

Content Manager OnDemand for z/OS
Version 10 Release 5

Introduction and Planning Guide



Note

Before using this information and the product it supports, read the information in [“Notices” on page 63.](#)

This edition applies to the following products and to all subsequent releases and modifications until otherwise indicated in new editions:

- Version 10 Release 5 of IBM® Content Manager OnDemand for z/OS (product number 5697-CM1)

© **Copyright 2017 - 2020 All Rights Reserved. UNICOM Systems, Inc. – a division of UNICOM Global.**

© **Copyright International Business Machines Corporation 1993, 2020.**

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

Contents

- ibm.com[®] and related resources..... V**
 - Contacting IBM..... v

- Chapter 1. Content Manager OnDemand overview..... 1**
 - System architecture..... 1
 - Indexers..... 2
 - Enhanced Retention Management (ERM)..... 3
 - Content Manager OnDemand Distribution Facility (ODF)..... 3
 - Content Manager OnDemand administrative client..... 4
 - Content Manager OnDemand concepts..... 4
 - Content Manager OnDemand applications..... 4
 - Content Manager OnDemand application groups..... 5
 - Content Manager OnDemand folders..... 5
 - Content Manager OnDemand cabinets..... 6
 - Indexing methods..... 7
 - Content Manager OnDemand documents..... 7
 - Content Manager OnDemand servers..... 8
 - Request manager..... 9
 - Database manager..... 10
 - Content Manager OnDemand storage manager..... 10
 - Data indexing and loading..... 10
 - Content Manager OnDemand management programs..... 11
 - Content Manager OnDemand Web Enablement Kit..... 12
 - Viewing and transforming documents..... 12

- Chapter 2. Preparing for Content Manager OnDemand..... 15**
 - Administrative roles and responsibilities..... 16
 - Content Manager OnDemand application programming interfaces..... 17
 - Client customization..... 17
 - Content Manager OnDemand server programs..... 17
 - Content Manager OnDemand server logging..... 18
 - Content Manager OnDemand user security exit..... 18
 - Unified login exit..... 19
 - Report Specifications Archive Definition exit..... 19
 - Client preview exit..... 19
 - Table space creation exit..... 20
 - JES Spool Data Capture facility..... 20
 - Indexer exits..... 21

- Chapter 3. System requirements..... 23**

- Chapter 4. Storage..... 25**
 - Storage objects..... 25
 - Other external storage managers..... 26

- Chapter 5. Reports and other data..... 27**
 - Collecting requirements..... 27
 - Input data formats..... 27
 - Indexing data..... 28
 - Document indexing..... 29

Report indexing.....	29
Indexing data with ACIF on z/OS.....	30
Using the Content Manager OnDemand 390 Indexer.....	31
Using the Content Manager OnDemand Generic Indexer.....	32
Using the Content Manager OnDemand XML Indexer.....	32
Indexing reports using date fields.....	33
Chapter 6. Content Manager OnDemand objects.....	35
Folders.....	35
Cabinets.....	36
Application groups.....	36
Applications.....	38
Users and groups.....	39
Permissions.....	39
Folder permissions.....	39
Application group permissions.....	40
Naming rules.....	41
Data types and field types.....	43
Chapter 7. Storage requirements overview.....	49
Storage hierarchy.....	50
Data compression.....	51
Calculating data storage requirements.....	52
Temporary space for indexing.....	52
Cache storage.....	52
Content Manager OnDemand database storage.....	53
Database log file storage.....	55
Server print storage space.....	55
Storage sizing examples.....	55
Calculating archive storage requirements.....	57
Report storage space.....	57
Migrated index storage space.....	58
Device storage.....	58
Storage sizing worksheets.....	58
A report that contains logical items.....	59
A report that contains transaction data.....	59
Chapter 8. Backup and recovery overview.....	61
Server software.....	61
Server information.....	61
Database backup.....	61
Reports.....	61
Cache storage.....	62
Archive storage.....	62
Notices.....	63
Trademarks.....	64
Terms and conditions for product documentation.....	65
IBM Online Privacy Statement.....	65
Index.....	67

ibm.com[®] and related resources

Product support and documentation are available from [ibm.com](https://www.ibm.com)[®].

Support and assistance

From [ibm.com](https://www.ibm.com), click **Support & downloads** and select the type of support that you need. From the Support Portal, you can search for product information, download fixes, open service requests, and access other tools and resources.

IBM Knowledge Center

See your online product information in IBM Knowledge Center at <https://www.ibm.com/support/knowledgecenter/SSQHWE>.

PDF publications

See the following PDF publications for your product at <https://www.ibm.com/support/pages/node/1079043>.

Contacting IBM

For general inquiries, call 800-IBM-4YOU (800-426-4968). To contact IBM customer service in the United States or Canada, call 1-800-IBM-SERV (1-800-426-7378).

For more information about how to contact IBM, including TTY service, see the Contact IBM website at <http://www.ibm.com/contact/us/>.

Chapter 1. Content Manager OnDemand overview

Content Manager OnDemand provides features that can benefit organizations of many different sizes.

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

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

Content Manager OnDemand fully integrates the capabilities of Advanced Function Presentation (AFP), including management of resources, indexes, and annotations, and supports full fidelity reprinting of documents to devices attached to a workstation, a Content Manager OnDemand server, or any other server on the network.

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

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

Some of the advantages that Content Manager OnDemand offers include:

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

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

Content Manager OnDemand does the following:

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

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

System architecture

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

[Figure 1 on page 2](#) shows an example.




Figure 1. Content Manager OnDemand system

Content Manager OnDemand client programs run on workstations attached to the network and communicate with Content Manager OnDemand servers. The Content Manager OnDemand library server manages a database of information about the users of the system and the reports stored on the system. A Content Manager OnDemand object server manages the reports on disk, optical, and tape storage devices. A Content Manager OnDemand system has one library server and one or more object servers. An object server can operate on the same system or node as the library server or on a different system or node than the library server.

Content Manager OnDemand client programs operate on personal computers running Windows. You can also configure most client functions to run through a Web browser on many other operating systems by using the Content Manager OnDemand Web Enablement Kit. With the client program, users can search for reports, retrieve documents from Content Manager OnDemand, view and print copies or pages of documents, and attach electronic notes to pages of a document. Users can also save their search criteria to reuse, similar to bookmarking a Web site in a browser.

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

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

Indexers

You can use indexers to define the characteristics of your documents to IBM Content Manager OnDemand.

The indexers in IBM Content Manager OnDemand are programs that provide the following functions:

- A way to load and store reports
- Determining where one document ends and the next document begins
- Determining which index values are to be associated with each document
- Using index values to identify and retrieve documents for viewing or printing
- Creating a resource file that contains all the resources that are needed to view and print a document

Usually, the index values are extracted from the content of the documents; however, they can also be created manually or by a custom application (an exit). The choice of an indexer depends on the operating system, the format of the documents, and the function that is needed. Content Manager OnDemand provides the following indexers:

- ACIF indexer

- 390 indexer
- 400 indexer
- PDF indexer
- Generic indexer
- XML indexer

Enhanced Retention Management (ERM)

The Content Manager OnDemand Enhanced Retention Management feature enables you to manage and enforce retention of documents.

In a Content Manager OnDemand system, you retain documents for a specific amount of time. This practice is commonly referred to as retention management. Records management describes the process of retaining and deleting documents under a set of circumstances that are not necessarily bounded by time, for example, until the end of litigation.

Without the Enhanced Retention Management feature, Content Manager OnDemand cannot implement records management due to the following limitations:

- Records management requires that you have control over individual documents. Content Manager OnDemand manages application groups instead of individual documents and works with a storage manager to delete (expire) documents.
- Records management requires flexibility in defining the time when documents are deleted. However, Content Manager OnDemand defines only the time when application groups with fixed time ranges are deleted, for example, five years after Content Manager OnDemand loads documents.

To overcome these limitations, you can purchase and install the Enhanced Retention Management feature. With the Enhanced Retention Management feature, you can control individual documents by introducing holds, a mechanism that identifies the documents that you want to keep for some time. To expire a document with a hold, you need to remove the hold. Holds give you flexibility to choose when to delete documents because you control when to remove a hold. By using holds, you control when to delete a document. You can manage holds through any of the following interfaces:

- Content Manager OnDemand client or administrative client
- IBM Content Navigator
- ARSDOC command
- ODWEK Java APIs
- FileNet P8 when you integrate it with Content Federation Services for Content Manager OnDemand by enabling the Content Federation Services for Content Manager OnDemand. The Content Federation Services for Content Manager OnDemand also enables you to federate Content Manager OnDemand repositories. This feature connects your Content Manager OnDemand content to business process management (IBM BPM) and records management features of FileNet P8.

Content Manager OnDemand Distribution Facility (ODF)

Content Manager OnDemand Distribution Facility (ODF) is the report distribution feature for Content Manager OnDemand for Multiplatforms, Content Manager OnDemand for z/OS®, and Content Manager OnDemand for i.

ODF provides an easy way to automatically group reports and portions of reports and distribute the reports to multiple users. ODF distributions can be printed, created as an output file, or emailed as an attachment. ODF can retrieve reports that are stored in a Content Manager OnDemand server on z/OS or IBM i, or any of the operating systems that are supported by Content Manager OnDemand for Multiplatforms.

Content Manager OnDemand also provides an ODF monitor utility. The ODF monitor utility is an interactive workstation client program that enables you to check the status of distributions that are submitted for processing and to monitor distribution activity.

Content Manager OnDemand administrative client

The Content Manager OnDemand administrative client helps you maintain your Content Manager OnDemand system.

The administrative client provides tools that enable you to:

- Define reports to the system
- Add and maintain users and groups
- Add and maintain server printers
- Add and maintain storage sets and storage nodes
- Add and maintain users, groups, applications, application groups, storage sets, folders, printers, holds, and cabinets
- Add and maintain OnDemand Distribution Facility elements including recipients, recipient lists, report IDs, report bundles, and distributions
- Add and maintain servers
- Set system parameters for servers and client programs
- Copy items from one server to another
- Track changes to the system

When you use the administrative client to add or update an object in Content Manager OnDemand, information about the object is saved in the system log. The information includes the changes that you made.

The administrative client also provides a report wizard that enables you to:

- Activate Full Text Search in the application group and to create FTS folder fields.
- Create an application, application group, and folder using a wizard for reports that are loaded by using the generic indexer format.
- Provide the ability to define a generic indexer report.
- Enable the function to add Document Size and Page Count fields to the application group and folder definitions that are created.

Content Manager OnDemand concepts

Content Manager OnDemand has unique definitions for the terms *application*, *application group*, *folder*, and *cabinet*.

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

Content Manager OnDemand applications

An application describes the physical characteristics of a report to Content Manager OnDemand.

Typically, you define an application for each program that produces output that will be stored in Content Manager OnDemand. The application includes the following information:

- Format of the data
- Orientation of data on the page
- Record length
- Code page of the data

The application also includes parameters that the indexing program uses to locate and extract index data and processing instructions that Content Manager OnDemand uses to load index data in the database and documents on storage volumes.

Content Manager OnDemand application groups

An *application group* contains the storage management attributes of and index fields for the data that you load into Content Manager OnDemand.

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

Content Manager OnDemand folders

A *folder* enables users to query and retrieve data stored in Content Manager OnDemand.

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

Figure 2 on page 5 illustrates applications, application groups, and folders.

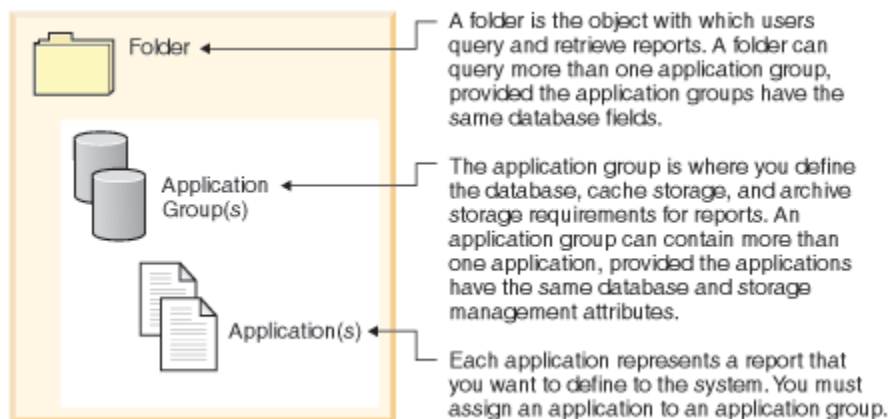


Figure 2. Folders, application groups, and applications (part 1 of 2)

Figure 3 on page 6 shows an example.

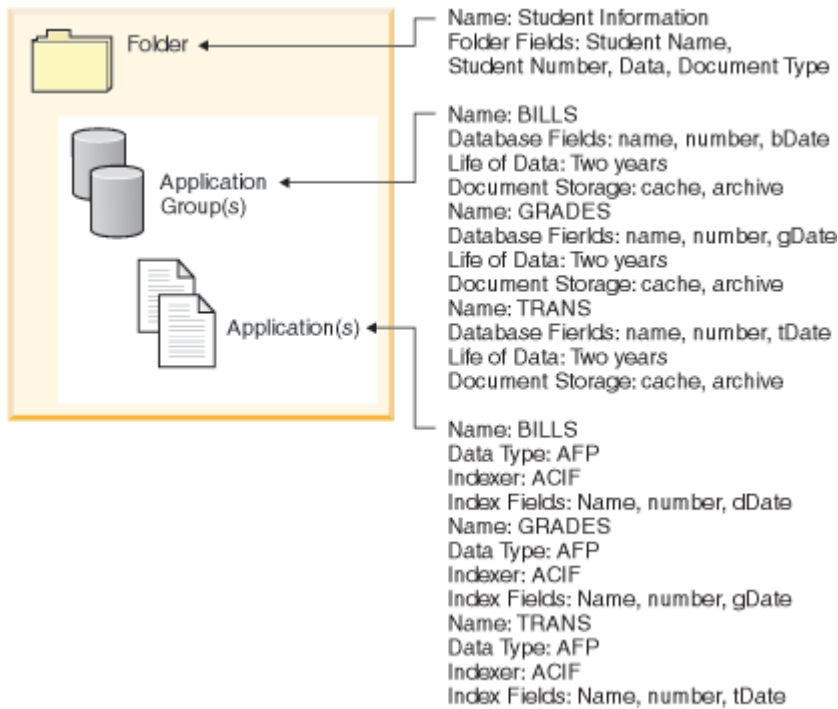


Figure 3. Folders, application groups, and applications (part 2 of 2)

Content Manager OnDemand cabinets

Cabinets are an optional feature that enable users to navigate to folders more easily.


If users have many folders, they might find it helpful to group their folders into cabinets.

Cabinets follow these rules:

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

Figure 4 on page 7 demonstrates how cabinets can be used to organize the folders that a user needs for generating different types of reports. In this example, the user needs to collect information on fund performance, fund balance, and fund transactions for a monthly report. The Monthly Report cabinet contains folders for each type of information that the user needs to collect. The user also needs to generate investment performance reports for clients. Investment performance reports include information on stock performance, bond performance, and fund performance. The Client Report cabinet contains folders for stocks, bonds, and funds. Both cabinets contain the folder for fund performance because the user needs information on fund performance to generate both reports.

