Content Manager OnDemand for Multiplatforms
Version 10 Release 1

Installation and Configuration Guide

GC19-3342-02

Note on notices

Before using this information and the product it supports, read the information in “Notices” on page 183.

This edition applies to Version 10 Release 1 of Content Manager OnDemand for Multiplatforms (product number 5724-J33) and to all subsequent releases and modifications until otherwise indicated in new editions.

This product includes software developed by the Apache Software Foundation (`http://www.apache.org/`).

Contents

ibm.com® and related resources

Product support and documentation are available from ibm.com®.

Support and assistance

From 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/SSEPCD_10.1.0/com.ibm.ondemandtoc.doc/ondemandmp_10.1.0.htm

PDF publications

See the following PDF publications for your product at http://www.ibm.com/support/docview.wss?uid=swg27050851.

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. Before you begin

Review the pre-installation tasks before installing the product.

About this task

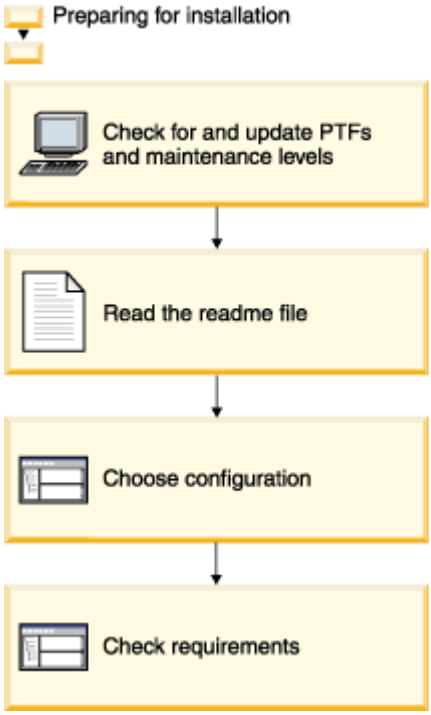

Figure 1: Pre-installation tasks

This section contains important information that is common to all Content Manager OnDemand installations. You will find specific information about your platform in the section devoted to it.

You should review the following topics:

- For hardware and software requirements, see http://www.ibm.com/support/docview.wss?uid=swg27049168 for details.
- For disk storage requirements see the *IBM Content Manager OnDemand for Multiplatforms: Introduction and Planning Guide*.

What you should know first

Before using the Content Manager OnDemand for Multiplatforms: Installation and Configuration Guide and installing and configuring your Content Manager OnDemand system, you should be familiar with some key concepts.

Key concepts to know:

- Administering the server operating system you plan to install Content Manager OnDemand on.
- The networking protocols that will be required for clients and servers to communicate.
- The devices and file systems that will be available to Content Manager OnDemand. Before you begin the installation, you must identify the devices and file systems that will be used for program files, the database, data downloaded from other systems, data indexing and loading, cache storage, temporary storage, and so forth. You must prepare the storage volumes and configure them to support the different components of the system.

- The database management product that you will be using with Content Manager OnDemand. Most customers should have an experienced database administrator available to help with the installation, configuration, and operation of the system.
- (Optional) The devices that will be available to Tivoli® Storage Manager. If you plan to use Tivoli Storage Manager to maintain Content Manager OnDemand data, you must install and configure the devices that will be used by Tivoli Storage Manager. This publication describes additional configuration required by the Content Manager OnDemand system.
- The operational requirements for the system. For example, you might need to configure maintenance tasks to run automatically on a regular schedule and you need to know what type of database backups should be taken and when.

Installing and configuring Content Manager OnDemand for Multiplatforms

Describes the information you need to know before you configure your Content Manager OnDemand for Multiplatforms system, including access to systems, resources, and people that have knowledge of those systems and resources.

Choosing a configuration

Before you install Content Manager OnDemand, you should learn about the Content Manager OnDemand system configuration and the various components that make up the system.

About this task

The *IBM Content Manager OnDemand for Multiplatforms: Introduction and Planning Guide* provides details about the system configuration and the required and optional components for each type of server. You should also print and read the README file from the Content Manager OnDemand product physical (CD or DVD) or unpacked electronic media (for example, `*.tar`). After you have that information, you can choose the components that you need to install and prepare to install the components by setting up your environment. Setting up your environment consists of two tasks:

Procedure

1. Planning for sufficient performance and capacity.
2. Installing the platform-specific prerequisites.

Results

Content Manager OnDemand supports several types of system configurations. However, most customers install the standard library/object server system. A standard library/object server system is a single workstation that includes a complete Content Manager OnDemand system and performs both the library server and object server functions.

In a production environment, you should dedicate each workstation to Content Manager OnDemand work and processes. That is, you should not run other applications on a workstation with a Content Manager OnDemand library or object server.

Library server

The Content Manager OnDemand library server uses a relational database manager to manage objects and provide data integrity by maintaining index information and controlling access to objects stored on one or more object servers.

Library servers run on AIX, Linux, and Windows servers, and can use IBM DB2 Universal Database, Microsoft SQL Server (on Windows servers), or Oracle to manage the library contents. A Content Manager OnDemand system has one library server.

The library server directs requests from clients to query, retrieve, and print items in the database, which contains object indexes and other information. The library server routes requests to the appropriate object server to store, retrieve, and delete objects.

Planning for capacity
The library server workstation is primarily a database machine. It builds search requests and transmits the results of searches to the client. In addition to reserving disk space for prerequisite software and the Content Manager OnDemand program files, you must allocate storage for the database as it grows.

Physically separating program directories, the database, and log file directories will improve performance and the time it takes to recover from problems.

Library server machines have high input/output workloads, and they need a powerful processor to accommodate concurrent requests from multiple users. Because the database lies at the core of the library server, good database administration is crucial to the efficient operation of the Content Manager OnDemand system.

LDAP servers
The first three LDAP servers in this list support anonymous bind. The Microsoft ADAM and AD servers do not.

The Content Manager OnDemand library server supports the following LDAP servers:

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

Object server

Content Manager OnDemand stores and retrieves objects that reside on an object server through requests routed by the library server.

An object server is the repository for objects stored on the system. The object server manages storage resources that are defined through the Content Manager OnDemand administrative programs. The object server supports attachment of disk and storage devices managed by Tivoli Storage Manager (also known as IBM Spectrum Protect), as well as a number of cloud storage managers. A Content Manager OnDemand system can have many object servers distributed across networks to provide convenient user access. Object servers run on AIX, Linux, and Windows servers.

Object servers work with the Content Manager OnDemand administrative programs to efficiently manage storage resources. This allows the Content Manager OnDemand administrator to specify how long documents reside on one media type before migrating them to another and how long Content Manager OnDemand maintains documents on the system.

Planning for capacity
To plan the capacity requirements for storing documents in your environment, you must consider a number of factors.

The *IBM Content Manager OnDemand for Multiplatforms: Introduction and Planning Guide* provides details, examples, and worksheets to help you estimate the amount of storage needed to support your environment.

- Prerequisite software and Content Manager OnDemand program files
- Staging areas for data download and indexing data
- Temporary spaces for loading and printing
- Cache (short-term) storage for documents
- Archive (long-term) storage for documents

Physically separating program directories, staging areas, temporary spaces, cache storage, and archive storage will improve performance and the time it takes to recover from problems.

Tivoli Storage Manager considerations for object servers
Content Manager OnDemand uses TSM to maintain object storage on media other than fixed disks.

To store the primary copy of a document, the object server writes to TSM using a TSM archive copy group. Because there is no need for multiple generations of the object server's objects, IBM recommends that you set the maximum number of backup copies to one. You can choose to maintain more than one copy of data in a storage pool. However, you must define the backup copy to Tivoli Storage Manager and define a schedule to automate the backup copy. Physically separating the primary copy and backup copy can improve the time it takes to recover from problems.

Chapter 2. Installing Content Manager OnDemand on AIX servers

This part of the IBM Content Manager OnDemand for Multiplatforms: Installation and Configuration Guide explains how to install and configure Content Manager OnDemand on an AIX server and how to install and configure related software to work with Content Manager OnDemand.

About this task

There are five basic phases to the installation, which are illustrated in Chapter 2, “Installing Content Manager OnDemand on AIX servers,” on page 5:

- Preparing for the installation
- Installing and configuring Content Manager OnDemand and related software
- Verifying the installation
- Preparing the system for use
- Adding optional software

You will find checklists for each of these phases in “Checklist for installation on AIX” on page 6.

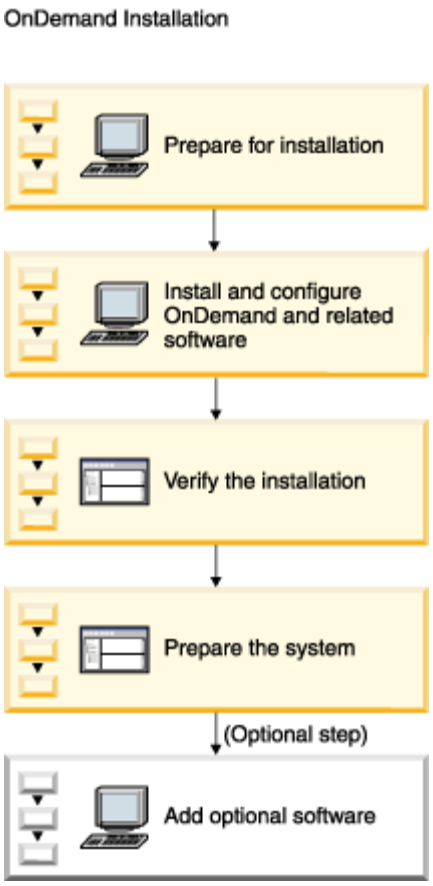

Figure 2: Installing Content Manager OnDemand on an AIX server

Checklist for installation on AIX

Review the pre-installation checklist before installing the product on AIX.

About this task

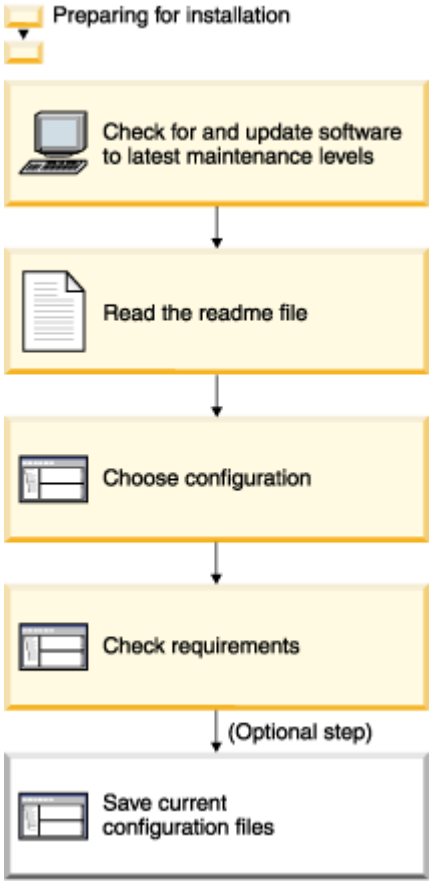

Figure 3: Pre-installation tasks

Procedure

Before beginning the installation, you should complete the following tasks:

1. Contact the IBM Support Center for the latest maintenance levels of DB2®, Content Manager OnDemand, and optionally, Tivoli Storage Manager and Infoprint Manager (Infoprint). If you are using Oracle instead of DB2, contact Oracle for information about the latest maintenance level of Oracle.
2. Obtain a copy of the latest Content Manager OnDemand README file. Print and read the entire file before you begin.
3. Check the Content Manager OnDemand prerequisites and verify the required and optional hardware and software products (see " AIX server requirements" on page 10).
4. Check the hardware and software requirements for all system components and features. See http://www.ibm.com/support/docview.wss?uid=swg27049168 for details.
5. Determine the type of system configuration that you need to install (see "Choosing a configuration" on page 2).
6. If you are upgrading to a new release of Content Manager OnDemand, save the configuration files used by the system (see "Saving configuration files on AIX" on page 11).

Results

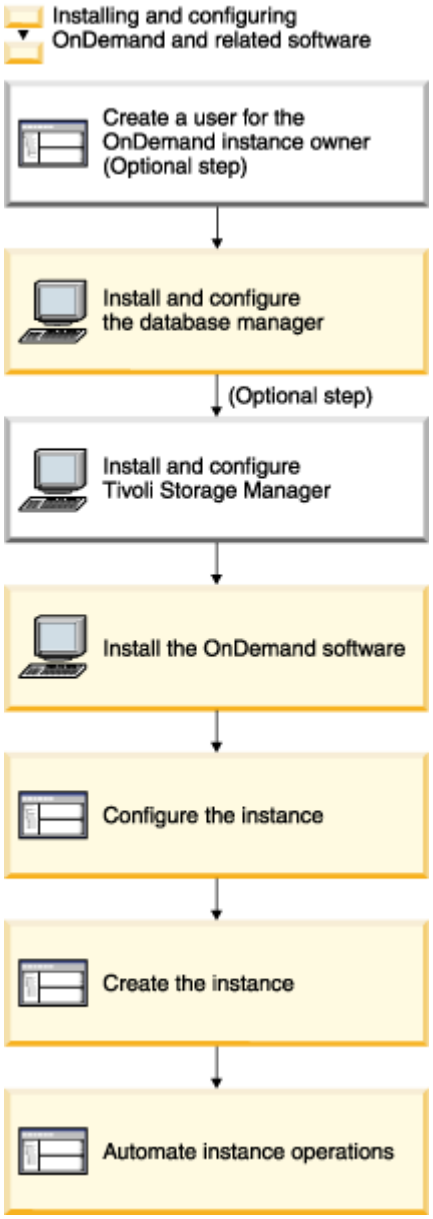

Figure 4: Installing Content Manager OnDemand and related software

Configuring a Content Manager OnDemand system typically requires that you do the following tasks:

1. (Optional) Create a user account for the Content Manager OnDemand instance owner (see "Creating a user for the Content Manager OnDemand instance owner on AIX" on page 12).
2. Install and configure the database manager product on the library server (see "Installing the database manager on AIX" on page 14).
3. Install the IBM Global Security Kit (GSKit) (see "Installing IBM® Global Security Kit on AIX®" on page 16).
4. If you plan on using SSL for security, set up SSL on the Content Manager OnDemand server and client (see "Setting up SSL on Content Manager OnDemand for AIX" on page 18).
5. To save Content Manager OnDemand passwords into encrypted files, create a stash file and save it in a directory with restricted access (see "Saving Content Manager OnDemand passwords into encrypted files" on page 21).

6. If you plan to maintain data in archive storage, install and configure Tivoli Storage Manager on the library server or on each object server that will be used to maintain data in archive storage (see “Installing and configuring Tivoli Storage Manager on AIX” on page 22).
7. Install the Content Manager OnDemand software on each workstation that is part of the Content Manager OnDemand system (see “Installing the Content Manager OnDemand software on AIX” on page 26).
8. Configure an instance of Content Manager OnDemand on each workstation that is part of the Content Manager OnDemand system (see “Configuring instances on AIX” on page 28). This step includes:
   a. Specify the instance in the `ARS.INI` file (see “ARS_DB_ENGINE parameter” on page 31)
   b. Specify the `ARS.CFG` file for the instance (see “Specifying the ARS.CFG file for the instance” on page 30)
   c. Specify the `ARS.DBFS` file for the instance (see “Specifying the ARS.DBFS file for the instance” on page 36)
   d. Specify the `ARS.CACHE` file for the instance (see “Creating the ARS.CACHE file for the instance” on page 38)
9. Create the instance of Content Manager OnDemand (see “To create an instance of Content Manager OnDemand” on page 40). This step includes the following tasks:
   a. Specify permissions for the database directories (see “Specifying permissions for the database directories” on page 40)
   b. Create the instance by running the ARSDB program (see “Creating a database instance” on page 41)
   c. Initialize the system logging facility by running the ARSSYSCR program (see “Initializing the system logging facility” on page 42)
   d. (Optional) Initialize the system migration facility by running the ARSSYSCR program (see “Initializing the system migration facility” on page 43)
10. Automate instance operations (see “Automating instance operations on AIX®” on page 44). This step includes the following tasks:
    a. Start the database on the library server (see “Starting the database” on page 44)
    b. Start the instance on the library server (see “Starting the instance on the library server” on page 44)
    c. Start the instance on an object server (see “Starting the instance on an object server” on page 45)
    d. Start the data loading programs (see “Starting the data loading programs” on page 45)
    e. Schedule application group maintenance on the library server (see “Scheduling application group maintenance on the library server” on page 47)
    f. Schedule application group maintenance on an object server (see “Scheduling application group maintenance on an object server” on page 47)
    g. Schedule system table maintenance (see “Scheduling system table maintenance” on page 47)
    h. Schedule a backup of the Content Manager OnDemand database (see “Scheduling the Content Manager OnDemand database backup” on page 48)

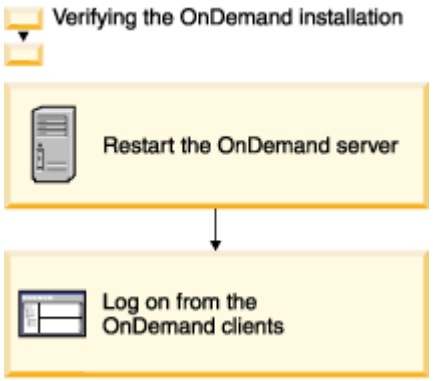

Figure 5: Verifying the installation

Verify the installation of Content Manager OnDemand (see “Verifying the installation” on page 137):

1. After installing and configuring each Content Manager OnDemand server, restart the system. The operating system reinitializes and starts the services required by Content Manager OnDemand.
2. Log on to the library server with a Content Manager OnDemand client program. To access the system, you must install at least one of the Content Manager OnDemand client programs on a PC running Microsoft Windows. See *IBM Content Manager OnDemand: Client Installation Guide* for installation information about the Content Manager OnDemand client. See *IBM Content Manager OnDemand for Multiplatforms: Administration Guide* for installation information about the administrative client.

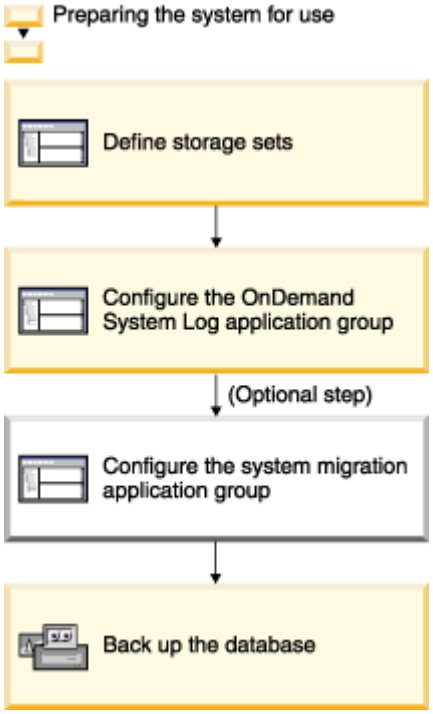

Figure 6: Preparing the system for use

Prepare the system for use:

1. Define storage sets (see “Define storage sets” on page 138). Before you add application groups or load data into the system, you must define storage sets.
2. Configure the System Log application group (see “Configuring the System Log application group” on page 138). Before you define reports to the system, load data, or let users access the system, you should configure the System Log application group.
3. Configure the System Load application group as described in “Configuring the System Load application group” on page 140.

4. If you plan to migrate index data to archive storage, configure the System Migration application group (see "Configure the System Migration application group" on page 142).
5. Back up the databases (see Chapter 7, "Backing up the Content Manager OnDemand database," on page 145). After configuring the system, you should create a full backup image of the Content Manager OnDemand database and the Tivoli Storage Manager database.

Installing and configuring optional software:

1. If you plan to use Download for the z/OS® feature (Download) to transmit data from z/OS systems to Content Manager OnDemand servers, you must install and configure Download. Follow the instructions in *PSF for z/OS: Download for z/OS* to plan, install, configure, and verify the installation of the Download software. Then configure Download on each Content Manager OnDemand server. Complete the following tasks:
   a. Obtain a copy of *PSF for z/OS: Download for z/OS*.
   b. Check the prerequisites and verify the z/OS and TCP/IP software levels for Download.
   c. Install and configure the Download software.
   d. Configure Download on each Content Manager OnDemand server that will receive data sets from an z/OS system. (see "Starting the data loading programs" on page 45).
2. If you plan to reprint documents using the Content Manager OnDemand server print function, you must install Infoprint on a workstation that belongs to the same network as the Content Manager OnDemand library server. Follow the instructions in the Infoprint documentation for your operating system to plan, install, configure, and verify the installation of the Infoprint software. Then configure the server print function on the library server. Complete the following tasks:
   a. Obtain a copy of the Infoprint documentation for the server operating system.
   b. Install and configure Infoprint.
   c. Verify that all of the resources and fonts that your organization requires to reprint the reports that you plan to store in Content Manager OnDemand are installed on the Infoprint server.
   d. Define the print queues and devices that Infoprint uses to manage the Content Manager OnDemand server print environment.
   e. Obtain the TCP/IP host name or IP address of the Infoprint server.
   f. On the library server, edit the ARSPRT file and insert the host name or IP address of the Infoprint server. The ARSPRT file can be found in the `/opt/IBM/ondemand/V10.1/bin` directory.
   g. Define a server printer on the Content Manager OnDemand library server with the administrative client.
3. If you need to customize and enhance the standard functionality within the product, see the user exit documentation in the Appendix of this publication. A user exit is a point during processing that enables you to run a user-written program and return control of processing after your user-written program ends. Content Manager OnDemand provides the following user exit points:
   a. Download user exit
   b. Report specifications archive definition user exit
   c. Retrieval preview user exit
   d. Security user exit
   e. System log user exit
   f. Table space creation user exit

AIX server requirements

The exact hardware and software configuration that you need for Content Manager OnDemand to support your organization depends on the volume of data that you plan to maintain on the system, the number of

concurrent users that the system must support, the backup and recovery requirements of your organization, and the performance levels that the system must meet.

About this task

At a minimum, you need one processor for a standard Content Manager OnDemand library/object server.

For all AIX® server requirements, see http://www.ibm.com/support/docview.wss?uid=swg27049168.

Saving configuration files on AIX

When you install software on a Content Manager OnDemand server, the installation programs copy program files, configuration files, and other types of files from the distribution media to directories on the server.

About this task

When you configure a server to meet the specific requirements of your environment, you make changes to configuration files and you might customize other files, such as user-defined files and font initialization files.

Before you upgrade to a new version of Content Manager OnDemand or upgrade the database manager software or other software related to Content Manager OnDemand, you should save a copy of the files listed in this section. You can save a copy of the files in a temporary directory, such as `/tmp`.

After you upgrade the software, you will probably need to reconfigure the files for your environment. To reconfigure the files, you can restore the copies of the files that you saved or make changes to the updated files, using the configuration information in the files that you saved as a guide.

Content Manager OnDemand files

You should save copies of all Content Manager OnDemand files needed to configure the product.

Save a copy of the Content Manager OnDemand configuration files.

Table 1: Content Manager OnDemand configuration files to save

File	Location	Purpose
`ars.cache`	`/opt/IBM/ondemand/V10.1/config`	Define cache storage file systems. Changes described in "Creating the ARS.CACHE file for the instance" on page 38.
`ars.cfg`	`/opt/IBM/ondemand/V10.1/config`	Content Manager OnDemand server configuration file. Changes described in "Creating the ARS.CACHE file for the instance" on page 38.
`ars.dbfs`	`/opt/IBM/ondemand/V10.1/config`	Define DB2 table space file systems. Changes described in "Specifying the ARS.DBFS file for the instance" on page 36.
`ars.ini`	`/opt/IBM/ondemand/V10.1/config`	Configure Content Manager OnDemand instances. Changes described in "Specifying instances in the ARS.INI file" on page 29.
`arslog`	`/opt/IBM/ondemand/V10.1/bin`	The System Log user exit program. Described in "System log user exit" on page 162.

Table 1: Content Manager OnDemand configuration files to save (continued)

File	Location	Purpose
arsprt	/opt/IBM/ondemand/V10.1/bin	Server print program.

Tivoli Storage Manager files

If you use Tivoli Storage Manager to maintain Content Manager OnDemand data in archive storage, save a copy of the Tivoli Storage Manager configuration files.

