` that start with ARS
- Configuration files

See Chapters 8 and 9 of the *IBM Content Manager OnDemand for i Common Server Planning and Installation Guide* for more information and a list of configuration files.

Recovery considerations

Contact your IBM Content Manager OnDemand support provider for instructions on recovering your OnDemand archives in the event of a disaster. Many factors can influence the recovery plan, depending on the frequency and extent of the backups you have available.

Reports

IBM Content Manager OnDemand can store copies of reports and resources on disk storage 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. The reports in archive storage can also be used as backup copies in the event that disk storage becomes corrupted or unavailable. Archive storage consists of optical or tape storage volumes.

Most customers configure the system to copy reports to disk storage and archive media 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 deleted from disk storage or if the copy on disk storage is unavailable. However, you must configure the system to maintain multiple copies of reports. You configure Content Manager OnDemand to use archive storage by defining migration policies and storage sets that identify archive storage, assigning application groups to the storage sets, and configuring other storage management parameters in application groups.

Tip: If you do not plan to copy reports to archive storage, then take regular backups of the data on disk. However, if a media failure occurs or disk storage becomes corrupted, users cannot retrieve reports until the backups are restored.

Installing the administrative client

System requirements

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

Installing the IBM Navigator for i interface

When you install IBM Content Manager OnDemand for i Version 7 Release 2, all the necessary software is included to support the Content Manager OnDemand component of IBM Navigator for i. IBM Navigator for i allows you to use an internet browser to perform many system functions for your IBM i system. No additional steps are required to install the Content Manager OnDemand component of IBM Navigator for i. Refer to the IBM Navigator for i documentation for additional information.

The Content Manager OnDemand component of IBM Navigator for i is required to work with the following Content Manager OnDemand objects:

- Disk Pools
- Network File System (NFS) Disk Pools
- Tivoli Storage Manager (TSM) Storage Nodes
- Optical Storage Groups
- Optical Volumes
- Tape Devices
- Tape Volumes
- Migration Policies
- Directory Monitors
- Output Queue Monitors

To start IBM Navigator for i, start a browser on your workstation and enter a URL similar to the following: `http://system_name:2001`

where *system_name* is the name of your IBM i server on which Content Manager OnDemand is installed, such as `mysystem.mycompany.com`.

You should delete the IBM Navigator for i log files periodically if you do not need the information in them for problem determination of any of the IBM Navigator for i components. To delete the log files, you must first end the HTTP Admin server. Do this at a time when the HTTP Admin server is not required for any other activity on your IBM i system.

To end the HTTP Admin server, enter the End TCP/IP Server command:

```
ENDTCPSVR SERVER(*HTTP) HTTPSVR(*ADMIN)
```

Then, delete the `lwistart.txt`, `lwistderr.txt`, and `lwistdout.txt` log files in the `/QIBM/UserData/OS/OSGi/LWISysInst/admin2/lwi/logs` directory on your IBM i system.

Then, to restart the HTTP Admin server, enter the Start TCP/IP Server command:

```
STRTCPSVR SERVER(*HTTP) HTTPSVR(*ADMIN)
```

To assist with problem determination, you can create a data area named QRLWTRACE to control verbose message logging for the Content Manager OnDemand component of IBM Navigator for i. Verbose message logging causes additional messages to be written to the `lwistdout.txt` log file. See "Using Content Manager OnDemand data areas" on page 147 for details about creating this data area. Remember to turn verbose message logging off when you no longer need it.

The Content Manager OnDemand component of IBM Navigator for i replaces the Content Manager OnDemand Archive plug-in of System i Navigator and the Content Manager OnDemand Administration component of IBM Navigator for i that were available in previous releases.

Installing the OnDemand Administrator client

The Setup program installs the IBM Content Manager OnDemand client software on the PC.

To download the Content Manager OnDemand clients, see: <http://www.ibm.com/support/docview.wss?rs=129&uid=swg21450570>, or search for 1450570 at <http://www.ibm.com>.

You can install all of the Content Manager OnDemand features at once, or individual features as you need them.

Running Setup

When you run the Setup program, the Setup screens show the names of the IBM Content Manager OnDemand directories so you know where the files are being placed.

Information about network installations is available in the *IBM Content Manager OnDemand: Windows Client Customization Guide*.

To install on a PC

About this task

To install from the CD-ROM or to install from a network file server, follow these instructions:

Procedure

1. Insert the CD-ROM in the appropriate drive or obtain the drive letter of the network drive on which the IBM Content Manager OnDemand software is located on the network file server.
2. From the Windows taskbar, click **Start**, and then choose **Run**.
3. Type `x:\client\win32\setup` (where *x* is the letter of your CD-ROM drive or the network drive).
4. Click **OK**.
5. After the Setup program starts, click **Next** to continue.
6. Read and accept the license agreement, and then click **Next** to continue.

7. Click **Next** to accept the destination drive and directory, or specify a destination and then click **Next**.
8. Select the Custom setup type and then click **Next** to continue.
9. Select **OnDemand Administrator** and select your language. (You can also select other components that you want to install on the PC at this time.)
10. Click **Next** to continue. Follow the instructions on the screen to complete the installation.

To install on a network file server

Information about network installations is available in the *IBM Content Manager OnDemand: Windows Client Customization Guide*.

To use automated install

Automated install allows administrators to standardize the IBM Content Manager OnDemand installation for all users in an organization. It also allows administrators to install Content Manager OnDemand clients without the presence of users at their PCs. To read more about automated install, see the *IBM Content Manager OnDemand: Windows Client Customization Guide*.

To uninstall

You can use **Add/Remove Programs** in **Control Panel** to remove the IBM Content Manager OnDemand client from your PC. **Add/Remove Programs** removes the directories for Content Manager OnDemand and any references to Content Manager OnDemand in system files.

To run the uninstall Procedure

1. From the Windows taskbar, click **Start**. Choose **Settings**, and then choose **Control Panel**.
2. Double-click **Add/Remove Programs**.
3. From the list, select **OnDemand32**.
4. Click **Add/Remove**.
5. Click **OK**.

About the administrative client

The administrative functions in IBM Content Manager OnDemand are set up using the Content Manager OnDemand Administrator client and IBM Navigator for i.

The Content Manager OnDemand Administrator client functions include:

- Defining reports to the system
- Adding and maintaining Content Manager OnDemand users and groups
- Adding and maintaining server printers
- Creating summaries about users, groups, applications, application groups, storage sets, folders, cabinets, holds, and printers
- Adding and maintaining servers
- Setting system parameters for Content Manager OnDemand servers and client programs
- Copying items from one Content Manager OnDemand server to another
- Tracking changes made to the system. When you add to or update the database, Content Manager OnDemand places a document in the system log which shows the changes that you made.

The items that are maintained by using the Content Manager OnDemand Administrator client are:

- Users
- Groups
- Printers
- Storage sets
- Application groups
- Applications
- Folders
- Cabinets
- Holds

The items that are maintained by using the Content Manager OnDemand component of IBM Navigator for i functions include:

- Migration policies
- Output queue monitors
- Directory monitors
- Disk pools
- Network File System (NFS) disk pools
- Tivoli Storage Manager (TSM) storage nodes
- Optical storage groups
- Optical volumes
- Tape devices
- Tape volumes

Getting started

- You can collapse and expand areas in the navigator pane (on the left) to make it easier to see. A plus sign next to an area means there are items inside.
- When you click on an area in the navigator pane, the items appear in the list pane (on the right).
- To make the panes narrower or wider, point to the vertical bar between the two panes of the window until the pointer turns into a two-headed arrow. Then click and hold the left mouse button and drag it in either direction.
- Use buttons on the toolbar to switch between the different ways to look at items: large or small icons, a list, or details.
- To query the server for a new list of items, press the **F5** key or select **Refresh List** from the **View** menu.
- After you log on to a server, the status bar shows the OnDemand user ID and the name and version number of the IBM Content Manager OnDemand server.

Using online help

Online help provides information to assist you with completing tasks. The administrative client online help contains information about the options, fields, and commands on the windows, dialog boxes, and property sheets that you see when using the program.

To display the online help, press the **F1** key any time that the administrative client is active in Windows. Help is available for dialog box commands and options. The main help topic for each dialog box usually contains information about the kinds of tasks you can perform. For example, the online help about Logical Views lets you learn how to create public and private logical views. The online help provides brief procedures rather than lengthy descriptions.

Adding a server

About this task

You can use the **New Server** command to add a server.

Procedure

1. Start the **OnDemand Administrator** client.
2. Click once on **OnDemand Servers** at the top of the list of servers in the left panel.
3. From the **File** menu, select **New Server** to open the **Add a Server** dialog box.
4. Type the name of the server in the **Server** field. The server name identifies the server in the navigator pane of the administrator window. You can use an alias, the actual computer or network name of the server, or any other identifier you choose. By default, the administrative client copies what you type to the **Host Name** field.
5. Verify the value of the **Host Name** field. (By default, the **Host Name** field contains the same value as the **Server** field.) The host name can be a host name alias, fully-qualified host name, or IP address of the server.
6. Verify the Protocol. Choose from TCP/IP or Local:
 - TCP/IP. Use TCP/IP (Transmission Control Protocol/Internet Protocol) as the network protocol. To use TCP/IP, the server and the client must include TCP/IP in the protocol stack.

- Local. Use the local file system manager to communicate with the server. This is most often used when your archived data is on a CD-ROM.
7. For TCP/IP, verify the Port Number that the server monitors for client program requests.
The default value, 0 (zero), means that the server monitors the port number of the Content Manager OnDemand TCP/IP service. By default, the server monitors port number 1445. If you plan to use a port number other than 1445, then you must enter a valid port number. If you have more than one instance defined to Content Manager OnDemand, then you would have separate server definitions for every instance with unique port numbers for each instance. The value range is from 0 to 65535. To see what ports are currently in use on your system, enter the Work with TCP/IP Network Status (WRKTCPSTS) command with OPTION(*CNN).
 8. If your logon panel includes the **Attempt Unified Logon** checkbox, confirm that it is not checked. (Unified Logon is used only for Windows servers.)
 9. If you are adding a Local server, specify a Directory and select an Operating System and a Database. See the online help for assistance.
 10. Click **OK** to add the server.

Logging on a server

About this task

Procedure

1. Point to the server and double-click the left mouse button to open the **Logon** dialog box.
2. Enter your OnDemand user ID and password in the spaces provided and click **OK**. For a Local server, the built-in user ID is admin; no password is assigned to the admin user ID.

Changing passwords

This section applies only if you are using IBM Content Manager OnDemand user IDs and passwords rather than IBM i user IDs and passwords. See the *IBM Content Manager OnDemand for i Common Server Planning and Installation Guide* section titled "OnDemand user relationship to IBM i user profiles" for important information about passwords.

Changing a password

About this task

To change a user's password:

Procedure

1. Select and expand the library server.
2. Select **Users**.
3. In the User ID list, point to the user ID and click the right mouse button.
4. From the pop-up menu, select **Update** to open the **Update a User** dialog box.
5. Type the new password in the **Password** field.

Important: When creating a password, the value that you specify can be a maximum of 20 characters. However, the password authentication that is built into IBM Content Manager OnDemand verifies only the first eight characters

that are entered by the user. The additional characters are provided for customers that choose to implement their own password security by using the logon user exit.

6. Verify the new password by retyping it into the **Verify Password** field.
7. Click **OK**. Content Manager OnDemand updates the database and returns to the main window.

Changing an expired password

IBM Content Manager OnDemand provides password expiration processing to help you manage security on the system. You can set a value that represents the time in days that passwords assigned to users remain valid. After a user's password reaches the value that you specify, the user must change the password. See "System parameters" for information about how Content Manager OnDemand expires passwords.

After a password reaches the expiration value, the next time the user ID is used to log on to a server, Content Manager OnDemand prompts the user to enter a new password.

The user must enter the current password for the user ID, a new password, and verify the new password by retyping the new password

System parameters

IBM Content Manager OnDemand system parameters allow you to establish the following operational settings for client programs and servers.

Maximum Password Age

Sets a time limit for passwords and determines when IBM Content Manager OnDemand prompts users to change passwords. The default setting is **Password Never Expires**, meaning that passwords do not expire and Content Manager OnDemand never prompts users to change passwords. If you select **Password Always Expires**, then users must change to new passwords each time that they log on to a server. To set a specific time limit for passwords, select **Expires In __ Days** and enter the number of days that passwords are valid in the space provided. The value can be from 1 (one) to 365.

If you are linking your Content Manager OnDemand user IDs to IBM i user profiles, then **Maximum Password Age** should be set to **Password Never Expires**. If you specify a value for **Maximum Password Age**, then Content Manager OnDemand may force a user to change their password before it is required by i.

Minimum Password Length

Determines whether passwords are required. If passwords are required, also determines the fewest number of characters that passwords can contain. The default value is At Least 8 Characters.

If you select **Permit Blank Password**, meaning that passwords are not required, then the valid password length is 0 (zero) to 20 characters.

If you are linking your IBM Content Manager OnDemand user IDs to IBM i user profiles, then **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 i to use its own rules.

To set a specific minimum password length, select **At Least __ Characters** and enter a number in the space provided. The value can be from 1 (one) to 20 characters.

Important: When creating a password, the value that you specify can be a maximum of 20 characters. 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 that 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.

Inactivity Time Out

Determines when IBM Content Manager OnDemand terminates sessions between inactive clients and the server. The default setting is Time Out in 60 Minutes. The value can be from 1 (one) to 1440 (24 hours).

The period of inactivity is measured between requests to a server. For example, when a user enters a query, Content Manager OnDemand searches the database and builds the document list. This completes a request to the server. If the user does not work with the items in the document list, open another folder, or invoke another query before the inactivity time out occurs, Content Manager OnDemand automatically terminates the session with the client.

Use caution when you set the inactivity time out. For example, assume that you set the inactivity time out to 10 (ten). You log on to Content Manager OnDemand to add an application group. Creating the application group takes you 15 minutes to complete. After entering all of the information about the application group, you click OK to create the application group. Content Manager OnDemand issues a message that shows a time out has occurred. You must logoff the server, and you cannot save the information you entered about the application group.

System Logging

Determines the messages that IBM Content Manager OnDemand saves in the system log. Content Manager OnDemand provides the system log to help you track activity and monitor the system. Content Manager OnDemand saves messages that are generated by the various commands, such as the ADDRPTOND command. Content Manager OnDemand can save a message in the system log when the following events occur:

- A user logs on to the system
- A user logs off the system
- A user logon fails
- Application group data is queried, retrieved, loaded, updated, deleted, or maintained

User Exit Logging

IBM Content Manager OnDemand provides a user exit at each of the four system log event points. These exits allow you to filter the messages that Content Manager OnDemand saves in the system log and take action when a particular event occurs. For example, you can provide a user exit program that sends a message to a security administrator when someone attempts and fails to log on to the system. You can also use a user exit to determine what information appears in the system log.

If you plan to migrate index data to archive storage, then we recommend that you configure the system to save application group messages in the system log and send them to a system log user exit program. You should design a system log user exit program to notify an administrator when a query for migrated data occurs. Before a query for migrated data can be completed, an administrator must import a copy of the table or tables that are required from archive storage to the database.

See the *IBM Content Manager OnDemand for i Common Server Planning and Installation Guide* for help with configuring the system log user exit.

Login Processing

The login processing system parameters allow you to specify whether user IDs and passwords are case sensitive.

Before continuing with this section, please refer to the OnDemand user ID relationship to IBM i user profiles section in the *IBM Content Manager OnDemand for i Common Server Planning and Installation Guide* for a detailed explanation of the significance of choosing to relate your Content Manager OnDemand users to existing i user profiles. It is important that you understand that concept before you make your choices for Login processing.

Now that you understand the relationship between Content Manager OnDemand users and your i users, you should note the following:

- If your Content Manager OnDemand user IDs are linked to your i user profiles (which is the default when Content Manager OnDemand is installed), and if your i security level is set to 0 or 1, you **SHOULD NOT** check the Password Case Sensitive checkbox.
- If your Content Manager OnDemand user IDs are linked to your i user profiles (which is the default), and if your i security level is set to 2 or 3, you **SHOULD** check the **Password Case Sensitive** checkbox.
- Regardless of your system security level, you should **NEVER** check the user ID Case Sensitive checkbox if your Content Manager OnDemand users are linked to your i user profiles.
- If you are using Content Manager OnDemand user IDs and passwords that are not linked to i user profiles (which is not the default), then you can set the two Login processing checkboxes as you choose.

By default, user IDs and passwords are case insensitive. When you add a user, Content Manager OnDemand converts lowercase letters in the user ID to uppercase. A person can type letters in a user ID in uppercase, lowercase, or mixed case letters. For example, if you add the user ID LaGuarde, a person can enter LAGUARDE, laguarde, or LaGuarde to log on to the server.

If you select **UserID Case Sensitive**, then a user must type the user ID exactly as it was entered when the user was added. For example, if you add the user ID LaGuarde, then the user must enter LaGuarde to log on to the server.

If you select **Password Case Sensitive**, then a user must type the password exactly as it was entered when the user was added. For example, if you set the password to Spring2Far, then the user must enter Spring2Far to log on to the system.

We strongly encourage you to decide whether you want user IDs and passwords to be case sensitive when you install the system, change the defaults if necessary, and do not change the settings again. Otherwise:

- If user IDs are initially case insensitive and you later choose **UserID Case Sensitive**, then user IDs that were added before you changed the parameter must be entered in uppercase. The same is true for passwords.
- If user IDs are initially case sensitive and you later clear **UserID Case Sensitive**, then the user IDs that were added before you changed the parameter that contain mixed or lowercase letters will no longer be valid. The same is true for passwords.

Annotations

This section specifies which types of annotations (referred to as "notes" in the OnDemand client) can be added by a user. This selection applies to all users with authority to add annotations in the system.

There are three types of annotations that a user can add:

Allow Public

Allows the user to add public annotations. Public annotations to a document can be viewed by anyone who opens that document.

Allow Private to User

Allows the user to add private annotations to a document, and those annotations can be viewed only by the user that created the note, application group administrators, and system administrators.

Allow Private to Group

Allows the user to add annotations to a document, and those annotations can be viewed only by a specific group of users.

The Default Annotation Type section specifies the annotation that is selected as the default.

System Log Comments

The system log comments specify whether the administrative client displays the System Log Comments window when you perform an add, update, or delete operation.

Enable comments

Select this option to display the System Log Comments window when you perform an add, update, or delete operation.

Require comments

This option requires the user to enter one or more characters in the Comments field.

LDAP Authentication

Specify whether you want to use LDAP (Lightweight Directory Access Protocol) authentication in your OnDemand server.

Select the **Enable** check box to use LDAP authentication. After LDAP authentication is enabled, OnDemand server makes an authentication request to the LDAP server every time it receives a login request from the client, and processes the client request only after the user information is verified by the LDAP server.

Clear the **Enable** check box to disable LDAP authentication.

Setting system parameters

About this task

To set the system parameters for a IBM Content Manager OnDemand server:

Procedure

1. Log on to the server.
2. Point to the server and click the right mouse button.
3. From the pop-up menu, select **System Parameters** to open the **System Parameters** dialog box.
4. To change the Minimum Password Age, select the appropriate option. If you select **Expires In**, enter the number of days in the space provided.
5. To change the Minimum Password Length, select the appropriate option. If you select **At Least**, enter the number of characters in the space provided.
6. To change the Inactivity Time Out, select the appropriate option. If you select **Time Out In**, enter the number of minutes in the space provided.
7. To choose a **System Logging**, **User Exit Logging**, or **Login Processing** option, select the check box next to the item.
8. Specify which types of annotations can be used by a user. In the Default Annotation Type section, specify the annotation that is selected as the default type.

Important: This section applies to all users with authority to add annotations to the system.

There are three types of annotations available:

Option	Description
Allow Public	Allows the user to add public annotations. Public annotations to a document can be viewed by anyone who opens that document.
Allow Private to User	Allows the user to add private annotations to a document, and those annotations can be viewed only by the user who created the note, application group administrators, and system administrators.
Allow Private to Group	Allows the user to add annotations to a document, and those annotations can be viewed only by a specific group of users.

9. Specify whether the administrative client displays the System Log Comments window when you perform an add, update, or delete operation.
10. Specify whether you want to use LDAP (Lightweight Directory Access Protocol) authentication in your Content Manager OnDemand server.
11. To generate a summary of the system parameters and display the information in a window where it can be viewed and printed, click **Summary**.
12. When you have finished making changes to the system parameters, click **Update**. (Click **Cancel** to close the **System Parameters** dialog box without saving your changes.) Content Manager OnDemand stores the changes in the database and returns to the administrator window.

Setting trace parameters

About this task

The levels of trace reporting are defined as:

- Error: Returns error messages
- Warning: Returns warning messages
- Information: Returns informational messages
- Flow: Documents entering and exiting of functions

You can set different trace levels for each component. For example, you can set your database to return informational messages, and your server to return error messages.

To set the trace parameters on an IBM Content Manager OnDemand server:

Procedure

1. From the OnDemand Administrator client, log onto the Content Manager OnDemand server for which you want to set trace parameters.
2. In the left panel, right-click the name of the server to which you just logged on.
3. Select **Trace Parameters**.
4. In the **System Trace Settings** dialog box select the **Activate System Trace** checkbox.
5. Select the Trace Level Reporting options you want. (You can select multiple options.) Refer to the online help for more information on each of the options.

Important: You must select the **Active System Trace**, at least one component, and one option for tracing information to actually be logged.

6. Click **Update** to save your selections.

Adding items to a server

This section explains how to add items to a server. You can use commands or a drag-and-drop operation to add items to a server.

When you use the administrative client to add or update the database, IBM Content Manager OnDemand adds a record to the system log that shows the changes you made.

New command

About this task

After logging on to a server, select the area, for example, **Users**. From the **File** menu, select the **New** command to open the **Add** dialog box.

Copy command

About this task

After logging on to a server, select the area. In the list pane, point to the item that you want to copy and click the right mouse button. From the pop-up menu, select **Copy** to open the **Add** dialog box. The fields in the **Add** dialog box contain information copied from the item you selected. Before you can add the item, you

must change the item name. Depending on the item you want to add, you may need to change other fields.

Export command

About this task

The **Export** command is like the **Copy** command, except IBM Content Manager OnDemand adds the item to a different server. You can use the **Export** command to export items from the source server, and add (Import) them to the destination server.

After logging on to the server that contains the item you want to export, select the area. In the list pane, point to the item that you want to export and click the right mouse button. From the pop-up menu, select **Export** to open the **Export** dialog box. Verify the destination server. Then click **Export** to add the item to the server.

If the item already exists on the destination server, the export fails.

When exporting or importing a Content Manager OnDemand application definition from one hardware platform to another (for example, from OnDemand for i to OnDemand for Multiplatforms or the reverse), you must update the Content Manager OnDemand application definition after the export, to ensure that the correct indexer is named. An application definition being exported from an IBM i server might have "OS400" selected for the Indexer value on the Indexer Information tab, which is not valid for any platform other than IBM i. Similarly, an application definition being imported to an IBM i server from Content Manager OnDemand for Multiplatforms might have "ACIF" selected as the indexer, which is not valid on IBM i.

Drag and drop operation

About this task

You can copy and export items using a drag-and-drop operation. For example, to export items from one server and add them to another, select one or more items from the list pane and, while holding the left mouse button down, point to the destination server. Then release the mouse button. If you are logged on to the destination server, then IBM Content Manager OnDemand opens the **Export** dialog box. If you are not logged on to the destination server, then Content Manager OnDemand opens the **Logon** dialog box. After verifying options in the **Export** dialog box, click **Export** to copy the items to the server.

If the item exists on the destination server, the export fails.

You can also use a drag-and-drop operation to copy an item. For example, to copy a user, select the user ID from the User ID list and, while holding the left mouse button down, point to the same server on which the user is listed. Then release the mouse button to open the **Add** dialog box.

Report Wizard

IBM Content Manager OnDemand provides user assistance and easy-to-use tools to help you administer Content Manager OnDemand. The Report Wizard helps you add a report to Content Manager OnDemand by asking questions, which allows you to progress in an organized manner toward completing an application group, application, and folder. Here are a few things to remember about the Report Wizard:

- You move through the Report Wizard by answering questions that appear on the screen.
- You can return to the previous screen at any time by clicking **Previous**.
- You can advance to the next screen at any time by clicking **Next**.
- You can advance to the final screen by clicking **Finish**. By choosing **Finish**, you permit the Report Wizard to make all remaining decisions for you.
- You can obtain online help for a screen at any time by clicking **Help** or pressing F1.

Important: The Report Wizard processes your own input files. You must select a sample input file to proceed. Then, the graphical indexer is invoked to allow you to mark the data to define your indexing parameters. The graphical indexer that is invoked through the Report Wizard is the same graphical tool that is invoked directly by selecting **Sample Data** and then clicking on the **Modify** button from the indexer information tab of a Content Manager OnDemand application definition.

You can use the Report Wizard to add an application group, application, and folder for a selected report. These actions include defining indexing information, defining database and folder fields, configuring data and storage management, specifying whether the application group can contain more than one application, and naming the application group, application, and folder.

You can also use the Report Wizard to add an application to an existing application group. This action includes defining indexing information, specifying storage information, and identifying the application within the application group. To add an application to an application group, the application group must have a database field to hold the values that uniquely identify an application within the application group. The field must contain at least one unassigned application identifier. See the Field Information page in application groups for detailed information about application identifiers.

Starting the Report Wizard

About this task

From the administrative client, log on to the server to which you want to add the report.

Procedure

1. To define a new application group, application, and folder, click the Report Wizard icon on the toolbar.
2. To add an application to an existing application group:
 - a. Under the server, select **Application Groups**
 - b. Select the name of the application group to which you want to add the application

- c. Click the Report Wizard icon on the toolbar
3. Follow the on-screen instructions to add the report.

Using the Report Wizard

The screens in the Report Wizard are described in the following. On most screens, standard options will already be selected for you. Unless you have a clear reason not to, accept the defaults.

Note: Depending on how you use the Report Wizard, you may not see all of the screens described below.

Introduction screen

Provides a brief explanation of the Report Wizard. First, choose the data type of the report you are defining. Click the **Select Sample Data** button to select a file that contains a sample of the actual report data. The Report Wizard lets you select a spooled file on the server (search by user profile or by output queue) and copies that sample data to your workstation for you to use for indexing.

When you click **OK** from the Select File panel, the Report Wizard reads the data into the Report window.

Report window

Displays the sample data file and provides easy-to-use tools to help you define indexing information, database fields, and folder fields. Press **F1** to display the online help for options and commands available from the Report window. Use the online help to learn how to define triggers, fields, and indexes, database fields, and folder fields.

Important: When you have finished defining the indexing, database, and folder information, be sure to save your changes when prompted.

Managing data screen

When you load a report into the system, you can specify that you want report data to be stored in Large Objects. You also need to specify how you want IBM Content Manager OnDemand to manage annotations that users attach to pages of the report.

Application identifier screen

When you use the Report Wizard to add an application to an existing application group, you must specify the name of the application and select a value that uniquely identifies the application within the application group.

Storage management screen

Determines where the storage manager maintains copies of reports, and how and when Content Manager OnDemand deletes report data from the system.

Applications in the application group screen

If the report that you are defining is one of several that will be stored in the same application group, you can use the Report Wizard to define an *application ID field*.

An application ID field is a database field that contains values that identify an application within the application group. IBM recommends that you always define an application ID field. See the figure for an example. You may not think that you need an application ID field at the time that you create your application group, if you are adding an application group that will hold only one application **at the current time**. However, if you decide later to add other applications to the application group (or if you want to maintain multiple versions of your application definitions) and you then need to define an application ID field for the application group, you will not be able to do so because an application ID field **cannot be added after the application group is created**. (All fields must be added during the original application group definition.) Also note that the application ID field can be hidden from users that do not require it to search for documents. See the online help for more information about the application ID field.

Name screen

Specify the names of the application group, application, and folder. After you enter the names, Content Manager OnDemand queries the server to make sure that the names are valid and unique.

Wizard complete screen

Confirms the selections you made for the report. Click **Display** to view details about the application group, application, and folder. From the detail report window, choose the **Print** icon from the toolbar to print a copy of the detail report.

Note: When you are satisfied with the selections you made for the report, click **Finish** to complete defining the report. Content Manager OnDemand adds the application group, application, and folder to the library server, closes the Report Wizard, and returns to the administrator window.

Concepts

You can use IBM Navigator for i to setup and maintain Content Manager OnDemand disk pools, network file system (NFS) disk pools, Tivoli Storage Manager (TSM) storage nodes, optical storage sets, optical volumes, tape devices, tape volumes, directory and output queue monitors, and migration policies.

You can use the Content Manager OnDemand Administrator client to maintain Content Manager OnDemand users, groups, printers, application groups, applications, folders, cabinets, and holds.

Migration policies

Migration policies contain migration and storage media characteristics for data archived using IBM Content Manager OnDemand. The information is used by the archive storage management process (ASM), also referred to as the migration process, to determine if and when archived data should be moved as it ages through a hierarchy of storage media having different performance and capacity characteristics. Examples of these media types include disk, optical, and tape storage. Each step in the movement of data through this storage hierarchy is referred to as a migration policy storage level, or simply, a storage level. Each policy must contain at least one storage level. Additional levels can be defined to meet your storage and retrieval requirements. The STRASMOND command is the command used to force the data to move to the next storage level in the migration policy. A report is produced when ASM is run and has a spooled file name of QPRLCASM1. The report provides a list of actions that the ASM process performed. This report should be checked each time ASM is run to ensure that processing of the data completed successfully. If you find a failure, you should check the job log for the STRASMOND job to determine the cause of the failure.

The migration policy also specifies:

- If separate archived files are to be aggregated, or combined, with other archived files having similar retention and migration characteristics. Enabling aggregation is usually recommended because it can improve performance by allowing Content Manager OnDemand to manage a smaller number of larger objects rather than a large number of small objects. However, it is important to understand that an aggregate must reach its maximum size or exceed its time period before the aggregated object can flow to the first level of the Migration Policy. Because of this, you should be cautious of large aggregate sizes. If you are aggregating many small objects, it is possible for the aggregate not to be migrated to the first level of the Migration Policy for many months. For this reason, you may prefer to aggregate after a specified time period rather than by size.
- If two copies of archived data are to be kept at some or all levels in the migration sequence.
- If a one-time tape backup is performed.

The migration policy name can be up to 60 characters long and must not be a duplicate of another policy within the same instance. If the Enable aggregation is selected, the archive storage management process combines individual archived objects on disk into larger objects to provide efficient processing. This process occurs prior to migration of the object from disk to the first storage level. The

aggregation process appends to the same file (aggregate) until the aggregate is 'closed.' The aggregate is closed when it either reaches a specified maximum size or a specified number of days. Storage levels can be added before or after a pre-existing storage level. Existing storage levels within a migration policy can also be changed or removed, however these changes will not affect migrated data already residing at this level. If you need to change the dates for migrated data already residing at a particular level, use the Change Policy Level Data (CHGPLDOND) command. If a one-time tape backup is requested, a tape media type must be specified.

Monitor definitions

Monitor definitions are used to specify what output queue or directory will be monitored for input files to be processed. If defining more than one monitor job, specify a unique job name for each monitor.

When using an output queue monitor, Content Manager OnDemand will only process spooled files that are in a ready (RDY) state. When the monitor job selects a spooled file from the selected output queue for archiving (also known as loading), it needs to determine which application group and application to associate with the spooled file so that the file can be archived correctly. Since the only data available to the monitor are the attributes of the selected spooled file, the application group name and application name must be derived from the contents of one of these attributes. Not all attributes are suitable for this purpose. Content Manager OnDemand will examine the contents of up to three of the following nine attributes, in the order specified in the Check first, Check next, Check last selections:

- Spooled file name
- Form type
- User data
- Job name
- User-defined options 1 through 4
- User-defined data

When using a directory monitor, Content Manager OnDemand will only process IFS files with specific file extensions. For a regular directory monitor, files ending in .IND or .PDF must exist in the directory to trigger the archive process. For a 'type 2' directory monitor, files ending in .ARD must exist in the directory to trigger the process. When the monitor job selects a file from the selected directory for archiving, it needs to determine which application group and application to associate with the file so that it can be archived correctly. Since the only data available to the monitor is the filename of the selected file, the application group name and application name must be derived from the filename. Content Manager OnDemand will examine the contents of the first, second, third, or fourth part of the filename, in the order specified in the Check first, Check next, Check last selections. For example, a file named INVOICES.NOVEMBER.pdf might be defined by specifying **FIRST** from the **Check first** pull down list for the application group name to cause Content Manager OnDemand to use the 'INVOICES' application group definition to archive the file.

The attribute selected from the **Check first** pull down list is examined first. If the value of this attribute does not match the name of an existing application group, Content Manager OnDemand examines the attribute selected from the **Check next**

pull down list, if specified. If the value of this attribute does not match an application group name, the attribute selected from the **Check last** pull down list, if specified, is checked.

For output queue monitors, if a valid application group is not determined using the above method, the spooled file is moved to another output queue designated as an 'error' queue which is defined in this monitor definition. If the spooled file is successfully archived, it can be moved to a 'processed' queue or deleted according to the specifications in the monitor definition. The same process is followed to find a valid application name, unless you specify that the application name is the same as the application group name.

For directory monitors, if a valid application group is not determined using the above method, the file remains in the directory and is renamed by appending '.ERR' to the filename. If the file is successfully archived, it can be renamed by appending '.PRC' to the filename or deleted according to the specifications in the monitor definition. The same process is followed to find a valid application name, unless you specify that the application name is the same as the application group name.

A monitor can be started manually, by a job scheduler, or started when the subsystem starts. A monitor can be ended manually, at a specific time of day, after a specified time period, after all queue entries are processed, or it can be specified when the monitor is started.

For output queue monitors, the first time you start a monitor for a particular output queue, it is best to do it when there are no spooled files in the output queue. When a monitor is started for the first time, an empty data queue with the same name as the output queue is created, which will receive entries for all spooled files that appear in the output queue in Ready status. Once the output queue monitor has been started (and therefore the data queue has been created), you can then begin moving the spooled files that you wish to capture into the output queue. The data queue entries that get created will trigger the monitor to process each spooled file.

Tape devices

When a tape backup is requested, or if you will use tape as an archive media, you must define a tape device to IBM Content Manager OnDemand. The tape device name you specify must match the name of an existing IBM i tape device description. A media type must be specified from the list of supported media types for either read or write operations. If the tape device being defined has an automatic cartridge loader (ACL), the number of cartridges can be specified. A value of zero specifies that this device does not have an ACL. Content Manager OnDemand can be used with an automated tape library; if a tape library will be used, the name must be specified. If a tape manager other than Content Manager OnDemand (such as BRMS) is used, a media library does not need to be specified.

Tape volumes

Tape volumes that can be used by IBM Content Manager OnDemand must be defined. The name of the volume specified must match the name that was used when the tape volume was initialized. The instance to which the tape volume belongs must be specified, along with the capacity and media type of the volume. The media type tells Content Manager OnDemand which tape device to use. A media device library can be specified if an automated tape library is used. Leave

this field blank if a tape manager other than Content Manager OnDemand (such as BRMS) is used. A tape volume can be marked full, preventing Content Manager OnDemand from writing any additional data to the volume. Content Manager OnDemand will automatically mark the volume full when it detects that the tape is full.

Optical storage groups

Optical storage groups are used by IBM Content Manager OnDemand to group sets of optical volumes together to store related data. A storage group lets you group together reports that have similar storage requirements such as days on disk or expiration days. By referring to a specific storage group in your migration policy, you can control which reports are stored on a particular set of optical volumes. If optical volumes are defined as rewritable, the space can be reused on the volume by defining a free space threshold percent and volume full reset. If the volume full reset is not defined, once the volume is marked full it remains full unless manually changed. The optical storage group can also be defined as the primary or backup storage group.

Optical volumes

Optical volumes that can be used by IBM Content Manager OnDemand must be defined. The name of the volume specified must match the name that was used when the optical volume was initialized. The OnDemand instance and optical storage group to which the volume belongs must be specified. The capacity and volume type (primary or backup) of the volume is also required. An optical volume can be marked full, preventing Content Manager OnDemand from writing any additional data to the volume. Content Manager OnDemand will automatically mark the volume full when it detects that it is full.

Disk pool storage groups

The disk pool storage group is used to identify an IBM i auxiliary storage pool (ASP) that the archive storage management process may use as storage media when migrating archived data. You must specify a pool number ranging from 1 to 32 which corresponds to an existing ASP. The type of data, primary or backup, that will be stored to the defined ASP must also be specified.

You can also use a mounted Network File System (NFS) exported directory as a disk pool. See “Using a Network File System (NFS) directory for document storage” on page 21 for detailed setup instructions.

Users

When you define an IBM Content Manager OnDemand user, you create a user ID with which a person in your organization logs on to the Content Manager OnDemand server. You can optionally add the user ID to folders and application groups permissions, which is one way to let the user open folders and access data.

Each person in your organization logs on to the server using a Content Manager OnDemand user ID. Content Manager OnDemand authenticates the user ID and determines the usage and administrative authority available to that person, based on the user ID. It is important that you understand the details in the OnDemand user ID relationship to IBM i user profiles topic in the *IBM Content Manager OnDemand for i: Common Server Planning and Installation Guide* before you continue with this section.

A Content Manager OnDemand user ID does not necessarily have to identify an individual by name. However, for accounting purposes, most customers assign a Content Manager OnDemand user ID to each person that will use the system. When you initialize the system, Content Manager OnDemand automatically creates the QONDADM user ID. 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 to Content Manager OnDemand, and creating, updating, and deleting application groups, applications, folders, migration policies and storage sets, and printers.

Remember: The QONDADM user ID has an initial password of QONDADM1. Because the QONDADM user ID has system authority, we force you to change the password the first time you sign on.

When naming Content Manager OnDemand users, the name that you specify:

- Can contain one to 10 or 128 characters (bytes) depending on whether you are linking i user profiles to your Content Manager OnDemand users. If you are linking the two, then the user names should match your i user profile names.
- Cannot include the ' (apostrophe), * (asterisk), % (percent), + (plus), _ (underscore), [(left bracket),] (right bracket), " (double quote), or blank character
- Must be unique to the server

You can specify a user ID in mixed case. By default, Content Manager OnDemand ignores the case (for example, LaGuarde is the same as laguarde). Content Manager OnDemand converts lowercase letters in a user name to uppercase (LaGuarde is stored as LAGUARDE). However, depending on how you configure the Login Processing system parameters, user ID processing on your system may be different (the case may be significant). See “System parameters” on page 48 for more information.

User types

When you add a user to IBM Content Manager OnDemand, you specify the User Type. The User Type and the Authority determines the types of tasks that the user can do when logged on to the system. You can choose from the following User Types:

User Users can log on to Content Manager OnDemand, open folders that they are authorized to access, and search for and retrieve data from application groups that they are authorized to access. Users can be given authority to do other things on the system.

User Administrator

A user that can also add, update, and delete users and user administrators. A user administrator can be given authority to do other things on the system.

Application Group/Folder Administrator

A user that can also add, update, and delete application groups, applications, and folders. An application group/folder administrator is automatically given Logical Views permission to all application groups. An application group/folder administrator can be given authority to do other things on the system.

System Administrator

A user that can also add, update, and delete any user, group, application

group, application, or folder on the system. A system administrator is automatically given Logical Views permission to all application groups. A system administrator is also the only user that can maintain storage sets and printers.

Restriction: When adding or updating a user, you are not permitted to set the User Type or Authority to a level that exceeds your own. For example, a user with Create Users and Create Groups authority cannot create a user with Create Folders authority.

By default, only the user, the user that created the user, user administrators, and system administrators can view or maintain the user. See the User Permissions page for more information.

Only a system administrator, an application group/folder administrator, a user with administrator authority for an application group, or a user with add document permission can store data in an application group.

Only a system administrator, an application group/folder administrator, or a user with delete document permission can delete data from an application group.

Users who need to run server commands such as those listed in “Command reference” on page 243 or server APIs from QSHHELL such as those listed in “API and user exit reference” on page 259 need to have QRDARSADM as the group profile (or a supplemental group) in their IBM i user profile.

Authority

The Authority options allow the user to do other things in IBM Content Manager OnDemand. For example, A User Type of User can be permitted to create users. The authority options that you can select depend on the User Type. For example, if the User Type is Application Group/Folder Administrator, then by definition, the user can create application groups and folders. Therefore, the only additional authorities that the user can be given are Create Users and Create Groups. Choose from the following:

Create Users

An optional authority for users and application group/folder administrators:

- If the User Type is User, lets the user create users with a User Type of User.
- If the User Type is Application Group/Folder Administrator, lets the user create users with a User Type of User or a User Type of Application Group/Folder Administrator.

Tip: Users with Create Users authority can maintain the users that they create, so long as they remain an administrator of the user.

Create Groups

An optional authority for users, user administrators, and application group/folder administrators. Lets the user create groups. Users with Create Groups authority can maintain the groups they create, so long as they remain a group owner.

Remember: These group definitions are not the same as IBM i group profiles, although the names may match if you find that easier to maintain.

Create Application Groups

An optional authority for users and user administrators. Lets the user create application groups. Users with Create Application Groups authority can maintain the application groups that they create, so long as they remain an application group administrator.

Create Folders

An optional authority for users and user administrators. Lets the user create folders. Users with Create Folders authority can maintain the folders that they create, so long as they remain a folder administrator.

Restriction: When adding or updating a user, you are not permitted to set the User Type or Authority to a level that exceeds your own. For example, a user with Create Users and Create Groups authority cannot create a user with Create Folders authority.

Groups

IBM Content Manager OnDemand groups are a means to organize users of the system by function, authorization, or any other purpose that you might require. You do not have to assign a user to a group, however doing so can simplify administration of users with similar requirements and capabilities.

When you define a group, you can add users to the group and specify folder and application group permissions that are common to all of the users that belong to the group. The permissions determine the types of actions users assigned to the group can perform on the system.

When naming groups, the name that you specify:

- Can contain one to 128 characters (bytes)
- Cannot include the ' (apostrophe), * (asterisk), % (percent), + (plus), _ (underscore), [(left bracket),] (right bracket), " (double quote), or blank character
- Can be mixed case. However, the case does not create a unique name (LaGuarde is the same as laguarde)
- Must be unique to the server

You can assign a group owner. The group owner can add users to and remove users from the group. To maintain a group's application group and folder permissions, the group owner must have administrator authority for the application groups and folders or be an application group/folder administrator or a system administrator. If you do not assign a group owner, only a system administrator user can maintain the group.

Remember: These group definitions are not the same as IBM i group profiles, although the names may match if you find that easier to maintain.

Printers

IBM Content Manager OnDemand supports two types of server print devices: a fax machine and a physical printer. A server print device always has an output queue on the server and is defined using the Content Manager OnDemand administrative client.

PSF/400 is required for formatting Advanced Function Presentation (AFP) print output from Content Manager OnDemand. (PSF/400 may also be required by some fax software as well.) PSF/400 allows you to use electronic forms, images, graphics, and typographical fonts in the documents that you print. PSF/400 accepts input data streams, such as AFP and line data and prepares the data for the destination print device.

When a user selects an item and chooses server print, Content Manager OnDemand retrieves a copy of the item and places it on the output queue associated with the server printer.

Storage sets

Storage sets are defined for application groups with similar storage management characteristics, such as the length of time that files are maintained in the application group and the type of media on which the files are stored. A storage set is created automatically when you create a migration policy, and the names are identical. You do not have to manually create a storage set.

If you plan to maintain application group data in archive storage, then you must specify a storage set name in the application group definition that matches the migration policy to be used. For more information on migrating and expiring documents, and recommendations for storage management criteria defined in your application groups, storage sets, and migration policies, see “Defining document storage management” on page 11.

Application groups

An application group represents the data that you store in IBM Content Manager OnDemand and the documents that users query, view, print, and fax using Content Manager OnDemand client programs. For example, the data can be reports generated by an application program, index data, and annotations created by users.

When you define an application group, you specify properties of the application group, such as the organization of the database and the storage characteristics for the files that are to be stored in the application group. You also define the database fields that will hold index data extracted from the reports that you store in the application group.

Content Manager OnDemand extracts index data from the reports that you load into an application group and places the data in the database fields that you define. Content Manager OnDemand uses the index values to identify the documents that meet the search criteria entered by a user.

When you define an application group, you can also select the types of application group messages that Content Manager OnDemand saves in the system log.

When you define an application group, you specify permissions that let users access and maintain the application group and application group data. You can identify the groups and users that can access data stored in the application group with Content Manager OnDemand client programs. You can specify the types of functions that users can perform, such as viewing, printing, and annotating reports. You can assign administrator authority to a user or a group. Administrator authority allows a user to update the application group, for example, to authorize other users to access data stored in the application group.

Content Manager OnDemand organizes information about an application group into tabs: General, Message Logging, Storage Management, Permissions, Field Definition, and Field Information. Each tab contains options, fields for you to enter information about the application group, and command buttons.

Applications

You typically define an application for each different type of report or source of data that you plan to store in IBM Content Manager OnDemand.

When you define an application, you assign the application to an application group and specify the physical and logical characteristics of the report. The physical characteristics of a report include the code page, the type of data found in the input file, and information about carriage control characters. The logical characteristics of a report include the different ways that you want to present the information contained in the report to your users.

The Content Manager OnDemand data indexing, loading, and viewing programs use the information that you provide to process the report. For example, you can specify the parameters that the Content Manager OnDemand indexing program uses to locate and extract index data from the report. You can create logical views for the application. Logical views represent different ways to display pages of the report. You can set up printing options, such as defining a default printer for users and printing options for AFP and line data documents.

Content Manager OnDemand organizes information about an application into tabs: General, View Information, Indexer Information, Load Information, Logical View Fields, Logical Views, and Print Options. Each tab contains options, fields for you to enter information about the application, and command buttons.

Folders

A folder provides users the means to access the reports that you store in IBM Content Manager OnDemand. A user opens a folder, constructs a query, and retrieves documents from the application groups that can be searched from the folder. The user can use the folder to view, print, annotate, fax, and email documents.

When you define a folder, you specify the properties of the folder, such as the name and description of the folder, create the search and display fields that appear when the user opens the folder, and map the folder fields to application group database fields.


You can also specify the groups and users that can open the folder with Content Manager OnDemand client programs and other folder permissions. For example, you can authorize a user to be the folder administrator. The folder administrator can authorize other users to open the folder and make changes to the folder fields.

Content Manager OnDemand organizes information about a folder into tabs: General, Permissions, Field Definition, Field Information, and Field Mapping. Each tab contains options, fields for you to enter information about the folder, and command buttons.

Cabinets

A cabinet is a container for folders. You can use cabinets to manage folders and enable users to navigate to folders more easily. A folder can belong to one or more cabinets. The following diagram describes the relationship between two cabinets and five folders. The Monthly Report cabinet contains three folders: Fund balance, Fund transactions, and Fund performance. The Client Report cabinet also contains three folders: Fund performance, Bond performance, and Stock performance. The two cabinets share the Fund performance folder. In this example, you can use the Client Report cabinet to quickly retrieve information about a client's portfolio, and use the Monthly Report cabinet to obtain fund information for internal analysis.

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. Cabinets and folders

Holds

In Content Manager OnDemand, documents are stored and retained for a specified period of time. After that specified period of time, the documents are removed from the system. In certain situations, you might need to keep one or more documents beyond the expiration date by using hold objects. You can put documents on hold in several ways:

- Placing a hold on a single document
- Placing a hold on a load of documents

You can also add one or more documents to existing hold objects.

About application groups, applications, and folders

Before you can store a report into IBM Content Manager OnDemand, you must create an application group and an application. Before users can search for and retrieve data, you must create a folder.

- Users open a folder to search for reports that you load into OnDemand. You define one or more search fields for the folder. A folder search field is mapped to an application group database field.
- Each database field that you define represents a category of information in the report, such as a customer name, invoice number, or balance. When you add an application group, Content Manager OnDemand creates a database table. The database fields that you define are columns in the table.
- You define an application for each report that you plan to store in Content Manager OnDemand. When you add an application to the system, you define information that the Content Manager OnDemand viewing, indexing, and loading programs use. When you define indexing information, you identify the name, location, and length of each category of index information that you want to extract from the report. When you define loading information, you map index fields in the report to application group database fields.
- When you load a report into the system, Content Manager OnDemand stores the index values extracted from the report into database fields in records that are added to an application group table.
- Content Manager OnDemand uses the index values to identify the documents that meet the search criteria that the user entered into the folder search fields.

When you want to define a report to Content Manager OnDemand, your first task is to identify the application group from which Content Manager OnDemand obtains information about the index fields and how documents are to be maintained on the system. When you define an application group, you specify how you want Content Manager OnDemand to structure information in the database and define the database fields. When you define an application group, you also specify how you want Content Manager OnDemand to maintain data on the system. For example, you might specify that report data should be maintained on disk storage for 60 days and in archive storage for five years. Content Manager OnDemand maintains all of the data stored in the application group the same way. Content Manager OnDemand maintains each report that you store in the application group for the same length of time.

- You can store the report in an existing application group. However, you must be able to index the report by using the database fields that are already defined in the application group. The storage management information for the application group must support the length of time that you want Content Manager OnDemand to maintain the report on the system and how and where that you want Content Manager OnDemand to store and maintain the report data. You can verify information about an existing application group with the Properties command. The General tab shows the database organization for the application group. The Storage Management tab shows the data migration information. The Field Definition tab shows the application group database fields.
- If there are no application groups defined to Content Manager OnDemand or there are no application groups that support the database and storage management requirements of the report, then you must add an application group to the system.

After you add an application group, you must define an application for the report. Most customers create a Content Manager OnDemand application for each different type of report or source of data that they plan to store in Content Manager OnDemand. When you create an application, you must assign it to an application group. The application group determines where Content Manager OnDemand will store the report data. When you create an application, you also

specify information that the Content Manager OnDemand client programs use to view and print pages of the report and you specify instructions for the data indexing and loading programs.

The last step in the process of adding a report to Content Manager OnDemand is to create a folder. Users open the folder to search for, display, and print reports. When you define a folder, you select the application group that contains the data that you want users to search when they open the folder. By creating folders that can search specific application groups, you can determine the reports that are available to users when they open a folder.

When you define a folder, you define search and display fields. You specify characteristics of the folder fields, such as default search operators and whether Content Manager OnDemand displays default search values for the fields when a user opens the folder. You also map the folder fields to application group database fields.

OnDemand permissions

Permissions are the means by which IBM Content Manager OnDemand determines who can open folders and search for documents stored in application groups. Content Manager OnDemand also uses permissions to determine who can maintain folders and application groups with the administrative client.

By default, only an application group/folder administrator, a system administrator, or the person who adds the folder can open and maintain the folder. By default, only an application group/folder administrator, a system administrator or the person who adds the application group can access data stored in the application group or maintain the application group.

Content Manager OnDemand provides several ways for you to specify permissions. You can specify permissions when you add or update a folder or an application group. You can also add, remove, or update a user's or group's folder or application group permissions when you add or update the user or group.

As both a convenience and security measure, you can assign a user to a group. Content Manager OnDemand groups allow you to organize users by function, authorization, or any other logical grouping that you might require. When you assign a user to a group, the user obtains the permissions that are in effect for the group. For example, suppose that 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. However, there are exceptions. See "Specifying permissions" on page 71 for details.

Remember: These group definitions are not the same as IBM i group profiles, although the names may match if you find that easier to maintain.

You can specify a default set of permissions for folders and application groups. Content Manager OnDemand uses the default permissions when users and groups do not have specific permissions for the folder or application group. If you specify permissions for a group, then the group permissions take precedence over the default permissions. If you specify permissions for a user, then the user permissions take precedence, regardless of any group that the user may belong to or the default permissions that you specified.

Folder permissions

You can specify default (*PUBLIC) folder permissions and folder permissions for specific groups and users. The default permissions provide every user and group defined to the server with the permissions that you specify. Permissions for a group provide the users that you add to the group with the permissions that you specify for the group. Permissions for a user provide the user with the permissions that you specify. By default, only an application group/folder administrator, a system administrator or the person who adds the folder can open and maintain the folder.

Refer to online help for details about setting folder permissions.

Application group permissions

You can specify default (*PUBLIC) application group permissions and application group permissions for specific groups and users. The default permissions provide every user and group defined to the server with the permissions that you specify. Permissions for a group provide the users that you add to the group with the permissions that you specify for the group. Permissions for a user provide the user with the permissions that you specify. By default, only an application group/folder administrator, a system administrator or the person who adds the application group can access data stored in the application group and maintain the application group.

Refer to online help for details about setting application group permissions.

Specifying permissions

To ease the administration of IBM Content Manager OnDemand, most customers organize their users into groups, add the groups to folders and application groups, and specify permissions for the groups. You should plan your groups before you begin creating them. After you start using the system, you may find it difficult to change the organization of your groups.

Remember: These group definitions are not the same as IBM i group profiles, although the names may match if you find that easier to maintain.

When you add a user to a group, the user automatically obtains the permissions that were specified for the group. When you add a user to more than one group, the user normally obtains the permissions of all of the groups. For example, using the group properties listed in the table, a user that belongs to both groups can open the Student Bills and Student Transcripts folders.

Table 2. Group permissions

Group	GID	Folders	Permission
Admissions	1080100	Student Transcripts	Access
Accounting	1080101	Student Bills	Access

Most situations involve adding a group to a folder, specifying permissions for the group, and then adding users to the group. However, there may be situations when you need to deny a group of users access to a folder. When you use groups to deny access to a folder, you must understand how Content Manager OnDemand determines folder permissions for a group (and users assigned to the group). For example, consider the group properties listed in the following table.

Table 3. Group permissions

Group	GID	Folders	Permission
Admissions	1080100	Student Bills	None
Admissions	1080100	Student Transcripts	Access
Accounting	1080101	Student Bills	Access
Accounting	1080101	Student Transcripts	Access

A user that belongs to both groups can not open the Student Bills folder.

When a user belongs to more than one group, Content Manager OnDemand uses the group identifier (GID) to determine the user's permissions. When two (or more) groups provide permissions for the same folder, the user obtains the permissions of the group with the lowest GID.