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



A folder can belong to more than one cabinet.

Figure 4. Cabinets are used to organize folders

Indexing methods

Content Manager OnDemand provides two ways to index data: document indexing and report indexing

Document Indexing

Used for reports that contain logical items such as policies and statements. Each of the items in a report can be individually indexed on values such as account number, customer name, and balance. Content Manager OnDemand supports up to 128 index values per item. With document indexing, the user does not necessarily need to know about reports or report cycles to retrieve a document from Content Manager OnDemand.

Report indexing

Used for reports that contain many pages of the same kind of data, such as a transaction log. Each line in the report usually identifies a specific transaction, and it is not cost effective to index each line. Content Manager OnDemand stores the report as groups of pages and indexes each group. When reports include a sorted transaction value (for example, invoice number), Content Manager OnDemand can index the data on the transaction value. This is done by extracting the beginning and ending transaction values for each group of pages and storing the values in the database. This type of indexing lets users retrieve a specific transaction value directly.

Content Manager OnDemand documents

Content Manager OnDemand documents represent indexed groups of pages.

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

Figure 5 on page 8 illustrates Content Manager OnDemand applications and documents. An administrator defines the BILLS application for a report that contains logical items, such as customer

statements. The BILLS application uses the document indexing method to divide the report into documents. Each statement in the report becomes a document in Content Manager OnDemand. Users can retrieve a statement by specifying the date and any combination of name and number.

An administrator defines the TRANS application for a report that contains lines of sorted transaction data. The TRANS application uses the report indexing method to divide the report into documents. Each group of 100 pages in the report becomes a document in Content Manager OnDemand. Each group is indexed by using the first and last sorted transaction values that occur in the group. Users can retrieve the group of pages that contains a specific transaction number by specifying the date and the transaction number. Content Manager OnDemand retrieves the group that contains the value entered by the user.

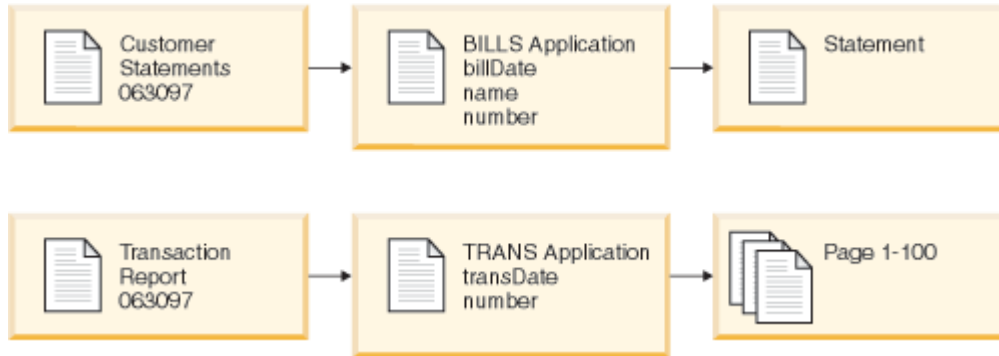


Figure 5. Applications and documents

Content Manager OnDemand servers

The Content Manager OnDemand server environment includes a library server and one or more object servers residing on one or more systems connected to a TCP/IP network.

The library server maintains a central database about the reports stored in Content Manager OnDemand. The database also contains information about the objects defined to the system, such as users, groups, printers, application groups, applications, folders, and storage sets. The database manager provides the database engine and utilities to administer the database. The library server processes client logons, queries, and print requests and updates to the database. The major functions that run on the library server are the request manager, the database manager, and the server print manager.

An object server maintains documents in cache storage volumes (HFS or ZFS) in addition to VSAM files and OAM objects. An object server loads data, retrieves documents, and expires documents. The major functions that run on an object server are the cache storage manager, Content Manager OnDemand data loading and maintenance programs and optionally, the archive storage manager (OAM or VSAM).

The basic Content Manager OnDemand configuration is a library server and an object server on the same z/OS system. This single library and object server configuration supports the database functions and cache storage on the same system. You can also configure your Content Manager OnDemand system with the library server on one system and one or more object servers on different systems. This configuration is known as a distributed library and object server system. The distributed library and object server configuration supports caching of documents on different servers.

The Content Manager OnDemand server environment contains the following components:

Request manager

Provides client, network, and operating system services with security and accounting. The request manager resides on the library server.

Database manager

Maintains the index data for the reports you store on the system. The database manager is the relational database management product, DB2® (Content Manager OnDemand). The database manager resides on the library server.

Database control information

Information about the users, groups, application groups, applications, folders, storage sets, and printers that you define on the system. The control information determines who can access the system, the folders that a user can open, and the application group data that a user can query and retrieve. The database resides on the library server.

Cache storage manager

Maintains documents in the Hierarchical File System (HFS or ZFS). Cache storage is for high-speed access to the most frequently used documents.

Archive storage manager

An optional part of the system. The archive storage manager is for the long-term storage of one or more copies of documents on archive media, in the form of either VSAM files or OAM objects.

Data indexing and conversion programs

Programs that extract index data from input files or generate index data and, depending on the indexer, optionally collect resources and transform input data from one format to another. Content Manager OnDemand provides several indexing programs:

- The AFP Conversion and Indexing Facility (ACIF) can be used to index line data, ASCII data, and AFP input files. ACIF can collect the resources that are required to view AFP documents and convert line data input into AFP data to be stored on the system.
- The 390 Indexer is particularly useful for users who are migrating their data from Content Manager OnDemand for OS/390®, Version 2.1. The 390 Indexer supports the Version 2.1 report types of DOC, PAGE, PDOC, NODX and AFP. The 390 Indexer also supports the Version 2.1 style of Input, Index and Anystore Exits.
- The IBM Content Manager OnDemand PDF Indexer for Multiplatforms can be used to create index data for Adobe Acrobat PDF input files.
- The IBM Content Manager OnDemand Generic Indexer can be used to create index data for almost any other type of data that users want to store on the system, such as Hypertext Markup Language (HTML) files and compressed and uncompressed Tagged Image File Format (TIFF) files.
- The IBM Content Manager OnDemand XML indexer can be used to index and load XML input data.

The indexing programs can run on any Content Manager OnDemand server, and if necessary, the index data can be transferred to the z/OS system.

Data loading programs

Programs that can be set up to automatically store report data into application groups and update the database. The data loading programs can run on any Content Manager OnDemand server.

Server print facility

A facility that allows users to reprint a large volume of documents at high speed.

Content Manager OnDemand management programs

Programs to maintain the Content Manager OnDemand database and documents in cache storage.

System logging facility

Facility that provides administrators with tools to monitor server activity and respond to specific events as they occur. The interface to the system logging facility is through the System Log folder and the System Log user exit.

Request manager

The request manager processes search requests from Content Manager OnDemand client programs.

When a user enters a query, the client program sends a request over the network to the request manager. The request manager works with the database manager to compile a list of the items that match the query and returns the list to the client program. When the user selects an item for viewing, the request manager sends a retrieval request to the appropriate storage manager: the cache storage manager, if the document resides in cache storage, or the archive storage manager, if the document resides in archive storage. The appropriate storage manager retrieves the document and, optionally, the resources associated with the item. The Content Manager OnDemand client program extracts and displays the document.

Content Manager OnDemand management programs include utilities that maintain the database and cache storage. The utilities can automatically migrate data from the database and cache storage volumes to archive storage. These programs use the services of the request manager to manage index data, documents, and resource files.

When a user logs on to the system, Content Manager OnDemand assigns a unique transaction number to that instance of the client program. All activity associated with that instance of the client program contains the same transaction number. The request manager records messages that are generated by the various Content Manager OnDemand programs in the system log, for example, logon, query, and print. These system log messages contain the transaction number, user ID, time stamp, and other information. Administrators can open the system log folder and view the messages. Content Manager OnDemand also provides a system log user exit so that you can run a user-defined program to process messages. For example, you can design a user-defined program to send an alert to an administrator when certain messages appear in the system log. The messages in the system log can also be used to generate usage and billing reports.

Database manager

Content Manager OnDemand uses DB2 to maintain index data and Content Manager OnDemand system tables.

Content Manager OnDemand uses a database management product, Db2, to maintain the index data for the reports that you load into the system. The database manager also maintains the Content Manager OnDemand system tables that describe the applications, application groups, storage sets, folders, groups, users, and printers that you define to the system. You should periodically collect statistics on the tables in the database to optimize the operation of the Content Manager OnDemand database.

Content Manager OnDemand storage manager

The Content Manager OnDemand cache storage manager maintains a copy of documents in the HFS and, optionally, the archive storage manager migrates a copy of the documents into VSAM files or OAM objects.

SMS is required to maintain VSAM data sets that are used for archive storage of Content Manager OnDemand documents.

Data indexing and loading

The reports that you store in Content Manager OnDemand must be indexed. Content Manager OnDemand supports several types of index data and indexing programs.

For example, you can use ACIF to extract index data from the reports that you want to store on the system. An administrator defines the index fields and other processing parameters that ACIF uses to locate and extract index information from reports. Content Manager OnDemand data loading programs read the index data generated by ACIF and load it into the Content Manager OnDemand database. The data loading programs obtain other processing parameters from the Content Manager OnDemand database, such as parameters used to segment, compress, and store report data in cache storage and on archive media. If you plan to index reports on a Content Manager OnDemand server, you can define the parameters with the administrative client. The administrative client includes a *report wizard* that helps you create ACIF indexing parameters by visually marking up sample report data. Content Manager OnDemand indexing programs also provide various exit points to capture other types of source data, such as TIFF images.

[Figure 6 on page 11](#) shows an overview of the data preparation process.




Figure 6. Data preparation, indexing, and loading

In Figure 6 on page 11, user-defined application programs generate printed reports and save report data to disk.

Figure 7 on page 11 shows an overview of the data indexing and loading process.




Figure 7. Data indexing and loading process

The Content Manager OnDemand data loading program first determines whether the report needs to be indexed. If the report needs to be indexed, the data loading program calls the appropriate indexing program. The indexing program uses the indexing parameters from the Content Manager OnDemand application to process the report data. The indexing program can extract and generate index data, divide the report into indexed groups, and collect the resources required to view and reprint the report. After indexing the report, the data loading program processes the index data, the indexed groups, and the resources that use other parameters from the application and application group. The data loading program works with the database manager to update the Content Manager OnDemand database with index data extracted from the report. The data loading program also works with the cache storage manager to segment, compress, and copy report data to HFS cache storage files.

Related information

[Indexing Content Manager OnDemand](#)

Content Manager OnDemand management programs

Content Manager OnDemand management programs help you optimize use of cache storage by managing the migration of data from cache storage to archive store.

When someone in your organization creates an application group, they specify parameters that Content Manager OnDemand programs use to maintain the report data that is stored in the application group. For example, when creating an application group, the administrator specifies the length of time that documents should be maintained on the system and whether index data should be migrated from the database to archive storage. Content Manager OnDemand programs use the information to migrate documents from cache storage to archive storage, delete documents from cache storage, migrate index data from the database to archive storage, and delete index data from the database. These functions are useful because Content Manager OnDemand can reclaim the database and cache storage space released by expired and migrated data. It is recommended that you configure your Content Manager OnDemand

system to automatically start these management programs on a regular schedule, usually once every night or once a week, depending on the workload on the system and other installation-dependent requirements.

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

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

Content Manager OnDemand Web Enablement Kit

The Content Manager OnDemand Web Enablement Kit (ODWEK) provides a set of programming interfaces that can search and retrieve documents from Content Manager OnDemand servers.

ODWEK allows users to access data that is stored in Content Manager OnDemand with IBM Content Navigator or a user-written program. An application that uses the Java APIs can verify permissions, manage hit lists and return data. For example, ODWEK verifies that the user information is valid on the Content Manager OnDemand server, such as permission to access the server and data stored in an application group. After the user submits a search, ODWEK Java API returns a list of the documents that match the query. The user selects a document to view and IBM Content Navigator or the user-written program ends the document to the browser.

The diagram shows a workstation with a web browser that is being used to access data from a Content Manager OnDemand server.



Figure 8. Accessing data stored in Content Manager OnDemand using ODWEK

ODWEK can search for and retrieve documents from the Content Manager OnDemand server.

ODWEK contains several components:

- The ODWEK Java Application Programming Interface (Java API). The Java API uses standard Content Manager OnDemand interfaces and protocols to access data stored in a Content Manager OnDemand server. No additional code is needed on the Content Manager OnDemand server to support ODWEK. The Java API provides a way to access Content Manager OnDemand data from a user-written program. The Java API also provides a way to access Content Manager OnDemand data from software such as IBM Content Navigator.
- The IBM Content Manager OnDemand AFP Web Viewer. The AFP Web Viewer lets users search, retrieve, view, navigate, and print AFP documents from a web browser.
- The HTML5 Line Data Viewer. The HTML5 Line Data Viewer lets users view line data documents from a web browser.

Viewing and transforming documents

To view other types of documents that are stored in Content Manager OnDemand, you must obtain and install the appropriate viewer.

For example, to view Adobe Portable Data Format (PDF) documents, it is recommended that you obtain the Adobe Acrobat viewer for the browsers that are used in your organization.

To convert AFP documents that are stored in Content Manager OnDemand into PDF documents, a transform is required.

Chapter 2. Preparing for Content Manager OnDemand

This section contains an outline that you might find helpful as you prepare your organization for the Content Manager OnDemand environment and perform a pilot release of the system.

About this task

- Work with a single department or group of end-users. Send a memo to the users to explain how Content Manager OnDemand will affect their daily work.
- Develop an end-user training course or contact IBM for help with training for Content Manager OnDemand.
- Establish a support plan for the users. The plan should include the names and phone numbers of persons to contact for assistance and a list of troubleshooting tips.
- Develop a set of evaluation and completion criteria that you can use to compare against the actual performance of the system.
- Choose a report or set of reports for an initial migration to Content Manager OnDemand. Obtain hardcopy of the reports.
- Review the reports and determine the type of indexing required. Then select the fields from the reports for index, filter, and display fields.
- Review the selections with the users. Verify that the index, search, and display fields allow the users to retrieve the data that they need.
- Determine the viewing requirements of your users.
- Identify the type of data contained in the report and determine how you will create index data.
- Determine whether you will use ACIF. If you plan to use ACIF to index the report, you must identify the resources used by the report. Resources are reusable objects found on pages of a report, such as overlays and page segments. Overlays contain constant data that is merged with variable report data during printing and viewing. Page segments are graphics and images that appear on pages of a report file, such as a company logo. Resources can be used by different applications in Content Manager OnDemand.
- Configure cache storage and archive storage (optical and tape storage devices) on the Content Manager OnDemand servers. Define and configure archive media devices to OAM or VSAM. Define storage management policies to OAM or VSAM to support the reports that you plan to store on the system. VSAM data sets must be managed by SMS.
- Use the Content Manager OnDemand administrative client to create the application groups and applications required to support your reports.
- Use the administrative client to define the folders that users open to access data stored on the system.
- Use the administrative client to define users and groups to Content Manager OnDemand.
- Index the reports.
- Load the report, resources, and index data into the application group.
- Survey the users about initial testing and index, search, and display fields.
- Collect additional information from users, report suppliers, production scheduling, and capacity planning. For example:
 - The frequency with which a report is generated and must be loaded into the system
 - The number of pages in a report
 - The number of indexed items, such as statements, contained in a report
 - The access frequency and patterns of your users
 - The length of time until a version of a report is out of date; the length of time that you want Content Manager OnDemand to maintain a report on the system

- The number of copies of a report that must be maintained on the system
- Use the administrative client to update Content Manager OnDemand with the information that you collect.
- Survey users about their satisfaction with Content Manager OnDemand. Compare the performance of the system with the evaluation and completion criteria that you established. Prepare a list of issues to resolve.
- Update your company's vital records list to include the hardware and software required by the Content Manager OnDemand system. Update your company's operations and recovery manuals with information required to operate, support, and backup the Content Manager OnDemand system.