Table 2: Tivoli Storage Manager configuration files to save

File	Location	Purpose
dsmserv.dsk	/opt/tivoli/tsm/server/bin64	Locations of the Tivoli Storage Manager database and recovery logs
history.dev	/opt/tivoli/tsm/server/bin64	Tivoli Storage Manager device history file
history.vol	/opt/tivoli/tsm/server/bin64	Tivoli Storage Manager storage volume history file
dsmserv.opt	/opt/tivoli/tsm/server/bin64	Tivoli Storage Manager server options file
dsm.sys	/opt/tivoli/tsm/server/bin64	Tivoli Storage Manager servers file
dsm.opt	/usr/tivoli/tsm/client/ba/bin64	Tivoli Storage Manager client options file
dsm.db2.opt	/usr/tivoli/tsm/client/api/bin64	Tivoli Storage Manager client options file for maintaining DB2 archived log files and back up files.
dsm.sys	/usr/tivoli/tsm/client/api/bin64	Tivoli Storage Manager client system options file

Creating a user for the Content Manager OnDemand instance owner on AIX

This publication was written assuming that OnDemand instances will be run under the root user. The information in this section is provided for customers who need to run instances of Content Manager OnDemand under a user other than the root user.

Those customers should print the information in this section and have it available to assist them as they continue with the installation and configuration process.

New installations (instances) of Content Manager OnDemand can be configured to run under a user other than the root user. If you plan to run an instance under a user other than root, you must do the following tasks:

- Create the user for the Content Manager OnDemand instance owner
- Set permissions for the cache storage file systems
- Set permissions for the Content Manager OnDemand configuration and script files
- Give the instance owner permission to write to the system console
- Specify the instance owner in the ARS.INI file

If you plan to run a distributed library/object server system, with one or more object servers on different workstations or nodes than the library server, then you should also configure Content Manager OnDemand on the object servers.

Configure the library server

Create a user that is a member of the database owner's group. This group has administrator authority for the database and the database file systems.

Give the Content Manager OnDemand instance owner the following authorities and permissions:

- Administrator authority for the database. You can do this by adding the Content Manager OnDemand instance owner to the database owner's group.
- Ownership of the cache storage file systems that are listed in the `ARS.CACHE` file. You can do this by running the Change Owner command for each file system that is listed in the `ARS.CACHE` file and specifying the user and group for the Content Manager OnDemand instance owner.
- Permission to read the Content Manager OnDemand configuration files. Make sure that the Content Manager OnDemand instance owner has permission to read the following files:
  - `ARS.CACHE`
  - `ARS.CFG`
  - `ARS.DBFS`
  - `ARS.INI`
- Permission to read and execute the Content Manager OnDemand script files. Make sure that the Content Manager OnDemand instance owner has permission to read and execute the following files:
  - `ARSLOG`
  - `ARSPRT`
- Permission to write to the console. Make sure that the Content Manager OnDemand instance owner has permission to write to the system console.

You should specify a different user for each instance that you create. This allows for easier error recovery if a system error occurs.

Important: You cannot set the permissions to read and execute Content Manager OnDemand files until you complete installation of the Content Manager OnDemand software. See "Installing the Content Manager OnDemand software on AIX" on page 26 for instructions on installing the Content Manager OnDemand software on AIX.

Configuring an object server

If you plan to run a distributed library/object server system, with one or more object servers on different workstations or nodes than the library server, then you should also configure Content Manager OnDemand on each of the object servers.

Procedure

To configure Content Manager OnDemand on the object servers, do the following tasks:

1. Create a user for the Content Manager OnDemand instance owner.
2. Give ownership of the cache storage file systems listed in the `ARS.CACHE` file to the user for the Content Manager OnDemand instance owner.
3. Give permission to read the following files to the Content Manager OnDemand instance owner:
   - `ARS.CACHE`
   - `ARS.CFG`
   - `ARS.INI`
4. Give permission to write to the console to the Content Manager OnDemand instance owner.

Installing the database manager on AIX

This section provides installation and configuration information specific to Content Manager OnDemand for both DB2 and Oracle.

Installing DB2®

You must install either DB2 or Oracle on the Content Manager OnDemand library server. This section describes how to install DB2

About this task

See "Installing Oracle" on page 15 for instructions about installing Oracle.

The DB2 Universal Database Enterprise Edition program physical or electronic media are provided with the Content Manager OnDemand program package. The README file explains how to locate the information that you need. Follow the instructions in *IBM DB2 Universal Database Quick Beginnings for DB2 Servers* to plan, install, configure, and verify the installation of DB2.

Procedure

To install DB2 on the library server:

1. Install DB2 Universal Database™ Enterprise Edition.
2. When prompted, select Typical as the installation type, to install all DB2 components required to support Content Manager OnDemand. You can take most default options (unless you have specific requirements of your own).
3. Create the DB2 instance for Content Manager OnDemand when you install DB2. Use the following values:

Parameter	Value
Instance Name or User	archive
Group Name	gname **Note:** The group must have SYSADM authority, and its name must be unique. The group name on your database might be something other than 'gname'. Ask your database administrator if you do not know the group name for your database.
Home Directory	`/home/archive`
Auto start DB2 instance at boot time	no
Create a sample database for DB2 instance	no

4. After you install the software, apply the latest fix pack for DB2.

   You can obtain the latest fix packs at http://www.ibm.com/support/docview.wss?uid=swg27007053. Print the README file. Follow the instructions in the README file to apply the service update. After installing a fix pack, you might need to update your database instances (for example, archive). See the DB2 README for details.

Adding the user to the DB2 instance owner group

After you install DB2 on the library server, you can configure the user and instance group.

After installing DB2 on the library server:

1. Add the user that owns the Content Manager OnDemand instance to the DB2 instance owner's group.

For example, if the DB2 instance owner's group is db2iadm1 and the Content Manager OnDemand instance owner is root, specified by the SRVR_INSTANCE_OWNER parameter in the ARS.INI file, add the root user to the db2iadm1 group.

2. Create links for the DB2 files. See the instructions in *IBM DB2 Universal Database Quick Beginnings for DB2 Servers* to create links to the DB2 files.
3. Optionally create a table space for the Content Manager OnDemand system tables. If you plan to store the system tables in their own table space, specify the name of the table space on the ARS_DB_TABLESPACE parameter in the ARS.CFG file.
4. Verify the value of the DB2INSTANCE parameter in the ARS.CFG file. The value of the DB2INSTANCE parameter is case-sensitive. This value must specify the name of the DB2 instance that you created for Content Manager OnDemand. The default value is archive.

Setting the DB2® operating environment
If you plan to use DB2 commands to work with the Content Manager OnDemand database, you must execute a script file to set the DB2 operating environment before you start the DB2 command line interface.

About this task

For Bourne or Korn shell users, run the DB2PROFILE script file. For C shell users, run the DB2CSHRC script file.

The script files can be found in the INSTHOME/sqllib directory, where INSTHOME is the home directory of the instance owner. If you installed and configured the system using the suggested defaults, the instance owner is archive and the script files reside in the sqllib directory under /home/archive.

You should add the script file to your .profile or .login file. For example: . /home/archive/sqllib/db2profile

After executing the script file, you can start the DB2 command line interface and connect to the database. For example:

```
$>db2
  .
  .
  .
```

To stop the DB2 command line interface, enter: db2 =>quit

Installing Oracle

You must install Oracle on the Content Manager OnDemand library server.

About this task
After you verify the installation of the Oracle software on the library server, you must configure it to work with Content Manager OnDemand.

Procedure

To configure Oracle to work with Content Manager OnDemand:

1. Configure login processing to run under the UID of the root user.
2. Create the Content Manager OnDemand database using the Oracle utilities. The name that you specify for the database should match the value that you specify for the SRVR_INSTANCE parameter in the ARS.INI file.
3. Create the user ID of the Content Manager OnDemand instance owner in Oracle.

   This user will own all tables that Content Manager OnDemand creates. If you want to have a default Oracle table space for the user, specify the table space when you create the user.

To create the Content Manager OnDemand user in Oracle:

```
CREATE USER root IDENTIFIED BY password ;
GRANT dba to root ;
```

Where *root* and *password* are the user ID and password stored in the stash file

4. Specify the base Oracle installation directory on the ARS_ORACLE_HOME parameter in the `ARS.CFG` file. The default value is `/oracle`.
5. Specify Oracle as the database manager on the ARS_DB_ENGINE parameter in the `ARS.CFG` file.
6. Optional: Create a table space for the Content Manager OnDemand system tables.

   If you plan to store the system tables in their own table space, specify the name of the table space on the ARS_DB_TABLESPACE parameter in the `ARS.CFG` file.

Installing IBM® Global Security Kit on AIX®

Determine whether you already have GSKit installed on your system and, if so, which version of GSKit.

About this task

You can install IBM Global Security Kit (GSKit) with one of the following tools:

- SMIT GUI
- `installp` command

Procedure

Before you can install GSKit, you must do the following steps:

1. If another product required that you install GSKit in one of the following default locations, run the `gsk8ver` command or the `gsk8ver_64` command to determine which version of GSKit is installed:
   - `/usr/opt/IBM/gsk8`
   - `/usr/opt/IBM/gsk8_64`

   If another product required that you install GSKit in a location other than the default location, continue to the next step to install GSKit in the default location.
2. Place the GSKit media in a location that the tools can access.

   The SMIT GUI tool can access the installation media through a read-only device, for example, a DVD; however, for faster installation, use the electronic image or copy the GSKit media to the local file system in either the `/var/spool/pkg` or `/tmp/gsk8` directory.
3. Determine whether you need to install the 32-bit or the 64-bit version of GSKit by reviewing the following list:
   - The Content Manager OnDemand server requires the 64-bit version of GSKit.
   - The ODWEK Java API can use the 32-bit or 64-bit version of GSKit.
4. Extract the corresponding GSKit media, which is distributed as a compressed TAR file, into your chosen file path by running one of the following commands:
   - For the 32-bit version:

     ```
     zcat gskcrypt32-8.0.14.45.aix.ppc.tar.Z | tar -xvf -
      zcat gskssl32-8.0.14.45.aix.ppc.tar.Z | tar -xvf -
     ```

   - For the 64-bit version:

     ```
     zcat gskcrypt64-8.0.14.44.aix.ppc.tar.Z | tar -xvf -
      zcat gskssl64-8.0.14.44.aix.ppc.tar.Z | tar -xvf -
     ```

Installing GSKit on AIX® by using the SMIT GUI tool

Create the table of contents that is needed by the SMIT GUI tool by running the following `inutoc` command.

Procedure

To install GSKit on AIX, do the following steps:

1. Creates the table of contents that is needed by the SMIT GUI tool by running the following `inutoc` command in the directory that contains the GSKit packages:

   ```
   /usr/sbin/inutoc /tmp/gsk8
   ```

2. Start the SMIT GUI tool by entering `smit`.
3. Select **Software Installation & Maintenance**.
4. Select **Install Software**.
5. Select **Install and Update Software by Package Name**.
6. On the device/directory window, specify the directory that contains the installation media.
7. Select the following packages from the list:
   - `GSKit8.gskcrypt64.ppc.rte`
   - `GSKit8.gskssl64.ppc.rte`
   - `GSKit8.gskcrypt32.ppc.rte`
   - `GSKit8.gskssl32.ppc.rte`
8. Select the file sets of the software package to install.
9. From the options window, mark the options that are appropriate to your installation requirements and set the **Install all prereqs** options to **Yes**.
10. Confirm to complete the installation.

Installing GSKit on AIX® by using the installp command-line tool

You can install the GSKit by using the command line.

Procedure

Run the following commands:

```
/usr/bin/inutoc /tmp/gsk8
installp -acgqw -d /tmp/gsk8 GSKit8.gskcrypt64.ppc.rte \
                               GSKit8.gskssl64.ppc.rte \
                             GSKit8.gskcrypt32.ppc.rte \
                               GSKit8.gskssl32.ppc.rte
```

SSL for Content Manager OnDemand

Secure Sockets Layer (SSL) and its successor, Transport Layer Security (TLS), encrypts all transmissions between the Content Manager OnDemand servers and many of the supported clients (for example, ODWEK Java API, the Windows client, and arsdoc).

The CICS® client does not support SSL connections.

Before you begin setting up SSL on Content Manager OnDemand

Because of possible problems with system performance, create SSL connections only for communications that require secure transmission. Consider adding additional processor resources on the Content Manager OnDemand server, client, or both to manage the increased processor usage.

GSKit provides the GSKCapiCmd tool, which helps you create and manage digital certificates and key databases. The instructions in "Setting up SSL on Content Manager OnDemand for AIX" on page 18 provide examples of how to run the GSKCapCmd tool; however, to view the complete syntax and understand the behavior of this tool, see ftp://ftp.software.ibm.com/software/webserver/appserv/library/v80/GSK_CapiCmd_UserGuide.pdf.

Choose the scenario from the following list that fits your requirements, then follow the instructions for that scenario:

- Content Manager OnDemand server listens only on a non-SSL port. You cannot set up SSL for this situation. Continue to the next Content Manager OnDemand installation task.
- Content Manager OnDemand server listens only on an SSL port. You must do the following tasks:
  - Set up SSL on Content Manager OnDemand.
  - Install GSKit on all clients.
  - Configure the clients to support SSL.
- Content Manager OnDemand server listens on both a non-SSL port and an SSL port. You must do the following steps:
  - Set up SSL on Content Manager OnDemand.
  - Install GSKit on the clients that connect to the SSL port.
  - Configure those clients to support SSL.

Setting up SSL on Content Manager OnDemand for AIX

You can set up Secure Sockets Layer (SSL) on Content Manager OnDemand.

Procedure

To set up SSL on Content Manager OnDemand:

1. Create the key database and store it in the `config` subdirectory of Content Manager OnDemand server installation directory: `/opt/IBM/ondemand/V10.1`.

   To create the key database, run a command similar to the following command:

```
gsk8capicmd_64 -keydb -create -db "ondemand.kdb" -pw "myKeyDBpasswd" -stash -
populate
```

   The following list describes why these parameters were chosen:

   -keydb -create -db "ondemand.kdb"

   Indicates that you want to create a key database called ondemand.kdb.

   -pw "myKeyDBpasswd" -stash

   Indicates that you want to create a stash file and store the password (`myKeyDBpasswd`) in that stash file. The GSKCapiCmd tool stores the stash file at the same path as the key database. You must remember this path because you must specify it in the ars.ini file. GSKCapiCmd creates the stash file with the same file name as the key database (`ondemand`), with the file extension of `.sth`. When Content Manager OnDemand starts, GSKit retrieves the password to the key database from this stash file.

   -populate

   Populates the key database with a set of predefined trusted certificate authority (CA) certificates. A trusted CA is a certificate authority root certificate is noted as trusted in the

key database. For the list of default trusted root certificates, see "Default GSKit trusted root certificates" on page 175.

2. Create a digital certificate. You can create a self-signed certificate, which is useful for testing. When you are ready to move to a production environment, create a CA-signed digital certificate. "Creating a self-signed certificate" on page 19 and "Creating a CA-signed digital certificate" on page 20
3. Configure the Content Manager OnDemand initialization file. Add the following lines to the `ARS.INI` file:

```
SSL_PORT=port_number
SSL_KEYRING_FILE=/opt/IBM/ondemand/V10.1/config/ondemand.kdb
SSL_KEYRING_STASH=/opt/IBM/ondemand/V10.1/config/ondemand.sth
SSL_KEYRING_LABEL=IBM Content Manager OnDemand
SSL_CLNT_USE_SSL=0
```

The following list describes these parameters:

SSL_PORT

Specify one of the following values:

port_number

The port number on the Content Manager OnDemand server dedicated to communicating with the SSL protocol.

0

No port on the Content Manager OnDemand server to communicate with the SSL protocol.

-1

All ports on the Content Manager OnDemand server to communicate only with the SSL protocol.

SSL_KEYRING_FILE

Specify the full path and file name of the key database that contains the digital certificates.

SSL_KEYRING_STASH

Specify the full path and file name of the stash file for the key database.

SSL_KEYRING_LABEL

Specify the name of the certificate in the key database.

SSL_CLNT_USE_SSL

Specify whether the server-side clients (for example, ARSDOC, ARSMAINT, or ARSLOAD) must communicate with the SSL protocol. Specify 0 to indicate that the clients do not communicate with the SSL protocol. Specify 1 to indicate that the clients must communicate with the SSL protocol.

4. Restart the Content Manager OnDemand server.

Because a trusted certificate authority provided the digital certificate, the Content Manager OnDemand server accepts the certificate. Both `ondemand.kdb` and `ondemand.sth` files need to be placed on the workstation where the Content Manager OnDemand clients are installed. Download both files to the `config` subdirectory under the client installation directory.

Creating a self-signed certificate

You can create a self-signed certificate by using `GSKCapiCmd`.

Procedure

To create a self-signed certificate, do the following steps:

1. Create a self-signed certificate by using `GSKCapiCmd`.
   The following example creates a self-signed certificate with the label `myselfsigned`:

```
gsk8capicmd_64 -cert -create -db "ondemand.kdb" –pw "myKeyDBpasswd" -label
"myselfsigned" \
```

```
–dn
"CN=myhost.mycompany.com,O=myOrganization,OU=myOrganizationUnit,L=Boulder,ST=CO,C=US
"
```

2. Extract the certificate to a file by using GSKCapiCmd.
   The following example extracts the certificate into a file called ondemand.arm:

```
gsk8capicmd_64 -cert -extract -db "ondemand.kdb" -pw " myKeyDBpasswd " -label
"myselfsigned" \
-target "ondemand.arm" -format ascii
```

3. Distribute the file you created to all computers that run clients that will establish SSL connections to your Content Manager OnDemand server.

Creating a CA-signed digital certificate

You create CA-signed digital certificate for an RSA private-public key pair and PKCS10 certificate request, which are stored in the key database in a file with the .rdb extension.

About this task

Specify the name of the file, with the -file option, that you send to the CA.

Procedure

To create a CA-signed digital certificate:

1. Create a Certificate Signing Request (CSR) by using GSKCapiCmd.
   The following example shows how to create a CSR that is stored in ondemand.kdb.

```
gsk8capicmd_64 -certreq -create -db "ondemand.kdb" -pw "myKeyDBpasswd" -label
"mycert" \
-dn
"CN=myhost.mycompany.com,O=myOrganization,OU=myOrganizationUnit,L=Boulder,ST=CO,C=US
" \
-file "mycertRequestNew"
```

2. Verify the contents of the CSR by using GSKCapiCmd.
   The following example shows how to display the contents of the CSR:

```
gsk8capicmd_64 -certreq -details -db "ondemand.kdb" -pw "myKeyDBpasswd" –label
"mycert"
```

   If you need to delete this CSR, run GSKCapiCmd similar to the following example:

```
gsk8capicmd_64 -certreq -delete -db "ondemand.kdb" -pw "myKeyDBpasswd" -label
"mycert"
```

3. Go to the website of a well-known CA (for example, Verisign) and follow their instructions for registering and obtaining a signed digital certificate. The instructions include paying the CA for their services and providing them with the file you specified with the -file option. In the following example and for the rest of these instructions, a trial version of a digital certificate is used.
4. Use a text editor (for example, vi) to save each certificate into a file. The CA sends you an email with the following information:
   - The MyCertificate.arm file, your trial signed digital certificate.
   - A link to download IntermediateCert.arm, the trial intermediate digital certificate.
   - A link to download RootCert.arm, the root digital certificate.
5. Add the trial root digital certificate to the key database.

The following example adds RootCert.arm to ondemand.kdb:

```
gsk8capicmd_64 -cert -add -db "ondemand.kdb" -pw "myKeyDBpasswd" -label
"trialRootCACert" \
-file RootCert.arm -format ascii
```

6. Add the trial intermediate certificate to the key database.
   The following example adds IntermediateCert.arm to ondemand.kdb:

```
gsk8capicmd_64 -cert -add -db "ondemand.kdb" -pw "myKeyDBpasswd" –label
"trialIntermediateCACert" \
-file IntermediateCert.arm -format ascii
```

7. Receive your signed digital certificate to the key database.
   The following example receives MyCertificate.arm to ondemand.kdb:

```
gsk8capicmd_64 -cert -receive -file MyCertificate.arm -db "ondemand.kdb" –pw
“myKeyDBpasswd” \
–format ascii
```

8. Verify that all the certificates were stored in the key database by using GSKCapiCmd.
   The following example lists the certificates stored in ondemand.kdb:

   ```
   gsk8capicmd_64 -cert -list all -db "ondemand.kdb" -pw "myKeyDBpasswd"
   ```

   GSKCapCmd displays the following result:

   ```
   Certificates found
   * default, - personal, ! trusted
   -! mycert
   ! trialIntermediateCACert
   ! trialRootCACert
   -! myselfsigned
   ```

Saving Content Manager OnDemand passwords into encrypted files

You can store user IDs and passwords in encrypted files (also called stash files). Storing passwords in a stash file can help you improve security because you do not need to specify the password on the command line, where the password might be visible to others.

About this task

You can store the user ID and password for the following situations in one stash file:

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

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

Procedure

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

1. Create a stash file by running the arsstash command. The command prompts you for the password.
2. Save the stash file in a directory and limit access to that file through file permissions.

Results

When you configure the Content Manager OnDemand instance, you modify the ARS.INI file and include the SRVR_OD_STASH parameter and specify the directory that you specified.

Installing and configuring Tivoli Storage Manager on AIX

This section explains how to set up Tivoli Storage Manager for Content Manager OnDemand on an AIX workstation.

About this task

Tivoli Storage Manager can be used with Content Manager OnDemand object servers to store report data on devices that are supported by Tivoli Storage Manager. Devices supported by Tivoli Storage Manager include optical libraries and tape media. The use of Tivoli Storage Manager is optional and is needed only if you want to provide long-term storage for your reports on devices other than the fixed disks attached to the object server. You can also use Tivoli Storage Manager facilities to maintain DB2 archived log files and back up image files.

You will need the *IBM Tivoli Storage Manager for AIX: Quick Start* publication to install and configure Tivoli Storage Manager. HTML and PDF versions of Tivoli Storage Manager publications, including the *Quick Start*, are available at http://publib.boulder.ibm.com/tividd/td/tdprodlist.html.

Planning for interoperability between Content Manager OnDemand and Tivoli Storage Manager

Content Manager OnDemand uses the Tivoli Storage Manager API client to store data into the Tivoli Storage Manager server.

Content Manager OnDemand supports Tivoli Storage Manager in the following configurations:

- Content Manager OnDemand library/object or object server plus Tivoli Storage Manager on one workstation. Install the Server, 64-bit Client API, Device Support Runtime, Server Runtime, and Licenses packages on the object server workstation.
- Content Manager OnDemand library/object or object server only (where Tivoli Storage Manager resides on a workstation other than the library or object server). Install the 64-bit Client API packages on the object server workstation.

The Tivoli Storage Manager server is managed and administered independently of Content Manager OnDemand. The Tivoli Storage Manager administrator must ensure that the following conditions are met:

- All the normal requirements for Tivoli Storage Manager storage are monitored and managed accordingly
- All required Tivoli Storage Manager policies, management classes, storage pools, and volumes are defined accordingly
- All required Tivoli Storage Manager storage pools and volumes are online
- All Tivoli Storage Manager storage pools and volumes have sufficient storage space to satisfy the needs of Content Manager OnDemand
- The Tivoli Storage Manager server is active when Content Manager OnDemand needs to read from or write to its storage repository

If your Tivoli Storage Manager configuration cannot support Content Manager OnDemand, system requests (that require Tivoli Storage Manager services) will fail. The Tivoli Storage Manager administrator should examine the system to ensure that it will support the storage and retrieval of data by Content Manager OnDemand.

Configuring Tivoli Storage Manager to maintain Content Manager OnDemand data in archive storage

Provides general guidance about how to configure Tivoli Storage Manager to maintain Content Manager OnDemand data in archive storage.

About this task

Tivoli Storage Manager can maintain the reports that you load into Content Manager OnDemand, can maintain migrated index data, and can maintain DB2 archived log files and back up image files.

Before you begin, familiarize yourself with the following sources available in the IBM Knowledge Center:

- For information about configuring and managing server storage, see the *IBM Tivoli Storage Manager for AIX: Administrator's Guide*.
- For detailed information about the Tivoli Storage Manager commands, see the *IBM Tivoli Storage Manager for AIX: Administrator's Reference*. This guide is useful as your primary reference.

If you encounter problems configuring Tivoli Storage Manager, see the Tivoli Storage Manager publications.

Procedure

Complete these tasks to set up Tivoli Storage Manager for Content Manager OnDemand on an AIX workstation. Use the *IBM Tivoli Storage Manager for AIX: Administrator's Guide* and *IBM Tivoli Storage Manager for AIX: Administrator's Reference* for specific instructions on how to complete each task:

1. Define the Tivoli Storage Manager server options.
2. Define the Tivoli Storage Manager client system options.

   When you update the servers file, add the following line to turn off compression: `COMPRESSION OFF`
3. Define the Tivoli Storage Manager client options.
4. Register Tivoli Storage Manager licenses.
5. Register Tivoli Storage Manager administrators.
6. Define other Tivoli Storage Manager server options.
7. Start, halt, and restart Tivoli Storage Manager server.
8. Increase Tivoli Storage Manager database and recovery log sizes.
9. Define a storage library.
10. Define policy domains.
11. Register client nodes.