It is important to note that you cannot change a GID after a group has been created. You can, however, modify the value that is presented as the suggested GID when the group is created.

In the example depicted in the second table, both groups have been added to the Student Bills folder. Since the Admissions group has a lower GID than the Accounting group, Content Manager OnDemand uses the permissions specified for the Admissions group to determine the permissions of a user that is assigned to both groups. Consequently, users assigned to both groups cannot access the Student Bills folder.

You're probably asking yourself, Why would I assign a user to more than one group? or Why would I create a group with no access to a folder? Perhaps some examples will help answer these questions (and clarify the Content Manager OnDemand permission hierarchy). As you review the examples, please remember the following rules:

- By default, only an application group/folder administrator, a system administrator, or the person who created the folder can access the folder
- You can use the *PUBLIC name to specify default permissions for all other users
- You can specify permissions for specific groups and users:
 - All of the users that belong to a group that you add to a folder will obtain the permissions that you specify for the group
 - A user that belongs to two (or more) groups that have been added to the same folder will obtain the permissions of the group that has the lowest GID
 - The permissions that you specify for a user override all other permissions, including any default permissions (*PUBLIC) and any groups to which the user belongs and that are added to the folder

Examples

The examples that follow show how to add groups to folders and specify folder permissions. The same considerations hold true for adding groups to application groups and specifying application group permissions.

Providing a group of users access to a folder:

About this task

Let's say that you want to provide a single group of users access to a folder. Complete the following steps:

Procedure

1. With *PUBLIC selected, clear all of the permissions check boxes (this is the default).
2. Add the group to the folder.
3. Select the **Access** check box.

Results

Users assigned to the group automatically obtain permission to open the folder.

Denying a group of users access to a folder:

About this task

Let's say that you want to prohibit a single group of users from accessing a folder, while allowing all other users defined to the server to open the folder. Complete the following steps:

Procedure

1. With *PUBLIC selected, select the Access check box (this lets all users open the folder).
2. Add the group to the folder.
3. Clear all the permissions check boxes.

Results

Users assigned to the group cannot open the folder.

Providing one group of users access and denying another group of users access to the same folder:

About this task

Let's say that you want to allow a group of users to access a folder. However, you need to prohibit certain users in the group from accessing the folder. You could exclude the users from the group that can access the folder. However, there may be other folders that you want the users to access as part of the group. To solve this, create two groups, one without access to the folder and the other with access to the folder, and assign the users to the respective groups. For example:

Procedure

1. Create the "no access" group. This group must have a lower GID than the "access" group. Add users to the group.
2. Create the "access" group. Add users to the group.
3. With *PUBLIC selected, clear all of the permissions check boxes (this is the default).
4. Add the "access" group to the folder.
5. Select the **Access** check box.
6. Add the "no access" group to the folder.
7. Clear all of the permissions check boxes.

Results

If you later need to deny other users access to the folder, simply add the users to the "no access" group. You can also move users from one group to the other.

Denying one user in a group access to a folder:

About this task

Let's say that you want to prohibit one user in a group from accessing a folder. After adding the group to the folder and specifying the access permission, all users assigned to the group can open the folder. To override the group permissions, we can add an individual user to the folder and set permissions at the user level. Complete the following steps.

Procedure

1. With *PUBLIC selected, clear all of the permissions check boxes (this is the default).
2. Add the group to the folder.
3. Select the **Access** check box.
4. Add the user to the folder.
5. Clear all the permissions check boxes.

Results

Even though the user belongs to the group, the user cannot open the folder.

Providing one user in a group administrator authority:

About this task

Let's say that you want to provide one user in a group the ability to administer the folder. Complete the following steps.

Procedure

1. With *PUBLIC selected, clear all of the permissions check boxes (this is the default).
2. Add the group to the folder.
3. Select the **Access** check box.
4. Add the user to the folder.
5. Select the **Administrator** check box.

Results

Only the user that you added can administer the folder; the other users in the group can open the folder.

Specifying default permissions:

The default (*PUBLIC) permissions that you specify for an application group or a folder will apply to every user or group defined to the server who is not provided with specific permissions.

For example, suppose that you specify **Access** as the default permission for an application group. Every user and group that is not provided with specific permissions can access the data that is stored in the application group. Then, you specify **Access** as the default permission for a folder. Every user and group that is not provided with specific permissions can open the folder. Later, you add a user, without specifying application group or folder permissions. The user can open the folder and access the data stored in the application group.

While default permissions do provide flexibility to maintain your system, you must make sure that using the default permissions does not circumvent your security strategy. Rather than specifying default permissions for application groups and folders, you may want to use groups as a means to implement your security strategy. For example, you can clear all of the permissions under *PUBLIC and then add groups to a folder and specify the appropriate permissions for each group. When you add a user to the system, you can assign the user to a group. The user automatically obtains the permissions of the group. If the group does not have access to a particular application group or folder, then neither does the user. With this strategy, until you assign the user to a group, or provide the user with specific permissions, the user cannot access the folder.

Hints and tips

1. To simplify the task of providing access to application groups and folders, give access to a group rather than a user. When a new user needs access, add the user to the group.
2. To allow an Application Group/Folder Administrator to see groups in the permissions list, add the Application Group/Folder Administrator to the groups that require access to application groups and folders.
3. To allow multiple users to administer the same groups, create a group of users and make that group the group owner for any groups that need to be administered by multiple users.
4. The Create Groups authority is most effectively used if it is combined with the Create Users authority or added to a User Administrator. Because the purpose of a group is to give a set of users permissions to another object, it is not very useful if the user that creates the group does not have access to any users. Otherwise, the user that creates a group must be given access to each user that needs to be added to the group.
5. If you need to reset the Content Manager OnDemand information related to your optical volumes, issue the following program call (with one parameter):

```
CALL PGM(QRDARS/QRLCASMROV) PARM(instancename)
```

where *instancename* is the name of the Content Manager OnDemand instance for which you need to reset your optical volumes.

Running this program updates the Optical Volume Capacity, recalculates the Optical Bytes Used, and sets the Optical Volume Full Flag to Y if it is currently set to N and the optical volume has less than 1 MB of space available. This program can only be run by a user profile with QONDADM group or supplemental group profile in their user profile.

You should not run this program while the Archive Storage Manager (ASM) or the Disk Storage Manager (DSM) with ASM(*YES) are running. The program might cause some or all optical volumes that are known to Content Manager OnDemand to be mounted to check for space.

LDAP (Lightweight Directory Access Protocol) authentication support

The Lightweight Directory Access Protocol (LDAP) is an open industry standard that has evolved to share information between distributed applications on the same network, organize information in a clear and consistent manner, and prevent unauthorized modification or disclosure of private information. In recent years, LDAP has gained wide acceptance as the directory access method of the Internet, and becomes strategic within corporate intranets.

You can use LDAP to manage basic login authentication directly on the server, in other words, you no longer need to use the user security exit.

Requirements

The following LDAP servers are supported:

- Novell eDirectory Version 8.8 SP2
- Sun Java System Directory Server Enterprise Edition 6.3
- IBM Tivoli Directory server (TDS)
- Microsoft Active Directory Application Mode (ADAM) server
- Microsoft Active Directory (AD)

How Content Manager OnDemand works with LDAP

The following diagram illustrates how Content Manager OnDemand works with LDAP:




Figure 2. How OnDemand works with LDAP

When a Content Manager OnDemand client makes a login request to the Content Manager OnDemand server, if you enabled LDAP authentication in the server, the Content Manager OnDemand server makes an authentication request to the LDAP through either an anonymous or credentialed bind.

This initial call accesses the LDAP server, searches for the user's credentials and finds the user's distinguished name (DN). If the user's DN is found, the Content Manager OnDemand server makes another call to the LDAP server using that DN to confirm that the password that was given by the user is correct. If the password

is correct, the LDAP server returns a mapped attribute in LDAP, which is usually the Content Manager OnDemand user ID. The Content Manager OnDemand server takes the attribute, and proceeds with its login.

Enabling LDAP authentication

To enable LDAP authentication, in the Content Manager OnDemand Administrator client, right-click your Content Manager OnDemand server, and select **System Parameters**. In the System Parameters window, under LDAP Authentication, select the **Enable** check box. Under Login Processing, select the **Password Case Sensitive** check box.

To disable LDAP authentication, clear the **Enable** check box.

You must also add information about the LDAP server and the LDAP attributes that are used for authentication to the ARS.CFG file for the instance. Then, after enabling LDAP support, you must stop and restart the Content Manager OnDemand server for the changes to take effect.

Other considerations

- OS400 security integration is not supported when you are using LDAP. When LDAP is enabled, you should disable OS400 security integration by editing the ARS.INI file. For the instance that is using LDAP, change `SRVR_FLAGS_SECURITY_EXIT=1` to `SRVR_FLAGS_SECURITY_EXIT=0`.
- After you disable OS400 security integration, the password for the administrative user QONDADM is blank. It is recommended that you immediately change the password for QONDADM. Do not delete QONDADM from the instance.
- After you disable OS400 security integration, it is recommended that you change the system parameters of the instance to set a minimum password length. To do this, log on to the Content Manager OnDemand Administrator client, right-click the instance with LDAP enabled, select System Parameters, and then set the Minimum Password Length.
- After you disable OS400 security integration, you must specify a password when you add a user to Content Manager OnDemand.
- If you enable LDAP on an existing instance, all existing users will have a blank password. Use the Content Manager OnDemand Administrator client to set a password for each user.
- Even when you are using LDAP, you must still add your users to the Content Manager OnDemand instance. The Content Manager OnDemand user ID must match the value that is returned by the LDAP server in the field that is mapped to `ARS_LDAP_MAPPED_ATTRIBUTE` in your ARS.CFG file.
- If the LDAP authentication fails, Content Manager OnDemand will attempt its normal logon process by using the user ID and password that was entered. This permits users that are not in the LDAP directory to access Content Manager OnDemand.

LDAP authentication processes

You can use LDAP (Lightweight Directory Access Protocol) to authenticate to the Content Manager OnDemand server.

You can logon to the Content Manager OnDemand server without the LDAP authentication. See the “Logging on a server” link for more information.

Related information:

 Collecting data for LDAP authentication problems in Content Manager OnDemand server

OnDemand logon with the LDAP authentication and anonymous binding

Some LDAP servers are set up to allow initial connection to the LDAP server or an anonymous bind.

An LDAP server that is set up to allow anonymous bind is open to the world. Anyone can connect to the LDAP server and search. It is like a building with its front doors unlocked. Anyone can walk in and look for a name from the directory in the lobby. In this example, even though the building is open, the individual offices might not be open. The individual offices are analogous to the entries in the LDAP server.

After a record is found, you might need to supply a password to access the information in that record (analogous to a key for a locked office). An example of an LDAP server that allows the anonymous bind is a company's intranet server. With this type of open LDAP server, the only logon information that an employee of that company would need is a user ID and password.

The anonymous bind LDAP authentication process:

- Connect to the LDAP server (called initial bind) without a user ID or password.
- If the initial bind is successful, search for an entry under the bind attribute name that is specified in **ARS_LDAP_BIND_ATTRIBUTE** at the location specified by **ARS_LDAP_BASE_DN**.
- If an entry is found, perform a second bind (logon) to it by using the user ID and password that was entered into the Content Manager OnDemand logon panel. If the second bind is successful, locate the value under the mapped attribute name that is specified in **ARS_LDAP_MAPPED_ATTRIBUTE**.
- If the value is found, return it to Content Manager OnDemand.

Configuring OnDemand with the anonymous bind set to true:

1. The user types the user ID and password in the **Logon to a Server** panel of the Content Manager OnDemand client.
2. The user ID and password are sent to the Content Manager OnDemand server to be authenticated.
3. The Content Manager OnDemand LDAP authentication component connects to the LDAP server. If the LDAP server is down or cannot be reached, the flow is based on the setting of **ARS_LDAP_OD_AUTHORITY_FALLBACK**.
 - If **ARS_LDAP_OD_AUTHORITY_FALLBACK=FALSE** (default) then the logon fails.
 - If **ARS_LDAP_OD_AUTHORITY_FALLBACK=TRUE**, then the logon reverts back to the Content Manager OnDemand logon and continues.
4. The Content Manager OnDemand LDAP authentication component searches the LDAP server for the user ID that is under the attribute name specified in the **ARS_LDAP_BIND_ATTRIBUTE** configuration parameter.

Remember: The user ID can be an e-mail address. If the attribute name does not exist on the LDAP server, the logon to the Content Manager OnDemand server fails.

- If the attribute name is correct but the user ID does not exist on the LDAP server, then the flow is based on the setting of **ARS_LDAP_OD_AUTHORITY_FALLBACK**.
 - If the user ID exists on the LDAP server and the password is wrong, the logon to the Content Manager OnDemand server fails.
 - If the user ID exists on the LDAP server and the password is correct, the LDAP server returns a value to the Content Manager OnDemand server.
5. You tell the LDAP server what to return by setting the value on the configuration parameter **ARS_LDAP_MAPPED_ATTRIBUTE** to an attribute or field name that is known to the LDAP server. If this attribute name does not exist on the LDAP server, the logon to the Content Manager OnDemand server fails.
- If the attribute name is correct, a value is returned. For example, if the ID that is entered on the Content Manager OnDemand logon screen is J12345, **ARS_LDAP_MAPPED_ATTRIBUTE** is set to odid, and the J12345 attribute odid is jasson1, then jasson1 is returned.
 - If the returned value matches a user ID that is stored in the Content Manager OnDemand server database, the user logs on. Otherwise, the logon fails. The Content Manager OnDemand password is not checked if the LDAP authentication is successful.

Remember: If the user ID and password case sensitivity is disabled (the default), the user ID and password are converted to uppercase before they are sent to the LDAP server. The returned user ID string from the LDAP server is then converted to uppercase and compared with the user IDs that are stored in the database.

If the Content Manager OnDemand user ID and password sensitivity is enabled, and the IDs that are stored in the Content Manager OnDemand database are in uppercase (for example, as a result of case insensitivity in the past), the ID that is returned by LDAP must be in uppercase otherwise the logon fails.

OnDemand logon with the LDAP authentication and non-anonymous binding

Some LDAP servers are set up to allow initial connection to the LDAP server or a non-anonymous bind.

In the previous scenario, the term "anonymous bind" was used. An LDAP server can also be set up to disallow the "anonymous bind", or to allow only the "non-anonymous bind." The Microsoft® Windows® Active Directory® (AD) server can only be set up using "non-anonymous bind.". When an LDAP server is set up this way, it is not "open" to the world. It is like a locked building that requires a key or badge to enter. In this example, the individual offices might or might not be locked. The individual offices are analogous to the entries in the LDAP server.

There are essentially two logons to an LDAP server that disallows the "anonymous bind" during the LDAP authentication process. The first logon (the initial bind) is required to gain connection to the LDAP server, and the second logon (the bind) is required to gain access to the entry or the record.

To connect to this type of LDAP server, you must supply a valid user ID and password at the initial bind time. It is like entering a locked building that requires a key or badge to enter. The initial bind ID and password are stored in the LDAP configuration under the parameters **ARS_LDAP_BIND_DN** and **ARS_LDAP_BIND_DN_PWD**,

and are not entered by the user at logon time. When you are connected, you can perform the search. When you find a record, you might need to supply a password to access the record's information.

The non-anonymous bind LDAP authentication process:

- Connect to the LDAP server (called initial bind) with a user ID and password that are specified under the parameters **ARS_LDAP_BIND_DN** and **ARS_LDAP_BIND_DN_PW**.
- If the initial bind is successful, search for an entry under the bind attribute name that is specified in **ARS_LDAP_BIND_ATTRIBUTE** at the location specified by **ARS_LDAP_BASE_DN**.
- If an entry is found, perform a second bind (logon) to it using the user ID and password entered into the Content Manager OnDemand logon panel.
- If the second bind is successful, locate the value under the mapped attribute name that is specified in **ARS_LDAP_MAPPED_ATTRIBUTE**.
- If a value is found, return it to Content Manager OnDemand.

Configuring OnDemand with the non-anonymous bind set to false:

1. The user types the user ID and password in the **Logon to a Server** panel of the Content Manager OnDemand client.
2. The user ID and password are sent to the Content Manager OnDemand server to be authenticated.
3. The Content Manager OnDemand LDAP authentication component connects to the LDAP server. If the LDAP server is down or cannot be reached, the flow is based on the setting of **ARS_LDAP_OD_AUTHORITY_FALLBACK**.
 - If **ARS_LDAP_OD_AUTHORITY_FALLBACK=FALSE** (default) then the logon fails.
 - If **ARS_LDAP_OD_AUTHORITY_FALLBACK=TRUE**, then the logon reverts back to the Content Manager OnDemand logon and continues.
4. If the LDAP server can be reached, the Content Manager OnDemand LDAP authentication component connects to the LDAP server with the LDAP user ID and password that are specified in the **ARS_LDAP_BIND_DN** and **ARS_LDAP_BIND_DN_PWD** parameters.
 - If the user ID and password are valid, the connection is established. Depending on the LDAP server setup, this user ID might need to have some administrative authority.
 - For the Windows Active Directory server, the default is that any valid user ID that is a member of the domain can be used in this initial logon to the LDAP server.
 - This is the first logon. If this logon fails, the logon to the Content Manager OnDemand server fails.
5. After the connection is made, the Content Manager OnDemand LDAP authentication component searches the LDAP server for the user ID that is under the attribute name specified in the **ARS_LDAP_BIND_ATTRIBUTE** configuration parameter.

Remember: The user ID can be an e-mail address. If the attribute name does not exist on the LDAP server, the logon to the Content Manager OnDemand server fails.

- If the attribute name that is specified in the **ARS_LDAP_BIND_ATTRIBUTE** parameter is correct but the user ID does not exist on the LDAP server, then the flow will be based on the setting of **ARS_LDAP_OD_AUTHORITY_FALLBACK**.

For example, if **ARS_LDAP_BIND_ATTRIBUTE=mail** and the user entered **jscott@us.ibm.com** in the Content Manager OnDemand logon panel, but this e-mail address does not exist on the LDAP server.

- If the user ID exists on the LDAP server and the password is wrong, the logon to the Content Manager OnDemand server fails.
- If the user ID exists on the LDAP server and the password is correct, the LDAP server returns a value to the Content Manager OnDemand server.
-
- If the attribute name is correct, a value is returned. If the returned value matches a user ID that is stored in the OnDemand server database, the user logs on. Otherwise, the logon fails. The Content Manager OnDemand password is not checked if the LDAP authentication is successful.

Remember: If the user ID and password case sensitivity is disabled (the default), the user ID and password are converted to uppercase before they are sent to the LDAP server. The returned user ID string from the LDAP server is then converted to uppercase and compared with the user IDs that are stored in the database.

If the Content Manager OnDemand user ID that is stored is in uppercase, if the user ID case-sensitivity is enabled, and if the returned user ID is in lowercase, then the logon fails.

Bypassing the LDAP server authentication

The LDAP server authentication can be bypassed. The logon process is reverted back to the normal Content Manager OnDemand logon.

The LDAP server authentication is bypassed if one of the following conditions is true. The logon process is reverted back to the normal Content Manager OnDemand logon if:

- The LDAP server is down or cannot be reached and **ARS_LDAP_OD_AUTHORITY_FALLBACK=TRUE**
- The user ID does not exist on the LDAP server for the bind attribute and **ARS_LDAP_OD_AUTHORITY_FALLBACK=TRUE**

The LDAP authentication might fail if one of the following conditions is true:

- The initial bind fails because of the incorrect user ID, password, or both.
- The second bind fails because of the incorrect user ID, password, or both.
- The attribute name that is specified in the **ARS_LDAP_BIND_ATTRIBUTE** configuration parameter does not exist on the LDAP server.
- The returned value from the LDAP server does not match any existing Content Manager OnDemand user ID.
- The returned value from the LDAP server matches an existing ID but it is in the wrong case and the Content Manager OnDemand user ID case sensitivity has been turned on.
- When the Content Manager OnDemand user ID sensitivity option is off, the returned value is converted to the uppercase characters. If the stored Content Manager OnDemand ID is in lower or mixed case, the logon can fail.

Depending on the LDAP server type and setup, the response you get from the LDAP authentication might not always be the same. The best way to diagnose the LDAP authentication problems is to get a detail trace and examine it. The trace can tell you which step went wrong in the LDAP authentication process.

Sample LDAP configurations

The Content Manager OnDemand user ID of administrators is not subjected to the LDAP authentication.

You can view LDAP parameters in the “Configuring the ARS.CFG file” section of the *IBM Content Manager OnDemand for i Common Server: Planning and Installation Guide*.

Anonymous Bind LDAP Server:

```
ARS_LDAP_SERVER=ldap1.yourcompany.com
ARS_LDAP_PORT=
ARS_LDAP_USE_SSL=FALSE
ARS_LDAP_BASE_DN=ou=yourgroup,o=yourcompany.com
ARS_LDAP_BIND_DN=
ARS_LDAP_BIND_DN_PWD=
ARS_LDAP_BIND_ATTRIBUTE=mail
ARS_LDAP_MAPPED_ATTRIBUTE=userid
ARS_LDAP_ALLOW_ANONYMOUS=TRUE
ARS_LDAP_BIND_MESSAGES_FILE=
```

Microsoft Active Directory (AD) server:

```
ARS_LDAP_SERVER=adserver.yourcompany.com
ARS_LDAP_PORT=
ARS_LDAP_USE_SSL=FALSE
ARS_LDAP_BASE_DN=dc=ondemand,dc=yourdomain,dc=local
ARS_LDAP_BIND_DN=CN=userid,OU=users,OU=yourunit,OU=yourcity,DC=dc1,DC=dc2,DC=local
ARS_LDAP_BIND_DN_PWD=password
ARS_LDAP_BIND_ATTRIBUTE=cn
ARS_LDAP_MAPPED_ATTRIBUTE=sAMAccountName
ARS_LDAP_ALLOW_ANONYMOUS=FALSE
ARS_LDAP_BIND_MESSAGES_FILE=
```

Microsoft Active Directory Application Mode (ADAM) server:

```
ARS_LDAP_SERVER=adamserver.yourcompany.com
ARS_LDAP_PORT=
ARS_LDAP_USE_SSL=FALSE
ARS_LDAP_BASE_DN=ou=yourlocation,o=yourcompany
ARS_LDAP_BIND_DN=cn=admin,o=yourcompany
ARS_LDAP_BIND_DN_PWD=password
ARS_LDAP_BIND_ATTRIBUTE=mail
ARS_LDAP_MAPPED_ATTRIBUTE=cn
ARS_LDAP_ALLOW_ANONYMOUS=FALSE
ARS_LDAP_BIND_MESSAGES_FILE=
```

IBM Tivoli Directory server (TDS) with SSL:

```
ARS_LDAP_SERVER=yourtds.yourcompany.com
ARS_LDAP_PORT=
ARS_LDAP_USE_SSL=TRUE
ARS_LDAP_KEYRING_FILE=/usr/lpp/ars/config/ondemand.kdb
ARS_LDAP_KEYRING_LABEL=LDAP Label
ARS_LDAP_BASE_DN=ou=yourlocation,o=yourcompany
ARS_LDAP_BIND_DN=cn=root
ARS_LDAP_BIND_DN_PWD=password
ARS_LDAP_BIND_ATTRIBUTE=email
ARS_LDAP_MAPPED_ATTRIBUTE=sn
ARS_LDAP_ALLOW_ANONYMOUS=FALSE
ARS_LDAP_BIND_MESSAGES_FILE=
```

Novel eDirectory server:

```
| ARS_LDAP_SERVER=yournds.yourcompany.com  
| ARS_LDAP_PORT=  
| ARS_LDAP_USE_SSL=FALSE  
| ARS_LDAP_BASE_DN=ou=yourlocation,o=yourcompany  
| ARS_LDAP_BIND_DN=cn=admin,ou=users,o=yourcompany  
| ARS_LDAP_BIND_DN_PWD=password  
| ARS_LDAP_BIND_ATTRIBUTE=mail  
| ARS_LDAP_MAPPED_ATTRIBUTE=cn  
| ARS_LDAP_ALLOW_ANONYMOUS=FALSE  
| ARS_LDAP_BIND_MESSAGES_FILE=
```

| **Sun Java Directory server (JDS)**

```
| ARS_LDAP_SERVER=yourjds.yourcompany.com  
| ARS_LDAP_PORT=  
| ARS_LDAP_USE_SSL=FALSE  
| ARS_LDAP_BASE_DN=ou=boulder,o=yourcompany  
| ARS_LDAP_BIND_DN=cn=Directory Manager  
| ARS_LDAP_BIND_DN_PWD=password  
| ARS_LDAP_BIND_ATTRIBUTE=mail  
| ARS_LDAP_MAPPED_ATTRIBUTE=cn  
| ARS_LDAP_ALLOW_ANONYMOUS=FALSE  
| ARS_LDAP_BIND_MESSAGES_FILE=
```

Examples

The examples that follow demonstrate how to use the Content Manager OnDemand component of IBM Navigator for i to:

- Create output queue monitor definitions which contain information about the output queue to be monitored, the handling of the spooled files being archived, and information about starting and ending the monitor job. The example uses an output queue monitor to archive spooled files, but you can also monitor IFS directories to archive IFS files as well.
- Create a disk pool storage group definition to identify an IBM i Auxiliary Storage Pool (ASP) that the archive storage management (ASM) process may use as storage media when migrating archived data.
- Create optical storage groups which are used to group sets of optical volumes together to store related data.
- Add optical volumes which are available for use by Content Manager OnDemand.
- Create tape devices that are available for use by Content Manager OnDemand if tape archival is selected.
- Add tape volumes which are available for use by Content Manager OnDemand
- Create migration policies which provide migration and storage media characteristics for data archived using Content Manager OnDemand.

The examples that follow also demonstrate how to use the Content Manager OnDemand Administrator client to:

- Add users and groups to the system. Users on the system obtain permission to open folders and access application group data from a group.
- Define a server printer that can be used as the default printer for an application.
- Add a report to the system. To add a report, create an application group, an application, and a folder.

System configuration

The example assumes that an external customer (Customer XYZ) accesses the system by using the OnDemand Web Enablement Kit (ODWEK) feature. However, for purposes of demonstrating how the system works and how to use the administrative client to implement the requirements, these users could easily represent another department within the company.

Here are the assumptions about the system:

- Two groups of users: the customer service department and users at Customer XYZ.
- A sample telephone bill report. The telephone report is generated by an application program running on an IBM i system with IBM Content Manager OnDemand installed.
- Disk storage. When a report is loaded into the system, Content Manager OnDemand stores a copy on disk and maintains the report on disk for 60 days.
- Archive storage. When a report is loaded into the system, Content Manager OnDemand stores a copy of the report in archive storage. The archive storage manager maintains the report on optical storage for five years.

- System parameters. Users must log on to the system with a password. Set the minimum password length to four characters. Accept the defaults for the other system parameters.

Creating monitor definitions

About this task

To create an output queue monitor definition:

Procedure

1. Access IBM Navigator for i.
2. Click Content Manager OnDemand, and verify the current instance.
3. Click **Output Queue Monitors** and then click **Add**. You can also create directory monitor definitions, but the example uses an output queue monitor to archive spooled files.
4. When the Output Queue Monitor Definition window appears, the general page will be displayed. Specify the name and library of the output queue to monitor. For this example, specify the output queue as MONITORQ in library MYLIB. The job name is automatically set to the same name as the monitored output queue you specified. This can be changed, but for the example leave it blank so that the job name is set to MONITORQ. You may also provide a brief description for the output queue monitor definition. For this example, leave the default job description name and library, or click **Browse** to select a different job description.

When the monitor job selects a spooled file from the specified output queue for processing, it needs to determine which application group to associate with the spooled file so that the file can be archived correctly. The only data available to the monitor are the attributes of the selected spooled file. The application group name must be derived from the contents of one of these attributes.

5. Use the pull down to select the attribute that the monitor should use for comparing to find the application group with the matching name. For the example, use Spooled file name. You can also specify what the monitor should check next and last if it does not find a match on check first. In this example, leave **Check next** and **Check last** with the default value of **(none)**.
6. You can check the box to have the monitor job determine the application name used or you can specify what the monitor should check for comparing to find the application name. In this example, check the box and let the monitor job determine the application name.
7. Once the general page has been completed you can click on the **Start and End Method** tab. There are several options for starting and ending the monitor job. Accept the defaults for this example which are to start and end the monitor manually. However you can specify whether the monitor will be started from a job scheduler or when a subsystem starts. You can also specify to have the monitor job end at a specific time of day, after a certain period of time, after all entries have been processed from the output queue, or you can specify the end time as a parameter when the monitor is started.

Tip: The first time you start a monitor for a particular output queue, it is best to do it when there are no spooled files in the output queue. When a monitor is started for the first time, an empty data queue with the same name as the output queue is created, which will receive entries for all spooled files that appear in the output queue in Ready status. Once the output queue monitor has been started (and therefore the data queue has been created), you can then

begin moving the spooled files that you wish to capture into the output queue. The data queue entries that get created will trigger the monitor to process each spooled file.

8. Next click on the **Actions To Take After Archiving** tab. This specifies what Content Manager OnDemand should do with a spooled file once it has been processed. If the spooled file is processed successfully, the options are to have it deleted from the system or moved to an output queue you specify. In the example, the spooled file will be moved to a processed output queue called PROCESSED in library QUSRRDARS. You must also specify where you want the spooled files to be sent should they fail to process. In the example, specify ERROR in library QUSRRDARS as the error queue.
9. Click **OK** to create the output queue monitor definition.

Creating disk pools

About this task

To create a Disk pool to use with IBM Content Manager OnDemand:

Procedure

1. Access IBM Navigator for i.
2. Click Content Manager OnDemand, and verify the current instance.
3. Click **Disk Pools** and then click **Add**.
4. Identify an IBM i Auxiliary Storage Pool (ASP) number that the archive storage management (ASM) process may use as storage media when migrating archived data. For the example, specify the **Pool number** as 3. You may also provide a brief description for the disk pool.
5. Select **Primary** or **Backup** to specify whether the disk pool storage group stores primary or backup copies of archived data. For the example, the type for the disk pool is **Primary**. Click **OK** to create the disk pool definition.

Creating optical storage groups

About this task

To create an optical storage group in IBM Content Manager OnDemand:

Procedure

1. Access IBM Navigator for i.
2. Click Content Manager OnDemand, and verify the current instance.
3. Click **Optical Storage Groups** and then click **Add**.
4. The optical storage group is used to group sets of optical volumes together for use by Content Manager OnDemand. For the example, call the optical storage group OPTSTG. You may also provide a brief description.
5. The optical volumes are rewritable and support the dynamic reuse of space without re-initializing the entire volume. So for this example, check the **Volume full reset** checkbox and set the **Free space threshold percent** at 40. This indicates when the volume full flag is reset. In the example, when there is 40 percent free space available on the volume, the full flag is reset and the volume can again be used for newly migrated data.
6. This optical storage group will contain primary volumes; mark the type as **Primary**.

7. Click **OK** to create the optical storage group.

Adding optical volumes

About this task

To add optical volumes to IBM Content Manager OnDemand:

Procedure

1. Access IBM Navigator for i.
2. Click Content Manager OnDemand, and verify the current instance.
3. Click **Optical Volumes** and then click **Add**.
4. The volume name specified must match the name used when the volume was initialized in the optical library. A single disk has two sides which are considered volumes. As you will see in a later step, the opposite side volume name must be specified. The name for this volume is OND00001.
5. If the volumes being defined will belong to a Primary optical storage group, then click **Primary**. If the volumes being defined will contain backup versions of the data, click **Backup**. In the example, the optical storage group was defined as **Primary**, which means that all the volumes assigned to that storage group must also be of type **Primary**. Therefore, mark this volume as **Primary**.
6. Enter the capacity, in megabytes, of one side of the disk (one volume).
7. To select the Optical media family, use the pull down and select the appropriate media type. Review the help text for explanations of the possible values. Then, click the pull down to select the Optical storage group to which the optical volumes belong, such as the OPTSTG storage group that was created earlier. The optical library name is required when the media family is LANR or LANW. It can be used for documentation purposes for all other media families.
8. Specify the name of the opposite side of this volume as OND00002.
9. Leave the **Volume is full** checkbox unchecked. This flag will automatically be set to full when the optical volume has reached its capacity.
10. Click **OK** to create the optical volume definitions.

Adding tape devices

About this task

To add a tape device to IBM Content Manager OnDemand:

Procedure

1. Access IBM Navigator for i.
2. Click Content Manager OnDemand, and verify the current instance.
3. Click **Tape Devices** and then click **Add**.
4. First, the name of the tape device you want to add to Content Manager OnDemand must be specified. The name must match the name of an existing IBM i tape device description. In this example, the name of the tape device is TAP01. You may also provide a brief description for your tape device.
5. From the list of supported media types, select the media type for this tape device. In this example, the media type is 8MM.
6. Click the **Add -->** button to add the 8MM media type to the **Read operations** list. Then click the **Add -->** button to add the 8MM media type to the **Write**

operations list. If the device you are using has an automatic cartridge loader, you can specify the number of cartridges the loader holds. In the example, leave this at 0. If you use a media library such as the IBM 3494 tape library with Content Manager OnDemand, specify the library name here. For this example, leave the media library blank.

7. Click **OK** to create the tape device.

Adding tape volumes

About this task

To add tape volumes to IBM Content Manager OnDemand:

Procedure

1. Access IBM Navigator for i.
2. Click Content Manager OnDemand, and verify the current instance.
3. Click **Tape Volumes** and then click **Add**.
4. First, specify the name of the tape volume that you want to make available to Content Manager OnDemand. The volume name must match the name you used when you initialized the tape volume on your IBM i server. For the example, the name of the volume is OND001.
5. Next, specify the total capacity of the tape volume in megabytes. For the example, specify a capacity of 2300 MB.
6. For media type, use the pull down to select the correct type. For the example, select 8MM, and leave media device library blank.
7. The migration policy, which will be created later, will request a tape backup. So, for this example, make the volume type Backup. If, instead, the plan was to use this tape volume for the primary archive media for some of the Content Manager OnDemand data, you would leave the volume type set to **Primary**. For this example, however, set the volume type to **Backup**. Make sure **Tape is full** is not checked.
8. Click **OK** to add the tape volume to Content Manager OnDemand.

Creating migration policies

About this task

To create a new migration policy:

Procedure

1. Access IBM Navigator for i.
2. Click Content Manager OnDemand, and verify the current instance.
3. Click **Migration Policies** and then click **Add**.
4. First you must specify the name of your migration policy. The name can be up to 60 characters long, but must not duplicate the name of another policy within the same instance. For this example, choose POLICY1 as the new policy name. You may also provide a brief description of the migration policy you are creating.
5. You can choose whether or not you want to enable aggregation. If aggregation is enabled, indicate the maximum size and maximum number of days before Content Manager OnDemand closes the aggregate. (See "Migration policies" on page 59 for more information on migration policies and aggregation.)

6. You may also specify if a tape backup is requested. Request a tape backup and select the **Media** type by clicking on the pull down and selecting **8MM**.
7. You set up storage levels within the migration policy to define the route that your archived data follows as it moves through different storage media in the migration process. At least one level must be defined. The order of the levels in the list determines the migration sequence.
 - a. To create the first level, click the **Add after** button. This opens the Migration Policy Storage Level Definition panel.
 - b. Assign a level identifier. This identifier must be unique within this migration policy. For this example, you might use 0010 as the level identifier. You may also provide a brief description of this level. If you are setting this up in advance of when you need it, you can check **Disabled** so that it is not used until you update the level and uncheck the **Disabled** checkbox.
 - c. For the primary media type and primary storage group, click the pulldown to select your preferred media and storage group for the level, and then set the duration for which your data should reside at this level. Leave the **Backup copy** and **Stage to disk** values unchanged.
 - d. Click **OK**. This will add the storage level in the storage levels window within the migration policy.
 - e. If you wish to add the next level, select the level that was just added and then click the **Add after** button. Continue as needed, following the same procedure.
 - f. For example, you might specify the level identifier as 0020 with a brief description. The media is **Optical** and the number of days is 1825 (5 years).
 - g. Choose the optical storage group that was created earlier, **OPTSTG**, as the primary storage group.
 - h. Click **OK**, and you should now see both storage levels listed for the migration policy. In both storage levels, the **Create backup copy** option was not selected. However, you can select this option and specify a backup storage group causing the migration process to create a duplicate copy of the archived data when it is moved to this level.
8. Click **OK** to create the migration policy. Clicking **OK** will also create a Content Manager OnDemand storage set of the exact same name. When you define an application group to Content Manager OnDemand, you can select the storage set name that matches this migration policy name, which will cause Content Manager OnDemand to archive that application group data as defined in this migration policy.

Adding users

Adding a user with the administrative client

You can use the administrative client to add a user.

Procedure

To add a user:

1. Choose a server and select **Users**.
2. Pick one of two ways to add a user.
3. Define the properties of the user by completing fields in the **Add a User** dialog box.

4. Optional: Add the user to groups. See the online help for more details.
5. Optional: Add the user to application groups and set application group permissions.
6. Optional: Add the group to folders and set folder permissions.
7. When finished, add the user by clicking **OK** in the **Add a User** dialog box.

What to do next

1. To add a user, the logon user ID must be a user with create users authority, a user administrator, or a system administrator.
2. To make problem determination easier, IBM recommends that each user profile that will be used to load reports (ADDRPTOND, STRMONOND, arslod, arsdod add) have a home directory in IFS on the IBM i server. If a home directory exists for a user profile running one of the load commands, any output files and temporary files will be placed in that user profile's home directory rather than "lost" within the entire root directory on the IBM i system. To create a home directory for a user profile, issue the following command:
MD '/home/usrprf/' where *usrprf* is the name of the user profile that will be used to load reports.

Choose a server

Procedure

1. On the left side of the administrator window, click the name of the server to which you want to add the user.
2. Expand the areas of the server. Double click the server name or click the + (plus) to the left of the server name.
3. Select **Users**.

Two ways to add a user

You can add a user by using the **New User** command. You can also add a user by copying an existing user definition.

New User command

From the **File** menu, select **New User** to open the **Add a User** dialog box.

Copy command

You can use the **Copy** command to add a user. In the User ID list, point to the user that you want to copy and click the right mouse button. Select **Copy** from the pop-up menu to open the **Add a User** dialog box. The fields in the dialog box contain information copied from the user you selected. At a minimum, you need to change the User ID. (User IDs must be unique to the server.)

Adding the user

About this task

In the **Add a User** dialog box, click **OK**. The administrative client adds the user to the database and returns to the main window.

Examples

First, review the requirements of the users that need to access the telephone bill reports.

- One set of user IDs for the customer service department. Identify one user as a user administrator. The user administrator can add other users.
- One user ID for Customer XYZ.

On the example system, users obtain permissions from groups. This means that the users are not added to application groups and folders. When groups are defined, users are added to the groups.

Important: The steps that follow do not show how to add all of the sample users to the system. Two users will be added; you can repeat the steps to add the others.

Adding the customer service users

About this task

Use the **New User** command to add a user.

Procedure

1. First, point to **Users** and click the right mouse button. From the pop-up menu, select **New User** to open the **Add a User Dialog** box.
2. In the **User ID** field, enter the name of the user: CSR1. If your IBM Content Manager OnDemand user IDs are linked to your IBM i user profiles, then this user ID must also exist as an i user profile.
3. Accept the UID generated by Content Manager OnDemand.
4. Set the user's initial password to the user ID. In the **Password** field, enter: CSR1. If your Content Manager OnDemand user IDs are linked to your IBM i user profiles, then this password will be ignored.
5. In **Verify Password** field, enter: CSR1
6. In the **Description** field, enter: Customer Service Representative
7. Under **User Type**, select **User Administrator**. The user will be able to maintain user IDs on the system.
8. Click the **User Information** tab.
9. Complete the fields on the User Information page, such as the **Name**, **Department**, and **Phone Number**.
10. Click the **General** tab.
The user will obtain application group and folder permissions from a group. Add the user to the group when the group is added to the system. Therefore, do not add the user to application groups, folders, or groups at this time.
11. At this point, the properties of the user meet the requirements. Click **OK** to add the user.

Adding another user:

About this task

Use the **Copy** command to add another user.

Procedure

1. Point to **CSR1** and click the right mouse button. From the pop-up menu, select **Copy** to open the **Add a User** dialog box.
2. In the **User ID** field, replace CSR1 with CSR2. If your IBM Content Manager OnDemand user IDs are linked to your IBM i user profiles, then this user ID must also exist as an i user profile.
3. Accept the UID generated by Content Manager OnDemand.

4. Set the user's initial password to the Content Manager OnDemand user ID. In the **Password** field, enter: CSR2. If your Content Manager OnDemand user IDs are linked to your i user profiles, then this password will be ignored.
5. In the **Verify Password** field, enter: CSR2
6. Under **User Type**, select **User**.
7. Click the **User Information** tab.
8. Replace the information in the fields on the User Information page.
9. Click the **General** tab.
10. At this point, the properties of the user meet the requirements. Click **OK** to add the user.

Adding groups

About this task

Important: To add a group, the logon user ID must be a user with create groups authority or a system administrator.

In general, here is how you work in the administrative client to add a group:

Procedure

1. Choose a server and select **Groups**.
2. Pick one of two ways to add a group.
3. Define the properties of the group by completing fields in the **Add a Group** dialog box.
4. Optional: Assign a group owner.
5. Optional: Add users to the group.
6. Optional: Add the group to application groups and set application group permissions.
7. Optional: Add the group to folders and set folder permissions.
8. When finished, add the group by clicking **OK** in the **Add a Group** dialog box.

Choose a server

Procedure

1. On the left side of the window, click the name of the server to which you want to add the group.
2. Expand the areas of the server. Double click the server name or click the + (plus) to the left of the server name.
3. Select **Groups**.

Two ways to add a group

This section explains how to add a group by using the **New Group** command. You can also add a group by copying an existing group definition.

New Group command

From the **File** menu, select **New Group** to open the **Add a Group** dialog box.

Copy command

You can use the **Copy** command to add a group. In the Name list, point to the group that you want to copy and click the right mouse button. Select **Copy** from the pop-up menu to open the **Add a Group** dialog box. The fields in the dialog box contain information copied from the group you selected. At a minimum, you need to change the group name (group names must be unique to the server).

Remember: These are not the same as group profiles in IBM i. However, the names may match if you find that easier to maintain.

Adding users

About this task

You can add one or more users to the group. Complete the following steps to add a user to a group.

Procedure

1. From the List of Users list, select the user.
2. Click **Add**. The administrative client moves the user to the Users in the Group list.
3. To remove a user from the group, select the user in the Users in the Group list and click **Remove**.

Adding the group

About this task

In the **Add a Group** dialog box, click **OK**. The administrative client adds the group to the database and returns to the main window.

Examples

Review the requirements of the groups that need to access the sample telephone bill reports.

- Customer service group. Users that belong to the group can open the telephone bill report folder and query documents stored in the telephone bill report application group. When the report is added to the system, the group will be added to the application group and the folder. Identify a group owner. The group owner can add new customer service users to the group and remove users from the group.
- Customer XYZ group. Users that belong to the group can also open the telephone bill report folder and query documents stored in the telephone bill report application group. However, you can limit access to documents that contain a specific customer name and account number. When the report is added to the system, add the group to the application group and the folder and specify the necessary restrictions.
- Users. Add the users that were defined in “Adding users” on page 90 to the groups.

Adding the customer service group

About this task

Use the **New Group** command to add the group.

Procedure

1. First, point to **Groups** and click the right mouse button. From the pop-up menu, select **New Group** to open the **Add a Group Dialog** box.
2. In the **Name** field, enter the name of the group: CustomerService
3. Accept the GID generated by IBM Content Manager OnDemand.
4. From the Group Owner list, select the user that can add users to and remove users from the group: CSR1

Important: The user that you select does not obtain permissions from the group unless you add the user to the group (see step 6). However, a group owner can add their user ID to the group at any time.

5. In the **Description** field, enter: Access to Telephone Bill Reports; call Leonard Little, x90565, for more information
6. From the List of Users list, select and add users to the groups: CSR1, CSR2, CSR3, CSR4, and CSR5
7. At this point, the properties of the group meet the requirements. Click **OK** to add the group.

Adding the Customer XYZ group About this task

Use the **Copy** command to add the group.

Procedure

1. Point to **Customer Service** and click the right mouse button. From the pop-up menu, select **Copy** to open the **Add a Group** dialog box.
2. In the **Name** field, replace CustomerService with CustomerXYZ
3. Accept the GID generated by IBM Content Manager OnDemand.
4. Use a system administrator to maintain the group. Therefore, do not assign a group owner. In the Group Owner list, replace CSR1 with *NONE.
5. Replace the contents of the **Description** field with: Access to Telephone Bill Reports by Customer XYZ
6. From the Users in the Group list, remove CSR1, CSR2, CSR3, CSR4, and CSR5.
7. From the List of Users list, add XYZ1.
8. Click **OK** to add the group.

Adding server printers

Before you begin

1. To add a server printer, the logon user ID must be a system administrator.
2. Server print is supported for spooled file data types of SCS, SCS-extended, AFPDS, and LINE.

About this task

In general, here is how you work in the administrative client to add a server printer:

Procedure

1. Choose a server and select **Printers**.
2. Pick one of two ways to add a server printer.

3. Define the properties of the server printer by completing fields in the **Add a Printer** dialog box. Search for *printer,adding* in the online help for more details.
4. When finished, add the server printer by clicking **OK** in the **Add a Printer** dialog box.

Choose a server

Procedure

1. On the left side of the main window, click the name of the server to which you want to add the server printer.
2. Expand the areas of the server. Double click the server name or click the + (plus) to the left of the server name.
3. Select **Printers**.

Two ways to add a server printer

This section explains how to add a server printer by using the **New Printer** command. You can also add a server printer by copying an existing server printer definition.

New Printer command

From the **File** menu, select **New Printer** to open the **Add a Printer** dialog box. The figure in “Adding server printers” on page 95 shows an example of the **Add a Printer** dialog box.

Copy command

You can use the **Copy** command to add a server printer. In the Name list, point to the server printer that you want to copy and click the right mouse button. Select **Copy** from the pop-up menu to open the **Add a Printer** dialog box. The fields in the dialog box contain information copied from the server printer you selected. At a minimum, you need to change the name (printer names must be unique to the server).

Adding the server printer

About this task

In the **Add a Printer** dialog box, click **OK**. The administrative client adds the server printer to the database and returns to the main window.

Examples

Add a server printer that can be selected as the default server printer for the application (added in “Adding a report” on page 97). The physical printer resides in the customer service department. By default, when users of the telephone bill report application select a document and choose the server printer command, IBM Content Manager OnDemand sends the document to this print device.

Adding the server printer

About this task

Use the **New Printer** command to add the server printer.

Procedure

1. First, point to **Printers** and click the right mouse button. From the pop-up menu, select **New Printer** to open the **Add a Printer** dialog box.
2. In the **Name** field, type the name of the server printer: Customer Service printer
3. In the **Description** field, type: Customer Service printer for the telephone bill report application
4. In the **Server Queue Name**, type: ip60cs If no library name is entered (as in the example), Content Manager OnDemand assumes *LIBL; otherwise, enter a value such as QUSRSYS/PRT04 to identify both the library name and the output queue name.
5. Accept the default Printer Type of **Printer**
6. At this point, the properties of the server printer meet the requirements. Click **OK** to add the server printer.

Adding storage sets

A storage set is automatically added each time you create a migration policy. The name of the storage set that is created is an exact match of the name of the migration policy. See “Creating migration policies” on page 89 for details on migration policies.

Adding a report

About this task

When you define a report to the system, you typically add an application group, an application, and a folder.

- The application group identifies database and storage management information.
- The application identifies viewing, indexing, loading, and printing information.
- The folder provides users the ability to search for, retrieve, view, and print report data.

In general, here is how you work in the administrative client to define a report to the system:

Procedure

1. Choose a server. On the left side of the main window, click the name of the server on which you want to define the report.
2. Select the area.
 - a. First expand the areas of the server.
 - b. Double click the server name or click the + (plus) to the left of the server name.
 - c. Then select the area.
3. Pick one of two ways to add an application group, an application, and a folder. For example, use one of the following methods to add an application group:
 - Add a new application group. From the **File** menu, select **New Application Group** to open the **Add an Application Group** dialog box.
 - Copy an existing application group. In the Name list, point to the application group that you want to copy and click the right mouse button. Select **Copy** from the pop-up menu to open the **Add an Application Group** dialog box. The fields in the dialog box contain information copied from the application

group you selected. At a minimum, you need to change the Name (application group names must be unique to the server).

Important:

- To add an application group, the logon user ID must be a system administrator, an application group/folder administrator, or a user with create application groups authority.
 - To add an application, the logon user ID must be a system administrator, an application group/folder administrator, or a user with create application groups authority.
 - To add a folder, the logon user ID must be a system administrator, an application group/folder administrator, or a user with create folders authority.
4. Define the properties of the application group, application, and folder.
 5. When finished, add the application group, application, and folder by clicking **OK** in the add dialog box.

An example

This section describes how to define a sample telephone bill report to the system.

About the report

A telephone bill report typically contains hundreds of pages of line data. The report is logically segmented into statements. Users search for statements using a date and any combination of account number and customer name.

Most queries about a statement occur in the first 60 days after it is mailed to the customer. Little or no activity occurs a year after a statement is generated. For legal reasons, a statement must be maintained on the system for five years. The system should maintain index information in the most efficient way possible.

Two groups of users need to access the telephone bill reports. The customer service department is responsible for handling queries from customers. They answer questions about the statements, attach annotations to statements, and reprint and fax a replica of original statements. As part of a customer service initiative, Customer XYZ is permitted to access the system and retrieve and view their statements.

About the application group

Before adding the application group, review the database and storage management requirements.

Database requirements

Database requirements can be grouped in two categories: database tables and the database fields.

Database tables

- A database table contains index data from one or more reports
- An annotation field is not required in the database
- The date field is the segment field for the application group

Database fields

- Three database fields: account number, customer name, and report date

- Account number is the index, data type string, 16 bytes. However, only 13 bytes should be stored in the database, removing the - (dash) characters from the account number string.
- Customer name is a filter, data type string, 30 bytes
- Report date is a filter, data type date

Storage management requirements

The storage management requirements determine where, how, and how long IBM Content Manager OnDemand maintains the report and index data.

- Maintain a report for five years
- Copy documents to disk and maintain them for 60 days
- Copy documents to archive storage when the report is loaded into the system
- Delete a table of index data at a time

Adding the application group

About this task

Use the **New Application Group** command to add the application group.

Procedure

1. First, point to **Application Groups** and click the right mouse button. From the pop-up menu, select **New Application Group**. The administrative client opens the **Add an Application Group** dialog box.
The pages of the **Add an Application Group** dialog box organize information about the application group into sections. The tabs show which page you are on: General, Message Logging, Storage Management, Permissions, Field Definition, and Field Information.
2. Start by completing the General page. In the **Name** field, type the name of the application group.
3. In the **Description** field, type up to 120 characters of descriptive information about the application group.
4. Based on the database requirements, accept the recommended defaults for the rest of the fields on the General page. (You can click **Advanced** to see the default options.)
5. Select the **Message Logging** tab to specify the types of application group messages that Content Manager OnDemand should save in the system log.
6. Select **Retrieval**, **Database Queries**, and **Server Printing**. Clear all of the other check boxes
7. Select the **Storage Management** tab to provide information that Content Manager OnDemand uses to manage data stored in the application group.
8. From the Storage Set Name list, select the name of the storage set that Content Manager OnDemand will use to maintain the report on disk storage and in archive storage. Select the storage set and migration policy that was defined in "Adding storage sets" on page 97.
9. Under **Cache Data**, select **Cache Data for __ Days** and enter 60 in the space provided. Content Manager OnDemand should maintain the report on disk storage for 60 days. After that time, when a user views or prints the report, Content Manager OnDemand retrieves the report from archive storage.
10. Under **Life of Data and Indexes**, select **Expire in ____ Days** and enter 1825 in the space provided. This is the number of days (1825, or five years) that Content Manager OnDemand should maintain index data, documents, and

resources related to the report. After that number of days, Content Manager OnDemand can delete the report from the system. Note that the value for Life of Data and Indexes can be changed even after you have loaded data into Content Manager OnDemand for this application group. If you change this value after you have already loaded data, the change will affect the existing data that is already loaded (if the Expiration Type is set to Load in the application group definition), as well as any new data you load after the change is made.

11. Because Content Manager OnDemand should delete a table of index data at a time, from the Expiration Type list, select **Segment**
12. Select the **Permissions** tab to specify the types of report and application group functions that users can perform. For example, you can let users query report data, create logical views, print and fax pages of the report, and maintain the application group. You can specify default permissions and permissions for specific groups and users. Unless you specify otherwise, the person that creates the application group is given all application group permissions; no other users can access report data or maintain the application group. For the example system, other users obtain permissions from a group. The groups were added in “Adding the customer service group” on page 94 and “Adding the Customer XYZ group” on page 95.
13. First, add the customer service group. From the User/Groups list, select **+CustomerService**.
14. Select the **Access** check box.
15. Click **Add**.
16. Next, add the Customer XYZ group. From the User/Groups list, select **+CustomerXYZ**.
17. In the Annotation area, clear the **View** check box. Content Manager OnDemand also clears the **Add** check box. Customer XYZ users can view, print, fax, and copy documents but do not have permission to use the annotation tools provided by Content Manager OnDemand.
18. Click **Add**.
19. Select the **Field Definition** tab to define the database fields.
20. For the telephone bill report, define the three database fields: **Account number**, **Customer name**, and **Report Date**
21. To define a database field, type the name of the field in the **Database Field Name** field and click **Add**.
22. Select the **Field Information** tab to define the attributes of the database fields.
23. First, define field information for the **Account number** field.
 - a. From the Type list, select **Index**
 - b. Type the string length: 13

Important:

IBM recommends that you always define an **Application ID** field for the application group. An **Application ID** field is a database field that contains values that identify an application within the application group.