Administrative roles and responsibilities

Content Manager OnDemand administrators are responsible for ensuring that the hardware and software required by Content Manager OnDemand is installed; defining, creating, and maintaining the components of Content Manager OnDemand; and controlling access to Content Manager OnDemand.

The following list summarizes the areas of responsibility of the Content Manager OnDemand administrator:

- Hardware includes the z/OS system, client workstations, printers and the networking equipment.
- Software includes the base operating system, prerequisite software, and client and server programs, configuration files and shell scripts.
- Administrators define Content Manager OnDemand applications and decide how Content Manager OnDemand manages data on the servers.
- Administrators define Content Manager OnDemand groups and users to the system and make sure that the client software is installed and operating properly.

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

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

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

- Installing and upgrading equipment
- Installing and maintaining Content Manager OnDemand programs and other software
- Defining and labeling storage volumes
- Monitoring the space used by the database and the space available on the system
- Monitoring the space used for cache storage and the space available on the system
- Monitoring the space used for archive storage and the space available on the system
- Scheduling jobs to maintain the database, cache storage, and archive storage
- Working with users to determine report indexing and retrieval requirements
- Defining storage sets and storage nodes
- Defining Content Manager OnDemand system printers

- Defining reports to the system
- Defining Content Manager OnDemand groups and users
- Loading reports on the system
- Managing the backup and recovery process for the database and other areas that contain data critical to the operation of the system
- Monitoring server activity and tuning system parameters
- Solving server, network, and application problems
- Answering end-user questions
- Establishing security and audit policies, for example set and maintain passwords and permissions; use OnDemand's audit facilities to monitor application group and user activity; develop, document, and maintain change control procedures to prevent unauthorized changes to the system

Content Manager OnDemand provides an administrative client to enable you to maintain Content Manager OnDemand objects through an easy-to-use, graphical user interface. The administrative client is available as a 32-bit or 64-bit Windows application. With the administrative client you can define and maintain application groups, storage sets, storage nodes, folders, cabinets, system printers, applications, groups, and users. The administrative client includes features that enable you to process sample report data and create ACIF indexing parameters and logical views by visually marking up a sample of a report.

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

Content Manager OnDemand application programming interfaces

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

Client customization

You can customize the Content Manager OnDemand Windows client to access data stored in Content Manager OnDemand.

The OnDemand client can be customized by specifying command line parameters, by invoking and manipulating the client from another application with the Dynamic Data Exchange (DDE) interface, by invoking another application's DLL file from the client, by retrieving related documents, by creating a Program Information File (PIF), and by auditing documents using the Document Audit Facility.

An example of invoking another application's DLL file is the integration of Monarch software with the OnDemand client. Users can view Content Manager OnDemand documents in Monarch and then perform complex data manipulation such as creating derived columns and generating charts and reports.

The OnDemand client (32-bit only) can also be customized by using the Object Linking and Embedding (OLE) APIs. The APIs are used by an application to make requests to the OnDemand client to perform functions such as logon to a server, open a folder, search for a document, view a document, and so on.

Another form of client customization is the use of the Content Manager OnDemand Web Enablement Kit (ODWEK) Java APIs. The Java APIs are used by a user-written program or IBM Content Navigator to search for, retrieve, and view documents.

Content Manager OnDemand server programs

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

- The ARSDOC program is a multipurpose document processing program. You can use the ARSDOC program to query the library server and generate a list of items that match a query; retrieve documents from the system; add, delete, and update documents; and send documents to the server print facility.

You can run the ARSDOC program in a variety of ways under Unix System Services and the z/OS environment.

- The ARSEXPDIR program can be used to process SMF records that indicate Content Manager OnDemand objects have expired and remove the associated index entries for those objects. The program uses the SMF type 65, type 85, and the user type records written by ARSSMFWR to determine which objects were deleted. The ARSEXPDIR program then instructs the server to remove the index entries.

If you do not want to configure SMS management policies to expire documents using the values that you specified in the Content Manager OnDemand application group, you can add an exit to the storage manager that writes a record that causes Content Manager OnDemand to clean up the indexes.

Use the ARSEXPDIR program and the collection of SMF data that reflects the object deletions. The installation must collect type 65 SMF records if you are using VSAM. If you are using OAM, collect type 85 and install ARSSMFWR as the CBRHADUX OAM auto-delete exit.

- The ARSTBLSP program can be run to change the table that Content Manager OnDemand loads data into. During normal operation, Content Manager OnDemand loads index rows into a table until the maximum rows value for the application group is reached. Such a table is said to be open for loading. When the maximum rows value is reached, the table is closed and a new table and table space are created. Under certain circumstances, an installation can close a table to loading before the maximum rows value is reached. For example, migration processing (by entering `arsmaint -e`) will not process a table that is open for loading, and the installation might migrate the table earlier than initially anticipated.

Related information

[Administering](#)

Content Manager OnDemand server logging

You can configure Content Manager OnDemand to save information, warning, and error messages in the system logging facility.

System logging facility

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

Content Manager OnDemand system log user exit

You can configure Content Manager OnDemand to send messages to the ARSLOG installation exit to issue WTOs, compile statistics, and generate accounting information.

Content Manager OnDemand can log messages about system activity, such as when users log on and log off the system and application group activity, such as when clients query and retrieve data. In addition, you can configure Content Manager OnDemand to send the messages to the ARSLOG installation exit. You can use the ARSLOG installation exit to issue WTOs, compile statistics, and generate accounting information. The ARSLOG installation exit is implemented as an MVS™ dynamic exit.

Content Manager OnDemand user security exit

The Content Manager OnDemand user security exit allows an installation to augment the processing of security-related activities or events.

The following list describes the security-related activities or events:

- Logon

- Change Password
- Add User ID or Delete User ID by using the Content Manager OnDemand administrative functions
- Access to a Content Manager OnDemand folder
- Access to a Content Manager OnDemand application group

An installation written exit routine (or set of exit routines) can interact with a security system such as RACF® to determine if the given activity is to be allowed or disallowed.

Unified login exit

The Content Manager OnDemand unified login exit (ARS.PTGN) enables a user to run the Content Manager OnDemand command line utilities (such as ARSLOAD) without needing to specify a user ID and password.

When enabled, if the user does not specify a Content Manager OnDemand user ID and password, the current user ID as returned by UNIX System Services will be used to log onto Content Manager OnDemand. If that fails, the user will be prompted for a user ID and password.

This facility to log on without a password uses the ability to specify a PassTicket as a password when using a RACROUTE REQUEST=VERIFY. Because the function to generate a Passticket is not a part of the SAF interface, Content Manager OnDemand implements the call to generate the PassTicket as an MVS dynamic exit. A sample exit (ARSPTGN) is provided that generates a PassTicket in the RACF environment.

If you use an external security product when you install Content Manager OnDemand, you will need to evaluate the supplied exit and possibly modify the exit for their environment.

Report Specifications Archive Definition exit

With the Report Specifications Archive Definition exit, you can modify some of the parameters used by Content Manager OnDemand when document data is being captured (loaded) by the ARSLOAD program.

The following parameters can be modified:

- The Application Group name
- The Application name
- The name of the Object Server to be used for data storage
- The name of the Storage Node to be used for data storage
- The indexer parameters set
- The input file control character type, logical record length and record format

See the *IBM Content Manager OnDemand for z/OS: Configuration Guide* for more information about the Report Specifications Archive Definition exit.

Client preview exit

The Content Manager OnDemand client preview exit allows an installation to process document data before the document is presented to the client. The client preview exit can be used to add, remove, or reformat data before the document is presented to the client. For example:

- Remove pages from the document, such as banner pages, title pages, all pages but the summary page, and so on.
- Remove specific words, columns of data, or other information from the document. That is, omit sensitive information such as salaries, social security numbers, and birth dates.
- Add information to the document, for example, a summary page, data analysis information, and confidential or copy statements.
- Reformat data contained in the document, for example, reorder the columns of data.

COBOL bindings are provided for the exit. This allows installations to use COBOL as the high-level language interface for the exit. The COBOL support requires Enterprise COBOL for z/OS Version 3.4 or

above because the exits can run in a POSIX thread environment and must use the **THREAD** compiler option.

See the *IBM Content Manager OnDemand for z/OS: Configuration Guide* for more information about the client preview exit.

Table space creation exit

The table space creation exit allows an installation to take action when Content Manager OnDemand is going to create a table space, table, or index tables that will be used to store application index data. The table space creation exit is not called for the Content Manager OnDemand system tables.

For table and index creation, the installation can alter the SQL statements that will be used to create the table or index.

COBOL bindings are provided for the exit. This allows installations to use COBOL as the high-level language interface for the exit. The COBOL support requires Enterprise COBOL for z/OS Version 3.4 because the exits can run in a POSIX thread environment and must use the **THREAD** compiler option.

See the *IBM Content Manager OnDemand for z/OS: Configuration Guide* for more information about the table space creation exit.

JES Spool Data Capture facility

The JES (Job Entry Subsystem) Spool Data Capture facility provides a means to collect and consolidate JES Spool (SYSOUT) data sets into files so the JES Spool data sets can be archived by Content Manager OnDemand.

This facility, known by its program name, ARSYSPIN, runs as a started task in its own address space. A control statement file is used to provide ARSYSPIN parameters. These parameters specify JES Spool file selection criteria (for example, from which SYSOUT classes output is to be selected) and other operational characteristics.

ARSYSPIN creates an intermediate output file that contains one or more Spool files from one or more jobs. Each captured spool file is bracketed by a pair of separator records. A "Begin" separator record is written before the first record of the spool file and an "End" separator record is written after the last record of the spool file data. The separator records contain information that can be used to construct index values to facilitate the retrieval and viewing of the captured data sets after they are stored by Content Manager OnDemand.

The intermediate output file is indexed and stored in Content Manager OnDemand by using the ARSLOAD program. ARSYSPIN invokes ARSLOAD when sufficient data is captured in the intermediate output file. ARSLOAD indexes the data by calling an indexer program (usually APKACIF) to extract the index values from the data and store them in an index file. ARSLOAD then adds the indexes to the database and loads the data onto archive media.

A sample APKACIF input exit is provided with ARSYSPIN that illustrates a technique for inserting a data record into the APKACIF input stream to provide additional indexing information. This sample exit scans the JES Job Log and System Messages spool files (if they are present) to locate critical system and job processing messages (for example, messages that indicate the execution of job steps that are terminated by ABEND or with nonzero completion codes), extract related values and place this information in the inserted record for subsequent indexing. After the captured Spool data is stored into Content Manager OnDemand, ARSYSPIN repeats the process of collecting and consolidating other spool files so that they too can be archived in Content Manager OnDemand. The cycle of collecting, consolidating, and loading data continues until the program is stopped by the MVS **STOP** command.

See the *IBM Content Manager OnDemand for z/OS: Configuration Guide* for more information about the JES Spool Data Capture facility.

Indexer exits

The 390 Indexer can be used to extract index data from and generate index data about line data and AFP reports. In addition, other data types, such as TIFF images, can be captured by using the Anystore Exit.

The 390 Indexer provides the following exits to allow you to run a user-written program to process report data before it is stored in the system.

- Anystore Exit. The Anystore Exit captures any type of data. The exit reads the data to be captured, breaking it into documents, and determining the index values. A sample Anystore Exit is provided that captures TIFF images using a pre-generated set of indexing instructions read from a separate file.
- Input Exit. The Input Exit processes the report input before the report is stored. This exit can be used only when the INDEXSTYLE is not set to AFP and when the ANYEXIT is not specified. The exit is called dynamically during the report capture process. The report capture routine calls the exit when the indexing parameters specify an input exit name in the **INPEXIT** parameter. The report administrator provides a program name for this parameter. There are no restrictions as to the type of processing that can be performed in an input exit with the exception that the exit must pass the standard parameter list back to the report capture program. Values must be supplied for all parameters.
- Index Exit. The Index Exit modifies the report indexes before insertion into the application group data table. This exit can be used with any type of report captured by the 390 Indexer. The exit is called dynamically during the capture process. The capture program calls the exit when the indexing instructions for the application include the **INDEXEXIT** parameter. The report administrator provides a program name for the Index Exit. There are no restrictions as to the type of processing that can be performed in an index exit with the exception that the exit must pass the standard parameter list back to the capture program. A sample exit is provided in member ARSEXNDX.

See the *Content Manager OnDemand Indexing Reference* for more information about the 390 Indexer exits.

Chapter 3. System requirements

Review your system specifications to help plan the disk storage required to support your system.

For hardware and software requirements, see <http://www.ibm.com/support/docview.wss?uid=swg27049168> or search for 27049168 at <http://www.ibm.com/>

Chapter 4. Storage

You can configure a Content Manager OnDemand system to maintain copies of reports in cache storage and in archive storage, also known as external storage.

The copies in archive storage are for long-term storage. You can use external storage managers such as OAM, VSAM, or cloud storage managers to maintain reports in archive storage. OAM and VSAM support a variety of DASD and tape storage devices. Other external storage managers provide support for cloud and external file system storage.

Storage objects

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

Figure 9 on page 25 shows the primary Content Manager OnDemand storage objects.

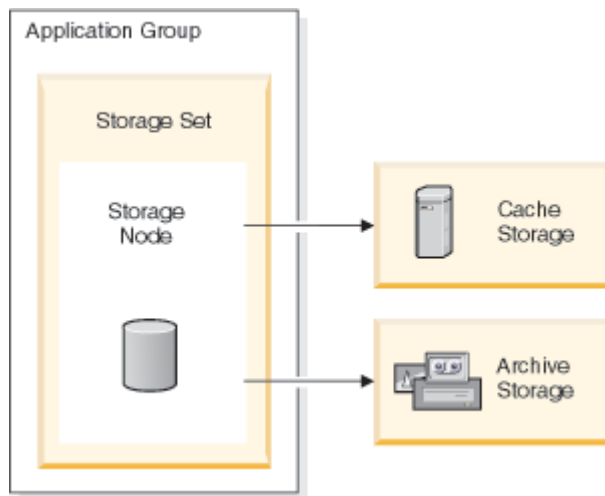


Figure 9. Content Manager OnDemand storage objects

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

When you load a report into Content Manager OnDemand, you assign it to an application group. The application group identifies a storage set. The storage set contains one or more storage nodes. A storage node identifies an object server on which data is stored. Content Manager OnDemand will automatically store a copy of the report in cache storage on the object server, unless you specify otherwise. A storage node can also identify an external storage manager for archive storage. If a storage node specifies OAM, VSAM, or another external storage manager, then Content Manager OnDemand automatically stores a copy of the report in archive storage, which is managed by the external storage manager that is specified in the storage node definition.