    You can use the information in “Registering client nodes” on page 24 to supplement the instructions provided by Tivoli Storage Manager.
12. Define archive copy groups.

    You can use the information in “Define archive copy groups” on page 24 to supplement the instructions provided by Tivoli Storage Manager.
13. Prepare storage pool volumes.
14. Optional: Configure Tivoli Storage Manager to maintain DB2 archived log files and back up image files.
15. Define a backup device for the Tivoli Storage Manager database.
16. Back up the Tivoli Storage Manager database and critical files.

Registering client nodes

A client node links clients and their data with storage volumes and devices. Before Content Manager OnDemand can store data in Tivoli Storage Manager storage, you must register at least one client node.

About this task

You must register at least one client node in each policy domain that will contain Content Manager OnDemand data. You can use the example that follows as a guide when registering client nodes. The example presents the procedure with a minimum of customization. If you want to do more, refer to the Tivoli Storage Manager documentation. Enter the command at the Tivoli Storage Manager server command line interface.

To register the client node PRI7YR and password `password`, assign the client node to the OD7YPD policy domain, and specify that the client node should be able to delete its own archive files from the server, enter:

```
    register node PRI7YR password domain=OD7YRPD archdel=yes contact='your
name'
```

The `archdel=yes` parameter is required for Content Manager OnDemand processing.

When you define a Content Manager OnDemand storage node (by using the Content Manager OnDemand facilities), specify a Tivoli Storage Manager client node and client node password to "link" the Content Manager OnDemand storage node to archive storage.

Define archive copy groups

Content Manager OnDemand stores data in Tivoli Storage Manager through its archive mechanism. The archive copy groups determine several Tivoli Storage Manager options for the data stored by Content Manager OnDemand, including the number of days that Tivoli Storage Manager maintains the files.

The archive copy group identifies the policy domain, policy set, and management class. The archive copy group also identifies the storage pool where Tivoli Storage Manager maintains the Content Manager OnDemand files and the length of time that Tivoli Storage Manager maintains them.

The archive copy also has two data retention types, creation-based and event-based. With creation-based retention, Tivoli Storage Manager expires data a fixed number of days after the data was loaded; this date is independent of settings defined in Content Manager OnDemand. Event-based retention keeps the data until an event is sent to Tivoli Storage Manager, in this case by Content Manager OnDemand, then the data expires based on its retention settings.

Content Manager OnDemand has the most control over data retention and expiration when the following conditions are met:

- Application groups in Content Manager OnDemand are set to expire by load.
- The archive copy group types used by Content Manager OnDemand are defined as event-based (RETINIT=EVENT) and the retain version (RETVER) and retain minimum (RETMIN) settings are both set to 0. This means the data expires immediately from Tivoli Storage Manager when Content Manager OnDemand expires its database information and cache data.
- Expiration on the Tivoli Storage Manager server is run regularly to ensure expired data is removed.

Configuring Tivoli Storage Manager to maintain DB2 files

You can use Tivoli Storage Manager to maintain DB2 archived log files and back up image files.

About this task

This capability means that you do not have to manually maintain these files on disk. The tasks in this section are optional, and are only recommended for customers who need to use Tivoli Storage Manager facilities to back up the Content Manager OnDemand database in DB2. For more information about using

Tivoli Storage Manager to manage DB2 files, see *IBM DB2 Universal Database: Data Recovery and High Availability Guide and Reference*, SC09-4831.

Do the following tasks to configure Tivoli Storage Manager to maintain DB2 files:

- Define server options
- Define client options
- Define storage objects
- Register the client node
- Set the client node password
- Review space requirements
- Review backup considerations

Protecting data with the data retention protection (DRP) protocol

To avoid the accidental erasure or overwriting of critical data, Content Manager OnDemand supports the Tivoli Storage Manager APIs related to data retention.

Data retention protection (DRP)
Prohibits the explicit deletion of documents until their specified retention criterion is met. Although documents can no longer be explicitly deleted, they can still expire.

Important: DRP is permanent. After it is turned on, it cannot be turned off.

Event-based retention policy
Retention based on an external event other than the storage of data. For Content Manager OnDemand, the retention event is the call to delete the data. A load, unload, application group delete, or expiration of data triggers the retention event.

Restriction: Content Manager OnDemand does not support deletion hold, which is a feature that prevents stored data from being deleted until the hold is released.

If you decide to use these policies in Tivoli Storage Manager, then the following scenarios result:

Table 3: Scenarios of using data retention protection

	Creation-based object expiration policy	**Event-based retention object expiration policy**
Data retention protection off	Content Manager OnDemand issues a delete object command through the Tivoli Storage Manager API. Objects are deleted during the next inventory expiration. If a Content Manager OnDemand application group is being deleted, a `delete filespace` command is issued, and the object file space is immediately deleted with the file space.	Content Manager OnDemand issues an event trigger command through the Tivoli Storage Manager API. The status of the objects that are affected are changed from PENDING to STARTED, and the objects are expired by Tivoli Storage Manager based on their retention parameters. If the retention parameters are set to `NOLIMIT`, then the objects never expire. If a Content Manager OnDemand application group is being deleted, a `delete filespace` command is issued instead, and the objects are immediately deleted along with the file space.

Table 3: Scenarios of using data retention protection (continued)		
	Creation-based object expiration policy	**Event-based retention object expiration policy**
Data retention protection on	Content Manager OnDemand issues no commands to Tivoli Storage Manager. The objects are effectively orphaned by Content Manager OnDemand and are expired by Tivoli Storage Manager based on their retention parameters. If the retention parameters are set to `NOLIMIT`, then the objects never expire.	Content Manager OnDemand issues an event trigger command through the Tivoli Storage Manager API. The event status of the objects that are affected are changed from PENDING to STARTED and the objects will be expired by Tivoli Storage Manager based on their retention parameters. If the retention parameters are set to `NOLIMIT`, then the objects never expire. If a Content Manager OnDemand application group is being deleted, then a `delete filespace` cannot be used with DRP enabled, therefore, the operation is treated the same as if a delete were indicated. The status of all the affected objects is changed from PENDING to STARTED, and they will be expired by Tivoli Storage Manager based on their retention parameters. Because this leaves the file space entries in TSM, you must manually delete these entries when the file space is empty (even with DRP enabled).

Recommendation: Set up the application groups to expire by load.

Installing the Content Manager OnDemand software on AIX

You must install a copy of the Content Manager OnDemand software on each workstation or node that is part of the Content Manager OnDemand system.

Before you begin

1. You need approximately 200 MB of free space in the `/usr` file system to install the software.
2. By default, the installation is carried out in the GUI mode, therefore, the X Windows support is required for the GUI install.

About this task

Complete the following steps to install the Content Manager OnDemand product files on an AIX workstation:

Procedure

1. Log in as the root user.
2. Go to the directory where `odaix.bin` is located.

3. Enter this command: `./odaix.bin`
4. Read the Welcome screen and then click **Next**. The License Agreement window appears.
5. Select **I accept the terms in the license agreement** to accept the license agreement. Click **Next**.
6. Accept the default directory name or, if you prefer a different directory name, type in the directory name. Click **Next**.

   If you have a version of Content Manager OnDemand older than Version 8.5 installed, the installation program removes the previous version before installing the new version.
7. When the process completes, this question,Would you like to display the product ReadMe file? appears. The location of the product readme file is displayed also. On AIX, the readme file is located in the `/opt/IBM/ondemand/V10.1` directory.
8. If you want to view the readme file now, click **Yes**. Otherwise, click **No**. Click **Next**.
9. Read the information in the window, and click **Next**.
10. Click **Finish**.
11. After installing the software, apply the latest service update for Content Manager OnDemand.

    You can obtain the latest service update from IBM service at http://www.ibm.com/eserver/support/fixes/.
12. After the installation completes successfully, eject the CD-ROM from the drive.

Results

Optionally, the installation can be performed in the character based console mode. To install the Content Manager OnDemand for AIX server in the console mode, enter the following command from the directory which contains the installer: `./odaix.bin -i console` and follow the instructions on the installation panels.

Installing optional Content Manager OnDemand software on AIX

Other software is available for installation in addition to Content Manager OnDemand software.

About this task

The command to install the Content Manager OnDemand PDF Indexing feature is:

```
./odpdfaix.bin
```

or

```
./odpdfaix.bin -i console
```

The command to install the IBM Content Manager OnDemand Distribution Facility feature is:

```
./ododfaix.bin
```

or

```
./ododfaix.bin -i console
```

The command to install the Content Manager OnDemand Full Text Search server feature is:

```
./odftsaix.bin
```

or

```
./odftsaix.bin -i console
```

To install the Content Manager OnDemand Enhanced Retention Management feature, see *Enhanced Retention Management Guide*.

Configuring instances on AIX

A Content Manager OnDemand instance is a logical server environment made up of a database, a library server, and one or more object servers.

An instance is defined in the ARS.INI file by naming the instance, identifying the name of the database used by the instance, and identifying the library server on which the database will be maintained. When you configure an object server, you identify its library server in the ARS.CFG file on the object server. An instance has its own table space file systems for the database and cache file systems. The table space file systems are defined in the ARS.DBFS file on the library server. The cache file systems are defined in the ARS.CACHE file on each object server. All of the servers that belong to an instance run in a single code page and on the same TCP/IP port number.

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

- To have separate test and production environments
- To have databases using different code pages

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

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

If you plan to run more than one instance on the same workstation:

- The ARS.INI file must contain one section for each instance. Each section identifies the ARS.CFG file, ARS.DBFS file, and ARS.CACHE file used by the instance.
- You must create a unique copy of the ARS.CFG file for each instance.
- You should maintain separate table space file systems and cache storage file systems for each instance, as in a ARS.DBFS file and ARS.CACHE file for each instance.
- Each instance must run on its own unique TCP/IP port number. The port for each instance is configured in the ARS.INI file.

Instances in the ARS.INI file

The ARS.INI file contains information about Content Manager OnDemand instances.

When you install the Content Manager OnDemand software, the ARS.INI file contains information about a default instance named archive. Most customers will use the default instance for their first or only instance of Content Manager OnDemand.

The information in the ARS.INI file is organized in sections with a header line that identifies each section. The header line can be identified by the brackets [ ] that delimit the beginning and end of the line.

The first section in the ARS.INI file contains information about the default instance. The following example shows the default instance as provided by IBM:

```
[@SRV@_ARCHIVE]
HOST=platte
PROTOCOL=2
PORT=0
SRVR_INSTANCE=archive
SRVR_INSTANCE_OWNER=root
SRVR_OD_CFG=/opt/IBM/ondemand/V10.1/config/ars.cfg
SRVR_DB_CFG=/opt/IBM/ondemand/V10.1/config/ars.dbfs
SRVR_SM_CFG=/opt/IBM/ondemand/V10.1/config/ars.cache
```

The HOST parameter identifies the host name alias, IP address, or fully qualified host name of the workstation on which the library server is running. The PROTOCOL parameter identifies the communications protocol used by the instance. The PORT parameter identifies the TCP/IP port number

that the instance monitors for client requests. The stanza name ([@SRV@_ARCHIVE]) identifies the name of the Content Manager OnDemand instance. The SRVR_INSTANCE parameter identifies the name of the Content Manager OnDemand database. The SRVR_INSTANCE_OWNER parameter identifies the userid of the Content Manager OnDemand instance owner. The SRVR_OD_CFG parameter identifies the ARS.CFG file used by the instance. The SRVR_DB_CFG parameter identifies the ARS.DBFS file used by the instance. The SRVR_SM_CFG parameter identifies the ARS.CACHE file used by the instance.

When adding an instance to the ARS.INI file, remember that each instance must specify a unique instance name. For example, to add an instance for testing new applications, you might add an instance named test. When you work with more than one instance, you must identify the instance name when you run Content Manager OnDemand programs (such as ARSDB, ARSLOAD, and ARSSOCKD) and database commands (such as connecting to the database). The following example shows a second instance in the ARS.INI file:

```
[@SRV@_TEST]
HOST=rhone
PROTOCOL=2
PORT=1444
SRVR_INSTANCE=test
SRVR_INSTANCE_OWNER=root
SRVR_OD_CFG=/opt/IBM/ondemand/V10.1/config/ars.test.cfg
SRVR_DB_CFG=/opt/IBM/ondemand/V10.1/config/ars.test.dbfs
SRVR_SM_CFG=/opt/IBM/ondemand/V10.1/config/ars.test.cache
```

The header line for the definition of the instance is TEST. The HOST parameter changes to rhone. The instance communicates over TCP/IP port number 1444. The name of the Content Manager OnDemand database is test. The name of the Content Manager OnDemand instance is test. The userid of the Content Manager OnDemand instance owner is root. The instance identifies its server configuration file (ARS.TEST.CFG), table space file systems file (ARS.TEST.DBFS), and cache file systems file (ARS.TEST.CACHE).

Specifying instances in the ARS.INI file

The ARS.INI file contains information about a default instance named archive. Most customers will use the default instance for their first or only instance of Content Manager OnDemand.

Procedure

To specify the instance in the ARS.INI file, follow these steps:

1. Log in to the server as the root user.
2. Change to the /opt/IBM/ondemand/V10.1/config directory.
3. Make a backup copy of the ARS.INI file provided by IBM.
4. Edit the ARS.INI file with a standard text editor such as vi.
5. Most customers will use the default instance named ARCHIVE for their first or only instance of Content Manager OnDemand.
6. **Note for distributed library/object servers:** Configure one copy of the ARS.INI file on each workstation that is part of the system. Verify that the information specified in the ARS.INI file is consistent on all workstations that are part of the instance. In addition:
   a) Ensure that the port number of the object server matches the port number of the library server.
   b) Verify that the HOST parameter on the object server must specify the host name alias, IP address, or fully qualified host name of the library server.
7. Save the file and exit the text editor.
8. You should control access to the ARS.INI file by changing the file permissions so that only the Content Manager OnDemand instance owner has read or write access to the file.

Verifying the default instance
Most customers will use the default instance named ARCHIVE for their first or only instance of Content Manager OnDemand.

Verify the following parameters and values:

- The header line contains a string that identifies the name of the instance. Unless you specify otherwise, the first or only instance is named ARCHIVE.
- The HOST parameter identifies the host name alias, IP address, or fully qualified host name of the library server.
- The PROTOCOL parameter identifies the communications protocol used by the instance. The number 2 identifies TCP/IP, and is the only valid value.
- The PORT parameter identifies the TCP/IP port number that the instance monitors for client requests (the number 0 means that the instance monitors port number 1445). If you use a port number other than 1445 on the library server, enter that number instead of 0 (zero). **For customers running more than one instance:** Each instance that runs on the same workstation must specify a different port number. If you configure a separate object server, ensure that the port number of the object server matches the port number of the library server.
- The stanza name ([@SRV@_ARCHIVE]) identifies the name of the Content Manager OnDemand instance. This value should match the name of the Content Manager OnDemand database (see "Installing DB2®" on page 14 or "Installing Oracle" on page 15). The instance name can be from one to eight characters in length, and can include the A through Z and 0 through 9 characters.
- The SRVR_INSTANCE_OWNER parameter identifies the user ID of the Content Manager OnDemand instance owner. This is the user ID that is permitted to run the Content Manager OnDemand server programs, such as ARSSOCKD, ARSLOAD, and ARSMAINT.
- The SRVR_OD_CFG parameter identifies the ARS.CFG configuration file used by the instance. See "Specifying the ARS.CFG file for the instance" on page 30.
- The SRVR_DB_CFG parameter identifies the ARS.DBFS table space file system file used by the instance. See "Specifying the ARS.DBFS file for the instance" on page 36.
- The SRVR_SM_CFG parameter identifies the ARS.CACHE cache file system file used by the instance. See "Creating the ARS.CACHE file for the instance" on page 38.
- The SRVR_OD_STASH parameter identifies the location of the stash file used by the instance and Content Manager OnDemand programs. See "Saving Content Manager OnDemand passwords into encrypted files" on page 21.

Specifying the ARS.CFG file for the instance

The ARS.CFG file contains information about the instance, such as identifying the object servers that belong to the instance, the language settings for the instance, and information that is used by database, storage, and print manager programs.

About this task

Before you create the Content Manager OnDemand database, start Content Manager OnDemand, use archive storage, use the server print function, migrate tables to table spaces, or import tables from archive storage to the database, you should review the parameters in the ARS.CFG file. The values that IBM provides are sufficient for most customers. However, you might need to change some of the values for your environment.

Procedure

To specify the ARS.CFG file for the instance, follow these steps:

1. Log in to the server as the root user.
2. Change to the /opt/IBM/ondemand/V10.1/config directory.
3. Make a backup copy of the file provided by IBM.
4. Edit the ARS.CFG file with a standard text editor such as vi.

5. **Note for distributed library/object servers:** Some parameters in the ARS.CFG file are not used on object servers. For example, an object server does not use the license parameters, server print parameters, and database parameters. See the sections that follow for more information. Configure one copy of the ARS.CFG file on each workstation that is part of the Content Manager OnDemand system. Set the ARS_SRVR parameter to the TCP/IP host name alias, fully qualified host name, or IP address of the library server and set the ARS_LOCAL_SRVR parameter to the TCP/IP host name alias, fully qualified host name, or IP address of the object server.
6. Save the file and exit the text editor.
7. You should control access to the ARS.CFG file by changing the file permissions so that only the Content Manager OnDemand instance owner has read or write access to the file.

ARS_DB_ENGINE parameter
The database manager product that you installed on the library server. You can specify DB2 or ORACLE. The default value is DB2. The ARS_DB_ENGINE parameter is ignored on object servers.

ARS_DB_IMPORT parameter
The method that Content Manager OnDemand uses to migrate index data to table spaces and import tables from archive storage to the database.

The default value is 0 (zero). The ARS_DB_IMPORT parameter is ignored on object servers.

If you are configuring a library server, then you must set the ARS_DB_IMPORT parameter to one of the following values:

0
Content Manager OnDemand uses the EXPORT and IMPORT commands to migrate table data. This method requires disk space to hold log records generated when exporting existing table data and importing data to the new table space. This is the default migration method.

1
Content Manager OnDemand uses the EXPORT and LOAD commands to migrate table data. This method requires disk space to hold log records generated when exporting existing table data. The LOAD command generates a backup image of the new table space. The image file is stored in Tivoli Storage Manager-managed storage. This is the recommended migration method. Before you can use Tivoli Storage Manager to manage DB2 backup image files, you must install and configure Tivoli Storage Manager. See "Installing and configuring Tivoli Storage Manager on AIX" on page 22 for details.

2
Content Manager OnDemand uses the EXPORT and LOAD commands to migrate the table data. This method requires disk space to hold log records generated when exporting existing table data. The LOAD command generates a backup image of the new table space. The image file is stored in the file system identified by the ARS_TMP parameter (see "ARS_TMP parameter" on page 36).

ARS_DB_PARTITION parameter
Determines whether you can partition the database across nodes or systems. By default, you cannot partition the database.

If the database manager product that you are using with Content Manager OnDemand supports partitioning, then you can specify that you want to partition the database by changing the value of this parameter to 1 (one). Content Manager OnDemand supports partitioning with DB2 Universal Database Extended Enterprise Edition only. To store application group index data in partitions, your application groups must specify a partition field. The ARS_DB_PARTITION parameter is ignored on object servers.

ARS_DB_TABLESPACE parameter
The name of the table space for the Content Manager OnDemand system tables.

The value of this parameter must match an existing table space name in the database. You must have created the table space in DB2 or Oracle.

ARS_DB_TABLESPACE_USEREXIT parameter
Determines if the Content Manager OnDemand table space creation exit will be invoked.

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

The following statement must exist in the ARS.CFG file that is associated with the instance so that the ARSUTBL DLL can be invoked:

```
ARS_DB_TABLESPACE_USEREXIT=absolute path name
```

For the sample ARSUTBL, you would specify the following statement in the ARS.CFG file:

```
ARS_DB_TABLESPACE_USEREXIT=/opt/IBM/ondemand/V10.1/bin/exits/arsutbl
```

"Table space creation user exit" on page 166 provides information about the exit point that gets invoked when Content Manager OnDemand creates table spaces, tables, and indexes for the Content Manager OnDemand data tables.

ARS_DB2_DATABASE_PATH parameter
If you are configuring the library server and you set the ARS_DB_ENGINE parameter to DB2 (the default), the base file system in which the Content Manager OnDemand database will reside.

You must make sure that the specified location contains enough space to hold the system tables, the USERSPACE1 table space, and any application group tables that are not stored in their own table spaces. The *IBM Content Manager OnDemand for Multiplatforms: Introduction and Planning Guide* can help you estimate the amount of space required to hold the database. The default value is /arsdb. The ARS_DB2_DATABASE_PATH parameter is ignored on object servers.

ARS_DB2_LOG_NUMBER parameter
If you are configuring the library server and you set the ARS_DB_ENGINE parameterto DB2 (the default), the number of primary log files.

The default value is 40. The ARS_DB2_LOG_NUMBER parameter is ignored on object servers.

The values of the ARS_DB2_LOGFILE_SIZE and ARS_DB2_LOG_NUMBER parameters determine the total amount of space available for DB2 to log changes to the database. The values that you specify must support the largest single report that you plan to load (or unload). DB2 will fail if there is not enough log file space available to hold the changes to the database. The default values allocate 160 MB of space. See the *IBM Content Manager OnDemand for Multiplatforms: Introduction and Planning Guide* for information about estimating the amount of storage space required to hold the DB2 log files.

ARS_DB2_LOGFILE_SIZE parameter
If you are configuring the library server and you set the ARS_DB_ENGINE parameter to DB2 (the default), the size of a log file, in 4 KB blocks. The default value is 1000. The ARS_DB2_LOGFILE_SIZE parameter is ignored on object servers.

The values of the ARS_DB2_LOGFILE_SIZE and ARS_DB2_LOG_NUMBER parameters determine the total amount of space available for DB2 to log changes to the database. The values that you specify must support the largest single report that you plan to load (or unload). DB2 will fail if there is not enough log file space available to hold the changes to the database. The default values allocate 160 MB of space. See the *IBM Content Manager OnDemand for Multiplatforms: Introduction and Planning Guide* for information about estimating the amount of storage space required to hold the DB2 log files.

ARS_DB2_PRIMARY_LOGPATH parameter
If you are configuring the library server and you set the ARS_DB_ENGINE parameter to DB2 (the default), the location that will hold the active archived log files.

The *IBM Content Manager OnDemand for Multiplatforms: Introduction and Planning Guide* can help you estimate the amount of space required to hold the active archived log files. The default value is /arsdb_primarylog. The ARS_DB2_PRIMARY_LOGPATH parameter is ignored on object servers.

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

ARS_LDAP_BASE_DN parameter
Specifies the base distinguished name to use.

This parameter is required for LDAP authentication.

Example 1:

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

Example 2:

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

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

This parameter is required for LDAP authentication.

Example:

```
ARS_LDAP_BIND_ATTRIBUTE=mail
```

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

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

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

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

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

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

Example:

```
ARS_LDAP_MAPPED_ATTRIBUTE=sAMAccountName
```

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

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

ARS_LOCAL_SRVR parameter
The name of the object server. The ARS_LOCAL_SRVR parameter is ignored on library servers.

However, if you are configuring a library server, you must either omit this parameter from the ARS.CFG file or set this parameter to a blank value, that is: ARS_LOCAL_SRVR= .

If you are configuring an object server, set this parameter to the TCP/IP host name alias, fully qualified host name, or IP address of the object server. If the object server is running on a node of a multi-processor workstation, then set this parameter to the external IP address of the node on which you installed the object server.

When you add a Content Manager OnDemand storage node to an object server, you must use the value of the ARS_LOCAL_SRVR parameter to name the storage node.

ARS_MESSAGE_OF_THE_DAY parameter
Use to show the message of the day. Set to the full path name of a file that contains the message that you want to show.

For example:

```
ARS_MESSAGE_OF_THE_DAY=/opt/IBM/ondemand/V10.1/tmp/message.txt
```

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

ARS_NUM_DBSRVR parameter
Determines the number of processes that Content Manager OnDemand starts on the library server to handle connections to the database. The ARS_NUM_DBSRVR parameter is ignored on object servers.

In addition to database connections by Content Manager OnDemand client programs, the value that you specify must support the number of active Content Manager OnDemand commands and daemons such as ARSLOAD, ARSDOC, ARSDB, ARSMAINT, and ARSADMIN.

Each connection to the database requires a database agent. Content Manager OnDemand can start a database agent for each connection. However, each agent requires its own private memory and some portion of application shared memory. You can use the ARS_NUM_DBSRVR parameter to optimize the way that Content Manager OnDemand handles the database load. For example, you can define ARS_NUM_DBSRVR so that Content Manager OnDemand starts a fixed number of database agents, regardless of the number of concurrent database requests. While this might appear restrictive, database requests typically process very quickly. For example, ten database agents can handle a heavy database request load, while balancing the impact on system resources.