You may not think that you need an application ID field at the time that you create your application group, if you are adding an application group that will hold only one application at the current time. However, if you decide later to add other applications to the application group (or if you want to maintain multiple versions of your application definitions) and you then need

to define an **Application ID** field for the application group, you will not be able to do so because an **Application ID** field cannot be added after the application group is created. (All fields must be added during the original application group definition.) Also note that the application ID field can be hidden from users that do not require it to search for documents. See the online help for more information about the application ID field.

24. Next, define field information for the **Customer name** field.
 - a. From the Name list, select **custname**
 - b. From the Type list, select **Variable**
 - c. In the Length field, type 30
25. Next, define field information for the **Report date** field.
 - a. From the Name list, select **rdate**
 - b. From the Data Type list, select **Date**
 - c. Select the **Segment** check box
26. At this point, the properties of the application group meet the requirements. Click **OK** in the Add an Application Group window. The administrative client adds the application group to the database and returns to the main window.

About the application

Before adding the application, review the viewing, indexing, and loading requirements.

Viewing requirements

- Source data stored in IBM Content Manager OnDemand as AFP data
- Format data into pages and enhance the appearance with images and fonts
- Retrieve statements of one or more pages
- Define a default printer for the application. The default printer is where Content Manager OnDemand sends documents when users select the server print command.

Indexing requirements

- Source data is EBCDIC, code page 500
- Segment report into groups of pages, one statement in each group
- Identify the beginning of a statement using:
 - Skip-to-channel one (X'F1')
 - PAGE 1 (X'D7C1C7C54040F1')
- Generate three indexes for each statement: statement date, account number, and customer name
- Collect resources

Loading requirements

- Compress and store data in the most efficient method possible
- Application group database field names and index names match
- Date format is Mth d, yyyy
- Remove embedded - (dash) character from account number before storing value in the database

Adding the application

About this task

Use the **New Application** command to add the application.

Procedure

1. First, point to **Applications** and click the right mouse button. From the pop-up menu, select **New Application** to open the **Add an Application** dialog box.
The pages of the **Add an Application** dialog box organize information about the application into sections. The tabs show which page you are on: General, View Information, Indexer Information, Load Information, Logical View Fields, Logical Views, and Miscellaneous Options.
2. Start by completing the General page. In the **Name** field, type the name of the application.
3. In the **Description** field, type information about the application.
4. Click **Select** to open the **Application Groups** dialog box.
5. From the Names list, select **Telephone Bill Reports**.
6. Click **OK**.
7. Select the **View Information** tab to specify information needed by Content Manager OnDemand client programs to display the telephone bill report. This information is also used by the indexing program.
8. From the Data Type list, select **Line**. Note that although this example shows Line data type, most IBM i spooled files will be either SCS or AFP.
9. Enter the correct **Code Page**. For Line, SCS, and SCS-Extended data types, this field determines the code page used by the Content Manager OnDemand client to display the data. For Line, the default value is 500. For SCS and SCS-Extended, the default value is the code page of the server.
10. In the RECFM area, select **Fixed**. The report contains fixed length records, 133 bytes in length.
11. Select the **Indexer Information** tab.
12. From the indexer list, select OS/400® Indexer.
13. Process a sample report using the graphical indexer. In the Parameters Source area, select **Sample Data**.
14. Click **Modify** to open the **Open** dialog box to select a file that contains a sample of the actual report data.
15. Select **OS/400 Spooled File** to work with spooled files on the server. Identify how to find your spooled file by entering a user profile or output queue name, and then click **Retrieve List** to generate a list of spooled files to choose from. Once you select a specific spooled file, Content Manager OnDemand copies that sample data to your workstation for you to use for indexing.
The following figure shows the Report window.




Figure 3. Report window

16. Define trigger number one.
 - a. First, select any blank column in the first record.
 - b. Click the **Add a Trigger** icon to open the **Add a Trigger** dialog box. (See online help for a description of trigger fields.)
17. In the Columns to Search area, select **Carriage Control**.
18. Click **OK** to add the trigger.
19. Define trigger number two.
 - a. First, select the string **PAGE 1**.
 - b. Click the right mouse button.
 - c. From the pop-up menu, select **Trigger** to open the **Add a Trigger** dialog box.
20. Click **OK** to add the trigger.
21. Define field number one.
 - a. First, select the string **Customer XYZ** and enough blank characters to the right of the string to reserve enough space to hold the largest index value (30 characters) the field can contain. (The selected string length guide, which appears above the field, displays the number of selected characters.)

- b. Then click the right mouse button.
 - c. From the pop-up menu, select **Field** to open the **Add a Field** dialog box.
22. From the Trigger list, select **Trigger2**.
 23. Click **OK** to add the field.
 24. Define field number two.
 - a. First, select the string **May 11, 1996**.
 - b. Then click the right mouse button.
 - c. From the pop-up menu, select **Field** to open the **Add a Field** dialog box.
 25. From the Trigger list, select **Trigger2**.
 26. Click **OK** to add the field.
 27. Define field number three.
 - a. First, select the string **303-555-1212-95B**.
 - b. Then click the right mouse button.
 - c. From the pop-up menu, select **Field** to open the **Add a Field** dialog box.
 28. From the Trigger list, select **Trigger2**.
 29. Click **OK** to add the field.
 30. Define the indexes.
 - a. First, clear any selected areas of the report.
 - b. To define the first index, click the **Add an Index** icon to open the **Add an Index** dialog box.
 31. From the Attribute list, select **custname**.
 32. In the Break area, select **No**.
 33. In the Fields list, double-click **Field1**.
 34. Now define the second index. From the Attribute list, select **rdate**.
 35. In the Break area, select **No**.
 36. In the Fields list, double-click **Field2**.
 37. Now define the third index. From the Attribute list, select **acct**.
 38. In the Fields list, double-click **Field3**. Leave the **Break** setting set to Yes for this index. This will cause Content Manager OnDemand to watch for a change in the value of the acct index by using the change to indicate an end to one document and the beginning of the next when breaking up the input file into separate documents. For more information about this **Break** setting, see the *IBM Content Manager OnDemand for i Common Server Indexing Reference*.
 39. Click **Done** to close the **Add an Index** dialog box.
 40. Close the Report window, saving the changes.
 41. Select the **Load Information** tab to specify information that Content Manager OnDemand uses to process the index data before storing it in the database.
 42. In the Application Group DB Name list, select **rdate**.
 43. From the Format list, select **%Y**.

If you need to specify a default value for **rdate**, you can specify it here in the **Default Value** field, or you can specify it using the DEFAULT keyword in the indexer parameters for this application. See the *IBM Content Manager OnDemand for i Common Server Indexing Reference* for important information about setting default values.
 44. In the Application Group DB Name list, select **acct**. To conserve space in the database, Content Manager OnDemand should remove the - (dash) character from index values before storing the values in the database. Also, it is best to define your numeric index fields (including date fields) so that leading,

trailing, and embedded blanks, dashes, currency symbols, thousands separators, etc. are removed by Content Manager OnDemand prior to being stored. Edits on these fields are sometimes strict and can cause a load to fail if non-numeric characters are found within the fields being defined.

45. In the **Embedded** field, type the - (dash) character.
46. If you require a postprocessor program to further process the index data that is extracted from the print page, enter the name of the symbolic link that points to your custom-written postprocessor in the **Postprocessor Parameters** field. For more information about writing a postprocessor program, see the *IBM Content Manager OnDemand for i Common Server Indexing Reference*.
47. Select the **Miscellaneous Options** tab to provide information that Content Manager OnDemand uses to print the report.
48. From the Default Server Printer list, select **Customer Service** printer. This is the printer that was added in “Adding server printers” on page 95.
49. At this point, the properties of the application meet the requirements. Click **OK** in the Add an Application window. The Content Manager OnDemand Administrator adds the application to the database and returns to the main window. If, however, you need to use an IBM i printer file to further define some of the server print parameters, see “Server printing and faxing” on page 255 for details.

About the folder

Before adding the folder, review the data access requirements, the types of permissions to specify, and the search and display fields to define.

Data access requirements

The folder allows users to access the telephone bill report application group and the telephone bill report application.

Permissions

Who needs access to the folder and what types of permissions do the users need?

- Users in the customer service department can open the folder to search for and retrieve statements.
- Users at Customer XYZ can open the folder to search for and retrieve statements that contain their account number and customer name.
- Define a set of folder fields for the Customer XYZ users. The folder fields will limit access to specific statements.

Search and display fields

Define two sets of folder fields:

- One set at the folder level. These folder fields allow users in the customer service department to access any statement in the database.
- One set for the CustomerXYZ group. These folder fields allow users at Customer XYZ to access specific statements.

Adding the folder

About this task

Use the **New Folder** command to add the folder.

Procedure

1. First, point to **Folders** and click the right mouse button. From the pop-up menu, select **New Folder** to open the **Add a Folder** dialog box.

The pages of the **Add a Folder** dialog box organize information about the folder into sections. The tabs show which page you are on: General, Permissions, Field Definition, Field Information, and Field Mapping.
2. Start by completing the General page. In the **Name** field, type the name of the folder.
3. In the **Description** field, type up to 120 characters of descriptive information about the folder.
4. Select the **Display Document Location** check box. This provides users with a visual clue about the type of media on which a statement is stored.
5. In the Application Groups list, select **Telephone Bill Reports**.
6. Select the **Permissions** tab to specify the types of folder functions that users can perform. For example, you can let users open the folder, create private named queries, and maintain folder fields. You can specify default permissions and permissions for specific groups and users. Unless you specify otherwise, the person that creates the folder is given all folder permissions; no other users can open or maintain the folder. On the example system, other users obtain permissions from a group. Add two groups to the folder. The groups were added in “Adding the customer service group” on page 94 and “Adding the Customer XYZ group” on page 95.
7. From the Users and Group list, select **+CustomerService**.
8. Select the **Access** check box
9. Click **Add**, to add the Customer Service group to the folder.
10. From the Users and Group list, select **+CustomerXYZ**.
11. Click **Add**, to add the CustomerXYZ group to the folder.
12. Select the **Field Definition** tab to define the folder fields. Define four folder fields to allow users to search for statements:
 - **Account Number**, a string field
 - **Customer Name**, a string field
 - **Report Date**, a date field
 - **Other Information**, a text search field
13. Complete the following steps to define a folder field:
 - a. In the **Name** field, type the name of the folder field.
 - b. In the **Description** field, type up to 120 characters of descriptive information about the folder field.
 - c. From the Field Type list, select the data type of the field.
 - d. For the other **Information** field, select a data type of text search.
 - e. Select the **Mapping Type**. All of the fields in this example use the **Single** mapping type.
 - f. Click **Add**.
14. Select the **Field Information** tab to specify the properties of the folder fields. Using the *PUBLIC identifier, you can specify field information that is used by all users that can open the folder. You can also specify field information for specific users and groups. The public field information will be used unless it is overridden by field information for a specific user or group. For the example folder, do the following:

- a. Specify public field information. For the **Account Number**, **Customer Name**, and **Other Information** fields, accept the default values. For the **Report Date** field, specify field information.
 - b. Specify field information for the CustomerXYZ group. For the **Other Information** and **Report Date** fields, accept the default values. For the **Account Number** and **Customer Name** fields, specify field information.
15. First, specify the public field information for the **Report Date** field. See online help for date format or time format values for the **Display Format** and **Defaults Format** fields.
 - a. From the Name list, select **Report Date**
 - b. From the ID list, select ***PUBLIC**
 - c. Select the **Default** check box
 - d. From the **Display Format** list, select **%Y**
 - e. From the **Defaults Format** list, select **%Y**
 - f. In the Interval area, select **Last**, type a 3 (three) in the entry field and select **Months**
 16. Next, specify the field information for the CustomerXYZ group. First, make a copy of the folder fields. Click the **Permissions** tab.
 17. From the Selected List, select **+CustomerXYZ**.
 18. In the **User/Group Fields** area, click **Yes**.
 19. Click the **Field Information** tab.
 20. Specify the Customer XYZ field information for the **Account Number** field:
 - a. From the Name list, select **Account Number**
 - b. From the ID list, select **+CustomerXYZ**
 - c. From the Default list, select **Equal**
 - d. Clear the **Like** check box
 - e. Select the **Default** check box
 - f. Select the **Fixed** check box
 - g. In the first **Defaults** entry field, type 1234567890123
 - h. Clear the **Append** check box
 21. Specify the Customer XYZ field information for the **Customer Name** field:
 - a. From the Name list, select **Customer Name**
 - b. From the Default list, select **Equal**
 - c. Clear the **Like** check box
 - d. Select the **Default** check box
 - e. Select the **Fixed** check box
 - f. In the first **Defaults** entry field, type Customer XYZ
 - g. Clear the **Append** check box
 22. Select the **Field Mapping** tab to map the folder fields to application group database fields. IBM Content Manager OnDemand uses the values that users type in folder fields to construct SQL queries. An SQL query uses the database field name.
 23. Map the following folder fields to their corresponding application group fields:
 - a. Account Number to acct
 - b. Customer Name to custname
 - c. Report Date to rdate

24. At this point, the properties of the folder meet the requirements. Click **OK** in the Add a Folder window. The administrative client adds the folder to the database and returns to the main window.

Summary

The example shows the basic requirements for adding a report to the system. Hopefully the scenario that was described and developed is similar to how you plan to use IBM Content Manager OnDemand at your company. The example should have enough variations to show the flexibility of Content Manager OnDemand to meet a range of business and operational requirements. Of course, there are several tasks that were not shown. For example, logical views of the report were not created. The system log user exit was not explored. And all of the ways to complete a given task or implement a specific requirement were not shown. As with most administrative software, there is more than one way to accomplish a task. Hopefully, the example showed you the most straightforward way to get things done with the administrative client.

You can use reference information provided with the product and the online help to find out more about how to use Content Manager OnDemand. If you have questions and cannot find the answers, contact the IBM support center. IBM also offers classes that further explore how to administer the system. Finally, you can let IBM know how well the information in this book was presented and if you found the book helpful. The section titled How to Send Your Comments explains how to let IBM know.

Adding a field to an existing application group

About the example

After working with the system for several months, the customer service department wants the following information displayed on the hit list, for each telephone bill:

- The number of pages
- The size, in bytes
- The format (for example, PDF or Line)

In addition, to help monitor data loads, the system administrator wants to add a Load Date & Time field that displays the date and time the telephone bill was loaded into the system. This information is available through the system log; however, adding it to the hit list makes it easier to find.

To implement these changes, you add the following fields to the Telephone Bill Reports application group:

Page Count

The number of pages in a telephone bill. Only new data loaded into this application displays a page count. Previously loaded data displays 0 for page count.

Document Size

The size (in bytes) of a telephone bill.

Load Date & Time

Display the time (according to the object server clock) that the telephone bill was loaded into the system.

Overview

In general, here is how you work in the administrative client to add fields to an existing application group. See the online help for details.

1. Update the Telephone Bill Reports application group to add the Page Count, Document Size, and Load Date & Time fields.
2. Update the Telephone Bill Reports application to indicate the format of the Load Date & Time field.
3. Make a new folder that displays these changes by copying the previous folder and adding new field definitions and mapping.

Note: To add an application group, the logon user ID must be one of the following types:

- A system administrator
- An application group/folder/cabinet administrator
- A user with create application groups authority

To add an application, the logon user ID must be one of the following types:

- A system administrator
- An application group/folder/cabinet administrator
- A user with create application groups authority

To add a folder, the logon user ID must be one of the following types:

- A system administrator
- An application group/folder/cabinet administrator
- A user with create folders authority

Update the application group

1. Click **Application Groups**.
2. Point to Telephone Bill Reports and click the right mouse button.
3. From the menu, select **Update**.
4. Click the **Field Definition** tab.
5. Add the following fields by typing in the name of the field in **Database Field Name**, then clicking Add:

Table 4. Values to add to fields.

Field name as it will appear in the hit list	Name to type into Database Field Name
Page Count	pagecnt
Document Size	docsize
Load Date & Time	loaddatetime

6. Click the **Field Information** tab.
7. For each field that you added, select the corresponding data type and values:

Table 5. Values to add and data types to select for fields.

Field name	Field Type	Data type	Select the following check box
pagecnt	Filter	Integer	Page Count
docsize	Filter	Integer	Document Size

Table 5. Values to add and data types to select for fields. (continued)

Field name	Field Type	Data type	Select the following check box
loaddatetime	Filter	Date/Time (TZ)	Log
<p>For loaddatetime, select Date/Time (TZ) to ensure that the application adjusts for differences in time zone.</p> <p>Always define the Field Type for new fields as Filter. Content Manager OnDemand always provides a single, default value for the new field as it updates existing application groups, as described in the following list:</p> <ul style="list-style-type: none"> • For numeric data types (for example, INTEGER), the default value is 0. • For fixed-length string data types, the default value is a string made of blanks. • For varying-length strings, the default value is a string of length 0. 			

8. Click **OK** to save the changes to the application group.

Update the application

1. Click **Applications**.
2. Point to Telephone Bill Reports and click the right mouse button.
3. From the menu, select **Update**.
4. Click the **Load Information** tab.
5. In the **Application Group Database Name** list, select loaddatetime.
6. In the **Default Value** field, type in a lowercase t. This field instructs the application to get the system date and time.
7. In the **Format** list, select %m/%d/%y %H:%M.
8. Click **OK** to save the changes to the application.

Create a folder

1. Click **Folders**.
2. Point to the folder you created when you added a report and click the right mouse button.
3. From the menu, select **Copy**.
4. Type in a new name in the **Name** field.
5. Click the **Field Definition** tab.
6. For each field, type in or select the values described in the following table and click **Add**:

Table 6. Values to add to fields.

Name	Description	Field Type	Mapping Type
Page Cnt.	Number of pages in the telephone bill.	Integer	Single
Doc Size	Number of bytes.	Integer	Single
Doc Type	Format of the telephone bill. (For example, PDF)	Document Type	
Load Date	Date and time that telephone bill was loaded into system.	Date/Time (TZ)	Single

7. Click the **Field Information** tab.

8. Select **Load Date** from the **Name** list.
9. Click the **Field Mapping** tab.
10. Select Page Cnt. from the **Name** list. Click pagecnt in the **Application Group Fields** box. Click **Add**.
11. Select Doc Size from the **Name** list. Click docsize in the **Application Group Fields** box. Click **Add**.
12. Select Load Date from the **Name** list. Click loaddatetime in the **Application Group Fields** box. Click **Add**.
13. Click **OK** to save this new folder.

Local server setup for offline administration

IBM Content Manager OnDemand supports two types of servers. The first type of server uses TCP/IP to communicate between the client programs and the server programs. The server programs can run on operating systems, such as AIX®, HP-UX, IBM i, OS/400, Sun Solaris, Windows , z/OS® and OS/390®. Depending on the server operating system, the databases supported by the server can be DB2®, Oracle (AIX, HP-UX, Sun Solaris, Windows platforms only), and Microsoft SQL server (Windows platform only). Definitions of users, groups, applications, for example, are stored in the database as well as the index values for the report data that is loaded into the system.

The second type of server is a *local server*. The local server is self-contained (no TCP/IP communication) and is defined using files contained in a directory located on a PC rather than in a server database. The files represent the system tables that define the various objects such as users, groups, and applications.

One of the uses of a local server is to enable Content Manager OnDemand administrators to work offline on administrative tasks. Another is to provide the ability to export definitions from a non-local server to a local server so the definitions can then be imported to a different non-local server. This is especially beneficial when a TCP/IP connection does not exist between two non-local servers. In either case, one of the limitations of exporting definitions from a non-local server to a local server is that the local server does not support any of the operating system-specific or database-specific parameters that are defined when using the administrative client.

When a Content Manager OnDemand administrator logs on to a server, the Content Manager OnDemand Administrator client determines whether the user has logged on to a local server or a non-local server. If the server is a non-local server, the operating system and the database types are determined. Based on the type of server, the operating system, and the database, default settings are established, entry fields are hidden or displayed, and values are added to or removed from selection lists. For definitions that have been exported from a non-local server to a local server, updating or viewing the definitions may not have the desired results because of the operating system differences between local and non-local servers. For example, when an application group and application are exported from an i server to a local server, the OS/400 Indexer that is specified in the application is not supported on a local server. When the application is viewed on the local server, the indexer field on the Indexer Information page will not have a value and the OS/400 Indexer will not be listed as a choice of indexers.

To correctly display the operating system- and database-specific parameters on a local server from parameters that have been exported from a non-local server, pull-downs for Operating System and Database appear on the Add a Server dialog

| box when you create a server with the Protocol parameter set to Local. Then, in the
| case of the application group and application that were exported from the IBM i
| server, the OS/400 Indexer will now be selected when the application is updated
| or viewed from the local server.

Restriction: Although a local server can resemble a non-local server, the following are system limitations for a local server:

- The user ID of a user that is defined on a local server cannot be updated.
- User and group permissions for users are not supported from the User Permissions page on Users dialog box.
- The Find function is not supported.
- Server Printers are not supported.
- Data Distribution Files and Groups are not supported.

For a local server, the default administrative user ID is admin. There is no initial password set for the admin user ID on a local server.

Loading spooled file data

This section provides an overview of the data loading process - the process of adding the index data to the database and loading the report data and resources into IBM Content Manager OnDemand. When a load process completes, you can view the messages that were saved in the system log. To complete the data loading process, you should backup the Content Manager OnDemand data on a regular basis.

Overview

You can archive the print output of your existing application programs in IBM Content Manager OnDemand without changing the print data stream or writing programs to process the data. In Content Manager OnDemand, the print output of an application program is called a report. Content Manager OnDemand provides programs that can index the reports, add the index data to the database, divide the input data into indexed groups of pages (documents), compress the documents, and copy the compressed documents into Content Manager OnDemand. After you archive a report in Content Manager OnDemand, your users can query, retrieve, and view or print pages of the report using the Content Manager OnDemand client program.

The Content Manager OnDemand data indexing and loading programs process input files that reside on the Content Manager OnDemand server. If you generate your reports on another system, then you would typically transfer the reports to the Content Manager OnDemand server and use the data indexing and loading programs to process them. When you index a report that contains AFP data, you must make sure that the data indexing program can access the resources required by the report. Resources include page segments and fonts. If the resources are not already on the server, you must transfer them before loading the report.

You can create up to 128 index fields for each type of report that you define to Content Manager OnDemand, providing many ways for users to query information contained in a report. The number of index fields that you define depends on the organization of the data in the report. For example, when you index a report that contains logical items, such as policies and statements, you might define index fields for the date, customer name, customer number, balance due, transaction number, and amount. When you index a report that contains transaction data, such as a general ledger, you might define index fields for the date and transaction number. After you determine what index fields you need and define them to the system, Content Manager OnDemand extracts the index values from a report during the load process and stores them in records that are added to the database.

Content Manager OnDemand compresses report data into storage objects, using information that you specify in the application. Depending on how you configure storage management for your application groups, Content Manager OnDemand can automatically copy the report to disk and archive storage.

The load process saves messages in the system log each time that you load an input file into the system. You can open the System Log folder and view the messages for information such as the name of the input file, the indexing information, and the number of rows that were added to the database.

Preparing to load reports

There are a number of things to consider when preparing to load reports.

Storage space

When you initially configure an IBM Content Manager OnDemand system, you calculate the total amount of disk, optical, and tape storage required to hold the Content Manager OnDemand database, database log files and reports. You also need to plan for the temporary space needed by Content Manager OnDemand programs. The amount of storage space that you need on your system is usually a factor of how much data that you plan to store in Content Manager OnDemand, how long you need Content Manager OnDemand to maintain the data, the compression ratio that you can expect to achieve on the report data, and the number of copies of reports that you need the system to maintain. If you plan to index your reports on the Content Manager OnDemand server, then you must allocate temporary space for the data indexing program. Temporary space is also required for the data loading program.

The *IBM Content Manager OnDemand for i Common Server: Planning and Installation Guide* provides information that can help you calculate your storage requirements.

Defining the application group

When you archive reports in IBM Content Manager OnDemand, the ADDRPTOND command adds index data to the database and compresses report data into storage objects and copies the storage objects into Content Manager OnDemand. The index data and the storage objects are associated with an application group. As part of defining a report to Content Manager OnDemand, you must specify the application group that Content Manager OnDemand uses to maintain the data. The application group specifies the database fields that hold the index data and the storage management information that determines where Content Manager OnDemand maintains the report data and how long it maintains the index data and the report data. You can use the Content Manager OnDemand administrative client to define an application group. See the online help for the administrative client for more information about defining database fields and specifying storage management information for application groups.

If you plan to maintain a copy of your reports in archive storage, then the application group must specify a migration policy and storage set that identify your storage management requirements.

Defining the application

Most customers define an application for each different report (or source of data) that they plan to archive in IBM Content Manager OnDemand. When you create an application, you specify information about the report, such as:

- The application group in which you want to store the report
- Physical information about the report, including the type of data found in the report
- The indexing parameters
- The processing that Content Manager OnDemand should do to the index data before adding it to the database

You can use the Content Manager OnDemand administrative client to define an application. See the online help for the administrative client for more information about defining applications.

Loading reports

You can load reports into IBM Content Manager OnDemand in a number of different ways, depending on your application requirements.

Running the ADDRPTOND command

The ADDRPTOND command is the primary IBM Content Manager OnDemand data indexing and loading command. The ADDRPTOND command determines if the input data needs to be indexed, and if it does, calls the indexing program. The ADDRPTOND command then processes the index data, adding it to the database, optionally compresses the report data into storage objects, and copies the storage objects to storage volumes.

You run the ADDRPTOND command each time that you want to load a report or set of reports into the system. You can either run the command from the command line or use the Content Manager OnDemand output queue monitor to periodically check for input data to process. See “Using an output queue monitor” for more detail.

Important: The Content Manager OnDemand server job must be running, otherwise the ADDRPTOND command will fail.

Using an output queue monitor

You can use the Start Monitor for OnDemand (STRMONOND) command to start a monitor program for any output queue to receive spooled files for processing. The monitor continuously checks an output queue for spooled files, and allows IBM Content Manager OnDemand to capture the spooled files as they arrive.

Important: IBM recommends using the monitors in a batch environment only.

You generally define and then manually store reports with the ADDRPTOND command during testing. Then, when the report is ready for production, you can automate storage with the Content Manager OnDemand monitor—if you use the *SPLFNAME, *FORMTYPE, *USERDATA, *JOBNAME, or *USRDFNxxxx spooled file attribute of the report to match the application group and application names you defined to Content Manager OnDemand.

The spooled files must be in ready status (RDY) on the output queue that you are monitoring. The first time a Monitor is started for an output queue, Content Manager OnDemand will create a data queue and attach it to the output queue being monitored. If there are spooled files already in that queue, the monitor will not detect them. If this happens, place these spooled files on hold, then release them. The Monitor should then detect and process them. This should only be necessary the first time a Monitor job is started for a particular output queue.

To begin capturing reports automatically when a spooled file arrives in a particular output queue, you can issue the STRMONOND command. The monitor runs continuously until the End time or Number of hours occurs. You can also use the End Monitor for OnDemand (ENDMONOND) command to stop processing.

The default job name for the monitor job is MONOUTQ. After the job has ended, the job log can be found with the User data (USRDTA) spooled file attribute set to MONOUTQ. For every spooled file that the monitor has processed, there will either be a successful load (message number 87) or a failure (message number 88) indicated in the System Log.

The monitor processes any spooled files that arrive in the named output queue in RDY (Ready) status. Spooled files with any other status are not affected by the monitor, and remain in the output queue until they are deleted or moved.

When a Common Server monitor job is active, the job status displayed by the Work with Active Jobs (WRKACTJOB) command changes a number of times during processing. The monitor job has a status of RUN when it is preparing to process a document (such as determining parameter values or running an exit program). It then spawns several jobs and goes to a status of TIMW while those jobs do the actual archiving. When the processing is complete, the status changes back to RUN while the monitor checks what needs to be done next. If there are no additional documents that are ready to be archived the monitor job goes into a DEQW status until the next document is ready to be archived or a request to end the monitor job is received.

You may wish to add the STRMONOND command to your system startup program so the monitor(s) start each time you IPL the system.

When starting a Content Manager OnDemand output queue monitor (using the STRMONOND command) from a job scheduler, you may be unsure of what job description to use. In most cases, STRMONOND will work best using the QOND400 job description. You may have special system needs that require the use of your own job description, but these two IBM-supplied job descriptions will work successfully for many customers.

See the online help for more information about these commands and their parameters.

Important: The Content Manager OnDemand server job must be running, otherwise the spooled files will fail to load.

Indexing reports

You must generate index data for a report before you load the report into IBM Content Manager OnDemand. If the report contains AFP data, then you need to store the AFP resources in Content Manager OnDemand. The resources are required to display and reprint pages of a report that contains AFP data. There are several programs provided with Content Manager OnDemand to help you generate index data for your reports:

- Content Manager OnDemand OS/400 Indexer - You can use the OS/400 Indexer to specify indexing parameters for SCS, SCS-Extended, Advanced Function Presentation (AFP), and Line spooled files.
- Content Manager OnDemand PDF Indexer - You can use the PDF Indexer to specify indexing parameters for Adobe PDF input files.
- Content Manager OnDemand Generic Indexer - You can use the Generic Indexer to specify index data for a wide variety other types of input files.

For details about indexing data, including how to use other indexers, see the *IBM Content Manager OnDemand for i Common Server: Indexing Reference*.

Processing the input data

After a report has been indexed, the ADDRPTOND command can process the index file to prepare the index data for loading into the database and prepare the report data and resource group files for the storage managers to load on storage volumes.

Processing index data

The index data is processed by the ADDRPTOND command before the index data is added to the database. The ADDRPTOND command extracts information from the application group and the application and performs the following processing:

- Determines the database field information from the application group.
- Determines the preprocessing information from the application, if needed.
- After preprocessing the index file, the ADDRPTOND command creates the database rows:
 - One row for every group of indexed pages in a report that contains a sorted transaction value
 - One row for every indexed item in a report that contains logical items, such as policies and statements
- Determines the postprocessing information from the Postprocessor Parameters specified on the Load Information tab of the application. For example, the ADDRPTOND command may need to drop duplicate index records. You specify exactly what processing you want the command to do when you define the application.
- Passes the index rows to the database manager. The rows consist of fields that contain the index values that the indexing program extracted from the report and other fields generated by IBM Content Manager OnDemand. An index row contains:
 - One column for each field defined in the application group
 - One or more columns of Content Manager OnDemand control information

Processing reports and resources

The ADDRPTOND command divides the input data into indexed groups of pages (documents) and compresses the documents into storage objects. Dividing a report into groups of pages improves the efficiency of queries and can improve the time required to retrieve and display the report. Compression improves the efficiency of the storage manager. In this step, the ADDRPTOND command:

- Determines the compression information from the application.
- Determines the storage management information from the application group. The storage management information determines the storage locations, such as disk storage and archive storage.
- Locates the AFP resources called for in the input data. If the resources are not available, the load will fail and the System Log will be updated with message number 88, indicating the failure. An ADDRPT job log will contain error messages.
- Compresses the documents into storage objects. IBM Content Manager OnDemand compresses report data into approximately 100 KB blocks (a default value) and places the blocks into a storage object. Content Manager OnDemand uses a 10 MB storage object (a default value) to improve storage efficiency and performance. Compressed data does not span storage objects. Content Manager OnDemand assigns unique file names to the storage objects and sequentially numbers them within an application group.

Loading index data

The ADDRPTOND command works with the database manager to add the index data to the database. Depending on the database organization that you specified when you defined the application group, IBM Content Manager OnDemand either creates a new table each time that you load a report or adds the index data to an existing database table.

Content Manager OnDemand uses a *segment table* as a high level index to the index data for an application group. Each row in the segment table identifies a specific table of application group index data. The fields in the segment table identify the application group and the dates found in a table. The dates represent the earliest and latest dates that can be found in that segment of application group index data. Content Manager OnDemand can use the segment table to limit a query to a specific table of application group index data.

The database manager updates the segment table if the beginning date in the report is earlier than a date already stored in the table.

Content Manager OnDemand limits the size of a table to improve performance and storage management. The number of rows in a table can be specified when you define the application group. The default size of a table is ten million rows. The database manager automatically closes a table and opens a new table when this threshold is reached. When closing a table, the database manager updates the segment table with the latest ending date found in the table.

Loading storage objects

The ADDRPTOND command calls the storage manager to copy storage objects to disk and archive storage. The storage manager extracts information from the application group to determine where and when to copy the storage objects.

IBM Content Manager OnDemand uses an object called a storage set to determine the locations that can hold report data. A storage set and its associated migration policy (of the exact same name) point to one or more storage locations. A storage set can write data to one and only one storage location at a time (the active storage level).

Disk storage

The primary purpose of disk storage is for short-term, high-speed retrieval of report data.

If you configure your application groups to copy data to disk storage, then the disk storage manager copies the storage object to disk. The **Cache Data for xx Days** setting on the Storage Management page determines whether IBM Content Manager OnDemand copies documents to disk storage.

Archive storage

A storage set can identify an archive storage media such as optical or tape. IBM Content Manager OnDemand uses its archive storage manager to maintain storage objects in archive storage for long-term storage and for backup copies of reports.

The storage manager can copy the storage object to archive storage when the report is initially loaded into the system or at a later time, depending on how you

configure your application groups. Most customers configure the system to copy report data to disk and archive storage at the same time.

Resources

IBM Content Manager OnDemand always stores resources (such as AFP overlays or page segments) on disk, to provide fast retrieval when a user selects an item for viewing. Content Manager OnDemand saves only one copy of a resource on the system, even if several reports use the same resource. When processing a resource group file, the ADDRPTOND command checks the resource identifier to determine if the resource is already present on the system.

Verifying processing

When you add a report into the system, the ADDRPTOND command saves a copy of the messages generated during the load process in the system log. After a load process completes, you can open the System Log folder and view the messages. Message number 87 indicates a successful load; message number 88 indicates a failure. A failed load will produce a job log that contains additional error messages. The user data of the job log will be ADDRPT. The System Log message reference the load program named ARSLOAD and include the date and time that the load process started and completed, the name of the input file, and the number of rows that were added to the database. For example:

```
|      arslload: Processing file
|      >/QIBM/USERDATA/ONDEMAND/QUSROND/TMP/SP_INVNEW1_DSP02_...
|      arslload: 01/26/09 12:29:40 -- Indexing started,
|      --UNKNOWN-- bytes to process
|      arslload: 01/26/09 12:29:41 Indexing completed
|      arslload: 01/26/09 12:29:41 -- Load started,
|      23306 bytes to process
|      Resource ./SP_INVNEW1_DSP02_DBRYANT_431445_000001_RDR400X_1090126_122921.res
|      will be added as resource >3-16-0<. Compression Type(0D77)
|      Original Size(10498) Compressed Size(5201)
|      OnDemand Load ID = >8495-53-0-1FAA-20130627000000-20130627000000-8496<
|      Loaded 4 rows into the database
|      arslload: 01/26/09 12:30:37 Loading completed
|      arslload: Processing successful for file >/QIBM/USERDATA/ONDEMAND/QUSRO...
```

You can verify the number of rows that IBM Content Manager OnDemand added to the database:

- For a report that contains transaction data that you have divided into indexed groups of pages, the number of pages in the report divided by the number of pages in an indexed group of pages should equal the number of rows added to the database. For example, if a report contains 150,010 pages and there are 100 pages in an indexed group of pages, then Content Manager OnDemand should have added 1,501 rows to the database.
- For a report that contains logical items, such as statements and policies, the number of rows added to the database should equal the number of indexed items in the report. For example, if a report contains 1,000 statements, Content Manager OnDemand should have added 1,000 rows to the database.

The OnDemand Load ID represents the data that the ADDRPTOND command stored in the system during a load process. The Load ID can be used to identify a specific load process. For example, you can run the RMRVPTOND command and specify the Load ID to delete the index data and documents that were created when the ADDRPTOND command processed an input file.

Beginning with version 7.2, the format of the Load Identifier changed. See online help for the RMVRPTOND or RPTRPTOND command for details about the new format. The format you use for the RPTID parameter of the RMVRPTOND and PRTRPTOND commands depends on the format of the Load Identifier that was generated when the data was loaded. If the Load Identifier that was generated when the data was loaded is in the old format, then you use the old Load Identifier format for RPTID. If the Load Identifier that was generated when the data was loaded is in the new format, then you use the new Load Identifier format. The following table shows an example of the fields in a Load ID at version 7.2:

Table 7. Example of an OnDemand Load ID

Example Load ID field value	Meaning
8495	Application group identifier
53	Primary storage node
0	Secondary storage node (always set to 0)
1FAA	The document identifier within the application group
20130627000000	The earliest date in the report (in YYYYMMDDHHMMSS format)
20130627000000	The latest date in the report (in YYYYMMDDHHMMSS format)
8496	Application identifier (optional)

Backing up databases

After you load reports into IBM Content Manager OnDemand, we recommend that you create a backup copy of the Content Manager OnDemand data.

- When you backup the Content Manager OnDemand database, you protect control information and index data that Content Manager OnDemand and the database manager need to support the system.
- When you backup the Content Manager OnDemand data in IFS, you protect your data that resides on disk that may not have been copied to archive media yet.

IBM recommends that you backup the data at least once a week, and more often if you load reports every day.

See “Backup and recovery” on page 37 for details about Content Manager OnDemand backup and recovery.

Loading image files

Overview

IBM Content Manager OnDemand provides support for storing and retrieving letters and other types of correspondence. The most straight forward way to store this type of information is to transform a printed copy of a letter into an image file using a scanner and image capture software. In addition to scanning the letter and saving it as an image file, you must create index data for the letter. You can then use the ARSLOAD program to load the index data into the database and store the letter file on the system. Your users can then query, retrieve, and view, print or send copies of the letter using one of the Content Manager OnDemand client programs.

Content Manager OnDemand provides the Generic Indexer so that you can index input files that contain data other than AFP data, line data, and PDF. You specify the index data for the input files that you want to index with the Generic Indexer in a parameter file. The parameter file contains the index field names and values and identifies the input files that you want to process. You can create up to 128 index fields for each input file that you want to load into the system, providing many ways for users to query and retrieve the letters. The number of index fields that you define usually depends on how your users retrieve documents from the system. In the example that follows, index fields were defined for the date on the letter, the name of the person that sent the letter, the company name, and the subject of the letter.

This section describes how to use a scanner and image capture software to create the image files. Before you can load the image files into the system, you must define an application group to manage the storage of the letters, an application to specify the physical and logical properties of the letters, and a folder to let users search for and retrieve the letters.

Defining the application group

When you load an input file into the system, IBM Content Manager OnDemand updates the database with the index data that you provide and stores the indexed groups of pages as documents in cache storage and archive storage. The application group contains the information that Content Manager OnDemand uses to store and maintain the index data and the documents on the system. This section provides information about some of the key properties of the application group. You can use the Content Manager OnDemand administrative client to define an application group.

Database Organization

The Database Organization determines how IBM Content Manager OnDemand organizes the index data that is stored in the application group.

Accept the default Database Organization of Multiple Loads per Database Table and Table Size of its maximum number of rows. Each time that a letter (input file) is loaded into the system, Content Manager OnDemand adds one row to a database table. When a table reaches its maximum number of rows, Content Manager OnDemand closes the table and creates a new table (unless you select

Single table for all loads). Content Manager OnDemand always adds index records to the open table; the closed tables can be queried.

Expiration Type

The Expiration Type determines how IBM Content Manager OnDemand deletes index data and documents from the application group.

Accept the default Expiration Type of Load. This means that Content Manager OnDemand deletes the index data and documents from one load process (one or more input files) at a time from the application group. Depending on the number of input files processed during a load process, one or more letters (input files) may be deleted at a time. For example, if you create a parameter file for the Generic Indexer that contains index data for several input files, then Content Manager OnDemand will delete all of the index data and documents from the application group at the same time.

Permissions

Application group permissions determine the users that can access information stored in the application group and determine the users that can do other types of tasks related to the application group.

Under the *PUBLIC identifier, specify the Access permission so that all of the users defined to the server can access data stored in the application group.

Field Definition

The Field Definition page is where you define the database fields for the application group. When you load an input file into the system, IBM Content Manager OnDemand stores the index data that is specified in the parameter file into fields in records that are added to the database. When a user searches for letters, Content Manager OnDemand compares the search criteria entered by the user with index data in the application group.

Define the following database fields. The fields allow users to locate letters based on different criteria, such as the date of the letter, the name of the person that sent the letter, and the subject of the letter.

ldate The date on the letter. Defined as a date field.

name The person that sent the letter. Defined as a string field that contains variable length data.

company
The person's company. Defined as a string field that contains variable length data.

subject
The subject of the letter. Defined as a string field that contains variable length data.

Defining the application

Most customers define an application for each different source of input data that they plan to load into IBM Content Manager OnDemand. This section provides information about some of the key properties of the application. You can use the Content Manager OnDemand administrative client to define an application.

Application Group

You must assign an application to an application group. Assign the application to the application group that was created in “Defining the application group” on page 121.

Data Format

The Data Format determines the kind of data that is found in the documents that are stored on the system. If the type of the data that is found in the original input file is different than the type of the data that is found in the documents that are stored on the system, then you should specify the type of the data that will be stored in Content Manager OnDemand. (In this example, the type of the data is the same, a TIFF image.) Selected TIFF as the Data Type. When saving the scanned image of a letter, save the data as a TIFF image.

Indexer

The indexer determines the indexing program that Content Manager OnDemand uses to index and convert input data. Select Generic as the indexer. To store TIFF images in the system, you must index them with the Generic Indexer program that is provided with Content Manager OnDemand.

Data Compression

The Data Compression determines whether IBM Content Manager OnDemand compresses the input files.

Selected Disable, so that Content Manager OnDemand does not attempt to compress the input files. In this example, the scanned image files are created as compressed TIFF images. Therefore, it is not necessary for the system to compress them. If you need to create scanned image files as uncompressed TIFF images, you should select one of the compression methods so that the system can compress the input files to improve storage efficiency and retrieval performance.

Defining the folder

You must define a folder so that users can search for and retrieve the input files that you load into Content Manager OnDemand. This section provides information about some of the key properties of the folder. You can use the Content Manager OnDemand administrative client to define a folder.

Application Group

A folder can be used to search one or more application groups. Select the application group that was created in “Defining the application group” on page 121. When users open the folder, they can search for and retrieve the input files that were loaded into the application group.

Permissions

Folder permissions determine the users that can open the folder and determine the users that can do other types of tasks related to the folder.

Under the *PUBLIC identifier, specify the Access permission so that all users defined to the library server can open the folder.

Field Definition

The Field Definition page is where you define the search fields for the folder. The search fields contain the search criteria entered by the user. For most folders, you probably want to define a search field for each field that you defined for the application group.

Define the following fields. The fields allow users to locate letters based on different criteria, such as the date of the letter, the name of the person that sent the letter, and the subject of the letter.

Letter Date

The date on the letter.

From The person that sent the letter.

Company

The person's company.

Subject

The subject of the letter.

Field Mapping

The Field Mapping page is where you *map*, or associate, the folder fields to the application group fields. This is how you specify that the search criteria that a user enters in a particular folder field should be used to search a specific application group field. Map each of the folder fields to their corresponding application group fields. For example, map the folder field named Letter Date to the application group field named Ldate.

Accessing the image files

The ARSLOAD program runs on the Content Manager OnDemand server and is the primary Content Manager OnDemand program for loading images and other user-defined data. The files that you want to load into the system by using the ARSLOAD program must be stored on the server or you must provide network access to the files.

In the example, the files were copied from a PC to the server as binary files with a file type of TIF.

Store the input files on the server in the location from which you plan to run the ARSLOAD program. Otherwise, you must specify the full path name of the input files in the Generic index file.

Creating index data

You must create index data for your input files before you can load them into IBM Content Manager OnDemand. If the input data is other than AFP, line, or PDF, then you must use the Generic Indexer to load it into the system. You specify index data for the Generic Indexer using a parameter file. You should use a standard text editor to create the parameter file. The following figure shows an example of a parameter file that can be used by the Generic Indexer to process image files. See the *IBM Content Manager OnDemand for i Common Server: Indexing Reference* for more information about the Generic Indexer and the parameter file.

```

CODEPAGE:819
COMMENT:      input file number 1
GROUP_FIELD_NAME:ldate
GROUP_FIELD_VALUE:09/01/95
GROUP_FIELD_NAME:name
GROUP_FIELD_VALUE:Mr. Earl Hawkins
GROUP_FIELD_NAME:company
GROUP_FIELD_VALUE:Soft Products
GROUP_FIELD_NAME:subject
GROUP_FIELD_VALUE:optical storage devices
GROUP_OFFSET:0
GROUP_LENGTH:0
GROUP_FILENAME:letter1.tif
COMMENT:      input file number 2
GROUP_FIELD_NAME:ldate
GROUP_FIELD_VALUE:09/01/95
GROUP_FIELD_NAME:name
GROUP_FIELD_VALUE:Hans G. Piker
GROUP_FIELD_NAME:company
GROUP_FIELD_VALUE:MBI Company
GROUP_FIELD_NAME:subject
GROUP_FIELD_VALUE:optical storage devices
GROUP_OFFSET:0
GROUP_LENGTH:0
GROUP_FILENAME:letter2.tif
COMMENT:      input file number 3
GROUP_FIELD_NAME:ldate
GROUP_FIELD_VALUE:09/16/95
GROUP_FIELD_NAME:name
GROUP_FIELD_VALUE:Laurie Unicolás
GROUP_FIELD_NAME:company
GROUP_FIELD_VALUE:Dove Properties
GROUP_FIELD_NAME:subject
GROUP_FIELD_VALUE:account balance due
GROUP_OFFSET:0
GROUP_LENGTH:0
GROUP_FILENAME:letter3.tif
COMMENT:      input file number 4
GROUP_FIELD_NAME:ldate
GROUP_FIELD_VALUE:10/01/95
GROUP_FIELD_NAME:name
GROUP_FIELD_VALUE:George VanLocal
GROUP_FIELD_NAME:company
GROUP_FIELD_VALUE:Express American
GROUP_FIELD_NAME:subject
GROUP_FIELD_VALUE:airline fares
GROUP_OFFSET:0
GROUP_LENGTH:0
GROUP_FILENAME:letter4.tif

```

Figure 4. Example of a Generic Indexer parameter file

Save the parameter file as LETTERS.IND on the server in the directory from which you plan to run the ARSLOAD program.

Configuring the ARSLOAD program

Important: The IBM Content Manager OnDemand server job must be running, otherwise the ARSLOAD program will fail.

The ARSLOAD program is the primary Content Manager OnDemand data indexing and loading program for image files. The ARSLOAD program determines if the input data needs to be indexed, and if so, it calls the indexing program. The ARSLOAD program then processes the index data, loading it into the database,

optionally compresses the input data into storage objects, and copies the storage objects to storage volumes. See “ARSLOAD” on page 294 for more information about running the ARSLOAD program within QSHHELL, including the parameters that you can specify to process input files.

You typically run the ARSLOAD program each time that you want to load a file or set of files into the system. You can either run the ARSLOAD program from the command line or configure it to run periodically to check for input data to process. To support a low volume scanning operation, most customers run the ARSLOAD program from the command line using QSHHELL. For example:

```
arsload -n -g Letters letters
```

Specify the following parameters to the ARSLOAD program:

-n Do not delete the input files.

-g Letters

The name of the application group to load.

letters

The name of the input file to process.

In the example, the ARSLOAD program must locate the input file LETTERS.IND in the current directory. The input file contains the index information that was created in “Creating index data” on page 124 (which is a parameter file for the Generic Indexer). The image files must be in the same directory.

Processing the input data

The ARSLOAD program processes the parameter file, loads the index data into the database, and loads the image files into Content Manager OnDemand.

Processing index data

The ARSLOAD program processes the parameters that were specified on the command line and the indexing parameter file before loading the index data into the database. The ARSLOAD program extracts information from the application group and the application and performs the following processing:

- Determines the database field information from the application group
- Determines the data type and other information from the application
- Creates one database row for each group that was specified in the parameter file. One group was specified for each image file.
- Passes the index rows to the database manager. The rows contain the index values and other fields generated by Content Manager OnDemand. An index row contains:
 - One column for each field that was defined for the application group
 - One or more columns of Content Manager OnDemand control information

Processing the image files

The ARSLOAD program processes each image file that was specified in the parameter file. The ARSLOAD program extracts information from the application group and the application and performs the following processing:

- Determines the compression information from the application

- Determines the storage management information from the application group. The storage management information determines the storage locations, such as disk and archive storage.
- Stores the image files in storage objects. Content Manager OnDemand uses a 10 MB storage object (the default value) to improve storage efficiency and performance. Content Manager OnDemand assigns unique names to the storage objects and sequentially numbers them within an application group.

Verifying processing

The ARSLOAD program saves a copy of the messages that were generated during a load process in the system log. (In the System Log folder, search for message number 87 for a successful load or 88 for an unsuccessful load.) After a load process completes, you can open the System Log folder and review the messages. The information in the messages includes the date and time that the load process started and completed, the name of the input file(s), and the number of rows that were added to the database. For example:

```
|
|      arslload: Processing file
|      >/images/IBMCM.ODKREL.Orders.PackingList.20090118.143129.ARD<
|      arslload: 01/26/09 12:25:42 -- Loading started, 201165 bytes to process
|      OnDemand Load Id = >8495-53-0-1FAA-20130627000000-20130627000000-8496<
|      Loaded 1 rows into the database
|      Document compression type used - Disable. Bytes Stored = >201281< Rows = >1<
|      arslload: 01/26/09 12:25:48 Loading completed
|      arslload: Processing successful for file
|      >/images/IBMCM.ODKREL.Orders.PackingList.20090118.143129.ARD<
```

You can verify the number of rows that Content Manager OnDemand added to the database. In the example, the number of rows added to the database should equal the number of groups (and image files) that was specified in the parameter file.

The OnDemand Load ID represents the data that the ARSLOAD program stored into the system during a load process. The Load ID can be used to identify a specific load process. For example, you can run the RMVRPTOND command and specify the Load ID to delete the index data and documents that were created when the ARSLOAD program processed a Generic Indexer parameter file.

Loading user-defined data

IBM Content Manager OnDemand provides support for storing and retrieving almost any type of data. For example, Content Manager OnDemand provides support for AFP, Line, PDF, GIF, JFIF (JPEG), PCX, and TIFF data. However, Content Manager OnDemand is not limited to maintaining these types of data. Content Manager OnDemand provides the User-Defined data type to support almost any other type of data that you want to store in the system. For example, you can configure the system to process Lotus® WordPro documents, so that when a user retrieves one of the files from the system, Content Manager OnDemand automatically starts Lotus WordPro to open the document.

To store user-defined data on the system, you must create index data for the input files and you must register the file type of the input file with Content Manager OnDemand. The file type determines the program that is started to open a file when a user retrieves one of the files from the system. The file type must also be registered with the client operating system. If your Content Manager OnDemand system supports client programs that run under different operating systems, then you must register the specified file type on all of the client operating systems.

Content Manager OnDemand provides the Generic Indexer so that you can index user-defined data. You specify the index data for the input files that you want to index with the Generic Indexer in a parameter file. The parameter file contains the index field names and values and identifies the input files that you want to process. You can create up to 128 index fields for each input file that you want to load into the system, providing many ways for users to query and retrieve documents. The number of index fields that you define usually depends on how your users retrieve documents from the system. For example, you might want to define index fields for the date, author, and subject or purpose of the user-defined data.

Before you can load user-defined data into the system, you must define an application group to manage the storage of the files, an application to specify the physical and logical attributes of the input files, and a folder to let users search for and retrieve the files. This section contains an example that shows how to define Lotus WordPro files to Content Manager OnDemand. It provides an overview of defining the application group, application, and folder. For details and a comprehensive example of defining input data to OnDemand, see the following sections:

Defining the application group

When you load an input file into the system, Content Manager OnDemand updates the database with the index data that you provide and stores the indexed groups of pages as documents in cache storage and archive storage. The application group contains the information that Content Manager OnDemand uses to store and maintain the index data and the documents on the system. This section provides information about some of the key properties of the application group. You can use the Content Manager OnDemand administrative client to define an application group.

Database organization

The Database Organization determines how Content Manager OnDemand organizes the index data that is stored in the application group.

Accept the default Database Organization of Multiple Loads per Database Table and Table Size of its maximum number of rows. Each time that a Lotus WordPro file is loaded into the system, Content Manager OnDemand adds one row to a database table. When a table reaches its maximum number of rows, Content Manager OnDemand closes the table and creates a new table. Content Manager OnDemand always adds index records to the open table; closed tables can be queried.

Expiration type

The Expiration Type determines how Content Manager OnDemand deletes index data and documents from the application group.

Accept the default Expiration Type of Load. This means that Content Manager OnDemand deletes the index data and documents from one load process (one or more input files) at a time from the application group. Depending on the number of input files processed during a load process, one or more Lotus WordPro files may be deleted at a time. For example, if you create a parameter file for the Generic Indexer that contains index data for several input files, Content Manager OnDemand will delete all of the index data and documents from the application group at the same time.

Permissions

Application group permissions determine the users that can access information stored in the application group and determine the users that can do other types of tasks related to the application group.

Under the *PUBLIC identifier, specify the Access permission so that all of the users defined to the server can access data stored in the application group.

Field definition

The Field Definition page is where you define the database fields for the application group. When you load an input file into the system, Content Manager OnDemand stores the index data that is specified in the parameter file into fields in records that are added to the database. When a user queries the system, Content Manager OnDemand compares the search criteria entered by the user with index data in the application group.

Define the following database fields. The fields allow users to locate files based on different criteria, such as the date, author, and subject of the file.

fdate The date associated with the file. For example, the date the file was created or the date the file was published. Defined as a date field.

author The author of the file. Defined as a string field that contains variable length data.

subject The subject or purpose of the file. Defined as a string field that contains variable length data.

Defining the application

Most customers define an application for each different source of input data that they plan to load into Content Manager OnDemand. This section provides information about some of the key properties of the application. You can use the Content Manager OnDemand administrative client to define an application.