One or more application groups can specify the same storage set. However, a storage set can write to only one (archive) storage node at a time. This means that all of the data that is written to a storage node will be maintained using the same *policy*, for example, the type of media, the devices, the length of time to maintain data on the system, and so forth.

If you use an external storage manager to maintain reports in archive storage, you should specify the same storage management criteria to Content Manager OnDemand and the external storage manager. For example, the Life of Data and Indexes, which is used by Content Manager OnDemand, and the retention period used by the external storage manager should be the same value.

Other external storage managers

Content Manager OnDemand supports external cloud storage managers such as Amazon Simple Storage Service (S3), Apache Hadoop Distributed File System (HDFS), Hitachi Content Platform, IBM Cloud Object Storage, Microsoft Azure, and OpenStack Swift. Content Manager OnDemand can also store data to external file systems.

Cloud storage options

The Content Manager OnDemand server can be configured to maintain copies of its stored data in both cache storage, managed by Content Manager OnDemand, and in archive storage, now referred to as external storage. Historically, OAM and VSAM have been the only options used by Content Manager OnDemand to maintain data stored in external storage. The addition of Amazon S3, Apache HDFS, Hitachi Content Platform, IBM Cloud Object Storage, Microsoft Azure, and OpenStack Swift support augments the storage capabilities of Content Manager OnDemand by providing multiple external cloud storage solution options. Cloud storage solutions allow Content Manager OnDemand users to leverage the advantages that such storage provides such as cost savings, data replication, and disaster recovery. This functionality is configured in Content Manager OnDemand and behaves much in the same way that communicating with OAM and VSAM does. This means that data in Content Manager OnDemand can be stored in cache as well as stored in Amazon S3, Apache HDFS, Hitachi Content Platform, IBM Cloud Object Storage, Microsoft Azure, or OpenStack Swift. The storing of data to any external cloud storage manager can take place at the same time that data is written to the Content Manager OnDemand cache or can be scheduled to migrate at a later date.

Amazon S3, Apache HDFS, Hitachi Content Platform, IBM Cloud Object Storage, Microsoft Azure, and OpenStack Swift storage options complement the functionality provided by OAM and VSAM. Content Manager OnDemand servers can be configured to use any combination of Amazon S3, Apache HDFS, Hitachi Content Platform, IBM Cloud Object Storage, Microsoft Azure, OpenStack Swift, OAM, and VSAM.

Additional information for each cloud storage solution can be found on the web at the following locations:

Amazon S3

<https://aws.amazon.com/s3/>

Apache HDFS

<https://hadoop.apache.org/>

Hitachi Content Platform

<https://www.hitachivantara.com/en-us/products/cloud-object-platform/content-platform.html>

IBM Cloud Object Storage

<https://www.ibm.com/cloud-computing/infrastructure/object-storage/>

Microsoft Azure

<https://azure.microsoft.com/>

OpenStack Swift

<http://docs.openstack.org/developer/swift/>

Using a file system for external storage

The Content Manager OnDemand server can be configured to maintain copies of its stored data in a file system accessible to Content Manager OnDemand. As with the cloud storage options, the storing of data to an external file system can take place at the same time that data is written to the Content Manager OnDemand cache or can be scheduled to migrate at a later date.

Detailed configuration information can be found in the *Content Manager OnDemand for z/OS Configuration Guide*.

Chapter 5. Reports and other data

This section contains information that can help you plan for the reports that you will store in Content Manager OnDemand. You can use the information to determine the hardware configuration that you need to support your Content Manager OnDemand system. This section lists questions that you might ask users of the reports, provides information about the types of data that you can store in Content Manager OnDemand, and provides information about indexing reports.

Collecting requirements

Planning for Content Manager OnDemand requires that you understand how the system will be deployed, who will use the system and how users will use it, and other requirements.

About this task

Answers to the following questions provide information that allows you to properly configure your Content Manager OnDemand system, including the storage and network configuration, to support your applications and users:

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

Input data formats

Content Manager OnDemand supports several types of input data:

- AFP print data streams (AFP or MO:DCA-P), including line data mixed with AFP structured fields and line data formatted with a page definition.
- Unformatted ASCII data that is typically generated in the workstation environment.
- Adobe Portable Data Format (PDF) files.
- Image files in the following formats:
 - BMP (Bitmap). A file that contains a bitmapped graphic.
 - GIF (Graphic Interchange Format). A bitmapped color graphics file format.

- JFIF (JPEG Format Image File). A file that contains image data compressed using the JPEG (Joint Photographic Experts Group) standard.
- PCX (Picture Exchange Format). A file that contains a graphic in the PCX file format, widely used by PC applications, such as the PC Paintbrush program. Compressed using PackBytes compression.
- TIFF (Tagged Image File Format). A bitmapped graphics image format for scanned images with resolutions up to 300 dots per inch. TIFF simulates gray-scale shading. Content Manager OnDemand supports single and multipage TIFF images that are uncompressed or are compressed using JPEG, CCITT Group 3, CCITT Group 3/2D, or CCITT Group 4 compression.
- Extensible Markup Language (XML) files.

In addition to the types of data described previously, with Content Manager OnDemand you can store almost any other type of data on the system. For example, you can define an application for HTML documents. When you define the application, you must identify the file type of the data. The file type determines the program that the client starts when the user retrieves a document. For example, if the file type is HTM, then the client can start a Web browser to view the document.

In the z/OS environment, Content Manager OnDemand allows application programs that produce 1403 or 3211 data streams to take advantage of overlays, page segments, and typographic fonts. This is done using a page definition that specifies how data are mapped on the page. The definition allows text to be moved to different positions on the page, fonts to be changed, and conditional processing. When combined with a form definition, the page definition allows sophisticated pages to be produced by existing line data applications without changing the application that generates the data.

You can use ACIF to convert line data to AFP data before loading it into the system. The resulting AFP data can add color or an electronic form to line data, making presentation of the information more effective. However, archiving line data without conversion usually results in much higher compression ratios.

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

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

Indexing data

One of the main operations that you do with Content Manager OnDemand is to index reports.

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

Index information can be added to reports at the same time that the application program generates the print data or, more typically, the output print data can be processed by one of the indexing programs that are supported by Content Manager OnDemand.

When you index a report, you can divide a large report into smaller, uniquely identifiable units of information. For example, when an application program generates customer bills, it can produce a large print stream made up of thousands of individual customer bills. With Content Manager OnDemand, you can identify the individual customer bills within the report as smaller, separate information units, or logical items (known as documents in Content Manager OnDemand). Your users can search for and retrieve the logical items by using identifiers such as account number, customer name, and date.

Content Manager OnDemand supports two general methods of indexing:

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

If a report does not contain logical items or sorted line data, it can usually be indexed by using the report indexing method.

See the *Content Manager OnDemand Indexing Reference*, for details about and examples on using the indexing programs that are provided with Content Manager OnDemand.

Document indexing

Document indexing indexes reports that are made up of logical items or indexes reports that contain unique values such as an account number or a customer name. When searching and retrieving these types of reports, Content Manager OnDemand returns a list of the items that match the user's query and transfers the individual items to the Content Manager OnDemand client program for viewing and printing. Content Manager OnDemand supports up to 128 fields as indexes or filters for document-type data. The fields do not have to be sorted and can contain numeric or text information. The fields are stored in the database as indexes or filters. [Figure 10 on page 29](#) shows an example of a report file and document indexing.

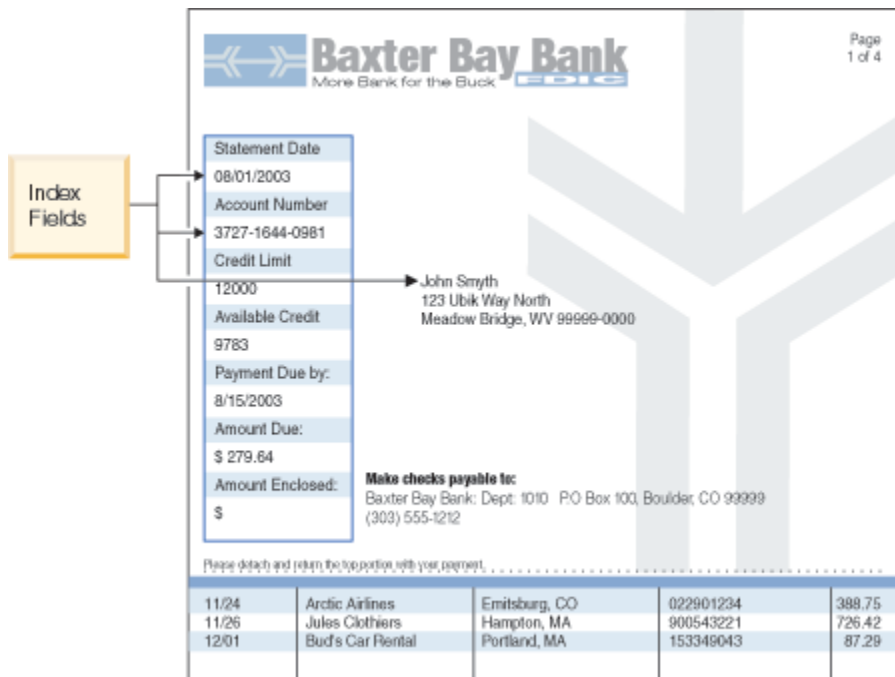


Figure 10. Document indexing method

Report indexing

Report indexing allows users to search sorted report data and retrieve the first occurrence of the value that they specified in the query. Content Manager OnDemand divides the report data into groups of pages and stores the first and last index values contained in each group of pages in the database. When the user enters a query, Content Manager OnDemand returns a list of the items that match the query. When the user selects an item for viewing, Content Manager OnDemand searches within the item for the value specified by the user. The Content Manager OnDemand client program displays the first page that contains the value specified by the user. Content Manager OnDemand uses a single, unique sorted index value for the retrieval of the report data, for example, an invoice number or a transaction identifier. [Figure 11 on page 30](#) shows an example of a report file and report indexing.

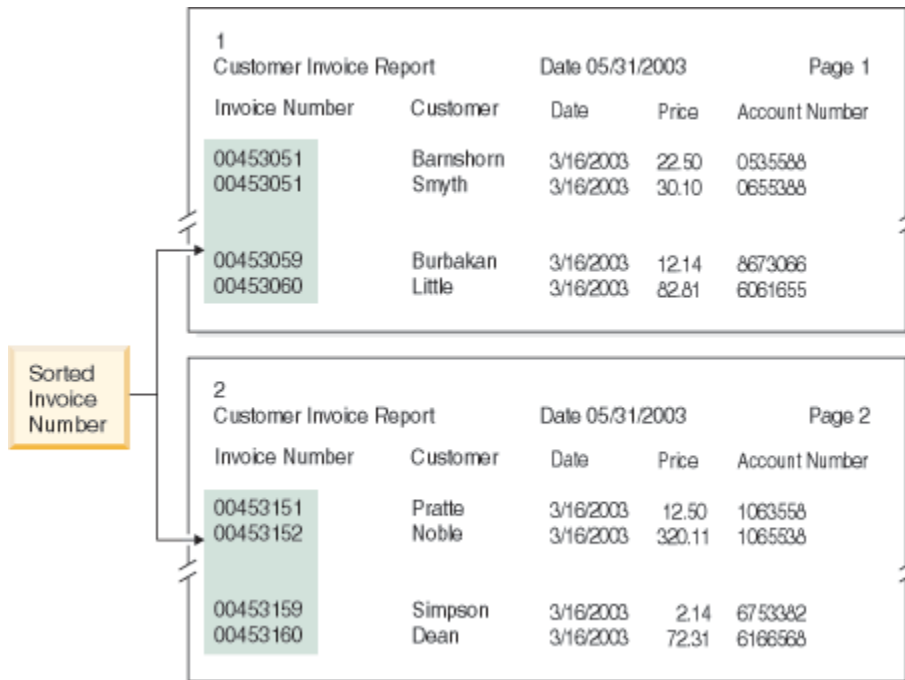


Figure 11. Report indexing method

Indexing data with ACIF on z/OS

About this task

ACIF is a powerful tool for indexing the print data streams of z/OS application programs. ACIF indexes reports based on the organization of the data in the report. You can optionally convert line data print streams into AFP data. ACIF processes three input sources:

- Indexing parameters that specify how the data should be indexed. You can create the indexing parameters when you define a Content Manager OnDemand application.
- AFP resources required to view and print the data if the data was created by an AFP application.
- The print data stream.

The output of ACIF is either a fully composed AFP data stream or the original line data input. ACIF can convert line data input to AFP data, can produce an index file that Content Manager OnDemand uses to create index data for the database, and optionally can collect resources into a resource group file.

ACIF produces a resource group file for AFP data. To create a resource group file, ACIF must have access to the resources required by the input data stream. Content Manager OnDemand usually stores the resources in cache storage and retrieves the resources associated with a specific document when a user selects the document for viewing.

ACIF can logically divide reports into individual items, such as statements, policies, and bills. You can define up to 128 index fields for each item in a report.

ACIF is designed to index reports that contain line data with a consistent structure and format. You can also use ACIF to index AFP input files that contain indexing controls and information.

Generating index data in applications

About this task

As an alternative to using ACIF to generate the index data, you can create index information in the application that generates the report and then process the report with ACIF. Some application programs already provide support to add indexing information. However, you might need to modify your application to add indexing functions.

Generating index data with AFP application programs

About this task

The IBM Document Composition Facility (DCF) is a product that can be used to create indexed AFP data. DCF prepares and formats documents for printing. DCF also can add both group-level and page-level indexing tags. DCF allows specific indexing information to be included in the output print data stream. You can process the output file created by DCF with ACIF to create an index file that can be processed with the Content Manager OnDemand data loading program.

In addition to DCF, several popular third-party programs can produce indexed AFP data.

Generating index data with the AFP API

About this task

The AFP Application Programming Interface (AFP API) is a product that can be used to index print data. If you know COBOL or PL/I, you can use the AFP API to format complex output without knowing the syntax and semantics of MO:DCA-P. Using the AFP API, you can index AFP files with both group-level and page-level indexing tags, which allows more specific information to be included in the output file. The indexing information is added at the same time that the application program generates the print data. You then process the output file with an indexing program, such as ACIF, to create the index data that the Content Manager OnDemand data loading program stores in the database.

Inserting AFP records in a data stream

About this task

A common way of indexing unstructured, mixed-mode data is to add no operation (NOP) structured fields to the data stream. ACIF can then be used to process the data stream and locate the NOP fields and extract the index values.

Using the Content Manager OnDemand 390 Indexer

The 390 Indexer can index individual documents, ranges of pages containing a sorted column, ranges of pages containing a sorted column within a higher level grouping value and allow the indexer to generate indexes when the report has no obvious index values.

About this task

The 390 Indexer is particularly useful if you are migrating from Content Manager OnDemand for z/OS Version 2.1, as it supports the Version 2.1 report types of DOC, PAGE, PDOC, NODX, and AFP. The 390 Indexer also supports the Version 2.1 style of Input, Index, and Anystore Exits.