You should specify a value for the ARS_NUM_DBSRVR parameter that supports the peak number of concurrent database connections that you expect the library server to handle. A low value limits access to the database during periods of high database activity. A high value requires more system resources during periods of high database activity. The value that you choose also depends on the characteristics of the queries. For example, general queries typically keep a connection open longer than a more specific query.

ARS_ORACLE_HOME parameter
Use to specify the base installation directory for Oracle.

The default value is:

```
ARS_ORACLE_HOME=/oracle
```

Replace the string `/oracle` with the name of the directory in which Oracle was installed.

ARS_PRINT_PATH parameter
The location where the Content Manager OnDemand server print function temporarily stores print data.

You must make sure that there is enough space in the specified location to hold the print files for the maximum number of concurrent print requests that the server will handle. The default value is `/tmp`. The ARS_PRINT_PATH parameter is ignored on object servers.

You should dedicate a file system to hold the print files. The file system contain at least 500 MB of free space at all times. If your storage configuration permits, you should allocate 1 GB or more of free space to the specified file system.

The permissions for the file system must be `drwxrwxrwt`. You can use the `CHMOD` command to set the permissions. For example, the command `chmod 1777 /tmp` sets the permissions for the `/tmp` file system.

ARS_SRVR parameter
The name of the library server. The ARS_SRVR parameter is ignored on library servers.

However, if you are configuring a library server, you must either omit this parameter from the `ARS.CFG` file or set this parameter to a blank value, that is: `ARS_SRVR= .`

If you are configuring an object server, set the ARS_SRVR parameter to the TCP/IP host name alias, fully qualified host name, or IP address of the library server. If the library server is running on a node of a multi-processor workstation, then set this parameter to the external IP address of the node on which you installed the library server.

ARS_STORAGE_MANAGER parameter
Determines whether the server program is linked to a cache-only storage manager or an archive storage manager. You must specify this parameter on library and object servers.

You can specify one of the following values:

CACHE_ONLY

Link the server program to a cache-only storage manager.

TSM

Link the server program to an archive storage manager. This is the default value in the `ARS.CFG` file that is provided by IBM. Before Content Manager OnDemand can work with an archive storage manager to maintain data, you must install and configure the archive storage manager software.

ADSM

Deprecated. This option has been replaced by TSM. ADSM is still supported for existing customers.

ARS_SUPPORT_CFSOD parameter
If you plan to use Content Federation Services for Content Manager OnDemand, you must set this parameter equal to 1.

ARS_SUPPORT_HOLD parameter
To use enhanced retention management, you must set this parameter to 1.

ARS_TMP parameter
The location where Content Manager OnDemand programs temporarily store data.

You must allocate sufficient free space in the specified file system to support tasks such as migrating and importing index data. The default value is: `/tmp`. You must specify the ARS_TMP parameter on the library server and on all object servers.

You should dedicate a file system to temporary storage. The file system should contain at least 500 MB of free space at all times. If your storage configuration permits, you should allocate 1 GB or more of free space to the specified file system.

The permissions for the file system must be `drwxrwxrwt`. You can use the `CHMOD` command to set the permissions. For example, the command `chmod 1777 /tmp` sets the permissions for the `/tmp` file system.

DB_ENGINE parameter
Deprecated. This parameter has been replaced by ARS_DB_ENGINE.

However, the DB_ENGINE parameter is still supported for existing customers.

DB2INSTANCE parameter
If you are configuring the library server and you set the ARS_DB_ENGINE parameter to DB2 (the default), the name of the database instance owner that you created when you installed DB2.

The default value is `archive`. The DB2INSTANCE parameter is ignored on object servers.

DSMI_CONFIG parameter
If you plan to use Tivoli Storage Manager, the full path name of the Tivoli Storage Manager API options file.

For example: `/usr/tivoli/tsm/client/api/bin64/dsm.opt`

You must set the DSMI_CONFIG parameter on each object server that uses Tivoli Storage Manager to maintain Content Manager OnDemand data.

DSMI_DIR parameter
If you plan to use Tivoli Storage Manager, the directory that contains the Tivoli Storage Manager API files.

For example: `/usr/tivoli/tsm/client/api/bin64`

You must set the DSMI_DIR parameter on each object server that uses Tivoli Storage Manager to maintain Content Manager OnDemand data.

DSMI_LOG parameter
If you plan to use Tivoli Storage Manager, the directory in which Tivoli Storage Manager stores the Tivoli Storage Manager API error log.

The default value is `/tmp`. You must set the DSMI_LOG parameter on each object server that uses Tivoli Storage Manager to maintain Content Manager OnDemand data.

Specifying the ARS.DBFS file for the instance

The `ARS.DBFS` file lists the file systems on the library server that can be used by the database manager to maintain index data in table spaces. The rules for using table space file systems are:

- You should store only Content Manager OnDemand application group data in the table space file systems.

- You should define a minimum of two table space file systems. (In general, the more table space file systems that you define, the better for performance and recovery.)
- You should allocate equal amounts of disk space to each table space file system. If you increase the amount of space in one table space file system, you should increase the amount of space in the other table space file systems by an equal amount.
- For DB2, if you are planning on using Automatic Storage for tablespaces, you do not need to define any file systems for the database tablespace containers in the ARS.DBFS file.

Each line in the ARS.DBFS file identifies the name of a file system that Content Manager OnDemand can use to store table spaces and specifies the type of table spaces created in the file system.

When naming table space file systems, you should use the following convention:

```
/filesystem SMS
```

Where `filesystem` is the name of the file system and SMS indicates the type of table spaces created in the file system. The name of the file system should identify the type of table spaces that can be created in the file system. For example, the following line identifies an SMS table space file system:

```
/arsdb/db1/SMS SMS
```

The following example shows an ARS.DBFS file that defines three SMS table space file systems:

```
/arsdb/db1/SMS SMS
/arsdb/db2/SMS SMS
/arsdb/db3/SMS SMS
```

Creating the ARS.DBFS file for the instance

The ARS.DBFS file lists the file systems on the library server that can be used by the database manager to maintain index data in table spaces.

Procedure

To create (or edit) the ARS.DBFS file for the instance:

1. Log in to the server as the `root` user.
2. Change to the `/opt/IBM/ondemand/V10.1/config` directory.
3. Create (or edit) the ARS.DBFS file using a standard text editor such as vi.
4. Add one line for each file system that Content Manager OnDemand can use for table spaces.
5. Save the file and exit the editor.
6. A table space file system must be owned by the database instance owner and group. The suggested defaults are `archive` (instance owner) and `db2iadm1` (group). You specified the instance owner and group when you installed the database manager product (see "Installing the database manager on Linux™" on page 58). Make sure that the user and group file permissions are set correctly.
   For example:

   ```
   drwxrws---  3 archive db2iadm1  512 May 17 12:58 /arsdb/db1/SMS
   ```

   You can use the CHOWN command to set the ownership permissions. For example, the following command changes the owner of all file systems in the `/arsdb` tree to the `archive` user and the `db2iadm1` group: `chown -R archive:db2iadm1 /arsdb*`
7. You can use the CHMOD command to set the file permissions.
   For example, the following commands set the correct permissions for the `/arsdb/db1/SMS` filesystem:

   ```
   chmod 2770 /arsdb/db1/SMS
   chmod g+s /arsdb/db1/SMS
   ```

Creating the ARS.CACHE file for the instance

About this task

The ARS.CACHE file lists the file systems on the object server that can be used by Content Manager OnDemand for cache storage.

If there are multiple file systems in the ARS.CACHE file, Content Manager OnDemand uses the file system with the greatest amount of space free to store the objects.

The following example shows an ARS.CACHE file that defines five cache storage file systems:

```
/arscache/cache1
/arscache/cache2
/arscache/cache3
/arscache/cache4
/arscache/cache5
```

Note: For a distributed library / object server system, configure one copy of the ARS.CACHE file on each server that is part of the Content Manager OnDemand system.

Procedure

To create the ARS.CACHE file for the instance:

1. Log in to the server as the `root` user.
2. Change to the `/opt/IBM/ondemand/V10.1/config` directory.
3. Create (or edit) the ARS.CACHE file using a standard text editor such as vi.
4. Insert one line in the file for each file system on the server that Content Manager OnDemand can use for cache storage.

   Important: The first entry in the ARS.CACHE file identifies the base cache storage file system. Content Manager OnDemand stores control information in the base cache storage file system. After you define the base cache storage file system to Content Manager OnDemand, you cannot add or remove it from Content Manager OnDemand. It must remain as the first entry.
5. Save the file and exit the editor.
6. Cache file systems must be owned by the Content Manager OnDemand instance owner and the `system` group.

   Make sure that only the user file permissions are set, not the group or other file permissions.

   For example:

   ```
   drwx------  3 root system  512 Sep 22 13:08 /arscache/cache1
   ```
7. Use the CHOWN command to set the ownership permissions.
   The following example shows how to change the user and group ownership of all file systems in the `/arscache` tree: `chown -R root:system /arscache*`
8. Use the CHMOD command to set the file permissions.
   For example, the following commands set the correct permissions for the `/arscache/cache1` file system:

   ```
   chmod 700 /arscache/cache1
   chmod g-s /arscache/cache1
   ```

Results

Content Manager OnDemand cache storage files and subdirectories should have the following permissions:

```
drwx------ for every subdirectory (700)
-r-------- for every object that has been migrated to archive storage (400)
-rw------- for every object that has not yet been migrated (600)
-rwxrwxrwx for every symbolic link under the retr and migr directories (777)
```

Specifying the ARSLDAP.INI file

The ARS_LDAP_BIND_MESSAGES_FILE parameter enables Content Manager OnDemand to customize message text returned from an LDAP server that is used to alert users that their LDAP password is about to expire or their LDAP account is locked.

The messages displayed to users are contained in the file referenced by this parameter. To enable this user-configurable message functionality, create a file with the appropriate message strings, and set ARS_LDAP_BIND_MESSAGES_FILE to the full path of the file. The ARSLDAP.INI file is provided with example message strings that can be used by the ARS_LDAP_BIND_MESSAGES_FILE parameter.

The ARSLDAP.INI file contains the following three sections:

```
[BIND_MESSAGES]
PASSWORD_EXPIRED="/opt/IBM/ondemand/V10.1/config/password_expired.txt"
ACCOUNT_LOCKED="/opt/IBM/ondemand/V10.1/config/account_locked.txt"

[PASSWORD_EXPIRED]
TDS6="Password has expired"
AD="data 532"
UDEF1=
UDEF2=
UDEF3=

[ACCOUNT_LOCKED]
TDS6="Account is locked"
AD="data 775"
UDEF1=
UDEF2=
UDEF3=
```

The BIND_MESSAGES section specifies the path to the files containing the user-configurable message text that is displayed to users when their LDAP password is about to expire, or their LDAP account is locked. Generic files are supplied, and should be customized to reflect your actual Content Manager OnDemand environment.

An example message that would be displayed to a user:

```
Your LDAP password has expired and needs to be changed.
Log into <company intranet> for password setting instructions.
```

The entries in the PASSWORD_EXPIRED and ACCOUNT_LOCKED sections are for Tivoli Directory Server Version 6.x and Microsoft Active Directory (AD). These sections also contain three user-defined entries (UDEFx), allowing you to enter your own pattern strings for LDAP servers that are not directly supported.

The LDAP server may return additional information when the user's bind operation fails. When an error is returned from the LDAP server, Content Manager OnDemand reads the file referenced by the ARS_LDAP_BIND_MESSAGES_FILE parameter and searches under the two stanzas, [PASSWORD_EXPIRED] and [ACCOUNT_LOCKED], for user-defined text that matches the LDAP server error. If a match is found, Content Manager OnDemand will display the text found in the files defined under the [BIND_MESSAGES] stanza.

If the ARS_LDAP_BIND_MESSAGES_FILE parameter is not defined, has no file referenced, or the PASSWORD_EXPIRED or ACCOUNT_LOCKED files do not exist, the user will receive a default 'The server failed while attempting to logon' message.

Currently only two error conditions can be handled: PASSWORD_EXPIRED and ACCOUNT_LOCKED. The section titles for these two conditions cannot be changed, but you can change the pattern strings and message text presented to the user to define any two error conditions.

To create an instance of Content Manager OnDemand

You create an instance of Content Manager OnDemand by running the ARSDB program on the library server.

About this task

The ARSDB program initializes the base system tables that are required by Content Manager OnDemand. You initialize other system tables by running the ARSSYSCR program on the library server. The ARSSYSCR program initializes the system tables that are required to support the system log, system migration, and other Content Manager OnDemand functions.

Prerequisites

Ensure that DB2 or Oracle have been installed and the database has been created.

Procedure

Before you create the instance, you must have completed the following tasks:

1. Installed and configured the database software (DB2 or Oracle), and created a database instance (DB2) or database (Oracle) for Content Manager OnDemand.
2. Installed and configured the Content Manager OnDemand software, including the following files:
   - `ARS.INI`
   - `ARS.CFG`
   - `ARS.DBFS`
   - `ARS.CACHE`

Creating an instance of Content Manager OnDemand on AIX®

Procedure

To create the instance, follow these steps:

1. Specify permissions for the database directories.
2. Create the instance by running the ARSDB program.
3. Initialize the system logging facility by running the ARSSYSCR program.
4. (Optional) Initialize the system load logging facility by running the ARSSYSCR program.
5. (Optional) Initialize the system migration facility by running the ARSSYSCR program.

Specifying permissions for the database directories

The group that the DB2 instance owner belongs to must have write access to the database directory names that are specified in the `ARS.CFG` file (the ARS_DB2_DATABASE_PATH and ARS_DB2_PRIMARY_LOGPATH parameters).

About this task

You created the DB2 instance owner and group when you installed the database manager.

Changing owners of directories

You can change the owners of database directories with the CHOWN command.

Procedure

To change the owner of the directories:

1. Log in to the server as the `root` user.
2. Use the CHOWN command to change directory ownership.
   For example, to change the owner and group of all file systems in the `/arsdb` tree to the `archive` owner and the `db2iadm1` group, enter the following command: `chown -R archive:db2iadm1 /arsdb*`

   Run the CHOWN command once to change the ownership of each of the database directories that are specified in the `ARS.CFG` file (the ARS_DB2_DATABASE_PATH, ARS_DB2_PRIMARY_LOGPATH, and ARS_DB2_ARCHIVE_LOGPATH parameters).

To create the database instance

You should use the ARSDB program to create the instance.

The ARSDB program completes the following tasks to create the instance:

- Updates the database configuration
- Verifies the directories for the primary and archived log files
- Creates a link to the database user exit program

  If the database user exit program encounters errors when copying files, it creates the `db2uexit.err` file in the temporary data directory. If this file exists, it usually means that you did not set the correct permissions for the log file directories or there is not enough free space to hold the archived log files. See your operating system documentation for information about increasing the size of a file system.
- Creates a backup of the database
- Builds the Content Manager OnDemand system tables and indexes

  The ARSDB program can build the Content Manager OnDemand system tables and indexes into a default table space or user-defined table spaces. If you want to use the default table space, continue with the instructions in this topic. If you want to use user-defined tables spaces, follow the instructions in Chapter 13, “Creating Content Manager OnDemand system tables into user-defined table spaces,” on page 177 before continuing with the instructions in this topic.
- Binds the database to Content Manager OnDemand

The ARSDB program creates the database using standard SQL commands. See the documentation provided with the database manager product for information about the SQL commands issued by the ARSDB program and messages printed at the console.

Creating a database instance

The ARSDB program creates the instance.

About this task

Oracle users: You must create the database by using the Oracle utilities before you create the Content Manager OnDemand instance.

Procedure

To create an instance of Content Manager OnDemand:

1. Log in to the server as the `root` user.

Option

DB2 Type the following command at the prompt:

```
/opt/IBM/ondemand/V10.1/bin/arsdb -I archive -gcv
```

Where `archive` is the name of the Content Manager OnDemand instance.

Oracle Type the following command at the prompt:

```
/opt/IBM/ondemand/V10.1/bin/arsdb -I archive -rtv
```

Where `archive` is the name of the Content Manager OnDemand instance.

2. Press the Enter key.
3. The ARSDB program prompts you before creating a link to the database user exit program:
   - If you maintain DB2 archived log files on disk, enter 1 when prompted
   - If you use Tivoli Storage Manager to maintain the DB2 archived log files, enter 2 when prompted

   Content Manager OnDemand creates the instance, makes a backup image of the database, and restores the Content Manager OnDemand system tables to the database. This process will take several minutes.

   The ARSDB program generates a series of messages. For example:

```
Creating the DB2 ARSDBASE database
Creating table ARSSERVR.arsag
Creating index ARSSERVR.arsag_name_idx
Creating index ARSSERVR.arsag_agid_idx
..... .. . . . . . . . . . . . . . . . . . .
..... .. . . . . . . . . . . . . . . . . . .
..... .. . . . . . . . . . . . . . . . . . .
Updating runstat statistics for table ARSSERVR.arsusrgrp
Creating table ARSSERVR.arsusrgrpid
Creating index ARSSERVR.arsusrgrpid_idx
Updating runstat statistics for table ARSSERVR.arsusrgrpid
```

Initializing the system logging facility

After you have successfully created the instance of Content Manager OnDemand, run the ARSSYSCR program to initialize the Content Manager OnDemand system logging facility for the instance.

Procedure

To initialize the Content Manager OnDemand system logging facility:

1. Log in to the server as the `root` user.
2. Type the following command at the prompt: `/opt/IBM/ondemand/V10.1/bin/arssyscr -I archive -l`

   Where `archive` is the name of the Content Manager OnDemand instance.

Initialize the system load logging facility

Content Manager OnDemand provides an optional logging facility to enable tracking Content Manager OnDemand loading activity.

About this task

When you enable load logging, Content Manager OnDemand stores the messages that are generated by Content Manager OnDemand load programs in the system load log. You use one of the Content Manager OnDemand client programs to search for and filter messages by load date, application group name, load ID, input file name, and other parameters.

Procedure

To initialize the Content Manager OnDemand system load logging facility:

1. Log in to the server as the `root` user.
2. Type the following command at the prompt: `/opt/IBM/ondemand/V10.1/bin/arssyscr -I archive -a`
3. Press the Enter key.
4. Content Manager OnDemand creates the tables that support the system load logging facility. This process may take several minutes.
   The ARSSYSCR program generates a series of messages. For example:

```
arssyscr: Updating ARSSERVR.ARSSYS
arssyscr: Adding to ARSSERVR.ARSAG with Storage Set Id = 0
arssyscr: Adding to ARSSERVR.ARSAGPERMS
arssyscr: Adding to ARSSERVR.ARSAGFLD
arssyscr: Adding to ARSSERVR.ARSAGFLDALIAS
arssyscr: Adding to ARSSERVR.ARSAG2FOL
arssyscr: Adding to ARSSERVR.ARSAPPUSR
arssyscr: Adding to ARSSERVR.ARSAPP
arssyscr: Adding to ARSSERVR.ARSFOL
arssyscr: Adding to ARSSERVR.ARSFOLPERMS
arssyscr: Adding to ARSSERVR.ARSFOLFLD
arssyscr: Adding to ARSSERVR.ARSFOLFLDUSR
arssyscr: Creation of System Load information was successful
```

Initializing the system migration facility

The system migration facility is required only by customers who plan to migrate application group index data from the database to archive storage.

About this task

After you have successfully created the instance of the Content Manager OnDemand, run the ARSSYSCR program to initialize the Content Manager OnDemand system migration facility for the instance.

Procedure

To initialize the Content Manager OnDemand system migration facility:

1. Log in to the server as the `root` user.
2. Type the following command at the prompt: `/opt/IBM/ondemand/V10.1/bin/arssyscr -I archive -m`

   Where `archive` is the name of the Content Manager OnDemand instance.
3. Press the Enter key.
   Content Manager OnDemand creates the tables that support the system migration facility. This process may take several minutes.

   The ARSSYSCR program generates a series of messages. For example:

```
arssyscr: Updating ARSSERVR.ARSSYS
arssyscr: Adding to ARSSERVR.ARSAG with Storage Set Id = 0
arssyscr: Adding to ARSSERVR.ARSAGPERMS
arssyscr: Adding to ARSSERVR.ARSAGFLD
arssyscr: Adding to ARSSERVR.ARSAGFLDALIAS
arssyscr: Adding to ARSSERVR.ARSAG2FOL
arssyscr: Adding to ARSSERVR.ARSAPPUSR
arssyscr: Adding to ARSSERVR.ARSAPP
arssyscr: Adding to ARSSERVR.ARSFOL
arssyscr: Adding to ARSSERVR.ARSFOLPERMS
arssyscr: Adding to ARSSERVR.ARSFOLFLD
```

```
arssyscr: Adding to ARSSERVR.ARSFOLFLDUSR
arssyscr: Creation of System Migration information was successful
```

Automating instance operations on AIX®

This section describes how to use operating system facilities to automatically start or schedule instance operations.

Procedure

You can automatically start these instance operations whenever the system is started:

1. Start the database on the library server.
2. Start the instance on the library server.
3. Start the instance on an object server.
4. Start the data loading programs.

Results

You can schedule these instance operations to begin automatically on a regular schedule:

1. Schedule application group maintenance on the library server.
2. Schedule application group maintenance on an object server.
3. Schedule system table maintenance.
4. Schedule a backup of the Content Manager OnDemand database.
5. Schedule a backup of the Tivoli Storage Manager database.

Starting the database

You can start the database on the library server using the ARSDB program.

Procedure

- To update the command, enter: `opt/IBM/ondemand/V10.1/bin/arsdb -gv >> /tmp/arsdb.log 2>&1`

  Alternatively, you can start DB2 manually with the db2start command.

  The following example shows an INIT record to automatically start the database when the operating system is initialized on the library server:

  ```
  ars2:2:wait:/opt/IBM/ondemand/V10.1/bin/arsdb -gv >> /tmp/arsdb.log 2>&1
  ```

  Important: If the DB2 installation program adds a record to the INIT facility to automatically start the DB2 services, make sure that you place the ARSDB record after the record that starts the DB2 services.

Starting the instance on the library server

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

About this task

The ARSSOCKD program controls a Content Manager OnDemand instance on the library server. The ARSSOCKD program runs on the library server. The data loading program (ARSLOAD) and the maintenance programs (such as ARSADMIN and ARSMAINT) will fail and clients will be unable to connect to the instance if the ARSSOCKD program is not running on the library server.

Procedure

Enter the command: `/opt/IBM/ondemand/V10.1/bin/arssockd archive`

Example

The following example shows an INIT record that automatically starts the instance named `archive` when the operating system is initialized on the library server:

```
ars3:2:once:/opt/IBM/ondemand/V10.1/bin/arssockd archive
```

Starting the instance on an object server

The ARSOBJD program controls a Content Manager OnDemand instance on an object server.

About this task

Content Manager OnDemand programs that work with an instance on an object server will fail if the ARSOBJD program is not running on the object server.

The ARSOBJD program should be started only on object servers that are running on some other workstation than the library server.

Procedure

Enter the command: `/opt/IBM/ondemand/V10.1/bin/arsobjd archive`

Example

The following example shows an INIT record that automatically starts the instance named `archive` when the operating system is initialized on an object server:

```
ars4:2:once:/opt/IBM/ondemand/V10.1/bin/arsobjd archive
```

Starting the data loading programs

This section describes how to use operating system facilities to automatically start the Content Manager OnDemand data loading programs.

About this task

The Content Manager OnDemand data loading programs are:

- ARSJESD, to receive data from z/OS systems and store the data in file systems on the server
- ARSLOAD, to create index data and load the data into the system

ARSJESD data load program
The ARSJESD program is the Content Manager OnDemand program that monitors a TCP/IP port for data transmitted to the Content Manager OnDemand server by Download for the z/OS feature from a host system.

About this task

The ARSJESD program receives the data transmitted by Download for the z/OS feature and stores the data in file systems on the server. See *PSF for z/OS: Download for z/OS* for details about configuring and operating Download for the z/OS feature on the host system.

Example

The following example shows an INIT record that automatically starts the ARSJESD program during operating system initialization:

```
ars5:2:once:/opt/IBM/ondemand/V10.1/bin/arsjesd -p 6001 -d /arsacif/acif1
-d /arsacif/acif2 -d /arsacif/acif3 >> /tmp/arsjesd.log 2>&1
```

In the example, the ARSJESD program monitors TCP/IP port number `6001` and stores transmitted data in the specified directories. The ARSJESD program writes output messages to the `arsjesd.log` file in the `/tmp` directory.