Application Group

You must assign the application to an application group. Assign the application to the application group that was created in “Defining the application group” on page 129.

Data Format

The Data Format determines the kind of data that is found in the documents that are stored on the system. If the type of the data that is found in the original input file is different than the type of the data that is found in the documents that are stored on the system, then you should specify the type of the data that will be stored in Content Manager OnDemand.

Because there is not a supplied Data Format for Lotus WordPro files, you must select User Defined from the Data Type list.

File Extension

When you select User Defined from the Data Type list, you must also enter a value in the File Extension. The File Extension determines the program that is started to open a user-defined file when it is retrieved from the system. In the example, type the characters LWP, for Lotus WordPro.

The File Extension that you specify must also be registered on the client operating system. See your operating system information for help with registering file extensions.

Indexer

The indexer determines the indexing program that Content Manager OnDemand uses to index and convert input data. Select Generic as the indexer. To store user-defined files on the system, you must index them with the Generic Indexer program provided with Content Manager OnDemand.

Data Compression

The Data Compression determines whether Content Manager OnDemand compresses the input files.

On the Load Information page, accept the default Data Compression of 0D77 so that Content Manager OnDemand will compress input files before storing them on the server. Resources are not supported with user-defined input data.

Defining the folder

You must define a folder so that users can search for and retrieve the input files that you load into Content Manager OnDemand. This section provides information about some of the key properties of the folder. You can use the Content Manager OnDemand administrative client to define a folder.

Application Group

A folder can be used to search one or more application groups. Select the application group that was created in “Defining the application group” on page 129. When users open the folder, they can search for and retrieve the Lotus WordPro files that were loaded into the application group.

Permissions

Folder permissions determine the users that can open the folder and determine the users that can do other types of tasks related to the folder.

Under the *PUBLIC identifier, specify the Access permission so that all users defined to the server can open the folder.

Folder fields

The Field Definition page is where you define the search fields for the folder. The search fields contain the search criteria entered by the user. For most folders, you probably want to define a search field for each field that you defined for the application group.

Define the following folder fields. The fields allow users to locate files based on different criteria, such as the date, the author, and the subject.

File Date

The date associated with the file. For example, the date the file was created or the date that the file was published.

Author

The person that created the file.

Subject

The subject or purpose of the file.

Field Mapping

The Field Mapping page is where you *map*, or associate, the folder fields to the application group fields. This is how you specify that the search criteria that a user enters in a particular folder field should be used to search a specific application group field. Map each of the folder fields to their corresponding application group fields. For example, map the folder field named File Date to the application group field named fdate.

Accessing the input files

The ARSLOAD program runs on the Content Manager OnDemand server and is the primary data loading program for user-defined data. The files that you want to load into the system by using the ARSLOAD program must be stored on the server or you must provide network access to the files.

In the example, the files were copied from a PC to the server as binary files with a file type of LWP.

Store the input files on the server in the location from which you plan to run the ARSLOAD program. Otherwise, you must specify the full path name of the input files in the Generic index file.

Creating the index data

You must create index data for your input files before you can load them into the system. If the input data is other than AFP, line, or PDF, then you must use the Generic Indexer to load it into the system. You must specify index data for the Generic Indexer using a parameter file. You should use a standard text editor to create the parameter file. The following figure shows an example of a parameter file that can be used by the Generic Indexer to process the Lotus WordPro files. See the *IBM Content Manager OnDemand for i Common Server: Indexing Reference* for more information about the Generic Indexer and the parameter file.

```
CODEPAGE:819
COMMENT:      input file number 1
GROUP_FIELD_NAME:fdate
GROUP_FIELD_VALUE:12/18/95
GROUP_FIELD_NAME:author
GROUP_FIELD_VALUE:Jessica Hawkins
GROUP_FIELD_NAME:subject
GROUP_FIELD_VALUE:optical storage devices
GROUP_OFFSET:0
GROUP_LENGTH:0
GROUP_FILENAME:letter1.lwp
COMMENT:      input file number 2
GROUP_FIELD_NAME:fdate
GROUP_FIELD_VALUE:12/18/95
GROUP_FIELD_NAME:author
GROUP_FIELD_VALUE:Paul Garveys
GROUP_FIELD_NAME:subject
GROUP_FIELD_VALUE:optical storage devices
GROUP_OFFSET:0
GROUP_LENGTH:0
GROUP_FILENAME:letter2.lwp
COMMENT:      input file number 3
GROUP_FIELD_NAME:fdate
GROUP_FIELD_VALUE:12/18/95
GROUP_FIELD_NAME:author
GROUP_FIELD_VALUE:Randy Perkinen
GROUP_FIELD_NAME:subject
GROUP_FIELD_VALUE:account balance due
GROUP_OFFSET:0
GROUP_LENGTH:0
GROUP_FILENAME:letter3.lwp
COMMENT:      input file number 4
GROUP_FIELD_NAME:fdate
GROUP_FIELD_VALUE:12/18/95
GROUP_FIELD_NAME:author
GROUP_FIELD_VALUE:Georgia July
GROUP_FIELD_NAME:subject
GROUP_FIELD_VALUE:airline fairs
GROUP_OFFSET:0
GROUP_LENGTH:0
GROUP_FILENAME:letter4.lwp
```

Figure 5. Example of a Generic Indexer parameter file

Save the parameter file as LWP.IND on the server in the directory from which you plan to run the ARSLOAD program.

Configuring the ARSLOAD program

Important: The Content Manager OnDemand server job must be running, otherwise the ARSLOAD program will fail.

The ARSLOAD program is the primary Content Manager OnDemand data indexing and loading program for user-defined data. The ARSLOAD program determines if the input data needs to be indexed, and if so, it calls the indexing program. The ARSLOAD program then processes the index data, loading it into the database, optionally compresses the input data into storage objects, and copies the storage objects to storage volumes. See “ARSLOAD” on page 294 for more information about running the ARSLOAD program within QSHELL, including the parameters that you can specify to process input data.

You typically run the ARSLOAD program each time that you want to load a file or set of files into the system. You can either run the ARSLOAD program from the command line or configure it to run as a daemon (UNIX servers) or service (Windows servers) to periodically check for input data to process. To store a few WordPro files at a time, most customers run the ARSLOAD program from the command line using QSHELL. For example:

```
arsload -n -g 'Lotus WordPro Documents' lwp
```

Specify the following parameters to the ARSLOAD program:

- n** Do not delete the input files.
- g 'Lotus WordPro Documents'**
The name of the application group to load.
- lwp** The name of the input file to process.

In the example, the ARSLOAD program must locate the input file LWP.IND in the current directory. The input file contains the index information that was created in “Creating the index data” on page 133 (which is a parameter file for the Generic Indexer). Because the full path name of the input file was not specified, they must reside in the same directory.

Processing the input data

The ARSLOAD program processes the parameter file, loads the index data into the database, and loads the Lotus WordPro files into Content Manager OnDemand.

Processing index data

The ARSLOAD program processes the parameters that were specified on the command line and the indexing parameter file before loading the index data into the database. The ARSLOAD program extracts information from the application group and the application and performs the following processing:

- Determines the database field information from the application group
- Determines the data type and other information from the application
- Creates one database row for each group that was specified in the parameter file. One group was specified for each Lotus WordPro file.
- Passes the index rows to the database manager. The rows contain the index values and other fields generated by Content Manager OnDemand. An index row contains:
 - One column for each field that was defined for the application group
 - One or more columns of Content Manager OnDemand control information

Processing the Lotus WordPro files

The ARSLOAD program processes each Lotus WordPro file that was specified in the parameter file. The ARSLOAD program extracts information from the application group and the application and performs the following processing:

- Determines the compression information from the application
- Determines the storage management information from the application group. The storage management information determines the storage locations, such as disk and archive storage.
- Stores the Lotus WordPro files in storage objects. Content Manager OnDemand uses a 10 MB storage object (the default value) to improve storage efficiency and performance. Content Manager OnDemand assigns unique names to the storage objects and sequentially numbers them within an application group.

Verifying processing

The ARSLOAD program saves a copy of the messages that were generated during a load process in the system log. After a load process completes, you can open the System Log folder and review the messages. Search for message number 87 for a successful load or 88 for an unsuccessful load. The information in the messages includes the date and time that the load process started and completed, the name of the input file, and the number of rows that were added to the database. For example:

```
|      arsload: Processing file
|      >/wordproc/letter.lwp.ARD<
|      arsload: 01/26/09 12:25:42 -- Loading started, 201165 bytes to process
|      OnDemand Load Id = >8495-53-0-1FAA-20130627000000-20130627000000-8496<
|      Loaded 1 rows into the database
|      Document compression type used - Disable. Bytes Stored = >201281< Rows = >1<
|      arsload: 01/26/09 12:25:48 Loading completed
|      arsload: Processing successful for file
|      >/wordproc/letter.lwp.ARD<
```

You can verify the number of rows that Content Manager OnDemand added to the database. In the example, the number of rows added to the database should equal the number of groups (and Lotus WordPro files) that were specified in the parameter file.

The OnDemand Load ID represents the data that the ARSLOAD program stored into the system during a load process. The Load ID can be used to identify a specific load process. For example, you can run the RMRPTOND command and specify the Load ID to delete the index data and documents that were created when the ARSLOAD program processed a Generic Indexer parameter file.

Restarting a load process

The ADDRPTOND command will terminate if an unrecoverable error occurs during index, database, or storage manager processing. Termination processing includes setting a return code and saving error messages in the system log.

To start problem determination, open the System Log folder and view the messages that the command generated during the load process. The message log will contain normal processing messages, return codes, and error messages, such as message number 87 for a successful load or 88 for an unsuccessful load. See “System log messages” on page 249 or *IBM Content Manager OnDemand: Messages and Codes*, SC27-1379 for other possible messages.

If the command failed during indexing, correct the problem and then restart the load process from the beginning. Common causes of problems during indexing include invalid input files or indexing parameter files and insufficient temporary space.

Tip: The messages in the system log will actually refer to a program named ARSLOAD.

If the command failed during database processing or storage manager processing:

- Determine and correct the problem.
- If a Load ID is listed in the message log that the ADDRPTOND command saved in the system log, then you can use the RMVRPTOND command to unload the data. See “Deleting a report” on page 139 for information about unloading data from Content Manager OnDemand.
- Restart the load process from the beginning.

Deleting a report

You can use the RMVRPTOND command to delete the index data and documents that the ADDRPTOND command stored in the system during a load process. To use the delete command, you must specify the application group name and Load ID that the ADDRPTOND command generated during the load process. The Load ID represents the index data that was added to the database and the storage objects that were copied to disk and archive storage. The complete Load ID can be found in the message that the ADDRPTOND command saved in the system log. A partial Load ID (which is all that RMVRPTOND actually requires) can be found by viewing the Properties of a document from the document list of the OnDemand client. From the Search Criteria and Document List panel, right-click a document and select Properties. The Document Properties window displays, which contains the Partial Load ID information.

See the online help for RMVRPTOND for more information about the command and parameters.

When the RMVRPTOND command completes, you should open the system log folder to view the messages that were generated during the delete process. The messages will reference a program named ARSADMIN.

Important: The Content Manager OnDemand server job must be running. Otherwise, the RMVRPTOND command fails.

Managing the server

This chapter provides information for a number of topics related to the IBM Content Manager OnDemand server.

Errors and alerts

During normal processing, IBM Content Manager OnDemand programs, including the client programs, generate messages. Content Manager OnDemand saves the messages in the System Log and sends a copy of each message to the System Log user exit program. Content Manager OnDemand assigns a severity to each message. Messages that are assigned a severity of alert or error are also sent to the IBM i QSYSOPR message queue and logged in the server job log. Additionally, when a user runs a query that requires a table of index data that has been migrated to archive storage, Content Manager OnDemand sends a message to QSYSOPR and the job log.

Content Manager OnDemand provides the System Logging facility to help you identify and resolve any alerts and errors that you may receive. You can open the System Log folder to display the messages that are saved in the System Log. See “System logging facility” for more information; see “System log messages” on page 249 or *IBM Content Manager OnDemand: Messages and Codes*, SC27-1379 for a listing of message numbers and text.

System logging facility

IBM Content Manager OnDemand provides a logging facility to help administrators track Content Manager OnDemand activity and monitor the system. When you enable logging for system events, user events, and application group events, Content Manager OnDemand stores the messages that are generated by the various Content Manager OnDemand programs in the System Log. You can use the Content Manager OnDemand client program to search for and filter messages by time stamp, severity, message number, and user name.

Searching for and viewing messages

To search for and view the messages that are stored in the System Logging facility, log on to IBM Content Manager OnDemand with an OnDemand client program and open the System Log folder. Enter search criteria in one or more of the search fields. When you choose the **Search** command, Content Manager OnDemand retrieves the messages from the database that match the search criteria that you specified.

You can specify a date and time value to search for and retrieve messages. You can also specify other search criteria, such as:

Userid

The OnDemand user ID

Account

Accounting information; the information specified in the Account field for the user.

Log Id

Each time that a client logs on to the server, Content Manager OnDemand assigns a number to that session. All messages that are generated during that session include the same Log Id.

Severity

Content Manager OnDemand assigns a severity to each message: Alert, Error, Warning, Info, and Debug

View Depending on the type of message in the System Log, you may be able to view other information that is related to or associated with the message. For example:

- You can display the message log that was generated during a load process by selecting an ARSLOAD message and then choosing the **View all Selected** command. You might see ARSLOAD messages after issuing the ADDRPTOND command.

While most processes do not generate other information that can be stored in the System Log, you could write a user exit program to process the messages and generate your own information about the events. For example, you could write a user exit program to generate a report that lists the number of users that are logged on to the system in 30 minute increments. Content Manager OnDemand provides a System Log user exit so that you can process any message that is stored in the System Log and take the action that you require. See below for more information about the System Log user exit.

- Other messages in the System Log do not provide additional records or other data associated with an event. For example, the Logon and Logoff events each generate a single message, with no additional information that you can view.

Msg Num

The message number that is assigned by Content Manager OnDemand

Message

The text of the message that Content Manager OnDemand uses to restrict a search. For example, if you type Login, Content Manager OnDemand searches for and displays the messages issued by the **Logon to a Server** command.

System Log user exit

When you enable logging for system, user, and application group events, IBM Content Manager OnDemand sends a copy of each message that is generated by the system to the System Log user exit program. The System Log user exit program is named `arslog` and resides in the Content Manager OnDemand library named QRDARS.

The System Log user exit program that is provided by IBM does not perform any functions. However, you can replace the program that is provided by IBM with your own program that does user-defined processing. For example, you could create a program that checks for certain message numbers or severity and takes whatever action you deem appropriate.

You configure Content Manager OnDemand to send messages to the System Log user exit by selecting User Exit Logging options with the System Parameters command. See the online help for the administrative client for more information about the User Exit Logging options and the System Parameters command.

See the *IBM Content Manager OnDemand for i Common Server Planning and Installation Guide* for more information about the System Log user exit.

Monitoring users

The IBM Content Manager OnDemand server generates System Log messages to help you track the number of users that are logged on to the server:

- Content Manager OnDemand stores message number 201 in the System Log approximately every 30 minutes if at least one activity is detected. This message contains the current number of users that are logged on to the server.
- Content Manager OnDemand stores message number 202 in the System Log every time that the number of concurrent users exceeds the previous maximum number of concurrent users. The number of concurrent users is reset each time that you restart the Content Manager OnDemand server processes.

Generating usage statistics

IBM Content Manager OnDemand administrators can run queries against the System Log database file to gather usage statistics about Content Manager OnDemand, provided they have checked the appropriate logging options in the Content Manager OnDemand Application Group definitions for the information they want to capture.

The initial System Log file name is SL2. This System Log file is found in the library whose name matches the instance name. For example, the initial System Log file for the QUSROND instance is named SL2 in library QUSROND on the IBM i server. When the initial System Log file becomes full, a new file is created, and its name is incremented by one; the name thus becomes SL3.

For example, on the **Message Logging** tab of an Application Group definition, the Content Manager OnDemand administrator can select the **Retrieval** checkbox, which provides messages in the System Log for every retrieval from the particular Application Group.

As another example, the Content Manager OnDemand administrator can set up Application Group fields such that when the end user searches on that particular field, the information is provided in the message placed in the System Log. To do this, select the **Log** checkbox on the **Field Information** tab of the Application Group for the field for which you want to capture information.

Two of the most common message numbers that can be queried from the System Log file are:

- Number 65 (the actual search)
- Number 66 (the retrieval with field information and OnDemand Application name)

ARSSUPPORT utility

You can use ARSSUPPORT, a Java based tool that runs on your IBM i server to gather diagnostic information such as log entries. This tool is especially helpful when you need to report problems to IBM support.


The ARSSUPPORT utility is delivered in the `arssupport.jar` file. To invoke the utility, use this QSHELL command (entered all on one line as one command from the QSHELL command prompt) on the IBM i system that is running Content Manager OnDemand:

```
java -Djava.library.path=/qsys.lib/qrdars.lib -cp /qibm/proddata/ondemand/support/arssupport.jar
com/ibm/cm/od/arssupport -I instance -l -m 28 -u admin
```

Prerequisites

- Ensure that you are logged on to the operating system using an ID that is also defined to Content Manager OnDemand as a system administrator.
- Run ARSSUPPORT from the QSHELL command prompt.
- To retrieve system log entries, ensure that the Content Manager OnDemand server is running.
- The data is collected from the system where ARSSUPPORT is run.

Syntax



Parameters

- h** Use this parameter to display help and usage information about this tool.
- I instance name**
Specify the instance name to collect the instance information. If you do not specify this option, QUSROND is used as the instance name.
- l** Specify this parameter to retrieve system log entries. If you do not specify this option, the log entries for the past 60 minutes are retrieved. Ensure that you use this parameter with the -u parameter.
- m minutes**
Specify how many past minutes of the system log entries to retrieve from the server. The maximum is 6000 minutes.
- o outputpath**
Specify the output directory name. If the output directory is not specified, the output directory is the current directory.
- p password**
Password that you use to access the server.
- u userid**
The Content Manager OnDemand user ID that you use to access the server. If the -l parameter is specified, the -u parameter is required.
- v** Verbose output while running.

ARSSUPPORT generates information about a Content Manager OnDemand server including information about its configuration and system environment.

ARSSUPPORT archives all files into one compressed file, arssupport.zip, and places this file in the odsupport subdirectory of the output directory.

Examples (shown as three separate QSHELL commands)

```
java -Djava.library.path=/qsys.lib/qrdars.lib
-cp /qibm/proddata/ondemand/support/arssupport.jar
com/ibm/cm/od/arssupport -h
```

```
java -Djava.library.path=/qsys.lib/qrdars.lib
-cp /qibm/proddata/ondemand/support/arssupport.jar
com/ibm/cm/od/arssupport -o /directory_name -l -u admin -p password
java -Djava.library.path=/qsys.lib/qrdars.lib
-cp /qibm/proddata/ondemand/support/arssupport.jar
com/ibm/cm/od/arssupport -I instance -l -m 28 -u admin -p password
```

Finding or changing the server job and its attributes for a particular instance

The name of the server job for a particular server matches the name of the IBM Content Manager OnDemand instance for which it is running. For example, the server job name is QUSROND for the default Content Manager OnDemand instance named QUSROND. You can use the command WRKACTJOB JOB(QUSROND) to find this default server job.

In addition, Common Server spawns many QSQSRVR jobs during its processes. To determine which specific QSQSRVR job to view if you encountered a problem with a Content Manager OnDemand process, look at the instance job log (for example, the job log for the QUSROND job) and compare the time stamps of the error messages to the time stamps of the QSQSRVR job logs to determine which spawned QSQSRVR job or jobs are related to the particular process you are investigating.

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

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

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

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

Controlling the run priority of instance server jobs

The run priority of Common Server instance server jobs is by default set by the *ANY routing entry in the QSYSWRK subsystem description. Normally this defaults to priority 50. (It depends on the class that is specified for the *ANY routing entry in subsystem QSYSWRK.) If you wish to change this for all instance server jobs, you may add a routing entry to the QSYSWRK subsystem description using the Add Routing Entry (ADDRTGE) command as follows:

```
ADDRTGE SBSD(QSYS/QSYSWRK) SEQNBR(nnnn) CMPVAL('QRLMSERVER') +
      PGM(QSYS/QCMD) CLS(class)
```

Where:

nnnn is an unused routing sequence number in the QSYSWRK subsystem, which is lower than the *ANY routing entry sequence number.

class is the name of an IBM i class object (OBJTYPE(*CLS)) that contains the attributes you want to use.

For example:

```
ADDRTGE SBSD(QSYS/QSYSWRK) SEQNBR(1000) CMPVAL('QRLMSERVER') +
      PGM(QSYS/QCMD) CLS(QSYS/QSYSCLS35)
```

The system ships with the following classes that you might consider using:

QSYS/QSYSCLS25 (Run priority of 25)

QSYS/QSYSCLS35 (Run priority of 35)

If you want to specify a different run priority, for example 45, then you need to create your own class. An example command that does this is:

```
CRTCLS CLS(QGPL/ONDSVR45) RUNPTY(45) TIMESLICE(2000) +
      PURGE(*YES) DFTWAIT(30) CPUTIME(*NOMAX) +
      MAXTMPSTG(*NOMAX) +
      TEXT('OnDemand Common Server run priority 45 class')
```

After creating the class, specify it as the class name in the routing entry for QRLMSERVER in subsystem QSYSWRK.

Using Content Manager OnDemand data areas

Data areas can be used to customize the way certain processes run. The following table lists the data areas that can be created and used with Content Manager OnDemand

Table 8. Content Manager OnDemand data areas

Data area name	Library	Scope	Description	Details
QDFTINST	QUSRRDARS or any other library in library list when command is run	Global (if in QUSRRDARS or any other library in library list when command is run)	<p>Default instance name used when INSTANCE(*DFT) is specified on an OnDemand command. The QDFTINST data area can exist in multiple libraries. A search for the QDFTINST data area is performed using the library list. The first data area found will be used. If the QDFTINST data area is not found using the library list, the QUSRRDARS library will be searched for the data area. If the data area exists in the QUSRRDARS library it will be used. If the QDFTINST data area is not found, instance QUSROND is used.</p> <p>This data area does not apply to the STRTCPSVR and ENDTCPSVR commands. See the STRTCPSVR data area instead.</p>	*CHAR with a length of 10. Contains a specific instance name, left justified.
STRTCPSVR	QUSRRDARS	Global	<p>Default instance name for use with the STRTCPSVR SERVER(*ONDMD) INSTANCE(*DFT) and ENDTCPSVR SERVER(*ONDMD) INSTANCE(*DFT) commands.</p>	*CHAR with a length of 10. Contains a specific instance name, left justified.

Table 8. Content Manager OnDemand data areas (continued)

Data area name	Library	Scope	Description	Details
QRLWTRACE	QUSRRDARS	Global	Controls verbose message logging in the Content Manager OnDemand component of IBM Navigator for i.	*CHAR with a length of 1. Possible values are '0' (which turns verbose logging off) and '1' (which turns verbose logging on). Default without data area is '0' (off).
QRLMFKPORT	QUSRRDARS	Global	Port number for FNDKEYOND command. Used when two or more Content Manager OnDemand servers exist in the network.	*CHAR with a length of 10. Contains port number for Content Manager OnDemand instance in positions 1 through 5. Positions 6 through 10 are no longer used and can be left blank.
QRLMMONQ	QUSRRDARS or instance library	Global (if in QUSRRDARS) or instance (if in instance library)	When this data area exists, the monitor will execute ILE monitor exit programs. If the data area does not exist, only OPM monitor exit programs will be executed. (Checks both QUSRRDARS and instance library.)	*CHAR with a length of 1. Can be left blank.
QRLCTAPE	QRDARS	Global	Controls end-of-tape option (ENDOPT) during Content Manager OnDemand tape processing, such as ASM.	*CHAR with a length of 7. Possible values are *LEAVE, and *UNLOAD, left justified. Default without data area is to unload the tape. Note that *REWIND is not supported, since it would not be logical to rewind the tape and leave it loaded.

Restarting journaling

If you receive a message on the server indicating that journaling needs to be started for your IBM Content Manager OnDemand database files, issue the following command from a command line while signed on with sufficient authority: **CALL QRLCSTRJ PARM(RLC)**

Important: The parameter value (RLC) must be entered in uppercase.

Importing and exporting objects through batch administration

Overview

You can use an XML interface to import and export administrative objects into and out of an IBM Content Manager OnDemand system. The administrative objects that can be imported and exported include:

- users
- groups
- applications
- application groups
- storage sets
- folders
- printers
- cabinets
- holds

This XML interface expands the functionality and enables you to export all administrative objects into a single XML file, and later import them into the same Content Manager OnDemand system or another system.

Also, you can create an XML file from scratch through a user application or Web interface according to the defined specifications, and import it into the system.

If you create an XML file from scratch, ensure that you include objects following this order:

1. users
2. groups
3. printers
4. storage sets
5. application groups
6. applications
7. folders
8. cabinets
9. holds

Installing batch administration

You can install and setup the prerequisites for Content Manager OnDemand batch administration. Information presented here explains how to do this, and describes a short installation verification process.

Prerequisites

Prior versions of the IBM Content Manager OnDemand batch administration software used the Java version of the Xerces XML parser. Content Manager OnDemand version 7.2 uses the C++ version of the Xerces XML parser which is included with the Content Manager OnDemand licensed program product. You do

not need to acquire or install any additional parser software to use the ARSXML API to perform batch administration functions.

The following files are included with the Content Manager OnDemand batch system administration:

- /QIBM/ProdData/OnDemand/bin/arsxml
- /QIBM/ProdData/OnDemand/bin/xml/ODAdmin.jar
- /QIBM/ProdData/OnDemand/bin/xml/ondemand.xsd - (OnDemand XML Schema file)
- /QIBM/ProdData/OnDemand/bin/xml/samples/addgroups.xml
- /QIBM/ProdData/OnDemand/bin/xml/samples/addusers.xml
- /QIBM/ProdData/OnDemand/bin/xml/samples/deletegroups.xml
- /QIBM/ProdData/OnDemand/bin/xml/samples/deleteusers.xml
- /QIBM/ProdData/OnDemand/bin/xml/samples/exportgroups.xml
- /QIBM/ProdData/OnDemand/bin/xml/samples/exportusers.xml
- /QIBM/ProdData/OnDemand/bin/xml/samples/updategroups.xml
- /QIBM/ProdData/OnDemand/bin/xml/samples/updateusers.xml

Installation verification

About this task

Perform the following steps:

Procedure

1. Log onto your IBM i system with a user profile that is defined to Content Manager OnDemand as an Administrator, and that has sufficient authority to create user profiles.
2. Create five user profiles, to be used for installation verification, called SAMPLEUSR0, SAMPLEUSR1, SAMPLEUSR2, SAMPLEUSR3, and SAMPLEUSR4 on your i system. (These user profiles will be deleted when you have completed the installation verification.)

What to do next

You should not add these five user profiles to Content Manager OnDemand at this point. You will do this later, using the batch administration function.

To run the ARSXML program, first start QSHELL. You do this by using the Start QSH (STRQSH) command.

Run this command:

```
arsxml add -h QUSROND -i /qibm/proddata/ondemand/bin/xml/samples/addusers.xml -v
```

where *QUSROND* is the name of the Content Manager OnDemand instance into which you want to add new users.

Important: Omitting the **-u** and **-p** parameters causes Content Manager OnDemand to use the user profile with which you are currently logged on.

The XML file contains user passwords that are six to eight characters. If the Content Manager OnDemand system where these users are created has different password restrictions, these passwords might need to be changed before you run the command.

After you run the command, you should receive several messages stating that a printer and five users were added successfully.

If the command is properly run, you should be able to use the Content Manager OnDemand Administrator client to view the newly added users and printer.

To remove the newly added objects, run this command:

```
arsxml delete -h <hostname> -u <user> -p <password> -i deleteusers.xml -v
```

where *QUSROND* is the instance to which you previously added users. Then, delete the five user profiles called SAMPLEUSR0, SAMPLEUSR1, SAMPLEUSR2, SAMPLEUSR3, and SAMPLEUSR4 from your IBM i system.

You can use several other sample XML files. However, these files depend on the users that are created by the `addusers.xml` file.

Common problem during installation verification

This section discusses a common error that you might encounter during installation verification.

A parsing error occurred in file xxxx/samples/addusers.xml, Line 3, Column 62: cvc-elt.1: Cannot find the declaration of element onDemand.

This error usually indicates that the IBM Content Manager OnDemand schema file, `ondemand.xsd`, was not found. See “Installation verification” on page 152 and ensure that the ARSXML command runs from the correct directory.

The sample files require that the Content Manager OnDemand schema file be located in the directory “above” where the ARSXML command is run. If changes are made to the sample files, it might be necessary to update the location of the schema file within the sample files to the full path name of the `ondemand.xsd` file. Currently, the sample files include this line: `xsi:noNamespaceSchemaLocation=“../ondemand.xsd”>`. The full path name would be specified as this:
`xsi:noNamespaceSchemaLocation=“/QIBM/ProdData/OnDemand/bin/xml/ondemand.xsd”>`.

Importing an XML file into an OnDemand system

About this task

Importing an XML file that contains administrative objects into a IBM Content Manager OnDemand system is a two-step process:

Procedure

1. Preparing an XML file for the import process
2. Importing the XML file by using the ARSXML API

Preparing an XML file for the import process

Before importing data into an IBM Content Manager OnDemand system, you need to have an XML file that contains all the data that needs to be imported. The XML file can be either created during a previous XML export process or created manually. If you decide to develop an XML file manually, you must follow the format of the OnDemand XML schema file.

The OnDemand XML schema file defines the syntactic format for all OnDemand XML files, and is used during the import process to validate the contents of the import XML file. For different objects, the schema file specifies which fields are required and which fields are optional. Also, the schema file can establish a list of valid values for certain fields.

The following sample is a portion of the OnDemand schema file for the object *user*:

```
<xs:element name="user">
  <xs:complexType>
    <xs:choice maxOccurs="unbounded">
      <xs:element name="userPermission" minOccurs="0" maxOccurs="unbounded">
        <xs:complexType>
          <xs:attribute name="task" type="taskString" use="optional"/>
          <xs:attribute name="user" type="nameString" use="optional"/>
          <xs:attribute name="group" type="nameString" use="optional"/>
          <xs:attribute name="userAuthority" type="authString" use="optional"
            default="Access"/>
        </xs:complexType>
      </xs:element>
    </xs:choice>

    <xs:attribute name="name" type="nameString" use="required"/>
    <xs:attribute name="uid" type="xs:integer" use="optional"/>
    <xs:attribute name="description" type="descString" use="optional"/>
    <xs:attribute name="password" type="passwordString" use="optional"/>

    <xs:attribute name="fullName" type="userMiscString" use="optional"/>
    <xs:attribute name="acctInfo" type="userMiscString" use="optional"/>
    <xs:attribute name="company" type="userMiscString" use="optional"/>
    <xs:attribute name="title" type="userMiscString" use="optional"/>
    <xs:attribute name="addr1" type="userMiscString" use="optional"/>
    <xs:attribute name="addr2" type="userMiscString" use="optional"/>
    <xs:attribute name="addr3" type="userMiscString" use="optional"/>
    <xs:attribute name="addr4" type="userMiscString" use="optional"/>
    <xs:attribute name="dept" type="userMiscString" use="optional"/>
    <xs:attribute name="building" type="userMiscString" use="optional"/>
    <xs:attribute name="room" type="userMiscString" use="optional"/>
    <xs:attribute name="phone" type="phoneString" use="optional"/>
    <xs:attribute name="fax" type="phoneString" use="optional"/>
    <xs:attribute name="coverPage" type="userMiscString" use="optional"/>
    <xs:attribute name="printer" type="xs:string" use="optional"/>
    <xs:attribute name="timeOut" type="xs:string" use="optional"/>
    <xs:attribute name="email" type="emailString" use="optional"/>

    <xs:attribute name="userType" type="userTypeString"
      use="optional" default="User"/>
    <xs:attribute name="createFoldersAuth" type="yesnoString"
      use="optional" default="No"/>
    <xs:attribute name="createUsersAuth" type="yesnoString"
      use="optional" default="No"/>
    <xs:attribute name="createGroupsAuth" type="yesnoString"
      use="optional" default="No"/>
    <xs:attribute name="createAppGroupsAuth" type="yesnoString"
      use="optional" default="No"/>
  </xs:complexType>
</xs:element>
```

The following sample is a portion of the OnDemand schema file for the object *group*:

```
<xs:element name="group">
  <xs:complexType>
    <xs:choice maxOccurs="unbounded">
      <xs:element name="user" maxOccurs="unbounded">
        <xs:complexType>
          <xs:attribute name="task" type="taskString" use="optional"/>
          <xs:attribute name="name" type="nameString" use="required"/>
        </xs:complexType>
      </xs:element>
    </xs:choice>

    <xs:attribute name="name" type="nameString" use="required"/>
    <xs:attribute name="gid" type="xs:integer" use="optional"/>
    <xs:attribute name="description" type="descString" use="optional"/>
    <xs:attribute name="ownerUser" type="nameString" use="optional"/>
    <xs:attribute name="ownerGroup" type="nameString" use="optional"/>
  </xs:complexType>
</xs:element>
```

The following example is a portion of the OnDemand XML file, which contains a Content Manager OnDemand administrative user named *bill* and a user group named *SpecialGroup*. *SpecialGroup* is owned by the user *admin*, and contains a single user *bill*.

```
<user name="bill"
  password=""
  description="This is an Admin user."
  email="Sample@us.ibm.com"
  userType="User Admin">
</user>

<group name="SpecialGroup"
  description="A Special Group for special users"
  ownerUser="admin">
  <user name="bill"/>
</group>
```

Creating an XML file

Occasionally, you need to manually create an OnDemand XML file. To do this you need to understand the syntax of XML and the structure of the XML objects. These objects are used by IBM Content Manager OnDemand.

Overall file structure

An OnDemand XML file is an unformatted text file that can be created by the user.

Every OnDemand XML file contains the following elements:

xml identifier tag

The XML identifier tag specifies the version of XML that is used and the encoding that is used. The standard XML identifier is:

```
<?xml version="1.0" encoding="UTF-8"?>
```

onDemand identifier tag

The onDemand identifier tag specifies what the XML file is used for and what schema file to use. The standard Content Manager OnDemand identifier is:

```
<onDemand xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xsi:noNamespaceSchemaLocation="../ondemand.xsd">
```

The name of the `ondemand.xsd` file, shown here as `ondemand.xsd`, should be changed if you have changed the name of your Content Manager OnDemand XML schema file. See “Specifying the `ondemand.xsd` schema file” for important details on the location and specification of the `ondemand.xsd` file.

OnDemand objects

See “XML objects in the Content Manager OnDemand XML file” on page 165 for objects and data model used in the XML file.

onDemand ending tag

The `onDemand` ending tag indicates the end of the Content Manager OnDemand XML file. The ending tag is:

```
</onDemand>
```

Every standard Content Manager OnDemand XML file looks like this:

```
<?xml version="1.0" encoding="UTF-8"?>
<onDemand xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xsi:noNamespaceSchemaLocation="ondemand.xsd">
  ... OnDemand Objects ...
</onDemand>
```

Specifying the `ondemand.xsd` schema file

Input XML files used with ARSXML require the location of the `ondemand.xsd` schema file to be specified on the OnDemand identifier tag, which is typically the second line of the input file. By default, the `ondemand.xsd` schema file is located in IFS on your IBM i server in the `/QIBM/ProdData/OnDemand/bin/xml` directory. However, there are a number of choices for the location of the schema file and the corresponding method to identify its location in your input XML file. You can choose to copy the IBM supplied `ondemand.xsd` file to one or more other locations within IFS. Or, you can create symbolic links to the IBM supplied file, which has the added advantage in that there is no need to synchronize any copies you make of the `ondemand.xsd` file with the latest version supplied by IBM.

Within the input XML file on the OnDemand identifier tag, the `xsi:noNamespaceSchemaLocation` keyword specifies the location of the `ondemand.xsd` schema file. The path name specified with this keyword might need to be modified in your XML, depending on where you locate your `ondemand.xsd` schema file. The examples below show multiple ways to specify this keyword and schema file location.

Absolute path name

One alternative is to specify a fully qualified path name to the `ondemand.xsd` file in your input XML file. For example, specifying the absolute path name to the IBM supplied `ondemand.xsd` is shown in this example:

```
<onDemand xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xsi:noNamespaceSchemaLocation="/QIBM/ProdData/OnDemand/bin/xml/ondemand.xsd"
```

Relative path name

The path name can also be specified as a location relative to the location of your input XML file. In this example, the input file that contains the XML you want to run is called `users.xml`. If the locations of `users.xml` and `ondemand.xsd` files are as shown:

```
/arstest/arsxml/xml/users.xml
/arstest/arsxml/ondemand.xsd
```

Then the corresponding entry for the **xsi:noNamespaceSchemaLocation** keyword in the users.xml file would be:

```
<onDemand xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
xsi:noNamespaceSchemaLocation="../ondemand.xsd"
```

Similarly, if the locations of users.xml and ondemand.xsd are:

```
/arstest/arsxml/xml/users.xml
/arstest/arsxml/schema/ondemand.xsd
```

Then the corresponding entry in users.xml would be:

```
<onDemand xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
xsi:noNamespaceSchemaLocation="../schema/ondemand.xsd"
```

Same directory

If both the users.xml and ondemand.xsd files are located in the same directory, then the corresponding entry in the users.xml file does not require any path name to be specified. The entry would be:

```
<onDemand xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
xsi:noNamespaceSchemaLocation="ondemand.xsd"
```

Symbolic link

You might prefer to create a symbolic link in the directory where your input XML file is located which points to the ondemand.xsd file provided by IBM. This method has the advantage of pointing directly to the IBM supplied file, so that any updates provided by IBM are immediately in effect. There is no need to remember to copy a newly updated ondemand.xsd file to any other directories.

The following command creates a symbolic link in directory /arstest/arsxml/xml that links to the IBM supplied ondemand.xsd in /QIBM/ProdData/OnDemand/bin/xml:

```
ADDLNK OBJ('/QIBM/ProdData/OnDemand/bin/xml/ondemand.xsd')
NEWLNK('/arstest/arsxml/xml/ondemand.xsd') LNKTYPE(*SYMBOLIC)
```

The corresponding entry in users.xml would be:

```
<onDemand xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
xsi:noNamespaceSchemaLocation="ondemand.xsd"
```

Special considerations

- If you run the ARSXML export function, the resulting XML from the export always specifies the **xsi:noNamespaceSchemaLocation** as shown:

```
<onDemand xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
xsi:noNamespaceSchemaLocation="ondemand.xsd"
```

You might need to modify the **xsi:noNamespaceSchemaLocation** entry before you attempt to run the exported XML, depending on which approach you have chosen to locate your ondemand.xsd schema file.

- See “Tip 3: Parsing error while running ARSXML” on page 314 for information about the error you might receive when running ARSXML if the ondemand.xsd schema file cannot be found.

Objects

The basic building blocks for the XML file are referred to as objects.

The following objects can be included in the XML file:

- user

- group
- printer
- storageSet
- applicationGroup
- application
- folder
- cabinet
- hold

These objects as well as all of the child objects are shown in detail in the XML data tables in “XML objects in the Content Manager OnDemand XML file” on page 165.

Ensure that you include the objects in this order:

1. user
2. group
3. printer
4. storageSet
5. applicationGroup
6. application
7. folder
8. cabinet
9. hold

Each object that is added to the XML file contains the following information:

- A Start of Object tag
- A list of object attributes
- An optional list of child objects
- An End of Object tag

Start of object: An object is specified in the XML file by placing the object name after a < symbol. For example,

```
<user
```

starts a user object.

Important: The capitalization of the object names is important and should be used exactly as shown in the data tables in “XML objects in the Content Manager OnDemand XML file” on page 165.

Object attributes: All of the information about the object is contained within the object attributes. To add an attribute, specify the attribute name followed by an equal symbol followed by the value of the attribute in quotation marks. (All attribute values need to be enclosed in a set of double quotes even if the value is a numeric value or a single character.) Following the attribute values, the greater-than symbol is used to indicate the end of the object tag. For example:

```
<user name="SAMPLEUSER" phone="(212) 555-1212" timeOut="4" >
```

This tag indicates a user with the name SAMPLEUSER, the phone number (212) 555-1212, and the time out value 4 minutes. All the attributes that can be specified for each object, as well as the possible values and default values, are shown in the XML data tables.

Important:

1. The capitalization of the object names is important and should be used exactly as shown in the data tables in “XML objects in the Content Manager OnDemand XML file” on page 165.
2. When you add a user, IBM Content Manager OnDemand converts lowercase letters in the user ID to uppercase. You can type the user ID in uppercase, lowercase, or mixed case letters. In the above example, whether you enter the user ID as SAMPLEUSER, sampleuser, or SampleUser, Content Manager OnDemand automatically converts it to SAMPLEUSER.

Child objects: Child objects are constructed the same way as the eight main OnDemand objects. All child objects must occur after the object to which they are associated, and before the end of the object tag. In this example, there are two user permission children for the SampleUser user:

```
<user name="SampleUser" phone="(212) 555-1212" timeOut="4" >
<permission user="APP1" adminAuthority="Yes" />
<permission user="APP2" adminAuthority="Yes" />
```

Important: If an object cannot contain any children, such as the permission objects shown above, it must be ended with a slash / and the > symbol.

See “XML objects in the Content Manager OnDemand XML file” on page 165 for detailed information about the types of child objects that each object can have and, in some cases, the maximum number of child objects that can be created.

Some of the child objects themselves can have children. These are defined in the same way as above. See the second example in the Examples section.

End of object: The last item the user object needs is the end of object tag. This is indicated by placing the name of the object between </ and >. For example, </user>

indicates the end of the user object.

As with any object if the object does not contain child objects, the object can be ended by placing a /> at the end of the object definition. So the following,

```
<group name="Sample" gid="84000">
</group>
```

is equivalent to:

```
<group name="Sample" gid="84000"/>
```

Examples:

The following example shows a complete OnDemand XML file that contains two users and a user group:

```
<?xml version="1.0" encoding="UTF-8"?>
<onDemand xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xsi:noNamespaceSchemaLocation="ondemand.xsd">

  <user
    name="SampleUsr1" email="Sample@us.ibm.com" printer="Sample"
    userType="User Admin"
    createFoldersAuth="Yes">
    <permission user="SampleUsr2" adminAuthority="Yes" />
    <permission user="SampleUsr3" adminAuthority="No" />
    <permission group="SampleGroup1"/>
  </user>
```

```

<user
  name="SampleUsr2" password="xxxxxxx" timeOut="No Limit"
  description="This is a description of Sample User Two">
</user>

<group
  name="SampleGroup1"
  description="A Sample Group Number One"
  ownerUser="Admin">
  <user name="SampleUsr2"/>
  <user name="SampleUsr3"/>
</group>

</onDemand>

```

The following example shows a folder with a field child, and the field child has a fieldInfo child.

```

<?xml version="1.0" encoding="UTF-8"?>
<onDemand xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xsi:noNamespaceSchemaLocation="ondemand.xsd">
  <folder name="FolderOne" searchType="Hit List">
    <field name="cost" fieldType="Decimal" >
      <fieldInfo user="SampleUser" sortOrder="2" greaterThan="Default" lessThan="Yes"/>
    </field>
  </folder>
</onDemand>

```

Creating OnDemand XML files for update: When you create an XML file to be used for updating information, you should use the same syntax and layout as discussed in the previous section. However, there are two major differences.

- You only need to specify those fields that you intend to be updated.
- Child objects might contain the additional attribute task. The task attribute indicates the task that is performed by the child object.

Each object that you update must contain the name of the object. Any other attributes that are specified are updated to the specified value. If an object that is updated requires a change to another object, that object must be included in the XML file before the object that references it. For example, to change the user SampleUser to timeout after 10 minutes, you need to put the following XML code into an XML file.

```
<user name="SampleUser" timeOut="10"/>
```

For all of the main objects, the name field can be specified as `_ALL` to update all of the objects of that type. For example, to remove the authority to create folders from all of the defined users, you can use the following XML code:

```
<user name="_ALL" createFoldersAuth="No"/>
```

When you specify some special characters in XML, you need to use the XML specification. For example:

Table 9. XML specification for special characters

XML code	Character
&	&
'	'
"	"
<	<

Table 9. XML specification for special characters (continued)

XML code	Character
>	>

Your input file might look like this:

```
...
<user
  name="DBRYANTDEU"
  acctInfo="Engraving & Printing"
```

To rename some objects, you should use the `newName` attribute. For example, to change the name of a printer from Boston to New York, use the following XML code during an update:

```
<printer name="Boston" newName="NewYork"/>
```

Most objects might contain a set of similar child objects. To update these objects, you should use the `task` attribute. The `task` attribute can have three values:

add Indicates that the child object will be added to the parent object. The default value for the `task` attribute is `add`. However, some child objects do not contain a `task` attribute. For those objects without a `task` attribute, the action taken for the object is the same action as the parent object.

update

This indicates that the child object already exists and should be updated with the attributes provided.

delete This indicates that the child object already exists and should be removed from the parent object.

The `task` attribute is examined only during an update process. To add, update or delete a child object, the parent object must be specified, and must contain the child object to be updated.

The default value for the `task` attribute is `add`. However, some child objects do not contain a `task` attribute. For those objects without a `task` attribute, the action taken for the object will be the same action as the parent object.

When you update or delete child permission objects, you might specify the value of `_ALL` for the user or the group, to indicate that the update or removal should pertain to all of the users or groups that are already defined for the parent object.

Examples:

To add a permission for the user `SampleUser` to the folder `FolderOne`, the following code can be used during an update:

```
<folder name="FolderOne">
<permission user="SampleUser" adminAuthority="Yes" maxHits="No Limit" />
</folder>
```

To remove the administration authority permission from `SampleUser` on the folder `FolderOne`, you need to update the permission child. The following code is an example:

```
<folder name="FolderOne">
<permission task="update" user="SampleUser" adminAuthority="No" />
</folder>
```

To remove the SampleUser permission from the folder, you need to delete the permission child.

```
<folder name="FolderOne">  
<permission task="delete" user="SampleUser"/>  
</folder>
```

To remove all user permissions that are currently defined for the folder FolderOne, the following XML code can be used during an update:

```
<folder name="FolderOne">  
<permission task="delete" user="_ALL"/>  
</folder>
```

Creating XML files for delete and export: When you create an XML file to use for exporting or deleting objects, you should use the same syntax and layout as described in previous sections. However, there are two major differences:

- The only attribute field that is examined is the name field.
- All child objects are ignored.

When you export or delete objects, the only information that is required is the name of the objects. All other attributes might be present in the XML file, but are ignored. It is possible to delete only child objects through the update process. During a delete process all child objects are ignored.

When you specify an application for delete or export, you must also specify the name of the application group in which the application is contained.

The name of _ALL can be used during export and delete to indicate that all objects of that type should be exported or deleted.

To export the users SampleOne, SampleTwo, and SampleThree, you can use the following OnDemand XML file:

```
<?xml version="1.0" encoding="UTF-8"?>  
<onDemand xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"  
  xsi:noNamespaceSchemaLocation="ondemand.xsd">  
  
  <user name="SampleOne"/>  
  <user name="SampleTwo"/>  
  <user name="SampleThree"/>  
</onDemand>
```

To delete the printer, OldPrinter, and the folder, MyFolder, use the following XML code:

```
<printer name="OldPrinter"/>  
<folder name="MyFolder"/>
```

To delete all defined printers, the following XML code could be used during a delete process.

```
<printer name="_ALL"/>
```

Application Index Parameter Object: The indexParm child object of the application object has a slightly different syntax than all other objects. Owing to the type of the data associated with the Index Parameters, you cannot use an object attribute to specify this information. Thus, for this object (and only this object), the data associated with the object is simply the character data that appears between the <indexParm> and the </indexParm> flags.

For example, the following code can be used to set the index parameters associated with the SampleApp application:

```
<application name="SampleApp" .... >
<indexParm>
*dummy index parameters
parm1=value;
parm2=value;
parm3=value;
</indexParm>
</application>
```

Importing the XML file by using the ARSXML command

Use the ARSXML program to import the XML file into the Content Manager OnDemand system. For syntax, description, and parameters of this command, see “ARSXML” on page 304.

Exporting OnDemand administrative objects to an XML file

You can use the ARSXML command to export objects from Content Manager OnDemand into an XML file in several ways:

- You can export a single object, such as a single user.
- You can export all the defined objects of a particular type, for example, all user groups
- You can export any combination of objects, for example, a user group and all users in that group

This feature can be used to back up part of a Content Manager OnDemand system, copy objects from one system to another, or store objects into an XML file so that they can be processed by another application.

Use the ARSXML command to export administrative objects into an XML file. For syntax, description, and parameters of this command, see “ARSXML” on page 304.

If you are exporting objects that have dependencies on other objects: Some objects in Content Manager OnDemand might have dependencies on other objects. For example, a group object has a dependency on all of the users that are defined within the group. When you export these objects, you can use the `-r` parameter with a value of `d` to include in the XML file all of the dependent objects that the exported objects might have.

See “ARSXML” on page 304 for more example XML files that are generated in different export scenarios.

Objects and data model used in the XML file

XML objects in the Content Manager OnDemand XML file

You can use an XML interface to import data into a Content Manager OnDemand system. Before importing data, you need to have an XML file that contains all the data that needs to be imported. The XML file can be either created during a previous XML export process or written from scratch. This section helps you to understand the objects in the XML file.

The following diagram illustrates the overall layout of the Content Manager OnDemand XML file.




Figure 6. Overall layout of the Content Manager OnDemand XML file

The following tables provide detailed descriptions about the objects in the Content Manager OnDemand XML file. The objects include:

- | • user
- | • group
- | • printer
- | • storageSet
- | • applicationGroup
- | • application
- | • folder
- | • cabinet
- | • hold

If you create an input XML file, ensure that you include the objects following this order:

1. user
2. group
3. printer
4. storageSet
5. applicationGroup
6. application
7. folder
8. cabinet
9. hold

For each object, the following information is provided:

Name The name of the attribute. An attribute name in bold indicates that the attribute is a key attribute, which is used to uniquely identify the object. For all top-level objects, this is the Name attribute.

DataType

The XML data type of the attribute. Types in italics indicates an XML data type that is defined in the Content Manager OnDemand XML schema file.

Required

Whether this attribute is required. Dependent means that the attribute may be required based on the settings of other attributes.

Default Value

The value that is used internally if the optional attribute is not specified in the XML file.

Possible Values

Lists any limitations on the value for attributes. If there are no limitations, this is blank.

Updateable

Whether this attribute can be updated.

For each child object, the following information is provided:

Name The name of the child object. An attribute name in bold indicates that the attribute is a "key" attribute which is used to uniquely identify the object.

Min/Max Number

Indicates the minimum and maximum number of children of this type that are allowed.

Deleteable

Whether this child object can be deleted.

Updateable

Whether this child object can be updated.

For all XML objects, the data that is associated with an object is only in the form of attribute data and child objects. Not all attributes in the tables are supported on all Content Manager OnDemand platforms.

Important: The capitalization of the object names is important and should be used exactly as shown in the data tables.

User object

Table 10. Attributes for the user object

Name	Description	Data Type	Required	Default value	Possible values	Updateable
name	Name of the user	<i>nameString</i>	Yes		<ul style="list-style-type: none"> User name "_ALL" 	Yes
uid	User ID value	Integer	No	0	<ul style="list-style-type: none"> Integer value 1–1,080,000 0 to indicate system generated values 	No
description	Description of the user	<i>descString</i>	No	""		Yes
password	User's password	<i>passwordString</i>	No	""		Yes
userType	Type of user	<i>userTypeString</i>	No	"User"	<ul style="list-style-type: none"> "User" "User Admin" "AG/Folder/Cabinet Admin" "System Admin" "Hold Admin" 	Yes
timeOut	Time out length in minutes	<i>timeOutType</i>	No	"Use System Value"	<ul style="list-style-type: none"> "Use System Value" "No Limit" Any positive integer value 	Yes
fullName	User's full name	<i>userMiscString</i>	No	""		Yes
acctInfo	Accounting for the user	<i>userMiscString</i>	No	""		Yes
company	Name of the company	<i>userMiscString</i>	No	""		Yes
title	User's title of job responsibility	<i>userMiscString</i>	No	""		Yes
addr1	User's address	<i>userMiscString</i>	No	""		Yes
addr2	Address, line number 2	<i>userMiscString</i>	No	""		Yes
addr3	Address, line number 3	<i>userMiscString</i>	No	""		Yes
addr4	Address, line number 4	<i>userMiscString</i>	No	""		Yes

Table 10. Attributes for the user object (continued)

Name	Description	Data Type	Required	Default value	Possible values	Updateable
dept	User's department name or number	<i>userMiscString</i>	No	""		Yes
building	User's building name or number	<i>userMiscString</i>	No	""		Yes
room	User's office name or number	<i>userMiscString</i>	No	""		Yes
phone	User's phone number	<i>phoneString</i>	No	""		Yes
fax	User's fax number	<i>phoneString</i>	No	""		Yes
email	User's email address	<i>emailString</i>	No	""		Yes
coverPage	Default Fax cover page for the user	<i>userMiscString</i>	No	""		Yes
printer	User's default printer	<i>nameString</i>	No	"NONE"		Yes
createCabinetsAuth	Whether this user or group has authority to create cabinets	<i>yesnoString</i>	No	"No"	<ul style="list-style-type: none"> • "Yes" • "No" 	Yes
createFoldersAuth	Whether this user have authority to create folders	<i>yesnoString</i>	No	"No"	<ul style="list-style-type: none"> • "Yes" • "No" 	Yes
createHoldsAuth	Whether this user has authority to create holds	<i>yesnoString</i>	No	"No"	<ul style="list-style-type: none"> • "Yes" • "No" 	Yes
createUsersAuth	Whether this user has authority to create users	<i>yesnoString</i>	No	"No"	<ul style="list-style-type: none"> • "Yes" • "No" 	Yes
createGroupsAuth	Whether this user have authority to create group	<i>yesnoString</i>	No	"No"	<ul style="list-style-type: none"> • "Yes" • "No" 	Yes
createAppGroupsAuth	Whether this user have authority to create application groups	<i>yesnoString</i>	No	"No"	<ul style="list-style-type: none"> • "Yes" • "No" 	Yes

Table 10. Attributes for the user object (continued)

Name	Description	Data Type	Required	Default value	Possible values	Updateable
disableUser	Specifies whether this user should be disabled	<i>yesnoString</i>	No	"No"	<ul style="list-style-type: none"> "Yes" "No" 	Yes
passwordExp	Specifies when the user's password expires	<i>pwExpValue</i>	No	"Use System Value"	<ul style="list-style-type: none"> "Use System Value" "Never Expires" An integer between 1 and 365 	Yes
printerClass	The printer class	<i>char</i>	No	""		Yes
printerBanner	Specifies whether there should be a printer banner	<i>yesnoString</i>	No	"No"	<ul style="list-style-type: none"> "Yes" "No" 	Yes
printerDest	The printer destination	<i>valString</i>	No	""		Yes
printerNode	The printer node	<i>valString</i>	No	""		Yes
printerWriter	The printer writer	<i>valString</i>	No	""		Yes
printerPagedef	The printer pagedef	<i>valString</i>	No	""		Yes
printerFormdef	The printer formdef	<i>valString</i>	No	""		Yes
printerForms	The printer forms	<i>valString</i>	No	""		Yes
printerRouting	The printer routing	<i>routingString</i>	No	""		Yes

Table 11. Child object under the user object

Name	Minimum number	Maximum number	Deleteable	Updateable
permission	0	Unbounded	Yes	Yes

Table 12. Attributes for the permission object under the user object

Attributes	Description	Datatype	Required	Default Value	Possible Values	Updateable
task	The task to perform on this object	<i>taskString</i>	No	"add"	<ul style="list-style-type: none"> "add" "delete" "update" 	N/A
user	The name of a user that can view and maintain this user	<i>nameString</i>	Yes	""	<ul style="list-style-type: none"> User name 	No

Table 12. Attributes for the permission object under the user object (continued)

Attributes	Description	Datatype	Required	Default Value	Possible Values	Updateable
group	The name of a group that can view and maintain this user	<i>nameString</i>	Yes	""	<ul style="list-style-type: none"> Group name 	No
adminAuthority	Whether this user or group has authority to update and delete the user	<i>yesnoString</i>	No	"No"	<ul style="list-style-type: none"> "Yes" "No" 	Yes

Remember: The "_ALL" value can only be used in an ARSXML export operation.