The 390 Indexer processes two input sources:

- Indexing parameters that specify how the data should be indexed. You can create the indexing parameters when you define a Content Manager OnDemand application. The parameters are of the same form as used by ACIF but have additional extensions that are unique to the 390 Indexer.
- The print data stream

The 390 Indexer indexes input data based on the organization of the data:

- Document organization. For reports made up of logical items, such as statements, policies, and invoices. The 390 Indexer can generate index data for each logical item in the report.
- Report organization. For reports that contain line data with sorted values on each page, such as a transaction log or general ledger. The 390 Indexer can divide the report into groups of pages and generate index data for each group of pages.
- Anystore Exit. This exit point allows for generating your own index values for each document.
- AFP Data Stream. Fully resolved AFP Data Streams are indexed by using in-stream Tagged Logical Element (TLE) or NOP records for each Page Group.

The 390 Indexer can logically divide reports into individual items, such as statements, policies, and bills. You can define up to 128 index fields for each item in a report.

See the *Content Manager OnDemand Indexing Reference*, for more information about the 390 Indexer.

Using the Content Manager OnDemand Generic Indexer

About this task

Content Manager OnDemand provides the Generic Indexer so that you can specify indexing information for input data that you cannot or do not want to index with ACIF or the other indexing programs. For example, if you want to load word processing documents into the system, the documents can be stored in Content Manager OnDemand in the same format in which they were created. The documents can be retrieved from Content Manager OnDemand and viewed with the word processor. However, because the documents do not contain AFP data, line data, or PDF data, you cannot index them with ACIF, the 390 Indexer, or the PDF Indexer. However, you can specify index information about the documents to the Generic Indexer and load the documents into the system. Users can then search for and retrieve the documents by using one of the Content Manager OnDemand client programs.

To use the Generic Indexer, you must specify all of the index data for each input data set that you want to store in and retrieve from Content Manager OnDemand. You specify the index data in a parameter data set. The parameter data set contains the index fields, index values, and information about the input data sets or documents that you want to process. The Generic Indexer retrieves the indexing information from the parameter data set and generates the index data that is loaded into the database. Content Manager OnDemand creates one index record for each input data set (or document) that you specify in the parameter data set. The index record contains the index values that uniquely identify a data set or document in Content Manager OnDemand.

The Generic Indexer supports group-level indexes. Group indexes are stored in the database and used to search for documents. You must specify one set of group indexes for each data set or document that you want to process with the Generic Indexer. You can define up to 128 index fields for each data set or document.

See the *IBM Content Manager OnDemand for z/OS: Indexing Reference* for more information about the Content Manager OnDemand Generic Indexer.

Using the Content Manager OnDemand XML Indexer

Content Manager OnDemand provides the XML indexer to allow you to specify indexing information for XML input data that you want to be stored in the system.

About this task

You transform your XML documents to a common format defined by IBM Content Manager OnDemand so that the data can be indexed correctly. The files can be stored in IBM Content Manager OnDemand in the same format in which they were created.

To use the XML indexer, you must specify all of the index data for each input file or document that you want to store in and retrieve from the system. The XML indexer requires that a transformation be applied to the source file in order to make certain all your XML data can be indexed. Your XML input data is transformed into a common XML file format using tools such as XSLT and XQuery. XSLT (Extensible Stylesheet Language Transformations) is a language of transforming XML documents into other XML documents, or even plain text. During the transformation the original document is not changed, rather, a new document is created based on the content of an existing one. The basic processing paradigm is pattern matching. The XSLT stylesheet defines what patterns to process and how to process them for output. While there are many processor implementations of XSLT, Saxon and Xalan are two of the more popular open source versions.

For more information, see the *Content Manager OnDemand Indexing Reference*.

Indexing reports using date fields

You can use the date that appears in the report, such as the run date, a transaction date, or the statement date.

About this task

If the data that you want to store in Content Manager OnDemand does not contain a date, you can use the date that the report was loaded into the system.

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

Chapter 6. Content Manager OnDemand objects

When you install and configure the Content Manager OnDemand software, you create and initialize a set of database tables that form the internal framework of the system.

When you define reports to the system, Content Manager OnDemand adds an application group table structure and other control information to the database.

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

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

A *cabinet* can be used to organize folders. Cabinets are an optional feature that enable users to navigate to folders more easily. For example, a cabinet can be used to group folders that a user needs to perform a certain task.

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

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

You must assign unique names to each object within each type of object. For example, each application group object must have a unique name. However, objects of different types can have the same name. For example, a folder object and an application object can have the same name.

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

Folders

A *folder* provides users the means to search for and retrieve related reports stored on the system. Users open folders, construct queries, and retrieve reports from application groups. (However, users do not need to know about application groups.) When you create a folder, you define the search and display fields that appear when the user opens the folder. You map the folder fields to database fields in the application groups referenced by the folder. The database fields contain index values extracted from the reports that are loaded into the application groups. For example, the folder search field **Customer Account Number** can be mapped to the acct# application group database field. Content Manager OnDemand creates database records that include the index values for the acct# field when you load a

report into the application group. When the user enters a query, Content Manager OnDemand retrieves records from the database if the values of the acct# database field match the value that the user typed in the **Customer Account Number** search field.

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

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

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

Cabinets

Cabinets are used to organize folders into useful groups. For example, if users need to retrieve reports from multiple folders to complete a task, you can create a cabinet to help them find task-related folders more easily. When you create a cabinet, you select which folders are grouped together and which users have access to the cabinet.

Cabinets are useful when users work with a large number of folders. Additionally, if users need the same folder to complete multiple tasks, the folder can be added to multiple cabinets.

Application groups

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

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

Content Manager OnDemand supports up to 128 *index* and *filter* fields for each application group:

- Index fields allow fast access to a specific record that uses a key, but generally require a large amount of disk storage to implement and require more time to load data into the application group. Content

Manager OnDemand uses index fields to locate the records in the database that meet the search criteria entered by the user. The index record contains the physical location of an item on a storage volume.

- Filter fields are used to refine queries, retrieving only a subset of the records found with an index field. Filter fields are generally used with an index field to identify a specific item in the application group. Filter fields can also be used to display additional information in the document list, for example, an address.

Content Manager OnDemand requires a *segment field* for each application group that you define. Content Manager OnDemand uses the segment field to organize and maintain application group data and to locate items that match a query. The segment field must be one of the following date or date/time fields:

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

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

- Some reports are comprised of logical groups of information, such as statements, invoices, and policies. These groups, or logical items, can contain one or more pages of information. Content Manager OnDemand can index, store, and retrieve the logical items contained in a report. Each logical item can be indexed on up to 128 values, for example, account number, customer name, and balance. Content Manager OnDemand creates a row in the database for each logical item it finds in the report.
- Other reports can be organized differently, and might not necessarily contain logical items. For example, a report can contain thousands of pages of transaction or general ledger data. Content Manager OnDemand can index, store, and retrieve information from these types of reports using index values such as date, page number, and a sorted value, such as transaction number. Content Manager OnDemand divides these types of reports into groups of pages and indexes each group of pages. While these types of reports might contain logical items, it probably would not be cost effective to index every item in the report. That is, indexing every item in these types of reports would probably result in thousands of index records being added to the database each time that a report is loaded into the application group.

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

Multiple Loads per Database Table

With this method, each time that you load a report into the application group, 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 logical 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 specific size. Content Manager OnDemand maintains a segment table for each application group. The segment table provides faster query performance by limiting searches to a specific table of application group data, using a date value to construct the query. It is recommended that you use this method to organize your database when the users that search for data stored in the application group do not necessarily know or care what version of a report generated the information that they need.

Single Load per Database Table

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

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

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

Database expiration processing

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

Cache migration processing

Cache migration is the process of copying reports from cache storage to archive storage. You specify when a report is copied from cache storage to archive storage when you create an application group. Content Manager OnDemand provides a utility that you can use to copy reports to archive storage. You typically set up the utility to run automatically on a regular schedule. Cache migration optimizes the use of cache storage and increases performance for short-term retrievals of reports. As a report ages, and users no longer search for the report, Content Manager OnDemand can automatically copy the report to long-term (archive) storage. You can also use cache migration to defer the loading of reports to archive storage to a time when there is little or no other system activity.

Cache expiration processing

Cache expiration is the process of deleting reports from cache storage. You specify how long a report should remain in cache storage when you create an application group. Content Manager OnDemand provides a utility that you can use to delete reports from cache storage. You typically set up the utility to run automatically on a regular schedule. Cache expiration reclaims cache storage space taken by expired reports so that the system has space for newer versions of reports.

Applications

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

You can specify default settings for viewing and printing pages of a report at the Content Manager OnDemand application level. For example, if you select a default printer for the application, when a user prints a document that is associated with the application, Content Manager OnDemand sends the document to the printer that you specified. Typically you define an application for each different report that you plan to load into the system.

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

For Content Manager OnDemand applications, you specify information to the indexing and data loading programs, such as the technique that Content Manager OnDemand uses to compress the report file, the parameter used to index the data, and information that Content Manager OnDemand uses to process index data before loading index records into the database. Content Manager OnDemand uses the indexing parameters, options, and data values that you specify to locate index data in and extract index data from the report.

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

Users and groups

Each user logs on to Content Manager OnDemand with a user ID. Content Manager OnDemand authenticates user IDs and determines the usage and administrative authority available to the user based on the user ID. A Content Manager OnDemand user ID does not necessarily have to identify an individual user. However, for accounting purposes, you might want to assign a Content Manager OnDemand user ID to each person that uses the system.

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

Content Manager OnDemand groups organize users by function, authorization, or any other purpose you might require. When you define a Content Manager OnDemand group, you can organize users by department or function and set folder and application group permissions that are common to all of the users assigned to the group. The permissions determine the types of actions that users assigned to the group can perform. You do not need to assign a user to a group, but doing so can simplify administration of users with similar requirements and capabilities.

Permissions

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

If you assign a user to more than one group, the user normally obtains the permissions of all of the groups. For example, using the group settings described in [Table 1 on page 39](#), a user assigned to both groups can access the Student Bills and Student Transcripts folders.

Group	Folders
Accounting	Student Bills
Admissions	Student Transcripts

However, there are exceptions to this rule. See information about permissions in the *IBM Content Manager OnDemand for z/OS: Administration Guide* for details.

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

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

Folder permissions

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

permissions. Setting permissions at the user level provides a specific user with the permissions that you define. User level permissions override group level permissions and folder level permissions.

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

You can set the following types of folder permissions:

Access

Users can open the folder with Content Manager OnDemand client programs and search for and retrieve data from the application groups referenced in the folder.

To search for and retrieve items, users must have access permission for the folder, and access permission to one or more of the application groups referenced in the folder.

Fields

Users can open the folder with Content Manager OnDemand client programs and can modify the folder field information with the administrator interface. Content Manager OnDemand maintains a set of folder fields for each user given fields permission for the folder.

Named Queries

A named query is a set of search criteria, saved by name, that can be selected and loaded into folder search fields. Content Manager OnDemand supports two types of named queries:

- Public - A named query that is available to all users that can open the folder.
- Private - A named query available only to the user that created the named query. Users can be given authority to view, create, modify, and delete named queries.

Administrator

A folder administrator can modify and delete the folder. A folder administrator can change user and group permissions, add and remove users and groups from the folder, and make changes to the folder field information.

Application group permissions

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

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

You can set the following types of application group permissions:

Access

Users can search for and retrieve data that is stored in the application group by using Content Manager OnDemand client programs.

Document

Determines the types of document functions that users can perform. The default document permissions are view, print, and copy.

Annotation

Determines the types of annotation functions users can perform. The default annotation permissions are view, print, and copy.

Logical Views

Determine how Content Manager OnDemand displays report file pages. Users can define their own logical views with Content Manager OnDemand client programs.

Administrator

Modifies and deletes the application group. An application group administrator can change user and group permissions, add and remove users and groups from the application group, change message logging options, update the storage management settings for the application group, and make changes to the application group field information.

Query restriction

Limits access to application group data. You typically set up a query restriction to limit the data that a specific user or group of users can access.

Naming rules

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

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

User naming rules

When naming a user, the name that you specify:

- Can contain from 1 to 128 characters (bytes)
- Cannot include blank characters or any of the following characters: ' * % + _ [] “
- Must be unique to the library server
- By default, Content Manager OnDemand converts lowercase letters in a user name to uppercase (for example, laguarde is stored as LAGUARDE)

If your organization implements the User Security Exit, you can determine the characteristics of user IDs on your system.

Password naming rules

When you create a password, the value that you specify:

- Can contain from 1 to 128 characters (bytes)
 - Content Manager OnDemand security verifies only the first eight characters that are entered by the user. The additional characters are provided for customers who choose to implement their own password security by enabling the User Security Exit.
 - If your organization enables the User Security Exit, you should set the Minimum Password Length option to Permit Blank Password so that Content Manager OnDemand security does not validate passwords that are entered by your users (when they set or change a password). Also, Content Manager OnDemand security ignores the Maximum Password Age option when you enable the User Security Exit.
 - Unless your organization enables the Security User Exit, it is recommended that you specify a value of 8 or fewer for the Minimum Password Length option.
- By default, Content Manager OnDemand converts lowercase letters in a password to uppercase (for example, laguarde is stored as LAGUARDE)

If your organization implements the User Security Exit, then you can determine the characteristics of passwords on your system.

Group naming rules

When you name a group, the name that you specify:

- Can contain from 1 to 128 characters (bytes)
- Cannot include blank characters or any of the following characters: ' * % + _ [] “

- Must be unique to the library server
- Can be mixed case; however, Content Manager OnDemand ignores the case (for example, LaGuarde is the same as laguarde)

Application group naming rules

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

- Can contain from 1 to 60 characters (bytes), including embedded blanks
- Cannot include any of the following characters: ' % _ [] “
- Can be mixed case; however, Content Manager OnDemand ignores the case (for example, LaGuarde is the same as laguarde)
- An application name must be unique to the application group where you assign the application
- An application group or folder name must be unique to the library server

Database field naming rules

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

- Can contain from 1 to 18 characters (bytes)
- Must begin with the letters A through Z
- Can include the letters A through Z, the numbers 0 through 9, and any of the following characters: @ \$ _ #
- Can be mixed case; however, Content Manager OnDemand does not create a unique name (for example, rDate is the same as rdate)
- Must be unique to the application group
- Cannot be any of the Content Manager OnDemand reserved words: annot, comp_len, comp_off, comp_type, doc_len, doc_name, doc_off, doc_type, prt_nid, resource, res_comp_type, sec_nid.
- Cannot be any of the words reserved by the database manager. (For a list of reserved words, see the documentation provided with your database manager product.)

Logical view naming rules

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

- Can contain from 1 to 30 characters (bytes)
- Can be mixed case
- A public view must be unique to the application
- A private view must be unique to the user

Folder field naming rules

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

- Can contain from 1 to 60 characters (bytes), including embedded blanks
- Cannot include any of the following characters: ' % _ [] “
- Can be mixed case
- Must be unique to the folder

Storage set naming rules

When you name a storage set, the name that you specify:

- Can contain from 1 to 60 characters (bytes)

- Can be mixed case; however, Content Manager OnDemand ignores the case (for example, LaGuarde is the same as laguarde)
- Must be unique to the library server

Primary storage node naming rules

When you name a primary storage node, the name that you specify:

- Can contain from 1 to 128 characters (bytes)
- Can be mixed case; however, Content Manager OnDemand ignores the case (for example, LaGuarde is the same as laguarde)
- Must be unique to the storage set

Server printer naming rules

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

- Can contain from 1 to 60 characters (bytes)
- Can be mixed case; however, Content Manager OnDemand ignores the case (for example, LaGuarde is the same as laguarde)
- Must be unique to the library server

Server print queue naming rules

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

- Can contain from 1 to 60 characters (bytes)
- Must be a valid printer queue name on the library server

Data types and field types

When you define an application group, Content Manager OnDemand creates a structure for a database table with the index and filter fields that you define.