You must verify the TCP/IP port number that the ARSJESD program monitors. Replace the string `6001` with the port number that is valid on the server that you are configuring. The ARSJESD program and Download on the z/OS system must specify the same port number. The port number that the ARSJESD program monitors is different than the TCP/IP port number that the Content Manager OnDemand server uses to communicate with clients.

You must verify the names of the directories in which the ARSJESD program can put the data. Replace the strings `/arsacif/acif1`, `/arsacif/acif2`, and `/arsacif/acif3` with the names of directories that are valid on the server that you are configuring.

See the ARSJESD command reference in the *IBM Content Manager OnDemand for Multiplatforms: Administration Guide* for more information about the options and parameters that you can specify.

ARSLOAD data loading program

The ARSLOAD program is the main Content Manager OnDemand data loading and indexing program.

You can configure the ARSLOAD program to monitor specific file systems for report data downloaded from other systems. If the data needs to be indexed, then the ARSLOAD program calls the indexing program that is specified in the Content Manager OnDemand application. The ARSLOAD program then works with the database manager to load the index data into the database and works with the storage manager to load the report data and resources on to storage volumes.

The Content Manager OnDemand instance (started by using ARSSOCKD or ARSOBJD) must be running, otherwise the ARSLOAD program will fail.

Automating the ARSLOAD program

In the example, the ARSLOAD program checks for input files in the specified directories every ten minutes (the default polling time). An input file must have a file type of .ARD or .PDF to initiate the load process. If an input file needs to be indexed, the ARSLOAD program stores the index data in the specified index directory.

Example

The following example shows an INIT record that automatically starts the ARSLOAD program for the instance named `archive` during operating system initialization:

```
ars6:2:once:/opt/IBM/ondemand/V10.1/bin/arsload -v -c /arsacif/acif4 -d /arsacif/acif1
-d /arsacif/acif2 -d /arsacif/acif3 -I archive
```

You must verify the names of the directories. Replace the strings `/arsacif/acif1`, `/arsacif/acif2`, `/arsacif/acif3`, and `/arsacif/acif4` with the names of directories that are valid on the server that you are configuring.

After indexing the data, the ARSLOAD program deletes the input files, unless you specify otherwise. Any output or error messages that are generated by the ARSLOAD program are written to `stdout`, `stderr`, and the system log.

See the *IBM Content Manager OnDemand for Multiplatforms: Administration Guide* for more information about the ARSLOAD program.

Scheduling application group maintenance on the library server

You can run the ARSMAINT program on the library server to maintain application group data in the database and cache storage.

About this task

See the *IBM Content Manager OnDemand for Multiplatforms: Administration Guide* for more information about the ARSMAINT program.

The instance must be started by the ARSSOCKD program, otherwise the ARSMAINT program will fail.

Example

The following example is a CRON record that automatically starts the ARSMAINT program every day at 4 am for the instance named `archive`. The ARSMAINT program will migrate and delete application group index data, optimize application group index data, copy report data from cache storage to archive media, delete report data from cache storage, and inspect and verify the cache file systems. This format of the command is typically used for a library/object server with Tivoli Storage Manager on one workstation.

```
00 4 * * * /opt/IBM/ondemand/V10.1/bin/arsmaint -cdeimrsv -I archive
```

Scheduling application group maintenance on an object server

You can run the ARSMAINT program on an object server to maintain application group data in cache storage.

About this task

See the *IBM Content Manager OnDemand for Multiplatforms: Administration Guide* for more information about the ARSMAINT program.

The instance must be started by the ARSOBJD program, otherwise the ARSMAINT program will fail.

Example

The following example is a CRON record that automatically starts the ARSMAINT program every day at 4 am for the instance named `archive`. The ARSMAINT program will maintain application group data in cache storage, including copying report data to archive storage. This format of the command is typically used for an object server with Tivoli Storage Manager on some other workstation than the library server.

```
00 4 * * * /opt/IBM/ondemand/V10.1/bin/arsmaint -cmsv
```

Scheduling system table maintenance

You can run the ARSDB program to maintain the Content Manager OnDemand system tables on the library server.

About this task

See the *IBM Content Manager OnDemand for Multiplatforms: Administration Guide* for more information about the ARSDB program.

The instance must be started by the ARSSOCKD program, otherwise the ARSDB program will fail.

Example

The following example is a CRON record that automatically starts the ARSDB program to maintain the Content Manager OnDemand system tables for the instance named `archive`. The ARSDB program will run twice a month, on the 7th and 14th of each month, beginning at 5 am.

```
00 5 7,14 * * /opt/IBM/ondemand/V10.1/bin/arsdb -mv -I archive >> /tmp/
arsdb.log 2>&1
```

Scheduling the Content Manager OnDemand database backup

Use the backup utilities that come with your database product to schedule backups of the Content Manager OnDemand database. Otherwise, you can use the ARSDB program to create a backup image of the Content Manager OnDemand database.

About this task

The ARSDB program supports table space backups and full database backups, online backups and offline backups, and the use of Tivoli Storage Manager to maintain the backup image files. See the *IBM Content Manager OnDemand for Multiplatforms: Administration Guide* for more information about the ARSDB program.

Example

The following example is a CRON record that automatically starts the ARSDB program to create a full online backup image of the Content Manager OnDemand database for the instance named `archive` every day beginning at 5:30 am. The backup image is written to a tape in the device `/dev/rmt0`. A tape must be mounted in the device before the ARSDB program begins.

```
30 5 * * * /opt/IBM/ondemand/V10.1/bin/arsdb -v -z /dev/rmt0 -I archive >> /tmp/arsdb.log 2>&1
```

Next steps on AIX®

After you have installed the Content Manager OnDemand and related software on the system, configured the instance of Content Manager OnDemand, created the instance, and automated instance operations, you are now ready to verify the installation on Content Manager OnDemand.

Related tasks

Verifying the installation

After you have completed installation and configuration of the database manager, Content Manager OnDemand software, and Tivoli Storage Manager software, and have configured and initialized the system, perform the following tasks.

Chapter 3. Installing Content Manager OnDemand on Linux™ servers

Install and configure Content Manager OnDemand on a Linux server. Install and configure related software that works with Content Manager OnDemand.

About this task

There are five basic phases to the installation:

- Preparing for the installation
- Installing and configuring Content Manager OnDemand and related software
- Verifying the installation
- Preparing the system for use
- Adding optional software

You will find checklists for each of these phases in “Checklist for installation on Linux™” on page 50.

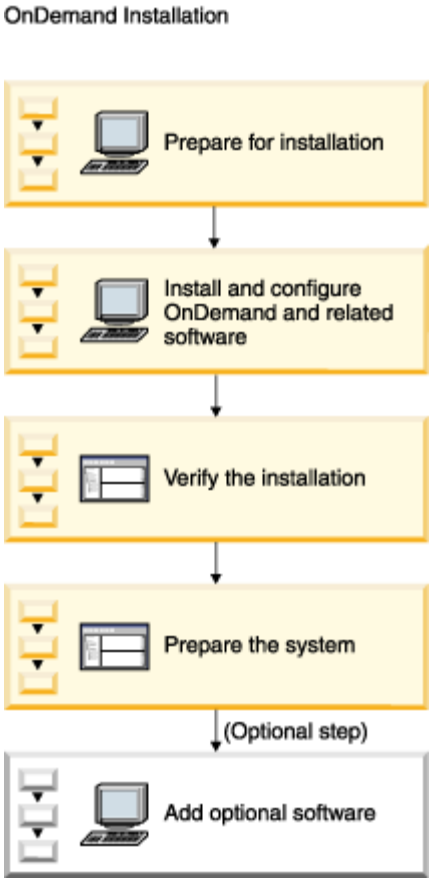

Figure 7: Installing Content Manager OnDemand on a Linux server

Checklist for installation on Linux™

You should review the pre-installation instructions checklist before installing the product on Linux.

About this task

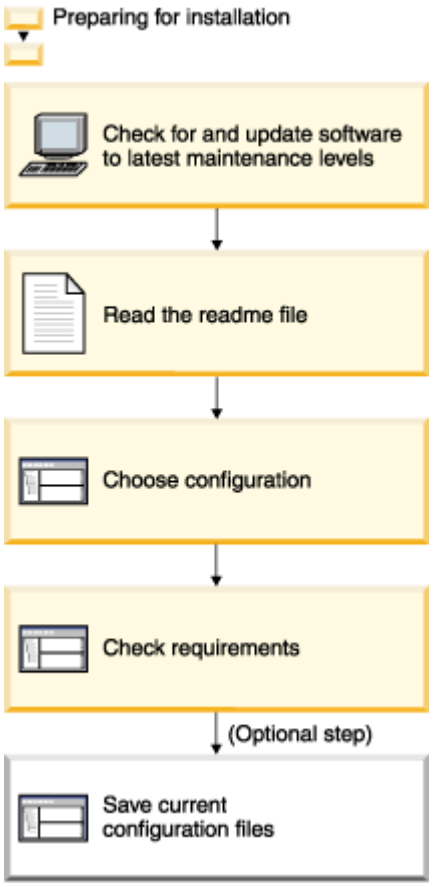

Figure 8: Pre-installation tasks

Before beginning the installation, you should complete the following tasks:

Procedure

1. Contact the IBM Support Center for the latest maintenance levels of DB2, Content Manager OnDemand, and optionally, Tivoli Storage Manager and Infoprint Manager (Infoprint).
2. Obtain a copy of the latest Content Manager OnDemand README file. Print and read the entire file before you begin.
3. Check the Content Manager OnDemand prerequisites and verify the required and optional hardware and software products (see “Linux™ server requirements” on page 54).
4. Check the hardware and software requirements for all system components and features. See http://www.ibm.com/support/docview.wss?uid=swg27016455 for details.
5. Determine the type of system configuration that you need to install (see “Choosing a configuration” on page 2).
6. If you are upgrading to a new version of Content Manager OnDemand, save the configuration files used by the system (see “Saving configuration files on Linux™” on page 55).

Results

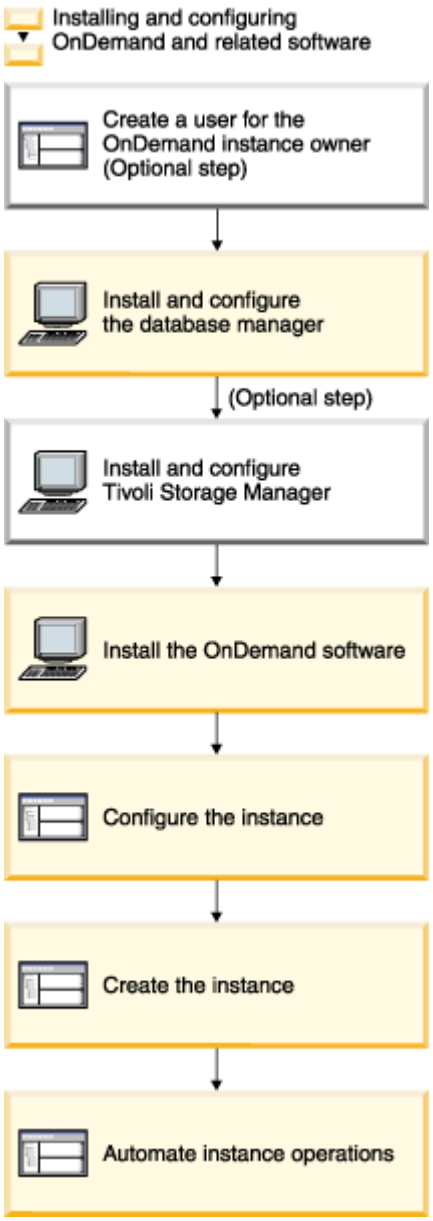

Figure 9: Installing Content Manager OnDemand and related software

Configuring a Content Manager OnDemand system typically requires that you do the following tasks:

1. (Optional) Create a user account for the Content Manager OnDemand instance owner (see "Creating a user for the Content Manager OnDemand instance owner on Linux" on page 57).
2. Install and configure the database manager product on the library server (see "Installing the database manager on Linux™" on page 58).
3. Install the IBM Global Security Kit (GSKit) (see "Installing IBM® Global Security Kit on Linux™" on page 60).
4. If you plan on using SSL for security, set up SSL on the Content Manager OnDemand server and client (see "Setting up SSL on the Content Manager OnDemand for Linux™ server" on page 61).
5. If you plan to maintain data in archive storage, install and configure Tivoli Storage Manager on the library server or on each object server that will be used to maintain data in archive storage (see "Installing and configuring Tivoli Storage Manager on Linux™" on page 65).

6. Install the Content Manager OnDemand software on each workstation that is part of the Content Manager OnDemand system (see “Installing the Content Manager OnDemand software on Linux” on page 70).
7. Configure an instance of Content Manager OnDemand on each workstation that is part of the Content Manager OnDemand system (see “Configuring instances on Linux™” on page 71). This step includes the following tasks:
   a. Specify the instance in the ARS.INI file (see “Specifying instances in the ARS.INI file” on page 73)
   b. Specify the ARS.CFG file for the instance (see “Specifying the ARS.CFG file for the instance” on page 74)
   c. Specify the ARS.DBFS file for the instance (see “Specifying the ARS.DBFS file for the instance” on page 79)
   d. Specify the ARS.CACHE file for the instance (see “Creating the ARS.CACHE file for the instance” on page 80)
8. Create the instance of Content Manager OnDemand (see “Creating an instance of Content Manager OnDemand on Linux” on page 82). This step includes the following tasks:
   a. Specify permissions for the database directories (see “Specifying permissions for the database directories” on page 82)
   b. Create the instance by running the ARSDB program (see “Creating an instance of Content Manager OnDemand on Linux™” on page 82)
   c. Initialize the system logging facility by running the ARSSYSCR program (see “Initializing the system logging facility” on page 84)
   d. (Optional) Initialize the system migration facility by running the ARSSYSCR program (see “Initializing the system migration facility” on page 85)
   e. (Optional) Initialize the system load logging facility by running the ARSSYSCR program (see “Initialize the system load logging facility” on page 42).
9. Automate instance operations (see “Automating instance operations on Linux™” on page 86). This step includes the following tasks:
   a. Start the database on the library server (see “Starting the database” on page 86)
   b. Start the instance on the library server (see “Starting the instance on the library server” on page 87)
   c. Start the instance on an object server (see “Starting the instance on an object server” on page 87)
   d. Start the data loading programs (see “Starting the data loading programs” on page 87)
   e. Schedule application group maintenance on the library server (see “Scheduling application group maintenance on the library server” on page 89)
   f. Schedule application group maintenance on an object server (see “Scheduling application group maintenance on an object server” on page 89)
   g. Schedule system table maintenance (see “Scheduling system table maintenance” on page 89)
   h. Schedule a backup of the Content Manager OnDemand database (see “Scheduling the Content Manager OnDemand database backup” on page 90)

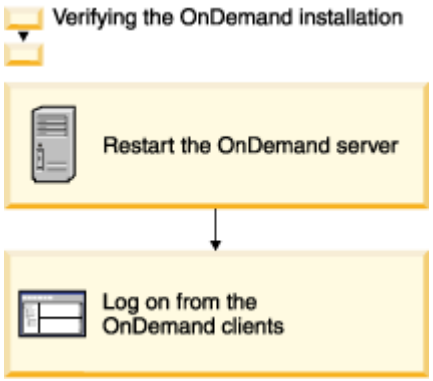

Figure 10: Verifying the installation

Verify the installation of Content Manager OnDemand (see “Verifying the installation” on page 137):

1. After installing and configuring each Content Manager OnDemand server, restart the system. The operating system reinitializes and starts the services required by Content Manager OnDemand.
2. Log in to the library server with a Content Manager OnDemand client program. (To access the system, you must install at least one of the Content Manager OnDemand client programs on a PC running Microsoft Windows. See the *IBM Content Manager OnDemand: Client Installation Guide* for installation information about the Content Manager OnDemand client or the *IBM Content Manager OnDemand for Multiplatforms: Administration Guide* for installation information about the Content Manager OnDemand administrative client.)

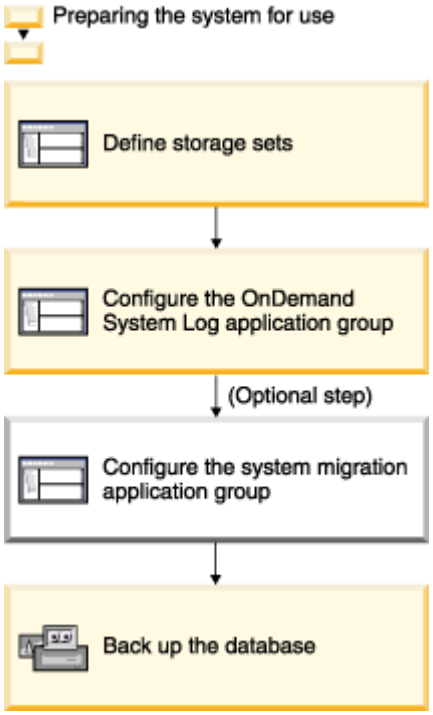

Figure 11: Preparing the system for use

Prepare the system for use:

1. Define storage sets (see “Define storage sets” on page 138). Before you add application groups or load data into the system, you must define storage sets.
2. Configure the System Log application group (see “Configuring the System Log application group” on page 138). Before you define reports to the system, load data, or let users access the system, you should configure the System Log application group.

3. If you plan to migrate index data to archive storage, configure the System Migration application group (see "Configure the System Migration application group" on page 142).
4. Back up the databases (see Chapter 7, "Backing up the Content Manager OnDemand database," on page 145). After configuring the system, you should create a full backup image of the Content Manager OnDemand database and the Tivoli Storage Manager database.

Installing and configuring optional software:

1. If you plan to use Download for the z/OS feature (Download) to transmit data from z/OS systems to Content Manager OnDemand servers, then you must install and configure Download. Follow the instructions in *PSF for z/OS: Download for z/OS* to plan, install, configure, and verify the installation of the Download software. Then configure Download on each Content Manager OnDemand server. Complete the following tasks:
   a. Obtain a copy of *PSF for z/OS: Download for z/OS*.
   b. Check the prerequisites and verify the z/OS and TCP/IP software levels for Download.
   c. Install and configure the Download software.
   d. Configure Download on each Content Manager OnDemand server that will receive data sets from an z/OS system. (see "Starting the data loading programs" on page 87).
2. If you plan to reprint documents using the Content Manager OnDemand server print function, you must install Infoprint on a workstation that belongs to the same network as the Content Manager OnDemand library server. Follow the instructions in the Infoprint documentation for your operating system to plan, install, configure, and verify the installation of the Infoprint software. Then configure the server print function on the library server. Complete the following tasks:
   a. Obtain a copy of the Infoprint documentation for the server operating system.
   b. Install and configure Infoprint.
   c. Verify that all of the resources and fonts that your organization requires to reprint the reports that you plan to store in Content Manager OnDemand are installed on the Infoprint server.
   d. Define the print queues and devices that Infoprint uses to manage the Content Manager OnDemand server print environment.
   e. Obtain the TCP/IP host name or IP address of the Infoprint server.
   f. On the library server, edit the ARSPRT file and insert the host name or IP address of the Infoprint server. The ARSPRT file can be found in the `/opt/ibm/ondemand/V10.1/bin` directory.
   g. Define a server printer on the Content Manager OnDemand library server with the administrative client.
3. If you need to customize and enhance the standard functionality within the product, see the user exit documentation in the Appendix of this publication. A user exit is a point during processing that enables you to run a user-written program and return control of processing after your user-written program ends. Content Manager OnDemand provides the following user exit points:
   a. Download user exit
   b. Report specifications archive definition user exit
   c. Retrieval preview user exit
   d. Security user exit
   e. System log user exit
   f. Table space creation user exit

Linux™ server requirements

The exact hardware and software configuration that you need for Content Manager OnDemand to support your organization depends on the volume of data that you plan to maintain on the system, the number of

concurrent users that the system must support, the backup and recovery requirements of your organization, and the performance levels that the system must meet.

At a minimum, you need one processor for a standard Content Manager OnDemand library/object server.

Restriction: Linux is designed as a server-only platform, and requires Windows to run the system administration client.

For all Linux server requirements, see http://www.ibm.com/support/docview.wss?uid=swg27049168

Saving configuration files on Linux™

When you install software on a Content Manager OnDemand server, the installation programs copy program files, configuration files, and other types of files from the distribution media to directories on the server.

About this task

When you configure a server to meet the specific requirements of your environment, you make changes to configuration files and you might customize other files, such as user-defined files and font initialization files.

Before you upgrade to a new version of Content Manager OnDemand or upgrade the database manager software or other software related to Content Manager OnDemand, you should save a copy of the files listed in this section. You can save a copy of the files in a temporary directory, such as `/tmp`.

After you upgrade the software, you will probably need to reconfigure the files for your environment. To reconfigure the files, you can restore the copies of the files that you saved or make changes to the updated files, using the configuration information in the files that you saved as a guide.

Content Manager OnDemand files

Save a copy of the Content Manager OnDemand configuration files.

Table 4: Content Manager OnDemand configuration files to save

File	Location	Purpose
`ars.cache`	`/opt/ibm/ondemand/V10.1/config`	Define cache storage file systems. Changes described in "Creating the ARS.CACHE file for the instance" on page 80.
`ars.cfg`	`/opt/ibm/ondemand/V10.1/config`	Content Manager OnDemand server configuration file. Changes described in "Specifying the ARS.CFG file for the instance" on page 74.
`ars.dbfs`	`/opt/ibm/ondemand/V10.1/config`	Define DB2 table space file systems. Changes described in "Specifying the ARS.DBFS file for the instance" on page 79.
`ars.ini`	`/opt/ibm/ondemand/V10.1/config`	Configure Content Manager OnDemand instances. Changes described in "Specifying instances in the ARS.INI file" on page 73.
`arslog`	`/opt/ibm/ondemand/V10.1/bin`	The System Log user exit program. Described in "System log user exit" on page 162.
`arsprt`	`/opt/ibm/ondemand/V10.1/bin`	Server print program.

Configuring the library server

Create a user that is a member of the database owner's group. This group has administrator authority for the database and the database file systems.

About this task

Give the Content Manager OnDemand instance owner the following authorities and permissions:

- Administrator authority for the database. You can do this by adding the Content Manager OnDemand instance owner to the database owner's group.
- Ownership of the cache storage file systems that are listed in the `ARS.CACHE` file. You can do this by running the Change Owner command for each file system that is listed in the `ARS.CACHE` file and specifying the user and group for the Content Manager OnDemand instance owner.
- Permission to read the Content Manager OnDemand configuration files. Make sure that the Content Manager OnDemand instance owner has permission to read the following files:
  - `ARS.CACHE`
  - `ARS.CFG`
  - `ARS.DBFS`
  - `ARS.INI`
- Permission to read and execute the Content Manager OnDemand script files. Make sure that the Content Manager OnDemand instance owner has permission to read and execute the following files:
  - `ARSLOG`
  - `ARSPRT`
- Permission to write to the console. Make sure that the Content Manager OnDemand instance owner has permission to write to the system console.

Important: You cannot set the permissions to read and execute OnDemand files until you complete installation of the Content Manager OnDemand software. See "Installing the Content Manager OnDemand software on Linux" on page 70 for instructions on installing the Content Manager OnDemand software on Linux.

You should specify a different user for each instance that you create. This allows for easier error recovery if a system error occurs.

Configuring an object server

If you plan to run a distributed library/object server system, with one or more object servers on different workstations or nodes than the library server, then you should also configure Content Manager OnDemand on each of the object servers.

Procedure

To configure Content Manager OnDemand on the object servers, do the following tasks:

1. Create a user for the Content Manager OnDemand instance owner.
2. Give ownership of the cache storage file systems listed in the `ARS.CACHE` file to the user for the Content Manager OnDemand instance owner.
3. Give permission to read the following files to the Content Manager OnDemand instance owner:
   - `ARS.CACHE`
   - `ARS.CFG`
   - `ARS.INI`
4. Give permission to write to the console to the Content Manager OnDemand instance owner.

Tivoli Storage Manager files

If you use Tivoli Storage Manager to maintain OnDemand data in archive storage, save a copy of the Tivoli Storage Manager configuration files.

Table 5: Tivoli Storage Manager configuration files to save

File	Location	Purpose
`dsmserv.dsk`	`/opt/tivoli/tsm/server/bin`	Locations of the Tivoli Storage Manager database and recovery logs
`history.dev`	`/opt/tivoli/tsm/server/bin`	Tivoli Storage Manager device history file
`history.vol`	`/opt/tivoli/tsm/server/bin`	Tivoli Storage Manager storage volume history file
`dsmserv.opt`	`/opt/tivoli/tsm/server/bin`	Tivoli Storage Manager server options file
`dsm.opt`	`/opt/tivoli/tsm/client/ba/bin`	Tivoli Storage Manager client options file
`dsm.db2.opt`	`/usr/tivoli/tsm/client/api/bin64`	Tivoli Storage Manager client options file for maintaining DB2 archived log files and back up files.
`dsm.sys`	`/usr/tivoli/tsm/client/api/bin64`	Tivoli Storage Manager client system options file

Creating a user for the Content Manager OnDemand instance owner on Linux

About this task

This publication was written assuming that OnDemand instances will be run under the root user. The information in this section is provided for customers who need to run instances of Content Manager OnDemand under a user other than the root user. Those customers should print the information in this section and have it available to assist them as they continue with the installation and configuration process.