In the permission child, you must specify either a user or a group object, however, you cannot specify both.

Group object

Table 13. Attributes for the group object

Name	Description	Data Type	Required	Default value	Possible values	Updateable
name	Name of the group	<i>nameString</i>	Yes		<ul style="list-style-type: none"> Group name "_ALL" 	Yes
gid	Group ID value	Integer	No	0	<ul style="list-style-type: none"> Integer value greater than or equal to 1080001 0 to indicate a system value 	No
description	Description of this group	<i>descString</i>	No	""		Yes
ownerUser	Name of the user that is the owner of this group	<i>nameString</i>	No	"NONE"		Yes
ownerGroup	Name of the group that is the owner of this group	<i>nameString</i>	No	"NONE"		Yes

Note: You can specify an ownerUser or an ownerGroup attribute, however, you cannot specify both.

Table 14. Child object under the group object

Name	Minimum number	Maximum number	Deleteable	Updateable
user	0	Unbounded	Yes	No

Table 15. Attributes for the user object under the group object

Attributes	Description	Datatype	Required	Default Value	Possible Values	Updateable
task	The task to perform on this object	<i>taskString</i>	No	"add"	<ul style="list-style-type: none"> • "add" • "delete" 	N/A
name	The name of a user that is a member of this group	<i>nameString</i>	Yes			No

Remember: The "_ALL" value can only be used in an ARSXML export operation.

Application object

Table 16. Attributes for the application object

Name	Description	Data Type	Required	Default value	Possible values	Updateable
name	Name of the application	<i>nameString</i>	Yes		<ul style="list-style-type: none"> Application name "_ALL" 	Yes
description	Description of the application	<i>descString</i>	No	""		Yes
appGroup	The application group that this application belongs to	<i>nameString</i>	Yes			No
identifier	The application group identifier	<i>dbString2</i>	Yes	""	Determined by the application group	Yes
dataType	The data type of this application	<i>appTypeString</i>	No	<ul style="list-style-type: none"> "AFP" "SCS" 	<ul style="list-style-type: none"> "AFP" "Line" "BMP" "GIF" "PCX" "TIFF" "PDF" "JPEG" "PNG" "User Defined" "OD Defined" "E-Mail" "None" "SCS" SCS-Extended Global DJDE 	No
indexer	The name of the indexer program to use with this application	<i>indexerString</i>	No	"None"	<ul style="list-style-type: none"> "None" "ACIF" "Generic" "PDF" OS/390 OS/400 	Yes
dataCompression	The compression technique that is used for the data	<i>compressionString</i>	No	"OD77"	<ul style="list-style-type: none"> "OD77" "LZW12" "LZW16" "Disable" "None" 	Yes
resCompression	The compression technique that is used for the resources	<i>compressionString</i>	No	"OD77"	<ul style="list-style-type: none"> "OD77" "LZW12" "LZW16" "Disable" "None" 	Yes

Table 16. Attributes for the application object (continued)

Name	Description	Data Type	Required	Default value	Possible values	Updateable
resourceComparison	The number of resource comparisons	<i>resCompValue</i>	No	50	0-9999	Yes
largeObject	Whether this application a Large Object	<i>yesnoString</i>	No	"No"	<ul style="list-style-type: none"> • "Yes" • "No" 	Yes
size	The number of pages or size of the object in kilobytes that is based on the setting of the largeObject attribute	Integer	No	100	Positive integer representing either object size or number of pages based on setting of largeObject	Yes
pageIdentifiers	Whether page identifiers should be used	<i>yesnoString</i>	No	"No"	<ul style="list-style-type: none"> • "Yes" • "No" 	Yes
postProcessorString	The post processor string	<i>postProcString</i>	No	""		Yes
printAmount	The default print option to specify which page or pages to print	<i>printAmString</i>	No	"All Pages"	<ul style="list-style-type: none"> • "All Pages" • "Current Pages" 	Yes
defaultPrinter	Default printer	<i>nameString</i>	No	"*NONE"		Yes
cicsJCLParms	Values that are used by the CICS® or ESA client program to print documents	string	No	"None"		Yes
printParms	Print parameters that are used by the server print manager	string	No	""		Yes
afpToLine	Options that OnDemand uses to convert AFP data to line data	string	No	""		Yes
usePreview	Whether to use the preview user exit	<i>yesnoString</i>	No	"No"	<ul style="list-style-type: none"> • "Yes" • "No" 	Yes
previewParms	Whether the preview parameters should be used	string	No	""		Yes

Table 16. Attributes for the application object (continued)

Name	Description	Data Type	Required	Default value	Possible values	Updateable
agQueryName	Name of the application group that has all the applications you want to export. When you specify agQueryName during export, the values for the name and appGroup attributes are ignored. If you specify agQueryName, you cannot specify folderQueryName.	string	No	""		No
folderQueryName	Name of folder that has the applications you want to export. When you specify folderQueryName during export, the values for the name and appGroup attributes are ignored. If you specify folderQueryName, you cannot specify agQueryName.	string	No	""		No

Note:

1. If the application group contains more than one field object, identifier is required.
2. The identifier object can be updated to an unused field value only.

Table 17. Child objects under the application object

Name	Minimum number	Maximum number	Deleteable	Updateable
lineData	0	1	No	Yes
SCSData	0	1	No	Yes
afpData	0	1	No	Yes
pdfData	0	1	No	Yes
imageData	0	1	No	Yes
udData	0	1	No	Yes
indexParm	0	1	Yes	Yes
preprocessParm	0	Dependent	No	Yes
imageOverlay	0	Unbounded	Yes	No

Important: The udData child is required if the application dataType is User Defined.

Table 18. Attributes for the SCSData object under the application object

Attributes	Description	Data Type	Required	Default value	Possible values	Updateable
codePage	The code page to use	Integer	No	37	Nonnegative integer	Yes
prMode	Processing Mode	<i>prmodeString</i>	No	"None"	<ul style="list-style-type: none"> • "None" • "SOSI1" • "SOSI2" • "SOSI3" 	Yes

Table 19. Child objects under the SCSData object

Name	Minimum number	Maximum number	Deleteable	Updateable
IvHeader	0	1	No	Yes
IvField	0	30	Yes	Yes
logView	0	20	Yes	Yes

See Table 25 on page 190, Table 26 on page 191, and Table 27 on page 192 for descriptions of the lvHeader, lvField, and logView objects.

Table 20. Attributes for the SCSExData object under the application object

Attributes	Description	Data Type	Required	Default value	Possible values	Updateable
codePage	The code page to use	Integer	No	37	Nonnegative integer	Yes
orientation	The orientation or rotation angle, in degrees, of the document	<i>orientString</i>	No	0	<ul style="list-style-type: none"> • 0 • 90 • 180 • 270 	Yes
paperSize	The output paper size	<i>paperString</i>	No	"Autosize"	See the list that follows this table for possible values for the paperSize attribute.	Yes

Possible values for the paperSize attribute are:

- A3
- A4 Wide
- A4
- A5
- Autosize
- B4 Wide
- B4
- B5
- Cut Sheet
- Euro Fanfold
- Executive 1
- Executive 2
- Executive 3
- Folio
- Ledger
- Legal Wide
- Legal
- Letter Wide
- Letter
- Narrow
- Quarto
- Short
- Statement
- Stationery
- Tabloid
- Wide
- 1403W
- 1403WS
- 3800N
- 3800NS
- 3800W
- 3800WS

Table 21. Child object under the SCSExData object

Name	Minimum number	Maximum number	Deleteable	Updateable
logView	0	20	Yes	Yes

Table 22. Attributes for the *logView* object under the *SCSExData* object

Attributes	Description	Data Type	Required	Default value	Possible values	Updateable
task	The task to perform on this object	<i>taskString</i>	No	"add"	<ul style="list-style-type: none"> "add" "delete" "update" 	N/A
name	Name of this logical view	<i>nameString</i>	Yes			Yes
user	User for this logical view	<i>nameString</i>	Yes		<ul style="list-style-type: none"> User name "*PUBLIC" 	No
defaultView	Whether this is the default view	<i>yesnoString</i>	No	"No"	<ul style="list-style-type: none"> "Yes" "No" 	Yes
zoom	Zoom Factor	Integer	No	100	Positive integer	Yes
selAreaColor	Selected Area Color	<i>saColorString</i>	No	"Yellow"	<ul style="list-style-type: none"> "White" "Black" "Red" "Blue" "Green" "Yellow" "Grey" "Custom" 	Yes
bgColor	Background	<i>bgColorString</i>	No	"White"	<ul style="list-style-type: none"> "White" "Black" "Red" "Blue" "Green" "Yellow" "Grey" "Green Bar" "Custom" 	Yes

Table 23. Attributes for the *lineData* object

Attributes	Description	Data Type	Required	Default value	Possible values	Updateable
lineCount	The line count	Integer	No	0	Nonnegative integer	Yes
codePage	The code page to use	Integer	No	500	Positive integer	Yes
recFM	The record format	<i>recfmString</i>	No	"Variable"	<ul style="list-style-type: none"> "Fixed" "Variable" "Stream" 	Yes
LRECL	The record length	Integer	No	133	Positive integer	Yes
delimiter	The line delimiter (specified in hex)	<i>hexString</i>	No	"0A"	Hex character string	Yes
prMode	Processing mode	<i>prmodeString</i>	No	"None"	<ul style="list-style-type: none"> "None" "SOSI1" "SOSI2" "SOSI3" 	Yes
trc	Whether the data contains table reference characters	<i>yesnoString</i>	No	"No"	<ul style="list-style-type: none"> "Yes" "No" 	Yes
CCType	The type of the carriage control character	<i>yesnoString</i>	No	"ANSI"	<ul style="list-style-type: none"> "ANSI" "Machine" "None" 	Yes

Table 24. Child objects under the lineData object

Name	Minimum number	Maximum number	Deleteable	Updateable
lvHeader	0	1	No	Yes
lvField	0	30	Yes	Yes
logView	0	20	Yes	Yes

Table 25. Attributes for the *IvHeader* object

Attributes	Description	Data Type	Required	Default value	Possible values	Updateable
pageHdrRows	Number of page header rows	Integer	No	0	Nonnegative integer	Yes
fieldHdrRows	Number of field header rows	Integer	No	0	Nonnegative integer	Yes
validationRow	The validation row	Integer	Dependent	0	Nonnegative integer	Yes
validationCol	The validation column	Integer	Dependent	0	Nonnegative integer	Yes
validationString	The validation string	String	Dependent	""		Yes

Table 26. Attributes for the IvField object

Attributes	Description	DataType	Required	Default value	Possible values	Updateable
task	The task to perform on this object	<i>taskString</i>	No	"add"	<ul style="list-style-type: none"> "add" "delete" "update" 	N/A
name	The name of the logical view field	<i>nameString</i>	Yes			Yes
dbName	The application group database name	<i>nameString</i>	No	"*NONE"	database field name or "*NONE"	Yes
startCol	The start column for this field definition	Integer	Yes	0	An integer value between 1 and 999	Yes
endCol	The end column for this field definition	Integer	Yes	0	an integer value between 1 and 999	Yes

Table 27. Attributes for the *logView* object

Attributes	Description	Data Type	Required	Default value	Possible values	Updateable
task	The task to perform on this object	<i>taskString</i>	No	"add"	<ul style="list-style-type: none"> "add" "delete" "update" 	N/A
name	Name of this logical view	<i>nameString</i>	Yes			Yes
user	user for this logical (or PUBLIC) view	<i>nameString</i>	Yes		<ul style="list-style-type: none"> User name "*PUBLIC" 	No
defaultView	Whether this is the default view	<i>yesNoString</i>	No	"No"	<ul style="list-style-type: none"> Yes No 	Yes
zoom	Zoom factor	Integer	No	100	Positive integer	Yes
selAreaColor	Selected area color	<i>saColorString</i>	No	"Yellow"	<ul style="list-style-type: none"> "White" "Black" "Red" "Blue" "Green" "Yellow" "Grey" "Custom" 	Yes
bgColor	Background color	<i>bgColorString</i>	No	"White"	<ul style="list-style-type: none"> "White" "Black" "Red" "Blue" "Green" "Yellow" "Grey" "Green Bar" "Custom" 	Yes

Table 27. Attributes for the *logView* object (continued)

Attributes	Description	Data Type	Required	Default value	Possible values	Updateable
headingColor	Heading color	<i>hdgColorString</i>	No	"BG Color"	<ul style="list-style-type: none"> "BG Color" "White" "Black" "Red" "Blue" "Green" "Yellow" "Grey" "Custom" 	Yes
textFidelity	Text fidelity	<i>fidelityString</i>	No	"Line"	<ul style="list-style-type: none"> "Line" "Word" "Char" "Draft" "240-pel" 	Yes
overstrike	Text overstrike option	<i>overstrikeString</i>	No	"Separate Lines"	<ul style="list-style-type: none"> "Separate Lines" "Merge Lines" "Overstrike Lines" 	Yes
fieldOrder	A string showing the order in which the fields are to be displayed	String	No	""		Yes
lockHeader	Whether the header lines should be locked in place during scrolling	<i>yesnoString</i>	No	"No"	<ul style="list-style-type: none"> "Yes" "No" 	Yes
lockField	Number of fields that you want to lock in place	Integer	No	0	Nonnegative integer	Yes

Table 28. Attributes for the *afpData* object

Attributes	Description	Data Type	Required	Default value	Possible values	Updateable
orientation	The orientation or rotation angle, in degrees, of the document	<i>orientString</i>	No	0	<ul style="list-style-type: none"> • "0" • "90" • "180" • "270" 	Yes
paperSize	The output paper size	<i>paperString</i>	No	"Autosize"	See the list that follows this table for possible values for the paperSize attribute.	Yes

The following are possible values for the paperSize attribute:

- A3
- A4 Wide
- A4
- A5
- Autosize
- B4 Wide
- B4
- B5
- Cut Sheet
- Euro Fanfold
- Executive 1
- Executive 2
- Executive 3
- Folio
- Ledger
- Legal Wide
- Legal
- Letter Wide
- Letter
- Narrow
- Quarto
- Short
- Statement
- Stationery
- Tabloid
- Wide
- 1403W
- 1403WS
- 3800N
- 3800NS
- 3800W

Table 29. Child objects under the *afpData* object

Name	Minimum number	Maximum number	Deleteable	Updateable
logView	0	20	Yes	Yes

Table 30. Attributes for the *logView* object under the *afpData* object

Attributes	Description	DataType	Required	Default value	Possible values	Updateable
task	The task to perform on this object	<i>taskString</i>	No	"add"	<ul style="list-style-type: none"> • "add" • delete • "update" 	N/A
name	Name of this logical view	<i>nameString</i>	Yes			Yes
user	User for this logical view (or PUBLIC)	<i>nameString</i>	Yes		<ul style="list-style-type: none"> • User name • "*PUBLIC" 	No
defaultView	Whether this is the default view	<i>yesnoString</i>	No	"No"	<ul style="list-style-type: none"> • Yes • No 	Yes
zoom	Zoom factor	Integer	No	100	Positive integer	Yes
selAreaColor	Selected area color	<i>saColorString</i>	No	"Yellow"	<ul style="list-style-type: none"> • "White" • "Black" • "Red" • "Blue" • "Green" • "Yellow" • "Grey" • "Custom" 	Yes
bgColor	Background color	<i>bgColorString</i>	No	"White"	<ul style="list-style-type: none"> • "White" • "Black" • "Red" • "Blue" • "Green" • "Yellow" • "Grey" • "Green Bar" • "Custom" 	Yes
imageColor	Image color	<i>imgColorString</i>	No	"Black"	<ul style="list-style-type: none"> • "Black" • "Red" • "Blue" • "Green" • "Yellow" • "Magenta" • "Cyan" • "Custom" 	Yes
textFidelity	Text fidelity	<i>fidelityString</i>	No	"Line"	<ul style="list-style-type: none"> • "Line" • "Word" • "Char" • "Draft" • "240-pel" 	Yes
imageIntensity	Image intensity	<i>intensityString</i>	No	"Normal"	<ul style="list-style-type: none"> • "Normal" • "Light" • "None" 	Yes
copyGroup	The copy group for the document	Integer	No	1	Nonnegative integer	Yes

Table 31. Attributes for the pdfData object

Attributes	Description	Data Type	Required	Default value	Possible values	Updateable
orientation	The orientation or rotation angle, in degrees, of the document	<i>orientString</i>	No	0	<ul style="list-style-type: none"> • 0 • 90 • 180 • 270 	Yes

Table 32. Child objects under the pdfData object

Name	Minimum number	Maximum number	Deleteable	Updateable
logView	0	Unbounded	Yes	Yes

Table 33. Attributes for the *logView* object under *pdfData*

Attributes	Description	Data Type	Required	Default value	Possible values	Updateable
task	The task to perform on this object	<i>taskString</i>	No	"add"	<ul style="list-style-type: none"> "add" delete "update" 	N/A
name	Name of this logical view	<i>nameString</i>	Yes			Yes
user	User for this logical view (or PUBLIC)	<i>nameString</i>	Yes		<ul style="list-style-type: none"> User name "PUBLIC" 	No
defaultView	Whether this is the default view	<i>yesnoString</i>	No	"No"	<ul style="list-style-type: none"> Yes No 	Yes
zoom	Zoom factor	Integer	No	100	Positive integer	Yes
selAreaColor	Selected area color	<i>saColorString</i>	No	"Yellow"	<ul style="list-style-type: none"> "White" "Black" "Red" "Blue" "Green" "Yellow" "Grey" "Custom" 	Yes

Table 34. Attributes for the *imageData* object

Attributes	Description	DataType	Required	Default value	Possible values	Updateable
orientation	The orientation or rotation angle, in degrees, of the document	<i>orientString</i>	No	0	<ul style="list-style-type: none"> • 0 • 90 • 180 • 270 	Yes

Table 35. Child objects under the *imageData* object

Name	Minimum number	Maximum number	Deleteable	Updateable
logView	0	20	Yes	Yes

Table 36. Attributes for the logView object under imageData

Attributes	Description	DataType	Required	Default value	Possible values	Updateable
task	The task to perform on this object	<i>taskString</i>	No	"add"	<ul style="list-style-type: none"> • "add" • delete • "update" 	N/A
name	Name of this logical view	<i>nameString</i>	Yes			Yes
user	User for this logical view (or PUBLIC)	<i>nameString</i>	Yes		<ul style="list-style-type: none"> • User name • "*PUBLIC" 	No
defaultView	Whether this is the default view	<i>yesnoString</i>	No	"No"	<ul style="list-style-type: none"> • "Yes" • "No" 	Yes
zoom	Zoom factor	Integer	No	100	Integer value 10 — 1000	Yes
contrast	Image contrast value	Integer	No	128	Integer value 0–255	Yes
brightness	Image brightness value	Integer	No	128	Integer value 0–255	Yes
greyScale	Whether the image should use grey scales	<i>yesnoString</i>	No	"No"	<ul style="list-style-type: none"> • "Yes" • "No" 	Yes

Table 37. Attributes for the *udData* object

Attributes	Description	DataType	Required	Default value	Possible values	Updateable
fileExt	File extension	<i>fileExtString</i>	Yes			Yes

Table 38. Attributes for the *indexParm* object

Attributes	Description	DataType	Required	Default value	Possible values	Updateable
task	The task to perform on this object	<i>taskString</i>	No	"add"	<ul style="list-style-type: none"> • "add" • "delete" • "update" 	n/a
<character data>	The character data that represents the index parameters for this application	<i>string</i>	Yes			Yes

Important: The index parameters are not represented by any attribute data. They are represented by the character data field between the <indexParm> and </indexParm> fields.

Table 39. Attributes for the preprocessParm object

Attributes	Description	Data Type	Required	Default value	Possible values	Updateable
dbName	Application group database name	dbString	Yes			No
loadIDName	Load ID name	namestring	No	dbName		Yes
defaultValue	The default value	string	No	""		Yes
embedded	The embedded characters to remove	string	No	","		Yes
leading	The leading characters to remove	string	No	""		Yes
trailing	The trailing characters to remove	string	No	""		Yes
divideBy	The value to divide by if the field is numeric	Integer	No	1 or 100 based on data type	Positive integer	Yes
format	The date format to use	string	No	%m/%d/%Y		Yes

Table 40. Attributes for the imageOverlay object

Attributes	Description	Data Type	Required	Default value	Possible values	Updateable
task	The task to perform on this object	taskString	No	"add"	<ul style="list-style-type: none"> "add" "delete" 	N/A
id	The name of the image overlay that is added to the application	string	Yes			No
fileName	The full path name of the DLL file that contains the image overlay bitmap that is added to the application	string	Yes			No

Remember: The "_ALL" value can only be used in an ARSXML export operation.

Application group object

Table 41. Attributes for the applicationGroup object

Name	Description	Data Type	Required	Default value	Possible Values	Updateable
autoFTI	All data loaded into the OnDemand server is processed by the full text indexer.	yesnoString	No	""	<ul style="list-style-type: none"> "Yes" "No" 	Yes
name	Name of the application group	nameString	Yes		<ul style="list-style-type: none"> Application group name "_ALL" 	Yes
database	Name of the database	nameString	No	""		Yes
description	Description of the application group	descString	No	""		Yes

Table 41. Attributes for the applicationGroup object (continued)

Name	Description	DataType	Required	Default value	Possible Values	Updateable
storageSet	Name of the storage set to use for this application group	<i>nameString</i>	No	""		Yes
cacheDataLen	How many days to cache the data for	<i>cacheDataLenValue</i>	No	90	<ul style="list-style-type: none"> "No Cache" "Search Cache" A positive integer 	Yes
CFSOD	Use CFSOD	<i>yesnoString</i>	No	"No"	<ul style="list-style-type: none"> "Yes" "No" 	Yes
declareCFSOD	Use to declare CFSOD	<i>yesnoString</i>	No	"No"	<ul style="list-style-type: none"> "Yes" "No" 	Yes
expirationType	Determines how data is deleted from the application group	<i>expTypeString</i>	No	"Load"	<ul style="list-style-type: none"> "Load" "Segment" "Document" z/OS only option: Storage Manager 	No
expirationDate	The number of days that OnDemand keeps documents, resources, and index data in the application group	<i>expDateValue</i>	No	90	<ul style="list-style-type: none"> "Never Expire" A positive integer 	Yes
federateCFSOD	Federate CFSOD	<i>yesnoString</i>	No	"No"	<ul style="list-style-type: none"> "Yes" "No" 	Yes
FTI	Whether to enable Full Text Indexing for the application group.	<i>yesnoString</i>	No	""	<ul style="list-style-type: none"> "Yes" "No" 	Yes
lastLoadDate	The date and time of the last load	<i>String</i>	No	""		No
lastQueryDate	The date and time of the last query	<i>String</i>	No	""		No
lastRetrieveDate	The date and time of the last retrieve	<i>String</i>	No	""		No
logRetrieval	Whether document retrieval messages should be logged	<i>yesnoString</i>	No	"Yes"	<ul style="list-style-type: none"> "Yes" "No" 	Yes
logBulkRetrieval	Whether document bulk retrieval messages should be logged	<i>yesnoString</i>	No	"Yes"	<ul style="list-style-type: none"> "Yes" "No" 	Yes
logDBQueriesAfter	Whether database queries (after the query) should be logged	<i>yesnoString</i>	No	"No"	<ul style="list-style-type: none"> "Yes" "No" 	Yes
logDBQueriesBefore	Whether document database queries should be logged	<i>yesnoString</i>	No	"No"	<ul style="list-style-type: none"> "Yes" "No" 	Yes
logServerPrinting	Whether document server printing messages should be documented	<i>yesnoString</i>	No	"Yes"	<ul style="list-style-type: none"> "Yes" "No" 	Yes
logIndexAdd	Whether document index add messages should be documented	<i>yesnoString</i>	No	"Yes"	<ul style="list-style-type: none"> "Yes" "No" 	Yes
logIndexUpdate	Whether document index update messages should be documented	<i>yesnoString</i>	No	"Yes"	<ul style="list-style-type: none"> "Yes" "No" 	Yes
logIndexDelete	Whether document index delete messages should be documented	<i>yesnoString</i>	No	"Yes"	<ul style="list-style-type: none"> "Yes" "No" 	Yes
logObjectStore	Whether object store messages should be documented	<i>yesnoString</i>	No	"Yes"	<ul style="list-style-type: none"> "Yes" "No" 	Yes

Table 41. Attributes for the applicationGroup object (continued)

Name	Description	Data Type	Required	Default value	Possible Values	Updateable
logObjectRetrieve	Whether object retrieve messages should be logged	<i>yesnoString</i>	No	"Yes"	<ul style="list-style-type: none"> "Yes" "No" 	Yes
logAnnotAdd	Whether annotation add messages should be logged	<i>yesnoString</i>	No	"Yes"	<ul style="list-style-type: none"> "Yes" "No" 	Yes
logAnnotUpdate	Whether annotation update messages should be logged	<i>yesnoString</i>	No	"Yes"	<ul style="list-style-type: none"> "Yes" "No" 	Yes
logAnnotDelete	Whether annotation delete messages should be logged	<i>yesnoString</i>	No	"Yes"	<ul style="list-style-type: none"> "Yes" "No" 	Yes
maxDBRows	The maximum number of database rows for this application group	Integer	No	10000000	<ul style="list-style-type: none"> Positive integer "Single table for all loads" 	Yes
loadsPerTable	The amount of loads per table	<i>loadsString</i>	No	"Multiple"	<ul style="list-style-type: none"> "Multiple" 	No
dbTablespaceType	The type of database table space to use for this application group	<i>dbtString</i>	No	"SMS"	<ul style="list-style-type: none"> "None" "SMS" "Automatic Storage" "Use Tablespace" 	Yes
dbCompress	For databases that support compression, OnDemand compresses the tables or table spaces. If the database does not support compression, OnDemand ignores this setting.	<i>yesnoString</i>	No	"No"	<ul style="list-style-type: none"> "Yes" "No" 	No
annot	Annotation flags in the document database tables	<i>yesnoString</i>	No	"No"	<ul style="list-style-type: none"> "Yes" "No" 	No
enhancedRetManagement	Use enhanced retention management	<i>yesnoString</i>	No	"No"	<ul style="list-style-type: none"> "Yes" "No" 	Yes
impliedHold	Use implied hold	<i>yesnoString</i>	No	"No"	<ul style="list-style-type: none"> "Yes" "No" 	Yes
objectSize	The object size in kilobytes	Integer	No	10,000	Positive integer	Yes
migrateCache	When to migrate data from cache	<i>migrateString</i>	No	"Load"	<ul style="list-style-type: none"> "Load" "Never" Migration A positive integer 	Yes
migrateIndex	When to migrate indexes from cache	<i>migrateIndString</i>	No	"Never"	<ul style="list-style-type: none"> "Never" A positive integer 	Yes
migrateIndexKeep	How long to keep imported migrated indexes	Integer	No	0	Nonnegative	Yes
parameterMarkers	Whether to use parameter markers	<i>yesnoString</i>	No	"Yes"	<ul style="list-style-type: none"> "Yes" "No" 	Yes
serverFTI	IP address of the full text indexer server	String	No	""		Yes
useFilegroups	Specifies whether this application group uses file groups	<i>yesnoString</i>	No	"No"	<ul style="list-style-type: none"> "Yes" "No" 	Yes
updateExistingTables	Specifies whether tables should be updated if an index changes	<i>yesnoString</i>	No	"No"	<ul style="list-style-type: none"> "Yes" "No" 	Yes

Table 41. Attributes for the applicationGroup object (continued)

Name	Description	Data Type	Required	Default value	Possible Values	Updateable
applQueryName	Name of the application that has all the application groups you want to export. When you specify applQueryName during export, the value for the name attribute is ignored. If you specify applQueryName, you cannot specify folderQueryName.	string	No	""		No
folderQueryName	Name of folder that has the application groups you want to export. When you specify folderQueryName during export, the values for the name attribute is ignored. If you specify folderQueryName, you cannot specify applQueryName.	string	No	""		No

Important:

1. The storageSet value can be updated only if it is currently set to "".
2. The default and possible dbTablespaceType values depend on the database that is used.

Remember: The "_ALL" value can only be used in an ARSXML export operation.

Table 42. Child objects for the applicationGroup object

Name	Minimum number	Maximum number	Deleteable	Updateable
permission	0	Unbounded	Yes	Yes
field	1	128	No	Yes
index	0	Unbounded	Yes	No

Table 43. Attributes for the index object under the applicationGroup object

Attributes	Description	Data Type	Required	Default value	Possible values	Updateable
task	Task to perform	taskString	No	"add"	<ul style="list-style-type: none"> • "add" • "delete" 	N/A
name	Name of the index	indexName	Yes		A positive integer	No
cluster	Whether it is a cluster index	yesnoString	No	"No"	<ul style="list-style-type: none"> • "Yes" • "No" 	No

Table 44. Child objects under the index object

Name	Minimum number	Maximum number	Deleteable	Updateable
field	2	16	Yes	No

Table 45. Attributes for the field object under the index object

Attributes	Description	Data Type	Required	Default	Possible values	Updateable
name	Name of the field	string	Yes			No

Table 45. Attributes for the field object under the index object (continued)

Attributes	Description	DataType	Required	Default	Possible values	Updateable
order	The field order	<i>orderString</i>	No	"Ascending"	<ul style="list-style-type: none"> "Ascending" "Descending" 	No

Table 46. Attributes for the permission object under the applicationGroup object

Attributes	Description	DataType	Required	Default	Possible values	Updateable
task	The task to perform on this object	<i>taskString</i>	No	"add"	<ul style="list-style-type: none"> "add" "delete" "update" 	N/A
user	Provide permissions to the application group to this user	<i>nameString</i>	Yes	""	<ul style="list-style-type: none"> User name "*PUBLIC" 	No
group	Provide permissions to the application group to this group	<i>nameString</i>	Yes	""	<ul style="list-style-type: none"> Group name "*PUBLIC" 	No
docAddPerm	Whether this user or group should be allowed to add documents to this application group	<i>yesnoString</i>	No	"No"	<ul style="list-style-type: none"> "Yes" "No" 	Yes
docDeletePerm	Whether this user or group should be allowed to remove documents from this application group	<i>yesnoString</i>	No	"No"	<ul style="list-style-type: none"> "Yes" "No" 	Yes
docUpdatePerm	Whether this user or group should be allowed to update documents in this application group	<i>yesnoString</i>	No	"No"	<ul style="list-style-type: none"> "Yes" "No" 	Yes
docPrintPerm	Whether this user or group should be allowed to print documents from this application group	<i>yesnoString</i>	No	"No"	<ul style="list-style-type: none"> "Yes" "No" 	Yes
docCopyPerm	Whether this user or group should be allowed to copy documents from this application group	<i>yesnoString</i>	No	"No"	<ul style="list-style-type: none"> "Yes" "No" 	Yes
docCFSODPerm	Whether this user or group should be allowed to add CFS-OD documents to this application group	<i>yesnoString</i>	No	"No"	<ul style="list-style-type: none"> "Yes" "No" 	Yes
docCMFedPerm	Whether this user or group should be allowed to add Content Manager documents to this application group	<i>yesnoString</i>	No	"No"	<ul style="list-style-type: none"> "Yes" "No" 	Yes
docFTIPerm	Whether this user or group should be allowed to add full text indexer documents to this application group	<i>yesnoString</i>	No	"No"	<ul style="list-style-type: none"> "Yes" "No" 	Yes
docHoldPerm	Whether this user or group should be allowed to add holds to documents for this application group	<i>yesnoString</i>	No	"No"	<ul style="list-style-type: none"> "Yes" "No" 	Yes

Table 46. Attributes for the permission object under the applicationGroup object (continued)

Attributes	Description	Data Type	Required	Default	Possible values	Updateable
annotViewPerm	Whether this user or group be allowed to view annotation in this application group	<i>yesnoString</i>	No	"No"	<ul style="list-style-type: none"> "Yes" "No" 	Yes
annotAddPerm	Whether this user or group be allowed to add annotation to this application group	<i>yesnoString</i>	No	"No"	<ul style="list-style-type: none"> "Yes" "No" 	Yes
annotDeletePerm	Whether this user or group be allowed to delete annotation from this application group	<i>yesnoString</i>	No	"No"	<ul style="list-style-type: none"> "Yes" "No" 	Yes
annotUpdatePerm	Whether this user or group be allowed to update annotation in this application group	<i>yesnoString</i>	No	"No"	<ul style="list-style-type: none"> "Yes" "No" 	Yes
annotPrintPerm	Whether this user or group be allowed to print annotation from this application group	<i>yesnoString</i>	No	"No"	<ul style="list-style-type: none"> "Yes" "No" 	Yes
annotCopyPerm	Whether this user or group be allowed to copy annotation from this application group	<i>yesnoString</i>	No	"No"	<ul style="list-style-type: none"> "Yes" "No" 	Yes
authority	Specifies the authority given	<i>authagString</i>	No	"Access"	<ul style="list-style-type: none"> "Access" "Administrator" "Logical Views" 	Yes
queryRes	A query restriction for the user or group	<i>queryString</i>	No	""		Yes

Important: Either a user or a group must be specified, but not both.

Table 47. Attributes for the field object under the applicationGroup object

Attributes	Description	Data Type	Required	Default value	Possible values	Updateable
task	The task to perform on the object	<i>taskString</i>	No	Same as the task attribute of the parent	<ul style="list-style-type: none"> "add" "Update" 	N/A
name	Database field name	<i>nameString</i>	Yes			No
type	Index type	<i>fieldTypeString</i>	No	<ul style="list-style-type: none"> "Filter" Index 	<ul style="list-style-type: none"> "Filter" "Index" "Not in Database" 	Yes, but only from "Filter" to "Index" or from "Index" to "Filter"
CFSOD	Whether this field is a CFSOD field	<i>yesnoString</i>	No	"No"	<ul style="list-style-type: none"> "Yes" "No" 	Yes

Table 47. Attributes for the field object under the applicationGroup object (continued)

Attributes	Description	Data Type	Required	Default value	Possible values	Updateable
dataType	The data type of the field	<i>dataTypeString</i>	No	"String"	<ul style="list-style-type: none"> • "Big Int" • "Date" • "Date (native)" • "Date/Time" • "Date/Time (native)" • "Date/Time (TZ)" • "Date/Time (TZ) (native)" • "Decimal" • "Integer" • "Small Int" • "String" • "Time" 	No
segment	If the expiration type is segment, value of this field is used by OnDemand to determine when to delete data from the application group	<i>yesnoString</i>	No	"No"	<ul style="list-style-type: none"> • "Yes" • "No" 	No
expireDate	Whether this field is used for the expiration date	<i>yesnoString</i>	No	"No"	<ul style="list-style-type: none"> • "Yes" • "No" 	No
lockdown	Whether this is a lockdown field	<i>yesnoString</i>	No	"No"	<ul style="list-style-type: none"> • "Yes" • "No" 	Yes
log	Whether to use the system log	<i>yesnoString</i>	No	"No"	<ul style="list-style-type: none"> • "Yes" • "No" 	Yes
userExit	Whether to send data to the user exit	<i>yesnoString</i>	No	"No"	<ul style="list-style-type: none"> • "Yes" • "No" 	Yes
partition	Used to partition the index data across the multiple nodes	<i>yesnoString</i>	No	"No"	<ul style="list-style-type: none"> • "Yes" • "No" 	Yes
cluster	Determines the clustering index for the application group	<i>yesnoString</i>	No	"No"	<ul style="list-style-type: none"> • "Yes" • "No" 	Yes
reference	Only for z/OS. Controls which reports are distributed from OnDemand Distribution Facility.	<i>yesnoString</i>	No	"No"	<ul style="list-style-type: none"> • "Yes" • "No" 	Yes
stringCase	The case to use if the data type is string	<i>caseString</i>	No	"Upper"	<ul style="list-style-type: none"> • "Upper" • "Lower" • "Mixed" 	No
stringType	The type of string	<i>stringType</i>	No	"Fixed"	<ul style="list-style-type: none"> • "Fixed" • "Variable" 	No

Table 47. Attributes for the field object under the applicationGroup object (continued)

Attributes	Description	Data Type	Required	Default value	Possible values	Updateable
stringLength	The length of the string	Integer	No	0	<ul style="list-style-type: none"> • For fixed strings: 1-254 • For variable length strings: 1-2000 	No
stringEmbedded	The embedded characters to remove from the string	string	No	""		Yes
stringLeading	The leading characters to remove from the string	string	No	""		Yes
stringTrailing	The trailing characters to remove from the string	string	No	""		Yes
appIDField	Whether this field is the application ID field	<i>yesnoString</i>	No	"No"	<ul style="list-style-type: none"> • "Yes" • "No" 	Yes
pageCount	Specifies whether this is a page count field	<i>yesnoString</i>	No	"No"	<ul style="list-style-type: none"> • "Yes" • "No" 	Yes
documentSize	Specifies whether this is a document size field	<i>yesnoString</i>	No	"No"	<ul style="list-style-type: none"> • "Yes" • "No" 	Yes
updateable	Whether this field can be updated	<i>yesnoString</i>	No	"Yes"	<ul style="list-style-type: none"> • "Yes" • "No" 	Yes
FTI	Whether to index this field through the Full Text indexer	<i>yesnoString</i>	No	""	<ul style="list-style-type: none"> • "Yes" • "No" 	Yes

The following restrictions apply to the Not in Database value:

- Do not use a field whose type attribute has a definition of Not in Database to hold information such as a report ID that applies to all documents that are returned in the hit list.
- The Not in Database value allows you to create a segment field that is based off the start/end dates that are found in the segment table without capturing another date in the application group data table.
- You can query on a value by using a folder field that is mapped to the application group field whose type attribute has a definition of "Not In Database". The system performs a check to determine whether the value matches the segment range. However, the value cannot be displayed because it is not stored anywhere.
- Use the "Not In Database" value for a database field that has been defined as segment field.
- You can use the "Not In Database" value to segment a report that is based on a value that you do not want to capture in the database as an index or filter. In other words, you only needed the field when you were indexing the report.

Table 48. Child objects under the field object

Name	Minimum number	Maximum number	Deleteable	Updateable
mapping	0	Unbounded	Yes	Yes
nls	0	26	Yes	Yes

Table 49. Attributes for the mapping object under field

Attributes	Description	Data Type	Required	Default value	Possible values	Updateable
task	The task to perform on this object	<i>taskString</i>	No	"add"	<ul style="list-style-type: none"> • "add" • "update" • "delete" 	N/A
dbValue	The database value	<i>dbString2</i>	Yes			No
displayedValue	The displayed value	<i>dvString</i>	Yes			Yes

Table 50. Attributes for the nls object under field

Attributes	Description	Data Type	Required	Default value	Possible values	Updateable
task	The task to perform on this object	<i>taskString</i>	No	"add"	<ul style="list-style-type: none"> • "add" • "update" • "delete" 	N/A
language	The natural language	string	Yes	""	<ul style="list-style-type: none"> • "English" • "French" • "German" • "Italian" • "Japanese" • "Norwegian" • "Portuguese (Brazil)" • "Spanish" • "Chinese (Simplified)" • "Chinese (Traditional)" • "Danish" • "Finnish" • "Swedish" • "Korean" • "Dutch" • "Arabic" • "Czech" • "Greek" • "Hebrew" • "Croatian" • "English" • "Hungarian" • "Polish" • "Russian" • "Slovakian" 	Yes

Table 51. Child objects under the nis object

Name	Minimum number	Maximum number	Deleteable	Updateable
mapping	0	Unbounded	Yes	Yes

Table 52. Attributes for the mapping object under nls

Attributes	Description	Data Type	Required	Default value	Possible values	Updateable
dbValue	The database value	string	Yes			No
displayedValue	The displayed value	string	Yes			Yes

Storage set object

Table 53. Attributes for the storageSet object

Attributes	Description	Data Type	Required	Default value	Possible values	Updateable
name	Name of the storage set	nameString	Yes		<ul style="list-style-type: none"> Storage set name "_ALL" 	Yes
description	A description of the storage set	descString	No	""		Yes
storageType	The load type of the storage set	storageTypeString	No	"Fixed"	<ul style="list-style-type: none"> "Fixed" "Local" 	Yes

Table 54. Child objects under the storageSet object

Name	Minimum number	Maximum number	Deleteable	Updateable
node	1	Unbounded	Yes (unless this is the only object left)	Yes

Table 55. Attributes for the node object under the storageSet object

Attributes	Description	Data Type	Required	Default value	Possible values	Updateable
task	The task to perform on this object	taskString	No	"add"	<ul style="list-style-type: none"> "add" "delete" "update" 	N/A
name	Name of the storage set node	nameString	Yes			No
description	A description of the storage set node	descString	No	""		Yes
server	The server where the node resides	string	No	"*ONDEMAND"		Yes
logon	Client node name	nameString	Dependent	""		Yes
password	Client node password	passwordString	Dependent	""		Yes
loadData	Whether this is a load data node	yesNoString	No	"No"	<ul style="list-style-type: none"> "Yes" "No" 	Yes
accessMethod	Determines the type of archive storage for the primary node	accessString	No	"TSM"	<ul style="list-style-type: none"> "OAM" "VSAM" "TSM" "Cache" 	Yes
configFile	Name of the TSM configuration file	string	No	""		Yes
reloadHoldData	Whether this node should be used to reload hold data	yesnoString	No	"No"	<ul style="list-style-type: none"> "Yes" "No" 	Yes

Remember: The "_ALL" value can only be used in an ARSXML export operation.

The logon and password attributes might be required based on the type of node.

Folder object

Table 56. Attributes for the folder object

Attributes	Description	Data Type	Required	Default value	Possible values	Updateable
name	Name of the folder	<i>nameString</i>	Yes		<ul style="list-style-type: none"> Folder name "_ALL" 	Yes
description	A description of the folder	<i>descString</i>	No	""		Yes
searchType	The note search type	<i>searchTypeString</i>	No	"Retrieve"	<ul style="list-style-type: none"> "Retrieve" "Hit List" "Note" 	Yes
displayDocHold	Whether the client displays the hold status for each document	<i>yesnoString</i>	No	"No"	<ul style="list-style-type: none"> "Yes" "No" 	Yes
displayDocLocation	Whether the client should show the storage location of each document in the document list	<i>yesnoString</i>	No	"No"	<ul style="list-style-type: none"> "Yes" "No" 	Yes
queryUnmappedFields	Whether this user or group should be allowed to add CFS-OD documents to this application group	<i>yesnoString</i>	No	"No"	<ul style="list-style-type: none"> "Yes" "No" 	Yes
applQueryName	Name of the application that has all the folders you want to export. When you specify <i>applQueryName</i> during export, the value for the <i>name</i> attribute is ignored. If you specify <i>applQueryName</i> , you cannot specify <i>agQueryName</i> .	string	No	""		No
agQueryName	Name of the application group that has all the folders you want to export. When you specify <i>agQueryName</i> during export, the value for the <i>identifier</i> attributes is ignored. If you specify <i>agQueryName</i> , you cannot specify <i>applQueryName</i> .	string	No	""		No

Table 57. Child objects under the folder object

Name	Minimum number	Maximum number	Deleteable	Updateable
permission	0	Unbounded	Yes	Yes

Table 57. Child objects under the folder object (continued)

Name	Minimum number	Maximum number	Deleteable	Updateable
applicationGroup	1	Unbounded	Yes	No
field	1	128	No	Yes
nls	0	26	Yes	Yes

Table 58. Attributes for the permission object under the folder object

Attributes	Description	Data Type	Required	Default value	Possible values	Updateable
task	The task to perform on this object	<i>taskString</i>	No	"add"	<ul style="list-style-type: none"> • "add" • "delete" • "update" 	N/A
user	Provide permission to the folder of this user	<i>nameString</i>	Yes	""	<ul style="list-style-type: none"> • User name 	No
group	Provide permission to the folder of this group	<i>nameString</i>	Yes	""	<ul style="list-style-type: none"> • Group name • "*PUBLIC" 	No
adminAuthority	Whether the user or group should have administrator authority	<i>yesnoString</i>	No	"No"	<ul style="list-style-type: none"> • "Yes" • "No" 	Yes
accessAuthority	Whether the user or group should have access authority	<i>yesnoString</i>	No	"No"	<ul style="list-style-type: none"> • "Yes" • "No" 	Yes
fieldsAuthority	Whether the user or group should have authority to change fields	<i>yesnoString</i>	No	"No"	<ul style="list-style-type: none"> • "Yes" • "No" 	Yes
publicNQAuthority	Whether the user or group should have public named query authority	<i>yesnoString</i>	No	"No"	<ul style="list-style-type: none"> • "Yes" • "No" 	Yes
privateNQAuthority	Whether the user or group should have private named query authority	<i>yesnoString</i>	No	"No"	<ul style="list-style-type: none"> • "Yes" • "No" 	Yes
viewNQAuthority	Whether the user or group should have named query view authority	<i>yesnoString</i>	No	"No"	<ul style="list-style-type: none"> • "Yes" • "No" 	Yes
maxHits	The maximum number of hits allowed	<i>maxHitsType</i>	No	"No Limit"	<ul style="list-style-type: none"> • "No Limit" • "None" • A positive integer value 	Yes
secondaryFolder	Whether this is a secondary folder	<i>yesnoString</i>	No	"No"	<ul style="list-style-type: none"> • "Yes" • "No" 	Yes
userGroupFields	Whether user/group fields should be created for this user/group	<i>yesnoString</i>	No	"No"	<ul style="list-style-type: none"> • "Yes" • "No" 	Yes
fullReportBrowse	Whether full report browse should be allowed for this user/group	<i>yesnoString</i>	No	"No"	<ul style="list-style-type: none"> • "Yes" • "No" 	Yes

Table 58. Attributes for the permission object under the folder object (continued)

Attributes	Description	Data Type	Required	Default value	Possible values	Updateable
serverBasedSorting	Whether sorting should be done on the server	<i>yesnoString</i>	No	"No"	<ul style="list-style-type: none"> "Yes" "No" 	Yes

Important: You must specify either a user or a group object, however, you cannot specify both.

Table 59. Attributes for the applicationGroup object under the folder object

Attributes	Description	Data Type	Required	Default value	Possible values	Updateable
task	The task to perform on this object	<i>taskString</i>	No	"add"	<ul style="list-style-type: none"> "add" "delete" 	N/A
name	Name of this object	<i>nameString</i>	Yes		Application group name	No

Table 60. Child objects under the applicationGroup object under the folder object

Name	Minimum number	Maximum number	Deleteable	Updateable
application	0	Unbounded	Yes	No

Table 61. Attributes for the application object under the applicationGroup object

Attributes	Description	Data Type	Required	Default value	Possible values	Updateable
task	The task to perform on this object	<i>taskString</i>	No	"add"	<ul style="list-style-type: none"> "add" "delete" 	N/A
name	Name of this object	<i>nameString</i>	Yes		Application name	No

Table 62. Attributes for the field object under the folder object

Attributes	Description	Data Type	Required	Default value	Possible values	Updateable
task	The task to perform on this object	<i>taskString</i>	No	"add"	<ul style="list-style-type: none"> "add" "delete" "update" 	N/A
name	Name of the folder field	<i>nameString</i>	Yes			Yes
description	A description of the folder field	<i>descString</i>	No	""		Yes

Table 62. Attributes for the field object under the folder object (continued)

Attributes	Description	Data Type	Required	Default value	Possible values	Updateable
fieldType	Field type of the folder object	<i>fieldString</i>	No	"String"	<ul style="list-style-type: none"> • "Ann Color Search" • "Ann Text Search" • "Application Group" • "Big Int" • "Date" • "Date (native)" • "Date/Time" • "Date/Time (native)" • "Date/Time (TZ)" • "Date/Time (TZ) (native)" • "DecFloat (16)" • "DecFloat (34)" • "Decimal" • "Document Type" • "FTI Highlight" • "FTI Score" • "FTI Search" • "FTI Summary" • "Hold" • "Integer" • "Segment" • "Small Int" • "String" • "Text Search" • "Time" 	No
mappingType	Mapping type of the folder field	<i>mappingString</i>	No	"Single"	<ul style="list-style-type: none"> • "Single" • "Range" • "Operator Or" 	No
applicationGroup	The Application Group to use for a segment field	<i>nameString</i>	No	""	""	Yes

Table 63. Child objects under the field object under the folder object

Name	Minimum number	Maximum number	Deleteable	Updateable
mapping	0	Unbounded	Yes	No
fieldInfo	1	Unbounded	Yes	Yes

Table 64. Attributes for the mapping object under the field object

Attributes	Description	Data Type	Required	Default value	Possible values	Updateable
task	The task to perform on this object	<i>taskString</i>	No	"add"	<ul style="list-style-type: none"> • "add" • "delete" 	N/A
dbName	The database field name	<i>nameString</i>	Yes			No
appGroup	The application group name	<i>nameString</i>	Yes			No

Table 65. Attributes for the fieldInfo object under the field object

Attributes	Description	Data Type	Required	Default value	Possible values	Updateable
task	The task to perform on this object	<i>taskString</i>	No	"add"	<ul style="list-style-type: none"> "add" "delete" "update" 	N/A
user	The user that is associated with this field definition	<i>nameString</i>	Yes		User name	Yes
group	The group that is associated with this field definition	<i>nameString</i>	Yes		<ul style="list-style-type: none"> Group name "*PUBLIC" 	Yes
defaultValue	The default value	string	No			Yes
defaultValue2	The second default value	string	No			Yes
displayOrder	The order of this field within the hit list	Integer	No	1	0 — Number of fields	Yes
queryOrder	The order of this field during a query	Integer	No	1	0 — Number of fields	Yes
sortOrder	The sort position of this field, if any	Integer	No	0	0 — Number of fields	Yes
sortType	Whether the documents should be sorted in an ascending or descending order	<i>sortTypeString</i>	No	"Ascending"	<ul style="list-style-type: none"> "Ascending" "Descending" 	Yes
equal	Whether the equal operator should be available	<i>compString</i>	No	"Default"	<ul style="list-style-type: none"> "Yes" "No" "Default" 	Yes
notEqual	Whether the not equal operator should be available	<i>compString</i>	No	"No"	<ul style="list-style-type: none"> "Yes" "No" "Default" 	Yes
lessThan	Whether the less than operator should be available	<i>compString</i>	No	"No"	<ul style="list-style-type: none"> "Yes" "No" "Default" 	Yes
ltOrEqual	Whether the less than or equal operator should be available	<i>compString</i>	No	"No"	<ul style="list-style-type: none"> "Yes" "No" "Default" 	Yes
greaterThan	Whether the greater than operator should be available	<i>compString</i>	No	"No"	<ul style="list-style-type: none"> "Yes" "No" "Default" 	Yes
gtOrEqual	Whether the greater than or equal operator should be available	<i>compString</i>	No	"No"	<ul style="list-style-type: none"> "Yes" "No" "Default" 	Yes
in	Whether the in operator should be available	<i>compString</i>	No	"No"	<ul style="list-style-type: none"> "Yes" "No" "Default" 	Yes

Table 65. Attributes for the fieldInfo object under the field object (continued)

Attributes	Description	DataType	Required	Default value	Possible values	Updateable
notIn	Whether the not in operator should be available	<i>compString</i>	No	"No"	<ul style="list-style-type: none"> • "Yes" • "No" • "Default" 	Yes
between	Whether the between operator should be available	<i>compString</i>	No	"No"	<ul style="list-style-type: none"> • "Yes" • "No" • "Default" 	Yes
notBetween	Whether the not between operator should be available	<i>compString</i>	No	"No"	<ul style="list-style-type: none"> • "Yes" • "No" • "Default" 	Yes
like	Whether the like operator should be available	<i>compString</i>	No	"No"	<ul style="list-style-type: none"> • "Yes" • "No" • "Default" 	Yes
notLike	Whether the not like operator should be available	<i>compString</i>	No	"No"	<ul style="list-style-type: none"> • "Yes" • "No" • "Default" 	Yes
default	Determines whether the field contains a default value when the user opens the folder	<i>yesnoString</i>	No	"No"	<ul style="list-style-type: none"> • "Yes" • "No" 	Yes
fixed	Determines whether the value that you type in the default fields can be changed by the user	<i>compString</i>	No	"No"	<ul style="list-style-type: none"> • "Yes" • "No" 	Yes
required	Determines whether the user must specify a search value for the field in order to type a query	<i>yesnoString</i>	No	"No"	<ul style="list-style-type: none"> • "Yes" • "No" 	Yes
viewTitle	Determines whether OnDemand client programs display the field name on the title bar of the viewing window when the user selects a document for viewing	<i>yesnoString</i>	No	"No"	<ul style="list-style-type: none"> • "Yes" • "No" 	Yes
internal	Determines whether OnDemand lists the displayed values or the database values for the folder search field	<i>yesnoString</i>	No	"No"	<ul style="list-style-type: none"> • "Yes" • "No" 	Yes

Table 65. Attributes for the fieldInfo object under the field object (continued)

Attributes	Description	Data Type	Required	Default value	Possible values	Updateable
wildCard	Determines whether OnDemand uses wildcards with field search values	<i>wildString</i>	No	"Append"	<ul style="list-style-type: none"> • "Prepend" • "Append" • "Both" • "None" 	Yes
min	For integer, DecFloat(16), DecFloat(34), decimal, date, and time fields, determines the minimum value that the user can type in the search field	Integer	No	Based on datatype		Yes
max	For integer, DecFloat(16), DecFloat(34), decimal, date, and time fields, determines the maximum value that the user can type in the search field	Integer	No	Based on datatype		Yes
decimalPrecision	Specifies the decimal precision for decimal values	Integer	No	2	Nonnegative integer	Yes
dateDisplayFormat	Determines the format that is used to display dates and times	string	No	""	Date or time format	Yes
dateDefaultFormat	Determines the format that OnDemand uses to validate date and time values	string	No	""	Date or time format	Yes
dateInterval	For date and time fields, determines whether the interval refers to the next or last interval	<i>intervalString</i>	No	"Last"	<ul style="list-style-type: none"> • "Last" • "Next" 	Yes
dateIntLength	For date and time fields, determines the number of units that OnDemand uses to compute a default search range	integer	No	0	Nonnegative integer	Yes
dateIntType	Units of measure for dateIntLength	<i>intTypeString</i>	No	"Days"	<ul style="list-style-type: none"> • "Days" • "Months" • "Years" • "Hours" • "Minutes" • "Seconds" 	Yes

Table 66. Attributes for the nls object under the folder object

Attributes	Description	Data Type	Required	Default value	Possible values	Updateable
task	The task to perform on this object.	<i>taskString</i>	No	"add"	<ul style="list-style-type: none"> • "add" • "update" 	No
language	The language employed by this object.	<i>langString</i>	Yes		<ul style="list-style-type: none"> • "English" • "French" • "German" • "Italian" • "Japanese" • "Norwegian" • "Portuguese (Brazil)" • "Spanish" • "Chinese (Simplified)" • "Chinese (Traditional)" • "Danish" • "Finnish" • "Swedish" • "Korean" • "Dutch" • "Arabic" • "Czech" • "Greek" • "Hebrew" • "Croatian" • "Hungarian" • "Polish" • "Russian" • "Slovakian" • "Slovenian" 	No

Table 67. Child objects under the nls object under the folder object

Name	Minimum number	Maximum number	Deleteable	Updateable
field	0	128	No	Yes

Table 68. Attributes for the field object under the nls object

Attributes	Description	Data Type	Required	Default value	Possible values	Updateable
name	Name of the folder field.	<i>stringLeng1-60</i>	Yes			No
langName	Name of the folder field written in the NLS language.	<i>stringLeng1-60</i>	Yes			Yes
langDescription	Description of the folder field written in the NLS language.	<i>stringLeng1-120</i>	No			Yes

Note:

1. You must specify either a user or a group, however, you cannot specify both.
2. The default values for displayOrder and queryOrder are based on their position within the file.