When you store a report in the application group, Content Manager OnDemand extracts index data from the report, places the index data into the database fields, and inserts rows into the application group table.

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

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

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

Table 2. Index field data types. The table lists the types of database fields supported by Content Manager OnDemand and the number of bytes required to hold a value in each type of index field.

Field Type	Description	Field Size (DB2)	Field Size (Oracle)	Field Size (SQL Server)
Small Int	Contains whole numbers between -32768 and 32,767.	2 bytes	21 bytes	2 bytes

Table 2. Index field data types. The table lists the types of database fields supported by Content Manager OnDemand and the number of bytes required to hold a value in each type of index field. (continued)

Field Type	Description	Field Size (DB2)	Field Size (Oracle)	Field Size (SQL Server)
Integer	Contains whole numbers between -2147483648 and 2147483647.	4 bytes	21 bytes	4 bytes
Big Int	Contains whole numbers between -9223372036854775808 and 9223372036854775807.	8 bytes	21 bytes	8 bytes
Decimal	Contains numbers between 10(307) and 10(308) with up to 15 significant digits.	8 bytes	21 bytes	8 bytes
DecFloat (16)	Decimal floating-point number with 16 digits of precision.	8 bytes	Not available	Not available
DecFloat (34)	Decimal floating-point number with 34 digits of precision.	16 bytes	Not available	Not available
String (Fixed)	Contains letters, numbers, special symbols, such as the % and #, and any other printable character.	1 - 254 bytes. <ul style="list-style-type: none"> • For SBCS string: 1 byte per character • For MBCS string: Up to 4 bytes per character 	1 - 254 bytes. <ul style="list-style-type: none"> • For SBCS string: 1 byte per character • For MBCS string: Up to 4 bytes per character 	1 - 254 bytes. <ul style="list-style-type: none"> • For SBCS string: 1 byte per character • For MBCS string: Up to 4 bytes per character

Table 2. Index field data types. The table lists the types of database fields supported by Content Manager OnDemand and the number of bytes required to hold a value in each type of index field. (continued)

Field Type	Description	Field Size (DB2)	Field Size (Oracle)	Field Size (SQL Server)
String (Variable)	Contains letters, numbers, special symbols, such as the % and #, and any other printable character.	1 - 2000 bytes. <ul style="list-style-type: none"> • For SBCS string: 1 byte per character plus 2 bytes of overhead • For MBCS string: Up to 4 bytes per character plus 2 bytes of overhead Unused bytes do not use storage. See your database documentation for more information on variable string type.	1 - 2000 bytes. <ul style="list-style-type: none"> • For SBCS string: 1 byte per character plus 2 bytes of overhead • For MBCS string: Up to 4 bytes per character plus 2 bytes of overhead Unused bytes do not use storage. See your database documentation for more information on variable string type.	1 - 2000 bytes. <ul style="list-style-type: none"> • For SBCS string: 1 byte per character plus 2 bytes of overhead • For MBCS string: Up to 4 bytes per character plus 2 bytes of overhead Unused bytes do not use storage. See your database documentation for more information on variable string type.
Date	Contains a valid date from January 1, 0001 to December 31, 9999.	4 bytes	7 bytes	3 bytes
Date (old style)	Contains a valid date from January 1, 1970 to December 31, 2058.	2 bytes	2 bytes	2 bytes
Time (old style)	Contains times of day, stored in three second increments, since midnight, and limited to 24 hours.	2 bytes	2 bytes	2 bytes
Date/Time	Contains a valid date from January 1, 0001 to December 31, 9999.	13 bytes	11 bytes	8 bytes
Date/Time (old style)	Contains both a date and time value. The date can be from January 1, 1970 to January 18, 2038. The time is stored in one second increments.	4 bytes	4 bytes	4 bytes

Table 2. Index field data types. The table lists the types of database fields supported by Content Manager OnDemand and the number of bytes required to hold a value in each type of index field. (continued)

Field Type	Description	Field Size (DB2)	Field Size (Oracle)	Field Size (SQL Server)
Date/Time (TZ)	Contains a valid date from January 1, 0001 to December 31, 9999.	13 bytes	11 bytes	8 bytes
Date/Time (TZ) (old style)	Contains both a date and time value. A Date/Time (TZ) field is exactly like a Date/Time field, but uses the time zone from the client.	4 bytes	4 bytes	4 bytes

Table 3 on page 46 lists additional types of fields that are supported in folders.

<i>Table 3. Additional folder field types</i>	
Field Type	Description
Annotation Color Search	Searches annotations to a document by specifying a color. A match occurs and an item is added to the document list if the color of the text in one or more of the annotations to a document is the same as the color that is specified in the search field. A folder can have one annotation text search field.
Annotation Text Search	Searches annotations to a document for the specified string. A match occurs and an item is added to the document list if one or more of the annotations to a document contain the text that is specified in the search field. A folder can have one annotation text search field.
Application Group	For a search field, contains a list of the application groups that can be searched from the folder. When you create a folder that contains more than one application group, you can define an application group field. If enabled for queries, users can select the name of the application group that Content Manager OnDemand searches, rather than searching all of the application groups contained in the folder (the default). For a display field, lists the name of the application group in which the document was found. A folder can have one application group field.
Segment	Contains a list of the tables of index data that are stored in the application groups that can be searched from the folder. Each item in the list represents a segment of application group data. OnDemand segments application group data by date. If enabled for queries, users can select a specific segment of application group data to search. A folder can have one segment field. Specifying a search value in the segment field can improve the performance of queries.

Table 3. Additional folder field types (continued)

Field Type	Description
Text Search	<p>Finds documents that contain a non-indexed word or phrase. A match occurs and an item is added to the document list when one or more lines in a document contain the word or phrase exactly as specified. The search string can contain letters, numbers, special symbols (such as % and #) and any other printable character. A folder can have one text search field.</p> <p>The (sequential) text search takes place on the server. A text search will delay the generation of the document list. Only documents that meet all of the criteria specified in the other folder fields will be searched for the specified word or phrase.</p> <p>A typical use of a text search field is to provide users an additional search field without incurring database overhead. For example, assume that a report is indexed on date and transaction number. A text search field allows users to optionally enter a customer's name, phone number, or any other information contained in the document or documents that the user needs to retrieve (the information is not contained in the database). However, a text search field has a direct impact on the generation of the document list and the performance of the server. A large number of users performing text searches at the same time can drain the resources of even the most powerful library server.</p>

Chapter 7. Storage requirements overview

To estimate storage requirements for a Content Manager OnDemand system, you must understand and document user requirements for storing and accessing data.

[Chapter 5, “Reports and other data,” on page 27](#) provides information that can help you gather user requirements.

Before you turn requirements into a storage subsystem to support your system, you must also review the various operational and performance issues. For example, Content Manager OnDemand supports up to 128 index fields for each report. However, users typically do not need many indexes to locate a specific version of a report or a document within a report. The number of index fields that you define has a direct impact on the amount of disk space that you need for your database.

In addition, the more indexes that you define for a report, the more time is required to load the report into the system. You must work with users and understand their data retrieval requirements. Define only the number of index fields that they need. You might have to balance user requirements with disk space, the amount of time required to load a report, and other performance issues.

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

You need several components to measure to determine the amount of disk, optical, and tape storage required to support a Content Manager OnDemand system. For example, the following components of the system require disk storage:

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

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

- Storage space for the database, which includes Content Manager OnDemand system tables (control information and objects that you define to Content Manager OnDemand) and application group tables (index data extracted from reports). The amount of database space that you should plan for a report is a

factor of the number of items contained in the report, the number of index fields that you define for the report, the number of versions of a report (or the frequency with which you load a report on the system), and how long you need to maintain a report on the system.

For reports that contain sorted transaction data, Content Manager OnDemand can divide the report into groups of a fixed number of pages and create one index row for each group of pages. (For sorted transaction data, the examples and calculations that follow assume that Content Manager OnDemand will create one indexed item for each group of 100 pages in a report. The number of pages in a group is a parameter that you can configure when you index a report with ACIF. The *Content Manager OnDemand Indexing Reference* provides more information. For reports that contain logical items, such as statements, and policies, Content Manager OnDemand can create one index row for each logical item in the report. Typically the database space required for indexing sorted transaction data is much less than the database space required for indexing reports that contain logical items. Also, index fields provide fast lookup, but require a significant amount of database space.

- Storage space for database log files. You should plan for disk space for active or primary log files and for log files that are not active but might still be needed for recovery (sometimes known as archived log files).
- Storage space for the database and logs used by the archive storage manager.
- Temporary storage space for server print.
- Temporary storage space for importing migrated indexes from archive media to the database.

The following components of the system require archive storage (optionally VSAM files or OAM objects):

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

When you calculate archive storage requirements, you should also determine the number of storage volumes and libraries that you need to support the data that will be stored on your system. Optical libraries are capable of holding a large amount of data, with the storage capacity usually expressed in amounts of uncompressed data.

Depending on the compression ratio achieved for your reports, an optical library might be able to hold more than the stated amount. For example, if Content Manager OnDemand can achieve a 6:1 compression ratio on the reports that will be stored in an IBM 3995 optical library, the library can hold multiple terabytes of report data, depending on hardware configuration.

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

Storage hierarchy

You can use several different storage management strategies with Content Manager OnDemand and most archive storage managers.

Content Manager OnDemand maintains a cache storage system independently of the archive storage manager. The cache storage system should contain the fastest storage devices for high-speed access to reports. When you load a report on the system, Content Manager OnDemand can automatically store one copy of the report in cache storage and another copy of the report in archive storage. Content Manager OnDemand also supports the option of storing reports in cache storage and then later migrating them to archive storage.

However, it is recommended that you always plan to copy reports to cache storage and archive media at the same time (when you load the report). Doing so usually eliminates the need for you to periodically

backup cache storage because a backup copy of your reports already exists on archive media. Copying reports to cache storage and archive storage at the same time also eliminates the need for you to migrate reports to archive media.

Reports *expire* (are eligible to be removed) from cache storage when they reach their cache storage expiration date. You specify the cache storage expiration date for a report when you create an application group. For example, you can specify that a report should expire from cache storage after it is stored there for 90 days. Content Manager OnDemand provides a utility that you can use to automatically remove expired reports from cache storage on a regular schedule. After you run expiration processing, Content Manager OnDemand reclaims the space taken by expired documents. If any reports have *holds* on them when you run expiration processing, Content Manager OnDemand does not remove any of those reports. The Enhanced Retention Management feature, which provides the holds functionality, requires you to disable the archive storage manager; therefore, the archive storage manager will not be able to delete these reports.

Content Manager OnDemand and the archive storage manager maintain documents independently of each other. For example, each uses its own criteria to determine when data expires and should be removed from the system. Each uses its own utilities to remove documents. However, for removal of documents from the system, you should specify the same criteria to Content Manager OnDemand and the archive storage manager. For example, the Life of Data and Indexes, which is used by Content Manager OnDemand, should specify the same length of time as the retention period for the VSAM or OAM archive file definitions.

Data compression

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

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

To properly estimate the amount of storage space required by a report, it is recommended that you measure the compression ratio achieved on a sample of the report. You can measure the compression ratio by using the ARSADMIN program. For example:

- For reports that contain logical items, such as statements and policies, use the following example:
`arsadmin compress -l 200000 -s inputFile -o outputFile`

Where *inputFile* is the report that you want to measure and *outputFile* is the compressed output.

To determine the compression ratio, divide the size of *outputFile* value by the length (`-l 200000`). For example, if the size of *outputFile* value is 66 000 bytes, then the compression ratio is 66000/200000 or 0.33 (3:1 compression).

- For reports that contain line data and include a sorted transaction value, such as a general ledger, first determine the size of an indexed group of pages, for example, 100 pages. Then, extract a group of pages from a larger report and process them with the ARSADMIN program, for example: `arsadmin compress -s groupPages -o outputFile`

Where *groupPages* is a file that contains a representative group of pages from a larger report and *outputFile* is the compressed output.

To determine the compression ratio, divide the size of *outputFile* value by *groupPages* value. For example, if the size of *outputFile* value is 40 000 bytes and the size of the group of pages is 200 000 bytes, then the compression ratio is 40000/200000 or 0.20 (5:1 compression).

To run the ARSADMIN program, the input file must exist in HFS. The output file will be written to the directory that contains the input file. See the *IBM Content Manager OnDemand for z/OS: Administration Guide* for more information about the ARSADMIN program.

Beginning with APAR PI41677, Content Manager OnDemand supports the zEnterprise Data Compression (zEDC) available with z/OS 2.1 for OD77 and OD77Lite compression methods. Content Manager OnDemand will automatically attempt to use the zEDC unless configured not to by using the ARS_USE_ZLIB_HW=0 parameter in the ARS.CFG file. In order for that attempt to succeed, the following conditions must occur:

1. The IFAPRDxx PARMLIB member must specify that the z/OS zEDC software feature is ENABLED;
2. The RACF users for which use of the zEDC hardware is allowed must have READ access to FPZ.ACCELERATOR.COMPRESSION in the FACILITY CLASS; and
3. Buffer sizes must meet the minimum sizes specified by the IQPPRM PARMLIB member. If the zEDC hardware cannot be used, software-based compression and decompression will be performed instead.

Additionally, the _HZC_COMPRESSION_METHOD=software environment variable can be used to force the zEDC to perform software-based compression and decompression.

The OD77, OD77Lite, and zEDC compression types implement the same compression and decompression algorithm. That is, something compressed with one can be decompressed by the other. However, all three have been tuned differently, so that compressed sizes are different. The OD77 compression tends to compress better than the zEDC, which in turn tends to compresses better than OD77Lite. Remember that different compression methods result in different CPU consumption, with greater compression consuming greater CPU. Experimentation should be performed with specific examples of data to determine what is the optimal compression to use.

Calculating data storage requirements

Temporary space for indexing

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

Use the following calculation to determine the amount of temporary space required to index reports: *Temporary space = Largest report file size * 1.5*. The value for *Largest report file size* is the size in bytes of the largest version of a report to be indexed or the total size of all of the reports that the server must index at the same time (if you index more than one report at a time).

For example, if the largest report is 400 MB and the report is indexed on the Content Manager OnDemand server, the temporary space required to index the report is 600 MB: 400 MB * 1.5 = 600 MB.

Cache storage

The amount of storage space that you dedicate to cache storage varies greatly based on requirements such as the number of reports that you store on the system, the compression ratio that Content Manager OnDemand can achieve, and the amount of time that you need to store a report in cache storage.

Most customers store reports in cache storage for a short time to provide the fastest retrieval for the most frequently used reports. As reports age, and retrieval requests for them are much less frequent. The reports can be retrieved from archive media.