New installations (instances) of Content Manager OnDemand can be configured to run under a user other than the root user. If you plan to run an instance under a user other than root, you must do the following:

- Create the user for the Content Manager OnDemand instance owner
- Set permissions for the cache storage file systems
- Set permissions for the Content Manager OnDemand configuration and script files
- Give the instance owner permission to write to the system console
- Specify the instance owner in the ARS.INI file

If you plan to run a distributed library/object server system, with one or more object servers on different workstations or nodes than the library server, then you should also configure Content Manager OnDemand on the object servers.

Installing the database manager on Linux™

The Content Manager OnDemand library server maintains system information and user-defined index data in a relational database.

About this task

You can use DB2 Universal Database or Oracle as the database manager. For either product, see the product documentation for complete installation instructions. This section provides installation and configuration information specific to Content Manager OnDemand for DB2 and Oracle.

Installing DB2®

You must install DB2 or Oracle on the Content Manager OnDemand library server.

About this task

This section describes how to install DB2. See "Installing Oracle" on page 59 for instructions on installing Oracle.

The DB2 Universal Database Enterprise Edition program DVDs or electronic images are provided with the Content Manager OnDemand program package. The DB2 technical information is available in HTML and PDF formats on separate DVDs or electronic images. The README file explains how to locate the information that you need. Follow the instructions in the *IBM DB2 Universal Database Quick Beginnings for DB2 Servers* to plan, install, configure, and verify the installation of DB2.

Procedure

To install DB2 on the library server:

1. Install DB2 Universal Database Enterprise Edition.
2. When prompted, select Typical as the installation type, to install all DB2 components required to support Content Manager OnDemand. You can take most default options (unless you have specific requirements of your own).
3. Create the DB2 instance for Content Manager OnDemand when you install DB2.

   Use the following values:

Parameter	Value
Instance Name or User	archive
Group Name	gname **Note:** The group must have SYSADM authority, and its name must be unique. The group name on your database might be something other than 'gname'. Ask your database administrator if you do not know the group name for your database.
Home Directory	`/home/archive`
Auto start DB2 instance at boot time	no
Create a sample database for DB2 instance	no

What to do next

After you install the software, apply the latest fix pack for DB2. You can obtain the latest fix packs at http://www.ibm.com/support/docview.wss?uid=swg27007053. Print the README file. Follow the instructions in the README file to apply the service update. After installing a fix pack, you might need to update your database instances (for example, archive). See the DB2 README for details.

Adding the user to the DB2 instance owner group

After you install DB2, you must add the user that owns the Content Manager OnDemand instance to the DB2 instance owner's group.

After installing DB2 on the library server:

1. Add the user that owns the Content Manager OnDemand instance to the DB2 instance owner's group.

   For example, if the DB2 instance owner's group is `db2iadm1` and the Content Manager OnDemand instance owner is `root`, specified by the SRVR_INSTANCE_OWNER parameter in the `ARS.INI` file, add the `root` user to the db2iadm1 group.
2. Create links for the DB2 files. For example: `/opt/IBM/db2/V10.1/cfg/db2ln`. See the instructions in *IBM DB2 Universal Database Quick Beginnings for DB2 Servers* to create links to the DB2 files.
3. Optionally create a table space for the Content Manager OnDemand system tables. If you plan to store the system tables in their own table space, specify the name of the table space on the ARS_DB_TABLESPACE parameter in the `ARS.CFG` file.
4. Verify the value of the DB2INSTANCE parameter in the `ARS.CFG` file. The value of the DB2INSTANCE parameter is case-sensitive. This value must specify the name of the DB2 instance that you created for Content Manager OnDemand. The default value is `archive`. .

Setting the DB2® operating environment

If you plan to use DB2 commands to work with the Content Manager OnDemand database, you must execute a script file to set the DB2 operating environment before you start the DB2 command line interface.

About this task

For Bourne or Korn shell users, run the `DB2PROFILE` script file. For C shell users, run the `DB2CSHRC` script file.

The script files can be found in the `INSTHOME/sqllib` directory, where `INSTHOME` is the home directory of the instance owner. If you installed and configured the system using the suggested defaults, the instance owner is `archive` and the script files reside in the `sqllib` directory under `/home/archive`.

Procedure

Add the script file to your `.profile` or `.login` file.
For example: `. /export/home/archive/sqllib/db2profile`

What to do next

After executing the script file, you can start the DB2 command line interface and connect to the database. For example:

```
$>db2
  .
  .
  .
```

To stop the DB2 command line interface, enter: `db2 =>quit`

Installing Oracle

You must install Oracle on the Content Manager OnDemand library server.

About this task

After you verify the installation of the Oracle software on the library server, you must configure it to work with Content Manager OnDemand.

Procedure

To configure Oracle to work with Content Manager OnDemand:

1. Configure login processing to run under the UID of the root user.
2. Create the Content Manager OnDemand database using the Oracle utilities. The name that you specify for the database should match the value that you specify for the SRVR_INSTANCE parameter in the ARS.INI file.
3. Create the userid of the Content Manager OnDemand instance owner in Oracle.

   This user will own all tables that Content Manager OnDemand creates. If you want to have a default Oracle table space for the user, specify the table space when you create the user.

   To create the Content Manager OnDemand user in Oracle:

   ```
   CREATE USER root IDENTIFIED BY password ;
   GRANT dba to root ;
   ```

   Where *root* and *password* are the user ID and password stored in the stash file.
4. Specify the base Oracle installation directory on the ARS_ORACLE_HOME parameter in the ARS.CFG file. The default value is /oracle.
5. Specify Oracle as the database manager on the ARS_DB_ENGINE parameter in the ARS.CFG file.
6. Optional: Create a table space for the Content Manager OnDemand system tables.

   If you plan to store the system tables in their own table space, specify the name of the table space on the ARS_DB_TABLESPACE parameter in the ARS.CFG file.

Installing IBM® Global Security Kit on Linux™

Determine whether you already have GSKit installed on your system and, if so, which version of GSKit.

Procedure

Before you can install GSKit, you must do the following steps:

1. If another product required that you install GSKit in one of the following default locations, run the gsk8ver command or the gsk8ver_64 command to determine which version of GSKit is installed, for example: /usr/local/ibm/gsk8.

   If another product required that you install GSKit in a location other than the default location, continue to the next step to install GSKit in the default location.
2. Log in to your system as the user root.

   Important: Do not use the nodeps flag unless instructed.
3. Determine whether you need to install the 32-bit or the 64-bit version of GSKit by reviewing the following list:

   - The Content Manager OnDemand server requires the 64-bit version of GSKit.
   - The ODWEK Java API can use the 32-bit or 64-bit version of GSKit.
4. Run the following rpm commands to install GSKit in the default location /usr/local.

   - For the 32-bit version:

     ```
     rpm -Uv gskcrypt32-8.0.14.44.linux.x86_64.rpm
     rpm -Uv gskssl32-8.0.14.44.linux.x86_64.rpm
     ```

   - For the 64-bit version:

     ```
     rpm -Uv gskcrypt64-8.0.14.44.linux.x86_64.rpm
     rpm -Uv gskssl64-8.0.14.44.linux.x86_64.rpm -
     ```

SSL with Content Manager OnDemand

Secure Sockets Layer (SSL) and its successor, Transport Layer Security (TLS), encrypts all transmissions between the Content Manager OnDemand servers and many of the supported clients (for example, ODWEK Java API, the Windows client, and arsdoc).

The following clients do not support SSL connections:

- ODWEK CGI
- ODWEK Java servlet
- CICS client

If you are not familiar with how a server and a client establish an SSL connection, see "Overview of the SSL handshake" on page 173.

Before you continue reading these instructions, verify that you understand the information in Chapter 12, "SSL, certificates, certificate authorities, and public-key cryptography," on page 173

Before you begin setting up SSL on Content Manager OnDemand for Linux

Because of possible problems with system performance, create SSL connections only for communications requiring secure transmission. Consider adding additional processor resources on the Content Manager OnDemand server, client, or both to manage the increased overhead.

GSKit provides the GSKCapiCmd tool, which helps you create and manage digital certificates and key databases. The instructions in "Setting up SSL on the Content Manager OnDemand for Linux™ server" on page 61 provide examples of how to run the GSKCapCmd tool; however, to view the complete syntax and understand the behavior of this tool, see ftp://ftp.software.ibm.com/software/webserver/appserv/library/v61/ihs/GSK7c_CapiCmd_UserGuide.pdf.

Choose the scenario from the following list that fits your requirements, then follow the instructions for that scenario:

- Content Manager OnDemand server listens only on a non-SSL port. You cannot set up SSL for this situation. Continue to the next Content Manager OnDemand installation task.
- Content Manager OnDemand server listens only on a SSL port. You must do the following tasks:
  - Set up SSL on Content Manager OnDemand for Linux server (see "Setting up SSL on the Content Manager OnDemand for Linux™ server" on page 61).
  - Install GSKit on all clients.
  - Configure the clients to support SSL (see "Setting up SSL on the Content Manager OnDemand client" on page 176).
- Content Manager OnDemand server listens on both a non-SSL port and a SSL port. You must do the following steps:
  - Set up SSL on Content Manager OnDemand for Linux server (see "Setting up SSL on the Content Manager OnDemand for Linux™ server" on page 61).
  - Install GSKit on the clients connecting to the SSL port.
  - Configure those clients to support SSL (see "Setting up SSL on the Content Manager OnDemand client" on page 176).

Setting up SSL on the Content Manager OnDemand for Linux™ server

You can set up SSL on the Content Manager OnDemand for Linux server and certify the file by using GSKCapiCmd tool.

Procedure

To set up SSL on the Linux server:

1. Create the key database and store it in the `config` subdirectory of Content Manager OnDemand server installation directory: `/opt/ondemand/config`

   To create the key database, run a command similar to the following command:

```
gsk8capicmd_64 -keydb -create -db "ondemand.kdb" -pw "myKeyDBpasswd" -stash -
populate
```

   The following list describes why these parameters were chosen:

   -keydb -create -db "ondemand.kdb"

   Indicates that you want to create a key database called ondemand.kdb.

   -pw "myKeyDBpasswd" -stash

   Indicates that you want to create a stash file and store the password (`myKeyDBpasswd`) in that stash file. The GSKCapiCmd tool stores the stash file at the same path as the key database. You must remember this path because you must specify it in the ars.ini file. GSKCapiCmd creates the stash file with the same file name as the key database (`ondemand`), with the file extension of `.sth`. When Content Manager OnDemand starts, GSKit retrieves the password to the key database from this stash file.

   -populate

   Populates the key database with a set of predefined trusted certificate authority (CA) certificates. A trusted CA is a certificate authority root certificate is noted as trusted in the key database. For the list of default trusted root certificates, see "Default GSKit trusted root certificates" on page 175.

2. Create a digital certificate. You can create a self-signed certificate, which is useful for testing. When you are ready to move to a production environment, create a CA-signed digital certificate.

3. Configure the Content Manager OnDemand for Linux server.

   Add the following lines to the `ARS.INI` file:

```
SSL_PORT=port_number
SSL_KEYRING_FILE=/opt/ondemand/config/ondemand.kdb
SSL_KEYRING_STASH=/opt/ondemand/config/ondemand.sth
SSL_KEYRING_LABEL=IBM Content Manager OnDemand
SSL_CLNT_USE_SSL=0
```

   The following list describes these parameters:

   SSL_PORT

   Specify one of the following values:

   port_number

   The port number on the Content Manager OnDemand server dedicated to communicating with the SSL protocol.

   0

   No port on the Content Manager OnDemand server to communicate with the SSL protocol.

   -1

   All ports on the Content Manager OnDemand server to communicate only with the SSL protocol.

   SSL_KEYRING_FILE

   Specify the full path and file name of the key database that contains the digital certificates.

   SSL_KEYRING_STASH

   Specify the full path and file name of the stash file for the key database.

   SSL_KEYRING_LABEL

   Specify the name of the certificate in the key database.

   SSL_CLNT_USE_SSL

Specify whether the server-side clients (for example, ARSDOC, ARSMAINT, or ARSLOAD) must communicate with the SSL protocol. Specify 0 to indicate the clients not communicate with the SSL protocol. Specify 1 to indicate the clients must communicate with the SSL protocol.

4. Restart the Content Manager OnDemand server. Because a trusted certificate authority provided the digital certificate, the Content Manager OnDemand server accepts the certificate. Communication between server and client can commence without updating the key database on Content Manager OnDemand client.

Creating a self-signed certificate

You can create a self-signed certificate by using GSKCapiCmd.

Procedure

To create a self-signed certificate, do the following steps:

1. Create a self-signed certificate by using GSKCapiCmd.
   The following example creates a self-signed certificate with the label myselfsigned:

```
gsk8capicmd_64 -cert -create -db "ondemand.kdb" –pw "myKeyDBpasswd" -label
"myselfsigned" \
–dn
"CN=myhost.mycompany.com,O=myOrganization,OU=myOrganizationUnit,L=Boulder,ST=CO,C=US
"
```

2. Extract the certificate to a file by using GSKCapiCmd.
   The following example extracts the certificate into a file called ondemand.arm:

```
gsk8capicmd_64 -cert -extract -db "ondemand.kdb" -pw " myKeyDBpasswd " -label
"myselfsigned" \
-target "ondemand.arm" -format ascii
```

3. Distribute the file you created to all computers that run clients that will establish SSL connections to your Content Manager OnDemand server.

Creating a CA-signed digital certificate

You create CA-signed digital certificate for an RSA private-public key pair and PKCS10 certificate request, which are stored in the key database in a file with the .rdb extension.

About this task

Specify the name of the file, with the -file option, that you send to the CA.

Procedure

To create a CA-signed digital certificate:

1. Create a Certificate Signing Request (CSR) by using GSKCapiCmd.
   The following example shows how to create a CSR that is stored in ondemand.kdb.

```
gsk8capicmd_64 -certreq -create -db "ondemand.kdb" -pw "myKeyDBpasswd" -label
"mycert" \
-dn
"CN=myhost.mycompany.com,O=myOrganization,OU=myOrganizationUnit,L=Boulder,ST=CO,C=US
" \
-file "mycertRequestNew"
```

2. Verify the contents of the CSR by using GSKCapiCmd.
   The following example shows how to display the contents of the CSR:

```
gsk8capicmd_64 -certreq -details -db "ondemand.kdb" -pw "myKeyDBpasswd" –label
"mycert"
```

If you need to delete this CSR, run GSKCapiCmd similar to the following example:

```
gsk8capicmd_64 -certreq -delete -db "ondemand.kdb" -pw "myKeyDBpasswd" -label
"mycert"
```

3. Go to the website of a well-known CA (for example, Verisign) and follow their instructions for registering and obtaining a signed digital certificate. The instructions include paying the CA for their services and providing them with the file you specified with the -file option. In the following example and for the rest of these instructions, a trial version of a digital certificate is used.
4. Use a text editor (for example, vi) to save each certificate into a file. The CA sends you an email with the following information:
   - The MyCertificate.arm file, your trial signed digital certificate.
   - A link to download IntermediateCert.arm, the trial intermediate digital certificate.
   - A link to download RootCert.arm, the root digital certificate.
5. Add the trial root digital certificate to the key database.
   The following example adds RootCert.arm to ondemand.kdb:

```
gsk8capicmd_64 -cert -add -db "ondemand.kdb" -pw "myKeyDBpasswd" -label
"trialRootCACert" \
-file RootCert.arm -format ascii
```

6. Add the trial intermediate certificate to the key database.
   The following example adds IntermediateCert.arm to ondemand.kdb:

```
gsk8capicmd_64 -cert -add -db "ondemand.kdb" -pw "myKeyDBpasswd" –label
"trialIntermediateCACert" \
-file IntermediateCert.arm -format ascii
```

7. Receive your signed digital certificate to the key database.
   The following example receives MyCertificate.arm to ondemand.kdb:

```
gsk8capicmd_64 -cert -receive -file MyCertificate.arm -db "ondemand.kdb" –pw
“myKeyDBpasswd” \
–format ascii
```

8. Verify that all the certificates were stored in the key database by using GSKCapiCmd.
   The following example lists the certificates stored in ondemand.kdb:

   ```
   gsk8capicmd_64 -cert -list all -db "ondemand.kdb" -pw "myKeyDBpasswd"
   ```

   GSKCapCmd displays the following result:

   ```
   Certificates found
   * default, - personal, ! trusted
   -! mycert
   ! trialIntermediateCACert
   ! trialRootCACert
   -! myselfsigned
   ```

Saving Content Manager OnDemand passwords into encrypted files for Linux

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

About this task

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

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

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

Procedure

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

1. Create a stash file by running the `arsstash` command. The command prompts you for the password.

   For a description of the syntax of the ARSSTASH command and examples, see *Content Manager OnDemand for Multiplatforms: Administration Guide*.
2. Save the stash file in a directory and limit access to that file through file permissions.

Results

When you configure the Content Manager OnDemand instance, you modify the ARS.INI file and include the SRVR_OD_STASH parameter and specify the directory that you specified.

Installing and configuring Tivoli Storage Manager on Linux™

Tivoli Storage Manager can be used with Content Manager OnDemand object servers to store report data on devices that are supported by Tivoli Storage Manager.

About this task

This section is for reference only. While optical libraries are used in the example, they are not supported on Linux. Devices supported by Tivoli Storage Manager include optical libraries and tape media. The use of Tivoli Storage Manager is optional and is needed only if you want to provide long-term storage for your reports on devices other than the fixed disks attached to the object server. You can also use Tivoli Storage Manager facilities to maintain DB2 archived log files and backup image files.

You will need the *IBM Tivoli Storage Manager for Linux: Quick Start* publication to install and configure Tivoli Storage Manager. HTML and PDF versions of Tivoli Storage Manager publications, including the Quick Start Guide, are available at http://publib.boulder.ibm.com/tividd/td/tdprodlist.html.

Planning for interoperability between Content Manager OnDemand and Tivoli Storage Manager

Content Manager OnDemand supports Tivoli Storage Manager in several different configurations with the library and object server.

Content Manager OnDemand supports Tivoli Storage Manager in the following configurations:

- Standard library/object server plus Tivoli Storage Manager on one workstation. Install the Server, Clients, 64-bit Client API, Device Support Runtime, Server Runtime, and Licenses packages on the workstation.
- Library server only (where Tivoli Storage Manager resides on some other workstation than the library server). Install the 64-bit Clients and 64-bit API packages on the library server workstation.
- Object server plus Tivoli Storage Manager on some other workstation than the library server. Install the Server, 64-bit Clients, 64-bit API, Device Support Runtime, Server Runtime, and Licenses packages on the object server workstation.

Content Manager OnDemand uses the Tivoli Storage Manager API client to store data into the Tivoli Storage Manager server. The Tivoli Storage Manager server is managed and administered independently

of Content Manager OnDemand. The Tivoli Storage Manager administrator must ensure that the following conditions are met:

- All the normal requirements for Tivoli Storage Manager storage are monitored and managed accordingly
- All required Tivoli Storage Manager policies, management classes, storage pools, and volumes are defined accordingly
- All required Tivoli Storage Manager storage pools and volumes are online
- All Tivoli Storage Manager storage pools and volumes have sufficient storage space to satisfy the needs of Content Manager OnDemand
- The Tivoli Storage Manager server is active when OnDemand needs to read from or write to its storage repository

If your Tivoli Storage Manager configuration cannot support Content Manager OnDemand, system requests (that require Tivoli Storage Manager services) will fail. The Tivoli Storage Manager administrator should examine the system to ensure that it will support the storage and retrieval of data by Content Manager OnDemand.

Configuring Tivoli Storage Manager to maintain Content Manager OnDemand data in archive storage

Provides general guidance about how to configure Tivoli Storage Manager to maintain Content Manager OnDemand data in archive storage.

About this task

Tivoli Storage Manager can maintain the reports that you load into Content Manager OnDemand, can maintain migrated index data, and can maintain DB2 archived log files and backup image files.

Before you begin, familiarize yourself with the following sources available in the IBM Knowledge Center:

- For information about configuring and managing server storage, see the *IBM Tivoli Storage Manager for Linux: Administrator's Guide*.
- For detailed information about the Tivoli Storage Manager commands, see the *IBM Tivoli Storage Manager for Linux: Administrator's Reference*. This guide is useful as your primary reference.

If you encounter problems configuring Tivoli Storage Manager, see the Tivoli Storage Manager publications.

Procedure

Complete these tasks to set up Tivoli Storage Manager for Content Manager OnDemand on a Linux workstation. Use the *IBM Tivoli Storage Manager for Linux: Administrator's Guide* and *IBM Tivoli Storage Manager for Linux: Administrator's Reference* for specific instructions on how to complete each task:

1. Define the Tivoli Storage Manager server options.
2. Define the Tivoli Storage Manager client system options.

   When you update the servers file, add the following line to turn off compression: `COMPRESSION OFF`
3. Define the Tivoli Storage Manager client options.
4. Register Tivoli Storage Manager licenses.
5. Register Tivoli Storage Manager administrators.
6. Define other Tivoli Storage Manager server options.
7. Start, halt, and restart the Tivoli Storage Manager server.
8. Increase Tivoli Storage Manager database and recovery log sizes.
9. Define a storage library.
10. Define policy domains.
11. Register client nodes.

You can use the information in “Registering client nodes” on page 67 to supplement the instructions provided by Tivoli Storage Manager.

12. Define archive copy groups.

    You can use the information in “Define the archive copy group” on page 67 to supplement the instructions provided by Tivoli Storage Manager.
13. Prepare storage pool volumes.
14. Optional: Configure Tivoli Storage Manager to maintain DB2 archived log files and backup image files.
15. Define a backup device for the Tivoli Storage Manager database.
16. Back up the Tivoli Storage Manager database and critical files.

Registering client nodes

About this task

A client node links clients and their data with storage volumes and devices. Before Content Manager OnDemand can store data in Tivoli Storage Manager storage, you must register at least one client node. You must register at least one client node in each policy domain that will contain Content Manager OnDemand data. You can use the example that follows as a guide when registering client nodes. The example presents the procedure with a minimum of customization. If you want to do more, refer to the Tivoli Storage Manager documentation. Enter the command at the Tivoli Storage Manager server command line interface.

To register the client node PRI7YR and password `password` and assign the client node to the `OD7YPD` policy domain, and specify that the client node should be able to delete its own archive files from the server, enter:

```
    register node PRI7YR password domain=OD7YRPD archdel=yes contact='your
name'
```

The `archdel=yes` parameter is required for Content Manager OnDemand processing.

Note: When you define a Content Manager OnDemand storage node (by using the Content Manager OnDemand facilities), specify a Tivoli Storage Manager client node and client node password to "link" the Content Manager OnDemand storage node to archive storage.

Define the archive copy group

The archive copy group determines several Tivoli Storage Manager options for the DB2 archived log files, including the number of days that Tivoli Storage Manager maintains the files. The DB2 archived log files must be maintained until they are no longer needed for database or table space recovery. Log files are valid between full, offline backup images of the database. When you create a full, offline backup image of the database, the log files created prior to the backup image can be deleted. For example, if you create a full, offline backup image of the database every thirty days, then you must keep log files for at least thirty days. If you do not create full, online backup images of the database, you should maintain the log files indefinitely.

The following example shows how to define an archive copy group. The archive copy group identifies the policy domain, policy set, and management class. The archive copy group also identifies the storage pool where Tivoli Storage Manager maintains the DB2 archived log files and the length of time that Tivoli Storage Manager maintains them. In the example, Tivoli Storage Manager maintains each log file stored in the storage pool for 366 days:

```
    define copygroup 1YRPD 1YRPS 1YRMG standard -
                     type=archive dest=ODSTGP2 retver=366
```

Configuring Tivoli Storage Manager to maintain DB2 files

About this task

You can use Tivoli Storage Manager to maintain DB2 archived log files and backup image files. This capability means that you do not have to manually maintain these files on disk. The tasks in this section are optional, and are only recommended for customers who need to use Tivoli Storage Manager facilities to backup and restore the Content Manager OnDemand database in DB2. For more information about using Tivoli Storage Manager to manage DB2 files, see *IBM DB2 Universal Database: Data Recovery and High Availability Guide and Reference*, SC09-4831.