Remember: The "_ALL" value can only be used in an ARSXML export operation.

Printer object

Table 69. Attributes for the printer object

Attributes	Description	Data Type	Required	Default value	Possible values	Updateable
name	Name of the printer	<i>nameString</i>	Yes		<ul style="list-style-type: none"> Printer name "_ALL" 	Yes
description	A description of the printer	<i>descString</i>	No	""		Yes
queue	The printer queue	<i>queueString</i>	Yes			Yes
printerType	The type of printer	<i>printerTypeString</i>	No	"Printer"	<ul style="list-style-type: none"> "Printer" "Fax" "Printer With Information" 	Yes

Table 70. Child objects under the printer object

Name	Minimum number	Maximum number	Deleteable	Updateable
permission	0	Unbounded	Yes	No

Table 71. Attributes for the permission object under the printer object

Attributes	Description	Data Type	Required	Default value	Possible values	Updateable
task	The task to perform on this object	<i>taskString</i>	No	"add"	<ul style="list-style-type: none"> • "add" • "delete" 	N/A
user	The user to give permission to	<i>nameString</i>	Yes	""	<ul style="list-style-type: none"> • User name 	No
group	The group to give permission to	<i>nameString</i>	Yes	""	<ul style="list-style-type: none"> • Group name • "*PUBLIC" 	No

Remember: The "_ALL" value can only be used in an ARSXML export operation.

Cabinet object

Table 72. Attributes for the cabinet object

Attributes	Description	Data Type	Required	Default value	Possible values	Updateable
name	Name of the cabinet	<i>nameString</i>	Yes		<ul style="list-style-type: none"> Cabinet name "_ALL" 	Yes
description	A description of the cabinet	<i>descString</i>	No	""		Yes

Table 73. Child objects under the cabinet object

Name	Minimum number	Maximum number	Deleteable	Updateable
folder	0	Unbounded	Yes	No
permission	0	Unbounded	Yes	Yes

Table 74. Attributes for the folder object under the cabinet object

Attributes	Description	Data Type	Required	Default value	Possible values	Updateable
name	Name of the folder	<i>nameString</i>	Yes			No
task	The task to perform on this object	<i>taskString</i>	No	"add"	<ul style="list-style-type: none"> • "add" • "delete" 	N/A

Table 75. Attributes for the permission object under the cabinet object

Attributes	Description	Data Type	Required	Default value	Possible values	Updateable
user	Name of the user	<i>nameString</i>	No		<ul style="list-style-type: none"> User name 	No
group	Name of the group	<i>nameString</i>	No		<ul style="list-style-type: none"> Group name "*PUBLIC" 	No
task	The task to perform on this object	<i>taskString</i>	No	"add"	<ul style="list-style-type: none"> "add" "delete" "update" 	N/A
accessAuthority	Whether this user or group has access authority	<i>yesnoString</i>	No	"No"	<ul style="list-style-type: none"> "Yes" "No" 	Yes
adminAuthority	Whether this user or group has administrative authority	<i>yesnoString</i>	No	"No"	<ul style="list-style-type: none"> "Yes" "No" 	Yes

Remember: The "_ALL" value can only be used in an ARSXML export operation.

Hold object

Table 76. Attributes for the hold object

Attributes	Description	Data Type	Required	Default value	Possible values	Updateable
name	Name of the hold	<i>nameString</i>	Yes		<ul style="list-style-type: none"> • Hold name • "_ALL" 	Yes
description	A description of the hold	<i>descString</i>	No	""		Yes

Table 77. Child objects under the hold object

Name	Minimum number	Maximum number	Deleteable	Updateable
permission	0	Unbounded	Yes	Yes

Table 78. Attributes for the permission object under the hold object

Attributes	Description	Data Type	Required	Default value	Possible values	Updateable
user	Name of a user that can view and maintain this hold	<i>nameString</i>	Yes	""	<ul style="list-style-type: none"> User name 	No
group	Name of the group that can view and maintain this hold	<i>nameString</i>	Yes	""	<ul style="list-style-type: none"> Group name 	No
task	The task to perform on this object	<i>taskString</i>	No	"add"	<ul style="list-style-type: none"> "add" "delete" "update" 	Yes
adminAuthority	Whether this user or group has administrative authority	<i>yesnoString</i>	No	"No"	<ul style="list-style-type: none"> "Yes" "No" 	Yes
accessAuthority	Whether this user or group has access authority	<i>yesnoString</i>	No	"No"	<ul style="list-style-type: none"> "Yes" "No" 	Yes

Remember: The "_ALL" value can only be used in an ARSXML export operation.

System Parameters object

Table 79. Attributes for the systemParameters object

Attributes	Description	Data Type	Required	Default value	Possible values	Updateable
allowPrivateUser	Allows the user to add private annotations to a document.	<i>yesnoString</i>	No	""	<ul style="list-style-type: none"> "No" "Yes" 	Yes
allowPublic	Allows the user to add public annotations.	<i>yesnoString</i>	No	""	<ul style="list-style-type: none"> "No" "Yes" 	Yes
allowGraphicAnnot	Allow the user to add graphic annotations.	<i>yesnoString</i>	No	""	<ul style="list-style-type: none"> "No" "Yes" 	Yes
allowTextAnnot	Allow the user to add text annotations.	<i>yesnoString</i>	No	""	<ul style="list-style-type: none"> "No" "Yes" 	Yes
defaultAnnot	Specify the default viewing scope for all annotations.	<i>String</i>	No	""	<ul style="list-style-type: none"> "Public" "Private to User" "Private to Group" 	
disableOrLockUser	Specify whether to disable or lock out a user after the user attempts and fails to log in for the specified number of times.	<i>String</i>	No	""	<ul style="list-style-type: none"> "Never" "Disable User" "Lock Out User" 	Yes
enableComments	Display the System Log Comments window when you perform an add, update, or delete operation.	<i>yesnoString</i>	No	""	<ul style="list-style-type: none"> "No" "Yes" 	
enableLDAP	Specify whether you want to use LDAP (Lightweight Directory Access Protocol) authentication in your OnDemand server.	<i>yesnoString</i>	No	""	<ul style="list-style-type: none"> "No" "Yes" 	
failedLogins	Specify the number of times a user can attempt to log in.	<i>integer</i>	No	""	An integer between 1 and 10	Yes

Table 79. Attributes for the systemParameters object (continued)

Attributes	Description	Data Type	Required	Default value	Possible values	Updateable
inactivityTO	Specify the number of minutes between server requests that must pass, after which OnDemand notifies user to log off.	<i>integer</i>	No		An integer between 1 and 1440.	
lifeOfAnnotations	Specify the number of days OnDemand keeps annotations	<i>integer or string</i>	No	""	<ul style="list-style-type: none"> • "Never expire" • An integer between 1 and 365 	Yes
lockoutMinutes	Specify the number of minutes to prevent a user from attempting to log in.	<i>integer</i>	No	""	An integer between 1 and 1440	Yes
logAGMsgs	Specify whether OnDemand saves a message in the system log when a user queries or retrieves application group data and other types of application group events.	<i>yesnoString</i>	No	""	<ul style="list-style-type: none"> • "No" • "Yes" 	
logFailedLogin	Specify whether OnDemand saves a message in the system log when there is an unsuccessful log on attempt.	<i>yesnoString</i>	No	""	<ul style="list-style-type: none"> • "No" • "Yes" 	
logLogin	Specify whether OnDemand saves a message in the system log when a user logs on the server.	<i>yesnoString</i>	No	""	<ul style="list-style-type: none"> • "No" • "Yes" 	
logLogoff	Specify whether OnDemand saves a message in the system log when a user logs off the server.	<i>yesnoString</i>	No	""	<ul style="list-style-type: none"> • "No" • "Yes" 	

Table 79. Attributes for the systemParameters object (continued)

Attributes	Description	Data Type	Required	Default value	Possible values	Updateable
logUEAGMsgs	Specify whether the OnDemand user exit saves a message in the system log when a user queries or retrieves application group data and other types of application group events.	<i>yesnoString</i>	No	""	<ul style="list-style-type: none"> • "No" • "Yes" 	
logUEFailedLogin	Specify whether the OnDemand user exit saves a message in the system log when there is an unsuccessful log on attempt.	<i>yesnoString</i>	No	""	<ul style="list-style-type: none"> • "No" • "Yes" 	
logUELogin	Specify whether the OnDemand user exit saves a message in the system log when a user logs on the server.	<i>yesnoString</i>	No	""	<ul style="list-style-type: none"> • "No" • "Yes" 	
logUELogoff	Specify whether the OnDemand user exit saves a message in the system log when a user logs off the server.	<i>yesnoString</i>	No	""	<ul style="list-style-type: none"> • "No" • "Yes" 	
notifyPasswordDays	Notify users that their password will expire within the specified number of days.	<i>integer or string</i>	No	""	<ul style="list-style-type: none"> • "Never notify" • An integer from 1 - 30 	Yes
passwordAge	Specify the number of days after which a password must be changed.	<i>integer or string</i>	No	""	<ul style="list-style-type: none"> • "Never Expires" • "Always Expires" • An integer from 1 - 365 	
passwordLen	Specify whether a password is required. If a password is required, specify the minimum length of a password.	<i>integer or string</i>	No	""	<ul style="list-style-type: none"> • "Permit Blank" • An integer from 1 - 128 	

Table 79. Attributes for the systemParameters object (continued)

Attributes	Description	Data Type	Required	Default value	Possible values	Updateable
previousPasswords	Specify whether users can reuse a password.	<i>integer or string</i>	No	""	<ul style="list-style-type: none"> "Never check" An integer from 1 - 10 	Yes
pwdCaseSensitive	Specify whether a password must be case sensitive.	<i>yesnoString</i>	No	""	<ul style="list-style-type: none"> "No" "Yes" 	
requireComments	Requires you to enter one or more characters in the Comments field when you perform an add, update, or delete operation.	<i>yesnoString</i>	No	""	<ul style="list-style-type: none"> "No" "Yes" 	
serverTraceOn	Enable or disable tracing for the OnDemand system.	<i>yesnoString</i>	Yes	"No"	<ul style="list-style-type: none"> "No" "Yes" 	Yes
serverTraceString	Specify parameters as instructed by IBM Software Support.	<i>String</i>	Yes	""		Yes
uidCaseSensitive	Specify whether the user ID must be case sensitive.	<i>yesnoString</i>	No	""	<ul style="list-style-type: none"> "No" "Yes" 	
updateDate		<i>String</i>	No	""		
userLoginInactivity	Specify whether to disable users that do not login after the specified number of days.	<i>integer or string</i>	No	""	<ul style="list-style-type: none"> "Never disable" An integer between 1 and 2147483647 	Yes

Data field limitations

Table 80. Data field limitations

Datatype	Minimum length	Maximum length
dbString	1	18
dbString2	1	254
defString	0	254
descString	0	120
dvString	1	254
emailString	0	254
fileExtString	1	20
formString	1	20
indexParm	0	32,700

Table 80. Data field limitations (continued)

Datatype	Minimum length	Maximum length
nameString	1	60
nameStringUG (Users and Groups)	1	128
passwordString	0	128
phoneString	0	32
postProcString	0	1000
queueString	1	60
routingString	0	24
valString	0	8
userMiscString	0	60

Command reference

This part contains reference information about the IBM Content Manager OnDemand server commands. The commands are presented in alphabetical order. Each command contains a description of its purpose. See the online help text for more details about the commands and command parameters.

Content Manager OnDemand server commands

These commands require that you be signed on to the server with a user profile that is also defined as a user in IBM Content Manager OnDemand. For commands that specify an application or application group name, if the value to be entered contains lower case letters, blanks, or special characters, it must be enclosed in apostrophes.

Many of these commands include an INSTANCE parameter which specifies the OnDemand instance name for which you are running the command. By default, the QUSROND default instance is used, and will produce the desired results for most systems. You can use an instance other than QUSROND as your default by defining the QDFTINST data area. You can also specify the instance name directly when you run the commands.

ADDRPTOND

The Add Report to OnDemand (ADDRPTOND) command allows you to load reports in IBM Content Manager OnDemand. During this process the report is broken into segments, indexed, compressed, and stored on disk for retrieval and later migration to optical or tape media if desired. Input can be in the form of a spooled file (*SPLF), a database file (*FILE), or a stream file (*STMF).

Important: Set the proper locale before issuing this command. See the chapter entitled "Defining a locale" in the *IBM Content Manager OnDemand for i Common Server Planning and Installation Guide* for important details.

After the ADDRPTOND command has run, message number 87 will appear in the System Log if the loading of data was successful; message 88 will appear if the data loading failed.

CHGPLDOND

| The Change Policy Level Date (CHGPLDOND) command changes the next level
| date to the new date for objects that are at the named level in the named migration
| policy. Only objects in the specified date range or for the specified report ID (also
| known as load ID) are changed.

CRTINSTOND

The Create Instance for OnDemand (CRTINSTOND) command creates an OnDemand instance.

An OnDemand instance is a logical server environment with its own library containing a unique set of database files. An instance is defined in the ARS.INI file by naming the instance (which identifies the name of the library used by the

instance). All of the database files that belong to an instance run in one and only one coded character set identifier (CCSID).

When you are creating an instance, your user profile must have its locale set to the locale of the instance you wish to create. Because the locale is set in the user profile, you may need to change your user profile, then sign off and back on before creating the instance. Use the Change User Profile (CHGUSRPRF) command to change (if necessary) your user profile. You should also make sure that other language-related parameters in your user profile are set correctly. You can use the Display User Profile (DSPUSRPRF) command to check the locale setting. The locale Job Attributes (SETJOBATR) parameter in your user profile is used to determine which values are obtained from the locale. For OnDemand, at a minimum, you must use SETJOBATR(*CCSID).

ENDMONOND

The End Monitor for OnDemand (ENDMONOND) command allows you to end a currently active monitor for an output queue or directory. Note that the *DIR monitor type on the ENDMONOND command will end either a *DIR or *DIR2 type of monitor, regardless of whether it was started as a *DIR or *DIR2 monitor.

FNDKEYOND

The Find Key for the OnDemand (FNDKEYOND) command allows you to search for a particular document that is available through the folder specified in the FOLDER parameter, and starts the IBM Content Manager OnDemand client to display the results of the search. The key fields entered must exist in the folder. This is intended as an API to start the Content Manager OnDemand client from an application running in a 5250 emulation session. See "5250 host connection to client viewer" on page 251 for more details.

MGRMEDRDAR

The Migrate Media (MGRMEDRDAR) command provides a tool to move OnDemand Spool File Archive data from one media type to another in an easy, recoverable way that can be stopped and restarted as needed. When the MGRMEDRDAR command runs, it updates all the necessary Spool File Archive files to point to the new location. The source media for the command can be an optical volume, a tape volume, or an individual report name. The target media can be disk for all types of source media, optical if the source media is optical, or *ASM (Archive Storage Manager) if the source media is disk and contains reports that have been migrated from the Spool File Archive environment to the Common Server environment.

A long running MGRMEDRDAR job can be ended if necessary. However, it must be ended in a controlled manner to prevent unexpected results. You can use the End Job (ENDJOB) command, specifying JOB(job-number/user-name/job-name) OPTION(*CNTRLD) DELAY(500) where job-number/user-name/job-name identifies the MGRMEDRDAR job.

Important: The MGRMEDRDAR command is included in this Common Server publication because the Spool File Archive data can be migrated to Common Server. For this reason, users of Common Server implementations might still be interested in using the MGRMEDRDAR command to move data that was originally stored in Spool File Archive from one medium to another. For detailed information on this media migration facility, refer to the IBM Content Manager

OnDemand for i support Web page at <http://www.ibm.com/software/data/ondemand/400/support.html> and search for MMF.

MRGSPLFOND

The Merge Spooled Files (MRGSPLFOND) command combines multiple SNA character stream (SCS) spooled files and writes the result to a single spooled file or database file member. The command can also store the combined file, if specified. Spooled files to be merged must be contained in a single output queue and must be in Ready (RDY) status. Spooled files that are not in a Ready (RDY) status and that are not SCS will be left in the source output queue.

The purpose of this command is to improve the archive performance of small SCS spooled files. To archive many small spooled files takes longer and uses more system resources than to archive one large spooled file.

PRTRPTOND

The Print Report from OnDemand (PRTRPTOND) command prints the specified report in its entirety. In version 7.2, the format of the report ID (also known as the load ID) field that specifies which report to print has changed. See online help for the PRTRPTOND command for details.

If *OUTQ is specified for the PRINTER parameter, you must also enter an output queue name and library name. An output queue name without a library qualifier will default to *LIBL for the library when PRINTER(*OUTQ) is specified. The instance server job's library list is then used to locate the output queue.

If a printer name is specified for the PRINTER parameter, and the Content Manager OnDemand printer definition names an output queue that has not been library qualified, the print will fail if the output queue is not found in the instance server job's library list.

Restriction: The COPIES and PAGERANGE parameters were removed from the PRTRPTOND command in Version 6 Release 1, and should be removed from any CL programs or job scheduler entries that might currently specify them.

PRTTXTOND

The Print Text for OnDemand (PRTTXTOND) command allows you to print all or part of a spooled file in a text-only format. This allows the report administrator to see what the report looks like to ADDRPTOND when it is indexed. Using this output (spooled file name QPRLMTXT), the administrator can determine how to index the report.

RMVRPTOND

The Remove Report from OnDemand (RMVRPTOND) command removes the specified report from IBM Content Manager OnDemand. In version 7.2, the format of the report ID (also known as the load ID) field that specifies which report to remove has changed. See online help for the RMVRPTOND command for details.

Restriction: The ERROPT parameter was removed from the RMVRPTOND command in version 7.2, and should be removed from any CL programs or job scheduler entries that might currently specify it. In previous releases, the ERROPT parameter permitted you to specify whether or not you wanted to remove a report based on the existence of a record of the report ID (load ID) in the Content

Manager OnDemand System Log. Beginning at version 7.2, the RMVRPTOND command always removes the report, regardless of whether or not the report ID is found in the System Log.

STRASMOND

The Start Archived Storage Mgmt (STRASMOND) command starts the Archived Storage Management (ASM) process which manages the movement of data within the ASM defined levels. This function can be canceled with a controlled cancel option if enough time is allowed to end what it is currently processing.

Important:

1. This command must only be run in batch (SBMJOB parameter set to *YES). Running this command interactively (with SBJOB(*NO)) may cause SQL errors.
2. By default, the QUSROND default instance is used, and will produce the desired results for most systems. If you need to run the STRASMOND command for multiple instances, you must issue the command separately for each instance. Note that if you initiate the archive storage manager by running the STRDSMOND command with RUNASM(*YES), then the instance name is passed from the disk storage manager and no further specifications are needed.
3. The ASM process produces multiple reports that list the actions that the ASM process performed. The reports are produced with spooled file names beginning with QPRLCASM, such as QPRLCASM, QPRLCASMA, QPRLCASMD, or QPRLCASMU. The process also loads the same reports into the Content Manager OnDemand System Log, automatically starting the Content Manager OnDemand server for you if the server is not started so that the reports can be loaded. These reports should be checked each time ASM is run to ensure that processing of the data completed successfully. If a failed condition is found, it is important to check the job log for the STRASMOND job to determine the cause of the failure. The STRASMOND command actually runs three sub-functions: one to process all previously unprocessed objects, one to process all aggregates with next level date less than or equal to today's date, and one to process all objects not in an aggregate with next level date less than or equal to today's date. You might see message RDR2798 one or more times in your STRASMOND job log if any of these three sub-functions had no objects to process.
4. The LOGSTS and SNDFAILMSG parameters have been removed from the STRASMOND command at Version 6 Release 1 and should be removed from any CL programs or job scheduler entries that may currently specify them.

STRDSMOND

The Start Disk Storage Management (STRDSMOND) command starts the Disk Storage Management (DSM) task which manages the movement of OnDemand data on disk and between disk and the Archived Storage Manager (ASM) or Tivoli Storage Manager (TSM). DSM also controls the expiration of data in OnDemand.

Although you can run multiple STRDSMOND commands for different application groups within the same instance or different instances, it is not recommended.

This process can be canceled if absolutely necessary. Note that spawned jobs that are part of DSM processing may continue to run even after you cancel the original DSM job. You may also receive many SQL messages in the job log of the instance server job. Also be aware that when the DSM job is canceled, the report that

provides information about the data that has been processed during the DSM job (system log message number 197) will not be created in the system log as is done when DSM runs to a normal completion.

Important:

1. If you run STRDSMOND for a specific application group (rather than the default of *ALL) and you set the Run ASM (RUNASM) parameter to *YES, be aware that ASM will run for ALL application groups, even though you have named a specific application group for DSM to use. You can, however, name a specific Policy for ASM to process, if desired.
2. When you specify RUNASM(*YES), Content Manager OnDemand will initiate a separate batch job for ASM.
3. When you specify RUNASM(*YES), the ASM process produces multiple reports that list the actions that the ASM process performed. The reports are produced with spooled file names beginning with QPRLCASM, such as QPRLCASM, QPRLCASMA, QPRLCASMD, or QPRLCASMU. The process also loads the same reports into the Content Manager OnDemand System Log, automatically starting the Content Manager OnDemand server for you if the server is not started so that the reports can be loaded. These reports should be checked each time ASM is run to ensure that processing of the data completed successfully. If a failed condition is found, it is important to check the job log for the STRASMOND job to determine the cause of the failure. The STRASMOND command actually runs three sub-functions: one to process all previously unprocessed objects, one to process all aggregates with next level date less than or equal to today's date, and one to process all objects not in an aggregate with next level date less than or equal to today's date. You might see message RDR2798 one or more times in your STRASMOND job log if any of these three sub-functions had no objects to process.

STRIMPOND

The Start Import into OnDemand (STRIMPOND) command allows you to import data into OnDemand. This command is used only if you have migrated your index data to an alternate media (such as optical or tape), which is not recommended, but may be necessary in some cases.

STRMONOND

The Start Monitor for OnDemand (STRMONOND) command allows you to specify the name of an output queue or IFS directory to monitor. When a spooled file is added to the output queue or a file is added to an IFS directory, the file is automatically processed by ADDRPTOND if it meets certain criteria as defined in the online help for the STRMONOND command. For example, spooled files must be in Ready (RDY) status to be processed. Files added to IFS directories must end in either a .IND extension (for *DIR monitor type) or a .ARD extension (for *DIR2 monitor type).

Important: Set the proper locale before issuing this command. See the chapter entitled "Defining a locale" in the *IBM Content Manager OnDemand for i Common Server: Planning and Installation Guide* for important details.

The application and application group names for spooled files are determined using spooled file attributes such as spooled file name or user data (or a number of others). You can alter these values that Content Manager OnDemand uses for application and application group by using the monitor user exit program as described in "API and user exit reference" on page 259.

The application and application group names for workstation files are determined by the file name itself. The monitor can use any one of the first four parts of the file name. For example, a file name of CHECKSTMTS.BIGSTMTS.data might be used to archive an application named CHECKSTMTS into an application group named BIGSTMTS.

After the STRMONOND command has processed an input file, message number 87 will appear in the System Log if the loading of data was successful; message 88 will appear if the data loading failed.

You can end the monitor by:

- Specifying a method for the monitor to end automatically.
- Running the ENDMONOND command.
- Ending the monitor job using the ENDJOB command. Specify OPTION(*CNTRLD) and DELAY(999999). The job will end as soon as the monitor finishes processing the current file.

System log messages

Overview

For a complete list of System Log messages, see *IBM Content Manager OnDemand Messages and Codes*, SC27-1379.

You should always keep the system logs. (See the application group and folder titled System Log.) For example, the System Log is an easy place to locate the Load Identifier for archived data. (You can also display a partial Load Identifier using the OnDemand client while viewing archived data.) Every time data is loaded into IBM Content Manager OnDemand, message number 87 is placed in the System Log and the Load Identifier is recorded as part of the message. The Load Identifier is also called the Report ID and is a required parameter (RPTID) for the Remove Report from OnDemand (RMVRPTOND) and Print Report from OnDemand (PRTRPTOND) commands. Beginning with version 7.2, the format of the Load Identifier changed. See online help for the RMVRPTOND or RPTRPTOND command for details about the new format. The format you use for the RPTID parameter of the RMVRPTOND and PRTRPTOND commands depends on the format of the Load Identifier that was generated when the data was loaded. If the Load Identifier that was generated when the data was loaded is in the old format, then you use the old Load Identifier format for RPTID. If the Load Identifier that was generated when the data was loaded is in the new format, then you use the new Load Identifier format.

When installed, the System Log application group is defined to never expire and IBM recommends that you do not change that setting. If you do change the setting, you should only expire the System Log data after all other application groups have expired their data. In other words, the value in the **Expire in x Days** field under Life of Data and Indexes on the Storage Management tab of the System Log application group should always be larger than the same value in any other application group. If any other application group is using the **Never Expire** setting, you should not change the System Log application group setting from the default.

5250 host connection to client viewer

This section describes the IBM Content Manager OnDemand 5250 Host Connection, which allows an IBM i application to send information to the Content Manager OnDemand client workstation viewer program (the viewer). The intent is to provide a mechanism by which a 5250 line-of-business application can use the viewer to display documents from the OnDemand database. This is done with little or no interaction between the user and the viewer.

This function is comprised of two parts:

- The FNDKEYOND command, which the line-of-business application will invoke. The FNDKEYOND command is a part of Option 10 (Common Server) of Content Manager OnDemand for i, and is used to retrieve documents archived using the Content Manager OnDemand Common Server feature. The output of the FNDKEYOND command will be displayed using the workstation viewer.
- A companion program that resides on the workstation, which is the interface between the FNDKEYOND command and the viewer.

Operational and environmental considerations

For this function to operate correctly, you must address several operational and environmental considerations. This section will describe these considerations.

Workstation installation tasks

For the FNDKEYOND command to operate correctly, the workstation companion program, QRLROCD.EXE, must be running on the workstation when the FNDKEYOND command is run on the server. This program is found in the OnDemand install directory, usually C:\Program Files\IBM\OnDemand32. The following paragraphs describe a suggested approach for simplifying the startup of this program.

After the OnDemand client software is installed on your workstation, you should establish a shortcut which points to the QRLROCD.EXE program. This shortcut can be placed on the workstation desktop, in the startup folder or in any other place which meets your needs. Placing the shortcut in your startup folder will start the program automatically when you start your workstation. This is appropriate if you use this interface frequently. If you end the program and need to restart it, or if you have no need for the program to be started automatically, you may want to place a shortcut on the desktop.

The command line parameters described in the Content Manager OnDemand Windows Client Customization Guide apply only to the OnDemand end-user client program named ARSGUI32.EXE. They are not supported by the QRLROCD.EXE program. The QRLROCD.EXE program supports two command line parameters: /p to specify port number, and /s to specify system name. Both can be used when the QRLROCD.EXE program is required to connect to multiple OnDemand instance servers from a single workstation as described in the section labeled Multiple OnDemand servers in the network.

In most cases, you will only need to start the companion program once. If, however, the workstation is attached to multiple OnDemand server systems, you may need to start the program multiple times. In this case, you should create a

shortcut for each one, for reasons explained in the section labeled Multiple OnDemand servers in the network. Follow the directions in the sections below, depending on your configuration. The term "multiple OnDemand server systems" refers to a network where there are two or more OnDemand server systems from which you need to view FNDKEYOND output concurrently using the viewer.

One OnDemand server in the network

If you will be using FNDKEYOND and there is only one OnDemand server where the command will be used, you may allow the command and the companion program to use the default port number (3005). In this case, it is not necessary to modify the shortcut to specify a different port. Skip to the section titled Other workstation considerations. If the above does not describe your configuration, continue with the next section.

Multiple OnDemand servers in the network

If there are multiple OnDemand servers where the FNDKEYOND command may be run, a separate workstation companion program must be started on the workstation for each system.

Each instance of the companion program must be configured to use a different port. After you have created the shortcuts you need, you should alter the properties of the shortcuts to specify a port number and, optionally, a system name. The system name, if specified, will appear as the window title for the viewer which the shortcut will start.

To specify a port and system for a shortcut, right-click the shortcut. From the pop-up menu, select Properties. On the Shortcut tab, you will see an entry box labeled Target. This should contain the path for the QRLROCD.EXE program. At the end of the path, after QRLROCD.EXE, add at least one space, followed by:

```
/p=nnnn /s=systemName
```

where nnnn is the port number and systemName is the name of the OnDemand server.

Example:

```
.../qr1rocd.exe /p=3006 /s=ACCOUNTING
```

You may also use uppercase P and S (...qr1rocd.exe /P=3007 /S=BILLING).

After adding the port and system parameters, press OK to save the new properties of the shortcut. You also should change the title which displays beneath the shortcut to something meaningful, such as "OnDemand Viewer for ACCOUNTING reports".

Make a note of the port numbers and system names you have specified for the shortcuts. It will be necessary to configure the server systems so that FNDKEYOND will use the matching port numbers instead of the default. This is discussed in the section "Server Configuration" below.

Other workstation considerations

For FNDKEYOND output to be displayed correctly using the IBM Content Manager OnDemand viewer, the FNDKEYOND command on the server system must be able to determine the TCP/IP address of the workstation. You should access your line-of-business applications which invoke FNDKEYOND through a terminal emulation session. The connection method for the session should be

TCP/IP. If the FNDKEYOND command is run in a job attached to this session, the address of the workstation can be determined by the FNDKEYOND command, and nothing further needs to be done.

If the emulation session does not use TCP/IP, the FNDKEYOND command searches for the presence of a data area, located in QUSRRDARS library, which has the same name as the device name associated with the emulation session. If the data area exists, the command looks for a non-blank value in positions 1-15 of the data area and uses this value as the workstation IP (Internet Protocol) address. It is the responsibility of the user to ensure that this data area exists and contains the IP address. You can use the CRTDTAARA and CHGDTAARA commands for this purpose.

If the FNDKEYOND command detects errors when sending the search request to the workstation companion program, it will assign a three-digit error code and will display this information in message RDR2882.

The following table lists the error codes and their meanings for exception RDR2882.

Table 81. Exception RDR2882 error codes

Error code	Error description
002	Environment error. A connection could not be established with the workstation companion program because the workstation IP address could not be determined.
091	The TCP/IP socket could not be created.
092	Connect failed for TCP/IP socket. The most likely cause is that the workstation companion program is not running, or there is a port number mismatch between the FNDKEYOND command and the companion program.
093	Socket Close operation failed.

Server configuration

The information in this section is pertinent only if you created more than one shortcut. If you created only one shortcut, the FNDKEYOND command will function correctly using the default port number, and nothing further needs to be done.

If you created more than one shortcut, you should also have specified different port numbers for the shortcuts. Since the FNDKEYOND command must use the same port numbers which you specified on the shortcuts, it is necessary to configure the command on the server systems to use the specified ports.

The FNDKEYOND command checks for the existence of a data area named QRLMFKPORT, in library QUSRRDARS. If this data area does not exist, the commands will use the default port number (3005). If the data area exists, the FNDKEYOND command will use the port number found in the data area contents. The first five characters of the data area contents represent the port number which the FNDKEYOND command should use. The next five characters are no longer used and will be ignored. To create this data area, enter the following:

```
CRTDTAARA QUSRRDARS/QRLMFKPORT TYPE(*CHAR) LEN(10) VALUE('AAAAABBBBB' )
where AAAAA is the FNDKEYOND port number and BBBBB is no longer used
and will be ignored.
```

For example, to specify that FNDKEYOND, on this system, should use port 3007, specify VALUE('0300700000'). It is important that the five-position port number be right-justified within the first five positions of the data area value, with leading zeros if necessary. Enter 00000 for the last five positions that are no longer used.

If there are multiple OnDemand servers in your network, and there are workstations which will display FNDKEYOND output from these systems, you must configure the FNDKEYOND command on these systems to use unique port numbers. Create the data area on the systems where it is necessary to change the port assignments for the FNDKEYOND command. For example, on SYSTEMA, specify VALUE('0300500000'), and on SYSTEMB specify VALUE('0300700000'). It is imperative that these commands on multiple systems in the network use different port numbers.

Server printing and faxing

Using an IBM i printer file to define server print parameters

You can use an IBM i printer file to define some of the server print parameters in your Common Server application definition by specifying the following in the Print Parameters field on the Print Options tab of your application:

`PRTF=LIBRARY/PRINTERFILENAME`

where *LIBRARY* is the name of the library that contains your printer file and *PRINTERFILENAME* is the name of the printer file.

Important:

1. The entire line must be entered in uppercase.
2. The printer file cannot be an externally-described printer file. An externally-described printer file is created using the Create Printer File (CRTPRTF) command, where a source file is named in the keyword SRCFILE that contains DDS that you have created. If you have an externally-described printer file, then create another printer file using the CRTPRTF command. Specify the same parameters (LPI, CPI, page size, overlays, and so on) as found in the externally-described printer file except for the Source File parameter value (keyword SRCFILE). You should also compare the attributes of the original spooled file with the parameters you specify in the new printer file to confirm that all parameters are set correctly.

This printer file specification is also used when faxing from the server to produce the temporary IBM i spooled file which is then faxed.

The following parameters, even though specified in the printer file identified in the PRTF parameter, are overridden as shown in the following table.

Table 82. Parameter Overrides

Printer file parameter	Changed to	Where specified	When specified
DEVTYPE	*AFPDS if AFP data.*SCS if SCS or SCS-extended data.*LINE if Line data and an EBCDIC code page is used; otherwise it is set to *SCS and the data is converted from ASCII to EBCDIC.	Application definition - View Information tab - Data Type field	When defining the application
OUTQ	Server Queue Name	Printer definition	When printing the document, you select the Server Printer definition to use
COPIES	Number of Copies	Client Print window	When printing the document

Table 82. Parameter Overrides (continued)

Printer file parameter	Changed to	Where specified	When specified
USRDTA	The first 10 characters of your OnDemand user ID	When you logon to OnDemand	When you logon to OnDemand
USRDFNDTA	Application name	Application definition - General tab - Name field	When defining or updating the application
PAGERANGE	Pages value	Client Print window, but only if already viewing the document	When printing the document

If no PRTF parameter is specified, printer file QSYSPRT is used and in addition to the printer file parameters specified above, the parameters in the following table are also set.

Table 83. QSYSPRT Parameters

Printer file parameter	Changed to	Where specified	When specified
CTLCHAR	*FCFC for SCS and Line data; otherwise use the CTLCHAR value specified in the QSYSPRT printer file.	Application definition - View Information tab - Data Type field	When defining or updating the application
PAGRTT	Orientation (for AFP)	Application definition - View Information - Orientation field	When defining or updating the application

Server fax setup

When defining a printer for use with server fax functions, a Server Queue Name is required in the printer definition within the IBM Content Manager OnDemand Administrator. You can either specify the name of a library/output queue or *NONE. If *NONE is specified, the output is temporarily spooled to QUSRRDARS/QRDARS400 output queue before it is sent to fax using the QRLMSFAX program. If you enter a library/output queue name, then Content Manager OnDemand will use that output queue to temporarily spool the data.

If you use the SNDFAX command in the QRLMSFAX program (which is the default as shipped), the restrictions in the following table apply to the fax information. If you use a different server fax product, see "Facsimile user exit program" on page 318.

Table 84. SNDFAX command fax information restrictions

OnDemand client prompt field	SNDFAX maximum length	SNDFAX keywords used	Length passed to QRLMSFAX
Recipient Attention	40	TO position 2	100
Recipient Company	40	TO position 3	100

Table 84. S NDFAX command fax information restrictions (continued)

OnDemand client prompt field	S NDFAX maximum length	S NDFAX keywords used	Length passed to QRLMSFAX
Recipient Fax Number	32	TO position 1, TO position 4	100
Sender Name	40	FROM position 1	100
Sender Company	40	FROM position 2	100
Sender Tel Number	40	FROM position 3	100
Sender Fax Number	–	(not used)	100
Sender Cover Page	10	CRTCVRP - can be blank, *NO or *YES; default is *YES	10
Subject	40	TITLE	100
Notes	40	COMMENT	100

API and user exit reference

API reference

This section contains reference information about the IBM Content Manager OnDemand server Application Programming Interfaces (APIs).

The APIs contained in this reference include APIs for various Content Manager OnDemand functions beyond the standard command set described in “Command reference” on page 243. The APIs are presented in alphabetical order. Each API contains a description of its purpose and syntax (including descriptions of the parameters that can be used). Examples and general information about using the APIs have also been included.

Important: These APIs require that you be signed on to the server with a user profile that is also defined as a user in Content Manager OnDemand.

Using quotes when executing the APIs

When you execute an IBM Content Manager OnDemand program from the IBM i command line interactively, in batch using the SBMJOB command or the QSHELL environment and you specify parameter values that contain a null (blank) character or some other special character (such as the parenthesis), you must delimit the parameter value with double quote characters. For example, when executing the ARSDOC GET program from the i command line and you specify the **-o** parameter to specify one or more application group field names, each field name is enclosed with parenthesis and the entire string must be surrounded by double quote characters, for example: `-o "(sdate)(student)"`. See the operating system documentation for more information about the use of quotes when executing the APIs. Note that if the same string is included in a parameter file using the **-F** parameter on the ARSDOC GET command, it would be enclosed in brackets and the double quotes are not required, for example: `[-o (sdate)(student)]`.

Using the QSHELL environment

All ARSxxxxxx API programs must be executed using the QSHELL environment. QSHELL is a command interpreter that allows IBM i to execute AIX commands on i. These commands can also be run from a script file. More information about QSHELL can be found in the IBM i Information Center. Option 30 for 5770-SS1 (IBM i) installs the QSHELL interpreter. The QSHELL environment is started by the **QSH** or **STRQSH** IBM i command.

Multiple ARSxxxxxx programs may be issued in one QSHELL session. After the completion of a QSHELL command, a \$ is displayed. This notifies the user that the QSHELL command has completed. This does not mean that the command actually ran successfully and performed the desired function.

All command examples assume that the QSHELL environment has been started and the ARSxxxxxx programs are in a directory specified in the PATH variable. IBM Content Manager OnDemand places a symbolic link to the commands in the /usr/bin directory during installation. To exit the QSHELL environment, press the **F3** key.

The **-v** parameter will provide detailed program logging information. The **-u** parameter specifies a valid Content Manager OnDemand user that exists with the proper authority to perform the functions requested. The **-p** parameter specifies the Content Manager OnDemand password for the Content Manager OnDemand user specified in the **-u** parameter. The user ID and password are not normally needed when running these programs in i. By default, the current i user profile and password are used as the Content Manager OnDemand user ID. If your i user profile does not exist in Content Manager OnDemand, you must specify a valid Content Manager OnDemand user ID and password to use these programs.

The **-h** is the Content Manager OnDemand instance name where the program is to execute. For the purposes of this document, the administrative user name will be `testadmin`, the password will be `ondemand` and the host name is `QUSROND`. User `testadmin` must have the appropriate authority to perform the function being requested by the `ARSxxxxxx` command.

Calling QSHELL commands from an IBM i command line

When executing the `ARSxxxxxx` programs from an IBM i command line, the `i QSH` command is used. When embedded single quotes exist within the program parameter, they must be doubled. For example:

```
QSH CMD('arsdoc get ..... "-i WHERE Account#='1234567"' "....-u testadmin -p ondemand')
```

Since the QSHELL environment will log messages to the terminal, you will want to control this by setting environment variable `QIBM_QSH_CMD_OUTPUT`. This variable will control where the messages are logged. This does not interfere with the messages that get logged to the System Log. The QSHELL environment logging can be sent to the terminal session, to a file in the IFS directory structure, or you can choose not to log the messages. The `ADDENVVAR`, `CHGENVVAR`, `RMVENVVAR` and `WRKENVVAR` commands can be used to manipulate the QSHELL environment. The environment parameters can be set for the entire system and for the job. In most cases, you will not want to change the environment parameter at the system level, because it will affect all users of QSHELL on the system.

- Set the job environment to log messages to the terminal
`ADDENVVAR ENVVAR(QIBM_QSH_CMD_OUTPUT) VALUE(STDOUT) LEVEL(*JOB)`
- Set the system environment to not log messages.
`ADDENVVAR ENVVAR(QIBM_QSH_CMD_OUTPUT) VALUE(NONE) LEVEL(*SYS)`
- Set the job environment to log messages to the a file in the IFS. `mydirectory` must exist in the IFS.
`ADDENVVAR ENVVAR(QIBM_QSH_CMD_OUTPUT)VALUE('FILE=/mydirectory/QSHELL_output') LEVEL(*JOB)`

The following is an example of the messages that might be displayed to the terminal or logged in the output file.

```
OnDemand Load Id = >8495-53-0-1FAA-20130627000000-20130627000000-8496<
Loaded 4 rows into the database
Document compression type used - 0D77. Bytes Stored = >10240<
```

Message `QSH0005` will be issued in the job log when running the `QSH` command. The message text is `Command ended normally with exit status &1`. The possible statuses returned are shown in the following table.

Table 85. Message QSH0005 status codes

Status	Description
0	Completed Successfully
1	Command Failure
2	Folder does not exist or do not have authority to folder
3	User ID or Password is not valid. Cannot establish communication to server
127	Command Not Found
254	No Hits Match Query

How to read a syntax diagram

A syntax diagram shows you how to specify an API program so that the operating system can correctly interpret what you type.

Read a syntax diagram from left to right and from top to bottom, following the horizontal line (the main path). If the line ends with an arrowhead, the API syntax is continued and the next line starts with an arrowhead. Facing arrowheads mark the end of the API syntax.

When you type an API from the syntax, be sure to include punctuation, such as commas and equal signs.

Parameters are classified as keywords or variables:

- Keywords represent constants and are shown (in syntax) in uppercase letters; however, you can enter keywords in either uppercase or lowercase.
- Variables represent names or values you supply and are shown (in syntax) in lowercase letters; however, you can enter variables in either uppercase or lowercase unless the syntax directions explicitly state the case restrictions. An example of a variable is a file name.

A parameter can be a combination of a keyword and a variable.

Required parameters are displayed on the main path.



Optional parameters are displayed below the main path.



A stack of parameters, with the first parameter displayed on the main path, shows that you must choose one of the parameters.



A stack of parameters, with the first parameter displayed below the main path, shows that you can choose one of the parameters.



An arrow returning to the left, above the path, shows that items can be repeated following these conventions:

- If the repeat arrow contains a break, the item can be repeated in a list with the items separated by blank spaces.




- If the repeat arrow contains a comma, the item can be repeated in a list with the items separated by commas.




You can repeat items from parameter stacks following the stack conventions for required and optional parameters described previously.

Some syntax diagrams contain parameter stacks within other parameter stacks. You can only repeat items from stacks according to the conventions described previously. That is, if an inner stack does not have a repeat arrow above it but an outer stack does, you can choose only one parameter from the inner stack and combine it with any parameter from the outer stack, and that combination can be repeated. For example, the following diagram shows that you could combine parameter **choice2a** with parameter **choice2** and then you can repeat that combination again (choice2 plus choice2a).



Some APIs are preceded by an optional path parameter.



If you do not supply the path parameter, the system searches the current directory for the API. If the API is not in the current directory, the system continues to search for the API using the directories defined in the PATH environment variable.

Some APIs in this section have several formats that accomplish the same task. These APIs appear (in syntax) similar to the following:



The description of the API directs you to the correct format to use.


ARSDATE

Purpose

For Date (old style), Date/Time (old style), Date/Time (TZ) (old style), and Time (old style) fields, do one of the following tasks:

- Content Manager OnDemand displays the internal database value for the specified date and time string.
- Content Manager OnDemand displays the date and time string for the specified internal database value.

Syntax



Description

The ARSDATE program displays the IBM Content Manager OnDemand internal database value for the specified date and time string or displays the date and time string for the specified Content Manager OnDemand internal database value.

Values whose data types are any of the following are represented in an internal form that is not apparent to the casual user of Content Manager OnDemand:

- Date (old style)
- Date/Time (old style)
- Date/Time (TZ) (old style)
- Time (old style)

Casual users enter date and time values the same way that they are displayed in a report. However, to search the database with an SQL string, a user must enter the internal form of the value. The ARSDATE program lists the internal value of a date or time string.

Parameters

-a Use to display database values and date strings for Date fields. For example, to display the database value for the date 9/1/99, enter:

```
arsdate -a 9/1/99
```

The ARSDATE program returns:

```
9/1/99 -> 10836
```

To display the date string for the database value 10836 in the default date format, enter:

```
arsdate -a 10836
```

The ARSDATE program returns:

```
10836 -> 9/1/99
```

-i Use to display database values and date and time strings for Date (old style), Date/Time (old style), Date/Time (TZ) (old style), and Time (old style) fields. The time part of the Date/Time (old style) and Date/Time (TZ) (old style) fields is not adjusted for the local time zone. You typically use the **-i** parameter to find out one of two things:

- Given a date and time printed on a report, what value did Content Manager OnDemand store in Date (old style), Date/Time (old style), Date/Time (TZ) (old style), or Time (old style) database field? You can use the result to search a date or time field with an SQL string.
- Given a value stored in a Date (old style), Date/Time (old style), Date/Time (TZ) (old style), or Time (old style) database field, what would be the date and time printed on a report?

-t Use to display database values and time strings for Time fields. The time is not adjusted for the local time zone. For example, to display the database value for the time 04:00:00, enter:

```
arsdate -t 04:00:00
```

The ARSDATE program returns:

```
04:00:00 -> 4800
```

To display the time string for the database value 4800, enter:

```
arsdate -t 4800
```

The ARSDATE program returns:

```
4800 -> 04:00:00
```

-z Use to display database values and date and time strings for Date/Time (TZ) fields. The time part of a Date/Time (TZ) field is adjusted for the local time zone. If you run the ARSDATE program with the **-z** parameter on systems in different time zones and you specify the same date and time value, the result will be different. For example, suppose that you need to determine the value stored in a Date/Time (TZ) field for "09/01/00 04:00:00". The command:

```
arsdate -z "09/01/00 04:00:00"
```

When run on a server in the Eastern time zone will return:

```
09/01/00 04:00:00 -> 936187200
```

If you run the same command on a server in the Mountain time zone, then result will be:

```
09/01/00 04:00:00 -> 936180000
```

A typical use of the **-z** parameter is to determine a database value with which to search the system log. You can use the result to search the Date/Time (TZ) field of the system log with an SQL string. For example, suppose a user in New York logs on to a server in Denver. To retrieve the log on messages with an SQL string, you must specify the date and time part of the query using the local time of the client that is running the query. If the user logged on to the server at 4 a.m. Eastern time, then a query that is run in Denver must specify 2 a.m. to retrieve the message.

-g Display the current system date.

-d days

Add the specified number of days to the specified internal value.

-h hours

Add the specified number of hours to the specified internal value.

-m months

Add the specified number of months to the specified internal value.

-n minutes

Add the specified number of minutes to the specified internal value.

-s seconds

Add the specified number of seconds to the specified internal value.

-y years

Add the specified number of years to the specified internal value.

-f "format"

Determines the format of the date and time string that Content Manager OnDemand displays.

Table 86 lists the standard date and time formats that are supported by Content Manager OnDemand. If the input data contains a date or time format that is not listed in the table, you can specify the format to Content Manager OnDemand. However, when specifying a format, you can only use values and separators from the standard formats that are listed in the table. Date and time formats may also be specified on the Load Information page in the application and the Field Information page in the folder.

The format consists of a set of values (for example, %m) and separators (such as the blank character). Some of these formats may require the removal of leading or imbedded blanks, or other characters.

Table 86. Date and time format specifications

Date Format Specifier	Date Format	Example
%m/%d/%y	mm/dd/yy	01/31/95
%d/%m/%y	dd/mm/yy	31/01/95
%f/%e/%y	m/d/yy	1/31/95
%e/%f/%y	d/m/yy	31/1/95
%m-%d-%y	mm-dd-yy	01-31-95
%d-%m-%y	dd-mm-yy	31-01-95

Table 86. Date and time format specifications (continued)

Date Format Specifier	Date Format	Example
%m%d%y	mmddyy	013195
%m%d%Y	mmddyyyy	01311995
%y%m%d	yymmdd	950131
m%d	yyyymmdd	19950131
%b %e, %Y	Mth d, yyyy	Jan 31, 1995
%B %e, %Y	Month d, yyyy	January 31, 1995
%m/%d/%y %H:%M	mm/dd/yy hh:mm	01/31/95 10:50
%H:%M	hh:mm	10:50
%T	hh:mm:ss	10:50:59

Important: Alpha dates, for example, the *Mth* (%b) and *Month* (%B) formats are supported for English only. In other words, you must only use those formats for months specified in English, such as Jan or January.

internalValue

The internal date and time value from the Content Manager OnDemand database. Enter one or more internal values when you want to display formatted date strings.

"dateString"

The date and time string. Enter one or more strings when you want to display internal date values.

Examples

1. The following example shows how to determine the database value for the specified date and time string. The data type of the database field is **Date/Time (old style)**.

```
arsdate -i "09/01/99 04:00:00"
09/01/99 04:00:00 -> 936158400
```

In the example, you could use the database value to search a **Date/Time (old style)** field in the database with an SQL string. For example:

```
arsdoc get -i "WHERE somedate=936158400" ...
```

2. The following example shows how to determine the date and time string for the specified database value. The data type of the database field is **Date/Time (old style)**. The result is shown using the default display format.

```
arsdate -i 936158400
936158400 -> 09/01/99 04:00:00
```

3. The following example shows how to determine the database value for the specified date and time string. The data type of the database field is **Date/Time (TZ) (old style)**. The ARSDATE program adjusts the time part of the result for the local time zone.

```
arsdate -z "09/01/99 04:00:00"
09/01/99 04:00:00 -> 936180000
```

If you were to run the same command on a server in the Eastern time zone, the result would be:

```
arsdate -z "09/01/99 04:00:00"  
09/01/99 04:00:00 -> 936187200
```

Notes

1. The ARSDATE program works with the following data types to display their internal values in a date or time format that you can understand:

- Date (old style)
- Date/Time (old style)
- Date/Time (TZ) (old style)
- Time (old style)

The following data types are stored in a date and time format that you can understand; therefore, you do not need to use the ARSDATE program to interpret the values:

- Date
- Date/Time
- Date/Time (TZ)

2. When displaying the internal value for a given date string, by default, the ARSDATE program expects you to specify the date string using the mm/dd/yy format. If you want to specify the date string using a different format, then you must specify the date format with the **-f** parameter. For example:

```
arsdate 11/12/99  
11/12/99 -> 10908  
  
arsdate 11/12/1999  
11/12/1999 -> -1 (Error)  
  
arsdate -f "%m/%d/%Y" 11/12/1999  
11/12/1999 -> 10908
```

3. The upper limit of the date is September 17, 2059. To format the output for a date, specify the **-f** parameter as follows:

```
arsdate -f "%m/%d/%Y" 32767  
32767 -> 09/17/2059
```

IFS Location

/usr/bin/arsdate

The IBM i executable program.