You might also keep records in cache storage when many users access them at the same time. Because the archive storage manager might require from six to sixty seconds to mount an optical or tape storage volume and retrieve a report, it is usually not possible to support a high transaction rate for reports stored on archive media.

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

Use the following calculation to determine the amount of disk space required for cache storage: $Cache\ Storage = Size\ of\ Data\ per\ week * Number\ of\ Weeks\ to\ cache * Data\ Compression\ ratio * 1.1$.

For example, if you plan to load 2 GB of report data on the system each week, the reports must be maintained in cache storage for 26 weeks, and the compression ratio is 8:1 (0.13), then the storage space required for cache storage space can be calculated as follows: $Cache\ Storage = 2\ GB * 26 * .13 * 1.1 = 7\ 436\ 000\ 000$

Content Manager OnDemand database storage

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

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

The following values determine the amount of storage space that is required for the Content Manager OnDemand database:

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

The types of index fields supported by Content Manager OnDemand are:

- Small Integer
- Integer
- Decimal
- String (Fixed)
- String (Variable)
- Date
- Date (old style)
- Time
- Date/Time
- Date/Time (old style)
- Date/Time (TZ)
- Date/Time (TZ) (old style)

See the DB2 Universal Database for information on the field sizes of the supported field types.

Calculating the size of the database

About this task

You can use the following calculations to determine the space required in the Content Manager OnDemand database to hold the index data for a report. In general, it is recommended that you add 10 to

20 percent to the calculated space requirements. The calculations can be used for reports that contain logical items and reports that contain a sorted transaction value.

Note: The formulas that follow were derived from information provided with the database manager products that work with Content Manager OnDemand. See the product information for details.

- $TableSize = (\text{Sum of column lengths})$
- $IndexSize = (Index\ 1\ length + 8) + (Index\ 2\ length + 8) + \dots$
- $DatabaseSize = ((TableSize + 40) * 1.5) + (IndexSize * 1.5) * \text{Number of indexed items per month} * \text{Number of months to keep index in database}$

Where:

- *Index n length* is the size of a database field for which you want Content Manager OnDemand to build an index. For example, a date field requires 4 bytes to hold the date value. DB2 requires an additional 8 bytes for each index that you define.
- Content Manager OnDemand adds approximately 40 bytes of control information to each row in a table.
- When the report contains logical items, the *Number of indexed items per month* is the number of statements, policies, and so forth.
- When the report contains a sorted transaction value, the *Number of indexed items per month* is the number of groups of indexed pages (by default, the system indexes a report in groups of 100 pages). You can specify the size of an indexed group of pages when you index a report with ACIF.

Example: Calculating the Db2 storage space for logical items

The following example illustrates how to calculate the Db2 database storage space required for a report that contains logical items, such as statements. In this example, the system indexes 1 million items per month and keeps the index data in the database for 36 months. [Table 4 on page 54](#) lists information about the database fields.

<i>Table 4. Database storage for a report that contains logical items</i>			
Field Name	Field Type	Field Size	Index or Filter
Account Number	Fixed String	12 bytes	Index
Report Date	Date	4 bytes	Filter
Customer Name	Variable String	20+4 bytes	Filter

The following list shows the results of the calculations for the example values:

- $TableSize = (12 + 4 + (20 + 4)) = 40$
- $IndexSize = (12 + 8) = 20$
- $DatabaseSize = ((40 + 40) * 1.5) + (20 * 1.5) = 150$

In this example, Content Manager OnDemand requires approximately 5.4 GB of storage space to store 36 months of report index data in the database.

Example: Calculating the Db2 storage space for line data

The following example illustrates how to calculate the database storage space required for a report that contains line data with a sorted transaction value. Because only one database row is generated for each indexed group of pages in the report, in general, significantly less database storage space is required than for reports that contain logical items.

Reports that contain line data with a sorted transaction value use a fixed type of indexing, where each database row contains the beginning value, the ending value, and the beginning page number for the group of pages. Content Manager OnDemand maintains the beginning and ending values as indexes and the page number as a filter. The main parameters for the calculation are:

- Length, in bytes, of the sorted transaction value

- Number of pages generated in a month
- Size of a group of indexed pages
- Number of months that Content Manager OnDemand maintains the index data in the database

In this example, the system indexes 1 million pages per month, in groups of 100 pages, and keep the index in the database for 24 months. [Table 5 on page 55](#) lists information about the database fields.

<i>Table 5. Database storage for a report that contains a sorted transaction value</i>			
Field Name	Field Type	Field Size	Index or Filter
Begin Transaction Value	Fixed String	10 bytes	Index
End Transaction Value	Fixed String	10 bytes	Index
Report Date	Date	4 bytes	Filter
Page Number	Integer	4 bytes	Filter

The following list shows the results of the calculations for the example values:

- $TableSize = (10 + 10 + 4 + 4) = 28$
- $IndexSize = ((10 + 8) + (10 + 8)) = 36$
- $DatabaseSize = ((28 + 40) * 1.5) + (36 * 1.5) = 156 * (1,000,000/100) = 1560000 * 24 = 37440000$

In this example, Content Manager OnDemand requires 37.44 MB of storage space to store 24 months of report index data in the database.

Database log file storage

The DB2 database requires storage space for recovery logs. Contact your database administrator to ensure that sufficient storage space for primary and secondary log files is defined and that database and log file backup procedures are established for your installation.

Server print storage space

Content Manager OnDemand requires temporary work space to process requests for the server print manager. You must allocate enough storage space to support the maximum number of concurrent print requests that the server must manage. It is recommended that you define a dedicated file system on which Content Manager OnDemand can store the temporary print files. At least 500 MB of free space must be available in this file system at all times. If your storage configuration permits, you should allocate 1 GB or more of free space to this file system.

Storage sizing examples

The following examples illustrate how to estimate storage requirements for two types of reports:

- A report that contains logical items, such as statements or policies
- A report that contains sorted transaction data

A report that contains logical items

<i>Table 6. Report that contains logical items. Database columns. Part 1 of 4.</i>			
Column number	Name	Index or filter	Bytes
1	Account number	Index	12
2	Report date	Filter	4
3	Customer name	Filter	24

Table 7. Report that contains logical items. Report profile. Part 2 of 4.

Report characteristic	Report estimate
Volume of data per month (bytes)	8 000 000 000
Average statement size (bytes)	8000
Number of statements per month	1 000 000
Number of cycles per month	20
Largest cycle data size (bytes)	400 000 000
Largest single report file size (bytes)	400 000 000
Largest cycle (number of statements)	50 000
Number of database columns from Table 6 on page 55	4
Life of data (months)	36
Number of weeks to cache data	26
Number of days to keep the index in the database	1098
Compression percentage (ratio)	0.13 (8:1) AFP

Table 8. Report that contains logical items. Disk storage requirements in bytes. Part 3 of 4.

Storage component	Storage requirement
Temporary storage	600 000 000
Cache storage	7 436 000 000
Content Manager OnDemand database	5 400 000 000
Total disk storage required (bytes)	13 436 000 000

Table 9. Report that contains logical items. Archive storage requirements in bytes. Part 4 of 4.

Storage component	Storage space requirement
Report data	41 184 000 000
Total archive storage required (bytes)	41 184 000 000

A report that contains transaction data

Table 10. Report that contains transaction data. Database columns. Part 1 of 4.

Column number	Name	Index or filter	Bytes
1	Beginning invoice number	Index	10
2	Ending invoice number	Index	10
3	Report date	Filter	4
4	Page number	Filter	4

Table 11. Report that contains transaction data. Report profile. Part 2 of 4.

Report characteristic	Report estimate
Volume of data per month (bytes)	5 000 000 000
Average page size (bytes)	5000
Number of pages per month	1 000 000
Number of cycles per month	20
Largest cycle (data size in bytes)	200 000 000
Largest cycle (number of pages)	50 000
Largest single report file size (bytes)	200 000 000
Group of indexed pages	100
Number of database columns from Table 10 on page 56	4
Life of data (months)	24
Number of weeks to cache data	0
Number of days to keep the indexes in the database	730
Compression percentage (ratio)	0.05 (20:1) Line data

Table 12. Report that contains transaction data. Disk storage requirements in bytes. Part 3 of 4.

Storage Component	Storage space requirement
Temporary storage	300 000 000
Cache storage	0
Content Manager OnDemand database	37 440 000
Total disk storage required (bytes)	337 440 000

Table 13. Report that contains transaction data. Archive storage requirements in bytes. Part 4 of 4.

Storage component	Storage space requirement
Report data	6 600 000 000
Total archive storage required (bytes)	6 600 000 000

Calculating archive storage requirements

Report storage space

When you estimate the amount of space required to store a report in archive storage, you must consider the size of the report, the compression ratio achieved, and the length of time that the archive storage manager maintains the report. Archive media can be optical storage or magnetic tape. Use the following calculation to estimate that amount of space required: $ArchiveStorageSpace = (Data\ per\ month * life\ of\ data\ in\ months) * compression\ ratio * 1.1$

For example, if you plan to store 8 GB of report data per month, the archive storage manager must maintain the data for 36 months, and Content Manager OnDemand can achieve a compression ratio of 8:1

(0.13), you need approximately 41 GB of archive storage space: $OpticalSpace = (8 \text{ GB} * 36) * 0.13 * 1.1 = 41184000000$

Migrated index storage space

Content Manager OnDemand supports automatic migration of indexes from the database to archive storage so that you can maintain seldom used indexes for long periods of time. However, you should migrate indexes only after you do not need to retrieve reports that refer to the indexes. For example, suppose that all of the queries for a report occur in the first 24 months after the report is loaded into the system, after that time, there are no more queries for the report. The indexes might be eligible to be migrated from the database to archive storage. Migration of index data is optional; you can choose to migrate indexes for all, some, or none of the application groups on your system. In addition, you determine the length of time that indexes stay in the database before Content Manager OnDemand migrates them to archive storage.

You can use the following calculation to determine the archive storage space that is required to hold migrated indexes: $ArchiveMediaDBSpace = (Database \text{ size per month} * \text{compression ratio}) * (\text{life of data} - \text{months before migrating data})$

For example, if the index data requires 1.5 MB of space in the database per month, you need to maintain the indexes for 84 months, and the indexes remain in the database for 24 months before being migrated, then the archive storage required to hold the migrated indexes is: $ArchiveMediaDBSpace = (1500000 * .33) * (60) = 29,700,000$

Device storage

Content Manager OnDemand has the ability to store reports by using different storage methods (HFS, VSAM, OAM) and different device types (DASD, Tape, Optical). Depending on the device type and storage method used, the exact amount of space required on the device for storing data will vary. For more accurate estimates of the amount of storage space required in your particular situation, see the appropriate access method and device specification manuals. Alternatively, the previously calculated numbers (which represent the amount of data that needs to be stored) can be used as a general estimate of the amount of storage space required.

Storage sizing worksheets

The following worksheets can help you estimate the storage requirements for two types of reports:

- A report that contains logical items, such as statements or policies
- A report that contains sorted transaction data

Make a copy of the worksheets on the following pages for each report that you want to store in Content Manager OnDemand. Complete the worksheets to calculate the storage requirements for the report. See [“Calculating data storage requirements” on page 52](#) and [“Calculating archive storage requirements” on page 57](#) for the formulas that you can use to calculate the storage requirements.

<i>Table 14. Report that contains logical items. Disk storage requirements in bytes. Part 3 of 4.</i>	
Storage component	Storage space requirement
Temporary storage	
Cache storage	
Content Manager OnDemand database	
Total disk storage required (bytes)	

Table 15. Report that contains logical items. Archive storage requirements in bytes. Part 4 of 4.

Storage component	Storage space requirement
Report data	
Total archive storage required (bytes)	

A report that contains logical items

Table 16. Report that contains logical items. Database columns. Part 1 of 4.

Column number	Name	Index or filter	Bytes
1			
2			
3			
4			

Table 17. Report that contains logical items. Report profile. Part 2 of 4.

Report characteristic	Report estimate
Volume of data per month (bytes)	
Average item size (bytes)	
Number of items per month	
Number of cycles per month	
Largest cycle data size (bytes)	
Largest single report file size (bytes)	
Largest cycle (number of items)	
Number of database columns from Table 16 on page 59	
Life of data (months)	
Number of weeks to cache data	
Number of days to keep the indexes in the database (default is life of data)	
Compression ratio: image, PDF 1; AFP 8:1 (0.13); Line data 20:1 (0:05)	

A report that contains transaction data

Table 18. Report that contains transaction data. Database columns. Part 1 of 4.

Column number	Name	Index or filter	Bytes
1			
2			
3			
4			

Table 19. Report that contains transaction data. Report profile. Part 2 of 4.

Report characteristic	Report estimate
Volume of data per month (bytes)	
Average page size (bytes)	
Number of pages per month	
Number of cycles per month	
Size of largest cycle (bytes)	
Size of largest cycle in pages	
Size of largest single report file (bytes)	
Pages in an indexed group (default is 100)	
Number of database columns from Table 18 on page 59	
Life of data (months)	
Number of weeks to cache data	
Number of days to keep the indexes in the database (default is life of data)	
Compression ratio: image, PDF 1; AFP 8:1 (0.13); Line data 20:1 (0:05)	

Table 20. Report that contains transaction data. Disk storage requirements in bytes. Part 3 of 4.

Storage component	Storage spacer Requirement
Temporary storage	
Cache storage	
Content Manager OnDemand database	
Total disk storage required (bytes)	

Table 21. Report that contains transaction data. Archive storage requirements in bytes. Part 4 of 4.

Storage component	Storage space requirement
Report data	
Total archive storage required (bytes)	

Chapter 8. Backup and recovery overview

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

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

Content Manager OnDemand supports storing index data in table spaces and the incremental backup of table spaces. Table spaces enhance the management of index data and provide improved performance, especially for database backups. An incremental table space backup finishes more quickly than a full database backup, providing you with increased flexibility in scheduling the loading of reports. Incremental backup images also require less storage space than full database backups.

Server software

You can usually use the original product media to restore the software programs.

Ensure that you store the original product media in a safe location. It is recommended that you register Content Manager OnDemand as part of your business recovery plan and store the original product media in the same place that you store the other programs and files that are vital to the operation of your systems.

Server information

When you installed and configured Content Manager OnDemand, you specified information that customized Content Manager OnDemand to operate in your environment. Some Content Manager OnDemand programs were stored in HFS while others were stored in MVS Libraries. Both are managed by SMS.

Also, if you periodically change the system, including the database, archive storage manager, and server print manager, you should back up the control files regularly, perhaps once a week. The Content Manager OnDemand control and data files are contained in the HFS.

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

Database backup

Contact your database administrator to ensure that database backup files are defined and backup procedures are established for your installation. See the *IBM Content Manager OnDemand for z/OS: Configuration Guide* for information about how to configure the system and maintain backup files.

Reports

Content Manager OnDemand can store copies of reports in cache storage and archive storage:

- Cache storage provides short-term, high-speed storage and retrieval of reports. Cache storage consists of disk storage volumes maintained by Content Manager OnDemand on one or more object servers.

- Archive storage provides long-term storage and retrieval of reports. Reports in archive storage can also be used as backup copies, in the event that cache storage becomes corrupted or unavailable. Archive storage consists of optical or tape storage volumes managed by the archive storage manager, optionally VSAM or OAM.