You need to do the following tasks to configure Tivoli Storage Manager to maintain DB2 files:

- Define server options
- Define client options
- Define storage objects
- Register the client node
- Set the client node password
- Determine space requirements
- Review backup considerations
- Protect data with the data retention protection (DRP) protocol

Protecting data with the data retention protection (DRP) protocol

To avoid the accidental erasure or overwriting of critical data, Content Manager OnDemand supports the Tivoli Storage Manager APIs related to data retention.

Data retention protection (DRP)
Prohibits the explicit deletion of documents until their specified retention criterion is met. Although documents can no longer be explicitly deleted, they can still expire.

Important: DRP is permanent. After it is turned on, it cannot be turned off.

Event-based retention policy
Retention based on an external event other than the storage of data. For Content Manager OnDemand, the retention event is the call to delete the data. A load, unload, application group delete, or expiration of data triggers the retention event.

Restriction: Content Manager OnDemand does not support *deletion hold*, which is a feature that prevents stored data from being deleted until the hold is released.

If you decide to use these policies in Tivoli Storage Manager, then the following scenarios result:

Table 6: Scenarios of using data retention protection		
	Creation-based object expiration policy	**Event-based retention object expiration policy**
Data retention protection off	Content Manager OnDemand issues a delete object command through the Tivoli Storage Manager API. Objects are deleted during the next inventory expiration. If a Content Manager OnDemand application group is being deleted, a `delete filespace` command is issued, and the object file space is immediately deleted with the file space.	Content Manager OnDemand issues an event trigger command through the Tivoli Storage Manager API. The status of the objects that are affected are changed from PENDING to STARTED, and the objects are expired by Tivoli Storage Manager based on their retention parameters. If the retention parameters are set to `NOLIMIT`, then the objects never expire. If a Content Manager OnDemand application group is being deleted, a `delete filespace` command is issued instead, and the objects are immediately deleted along with the file space.
Data retention protection on	Content Manager OnDemand issues no commands to Tivoli Storage Manager. The objects are effectively orphaned by Content Manager OnDemand and are expired by Tivoli Storage Manager based on their retention parameters. If the retention parameters are set to `NOLIMIT`, then the objects never expire.	Content Manager OnDemand issues an event trigger command through the Tivoli Storage Manager API. The event status of the objects that are affected are changed from PENDING to STARTED and the objects will be expired by Tivoli Storage Manager based on their retention parameters. If the retention parameters are set to `NOLIMIT`, then the objects never expire. If a Content Manager OnDemand application group is being deleted, then a `delete filespace` cannot be used with DRP enabled, therefore, the operation is treated the same as if a delete were indicated. The status of all the affected objects is changed from PENDING to STARTED, and they will be expired by Tivoli Storage Manager based on their retention parameters. Because this leaves the file space entries in TSM, you must manually delete these entries when the file space is empty (even with DRP enabled).

Recommendations:

- Set up the application groups to expire by load.
- Define the Tivoli Storage Manager archive copy groups to be event-based, and retain data for 0 days.

- Run the Tivoli Storage Manager inventory expiration regularly to ensure that expired data is removed.

Additionally, Content Manager OnDemand supports the following devices:

IBM DR450 and DR550
Disk-based system that contains a Tivoli Storage Manager that runs DRP.

EMC Centera
Disk-based system that is treated as a device by Tivoli Storage Manager. Tivoli Storage Manager must run DRP.

Installing the Content Manager OnDemand software on Linux

You must install a copy of the Content Manager OnDemand software on each workstation or node that is part of the Content Manager OnDemand system.

Before you begin

1. You need approximately 600 MB of free space in the `/opt` file system to install the software.
2. By default, the installation is carried out in the GUI mode, therefore, the X Windows support is required for the GUI install.

Procedure

Complete the following steps to install the Content Manager OnDemand product files on a Linux workstation:

1. Log in as the root user.
2. Go to the directory where `odlinux.bin` is located.
3. Enter this command: `./odlinux.bin`
4. Read the Welcome screen and then click **Next**. The License Agreement window appears.
5. Select **I accept the terms in the license agreement** to accept the license agreement. Click **Next**.
6. Accept the default directory name or, if you prefer a different directory name, type in the directory name. Click **Next**.

   If you have a version of Content Manager OnDemand older than Version 8.5 installed, the installation program removes the previous version before installing the new version.
7. When the process completes, this question **Would you like to display the product ReadMe file?** appears. The location of the product readme file is displayed also. On Linux, the readme file is located in the `/opt/ibm/ondemand/V10.1` directory.
8. If you want to view the readme file now, click **Yes**. After you finish reading the readme file, click **Next**.
9. Read the information in the window, and click **Next**.
10. Click **Finish**.
11. After installing the software, apply the latest service update for Content Manager OnDemand. You can obtain the latest service update from IBM service at http://www.ibm.com/eserver/support/fixes/.

Results

Optionally, the installation can be performed in the character based console mode. To install the Content Manager OnDemand for Linux server in the console mode, enter the following command from the directory which contains the installer:

```
./odlinux.bin -i console
```

and follow the instructions on the installation panels.

Installing optional Content Manager OnDemand software on Linux

Other software is available for installation in addition to Content Manager OnDemand software.

About this task

The command to install the Content Manager OnDemand PDF Indexing feature is:

```
./odpdflinux.bin
```

or

```
./odpdflinux.bin -i console
```

The command to install the IBM Content Manager OnDemand Distribution Facility feature is:

```
./ododflinux.bin
```

or

```
./ododflinux.bin -i console
```

The command to install the Content Manager OnDemand Full Text Search server feature is:

```
./odftslinux.bin
```

or

```
./odftslinux.bin -i console
```

To install the Content Manager OnDemand Enhanced Retention Management feature, see *Enhanced Retention Management Guide*.

Configuring instances on Linux™

A Content Manager OnDemand instance is a logical server environment made up of a database, a library server, and one or more object servers.

An instance is defined in the ARS.INI file by naming the instance, identifying the name of the database used by the instance, and identifying the library server on which the database will be maintained. When you configure an object server, you identify its library server in the ARS.CFG file on the object server. An instance has its own table space file systems for the database and cache file systems. The table space file systems are defined in the ARS.DBFS file on the library server. The cache file systems are defined in the ARS.CACHE file on each object server. All of the servers that belong to an instance run in a single code page and on the same TCP/IP port number.

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

- To have separate test and production environments
- To have databases using different code pages

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

Each additional instance requires additional system resources, such as virtual storage and disk space, and more administration. If you plan to run more than one instance on the same workstation:

- The ARS.INI file must contain one section for each instance. Each section identifies the ARS.CFG file, ARS.DBFS file, and ARS.CACHE file used by the instance.
- You must create a unique copy of the ARS.CFG file for each instance.

- You should maintain separate table space file systems and cache storage file systems for each instance, as in a ARS.DBFS file and ARS.CACHE file for each instance.
- Each instance must run on its own unique TCP/IP port number. The port for each instance is configured in the ARS.INI file.

Instances in the ARS.INI file

The ARS.INI file contains information about Content Manager OnDemand instances.

When you install the Content Manager OnDemand software, the ARS.INI file contains information about a default instance named archive. Most customers will use the default instance for their first or only instance of Content Manager OnDemand.

The information in the ARS.INI file is organized in sections with a header line that identifies each section. The header line can be identified by the brackets [ ] that delimit the beginning and end of the line.

The first section in the ARS.INI file contains information about the default instance. The following shows the default instance as provided by IBM:

```
[@SRV@_ARCHIVE]
HOST=platte
PROTOCOL=2
PORT=0
SRVR_INSTANCE=archive
SRVR_INSTANCE_OWNER=root
SRVR_OD_CFG=/opt/ibm/ondemand/V10.1/config/ars.cfg
SRVR_DB_CFG=/opt/ibm/ondemand/V10.1/config/ars.dbfs
SRVR_SM_CFG=/opt/ibm/ondemand/V10.1/config/ars.cache
```

The HOST parameter identifies the host name alias, IP address, or fully qualified host name of the workstation on which the library server is running. The PROTOCOL parameter identifies the communications protocol used by the instance. The PORT parameter identifies the TCP/IP port number that the instance monitors for client requests. The stanza name ([@SRV@_ARCHIVE]) identifies the name of the Content Manager OnDemand instance. The SRVR_INSTANCE parameter identifies the name of the Content Manager OnDemand database. The SRVR_INSTANCE_OWNER parameter identifies the userid of the Content Manager OnDemand instance owner. The SRVR_OD_CFG parameter identifies the ARS.CFG file used by the instance. The SRVR_DB_CFG parameter identifies the ARS.DBFS file used by the instance. The SRVR_SM_CFG parameter identifies the ARS.CACHE file used by the instance.

When adding an instance to the ARS.INI file, remember that each instance must specify a unique instance name. For example, to add an instance for testing new applications, you might add an instance named test. When you work with more than one instance, you must identify the instance name when you run Content Manager OnDemand programs (such as ARSDB, ARSLOAD, and ARSSOCKD) and database commands (such as connecting to the database). The following shows an example of a second instance in the ARS.INI file:

```
[@SRV@_TEST]
HOST=rhone
PROTOCOL=2
PORT=1444
SRVR_INSTANCE=test
SRVR_INSTANCE_OWNER=root
SRVR_OD_CFG=/opt/ibm/ondemand/V10.1/config/ars.test.cfg
SRVR_DB_CFG=/opt/ibm/ondemand/V10.1/config/ars.test.dbfs
SRVR_SM_CFG=/opt/ibm/ondemand/V10.1/config/ars.test.cache
```

The header line for the definition of the instance is TEST. The HOST statement identifies the host name alias of the library server. The instance communicates over TCP/IP port number 1444. The name of the Content Manager OnDemand database is `test`. The name of the Content Manager OnDemand instance is test. The userid of the Content Manager OnDemand instance owner is root. The instance identifies its server configuration file (ARS.TEST.CFG), table space file systems file (ARS.TEST.DBFS), and cache file systems file (ARS.TEST.CACHE).

Specifying instances in the ARS.INI file

The ARS.INI file contains information about a default instance named archive. Most customers will use the default instance for their first or only instance of Content Manager OnDemand.

Procedure

To specify the instance in the ARS.INI file, follow these steps:

1. Log in to the server as the root user.
2. Change to the /opt/ibm/ondemand/V10.1/config directory.
3. Make a backup copy of the file provided by IBM.
4. Edit the ARS.INI file with a standard text editor such as vi.
5. Most customers will use the default instance named ARCHIVE for their first or only instance of Content Manager OnDemand.
6. **Note for distributed library/object servers:** Configure one copy of the ARS.INI file on each workstation that is part of the system. Verify that the information specified in the ARS.INI file is consistent on all workstations that are part of the instance. In addition:
   a) Ensure that the port number of the object server matches the port number of the library server.
   b) Verify that the HOST parameter on the object server must specify the host name alias, IP address, or fully qualified host name of the library server.
7. Save the file and exit the text editor.
8. You should control access to the ARS.INI file by changing the file permissions so that only the Content Manager OnDemand instance owner has read or write access to the file.

Verifying the default instance

Most customers will use the default instance named ARCHIVE for their first or only instance of Content Manager OnDemand.

Verify the following parameters and values:

- The header line contains a string that identifies the name of the instance. Unless you specify otherwise, the first or only instance is named ARCHIVE.
- The HOST parameter identifies the host name alias, IP address, or fully qualified host name of the library server.
- The PROTOCOL parameter identifies the communications protocol used by the instance. The number 2 identifies TCP/IP, and is the only valid value.
- The PORT parameter identifies the TCP/IP port number that the instance monitors for client requests (the number 0 means that the instance monitors port number 1445). If you use a port number other than 1445 on the library server, enter that number instead of 0 (zero). **For customers running more than one instance:** Each instance that runs on the same workstation must specify a different port number. If you configure a separate object server, ensure that the port number of the object server matches the port number of the library server.
- The stanza name ([@SRV@_ARCHIVE]) identifies the name of the Content Manager OnDemand instance. This value should match the name of the Content Manager OnDemand database (see “Installing DB2®” on page 58). The instance name can be from one to eight characters in length, and can include the A through Z and 0 through 9 characters.
- The SRVR_INSTANCE_OWNER parameter identifies the user ID of the Content Manager OnDemand instance owner. This is the user ID that is permitted to run the Content Manager OnDemand server programs, such as ARSSOCKD, ARSLOAD, and ARSMAINT.
- The SRVR_OD_CFG parameter identifies the ARS.CFG configuration file used by the instance. See “Specifying the ARS.CFG file for the instance” on page 74.
- The SRVR_DB_CFG parameter identifies the ARS.DBFS table space file system file used by the instance. See “Specifying the ARS.DBFS file for the instance” on page 79.

- The SRVR_SM_CFG parameter identifies the ARS.CACHE cache file system file used by the instance. See "Creating the ARS.CACHE file for the instance" on page 80.
- The SRVR_OD_STASH parameter identifies the location of the stash file used by the instance and Content Manager OnDemand programs. See "Saving Content Manager OnDemand passwords into encrypted files for Linux" on page 64.

Specifying the ARS.CFG file for the instance

About this task

The ARS.CFG file contains information about the instance, such as identifying the object servers that belong to the instance, the language settings for the instance, and information that is used by database, storage, and print manager programs.

Before you create the Content Manager OnDemand database, start Content Manager OnDemand, use archive storage, use the server print function, migrate tables to table spaces, or import tables from archive storage to the database, you should review the parameters in the ARS.CFG file. The values that IBM provides are sufficient for most customers. However, you might need to change some of the values for your environment.

ARS_DB_ENGINE parameter

The database manager product that you installed on the library server. The default value is DB2. The ARS_DB_ENGINE parameter is ignored on object servers.

ARS_DB_IMPORT parameter

The method that Content Manager OnDemand uses to migrate index data to table spaces and import tables from archive storage to the database. The default value is 0 (zero). The ARS_DB_IMPORT parameter is ignored on object servers.

If you are configuring a library server, then you must set the ARS_DB_IMPORT parameter to one of the following values:

0

Content Manager OnDemand uses the EXPORT and IMPORT commands to migrate table data. This method requires disk space to hold log records generated when exporting existing table data and importing data to the new table space. This is the default migration method.

1

Content Manager OnDemand uses the EXPORT and LOAD commands to migrate table data. This method requires disk space to hold log records generated when exporting existing table data. The LOAD command generates a backup image of the new table space. The image file is stored in Tivoli Storage Manager-managed storage. This is the recommended migration method.

Tip: Before you can use Tivoli Storage Manager to manage DB2 backup image files, you must install and configure Tivoli Storage Manager. See "Installing and configuring Tivoli Storage Manager on Linux™" on page 65 for details.

2

Content Manager OnDemand uses the EXPORT and LOAD commands to migrate the table data. This method requires disk space to hold log records generated when exporting existing table data. The LOAD command generates a backup image of the new table space. The image file is stored in the file system identified by the ARS_TMP parameter (see "ARS_TMP parameter" on page 78).

ARS_DB_PARTITION parameter

Determines whether you can partition the database across nodes or systems. By default, you cannot partition the database. If the database manager product that you are using with Content Manager

OnDemand supports partitioning, then you can specify that you want to partition the database by changing the value of this parameter to 1 (one). Content Manager OnDemand supports partitioning with DB2 Universal Database Extended Enterprise Edition only. To store application group index data in partitions, your application groups must specify a partition field. The ARS_DB_PARTITION parameter is ignored on object servers.

ARS_DB_TABLESPACE parameter

The name of the table space for the Content Manager OnDemand system tables. The value of this parameter must match an existing table space name in the database. You must have created the table space in DB2.

ARS_DB_TABLESPACE_USEREXIT parameter

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

The following statement must exist in the ARS.CFG file that is associated with the instance so that the ARSUTBL DLL can be invoked:

```
ARS_DB_TABLESPACE_USEREXIT=absolute path name
```

For the sample ARSUTBL, you would specify the following statement in the ARS.CFG file:

```
ARS_DB_TABLESPACE_USEREXIT=/opt/ibm/ondemand/V10.1/bin/exits/arsutbl
```

"Table space creation user exit" on page 166 provides information about the exit point that gets invoked when OnDemand creates table spaces, tables, and indexes for the Content Manager OnDemand data tables.

ARS_DB2_DATABASE_PATH parameter

If you are configuring the library server and you set the ARS_DB_ENGINE parameter (see "ARS_DB_ENGINE parameter" on page 74) to DB2 (the default), the base file system in which the Content Manager OnDemand database will reside. You must make sure that the specified location contains enough space to hold the system tables, the USERSPACE1 table space, and any application group tables that are not stored in their own table spaces. The *IBM Content Manager OnDemand for Multiplatforms: Introduction and Planning Guide* contains information to help you estimate the amount of space required to hold the database. The default value is `/arsdb`. The ARS_DB2_DATABASE_PATH parameter is ignored on object servers.

ARS_DB2_LOG_NUMBER parameter

If you are configuring the library server and you set the ARS_DB_ENGINE parameter (see "ARS_DB_ENGINE parameter" on page 74) to DB2 (the default), the number of primary log files. The default value is 40. The ARS_DB2_LOG_NUMBER parameter is ignored on object servers.

The values of the ARS_DB2_LOGFILE_SIZE and ARS_DB2_LOG_NUMBER parameters determine the total amount of space available for DB2 to log changes to the database. The values that you specify must support the largest single report that you plan to load (or unload). DB2 will fail if there is not enough log file space available to hold the changes to the database. The default values allocate 160 MB of space. See the *IBM Content Manager OnDemand for Multiplatforms: Introduction and Planning Guide* for information about estimating the amount of storage space required to hold the DB2 log files.

ARS_DB2_LOGFILE_SIZE parameter

If you are configuring the library server and you set the ARS_DB_ENGINE parameter (see "ARS_DB_ENGINE parameter" on page 74) to DB2 (the default), the size of a log file, in 4 KB blocks. The default value is 1000. The ARS_DB2_LOGFILE_SIZE parameter is ignored on object servers.

The values of the ARS_DB2_LOGFILE_SIZE and ARS_DB2_LOG_NUMBER parameters determine the total amount of space available for DB2 to log changes to the database. The values that you specify must support the largest single report that you plan to load (or unload). DB2 will fail if there is not enough log file space available to hold the changes to the database. The default values allocate 160 MB of space. See the *IBM Content Manager OnDemand for Multiplatforms: Introduction and Planning Guide* for information about estimating the amount of storage space required to hold the DB2 log files.

ARS_DB2_PRIMARY_LOGPATH parameter

If you are configuring the library server and you set the ARS_DB_ENGINE parameter (see "ARS_DB_ENGINE parameter" on page 74) to DB2 (the default), the location that will hold the active archived log files. The *IBM Content Manager OnDemand for Multiplatforms: Introduction and Planning Guide* can help you estimate the amount of space required to hold the active archived log files. The default value is `/arsdb_primarylog`. The ARS_DB2_PRIMARY_LOGPATH parameter is ignored on object servers.

ARS_LDAP_ALLOW_ANONYMOUS parameter

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

ARS_LDAP_BASE_DN parameter

Specifies the base distinguished name to use. This parameter is required for LDAP authentication.

Example 1: `ARS_LDAP_BASE_DN=ou=mycity,o=xyzcompany`

Example 2: `ARS_LDAP_BASE_DN=dc=ondemand,dc=xyzcompany`

ARS_LDAP_BIND_ATTRIBUTE parameter

Specifies the attribute being bound and is the attribute name to be searched on the LDAP server. This parameter is required for LDAP authentication.

Example: `ARS_LDAP_BIND_ATTRIBUTE=mail`

ARS_LDAP_BIND_MESSAGES_FILE parameter

Specifies the location of a file containing the LDAP message strings the Content Manager OnDemand server looks for during login. This is used for issuing messages when the user's password is about to expire, or their LDAP account is locked. ARS_LDAP_BIND_MESSAGES_FILE is used in conjunction with the `ARSLDAP.INI` file to implement this functionality.

ARS_LDAP_IGN_USERIDS

This parameter specifies the user IDs that OnDemand ignores when you enable LDAP for authentication. If the parameter does not exist or you do not specify a value, OnDemand defaults to ADMIN. You can specify up to 10 user IDs, delimited by a comma. If you specify a list of user IDs and you want to include ADMIN, you must specify it on the list.

ARS_LDAP_MAPPED_ATTRIBUTE parameter

Specifies the attribute being returned to Content Manager OnDemand as a user ID. This is the attribute name to be returned from the LDAP server once the bind attribute name is found. It can be the same as the bind attribute or different. This parameter is required for LDAP authentication.

Example: `ARS_LDAP_MAPPED_ATTRIBUTE=sAMAccountName`

ARS_LDAP_PORT parameter

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

ARS_LDAP_SERVER parameter

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

ARS_LOCAL_SRVR parameter

The name of the object server. The ARS_LOCAL_SRVR parameter is ignored on library servers. However, if you are configuring a library server, you must either omit this parameter from the ARS.CFG file or set this parameter to a blank value, that is: ARS_LOCAL_SRVR=

If you are configuring an object server, set this parameter to the TCP/IP host name alias, fully qualified host name, or IP address of the object server. If the object server is running on a node of a multi-processor workstation, then set this parameter to the external IP address of the node on which you installed the object server.

When you add a Content Manager OnDemand storage node to an object server, you must use the value of the ARS_LOCAL_SRVR parameter to name the storage node.

ARS_MESSAGE_OF_THE_DAY parameter

Use to show the message of the day. Set to the full path name of a file that contains the message that you want to show. For example: ARS_MESSAGE_OF_THE_DAY=/opt/ibm/ondemand/V10.1/tmp/message.txt

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

ARS_NUM_DBSRVR parameter

Determines the number of processes that Content Manager OnDemand starts on the library server to handle connections to the database. The ARS_NUM_DBSRVR parameter is ignored on object servers.

In addition to database connections by Content Manager OnDemand client programs, the value that you specify must support the number of active Content Manager OnDemand commands and daemons such as ARSLOAD, ARSDOC, ARSDB, ARSMAINT, and ARSADMIN.

Each connection to the database requires a database agent. Content Manager OnDemand can start a database agent for each connection. However, each agent requires its own private memory and some portion of application shared memory. You can use the ARS_NUM_DBSRVR parameter to optimize the way that Content Manager OnDemand handles the database load. For example, you can define ARS_NUM_DBSRVR so that Content Manager OnDemand starts a fixed number of database agents, regardless of the number of concurrent database requests. While this might appear restrictive, database requests typically process very quickly. For example, ten database agents can handle a heavy database request load, while balancing the impact on system resources.

You should specify a value for the ARS_NUM_DBSRVR parameter that supports the peak number of concurrent database connections that you expect the library server to handle. A low value limits access to the database during periods of high database activity. A high value requires more system resources during periods of high database activity. The value that you choose also depends on the characteristics of the queries. For example, general queries typically keep a connection open longer than a more specific query.

ARS_ORACLE_HOME parameter
Use to specify the base installation directory for Oracle.

The default value is:

```
ARS_ORACLE_HOME=/oracle
```

Replace the string /oracle with the name of the directory in which Oracle was installed.

ARS_PRINT_PATH parameter

The location where the Content Manager OnDemand server print function temporarily stores print data. You must make sure that there is enough space in the specified location to hold the print files for the maximum number of concurrent print requests that the server will handle. The default value is /tmp. The ARS_PRINT_PATH parameter is ignored on object servers.

You should dedicate a file system to hold the print files. The file system contain at least 500 MB of free space at all times. If your storage configuration permits, you should allocate 1 GB or more of free space to the specified file system.

The permissions for the file system must be `drwxrwxrwt`. You can use the CHMOD command to set the permissions. For example, the command `chmod 1777 /tmp` sets the permissions for the `/tmp` file system.

ARS_SRVR parameter

The name of the library server. The ARS_SRVR parameter is ignored on library servers. However, if you are configuring a library server, you must either omit this parameter from the ARS.CFG file or set this parameter to a blank value, that is: ARS_SRVR=.

If you are configuring an object server, set the ARS_SRVR parameter to the TCP/IP host name alias, fully qualified host name, or IP address of the library server. If the library server is running on a node of a multi-processor workstation, then set this parameter to the external IP address of the node on which you installed the library server.

ARS_STORAGE_MANAGER parameter

Determines whether the server program is linked to a cache-only storage manager or an archive storage manager. You must specify this parameter on library and object servers.

You can specify one of the following values:

CACHE_ONLY

Link the server program to a cache-only storage manager.

TSM

Link the server program to an archive storage manager. This is the default value in the ARS.CFG file that is provided by IBM.

Before Content Manager OnDemand can work with an archive storage manager to maintain data, you must install and configure the archive storage manager software.

ADSM

Deprecated. This option has been replaced by TSM. ADSM is still supported for existing customers.

ARS_SUPPORT_CFSOD parameter

If you plan to use Content Federation Services for OnDemand, you must set this parameter equal to 1.