ARSDOC

Purpose

The ARSDOC program is a multi-purpose document processing program. You can use the ARSDOC program to do the following tasks:

- Query the database and generate a list of items that match a query
- Retrieve documents from the system
- Add, delete, and update documents
- Send documents to a server printer
- Add and release holds on documents
- Add or remove documents from full text indexing

- Federate documents to IBM FileNet® P8 through Content Federation Services for Content Manager OnDemand

Description and syntax

The ARSDOC program provides the following functions:

- ADD
- DELETE
- GET
- PRINT
- QUERY
- UPDATE
- HOLD_ADD
- HOLD_RELEASE
- CFSOD-FED
- FTI_ADD
- FTI_RELEASE

For each function, you can specify all of the required options on the command line or you can specify the name of a parameter file that contains the options. The syntax of each function is listed twice: first, when you specify the options on the command line; second, when you specify a parameter file.

ADD function:

Use to store data into the system by specifying the folder, application group, application, and database fields and values.

Important: Set the proper locale before issuing the ARSDOC ADD API. See the chapter entitled “Defining a locale” in the *IBM Content Manager OnDemand for i Common Server Planning and Installation Guide* for important details.

If you specify the **-0** parameter, then you do not have to specify all of the database fields; however, you must always specify fields related to date or time fields. When you specify the **-0** parameter, the ARSDOC program stores a default value in any database field that you omit. The default value for string fields is an empty string. The default value for numeric fields is 0 (zero). Numeric fields include integer and decimal (floating point) fields. When adding a document, you can specify the data in one of three ways:

- A document file
- An SQL query that contains clauses, database field names, index values, and operators
- A public named query

When you specify an SQL query or public named query, you are creating a database row that points to an existing document that has been identified by the query.

When you use the ADD function to add a row for an existing document, the row must be added to the application group and the application that contain the document. The application must be specified with the **-a** parameter. The application group must be specified with the **-g** parameter and must be one of the application groups referenced by the folder named with the **-f** parameter.

Important: The ADD function will fail unless the Database Organization for the application group named with the **-g** parameter is Multiple Loads per Database

Table and the Expiration Type is Segment or Document. If only index data is added, and no document data is added, the Expiration Type can be Segment or Document. You can add index only data for an existing document when the expiration type is Load.

Storage management for the application group needs to be configured to use cache. Cache is required to allow the OnDemand server to build a larger storage object and append to the data. Then this storage object can be migrated to TSM. Storing data directly into TSM is not recommended and can cause TSM performance issues.

```

>> arsdoc add --a application [-f folder] [-g applGroup] [-G applGroup]

```

```

>> -h instance [-i sqlQuery] [-l holdname] [-o docfilename] [-q namedQuery] [-t search_string] [-T search_string] [-n dbfield=value] [-_0]

```

```

>> [-p password] [-S startdate,enddate] [,-format] [-u userid]

```

```

>> [-U user_alias] [-v]

```

When you use the parameter file option, you must specify the **-F**, **-h**, **-u**, **-U**, and **-v** parameters on the command line. All other parameters must be specified in the parameter file. The following diagram shows the syntax of the ADD function when you use the parameter file option:

```

>> arsdoc add --F parmfile --h instance [-p password] [-u userid]

```

```

>> [-v]


```

The syntax of the parameter file for the ADD function is:

```

>> [--a application] [--f folder] [-g applGroup] [-G applGroup]

```




DELETE function:


Use to delete index records that point to individual documents in IBM Content Manager OnDemand. The data will still exist on disk or archive media, because the documents being deleted might represent only a few pages of an entire input file. However, the data will no longer be retrievable after the index records are deleted.

To identify the documents that you want to delete index records, you must enter an SQL query or specify the name of a public named query. The SQL query must contain clauses, database field names, index values, and operators. The DELETE function deletes index records for all documents that match the query.

If you want to delete an entire input file (load) of documents (such as an entire spooled file) from Content Manager OnDemand, use the Remove Report from OnDemand (RMVRPTOND) command instead of ARSDOC DELETE.




When you use the parameter file option, you must specify the **-F**, **-h**, **-u**, **-U**, and **-v** parameters on the command line. All other parameters must be specified in the parameter file. The following diagram shows the syntax of the DELETE function when you use the parameter file option:






The syntax of the parameter file for the DELETE function is:




GET function:

Use to retrieve documents and resources from the system. The GET function can also generate and save generic index data for the documents that match the query. You must identify the name of the IBM Content Manager OnDemand library server. You specify the application group or folder that you want to search. To query the database, you can enter an SQL query or specify the name of a public named query. The SQL query must contain clauses, database field names, index values, and operators. By default, the ARSDOC program sends a copy of the documents that match the query to the display (interactive) or the job log (batch). You can also choose to write the output to a file. To retrieve documents in a sorted order, you must specify the **-n** parameter.



When you use the parameter file option, you must specify the **-F**, **-h**, **-u**, **-U**, and **-v** parameters on the command line. All other parameters must be specified in the parameter file. The following diagram shows the syntax of the GET function when you use the parameter file option:



```

┌-u-userid┐ ┌-v┐
└────────┘ └──┘

```

The syntax of the parameter file for the GET function is:

```

┌-a-┐ ┌-A-value┐ ┌-c-┐ ┌-d-directory┐
└──┘ └────────┘ └──┘ └────────┘

```

```

┌-f-folder┐ ┌-g-┐ ┌-G-applGroup┐
└────────┘ └──┘ └────────┘

```

```

┌-i-sqlQuery┐ ┌-L-max#┐ ┌-n-┐ ┌-N-┐
├-q-namedQuery┤ └────────┘ └──┘ └──┘
├-t-search_string┤
├-T-search_string┤
└────────┘

```

```

┌-o-name┐ ┌-S-startdate,enddate┐ ┌-format┐
└────────┘ └────────┘ └────────┘

```

```

┌-x-loadId┐ ┌-X-loadID┐
└────────┘ └────────┘

```

PRINT function:

Use to send documents to an IBM Content Manager OnDemand server printer. You must name the Content Manager OnDemand library server and the folder that you want to search and specify the query to run. The items that match the query are sent to the server printer named with the **-P** parameter. The server printer must be defined to Content Manager OnDemand using the administrative client. To query the database, you can enter an SQL query or specify the name of a public named query. The SQL query must contain clauses, database field names, index values, and operators. You can limit the number of documents sent to the printer by using the **-L** parameter and specifying the maximum number of documents that should be retrieved, regardless of the number of documents that match the query. You can limit the number of database tables searched, and possibly increase the performance of a query, by specifying the **-S** parameter and specifying a start date and an end date. The PRINT function does not currently support server FAX.

```

┌-arsdoc print┐ ┌-f-folder┐ ┌-G-applGroup┐ ┌-h-instance┐
└────────┘ └────────┘ └────────┘ └────────┘

```

```

┌-i-sqlQuery┐ ┌-l-holdname┐ ┌-L-max#┐ ┌-p-password┐
├-q-namedQuery┤ └────────┘ └──┘ └────────┘
├-t-search_string┤
├-T-search_string┤
└────────┘

```

```

┌-P-printer┐ ┌-S-startdate,enddate┐ ┌-format┐ ┌-u-userid┐
└────────┘ └────────┘ └────────┘ └────────┘


```

```


┌-U-user_alias┐ ┌-v┐ ┌-X-loadID┐
└────────┘ └──┘ └────────┘

```

When you use the parameter file option, you must specify the **-F**, **-h**, **-u**, **-U**, and **-v** parameters on the command line. All other parameters must be specified in the parameter file. The following diagram shows the syntax of the PRINT function when you use the parameter file option:




The syntax of the parameter file for the PRINT function is:




QUERY function:


By default, the ARSDOC program sends the list of items that match the SQL query to the SQL query to the display (interactive) or the job log (batch). You can also choose to write the output to a file.




When you use the parameter file option, you must specify the **-F**, **-h**, **-u**, **-U**, and **-v** parameters on the command line. All other parameters must be specified in the parameter file. The following shows the syntax of the QUERY function when you use the parameter file option:



When you use the parameter file option, you must specify the **-F**, **-h**, **-u**, **-U**, and **-v** parameters on the command line. All other parameters must be specified in the parameter file. The following shows the syntax of the UPDATE function when you use the parameter file option:



The syntax of the parameter file for the UPDATE function is:





HOLD_ADD function:

Use to add documents to a hold. Use an SQL string or Named Query to query the database and determine the documents to add to the hold.


Ensure that enhanced retention management is turned on for the application group and you have permission to add documents to a hold before you attempt this function. Otherwise, a permission error will occur.

If a duplicate attempt is made to add the same documents to a hold, no error occurs, and ARSDOC displays a successful message. However, the system log might contain messages that indicate that the documents were already added.






When you use the parameter file option, you must specify the **-f**, **-h**, **-u**, **-U**, and **-v** parameters on the command line. All other parameters must be specified in the parameter file. The following diagram shows the syntax of the hold_add function when you use the parameter file option:



The syntax of the parameter file for the hold_add function is:



HOLD_RELEASE function:


Use to remove documents from a hold. You can query the database by using an SQL string or Named Query to determine the documents that should be removed from a hold.

If you selected the option in the application group to use implied hold, documents are placed on hold as soon as they are loaded, and there is not a hold name that is associated with the documents that are placed on hold. To remove documents from an implied hold, specify IMPLIED_HOLD as the hold name.


Ensure that enhanced retention management is turned on for the application group and you have permission to release documents from a hold before you attempt this function. Otherwise, a permission error will occur.

If a duplicate attempt is made to remove the same documents from a hold, no error occurs, and ARSDOC displays a successful message. However, the system log might contain messages that indicate that the documents were already released.






When you use the parameter file option, you must specify the **-f**, **-h**, **-u**, **-U**, and **-v** parameters on the command line. All other parameters must be specified in the parameter file. The following diagram shows the syntax of the `cfsod_fed` function when you use the parameter file option:



The syntax of the parameter file for the `cfsod_fed` function is:



Parameters

-a For the GET function, when retrieving AFP documents, specify this parameter to include resources with the documents that are retrieved. If documents from the same application have different resource groups, then the ARSDOC program creates separate output files for each resource group.

For the ADD function, you must specify the name of the IBM Content Manager OnDemand application. The application must belong to the application group named with the **-g** (or **-G**) parameter.

-A value

Use this parameter to retrieve annotations. The following are the basic values of this parameter:

- 0 Include public text annotations
- 1 Include private annotations
- 2 Include annotations that cannot be copied to another server
- 4 Include graphic annotations

You can also add up two or more of the basic parameter values to create new values. For example:

- 3 Retrieve all public and private text annotations
- 5 Retrieve public and private annotations (text and graphic) that can be copied to another server
- 6 Retrieve all public annotations (text and graphic)
- 7 Retrieve all text and graphic annotations

Table 87. Possible values for the ARSDOC GET function -A parameter

Flag value	Public	Text	Can be copied	Private	Cannot be copied	Graphic
-A 0	X	X	X			
-A 1	X	X	X	X		
-A 2	X	X	X		X	
-A 3	X	X	X	X	X	
-A 4	X	X	X			X
-A 5	X	X	X	X		X
-A 6	X	X	X		X	X
-A 7	X	X	X	X	X	X

See “Examples” on page 290 for examples of using the **-A** parameter with the ARSDOC GET function.

Do not use the **-A** parameter with the **-X** parameter. When you use the **-X** parameter, all of the documents in a load are retrieved. The documents in the load are identified by an index file that was created when the data was loaded. When the documents were loaded, no annotations existed. Therefore, no annotation information exists in the index file for the documents, and you cannot retrieve annotations by using this method.

-B *orderbystring*

For the QUERY and PRINT functions, used to specify which database field is used to sort the document list. The **-B** option cannot be specified if a load ID is specified.

-c

For the GET function, use to concatenate all of the documents that match the query into one output file. Name the output file with the **-o** parameter. However, even if you do specify the **-c** parameter, the ARSDOC program creates separate output files when any of the following conditions occur:

- If more than one application group is referenced by the folder. The ARSDOC program creates one output file for each application group that contains items that match the query.
- If more than one application is contained in an application group. The ARSDOC program creates one output file for each application that contains items that match the query.
- If documents from the same application have different resource groups, the ARSDOC program creates separate output files for each resource group.

For example, if a folder references two application groups, then the following specification:

```
-o student -c
```

Can result in file names such as: student.516 or student.517

Where 516 and 517 are application group identifiers. One file is created for each application group. Each file contains all of the items that match the query for that particular application group.

If you have difficulty viewing documents that are retrieved in the same file, retrieve the documents as individual documents.

-d *directory*

The name of the directory where the ARSDOC program writes the output files. The directory must exist before the ARSDOC program attempts to save the output files.

-D For the QUERY function, appends the document handle information to the end of each line. The document handle information consists of the following ten values, in the order listed:

1. Document name
2. Offset
3. Length
4. Compressed object offset
5. Compressed object length
6. Annotation type
7. Compression type
8. Resource ID
9. Primary node ID
10. Secondary node ID

The values are separated by a delimiter. The default delimiter is the comma character. You can specify a different delimiter with the **-e** parameter.

-e *delimiter*

For the QUERY function, specifies a one character delimiter to use as a separator between values. By default, Content Manager OnDemand separates values in the output with a comma.

-f *folder*

The name of the Content Manager OnDemand folder. The folder name must be specified exactly as it appears in Content Manager OnDemand. The case of the folder name is significant. For example, to query the System Log folder, you must enter:

```
-f "System Log"
```

If you are using a parameter file, then you must specify the **-f** parameter in the parameter file. If you are not using a parameter file and you do not specify the **-f** parameter, then the ARSDOC program prompts you for the folder name when you run the program.

For the GET and QUERY functions, you can omit the **-f** parameter and specify the **-G** parameter to search a specific application group.

For the UPDATE function, if the folder that is specified with the **-f** parameter contains only one application group, then you can omit the **-g** or **-G** parameter (you do not have to specify the name of the application group).

When you specify the **-X** parameter, you cannot specify the **-f** parameter.

Note: The following information applies only when an application group name is not provided.

A folder can be used to search one or more application groups. Because the ARSDOC program generates a single SQL query to search all of the application groups, the properties of the database fields must be the same for each application group. The properties include the field name, type, and length. For example, suppose that you define the following application groups and fields:

Application Group	Field Names
Student Bills	name, account, billDate
Student Grades	name, account, gradeDate
Student Transcripts	name, account, transcriptDate

You cannot query the application groups using the ARSDOC program because the name of the date field is not the same for each application group. However, if you were to define the application groups and fields as follows:

Application Group	Field Names
Student Bills	name, account, studentDate
Student Grades	name, account, studentDate
Student Transcripts	name, account, studentDate

Then you could query the application groups using the ARSDOC program because the names of the database fields are the same for each application group.

-F *parmFile*

Specifies the name of the file that contains the actions to run and other parameters, values, and options. You typically specify this option when you want to perform more than one action.

When you specify a parameter file, delimit the options and values with left [and right] brackets. The left and right brackets identify each parameter in the file, are required in the parameter file, and the parameter values cannot contain left or right brackets. In the following example, the parameter file for the ARSDOC query function is *parmfile.txt*:

```
arsdoc query -u oduser -p odpasswd -h odserver -v -F parmfile.txt
```

The following lines are examples of parameter values in the *parmfile.txt* parameter file:

```
[-f "Credit Card Statements"] [-i "where account = '000-000-000'"]
[-f "Credit Card Statements"] [-i "where account = '000-000-001'"]
```

If you want to use the left or right bracket as part of a parameter value, add a line to the beginning of the file that redefines the left and right delimiters with the keyword **DELIMS**. The keyword **DELIMS** is case-sensitive, must be on the first line of the file and it must start at the beginning of the line. For example, if you want to indicate that the left and right curly braces are the delimiters, add the line **DELIMS={}**. You can specify any two characters that are different, are not a space, and are not in any of the parameter values. The following example shows left and right curly braces defined as delimiters:

```
DELIMS={}
{-f "Credit Card Statements"} {-i "where account = '000-000-000' and name = 'Republic Bank [North]'" }
{-f "Credit Card Statements"} {-i "where account = '000-000-000' and name = 'Republic Bank [South]'" }
```

| An action (one or more input lines) can contain a maximum of 32767
| characters (bytes).

| You can use the \ (backslash) character to continue the parameters of an
| action to two or more lines.

| A parameter file can contain blank lines and comment lines. A comment
| line contains the # character in the first column.

-g For the GET function, use to generate Generic indexer data for the items that match the query.

See the *IBM Content Manager OnDemand for i Common Server: Indexing Reference* for details about the Generic indexer.

When you specify the **-g** parameter, you must specify the **-c**, **-N**, and **-o** parameters. However, you cannot specify database field names with the **-o** parameter.

The ARSDOC program uses the following convention to name the output files that are generated with the **-g** parameter:

-o.res_id.appl_group.appl.type

Where:

- *-o* is the value specified with the **-o** parameter
- *res_id* is the resource group identifier.
- *appl_group* is the name of the application group
- *appl* is the name of the application
- *type* is the file type:
 - out identifies a document file
 - ind identifies a generic index file
 - res identifies a resource file

In general, the number of files generated is dependent on the number of application groups in a folder, the number of applications in an application group, and the number of versions of resource groups in an application.

For the ADD and UPDATE functions, specifies the name of the Content Manager OnDemand application group. The application group that you specify is searched from the folder that is named with the **-f** parameter. For the UPDATE function, if the folder that is specified with the **-f** parameter contains only one application group, then you can omit the **-g** parameter (you do not have to specify the name of the application group).

-G *applGroup*

Use to specify the name of the application group.

For UPDATE: If the folder that is specified with the **-f** parameter contains only one application group, then you can omit the **-G** parameter (you do not have to specify the name of the application group).

For ADD: When the database query is run to retrieve the document that contains the data that is to be used in the add function, the search is limited to the specified application group, even if the folder named with the **-f** parameter can be used to search more than one application group. This ensures that only documents in the specified application group can be used for the add function. You can specify the name of the application group with the **-g** parameter or the **-G** parameter.

For DELETE: The **-G** parameter is an optional parameter. If specified, then the database query that is run to determine the document(s) to delete is limited to the specified application group. The addition of the **-G** parameter allows you to delete documents from a specific application group in folders that can search more than one application group. If you do not specify the **-G** parameter, then the query runs against all of the application groups that can be searched from the folder.

For GET: Specifies the application group to query and retrieve documents from. The **-G** parameter lets you retrieve documents from a specific application group from a folder that can search more than one application group. If you do not specify the **-G** parameter, then the query runs against all of the application groups that can be searched from the folder. You can omit the **-f** parameter and specify the **-G** parameter to search a specific application group. The **-G** parameter is required if you specify the **-X** parameter.

For PRINT: The **-G** parameter is an optional parameter. If specified, then the database query that is run to determine the document(s) to print is limited to the specified application group. The addition of the **-G** parameter allows you to print documents from a specific application group in folders that can search more than one application group. If you do not specify the **-G** parameter, then the query runs against all of the application groups that can be searched from the folder.

For QUERY: Specifies the application group to search. The **-G** parameter lets you search a specific application group from folders that can search more than one application group. If you do not specify the **-G** parameter, then the query runs against all of the application groups that can be searched from the folder. You can omit the **-f** parameter and specify the **-G** parameter to search a specific application group. The **-G** parameter is required if you specify the **-X** parameter.

For UPDATE: When the database query is run to determine the document(s) to update, the search is limited to the specified application group, even if the folder named with the **-f** parameter can search more than one application group. This guarantees that only documents in the specified application group can be updated. You can specify the name of the application group with the **-g** parameter or the **-G** parameter.

You can use the **-G** parameter with the **-i** parameter to query folders that can search more than one application group. For example, a folder contains three application groups; you want to query only one of the application groups. Use the **-G** parameter to specify the name of the application group that you want to query. Use the **-i** parameter to specify the application group's database field names. You can also use the **-G** and **-i** parameters when the application groups have different database field names. The following example shows how to search a folder and three application groups that have different database field names:

```
arsdoc get -f "Student Information" -G loans
-i "WHERE number LIKE '123456' AND loanDate = 10593"
arsdoc get -f "Student Information" -G grades
-i "WHERE number LIKE '123456' AND gradeDate = 10593"
arsdoc get -f "Student Information" -G transcripts
-i "WHERE number LIKE '123456' AND transDate = 10593"
```

You can use the **-G** parameter with the **-q** parameter to query folders that can search more than one application group. When you specify the **-G** parameter and you specify a public named query with the **-q** parameter, the ARSDOC program queries the application group named with the **-G**

parameter instead of the application group specified in the named query. (If you do not specify the **-G** parameter, then the query runs against the application group specified in the named query. If the named query does not identify an application group, then the query runs against all of the application groups that can be searched from the folder.)

-h *instance*

The name of the Content Manager OnDemand instance to process.

This is a required parameter.

Important: If you are running multiple instances of Content Manager OnDemand on the same workstation, always specify the **-h** parameter to identify the name of the instance that you want to process. Verify that the system is configured with the correct information for all instances of Content Manager OnDemand.

-H For the QUERY function, specify this parameter to generate a header record in the output. The default header record contains the application group field names. This parameter also generates a line that contains the names of the database fields. By default, the field names are delimited with the comma character. You can specify a delimiter of your choice with the **-e** parameter.

You can use the **-H** parameter to generate output that contains only the application group database field names. To do so, specify the **-H** parameter without specifying the **-i** or **-q** parameters. (You also must not specify the **-L**, **-n** or **-N** parameters.)

The ARSDOC program writes the database field names to the specified output file, or to the display (interactive) or the job log (batch).

-i *sqlQuery*

A valid SQL query, that includes the names of one or more application group database fields, index values, and operators. Content Manager OnDemand does not validate the string that you specify. See the SQL reference for your database manager product for an overview of SQL concepts and details about how to construct a query.

Restriction: If you specify the **-q** or the **-X** parameters, then you cannot specify the **-i** parameter.

For the DELETE or UPDATE functions, if more than one document meets the search criteria, then multiple documents will be deleted or updated. For an update, all of the documents will be updated with the same values.

To construct a query with a database field of type date, you must use the Content Manager OnDemand internal format of the date. That is, the number of days since January 1, 1970. You can use the ARSDATE program to list the internal format for a given date string. For example, the following shows how to use the ARSDATE program to obtain the internal date for July 21, 1995:

The ARSDATE program displays:

```
7/21/95 -> 9333
```

You would then enter 9333 as the index value for the date field.

-l *holdname*

For the GET, QUERY, PRINT, HOLD_ADD, and HOLD_RELEASE functions, this parameter specifies the hold name. Do not specify this

parameter when you specify the **-X** parameter. Do not specify this parameter if only database names are specified. Specifying the hold name limits the list of returned hits to only those hits that have been added to the specified hold. For example, an SQL query or a Named Query produces 10 hits. If 2 of the 10 hits have been added to a hold and the hold name is provided, the result contains only the two hits that are included in the hold.

-L *max#*

For the GET and PRINT functions, determines the maximum number of items retrieved from Content Manager OnDemand, regardless of the number of items that match the query.

For the QUERY function, determines the number of items included in the hit list, regardless of the number of items that match the query.

-n

For the GET function, use to retrieve items one at a time from the server. By default, the ARSDOC GET function uses a bulk retrieval method for high-speed retrieval of items from the server.

Important: The **-n** parameter is required when you work with data whose indexes are migrated from Spool File Archive and is still managed by Spool File Archive by using the Report Management Cycle (RMC).

Tip: Specify the **-n** and **-c** parameters if a sort order has been defined in the folder and it is a requirement that the documents be retrieved in the order specified by the sort order. If the **-c** parameter is not specified, the sort order will not be used when the documents are written to individual files. For more information about the sort order, see the online help on the Field Information page for the folder.

For the QUERY function, use to number the items in the output file. If you specify this option, the ARSDOC program sequentially numbers each line in the output file, beginning with 1 (one).

For the ADD and UPDATE functions, use to specify the application group database field names and their values using the form **-n dbfield=value**.

- Specify a null (blank) field value by using single quotes within double quotes. For example: **-n middle=""**
- Specify a string field value that contains a null (blank) or other special character by enclosing the field value in single quotes within double quotes. For example: **-n name=""Sally Smith""**

You can specify one or more field names and their values (by specifying the **-n** parameter one time for each database field name and its value). When adding a document, you must specify all of the application group fields unless you specify the **-0** parameter. When updating a document, you can specify one or more fields and their values. For a **date** field, you must specify the value using the Display Format from the Field Information page under folders.

-N "(dbfield1)(dbfield2)(dbfieldn)"

For the QUERY function, specify the order and names of the database fields to include in the output. For the GET function, when querying a folder that searches more than one application group or a folder that searches an application group that contains more than one application, specify this parameter to add the resource identifier, application group name, and application name to the output file name. When you specify the

-N parameter, you must specify the **-c** parameter. If you specify the **-g** parameter to generate generic index data, you must specify the **-N** parameter.

If the folder searches more than one application group or an application group contains more than one application and you do not specify the **-N** parameter, then the ARSDOC program adds the application group or application identifier to the output file name. For example, the following specification:

```
-o student -c
```

Can result in output file names such as: student.516 or student.517

Where 516 and 517 are application group identifiers. However, when you specify the **-N** parameter, the ARSDOC program uses the resource identifier, application group name, and application name to name the output file. For example, the following specification:

```
-o student -c -N
```

Can result in output file names such as: student.1.BILLS.1995 or student.1.BILLS.1996

Where 1 is the resource identifier, BILLS is the application group name, and 1995 and 1996 are application names.

The number of index files created is dependent on the number of application groups in a folder, the number of applications in an application group, and the number of resource groups in an application.

For the QUERY function, determines the application group fields that the ARSDOC program writes to the output file and the field names that appear in the header record. By default, the ARSDOC program writes all fields to the output file. You can specify one or more application group field names using the form **-N(dbfield)...(dbfield)**. Each field name that you specify must be delimited with parenthesis. When you run a query from the command line, you must delimit the entire string in double quote characters. For example, **-N"(dbfield)...(dbfield)"**.

-o name

For the GET function, use to write documents to one or more files and identify a user-defined string used to generate unique file names. For example, the following specification:

```
-o student -c
```

Can result in the following output file name: student

You can concatenate one or more of the database field names that you specify with the **-i** parameter to generate a unique file name. For example, the following specification:

```
-o "(sdate)(student)"  
-i "WHERE sdate='971025' AND student='001200340056'"
```

Can result in the file name: 971025.001200340056

When you use database field names to generate a unique file name:

- Content Manager OnDemand verifies that the field names that you specify are valid for the application groups that can be searched by the folder specified with the **-f** parameter.

- If the field name that you specify is a date field, the output format of the date is determined by the Format on the Load Information page under applications.
- The field names must be delimited with parenthesis.
- You can specify the fields in any order. The order that you specify determines the file name that the ARSDOC program generates.
- You cannot use a field name to represent a directory name. For example, the following is not valid:
-o "(field_1)/(field_2)
- You cannot specify the **-c** parameter to concatenate items in one output file. Each item that matches the query is stored in a separate output file.

If more than one item matches a query and you do not generate a unique file name using database field names, concatenate items in a single file with the **-c** parameter, or specify the **-g** parameter, then the ARSDOC program generates a unique file name for each item that matches the query by adding a *.n* extension to the file name. Where *n* is the number of the item that matched the query. For example, if you specify:

```
-o statements
```

And two items match the query, the ARSDOC program creates the following files: statements.1 and statements.2

You must specify the **-o** parameter when you specify the **-c** parameter.

For the QUERY function, determines the file name of the output file in which the ARSDOC program writes the list of items that match the query.

For the ADD function, determines the name of the input file that contains the document to be added. The value that you specify is not checked for valid characters. You can specify a full path name, including the back slash or forward slash characters that are part of a directory path. When adding a document, you can provide the input data by specifying the name of the input file that contains the data with the **-o** parameter, an SQL query with the **-i** parameter, or a public named query with the **-q** parameter. Only one document may be added at a time.

-O For the ADD function, you must specify this parameter if you intend to omit one or more database fields. However, you can never omit the following fields:

- Date
- Date (old style)
- Date/Time
- Date/Time (old style)
- Date/Time (TZ)
- Date/Time (TZ) (old style)

When you specify the **-O** parameter, the ARSDOC program stores a default value in any other database field that you omit. The default value for string fields is an empty (null) string. The default value for numeric fields is 0 (zero). Numeric fields include integer and decimal fields.

The **-O** option is not required in the case where an index row is being added for an existing document. For any index values that are not provided, the values from the existing document are used.

-p *password*

Specify one of the following options for *password*:

- The name of the stash (encrypted password) file that contains the password for the user ID specified with the **-u** parameter.
- The password of the Content Manager OnDemand user that you named with the **-u** parameter. If there is no password assigned to the user that you specify, use quotation marks to show a null password. That is, specify **-p ""**. If you omit the **-p** parameter, then the ARSDOC program retrieves the password for the user ID from the ARS.INI file for that instance. If there is no password assigned to the user that you specify, press the Enter key when prompted.

If you omit **-u** and **-p** parameters, Content Manager OnDemand uses the IBM i user profile of the user issuing ARSDOC as the Content Manager OnDemand user ID.

-P Indicates PDF files that are retrieved and should be stored in individual files even if concatenation has been requested.

-P *printer*

For the PRINT function, identifies the Content Manager OnDemand server printer to which you want to send the documents that match the query.

-q *namedQuery*

The name of a public named query for the folder named with the **-f** parameter. A named query is a set of search criteria previously saved on the library server that can be recalled by name to search a folder. A named query is typically defined to search a folder for a specific document or set of documents.

Restriction: If you specify the **-i** or the **-X** parameters, then you cannot specify the **-q** parameter.

-Q *SQLqueryfile*

Use this parameter to specify a file name that contains one or more query strings. Specify only one of the parameters, **-i**, **-q**, or **-Q**.

-s *seconds*

For the GET function, determines the number of seconds that the ARSDOC program waits between query requests when you specify more than one query with the **-F** parameter. If you do not specify this option, then the ARSDOC program does not wait between query requests. That is, the default is 0 (zero) seconds.

-t *search_str*

Search for the string *search_str* using the Content Manager OnDemand search.

-T *search_str*

Search for the string *search_str* using the full text index search.

-u *userid*

The Content Manager OnDemand user that is permitted to perform the specified function. The ARSDOC program verifies the following permissions:

- the user ID that you specify is a valid Content Manager OnDemand user for the library server that you name with the **-h** parameter
- the user ID is permitted to open the folder that you name with the **-f** parameter

- the user ID has permission in application groups to perform the specified function

If you omit the **-u** parameter, then the ARSDOC program prompts you to enter the user ID when you run the program.

If you do not specify the **-p** parameter with the **-u** parameter, the ARSDOC program retrieves the password for the user ID from the ARS.INI file for that instance.

-U *user_alias*

Identifies the users when multiple users share a common user ID. The maximum length for *user_alias* is 128.

-v Enables verbose mode, which displays all messages (informational and error). By default, the ARSDOC program displays error messages.

-x *loadId*

For the GET function, use to limit the documents that can be retrieved to the set of documents that were loaded into the system under the specified *loadId*.

For the QUERY function, use to limit the query to the set of documents that were loaded into the system under the specified *loadId*.

When you specify the **-x** parameter, use the **-f** and **-G** parameters as follows:

- Specify the **-f** parameter to search all application groups. You can specify the search using the **-i** parameter or the **-G** parameter.
- Specify the **-g** parameter to search a specific application group. You must specify the search using the **-i** parameter.
- Specify both the **-f** parameter and the **-G** parameter. The ARSDOC program verifies that the application group can be searched from the folder.

When you specify the **-x** parameter, you cannot specify the **-X** parameter.

-X *loadId*

For the GET function, retrieve documents by using the index file that was generated for the specified *loadId*.

Important: Content Manager OnDemand does not retrieve any documents with index values that you updated.

For the QUERY function, build a hit list from the index file that was generated for the specified *loadId*.

For the PRINT function, the **-i SQL** parameter can be specified when the **-X** flag is used. If it is specified, it is used only if the retrieve fails using the *loadID*. If the **-X** flag is used, the application group name must be provided using the **-G** flag.

For the HOLD_ADD function, add documents to a hold by using the index file that was generated for the specified *loadId*.

For the HOLD_RELEASE function, remove documents from a hold by using the index file that was generated for the specified *loadId*.

For the CFSOD-FED function, send documents to CFS-OD and make them available to IBM FileNet P8 features by using the index file that was generated for the specified *loadId*.

When you specify the **-X** parameter, you must specify the **-G** parameter and name the application group.

When you specify the **-X** parameter, you cannot specify the **-x** parameter, or the **-i**, **-q**, **-S**, and **-f** parameters.

Examples

1. The following shows how to use the GET function to retrieve documents and save them in a file in the current directory.

```
arsdoc get -h QUSROND -f "Student Information" -o student
-c -S 1/1/97,12/31/97 -i "WHERE student='001200340056'" -v
```

The ARSDOC program saves all of the documents that match the query in the following output file: student

2. The following shows how to use the GET function and a parameter file to run more than one query at a time. The parameter file is in the current directory. The output files are saved in the current directory.

```
arsdoc get -h QUSROND -F parmfile -v
```

The parameter file contains two queries:

```
[-f "Student Information"] [-S 1/1/97,12/31/97] [-o (student)(type)] \
[-i "WHERE student='123420010056' AND type='B' OR type='G' OR type='T'"]
```

```
[-f "Student Information"] [-S 1/1/97,12/31/97] [-o (student)(type)] \
[-i "WHERE student='123450011917' AND type='B' OR type='G' OR type='T'"]
```

Assuming that documents exist for all of the specified types, the ARSDOC program creates the following output files:

- 123420010056.Bills
 - 123420010056.Grades
 - 123420010056.Transcripts
 - 123450011917.Bills
 - 123450011917.Grades
 - 123450011917.Transcripts
3. The following shows how to use the GET function to retrieve documents, write the documents to a file, and generate and save generic index data for the documents that match the query. The example shows how to specify the name of a public named query that is valid for the specified folder.

```
arsdoc get -h QUSROND -f "Student Information"
-a -c -g -o student -q "3rd yr students GPA>3.5" -N -v
```

The number of output files that the ARSDOC program generates is a factor of the number of application groups queried, the applications contained in the application groups, whether the data is AFP and if so, the versions of resource groups in each application. At a minimum, for AFP data with one version of the resource group and one application group, the ARSDOC program generates three output files. For example:

- student.1.Student Information.TRANSSCRIPTS.ind
- student.1.Student Information.TRANSSCRIPTS.out
- student.1.Student Information.TRANSSCRIPTS.res

For AFP data, if there is more than one version of the resource group per application, then the ARSDOC program can generate additional output files. For example:

- student.1.Student Information.TRANSSCRIPTS.ind
- student.1.Student Information.TRANSSCRIPTS.out
- student.1.Student Information.TRANSSCRIPTS.res
- student.2.Student Information.TRANSSCRIPTS.out
- student.2.Student Information.TRANSSCRIPTS.ind
- student.2.Student Information.TRANSSCRIPTS.res

If the application group contains more than one application, then the ARSDOC program can generate additional output files. For example:

- student.1.Student Information.Bills.ind
- student.1.Student Information.Bills.out
- student.1.Student Information.Bills.res
- student.2.Student Information.Grades.out
- student.2.Student Information.Grades.ind
- student.2.Student Information.Grades.res

4. The following shows how to use the QUERY function to generate a list of items and save the list in a file in the current directory.

```
arsdoc query -h QUSROND -f "Student Information" -o query1.out -H
-S 1/1/97,12/31/97 -i "WHERE student='0012-0034-0056' AND type='B' OR
type='G' OR type='T'" -v
```

5. The following shows how to use the QUERY function and a parameter file to run more than on query at a time. The parameter file is in the current directory. The output files are saved in the current directory.

```
arsdoc query -h QUSROND -F parmfile -v
```

The parameter file contains three queries:

```
[-f "Student Information"] [-i "WHERE type='B'"] [-o query2.out] \
[-S 1/1/97,12/31/97] [-H] [-N (student)(id)(p_date)]
```

```
[-f "Student Information"] [-i "WHERE type='G'"] [-o query3.out] \
[-S 1/1/97,12/31/97] [-H] [-N (student)(id)(p_date)]
```

```
[-f "Student Information"] [-i "WHERE type='T'"] [-o query4.out] \
[-S 1/1/97,12/31/97] [-H] [-N (student)(id)(p_date)]
```

6. The following example shows how to use the QUERY function to search a specific application group:

```
arsdoc query -h QUSROND -i "where Date_Taken BETWEEN 9863 AND 11531"
-G load-scanned-images-jpeg -o jpeg -v
```

7. The following example shows how to use the ADD function to add a document to the Credit Card Statements folder:

```
arsdoc add -h QUSROND -o /newdata/crd.dat -n "crd_date=01/21/98"
-n "account='000-000-000'" -n balance=123.45 -n "name='John Watpole'"
-f "Credit Card Statements" -g CRD -a CRD -v
```

8. The following example shows how to use the UPDATE function to update a document in the Credit Card Statements folder, changing the balance from 123.45 to 0.00:

```
arsdoc update -h QUSROND -i "where account='000-000-000' and
name='John Watpole'" -n "balance=0.00" -S 1/1/97,12/31/97 -f
"Credit Card Statements" -g CRD -v
```

9. The following example shows how to use the DELETE function to delete a document from the Credit Card Statements folder:

```
arsdoc delete -h QUSROND -i "where account='000-000-000' and
name='John Watpole'" -f "Credit Card Statements" -S 1/1/97,12/31/97
-v
```

10. The following example shows how to use the PRINT function to send the documents that match a query to a server printer:

```
arsdoc print -h QUSROND -P
svrprt1 -i "where account='000-000-000' and name='John Watpole'"
-f "Credit Card Statements" -S 1/1/97,12/31/97 -v
```

11. You can use the ARSDOC program to use the ADD function to add an index that points to an existing document. For example, assume that you loaded a bank statement for account number 000-000-000, date 5/23/97, and account name Joe Smith. You want to add a new index, but point to the existing statement. The new index uses the same account number and date, but contains a different account name (for example, Sally Smith). After adding the index, if a query is run with account name Joe Smith or Sally Smith, the same bank statement will be retrieved. To add an index for an existing document:

```
arsdoc add -h QUSROND -i "where sdate=10005 and account='000-000-000'
and name='Joe Smith'" -n "sdate=5/23/97" -n "account=000-000-000" -n
"name='Sally Smith'" -f "Credit Card Statements" -S 1/1/97,6/31/97 -v
```

12. The following example shows how to use the QUERY function to limit a search to the documents that were loaded into the system under a specific load ID. For example, assume that the specified folder could be used to search several application groups; each application group contains more than one application; there are 500,000 documents in the application groups. By using the **-x** parameter, the query will be limited to the set of documents that was loaded into the system under the specified load ID. Without the **-x** parameter, the query is run against all 500,000 documents.

```
arsdoc query -h QUSROND -x 8495-53-0-1FAA-20130627000000-20130627000000-8496
-f load-scanned-images -q query -o out -v
```

13. The following example shows how to use the QUERY function to limit a search to the documents that were loaded into the system under a specific load ID. In the example, the load ID, application group name, and query string are provided. The search is limited to the application group and only those documents that were loaded into the system under the specified load ID.

```
arsdoc query -h QUSROND -x 8495-53-0-1FAA-20130627000000-20130627000000-8496
-i "where Date_Taken BETWEEN 9863 AND 11531" -G load-scanned-images-jpeg
-o jpeg -v
```

14. The following example shows how to use the QUERY function to limit a search to the documents that were loaded into the system under a specific load ID. In the example, the load ID, folder name, application group name, and named query are provided. The search is limited to the application group and only those documents that were loaded into the system under the specified load ID. Because a folder was specified, a named query can be used. (If an application group name is specified and a folder name is not specified, a named query cannot be used, because a named query is associated with a folder.)

```
arsdoc query -h QUSROND -x 8495-53-0-1FAA-20130627000000-20130627000000-8496
-f load_scanned-images -q query -G load-scanned-images-jpeg -o jpeg -v
```

15. The following example shows how to use the GET function and an index file to retrieve documents from the system. The index file was generated for the set of documents that was loaded into the system under the specified load ID. The example also shows the use of a load ID format prior to version 7.2. When using the **-X** parameter, the database is not queried; rather, documents are retrieved based on the information in the index file.

```
arsdoc get -h QUSROND -X 19867-025-0-3FAA-10136-10136
-G load-scanned-images-jpeg -o jpeg -v
```

16. **-A** parameter: The following are two examples of using the ARSDOC GET function with the **-A** parameter.

Tip: In the following examples, the **-u** and **-p** parameters are not required for IBM i.

- To retrieve public text annotations:

```
arsdoc get -u oduser -p passwd -h odserver -f "CRD" -q named_query
-o loaddata -a -g -c -N -A 0 -v
```

- An example of the messages that are generated by the ARSDOC program. New messages are highlighted:

```
03/10/2004 10:32:12: Starting arsdoc. Version: 8.4.1.0
03/10/2004 10:32:14: arsdoc get -u oduser -h odserver -f CRD -q named_query -o loaddata -a -g -c -N -A 0 -v
03/10/2004 10:32:14: Attempting login for userid 'oduser' on server 'odserver' ...
03/10/2004 10:32:20: Login successful
03/10/2004 10:32:20: Searching for folder 'CRD' ...
03/10/2004 10:32:27: Search successful
03/10/2004 10:32:27: Searching for documents in 'CRD' ...
03/10/2004 10:32:28: Querying database with SQL string 'where account = '000-000-000''
03/10/2004 10:32:56: Search successful
03/10/2004 10:32:56: 1 document(s) have been queried. Retrieving 1 document(s).
03/10/2004 10:34:05: (1): Retrieving document for userid 'oduser' ...
03/10/2004 10:34:05: Document successfully retrieved and stored in file 'loaddata.2.CRD.CRD.out'
03/10/2004 10:34:15: Writing generic indexer file(s).
03/10/2004 10:34:46: 1 annotations were written to file 'loaddata.2.CRD.CRD.ann'
03/10/2004 10:34:46: A total of 1 annotations were written to file 'loaddata.2.CRD.CRD.ann'
03/10/2004 10:34:53: Generic indexer file 'loaddata.2.CRD.CRD.ind' has been successfully created.
03/10/2004 10:34:54: arsdoc completed.
```

17. The following example is for the **-B** option. You want the command line to print all of the documents in the Credit Card Statements folder that have an account number of 000-000-000. The documents are to be printed in sorted ascending order based on the balance:

```
arsdoc print -u oduser -p odpasswd -h odserver -i "where account = '000-000-000'"
-f "Credit Card Statements" -P odprinter -B balance -v.
```

18. **-Q** parameter: The following is an example command and an example file that contains SQL statements.

Example command:

```
arsdoc get -h QUSROND -f "Labor Reports" -Q /home/dbryant/QSTRING.TXT -d /home/dbryant -o
REPORTS -a -g -c -N -v
```

Example file containing SQL statements:

```
# Query on one line
where EMPLNAME = 'B ROCKER'
# Query on multiple lines
where EMPLNAME\
= 'M VESPA'
```

19. **HOLD_ADD** function:

Tip: In the following example, the **-u** and **-p** parameters are not required for IBM i.

```
arsdoc hold_add -u oduser -p odpasswd -h odserver -l hold_audit -f
"Monthly Status Reports" -i "where code='TX' and sdate=14117" -v
```

20. **HOLD_RELEASE** function:

Tip: In the following example, the **-u** and **-p** parameters are not required for IBM i.

```
arsdoc hold_release -u oduser -p odpasswd -h odserver -l hold_audit -f
"Monthly Status Reports" -i "where code='TX' and sdate=14117" -v
```

21. CFSOD_FED function:

```
arsdoc cfsod_fed -u oduser -p odpasswd -h odserver -f
"Monthly Status Reports" -i "where code='TX' and sdate=14117" -v
```

22. FTL_ADD function:

```
arsdoc fti_add -u oduser -p odpasswd -h odserver -f
"Monthly Status Reports" -i "where code='TX' and sdate=14117" -v
```

Notes

The ADD function will fail unless the Database Organization for the application group named with the **-g** parameter is Multiple Loads per Database Table and the Expiration Type is Segment or Document.

The ADD function can be run without providing an input document from a file or by retrieving an existing document from the system. This means that you can add database field values without adding a document. To add database field values without adding a document, do not specify the **-o**, **-i**, or **-q** parameters; specify the database field names and their values using one or more **-n** parameters.

The ARSDOC program can print the PTF version number and the ARSDOC GET function can print the number of documents that were queried and retrieved and print a status message for each document that is retrieved. To enable the messages, specify the **-v** parameter.

The following shows an example of the new messages:

```
12/08/02 10:33:36: Starting arsdoc. Version: 7.1.0.12
12/08/02 10:33:36: arsdoc get -u admin -h instance -q 000-000-000 -f CC Stmts -o test -v -c
12/08/02 10:33:36: Attempting login for userid 'admin' on server 'instance' ...
12/08/02 10:33:37: Login successful
12/08/02 10:33:37: Searching for folder 'CC Stmts' ...
12/08/02 10:33:38: Search successful
12/08/02 10:33:38: Searching for documents in 'CC Stmts' ...
12/08/02 10:33:40: Search successful
12/08/02 10:33:40: 2 document(s) have been queried. Retrieving 2 document(s).
12/08/02 10:33:41: (1): Retrieving document for userid 'admin' ...
12/08/02 10:33:41: Document successfully retrieved and stored in file 'test'
12/08/02 10:33:42: (2): Retrieving document for userid 'admin' ...
12/08/02 10:33:42: Document successfully retrieved and stored in file 'test'
12/08/02 10:33:45: arsdoc completed.
```

IFS Location

/usr/bin/arsdoc

The IBM i executable program.

ARSLOAD


Purpose

The ARSLOAD program can be used to process the input files that you want to load into the system. The ARSLOAD program determines if the input data needs

to be indexed, and if so, calls the appropriate indexing program. The ARSLOAD program calls the storage manager programs to load report data on storage volumes and the database manager to update the IBM Content Manager OnDemand database with the index information that was extracted from or specified for the input file.

Important: Set the proper locale before issuing the ARSLOAD API. See the chapter entitled “Defining a locale” in the *IBM Content Manager OnDemand for i Common Server: Planning and Installation Guide* for important details.

Syntax



Description

The ARSLOAD program is the main IBM Content Manager OnDemand data indexing and loading program. The ARSLOAD program calls the indexing program if the input data needs to be indexed, creates input files for the storage manager, and updates the Content Manager OnDemand database. The ARSLOAD program saves processing messages in the system log. You can open the System Log folder and list the messages that were generated when an input file was processed.

You typically configure the ARSLOAD program to run as a monitor to periodically check specified IFS directories for input files to process.

The ARSLOAD program can use the following sources for input files to process:

- one or more file systems specified with one or more **-d** parameters,
- one or more load file names.

If you omit the load file name, the ARSLOAD program will run in monitor mode and attempt to load input data from the directories that are specified by the **-d** parameter. If you omit the load file name and do not specify the monitor mode parameter (**-d**), the ARSLOAD program will issue a usage note and exit.

Important: When running the ARSLOAD program in monitor mode, the **.ARD** or **.PDF** file name extensions are required to initiate a load process.

When you run the ARSLOAD program, you must provide the user ID of a Content Manager OnDemand user with administrator authority for the application group into which the input data will be loaded. There are several ways that you can provide the user ID:

- Do not specify the **-u** parameter. In this case, the current user profile of the job that is running ARSLOAD is used as the Content Manager OnDemand user ID. This is the recommended method.
- Use the **-u** parameter each time that you run the ARSLOAD program.
- Use the stash file.

Parameters

-a application

The name of the application to load. If the application group contains more than one application, then you must identify the application to load. Otherwise, the load will fail. If you plan to automate the loading of files into different application groups and applications, then use the **-A** parameter to specify the part of the file name that identifies the name of the application.

-A applID

Determines the part of the file name that identifies the application to load. If the application group contains more than one application, then you must identify the application to load; otherwise, the load will fail.

You typically use this parameter when you run the ARSLOAD program as a monitor to automate the loading of files into different application groups and applications. For example, a file transmitted from an OS/390 or z/OS system might use the following file naming convention:

```
MVS.JOBNAME.DATASET.FORMS.YYYYDDD.HHMMSS.ARD
```

Important: The **.ARD** file name extension is required to initiate a load process.

- Unless you specify otherwise, the ARSLOAD program uses the **FORMS** part of the file name to identify the application group to load. You can use the **-g** parameter to specify a different part of the file name that identifies the application group (**MVS™**, **JOBNAME**, or **DATASET**). For example, `arsload -G JOBNAME`.
- If the application group to load contains more than one application (source of data), then you must identify the application to load; otherwise, the load will fail. When you run the ARSLOAD program, you can use the **-A** parameter to specify the part of the file name that identifies the application (**MVS**, **JOBNAME**, **DATASET**, or **FORMS**). For

example, to use the DATASET part of the file name to identify the application, run the ARSLOAD program with the -A DATASET parameter.

-b "index_number,field"

Specify the name of an index field and a number that identifies the part of the input file name that the ARSLOAD programs assigns to that index field. For *index_number*, you can specify 1 - 8 and the number must match a part of the input file name identified by the **-B** parameter.

In the following example, assume that the input file name is App1Group1.App1.ARD:

```
arsload -b "1,field1" -B "AG.1APP.ARD" -g App1Group1 ...
```

During the load, the ARSLOAD program assigns the value App1 from the input file name to the index field "field1" of application group "App1Group1".

-B "format"

Use this parameter to do the following tasks:

- Define the file name formats for MVS download files and files processed by the ARSLOAD daemon.
- Identify the parts of the file name that the ARSLOAD program assigns to an index field specified by the **-b** parameter.

Use the following identifiers to define the file name format:

1 - 8 A number prefixed to a part of the file name that identifies that part as a value to assign to an index field.

AG Application group name

APP Application name

DAT Abbreviation for DATASET

EXT Extension

FLD Field place holder for indexing

FOR Abbreviation for FORMS

HH Abbreviation for HHMMM

IGN Ignore

JOB Abbreviation for JOBNAME

MVS MVS

WRI Abbreviation for WRITER

YY Abbreviation for YYDDD

The following example uses the full word of the JOB, DAT, FOR, YY, and HH identifiers to define the file name format:

```
MVS.JOBNAME.DATASET.FORMS.YYDDD.HHMMM.ARD
```

The following example shows the same file name format with the abbreviated words:

```
MVS.JOB.DAT.FOR.YY.HH.EXT
```

The following example shows the same file name format and indicates that the ARSLOAD program assigns the JOB part of file name to index field 1:

```
MVS.1JOB.DAT.FOR.YY.HH.EXT
```

The following examples show different ways that you can define the file name formats:

```
ARSLOAD -G FORMS -D "-" -B "WRITER-MVS-IGN-IGN-FORMS"  
ARSLOAD -G JOBNAME -B "WRITER.IGN.JOB.IGN.DAT.IGN.EXT"  
ARSLOAD -D "-" -B "IGN-IGN-AG-APP-IGN.EXT"  
ARSLOAD -B "APP.IGN.AG.IGN.IGN"
```

The following examples show how to identify the parts of the file name to assigns to an index field:

- Single index format, where the ARSLOAD daemon assigns the DAT part of the file name to index field `account_num`:
`arsload -b "1,account_num" -B "AG.APP.1DAT"`
- Multiple index format, where the ARSLOAD daemon assigns the different parts of the file name to different index fields:
`arsload -b "1,report_name" -b "2,run_date" -b "3,app_name" -B "AG.3APP.1FLD.2YY"`

The following table shows the assignments:

Index field name	Part of the file name
report_name	FLD
run_date	YY
app_name	APP

You can use the `-B` parameter or combine the `-g` and `-a` parameters, but you cannot use all three of them together. If you specify the `-B` parameter, you identify the application group name and the application name as parts of the file name.

-c indexDir

The file system in which IBM Content Manager OnDemand temporarily stores data created by the indexing program. The default location is the directory from which the ARSLOAD program was invoked.

-d dataDir

The directory that contains input files to process.

Any file with a file type extension of `.ARD` or `.PDF` will be processed (`.ARD` files are transmitted to the server by an OS/390 or z/OS download utility; `.PDF` files are created by Acrobat Distiller). The case of the file type extension is not significant.

You can specify this parameter one or more times. The ARSLOAD program will search for input files to load in each of the directories that you specify.

To specify more than one directory, specify the `-d` parameter multiple times. In the following example:

```
arsload ... -d dir1 -d dir2 -d dir3...
```

the ARSLOAD program will search for input files in the `dir1`, `dir2`, and `dir3` directories.

-f

Use to unload the data if the load process fails. If the database manager step fails, then Content Manager OnDemand should remove any index data that was added to the database. If the storage manager step fails, then Content Manager OnDemand should remove any storage objects that were copied to storage volumes.

Important: If an input file fails to load, you should review the message log that was created during the load process. You can retrieve a message log from the system log. If the message log contains a Load ID, then it means that for some reason, Content Manager OnDemand stored at least some of the input data in the application group. Before you attempt to reload the input data, you must remove the data that was created during the failed load process by using the RMVRPTOND command.

See “Deleting a report” on page 139 for help with removing the data that is created when a file is loaded into the system.