You can configure the system to copy reports to cache storage and archive storage at the same time, when you load a report into the system.

Content Manager OnDemand can retrieve a copy of a report from archive storage after the report is removed from cache storage or if the copy on cache storage is unavailable. However, you must configure the system to support multiple copies of reports. You configure Content Manager OnDemand to use archive storage by defining VSAM files or OAM objects in the administration client.

Tip: If you do not plan to copy reports to archive storage, you should back up the file systems that comprise cache storage regularly. However, if a media failure occurs or cache storage becomes corrupted, users cannot retrieve reports until the file systems are restored.

Cache storage

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

If you do not copy reports to archive storage when you store them in Content Manager OnDemand, you need to consider how you can recover the reports in the event that you need to do so (for example, if a device fails).

Cache storage can be protected by maintaining it on high-availability storage devices. If no high-availability storage is available, you should back up reports in cache storage (the HFS data sets) regularly. Consult your storage administrator for procedures to schedule and maintain backups of your installation.

Archive storage

Content Manager OnDemand retrieves the primary copy of the report from archive storage after the report is removed from cache storage. Customers with special business, legal, or performance reasons might want the system to maintain a backup copy of their reports in archive storage. The backup copy can be used if the primary copy becomes corrupted or unavailable.

You must configure the archive storage manager to maintain a backup copy of reports in archive storage.

See your storage administrator for details about defining and managing multiple copies of reports, backup and recovery of data, and scheduling operations.

Notices

This information was developed for products and services that are offered in the US. This material might be available from IBM in other languages. However, you may be required to own a copy of the product or product version in that language in order to access it.

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, MD-NC119
Armonk, NY 10504-1785
US*

For license inquiries regarding double-byte character set (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.
19-21, Nihonbashi-Hakozakicho, Chuo-ku
Tokyo 103-8510, Japan*

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 jurisdictions 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 websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.

IBM may use or distribute any of the information you provide 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 Director of Licensing
IBM Corporation
North Castle Drive, MD-NC119
Armonk, NY 10504-1785
US*

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.

The client examples cited are presented for illustrative purposes only. Actual performance results may vary depending on specific configurations and operating conditions.

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.

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 actual people or business enterprises 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.

Each copy or any portion of these sample programs or any derivative work must include a copyright notice as follows:

© (your company name) (year).

Portions of this code are derived from IBM Corp. Sample Programs.

© Copyright IBM Corp. _enter the year or years_.

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the web at "Copyright and trademark information" at www.ibm.com/legal/copytrade.shtml.

Adobe, the Adobe logo, PostScript, and the PostScript logo are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, and/or other countries.

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

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

Java™ and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

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

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

Terms and conditions for product documentation

Permissions for the use of these publications are granted subject to the following terms and conditions.

Applicability

These terms and conditions are in addition to any terms of use for the IBM website.

Personal use

You may reproduce these publications for your personal, noncommercial use provided that all proprietary notices are preserved. You may not distribute, display or make derivative work of these publications, or any portion thereof, without the express consent of IBM.

Commercial use

You may reproduce, distribute and display these publications solely within your enterprise provided that all proprietary notices are preserved. You may not make derivative works of these publications, or reproduce, distribute or display these publications or any portion thereof outside your enterprise, without the express consent of IBM.

Rights

Except as expressly granted in this permission, no other permissions, licenses or rights are granted, either express or implied, to the publications or any information, data, software or other intellectual property contained therein.

IBM reserves the right to withdraw the permissions granted herein whenever, in its discretion, the use of the publications is detrimental to its interest or, as determined by IBM, the above instructions are not being properly followed.

You may not download, export or re-export this information except in full compliance with all applicable laws and regulations, including all United States export laws and regulations.

IBM MAKES NO GUARANTEE ABOUT THE CONTENT OF THESE PUBLICATIONS. THE PUBLICATIONS ARE PROVIDED "AS-IS" AND WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING BUT NOT LIMITED TO IMPLIED WARRANTIES OF MERCHANTABILITY, NON-INFRINGEMENT, AND FITNESS FOR A PARTICULAR PURPOSE.

IBM Online Privacy Statement

IBM Software products, including software as a service solutions, (“Software Offerings”) may use cookies or other technologies to collect product usage information, to help improve the end user experience, to tailor interactions with the end user or for other purposes. In many cases no personally identifiable information is collected by the Software Offerings. Some of our Software Offerings can help enable you to collect personally identifiable information. If this Software Offering uses cookies to collect personally identifiable information, specific information about this offering’s use of cookies is set forth below.

This Software Offering does not use cookies or other technologies to collect personally identifiable information.

If the configurations deployed for this Software Offering provide you as customer the ability to collect personally identifiable information from end users via cookies and other technologies, you should seek your own legal advice about any laws applicable to such data collection, including any requirements for notice and consent.

For more information about the use of various technologies, including cookies, for these purposes, see IBM’s Privacy Policy at www.ibm.com/privacy and IBM’s Online Privacy Statement at www.ibm.com/privacy/details the section entitled “Cookies, Web Beacons and Other Technologies” and the “IBM

Software Products and Software-as-a-Service Privacy Statement” at www.ibm.com/software/info/product-privacy.

Index

Numerics

- 390 indexer
 - about [31](#)
- 390 indexer exits [21](#)

A

- ACIF
 - indexing [30](#)
- administrative
 - roles and responsibilities [16](#)
- AFP API
 - indexing [31](#)
- AFP Application Programming Interface
 - indexing [31](#)
- AFP application programs
 - indexing with [31](#), [32](#)
- AFP conversion and indexing facility (ACIF) [10](#)
- AFP data
 - indexing [30](#), [31](#)
 - input [27](#)
 - transforming [12](#)
- annotation color search field [43](#)
- annotation text search field [43](#)
- Anystore exit [21](#)
- APIs [17](#)
- application
 - naming [41](#)
- application group
 - database field types [43](#)
 - field types [43](#)
 - indexing [36](#)
 - naming [41](#)
 - permissions [40](#)
- application group field [43](#)
- application programming interfaces [17](#)
- application programs
 - indexing with [30](#)
- applications [4](#), [38](#)
- archive storage
 - migrated indexes [58](#)
 - reports on archive media [57](#)
 - requirements [57](#)
- archive storage manager [10](#)
- ARSDOC program [17](#)
- ARSEXP program [17](#)
- ARSTBLSP program [17](#)
- ARSYSPIN [20](#)

B

- backup and recovery
 - database [61](#)
 - documents [61](#)
 - reports [61](#)
- Bitmap (BMP)

- Bitmap (BMP) (*continued*)
 - indexing [32](#)
 - input [27](#)
- BMP (Bitmap)
 - indexing [32](#)
 - input [27](#)
- browsers
 - accessing OnDemand data with [12](#)

C

- cabinet [4](#), [6](#)
- cache expiration processing [36](#)
- cache migration processing [36](#)
- cache storage
 - requirements [52](#)
- cache storage manager [10](#)
- client
 - APIs [17](#)
 - customization [17](#)
- client preview exit [19](#)
- cloud
 - overview [26](#)
 - storage, managing [26](#)
- collecting requirements [27](#)
- commands [17](#)
- compression [51](#)
- Content Manager OnDemand administrative client [4](#)
- Content Manager OnDemand Distribution Facility [3](#)
- Content Manager OnDemand system architecture [1](#)
- conventions
 - names [41](#)

D

- data
 - compression [51](#)
 - formats [27](#)
 - indexing [28](#)
- data storage
 - requirements [52](#)
- database
 - backup and recovery [61](#)
 - expiration processing [36](#)
 - field names [41](#)
 - field types [43](#)
 - recovery [61](#)
 - storage requirements [49](#)
 - storage space requirements [53](#)
- database manager
 - expiration [11](#)
 - migration [11](#)
- database organization
 - multiple loads per database table [36](#)
 - single load per database table [36](#)
- date fields
 - database field type [43](#)

- date fields (*continued*)
 - folder field type [43](#)
 - indexing [33](#)
- Db2 [10](#)
- document
 - expiration [11](#)
 - indexing [10](#), [29](#)
 - migration [11](#)
- documents
 - backup and recovery [61](#)

E

- Enhanced Retention Management (ERM) [3](#)
- environment variables
 - time zone (TZ) [43](#)
 - TZ (time zone) [43](#)
- exits
 - 390 [21](#)
 - Anystore [21](#)
 - client preview [19](#)
 - index [21](#)
 - indexing [21](#)
 - input [21](#)
 - report specifications archive definition [19](#)
 - system log [18](#)
 - table space creation [20](#)
 - unified login [19](#)
 - user security [18](#)
- expiration [11](#)
- expiration processing [17](#), [36](#)
- Extensible Markup Language (XML) output
 - output [12](#)
- external
 - overview [26](#)
 - storage, managing [26](#)
- external storage [25](#)

F

- features [1](#)
- field names [41](#)
- field types [43](#)
- file formats supported by Content Manager OnDemand
 - AFP data [27](#)
 - Bitmap (BMP) [27](#)
 - BMP (Bitmap) [27](#)
 - GIF (Graphic Interchange Format) [27](#)
 - Graphic Interchange Format (GIF) [27](#)
 - JFIF (JPEG Format Image File) [27](#)
 - JPEG Format Image File (JFIF) [27](#)
 - line data [27](#)
 - PCX (Picture Exchange Format) [27](#)
 - PDF (Portable Data Format) [27](#)
 - Picture Exchange Format (PCX) [27](#)
 - Portable Data Format (PDF) [27](#)
 - Tagged Image File Format (TIFF) [27](#)
 - TIFF (Tagged Image File Format) [27](#)
 - user-defined data [27](#)
- file formats supported by OnDemand
 - Extensible Markup Language (XML) output [12](#)
 - HTML output [12](#)
 - XML output [12](#)

- folder
 - annotation color search field [43](#)
 - annotation text search field [43](#)
 - application group field [43](#)
 - field names [41](#)
 - field types [43](#)
 - naming [41](#)
 - permissions [39](#)
 - query field types [43](#)
 - search field types [43](#)
 - segment field [43](#)
 - text search field [43](#)

G

- generic indexer [32](#)
- GIF (Graphic Interchange Format)
 - indexing [32](#)
 - input [27](#)
- Graphic Interchange Format (GIF)
 - indexing [32](#)
 - input [27](#)
- groups
 - names [41](#)

H

- HTML
 - output [12](#)

I

- index exit [21](#)
- Indexers [2](#)
- indexes
 - archive storage space requirements [58](#)
- indexing
 - 390 indexer [31](#)
 - ACIF [30](#)
 - AFP API [31](#)
 - AFP Application Program Interface [31](#)
 - AFP application programs [31](#), [32](#)
 - AFP data [30](#), [31](#)
 - AFP records in the data stream [31](#)
 - application groups [36](#)
 - application programs [30](#)
 - Bitmap (BMP) [32](#)
 - BMP (Bitmap) [32](#)
 - data [28](#)
 - date fields [33](#)
 - documents [29](#)
 - generic indexer [32](#)
 - GIF (Graphic Interchange Format) [32](#)
 - Graphic Interchange Format (GIF) [32](#)
 - JFIF (JPEG Format Image File) [32](#)
 - JPEG Format Image File (JFIF) [32](#)
 - line data [30](#), [31](#)
 - methods [7](#)
 - PCX (Picture Exchange Format) [32](#)
 - Picture Exchange Format (PCX) [32](#)
 - programs [7](#)
 - reports [29](#)
 - space requirements [52](#)

indexing (*continued*)
 Tagged Image File Format (TIFF) [32](#)
 temporary space [52](#)
 TIFF (Tagged Image File Format) [32](#)
input data formats [27](#)
input exit [21](#)

J

Java API
 software requirement [12](#)
JES Spool data capture facility [20](#)
JFIF (JPEG Format Image File)
 indexing [32](#)
 input [27](#)
JPEG Format Image File (JFIF)
 indexing [32](#)
 input [27](#)

L

library server [8](#)
line data
 indexing [30, 31](#)
 input [27](#)
loading data [10](#)
log files [18](#)
logon
 unified login exit [19](#)
 user security exit [18](#)

M

management programs [11](#)
migrated indexes
 storage space requirements [58](#)
migration [11](#)
Monarch
 integrating with the client [17](#)
multiple loads per database table [36](#)

O

object server [8](#)
objects [35](#)
ODF [3](#)

P

password length [41](#)
passwords [41](#)
PCX (Picture Exchange Format)
 indexing [32](#)
 input [27](#)
PDF (Portable Data Format)
 input [27](#)
PDF transform [12](#)
permissions
 application group [40](#)
 folder [39](#)
Picture Exchange Format (PCX)
 indexing [32](#)
 input [27](#)

Portable Data Format (PDF)
 input [27](#)
printing
 storage space requirements [55](#)

Q

query field names [41](#)
query field types [43](#)

R

recovery
 database [61](#)
report
 expiration [11](#)
 formats [27](#)
 indexing [10, 28, 29](#)
 migration [11](#)
 storage on archive media [57](#)
 storage requirements [49](#)
report specifications archive definition exit [19](#)
reports
 backup and recovery [61](#)
request manager
 archive storage manager [9](#)
 cache storage manager [9](#)
requirements
 archive storage [57](#)
 cache storage [52](#)
 data storage [52](#)
 database storage [53](#)
 indexing reports [52](#)
 migrated index storage space [58](#)
 printing [55](#)
 reports on archive media [57](#)
 server print [55](#)
 storage [49](#)
 temporary space [52](#)
responsibilities
 administrative [16](#)
roles
 administrative [16](#)

S

search field names [41](#)
search field types [43](#)
segment field [43](#)
server commands [17](#)
server logging [18](#)
server print
 storage space requirements [55](#)
servers [8](#)
setting the time zone environment variable [43](#)
single load per database table [36](#)
Spool data capture facility (JES) [20](#)
storage
 archive storage requirements [57](#)
 cache space [52](#)
 data storage requirements [52](#)
 database [49, 53](#)
 estimating [49](#)

storage (*continued*)
 indexing reports [52](#)
 migrated indexes [58](#)
 planning for [49](#)
 printing [55](#)
 report [49](#)
 reports on archive media [57](#)
 requirements [49](#)
 server print [55](#)
 temporary space [52](#)
storage manager
 expiration [11](#)
 migration [11](#)
storage objects [25](#)
system log file [18](#)
system log user exit [18](#)
system requirements [23](#)

X

XML
 output [12](#)

T

table space creation exit [20](#)
Tagged Image File Format (TIFF)
 enter [27](#)
 indexing [32](#)
tasks
 administrative [16](#)
temporary space
 storage requirements [52](#)
text search field [43](#)
TIFF (Tagged Image File Format)
 enter [27](#)
 indexing [32](#)
time zone environment variable [43](#)
transaction number [9](#)
TZ environment variable [43](#)

U

unified login exit [19](#)
user exits
 390 indexer [21](#)
 Anystore [21](#)
 client preview [19](#)
 index [21](#)
 indexing [21](#)
 input [21](#)
 report specifications archive definition [19](#)
 system log [18](#)
 table space creation [20](#)
 unified login [19](#)
 user security [18](#)
user security exit [18](#)
user-defined data [27](#)
users
 names [41](#)
 passwords [41](#)

W

web browsers
 accessing OnDemand data with [12](#)



Product Number: 5697-CM1

SC19-3365-03