-F Trace output format

Optional parameter when `-1` is specified. The default output is text format. The possible values are C, T, and X:

C or CSV

The CSV format is a comma separated file that is useful in spreadsheets.

T or TEXT

Default output format.

X or XML

The XML format outputs XML structured data. Using the XML output requires a header and trailer to be added to the trace file to complete the XML syntax and create valid XML. After the XML structured data is combined with .XSL and .DTD files, it can be viewed in a browser or XML editor.

-g applGroup

The name of the application group to load. This parameter is required if you specify a load file name to process. This parameter is optional if you specify the `-d` parameter. However, if you specify the `-d` parameter, unless you specify otherwise, the ARSLOAD program uses the FORMS part of the file name to determine the name of the application group to load. If you plan to automate the loading of files into different application groups and applications, then you should use the `-g` parameter to specify the part of the file name that identifies the application group to load.

-G applGroupID

Determines the part of the file name that the ARSLOAD program uses to identify the name of the application group to load.

You typically use this parameter when you run the ARSLOAD program as a monitor to automate the loading of files into different application groups and applications. For example, a file transmitted from an OS/390 or z/OS download utility might use the following file naming convention:

```
MVS.JOBNAME.DATASET.FORMS.YYYYDDD.HHMMSS.ARD
```

Important: The .ARD file name extension is required to initiate a load process.

- Unless you specify otherwise, the ARSLOAD program uses the FORMS part of the file name to identify the application group to load. You can use the `-g` parameter to specify a different part of the file name that identifies the application group (MVS, JOBNAME, or DATASET). For example, `arsload -G JOBNAME`.
- If the application group to load contains more than one application (source of data), then you must identify the application to load; otherwise, the load will fail. When you run the ARSLOAD program, you

can use the **-A** parameter to specify the part of the file name that identifies the application (MVS, JOBNAME, DATASET, or FORMS). For example, to use the DATASET part of the file name to identify the application, run the ARSLOAD program with the **-A DATASET** parameter.

-h instance

The name of the Content Manager OnDemand instance to process. This is a required parameter.

Important: If you run multiple instances of Content Manager OnDemand, always specify the **-h** parameter to identify the name of the instance that you want to process. Verify that the system is configured with the correct information for all instances of Content Manager OnDemand. See *IBM Content Manager OnDemand for i: Planning and Installation Guide* for information on configuring instances.

-i Use to run the data indexing program only; do not copy report data to storage volumes or add the index data to the database.

-I instance

The name of the Content Manager OnDemand instance that you want to process. The name of the default instance on the IBM i is QUSROND. You must specify the **-I** parameter and name the instance, if you are running more than one instance on the same IBM i system and you want to process an instance other than the default instance (QUSROND).

-j parmFile

Use to specify the name of a parameter file that contains additional indexing parameters.

When you specify the **-j** parameter and name a parameter file, the ARSLOAD program adds the indexing parameters from the specified parameter file to the indexing parameters that it extracts from the application. (Indexing parameters are typically specified on the Indexing Information page in the application definition.) If an indexing parameter appears in both the application and the parameter file that you specify, unexpected results may occur.

-J File name delimiter

Use this parameter to define file name formats for MVS download files and files that are processed by the ARSLOAD daemon. By default, this parameter is ".".

-I fieldID

Indicates that the entire name of the input file should be captured and placed into the field identifier specified.

-L Trace level number

Specify a numeric value (0, 1, 2, 3, or 4), as directed by IBM Software Support.

-n Determines whether Content Manager OnDemand deletes the input files when the ARSLOAD program terminates.

- If you specify the **-N** parameter, then Content Manager OnDemand does not delete the input files when the ARSLOAD program ends.
- If you do not specify the **-N** parameter, then Content Manager OnDemand deletes the input files when the ARSLOAD program ends.

In either case, if the ARSLOAD program fails in the load step because of a device or system problem, then you can restart the load step after you

correct the problem by using the intermediate files that were created by the ARSLOAD program. The ARSLOAD program stores the intermediate files in the directory named with the **-c** parameter (or the directory from which you started the ARSLOAD program, if you did not specify the **-c** parameter). The intermediate files have the same file name as the original input file.

-p password

The password for the user that is specified with the **-u** parameter. If the user is not assigned a password, enter a null password (that is, specify **-p ""**).

-t seconds

Determines the polling time in seconds. This is the interval of time in which the ARSLOAD program checks the input directory (specified by the **-d** parameter) for input files to process. The default value is 600 seconds, which means that the ARSLOAD program checks the input directory every ten minutes.

-l Fully qualified trace file name

When a file is specified, trace is activated for ARSLOAD. If the file already exists, it is renamed with the current date and time, and a new file is created.

-u userid

The user ID of a Content Manager OnDemand user with administrator authority for the application group. The user must have permission to add documents to the application group.

-v Enables verbose mode, which displays all messages (informational and error). By default, the ARSLOAD program displays error messages.

-X indexer

Allows you to override the indexing program that was specified on the Indexer Information page for the application. The possible values are 4 and G:

- Specify **-X 4** to use OS/400 Indexer.
- Specify **-X G** to use the Content Manager OnDemand Generic Indexer.
- Specify **-X P** to use the Content Manager OnDemand PDF Indexer.

You must specify an indexing program if the input contains index data in a format other than the one supported by the indexing program that was specified for the application. For example, assume that you defined an application and specified OS/400 Indexer as the indexing program. Later, you used the ARSDOC program to extract documents from an application group. The ARSDOC program generates index data that is in the Generic indexer format. To load the index data into the application, you must specify:

```
arsload -X G . . .
```

loadFilename

Specifies an input file to process.

You may specify the names of one or more input files to process. If you specify more than one input file, separate the file names with a blank character.

The ARSLOAD program concatenates the following file type extensions to the file name that you specify: IND, OUT, and RES

- The IND file contains the index data
- The OUT file contains the report data
- The RES file contains the resource data

For example, if you specify `arsload -g BILLS po3510`, the ARSLOAD program processes the following files:

- `po3510.ind`
- `po3510.out`
- `po3510.res`

If the ARSLOAD program does not find a file with the IND file type extension, it automatically calls the indexing program to process the input file.

If you omit an input file name, the ARSLOAD program will run in daemon mode and attempt to load input data from the directories that are specified by the `-d` parameter. If you omit an input file name and do not specify the daemon mode parameter (`-d`), the ARSLOAD program will issue a usage note and exit.

Examples

1. The following shows how to run the ARSLOAD program to check the specified directory for input files to process. The input files must have a file type extension of `.ARD` or `.PDF`. The ARSLOAD program stores temporary work files in the location specified by the `-c` parameter. In this example, the ARSLOAD program uses the FORMS part of the file name to determine the application group to load and the application group contains only one application; the ARSLOAD program logs on to the system and accesses the application group with the user ID and password stored in the encrypted password file (stash file).

```
arsload -c /arsdir/dir1 -d /arsdir/dir2
```

2. The following shows how to run the ARSLOAD program to check the specified directory for input files to process. The input files must have a file type extension of `.ARD` or `.PDF`. The ARSLOAD program stores temporary work files in the location specified by the `-c` parameter. In this example, the ARSLOAD program uses the JOBNAME part of the file name to determine the application group to load and the DATASET part of the file name to determine the application to load; the ARSLOAD program logs on to the system and accesses the application group with the user ID and password stored in the encrypted password file (stash file).

```
arsload -c /arsdir/dir1 -d /arsdir/dir2 -A DATASET -G JOBNAME
```

3. The following shows how to run the ARSLOAD program to load the specified file into the specified application group. The ARSLOAD program logs on to the system and accesses the application group by using the user ID and password from the `-u` and `-p` parameters.

```
arsload -g BILLS -u bob P03510
```

4. The following shows how to run the ARSLOAD program to load several input files into the specified application group. The application group name contains an embedded blank character, and must be quoted. The ARSLOAD program logs on to the system and accesses the application group by using the user ID and password from the `-u` and `-p` parameters.

```
arsload -g "ABC Credit" -u bob RW7505 RW8505
```

Notes

The IBM Content Manager OnDemand server job must be running, otherwise the ARSLOAD program will fail.

IFS Location

`/usr/bin/arsload`

The IBM i executable program.

ARSSLU

Purpose

The ARSSLU command helps you audit your Content Manager OnDemand database for inconsistencies. Run this command only on the library server.

Syntax

```
▶ arsslucdFformat C T X -g Application group name
▶ -I od_inst -o Output file name -p -v -1 trace_file -2 level
```

Parameters

-a Audit the OnDemand database

This option checks the Content Manager OnDemand database for inconsistencies or potential problems.

-c Create message table and index

Creates a table and index for load data information which is extracted from the Content Manager OnDemand system log. This data is used for auditing purposes.

-d Drop message table and index

Deletes the table and index that are created with the `-c` parameter and any data contained. Do not combine this parameter with any other parameter except for `-v` and `-I`.

-F <format> Trace output information

Optional parameter when `-1` is specified. The default output is text format. The possible values are C, T, and X:

C or CSV

The CSV format is a comma separated file that is useful in spreadsheets.

T or TEXT

Default output format.

X or XML

The XML format outputs XML structured data. Using the XML output requires a header and trailer to be added to the trace file to complete the XML syntax and create valid XML. After the XML structured data is combined with .XSL and .DTD files, it can be viewed in a browser or XML editor. You can obtain these needed files from: <ftp://service.software.ibm.com/software/ondemand/utills/trace>

- | **-g <Application group name>**
| Specifies an application group name. If the parameter is not specified, the
| default is all application groups. This parameter is used with the -a
| parameter.
- | **-I <od_inst>**
| Specifies the database instance name.
- | **-L <level>**
| Specify a numeric value (0, 1, 2, 3, or 4), as directed by IBM Software
| Support.
- | **-o <Output file name>**
| This parameter is optional. Any output is sent to the console if it is not
| specified. This option provides a file name for any log output that is
| produced by the utility when you use the -a parameter. Output to a file is
| in XML format. Output to the console is in a human readable plain text
| format.
- | **-p Populate message table**
| Extracts load information from the System Log and populates the table
| created with the -c parameter. This parameter is used for auditing
| purposes.
- | **-l <trace_file> Fully qualified trace file name**
|
| When a file is specified, trace is activated for ARSSLU. If the file already
| exists, it is renamed with the current date and time, and a new file is
| created.
- | **-v Verbose**
| Prints additional status messages to the console.

Examples

Here are several examples of using the ARSSLU command:

```
arssl -I ARCHIVE -c -v
Creating table odadmin.arssl
Creating index odadmin.arssl.arsapp_idx
Updating runstat statistics for table odadmin.arssl

arssl -I ARCHIVE -p -v
System log data tables: 1
Table name: SL2 Messages present: Yes
Processing table SL2
Total log messages: 333
Log messages processed: 330
Log messages failed: 0
Log messages skipped: 3
Reorganizing table odadmin.arssl
Updating runstat statistics for table odadmin.arssl

arssl -I ARCHIVE -d -v
Dropping index odadmin.arssl.arsapp_idx
Dropping table odadmin.arssl
```


ARSXML

Purpose


The ARSXML program imports objects within an existing IBM Content Manager OnDemand XML file into a Content Manager OnDemand system, and exports existing Content Manager OnDemand objects into a Content Manager OnDemand XML file.

Syntax


Syntax for adding, updating, and deleting administrative objects in a Content Manager OnDemand system



Syntax for exporting existing Content Manager OnDemand objects into a Content Manager OnDemand XML file



Syntax for validating a Content Manager OnDemand XML file



Description

IBM Content Manager OnDemand includes an XML interface that imports and exports administrative objects in Content Manager OnDemand. In this model, all administrative objects are exported into an XML file, and can be imported into the same system or another system later. You can also create an XML file through a user application or Web interface according to the defined specifications, and import a single object or multiple objects into the system by importing the XML file.

The data import feature allows you to import a single object, a set of defined objects, or even an entire OnDemand system. This feature enables you to complete the following tasks:

- Update a system with newly defined objects
- Backup a complete system
- Copy a set of objects from one system to another system

The XML import file can be one of the following:

- A file which was previously created during a Content Manager OnDemand export process.
- A user created XML file which conforms to the Content Manager OnDemand XML schema.

Important:

- The syntax of each Content Manager OnDemand object that is supported by the import process is specified in the Content Manager OnDemand XML schema file. Each object within the XML file must conform to the exact syntax as specified in the OnDemand XML schema file that is shipped with the product. Objects must appear in the order described in the following document: <http://www.ibm.com/support/docview.wss?uid=swg27019614>.
- The import XML file references an `ondemand.xsd` file, which is provided with Content Manager OnDemand.
- The capitalization of the object names is important and should be used exactly as shown in the data tables in the “XML objects in the Content Manager OnDemand XML file” on page 165.

Parameters for ARSXML [addupdatedelete]

add This is the default action. It adds all of the objects in the input XML file into the specified IBM Content Manager OnDemand system. For objects that contain child objects, such as a list of users or a set of permissions, each child object that is found in the XML file is added to the current list of objects.

update

All the fields that are specified in the input XML file are updated for the object. For example, if the XML file contains a user with a name of harry and a description of The New Description specified, then only the description field of the user harry is updated. All other fields remain unchanged.

Restriction: Not all fields of all objects can be updated.

For objects which contain child objects, such as a list of users or a set of permissions, each child object that is found in the XML file are added to the list of objects. A special attribute on each child object can be set to delete for the cases where a member of a list need to be removed. For example, if you want to add the user `newUser` to a group, and delete the user `oldUser`, the following XML code can be used during an update:

```
...
<group name="MyGroup">
<user name="newUser" />
<user name="oldUser" task="delete" />
</group>
...
```

delete

All of the objects that are specified in the input XML file are deleted from the system. The only attribute that is examined in the XML file is the name attribute. While all other fields might be present in the XML file, they are ignored.

-d Specifies the directory that contains the XML file. If you do not use this parameter, ARSXML looks for the input XML file specified by the **-i** flag in the current working directory where the command is run.

-e error_handling

Controls how the import process handles any Content Manager OnDemand errors.

Important: The XML file syntax errors and other XML-specific errors are detected by the parsing code, and no objects are processed.

The **error_handling** parameter can take three values:

- a** Abort. This is the default value.
- c** Continue. If an error occurs because of problems in the objects or the Content Manager OnDemand system (for example, you attempt to add an object but that object already exists, or you attempt to delete an object that does not exist), the object containing the error is skipped, and the process continues. However, if an XML parsing error occurs, ARSXML stops regardless of whether or not this option is specified.
- u** Update. This value should only be used during an add action.

During an add action:

Abort If an error occurs during an add (for example, an invalid parameter, or the object already exists), the error is logged, the entire import process is stopped, and no further objects are added.

Continue

If an error occurs during an add action, the object containing the error is skipped and an error message is logged. However, the import process continues.

Update

If an object already exists, perform an update action instead of an add action.

During a delete action:

Abort If an error occurs during a delete (for example, there is an invalid object or the object does not exist), the error is logged, the entire delete process is stopped, and no further objects are deleted.

Continue

If an error occurs during an delete action, the object containing the error is skipped and an error message is logged. However, the import process continues.

During an update action:

Abort If an error occurs during a update (for example, there is an invalid

object or the object does not exist), the error is logged, the entire update process is stopped, and no further objects are updated.

Continue

If an error occurs during an update action, the object containing the error is skipped and an error message is logged. However, the import process continues.

-h instance

Specifies the name of the Content Manager OnDemand instance to process.

-i Specifies the input XML file.

-p password

The password is optional. If you omit the **-p** parameter but you do specify a user ID with the **-u** parameter, ARSXML prompts you to enter the password.

Specify one of the following options for *password*:

- The name of the stash (encrypted password) file that contains the password for the user ID specified with the **-u** parameter.
- The password assigned to the user ID you specify. If there is no password assigned to the user ID you specify, press Enter when prompted.

If you omit both the **-u** and the **-p** parameters, Content Manager OnDemand uses the IBM i user profile of the user running ARSXML as the Content Manager OnDemand user ID.

-u userid

The user ID is optional. If you omit this parameter, Content Manager OnDemand uses the current user profile of the job running ARSXML as the user ID. If that user profile does not match a Content Manager OnDemand user ID, ARSXML prompts you to enter a user ID.

-v Enables verbose mode, which displays all messages (informational and error). By default, the ARSXML program displays error messages.

-x Used to prevent prompting from occurring. If you do not use this option, during a delete operation, you are prompted whether you really want to complete the operation, for example:

```
The printer object named 'LabPrinter' is about to be deleted.  
Do you want to delete this object? (Y/N)
```

You need to respond with a *y* or *Y* to confirm the delete operation. Any other response terminates the operation. If you use this option, you do not get the prompt during an operation.

Important: While importing objects, if an object references another object in its definition but the referenced object cannot be found, and the import command uses the **-e c** parameter, an error message is issued, and the default value is used. For example, if a user is defined with a default printer and the default printer cannot be found, the user is added with a default printer of *NONE. If the **-e c** parameter is not used, the above action will not occur.

Parameters for ARSXML export

-d Use this option to specify the directory for the input XML file. If you do not use this parameter, the ARSXML command looks for the input XML file that is specified by the **-i** flag in the current working directory where the command is run.

-e error_handling

Specifies what to do if an error occurs during the export process. The **error_handling** parameter can take two values:

- a** Abort. This is the default value. The export process stops when an error occurs.
- c** Continue. If an error occurs due to problems in the objects or the IBM Content Manager OnDemand system (for example, you attempt to export an object but that object already exists), the object that contains the error is skipped, and the process continues. However, if an XML parsing error occurs, ARSXML stops regardless of whether this option is specified.

-h instance

The name of the Content Manager OnDemand instance you want to process.

- i** Use to specify the input XML file. All of the objects that are specified in the XML file are exported from the system into the output XML file. In this case, the only attribute in the XML file that is examined is the name attribute. While all other attributes are present in the XML file, they are ignored. If an object is specified with the name `_ALL`, all of the objects of that type defined on the system are exported. If the **-i** parameter is not present, all of the objects in the specified system are exported. However, you are prompted to confirm whether this is what is intended, because the export process could take a long time. You can use the **-x** option to turn off the prompting.

- o** Used to specify the name of the output XML file. If this parameter is not specified, the output is directed back to your display if you are running interactively, or to a spooled file if you are running in batch mode.

-p password

The password is optional. If you omit the **-p** parameter but you do specify a user ID with the **-u** parameter, ARSXML prompts you to enter the password.

Specify one of the following options for *password*:

- The name of the stash (encrypted password) file that contains the password for the user ID specified with the **-u** parameter.
- The password assigned to the user ID you specify. If there is no password assigned to the user ID you specify, press Enter when prompted.

If you omit both the **-u** and the **-p** parameters, Content Manager OnDemand uses the IBM i user profile of the user running ARSXML as the Content Manager OnDemand user ID.

-r range

Use this option to specify how much data to export. The range parameter can take four values:

- a** When an application group is exported, all of the applications that are contained within the application group are exported.
- d** Export the specified objects and all dependent objects. If you use this parameter, each object is examined for any dependent objects, and those objects are exported as well. For example, if only user groups are exported and the **d** option is used for the extent, then any users that the groups refer to are exported as well.

- l** Export the objects and any logical views that are associated with them.
- p** Export the objects and a list of permissions for each object.

More than one range option can be specified in the command line, in other words, the options may be combined. For example,

```
-r pl
```

If no range option is specified, then only the specified object or objects are exported.

-u userid

The user ID is optional. If you omit this parameter, Content Manager OnDemand uses the current user profile of the job running ARSXML as the user ID. If that user profile does not match a Content Manager OnDemand user ID, ARSXML prompts you to enter a user ID.

-v Enables verbose mode, which displays all messages (informational and error). By default, the ARSXML program displays error messages.

-w Specifies the encoding in which the resultant output XML file will be created. The default is UTF-8. The list of supported code pages can be found in the arscpcs.cfg file located under the locale directory for the server. For example: /QIBM/ProdData/OnDemand/locale/arscpcs.cfg

-x Used to prevent prompting from occurring. If you do not use this option, during an export operation, you are prompted whether you really want to complete the operation, for example:

```
No input file was specified.
Do you want to export all of the objects on the system? (Y/N)
```

You need to respond with a y or Y to confirm the process. Any other response terminates the process. If you use this option, you do not get the prompt during an operation.

-y Use this option to specify the directory for the output XML file. If this option is not used, the output file that is specified by the **-o** parameter is written to the current working directory where the command is run.

Parameters for ARSXML validate

-i file Specifies the input XML file.

-d directory

Specifies the directory that contains the XML file. If you do not use this parameter, ARSXML looks for the input XML file specified by the **-i** flag in the current working directory where the command is run.

Examples

Example: Adding users:

A user wants to add several users to an IBM Content Manager OnDemand system. That user has created an XML file called newusers.xml.

To complete this task, the following QSHHELL command line function is called:

```
arsxml add -h neptune.ny.ibm.com -u admin -i newusers.xml
```

The file newusers.xml might look like this:

```
<?xml version="1.0" encoding="UTF-8"?>
<onDemand xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xsi:noNamespaceSchemaLocation="ondemand.xsd">
  <user name="Bill" password="xxxxx" userType="User Admin"/>
  <user name="Erin" password="xxxxx" userType="User"/>
  <user name="Brie" password="xxxxx" userType="User"/>
</onDemand>
```

Example: Updating users:

A user wants to update the telephone numbers of the users in the system. This user created an XML file `phone.xml`, which contains the data for the users and their phone numbers.

To complete this task, the following QSHHELL command line function is called:

```
arsxml update -h neptune.ny.ibm.com -u admin -i phone.xml
```

The file `phone.xml` might look like this:

```
<?xml version="1.0" encoding="UTF-8"?>
<onDemand xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xsi:noNamespaceSchemaLocation="ondemand.xsd">
  <user name="Bill" phone="(212) 555-0919"/>
  <user name="Erin" phone="(212) 555-4295"/>
  <user name="Brie" phone="(212) 555-0301"/>
</onDemand>
```

Example: Deleting users:

A user wants to delete another user that no longer needs access to a printer on an IBM Content Manager OnDemand system. That user has created an XML file called `userExpDel.xml`.

To complete this task, the following QSHHELL command line function is called:

```
arsxml delete -h neptune.ny.ibm.com -u admin -p password -v-x-e c i userExpDel.xml
```

The file `userExpDel.xml` might look like this:

```
<?xml version="1.0" encoding="UTF-8"?>
<onDemand xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xsi:noNamespaceSchemaLocation="ondemand.xsd">
  <user name="Bill" password="xxxxx"/>
</onDemand>
```

Example: Exporting multiple objects:

A user wants to export the printer objects named `Prz1` and `Prz2`. This user also want to export the storage set named `FavoriteSS`. An XML file `exportlist.xml` has been created, and contains the information on these objects.

To complete this task, the following QSHHELL command line function is called:

```
>arsxml export -h jupiter.ibm.com -i exportlist.xml -o output.xml
```

The file `exportlist.xml` might look like this:

```
<?xml version="1.0" encoding="UTF-8"?>
<onDemand xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xsi:noNamespaceSchemaLocation="ondemand.xsd">
  <printer name="Prz1"/>
  <printer name="Prz2"/>
  <storageSet name="FavoriteSS"/>
</onDemand>
```

After the user runs the command, an output file named `output.xml` is created, and contains the information for the two printers and the storage set.

Example: Adding groups:

A user wants to add a group of users to an IBM Content Manager OnDemand system. That user has created an XML file called groupAdd.xml.

To complete this task, the following QSHHELL command line function is called:

```
arsxml add -h neptune.ny.ibm.com -u admin -i groupAdd.xml
```

The file groupAdd.xml might look like this:

```
<?xml version="1.0" encoding="UTF-8"?>
<onDemand xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
          xsi:noNamespaceSchemaLocation="ondemand.xsd">
  <group name="BigGroup" password="xxxxx" groupType="Group Admin"/>
</onDemand>
```

Example: Exporting a group:

A user wants to export a group named BigGroup and any users and user groups that are referenced by it. A file named exportbiggroup.xml has been created and looks like this:

```
<?xml version="1.0" encoding="UTF-8"?>
<onDemand xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
          xsi:noNamespaceSchemaLocation="ondemand.xsd">
  <group name="BigGroup"/>
</onDemand>
```

To complete the task, the following QSHHELL command should be used:

```
>arsxml export -h jupiter.ibm.com -i exportbiggroup.xml -r d -o output.xml
```

This command creates an XML file that contains the BigGroup object and all the users and groups that are referenced by it. Also, any users within the referenced groups are exported.

Example: Updating groups:

A user wants to update the name of an existing group. This user created an XML file groupUpdate.xml, which contains the data for the users and their user names.

To complete this task, the following QSHHELL command line function is called:

```
arsxml update -h neptune.ny.ibm.com -u admin -i groupUpdate.xml
```

The file groupUpdate.xml might look like this:

```
<?xml version="1.0" encoding="UTF-8"?>
<onDemand xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
          xsi:noNamespaceSchemaLocation="ondemand.xsd">
  <group name="BigGroup" name="SmallGroup"/>
</onDemand>
```

Example: Exporting all of the users:

A system administrator wants to export all of the users in an IBM Content Manager OnDemand system. A file named exportallusers.xml is created, and looks like this:

```
<?xml version="1.0" encoding="UTF-8"?>
<onDemand xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
          xsi:noNamespaceSchemaLocation="ondemand.xsd">
  <user name="_ALL"/>
</onDemand>
```

To complete the task, the following QSHHELL command should be used:

```
arsxml export -h jupiter.ibm.com -i exportallusers.xml -o users.xml
```

This command creates an XML file `users.xml` that contains the information on all of the users that are defined on the specified system.

Example: Validating input XML file schema:

A user created an XML file called `newusers.xml`.

```
<?xml version="1.0" encoding="UTF-8"?>
<onDemand xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xsi:noNamespaceSchemaLocation="ondemand.xsd">
  <user name="Bill" password="xxxxx" userType="User Admin"/>
  <user name="Erin" password="xxxxx" userType="User"/>
  <user name="Brie" password="xxxxx" userType="User"/>
</onDemand>
```

The user then uses this QSHHELL command to validate the `newusers.xml`:

```
arsxml validate -i newusers.xml
```

Example: Adding printers:

A user wants to add new printers to an IBM Content Manager OnDemand system. That user has created an XML file called `printerAdd.xml`.

To complete this task, the following QSHHELL command line function is called:

```
arsxml add -h neptune.ny.ibm.com -u -p password -v -i printerAdd.xml
```

The file `printerAdd.xml` might look like this:

```
<?xml version="1.0" encoding="UTF-8"?>
<onDemand xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xsi:noNamespaceSchemaLocation="ondemand.xsd">
  <printer name="Printer1"/>
  <printer name="Printer2"/>
  <printer name="Printer3"/>
</onDemand>
```

Tips for using ARSXML

This section includes tips for using ARSXML to import and export administrative objects.

Tip 1: Child objects must be created under parent objects:

Child objects can appear only under parent objects, but not vice versa. For example, the following XML file example is valid:

```
<group name="MyGroup">
  <user name="tom" />
  <user name="chuck" />
</group>
```

However, this example is not valid:

```
<user name ="tom">
  <group name ="MyGroup">
  </group>
</user>
```

The same is true for permissions for users and groups. They are child objects under the application group and folder objects. For example, to add a folder that contains permission for the user Bill, the XML file might look like this:

```
<folder name="MyFolder"...>
  <permission user="Bill" accessAuthority="Yes"
    viewNQAuthority="Yes" maxHits="No Limit" />
  ...
</folder>
```

Tip 2: Application group and folder authorities and permissions are not linked: If you specify `adminAuthority="Yes"`, you do not get the view, add, print, fax, and copy document permissions, or the view, add, and copy annotations permissions. You need to specify each item individually.

Tip 3: Parsing error while running ARSXML:

You might receive this error message while running ARSXML:

A parsing error occurred in file `fileName`, Line `nnn`, Column `nnn` :
cvc-elt.1: Cannot find the declaration of element 'onDemand'.

It usually indicates that the IBM Content Manager OnDemand schema file (`ondemand.xsd`) cannot be found. See “Specifying the `ondemand.xsd` schema file” on page 156 for details on the location and specification of the `ondemand.xsd` file.

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, `arsload`)

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.” For examples of `ARSSTASH` commands, see “Examples of `ARSSTASH` commands” on page 315.
2. Save the stash file in a directory and limit access to that file through file permissions.

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.

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`

User exit reference

Output queue or directory monitor user exit program

You can design a user exit program to alter the application group name or application name the monitor finds as it processes files in a monitored output queue or IFS directory.

Overview of the monitor

An output queue or directory monitor (started by the Start Monitor for OnDemand (STRMONOND) command or by using a monitor definition in IBM Navigator for i) automatically processes files from the specified output queue or IFS directory.

When you start the monitor, two parameters are used to determine which application group and application to use to archive the input files.

The Value for application name (APPSRC) parameter is used to determine the application name. Up to three sources can be specified. If a valid application name is not found that uses the first source, the second source is used. If a valid application name is not found that uses the second source, the third source is used.

The Value for application group (APPGRPSRC) parameter is used to determine the application group name. Up to three sources can be specified. If a valid application group name is not found that uses the first source, the second source is used. If a valid application group name is not found that uses the second source, the third source is used.

The same source or different sources can be specified for APPSRC and APPGRPSRC.

Valid sources for an output queue monitor are:

- *SPLFNAME
- *FORMTYPE
- *USERDATA
- *JOBNAME
- *USRDFNOPT1
- *USRDFNOPT2
- *USRDFNOPT3
- *USRDFNOPT4
- *USRDFNDTA

Valid sources for a directory monitor are:

- *FIRST
- *SECOND
- *THIRD
- *FOURTH

How the monitor exit works

Application name - first pass: When an input file is selected from an output queue or IFS directory to be processed, OnDemand checks for a user exit program in the monitor job's library list with a name that matches the first attribute specified in the APPSRC parameter of the STRMONOND command.

If a program is found, it is called, allowing the application name to be changed within the exit program as needed. OnDemand then uses the new application name to verify that an application definition by that name exists.

If a program is not found, OnDemand looks for an application name that matches the first attribute.

Application name - second pass: If an application is still not found, OnDemand now checks for a user exit program in the monitor job's library list with a name that matches the second attribute specified in the APPSRC parameter of the STRMONOND command.

If a program is found, it is called, allowing the application name to be changed as needed. OnDemand then uses the new application name to verify that an application definition by that name exists.

If a program is not found, OnDemand looks for an application name that matches the second attribute.

Application name - third pass: If an application is still not found, OnDemand now checks for a user exit program in the monitor job's library list with a name that matches the third attribute specified on the APPSRC parameter of the STRMONOND command.

If a program is found, it is called, allowing the application name to be changed as needed. OnDemand then uses the new application name to verify that an application definition by that name exists.

If a matching application name is still not found, error messages are placed in the monitor job log. For an output queue monitor, the spooled file is moved, in Ready (RDY) status, to the error output queue (ERROUTQ) specified in the STRMONOND command. For a directory monitor, files will be left in the directory and have '.ERR' added to end of the file name.

Application group name - first pass: Next, OnDemand checks for a user exit program in the monitor job's library list with a name that matches the first attribute specified in the APPGRPSRC parameter of the STRMONOND command.

If a program is found, it is called, allowing both the application group and application names to be changed as needed. OnDemand then uses the new application group name to verify that an application group definition by that name exists.

If a program is not found, OnDemand looks for an application group name that matches the first attribute.

Application group name - second pass: If an application group is still not found, OnDemand now checks for a user exit program in the monitor job's library list with a name that matches the second attribute specified in the APPGRPSRC parameter of the STRMONOND command.

If a program is found, it is called, allowing both the application group and application names to be changed as needed. OnDemand then uses the new application group name to verify that an application group definition by that name exists.

If a program is not found, OnDemand looks for an application group name that matches the second attribute.

Application group name - third pass: If an application is still not found, OnDemand now checks for a user exit program in the monitor job's library list with a name that matches the third attribute in the APPGRPSRC parameter of the STRMONOND command.

If a program is found, it is called, allowing both the application group and application names to be changed as needed. OnDemand then uses the new application group name to verify that an application group definition by that name exists.

If a program is not found, OnDemand looks for an application group name that matches the third attribute.

If a matching application group name is still not found, error messages are placed in the monitor job log. For an output queue monitor, the spooled file is moved, in Ready (RDY) status, to the error output queue (ERROUTQ) specified in the STRMONOND command. For a directory monitor, files will be left in the directory and have '.ERR' added to end of the file name.

Special case – APPSRC(*APPGRP): If the value of the APPSRC parameter is *APPGRP, the monitor exit is called only for the APPGRPSRC parameter. The application name can still be changed by the exit program.

Exit program details

When OnDemand calls the user exit, it passes several parameters, only two of which can be changed.

These two changeable fields are:

- Application group – 60 bytes
- Application – 60 bytes

For example, if the monitor finds PGM123 in the User Data spooled file attribute of the spooled file, but you want OnDemand to use the application name INVOICES, you can write your user exit program as follows:

```
CHGVAR      VAR(&APP) VALUE('INVOICES')
```

There is a sample monitor output queue user exit program in the source file QSAMPLES2 in the library QUSRRDARS, with member name PGM123.

Facsimile user exit program

The facsimile user exit program is designed to enable OnDemand server fax functions to work with IBM i facsimile (fax) software other than Facsimile Support from IBM. To accomplish this, modify and recompile the QRLMSFAX program source code to change the command that is issued when an end user requests a server fax during report retrieval. Simply change the line in the program that issues the SNDFAX command to use the command for the fax software you have installed on your system. (You could also call a program if no command is available.)

An end-user requesting a server fax causes the QRLMSFAX program to be called. If you do not change and recompile it, the standard program will run (which issues the Facsimile Support command). If you do change and recompile the QRLMSFAX program, then the changed program will run (which will issue **your** facsimile command).

The sample CL source code for this program (member name QRLMSFAX) can be found in source file QSAMPLES2 in libraries QRDARS and QUSRRDARS. (Any program source code that you modify should not be placed in QRDARS because that library is replaced during software upgrades. However, QUSRRDARS library is not replaced and can be used for your modified source.) IBM recommends that you copy this source code into a backup member in case you need to go back to the original function as shipped from IBM. When you recompile your program, be sure to preserve the QRLMSFAX program name and place the compiled program back in QRDARS library to replace the program that is shipped by IBM.

Automating ARSLOAD data loading

This section provides information to help you configure the ARSLOAD data loading API.

ARSLOAD

The ARSLOAD program is the main Content Manager OnDemand data loading and indexing API. You can configure the ARSLOAD API to monitor specific IFS directories for report data. If the data needs to be indexed, then the ARSLOAD program calls the indexing program that is specified in the Content Manager OnDemand application definition. The ARSLOAD program then works with the database manager to load the index data into the database and works with the storage manager to load the report data and resources on to storage volumes.

Note: The Content Manager OnDemand server job must be running, otherwise the ARSLOAD program will fail.

Automating ARSLOAD

The following shows an example of the ARSLOAD API, which is run in QSHHELL.

```
arsload -v -c /arsdir/dir4 -d /arsdir/dir1 -d /arsdir/dir2 -d /arsdir/dir3
```

In the example, the ARSLOAD program checks for data in the specified directories every ten minutes (the default polling time). If data needs to be indexed, then the ARSLOAD program stores the index data in the specified index directory.

You must verify the names of the directories. Replace the strings `/arsdir/dir1`, `/arsdir/dir2`, `/arsdir/dir3`, and `/arsdir/dir4` with the names of directories that are valid on the server that you are configuring.

Important: The ARSLOAD program uses a particular part of the input file name to determine the application group to load. You can use the `-G` parameter to specify a different part of the file name to identify the application group to load. If the application group contains more than one application, then you must identify the application to load. Otherwise, the load will fail. You can use the `-A` parameter to specify the part of the file name that identifies the application.

If a user ID and password are not specified on the ARSLOAD program, the current IBM i userid is used. See “API and user exit reference” on page 259 for more information about the ARSLOAD program and its userid and password parameters.

After indexing the data, the ARSLOAD program deletes the input files, unless you specify otherwise. Any output or error messages that are generated by the ARSLOAD program are written to the System Log. You can open the System Log folder and retrieve any messages that were generated by the ARSLOAD program. For example, you may see message number 87 for a successful load or message 88 for a failed load.

Accessibility features

This product includes a number of features that make it more accessible for people with disabilities. These features include:

- The ability to operate all features using the keyboard instead of the mouse.
- Support for enhanced display properties
- Options for audio and visual alert cues
- Compatibility with assistive technologies
- Compatibility with operating system accessibility features
- Accessible documentation formats

Keyboard input and navigation

Keyboard input

The OnDemand clients can be operated using only the keyboard. Menu items and controls provide access keys that allow users to activate a control or select a menu item directly from the keyboard. These keys are self-documenting, in that the access keys are underlined on the control or menu where they appear.

Keyboard focus

In Windows-based systems, the position of the keyboard focus is highlighted, indicating which area of the window is active and where the user's keystrokes will have an effect.

Features for accessible display

The clients have a number of features that enhance the user interface and improve accessibility for users with low vision. These enhancements include support for high-contrast settings and customizable font properties.

High-contrast mode

The clients support the high-contrast-mode option that is provided by the operating system. This feature assists users who require a higher degree of contrast between background and foreground colors.

Font settings

In Windows-based systems, you can specify display settings that determine the color, size, and font for the text in menus and dialog windows. The client allows you to select the font for the document list.

Non-dependence on color

You do not need to distinguish between colors in order to use any function of this product.

Alternative alert cues

In Windows-based systems, the SoundSentry feature can be used to provide visual feedback for general application and system alerts such as warning beeps.

Compatibility with assistive technologies

The clients are compatible with screen reader applications such as Narrator and Via Voice. The clients have properties required for these accessibility applications to make onscreen information available to visually impaired users.

Fully accessible alternatives for the line data viewer and the AFP plugin

To perform functions of the line data viewer and the AFP plugin, use the client or administrative client. At this time, functions of the line data viewer and the AFP plugin are not accessible through the Web administrator.

Accessible documentation

Documentation for the OnDemand product is available in HTML format. This allows users to view documentation according to the display preferences set in their browsers. It also allows the use of screen readers and other assistive technologies.

Notices

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

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

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant 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 document 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, the 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.

All statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

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 illustrate 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.

Trademarks

This topic lists IBM trademarks and certain non-IBM 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 <http://www.ibm.com/legal/copytrade.shtml>.

The following terms are trademarks or registered trademarks of other companies:

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

Intel, Intel logo, Intel Inside, Intel Inside logo, Intel Centrino, Intel Centrino logo, Celeron, Intel Xeon, Intel SpeedStep, Itanium, Pentium, and Xeon are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Java and all Java-based trademarks 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.

Microsoft, Windows, and the Windows logo are trademarks of Microsoft 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.

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

Index

Numerics

- 128 ApplGrp Segment Expire 3
- 14 DB Info Exported 6
- 166 ApplGrp Segment Export 6
- 168 ApplGrp Segment Not Available 7
- 196 Cache Expiration 17
- 197 Cache Migration 15
- 201 Concurrent Licenses 143
- 202 Concurrent Licenses 143
- 87 ApplGrp Load (System Migration) 6

A

- accessibility 321
- adding
 - application groups 97
 - applications 97
 - fields 108
 - folders 97
 - groups 93
 - primary storage nodes 97
 - printers 95
 - reports 97
 - server printers 95
 - servers 46
 - storage nodes 97
 - storage sets 97
 - users 90
 - users to a group 94
- administrative client
 - adding servers 46
 - assistants and wizards 55
 - changing a password 47
 - hardware requirements 41
 - installing 41
 - logging on a server 46
 - passwords 47
 - report wizard 55
 - software requirements 41
 - starting 46
 - using 45
 - wizards and assistants 55
- administrators
 - application group/folder administrator 63
 - folder administrator 63
 - system administrator 63
 - user administrator 63
- advanced system administration
 - native Lightweight Directory Access Protocol (LDAP) support 77
 - Enabling LDAP authentication 77
 - How Content Manager OnDemand works with LDAP 77
 - LDAP server requirements 77
- AFP data, indexing 102
- aggregation 59, 89
- annotation permission 71
- annotations 51

- APIs
 - how to read a syntax diagram 261
 - reference 259
- application group object 204
- application group/folder
 - administrator 63, 64
- application groups 296
 - access permission 71
 - adding 97, 99
 - administrator 63
 - administrator permission 71
 - annotation permission 71
 - concepts 66
 - defining 97, 99, 114, 121, 129
 - document permission 71
 - documents
 - maintaining 15, 17
 - expiration processing 3, 8, 17
 - fields
 - defining 99
 - FORMS parameter 296
 - importing and exporting 151
 - importing migrated index data 8
 - importing migrated indexes 7
 - index data
 - maintaining 3, 6, 7
 - logical views permission 71
 - maintenance 3, 6, 7, 15, 17
 - messages 3, 6, 7, 15, 17
 - migration processing 6, 7, 8, 15
 - overview 66
 - permissions 71, 99
 - query restriction 71
 - storage management 99
- application object 176
- Application Program Interfaces (APIs) 259
- applications
 - adding 97, 101
 - AFP data 102
 - concepts 67
 - DATASET parameter 296
 - defining 97, 101, 114, 122, 131
 - fields
 - defining 102
 - importing and exporting 151
 - indexes
 - defining 102
 - indexing parameters 102
 - load information 102
 - overview 67
 - print options 102
 - triggers 102
- archive storage manager, migration processing 8
- ARSDATE
 - reference 263
- ARSDOC
 - reference 267
- arsdoc command
 - parameters 278

- ARSLOAD
 - automating 319
 - Download for z/OS 296
 - reference 294
 - running 125, 133
 - starting 319
- ARSSLU
 - reference 303
- ARSSUPPORT utility 143
- ARSUUPDT
 - exit 296
 - report specification archive definition
 - exit 296
 - user exit 296
- ARSXML
 - reference 304
- assistants and wizards 55
- authority
 - create application groups 64
 - create folders 64
 - create groups 64
 - create users 64
 - requirements 90

B

- backup and recovery
 - database 120
 - reports and resources 38
- backup considerations 37
- batch administration 151
 - common problems during installation 153
 - installation verification 152
- installing 151
- prerequisites 151

C

- cabinet object 228
- cabinets
 - concepts 68
 - importing and exporting 151
 - overview 68
- commands
 - adding documents 267
 - ADDRPTOND 243
 - ARSDATE 263
 - ARSDOC 267
 - ARSLOAD 294
 - ARSSLU 303
 - ARSXML 304
 - CHGPLDOND 13, 59, 243
 - CRTINSTOND 243
 - data, loading 294
 - deleting documents 267
 - documents, adding, deleting, getting, printing, updating 267
 - documents, loading 294
 - ENDMONOND 243

- commands (*continued*)
 - exporting an XML file 304
 - FNDKEYOND 243
 - getting documents 267
 - importing an XML file 304
 - loading data, documents, reports 294
 - MGRMEDRDAR 244
 - MRGSPLFOND 245
 - printing documents 267
 - PRTRPTOND 243
 - PRTXTOND 243
 - reports, loading 294
 - RMVRPTOND 243
 - STRASMOND 243
 - STRDSMOND 243
 - STRIMPOND 243
 - STRMONOND 243
 - updating documents 267
 - XML file, exporting 304
 - XML file, importing 304
- concurrent users monitoring 143
- Content Manager OnDemand
 - date formats 263
 - internal date format 263
- conventions
 - group name 65
 - user name 62
- copying items 53
- create application groups authority 64
- create folders authority 64
- create groups authority 64
- create users authority 64
- creating an XML file 155

D

- data
 - adding 267
 - deleting 3, 15, 267
 - deleting a report 139
 - Download for z/OS 296
 - estimating storage space for 114
 - expiration processing 3, 15
 - importing migrated indexes 7
 - life of 3
 - loading 113, 121, 129, 294
 - migration processing 5, 7, 13
 - printing 267
 - restarting a load process 137
 - retrieving 267
 - unloading 139
 - updating 267
- data loading 319
- database
 - adding index data 118
 - backup 120
 - concepts 1
 - deleting index data 3, 7
 - expiration processing 3, 7
 - fields
 - defining 99
 - importing migrated index data 7
 - maintenance 3
 - migration processing 5, 7
 - segment table 118
 - updating 267
- database manager 1

- DATASET parameter
 - applications 296
 - ARSLOAD program 296
- date
 - Content Manager OnDemand internal
 - format 263
 - formats of 263
 - internal format 263
 - obtaining Content Manager
 - OnDemand internal format 263
- DB2
 - concepts 1
 - database 1
 - importing migrated index data 9
 - indexes 2
 - instances 1
 - journals 2
 - systems 1
 - tables 1
- defining
 - application groups 97
 - applications 97
 - fields 108
 - folders 97
 - groups 93
 - primary storage nodes 97
 - printers 95
 - reports 97
 - server printers 95
 - storage nodes 97
 - storage sets 97
 - users 90
- deleting a report 139
- deleting index data 3, 7
- disability 321
- document permission 71
- drag and drop operations 54

E

- Enabling LDAP authentication 77
- error messages 249
- estimating storage space 114
- exits
 - ARSUUPDT 296
 - report specification archive
 - definition 296
- expiration of passwords 48
- expiration processing 3, 7
- expired passwords 48
- exporting items 54

F

- facsimile parameters 256
- facsimile user exit program 318
- fax parameters 256
- fax user exit program 318
- fields 108
 - application group
 - defining 99
 - database
 - defining 99
 - defining 99, 105
 - folder
 - defining 105

- fields (*continued*)
 - indexing
 - defining 102
 - fields permission 71
- FILESTAT error codes 149
- folder administrator 63
- folder object 216
- folders
 - access permission 71
 - adding 97, 105
 - administrator 63
 - administrator permission 71
 - concepts 67
 - defining 97, 105, 123, 131
 - fields
 - defining 105
 - fields permission 71
 - importing and exporting 151
 - named queries permission 71
 - overview 67
 - permissions 71, 105
- FORMS parameter
 - application groups 296
 - ARSLOAD program 296

G

- generic indexer parameter file 124, 133
- graphical indexer 55
- group administration 93
- group object 171
- group owner 65
- groups
 - adding 93
 - adding users 94
 - assigning users 94
 - concepts 65
 - defining 93
 - importing and exporting 151
 - naming 65
 - overview 65
 - owner 65
 - permissions 71

H

- hold object 233
- holds
 - concepts 68
 - overview 68
- how to read a syntax diagram 261

I

- i Navigator
 - configuration vii
 - Management Central vii
 - system navigation vii
- image files
 - index data 124
 - loading 121
- importing migrated index data 7
- inactivity time out 49
- index data
 - about 2
 - adding 267

- index data *(continued)*
 - creating 124, 133
 - deleting 3, 7, 139, 267
 - expiration processing 3, 7
 - fields
 - defining 102
 - generic indexer 124, 133
 - importing 7
 - indexes
 - defining 102
 - life of 3
 - load information 102
 - loading 113, 121, 129, 294
 - migration processing 5, 7
 - processing with postprocessor programs 102
 - restarting a load process 137
 - retrieving 267
 - triggers 102
 - updating 267
- indexes
 - defining 102
- installing batch administration prerequisites 151
- installing the administrative client 41
- instances 1
- internal date format 263

J

- journal receivers 2

K

- keyboard 321

L

- LDAP 77
 - authentication 78, 79, 80, 82, 83
 - configurations 83
 - message examples 83
- LDAP authentication 51
- licenses, monitoring 143
- life of data and indexes 3, 5
- load user exit 296
- loading data 113, 121, 129, 319
- loading image files 121
- loading index data 113, 121, 129
- loading reports 113
- loading user-defined data 129
- logging on a server 47
- logical views permission 71
- login processing 50

M

- managing servers 243
- maximum password age 48
- messages
 - 128 ApplGrp Segment Expire 3
 - 14 DB Info Exported 6
 - 166 ApplGrp Segment Export 6
 - 168 ApplGrp Segment Not Available 7

- messages *(continued)*
 - 196 Cache Expiration 17
 - 197 Cache Migration 15
 - 201 Concurrent Licenses 143
 - 202 Concurrent Licenses 143
 - 87 ApplGrp Load (System Migration) 6
 - application groups 3, 6, 7, 15, 17
 - errors and alerts 141, 249
 - expiration processing 3, 17
 - licenses 143
 - migration processing 6, 7, 15
 - OnDemand errors and alerts 141, 249
 - System Log 141, 249
- MGRMEDRDAR 244
- migration processing 5, 7
- minimum password length 48
- monitor 115
- monitor exit program parameters 318
- monitoring
 - concurrent users 143
 - licenses 143
 - users 143

N

- named queries permission 71
- naming groups 65
- naming users 62

O

- objects 157
- obtaining Content Manager OnDemand
 - internal date format 263
- OnDemand object
 - child objects 159
 - end of object 159
 - object attributes 158
 - Start of object 158
- OnDemand XML files
 - creating for update 160
- ondemand.xsd file 156
- optical volume statistics reset 75
- OS/400 printer file 255
- Output queue monitors 115
- output queue or directory monitor user
 - exit program 315
- Output queue or directory monitor user
 - exit program
 - overview of function 315
- Output queue or directory monitor user
 - exit program processing 316

P

- password 50
- password length 47
- passwords 47, 48
 - specifying in ARSLOAD 296
- permissions
 - about 70
 - application groups 71, 99
 - concepts 70
 - examples 71
 - folders 71, 105

- permissions *(continued)*
 - overview 70
 - strategies 74
- postprocessor program 102
- prerequisites vii
- primary storage nodes
 - adding 97
 - concepts 66
 - defining 97
 - overview 66
- print parameters 255
- printer file 255
- printer object 224
- printers
 - adding 95
 - concepts 65
 - defining 95
 - importing and exporting 151
 - options 102
 - overview 65
- printing documents 255, 267
- programs
 - date formats 263
 - internal date format 263
 - obtaining Content Manager OnDemand internal format 263

Q

- query restriction 71

R

- receivers, journal 2
- recovery considerations 38
- related information vii
- report specifications archive definition
 - exit 296
- report wizard 55
- reports
 - adding 97, 267
 - backup and recovery 38
 - defining 97
 - deleting 139, 267
 - Download for z/OS 296
 - estimating storage space for 114
 - loading 113, 294
 - printing 267
 - restarting a load process 137
 - retrieving 267
 - unloading 139
 - updating 267
- requirements
 - hardware 41
 - software 41
- resetting optical volume statistics 75
- resources 119
- restarting a load process 137

S

- schema file 156
- security and user administration 90, 93
- segment table 118
- server fax parameters 256
- server fax user exit program 318

- server print parameters 255
- server printers
 - adding 95
 - concepts 65
 - defining 95
 - options 102
 - overview 65
- servers
 - adding 46
 - adding items 53
 - copying items 53
 - drag and drop operations 54
 - exporting items 54
 - logging on 47
 - system parameters 52
 - updating 46
- starting the administrative client 46
- storage
 - administration 97
 - for reports 114
 - management 99
- storage nodes
 - adding 97
 - concepts 66
 - defining 97
 - loading data 118
 - overview 66
- storage set object 215
- storage sets
 - adding 97
 - assigning to an application group 99
 - concepts 66
 - defining 97
 - importing and exporting 151
 - loading data 118
 - migration processing 8
 - overview 66
 - System Migration application group 8
- STRIMPOND command 9
- subsystem, changing server jobs 145
- syntax diagram, how to read 261
- system administrator 63
- System Log
 - error messages 249
 - message reference 249
- System Log messages
 - 128 ApplGrp Segment Expire 3
 - 14 DB Info Exported 6
 - 166 ApplGrp Segment Export 6
 - 168 ApplGrp Segment Not Available 7
 - 197 Cache Migration 15
 - 201 Concurrent Licenses 143
 - 202 Concurrent Licenses 143
 - 87 ApplGrp Load (System Migration) 6
 - application groups 3, 6, 7, 15, 17
 - enabling 49
 - expiration processing 3, 17
 - loading data 119, 127, 135
 - migration processing 6, 7, 15
 - monitoring 141
 - user exit programs 49
 - user messages 143
- System Log user exit program 7
- System Migration application group 8

- system parameters
 - about 48
 - annotations 51
 - inactivity time out 49
 - login processing 50
 - maximum password age 48
 - minimum password length 48
 - password 50
 - password age 48
 - password expiration 48
 - password length 48
 - passwords 48
 - setting 52
 - System Log 49
 - time out 49
 - user exit logging 49
 - userid 50
- systemParameters object 237
- systems 1

T

- tables 1
- time out 49
- trace parameters
 - error 53
 - flow 53
 - information 53
 - setting 53
 - warning 53
- triggers
 - defining 102

U

- unloading a report 139
- update servers dialog box 46
- user 63, 64
- user administration 90
- user administrator 63, 64
- user exits
 - ARSUUPDT 296
 - facsimile user exit program 318
 - fax user exit program 318
 - load user exit 296
 - migration processing 7
 - output queue or Directory monitor user exit program 315
 - reference 315
 - report specification archive definition 296
 - saving messages in the System Log 49
 - server fax user exit program 318
 - System Log 7
- user object 168
- user-defined data
 - index data 133
 - loading 129
- userid 50
 - specifying in ARSLOAD 296
- users
 - adding 90
 - adding to a group 94
 - application group/folder administrator 63

- users (*continued*)
 - assigning to a group 94
 - authority 64, 90
 - changing a password 47
 - concepts 62
 - defining 90
 - folder administrator 63
 - importing and exporting 151
 - licenses 143
 - monitoring 143
 - naming 62
 - overview 62
 - passwords 47
 - permissions 70, 71
 - system administrator 63
 - types 63
 - user 63
 - user administrator 63
- using the administrative client 45

W

- wizards and assistants 55

X

- XML file
 - creating 155
 - data model 165
 - exporting 304
 - importing 304
 - objects 165
- XML files
 - creating for delete and export 162
- XML objects
 - application 176
 - application group 204
 - cabinet 228
 - folder 216
 - group 171
 - hold 233
 - printer 224
 - storage set 215
 - system parameters 237
 - user 168
- XML schema file 154



Product Number: 5770-RD1

SC19-2792-01