ARS_SUPPORT_HOLD parameter

To use enhanced retention management, you must set this parameter to 1. See the document at http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP10137 8 for additional details about the set up and use of enhanced retention management.

ARS_TMP parameter

The location where Content Manager OnDemand programs temporarily store data. You must allocate sufficient free space in the specified file system to support tasks such as migrating and importing index data. The default value is: `/tmp`. You must specify the ARS_TMP parameter on the library server and on all object servers.

You should dedicate a file system to temporary storage. The file system should contain at least 500 MB of free space at all times. If your storage configuration permits, you should allocate 1 GB or more of free space to the specified file system.

The permissions for the file system must be `drwxrwxrwt`. You can use the CHMOD command to set the permissions. For example, the command `chmod 1777 /tmp` sets the permissions for the `/tmp` file system.

DB_ENGINE parameter

Deprecated. This parameter has been replaced by ARS_DB_ENGINE. However, the DB_ENGINE parameter is still supported for existing customers.

DB2INSTANCE parameter

If you are configuring the library server and you set the ARS_DB_ENGINE parameter (see "ARS_DB_ENGINE parameter" on page 74) to DB2 (the default), the name of the database instance owner that you created when you installed DB2 (see "Installing DB2®" on page 58). The default value is `archive`. The DB2INSTANCE parameter is ignored on object servers.

DSMI_CONFIG parameter

If you plan to use Tivoli Storage Manager, the full path name of the Tivoli Storage Manager API options file. For example: `/usr/tivoli/tsm/client/api/bin64/dsm.opt`.

You must set the DSMI_CONFIG parameter on each object server that uses Tivoli Storage Manager to maintain Content Manager OnDemand data.

DSMI_DIR

If you plan to use Tivoli Storage Manager, the directory that contains the Tivoli Storage Manager API files. For example: `/usr/tivoli/tsm/client/api/bin64`.

You must set the DSMI_DIR parameter on each object server that uses Tivoli Storage Manager to maintain Content Manager OnDemand data.

DSMI_LOG parameter

If you plan to use Tivoli Storage Manager, the directory in which Tivoli Storage Manager stores the Tivoli Storage Manager API error log. The default value is `/tmp`. You must set the DSMI_LOG parameter on each object server that uses Tivoli Storage Manager to maintain Content Manager OnDemand data.

Specifying the ARS.DBFS file for the instance

About this task

Creating the ARS.DBFS file for the instance

About this task

To create (or edit) the ARS.DBFS file for the instance:

Procedure

1. Log in to the server as the `root` user.
2. Change to the `/opt/ibm/ondemand/V10.1/config` directory.
3. Create (or edit) the ARS.DBFS file using a standard text editor such as vi.
4. Add one line for each file system that Content Manager OnDemand can use for table spaces.
5. Save the file and exit the editor.

   A table space file system must be owned by the database instance owner and group. The suggested defaults are `archive` (instance owner) and `db2iadm1` (group).
6. You specified the instance owner and group when you installed the database manager product (see "Installing the database manager on Linux™" on page 58). Make sure that the user and group file permissions are set correctly.
   For example:

```
drwxrws---  3 archive db2iadm1  512 May 17 12:58 /arsdb/db1/SMS
```

You can use the CHOWN command to set the ownership permissions. For example, the following command changes the owner of all file systems in the /arsdb tree to the archive user and the db2iadm1 group:

```
chown -R archive:db2iadm1 /arsdb*
```

7. You can use the CHMOD command to set the file permissions.
   For example, the following commands set the correct permissions for the /arsdb/db1/SMS filesystem:

```
chmod 2770 /arsdb/db1/SMS
chmod g+s /arsdb/db1/SMS
```

Creating the ARS.CACHE file for the instance

About this task

The ARS.CACHE file lists the file systems on the object server that can be used by Content Manager OnDemand for cache storage.

If there are multiple file systems in the ARS.CACHE file, Content Manager OnDemand uses the file system with the greatest amount of space free to store the objects.

The following example shows an ARS.CACHE file that defines five cache storage file systems:

```
/arscache/cache1
/arscache/cache2
/arscache/cache3
/arscache/cache4
/arscache/cache5
```

Note: For a distributed library / object server system, configure one copy of the ARS.CACHE file on each server that is part of the Content Manager OnDemand system.

Procedure

To create the ARS.CACHE file for the instance:

1. Log in to the server as the root user.
2. Change to the /opt/ibm/ondemand/V10.1/config directory.
3. Create (or edit) ARS.CACHE file using a standard text editor such as vi.
4. Insert one line in the file for each file system on the server that Content Manager OnDemand can use for cache storage.

   Important: The first entry in the ARS.CACHE file identifies the base cache storage file system. Content Manager OnDemand stores control information in the base cache storage file system. After you define the base cache storage file system to Content Manager OnDemand, you cannot add or remove it from Content Manager OnDemand. It must remain as the first entry.
5. Save the file and exit the editor.
6. Cache file systems must be owned by the Content Manager OnDemand instance owner and the system group. Make sure that only the user file permissions are set, not the group or other file permissions.
   For example:

```
drwx------  3 root root  512 Sep 22 13:08 /arscache/cache1
```

7. Use the CHOWN command to set the ownership permissions. The following example shows how to change the user ownership of all file systems in the /arscache tree: chown -R root:root /arscache*
8. Use the CHMOD command to set the file permissions.

For example, the following commands set the correct permissions for the /arscache/cache1 file system:

```
chmod 700 /arscache/cache1
chmod g-s /arscache/cache1
```

Results

Content Manager OnDemand cache storage files and subdirectories should have the following permissions:

```
drwx------ for every subdirectory (700)
-r-------- for every object that has been migrated to archive storage (400)
-rw------- for every object that has not yet been migrated (600)
-rwxrwxrwx for every symbolic link under the retr and migr directories (777)
```

Specifying the ARSLDAP.INI file

The ARS_LDAP_BIND_MESSAGES_FILE parameter enables Content Manager OnDemand to customize message text returned from an LDAP server that is used to alert users that their LDAP password is about to expire or their LDAP account is locked.

The messages displayed to users are contained in the file referenced by this parameter. To enable this user-configurable message functionality, create a file with the appropriate message strings, and set ARS_LDAP_BIND_MESSAGES_FILE to the full path of the file. The ARSLDAP.INI file is provided with example message strings that can be used by the ARS_LDAP_BIND_MESSAGES_FILE parameter.

The ARSLDAP.INI file contains the following three sections:

```
[BIND_MESSAGES]
PASSWORD_EXPIRED="/opt/ibm/ondemand/V10.1/config/password_expired.txt"
ACCOUNT_LOCKED="/opt/ibm/ondemand/V10.1/config/account_locked.txt"

[PASSWORD_EXPIRED]
TDS6="Password has expired"
AD="data 532"
UDEF1=
UDEF2=
UDEF3=

[ACCOUNT_LOCKED]
TDS6="Account is locked"
AD="data 775"
UDEF1=
UDEF2=
UDEF3=
```

The BIND_MESSAGES section specifies the path to the files containing the user-configurable message text that is displayed to users when their LDAP password is about to expire, or their LDAP account is locked. Generic files are supplied, and should be customized to reflect your actual Content Manager OnDemand environment.

An example message that would be displayed to a user:

```
Your LDAP password has expired and needs to be changed.
Log into <company intranet> for password setting instructions.
```

The entries in the PASSWORD_EXPIRED and ACCOUNT_LOCKED sections are for Tivoli Directory Server Version 6.x and Microsoft Active Directory (AD). These sections also contain three user-defined entries (UDEFx), allowing you to enter your own pattern strings for LDAP servers that are not directly supported.

The LDAP server may return additional information when the user's bind operation fails. When an error is returned from the LDAP server, Content Manager OnDemand reads the file referenced by the ARS_LDAP_BIND_MESSAGES_FILE parameter and searches under the two stanzas, [PASSWORD_EXPIRED] and [ACCOUNT_LOCKED], for user-defined text that matches the LDAP server error. If a match is found, Content Manager OnDemand will display the text found in the files defined under the [BIND_MESSAGES] stanza.

If the ARS_LDAP_BIND_MESSAGES_FILE parameter is not defined, has no file referenced, or the PASSWORD_EXPIRED or ACCOUNT_LOCKED files do not exist, the user will receive a default 'The server failed while attempting to logon' message.

Note: Currently only two error conditions can be handled: PASSWORD_EXPIRED and ACCOUNT_LOCKED. The section titles for these two conditions cannot be changed, but you can change the pattern strings and message text presented to the user to define any two error conditions.

Creating an instance of Content Manager OnDemand on Linux™

The ARSDB program initializes the base system tables that are required by Content Manager OnDemand. You initialize other system tables by running the ARSSYSCR program on the library server. The ARSSYSCR program initializes the system tables that are required to support the system log, system migration, and other Content Manager OnDemand functions.

Prerequisites

Ensure that DB2 or Oracle have been installed and the database has been created.

Before you create the instance, you must have completed the following tasks:

1. Installed and configured the database software (DB2 or Oracle), and created a database instance (DB2) or database (Oracle) for Content Manager OnDemand.
2. Installed and configured the Content Manager OnDemand software, including the following files:
   - `ARS.INI`
   - `ARS.CFG`
   - `ARS.DBFS`
   - `ARS.CACHE`

Creating an instance of Content Manager OnDemand on Linux

Procedure

To create the instance, follow these steps:

1. Specify permissions for the database directories.
2. Create the instance by running the ARSDB program.
3. Initialize the system logging facility by running the ARSSYSCR program.
4. (Optional) Initialize the system load logging facility by running the ARSSYSCR program.
5. (Optional) Initialize the system migration facility by running the ARSSYSCR program.

Specifying permissions for the database directories

About this task

The group that the instance owner belongs to must have write access to the database directory names that are specified in the ARS.CFG file (the ARS_DB2_DATABASE_PATH, ARS_DB2_PRIMARY_LOGPATH, and ARS_DB2_ARCHIVE_LOGPATH parameters). You created the instance owner and group when you installed the database manager (see “Installing the database manager on Linux™” on page 58). See “Specifying the ARS.CFG file for the instance” on page 74 for help with configuring the ARS.CFG file.

Changing owners of directories

You can change the owners of database directories with the CHOWN command.

Procedure

To change the owner of the directories:

1. Log in to the server as the `root` user.
2. Use the CHOWN command to change directory ownership.
   For example, to change the owner and group of all file systems in the `/arsdb` tree to the `archive` owner and the `db2iadm1` group, enter the following command: `chown -R archive:db2iadm1 /arsdb*`

   Run the CHOWN command once to change the ownership of each of the database directories that are specified in the `ARS.CFG` file (the ARS_DB2_DATABASE_PATH, ARS_DB2_PRIMARY_LOGPATH, and ARS_DB2_ARCHIVE_LOGPATH parameters).

To create the database instance

You should use the ARSDB program to create the instance.

The ARSDB program completes the following tasks to create the instance:

- Updates the database configuration
- Verifies the directories for the primary and archived log files
- Creates a link to the database user exit program

  If the database user exit program encounters errors when copying files, it creates the `db2uexit.err` file in the temporary data directory. If this file exists, it usually means that you did not set the correct permissions for the log file directories or there is not enough free space to hold the archived log files. See your operating system documentation for information about increasing the size of a file system.
- Creates a backup of the database
- Builds the Content Manager OnDemand system tables and indexes

  The ARSDB program can build the Content Manager OnDemand system tables and indexes into a default table space or user-defined table spaces. If you want to use the default table space, continue with the instructions in this topic. If you want to use user-defined tables spaces, follow the instructions in Chapter 13, “Creating Content Manager OnDemand system tables into user-defined table spaces,” on page 177 before continuing with the instructions in this topic.
- Binds the database to Content Manager OnDemand

The ARSDB program creates the database using standard SQL commands. See the documentation provided with the database manager product for information about the SQL commands issued by the ARSDB program and messages printed at the console.

Creating a database instance

The ARSDB program creates the instance.

About this task

Oracle users: You must create the database by using the Oracle utilities before you create the Content Manager OnDemand instance.

Procedure

To create an instance of Content Manager OnDemand:

1. Log in to the server as the `root` user.

Option

DB2 Type the following command at the prompt:

```
/opt/ibm/ondemand/V10.1/bin/arsdb -I archive -gcv
```

Where `archive` is the name of the Content Manager OnDemand instance.

Oracle Type the following command at the prompt:

```
/opt/ibm/ondemand/V10.1/bin/arsdb -I archive -rtv
```

Where `archive` is the name of the Content Manager OnDemand instance.

2. Press the Enter key.
   The ARSDB program prompts you before creating a link to the database user exit program:
   - If you maintain DB2 archived log files on disk, enter 1 when prompted
   - If you use Tivoli Storage Manager to maintain the DB2 archived log files, enter 2 when prompted
3. Content Manager OnDemand creates the instance, makes a backup image of the database, and restores the Content Manager OnDemand system tables to the database. This process will take several minutes.

   The ARSDB program generates a series of messages. For example:

```
Creating the DB2 ARSDBASE database
Creating table ARSSERVR.arsag
Creating index ARSSERVR.arsag_name_idx
Creating index ARSSERVR.arsag_agid_idx
..... .. . . . . . . . . . . . . . . . . . .
..... .. . . . . . . . . . . . . . . . . . .
..... .. . . . . . . . . . . . . . . . . . .
Updating runstat statistics for table ARSSERVR.arsusrgrp
Creating table ARSSERVR.arsusrgrpid
Creating index ARSSERVR.arsusrgrpid_idx
Updating runstat statistics for table ARSSERVR.arsusrgrpid
```

Initializing the system logging facility

After you have successfully created the instance of Content Manager OnDemand, run the ARSSYSCR program to initialize the Content Manager OnDemand system logging facility for the instance.

About this task

To initialize the Content Manager OnDemand system logging facility:

Procedure

1. Log in to the server as the `root` user.
2. Type the following command at the prompt: `/opt/ibm/ondemand/V10.1/bin/arssyscr -I archive -l`

   Where `archive` is the name of the Content Manager OnDemand instance.
3. Press the Enter key.
   Content Manager OnDemand creates the tables that support the system logging facility. This process may take several minutes.

   The ARSSYSCR program generates a series of messages. For example:

```
arssyscr: Updating ARSSERVR.ARSSYS
arssyscr: Adding to ARSSERVR.ARSAG with Storage Set Id = 0
arssyscr: Adding to ARSSERVR.ARSAGPERMS
arssyscr: Adding to ARSSERVR.ARSAGFLD
```

```
arssyscr: Adding to ARSSERVR.ARSAGFLDALIAS
arssyscr: Adding to ARSSERVR.ARSAG2FOL
arssyscr: Adding to ARSSERVR.ARSAPPUSR
arssyscr: Adding to ARSSERVR.ARSAPP
arssyscr: Adding to ARSSERVR.ARSFOL
arssyscr: Adding to ARSSERVR.ARSFOLPERMS
arssyscr: Adding to ARSSERVR.ARSFOLFLD
arssyscr: Adding to ARSSERVR.ARSFOLFLDUSR
arssyscr: Creation of System Log information was successful
```

Initialize the system load logging facility

OnDemand provides a logging facility to enable tracking OnDemand loading activity. When you enable load logging, OnDemand stores the messages that are generated by OnDemand load programs in the system load log. You use one of the Content Manager OnDemand client programs to search for and filter messages by load date, application group name, load ID, input file name, and other parameters.

About this task

Before you start OnDemand for the first time, you must initialize the system load logging facility:

Procedure

1. Log in to the server as the `root` user.
2. Type the following command at the prompt: `/opt/ibm/ondemand/V10.1/bin/arssyscr -I archive -a`
3. Press the Enter key.
   Content Manager OnDemand creates the tables that support the system load logging facility. This process may take several minutes.

   The ARSSYSCR program generates a series of messages. For example:

```
arssyscr: Updating ARSSERVR.ARSSYS
arssyscr: Adding to ARSSERVR.ARSAG with Storage Set Id = 0
arssyscr: Adding to ARSSERVR.ARSAGPERMS
arssyscr: Adding to ARSSERVR.ARSAGFLD
arssyscr: Adding to ARSSERVR.ARSAGFLDALIAS
arssyscr: Adding to ARSSERVR.ARSAG2FOL
arssyscr: Adding to ARSSERVR.ARSAPPUSR
arssyscr: Adding to ARSSERVR.ARSAPP
arssyscr: Adding to ARSSERVR.ARSFOL
arssyscr: Adding to ARSSERVR.ARSFOLPERMS
arssyscr: Adding to ARSSERVR.ARSFOLFLD
arssyscr: Adding to ARSSERVR.ARSFOLFLDUSR
arssyscr: Creation of System Load information was successful
```

Initializing the system migration facility

About this task

The system migration facility is required only by customers who plan to migrate application group index data from the database to archive storage.

After you have successfully created the instance of Content Manager OnDemand, run the ARSSYSCR program to initialize the Content Manager OnDemand system migration facility for the instance.

To initialize the Content Manager OnDemand system migration facility:

Procedure

1. Log in to the server as the `root` user.

2. Type the following command at the prompt: `/opt/ibm/ondemand/V10.1/bin/arssyscr -I archive -a`

   Where `archive` is the name of the Content Manager OnDemand instance.
3. Press the Enter key.
   Content Manager OnDemand creates the tables that support the system migration facility. This process may take several minutes.

   The ARSSYSCR program generates a series of messages. For example:

```
arssyscr: Updating ARSSERVR.ARSSYS
arssyscr: Adding to ARSSERVR.ARSAG with Storage Set Id = 0
arssyscr: Adding to ARSSERVR.ARSAGPERMS
arssyscr: Adding to ARSSERVR.ARSAGFLD
arssyscr: Adding to ARSSERVR.ARSAGFLDALIAS
arssyscr: Adding to ARSSERVR.ARSAG2FOL
arssyscr: Adding to ARSSERVR.ARSAPPUSR
arssyscr: Adding to ARSSERVR.ARSAPP
arssyscr: Adding to ARSSERVR.ARSFOL
arssyscr: Adding to ARSSERVR.ARSFOLPERMS
arssyscr: Adding to ARSSERVR.ARSFOLFLD
arssyscr: Adding to ARSSERVR.ARSFOLFLDUSR
arssyscr: Creation of System Migration information was successful
```

Automating instance operations on Linux™

This section describes how to use operating system facilities to automatically start or schedule instance operations.

Procedure

You can automatically start these instance operations whenever the system is started:

1. Start the database on the library server.
2. Start the instance on the library server.
3. Start the instance on an object server.
4. Start the data loading programs.

Results

You can schedule these instance operations to begin automatically on a regular schedule:

1. Schedule application group maintenance on the library server.
2. Schedule application group maintenance on an object server.
3. Schedule system table maintenance.
4. Schedule a backup of the Content Manager OnDemand database.
5. Schedule a backup of the Tivoli Storage Manager database.

Starting the database

You can start the database on the library server using the ARSDB program.

Procedure

- To update the command, enter: `opt/ibm/ondemand/V10.1/bin/arsdb -gv >> /tmp/arsdb.log 2>&1`

  Alternatively, you can start DB2 manually with the db2start command.

The following example shows an INIT record to automatically start the database when the operating system is initialized on the library server:

```
ars2:2:wait:/opt/ibm/ondemand/V10.1/bin/arsdb -gv >> /tmp/arsdb.log 2>&1
```

Important: If the DB2 installation program adds a record to the INIT facility to automatically start the DB2 services, make sure that you place the ARSDB record after the record that starts the DB2 services.

Starting the instance on the library server

About this task

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

The ARSSOCKD program controls a Content Manager OnDemand instance on the library server. The ARSSOCKD program runs on the library server. The data loading program (ARSLOAD) and the maintenance programs (such as ARSADMIN and ARSMAINT) will fail and clients will be unable to connect to the instance if the ARSSOCKD program is not running on the library server.

To manually start the archive instance, you can enter the command:

```
/opt/ibm/ondemand/V10.1/bin/arssockd -l archive
```

The following example shows an INIT record that automatically starts the instance named `archive` when the operating system is initialized on the library server:

```
ars3:2:once:/opt/ibm/ondemand/V10.1/bin/arssockd -l archive
```

Starting the instance on an object server

About this task

The ARSOBJD program controls a Content Manager OnDemand instance on an object server. Content Manager OnDemand programs that work with an instance on an object server will fail if the ARSOBJD program is not running on the object server.

The ARSOBJD program should be started only on object servers that are running on some other workstation than the library server.

To manually start the archive instance, you can enter the command:

```
opt/ibm/ondemand/V10.1/bin/arsobjd -I archive
```

The following example shows an INIT record that automatically starts the instance named `archive` when the operating system is initialized on an object server:

```
ars4:2:once:/opt/ibm/ondemand/V10.1/bin/arsobjd -I archive
```

Starting the data loading programs

About this task

This section describes how to use operating system facilities to automatically start the Content Manager OnDemand data loading programs.

The Content Manager OnDemand data loading programs are:

- ARSJESD, to receive data from z/OS and z/OS systems and store the data in file systems on the server
- ARSLOAD, to create index data and load the data into the system

ARSJESD

The ARSJESD program is the Content Manager OnDemand program that monitors a TCP/IP port for data transmitted to the Content Manager OnDemand server by Download for the z/OS feature from a host system. The ARSJESD program receives the data transmitted by Download for the z/OS feature and stores the data in file systems on the server. See *PSF for z/OS: Download for z/OS* for details about configuring and operating Download for z/OS feature on the host system.

The following example shows an INIT record that automatically starts the ARSJESD program during operating system initialization:

```
ars5:2:once:/opt/ibm/ondemand/V10.1/bin/arsjesd -p 6001 -d /arsacif/acif1
-d /arsacif/acif2 -d /arsacif/acif3 >> /tmp/arsjesd.log 2>&1
```

In the example, the ARSJESD program monitors TCP/IP port number `6001` and stores transmitted data in the specified directories. The ARSJESD program writes output messages to the `arsjesd.log` file in the `/tmp` directory.

You must verify the TCP/IP port number that the ARSJESD program monitors. Replace the string `6001` with the port number that is valid on the server that you are configuring. The ARSJESD program and Download on the z/OS system must specify the same port number. The port number that the ARSJESD program monitors is different than the TCP/IP port number that the Content Manager OnDemand server uses to communicate with clients.

You must verify the names of the directories in which the ARSJESD program can put the data. Replace the strings `/arsacif/acif1`, `/arsacif/acif2`, and `/arsacif/acif3` with the names of directories that are valid on the server that you are configuring.

See the ARSJESD command reference in the *IBM Content Manager OnDemand for Multiplatforms: Administration Guide* for more information about the options and parameters that you can specify.

ARSLOAD

The ARSLOAD program is the main Content Manager OnDemand data loading and indexing program. You can configure the ARSLOAD program to monitor specific file systems for report data downloaded from other systems. If the data needs to be indexed, then the ARSLOAD program calls the indexing program that is specified in the Content Manager OnDemand application. The ARSLOAD program then works with the database manager to load the index data into the database and works with the storage manager to load the report data and resources on to storage volumes.

The Content Manager OnDemand instance (started by using ARSSOCKD or ARSOBJD) must be running, otherwise the ARSLOAD program will fail.

Automating the ARSLOAD program

About this task

The following example shows an INIT record that automatically starts the ARSLOAD program for the instance named `archive` during operating system initialization:

```
ars6:2:once:/opt/ibm/ondemand/V10.1/bin/arsload -v -c /arsacif/acif4 -d /arsacif/acif1
-d /arsacif/acif2 -d /arsacif/acif3 -I archive
```

In the example, the ARSLOAD program checks for input files in the specified directories every ten minutes (the default polling time). An input file must have a file type of .ARD or .PDF to initiate the load process. If an input file needs to be indexed, the ARSLOAD program stores the index data in the specified index directory.

You must verify the names of the directories. Replace the strings `/arsacif/acif1`, `/arsacif/acif2`, `/arsacif/acif3`, and `/arsacif/acif4` with the names of directories that are valid on the server that you are configuring.

After indexing the data, the ARSLOAD program deletes the input files, unless you specify otherwise. Any output or error messages that are generated by the ARSLOAD program are written to `stdout`, `stderr`, and the system log.

See the *IBM Content Manager OnDemand for Multiplatforms: Administration Guide* for more information about the ARSLOAD program.

Scheduling application group maintenance on the library server

About this task

You can run the ARSMAINT program on the library server to maintain application group data in the database and cache storage. See the *IBM Content Manager OnDemand for Multiplatforms: Administration Guide* for more information about the ARSMAINT program.

The instance must be started by the ARSSOCKD program, otherwise the ARSMAINT program will fail.

The following is an example of a CRON record that automatically starts the ARSMAINT program every day at 4 am for the instance named `archive`. The ARSMAINT program will migrate and delete application group index data, optimize application group index data, copy report data from cache storage to archive storage, delete report data from cache storage, and inspect and verify the cache file systems. This format of the command is typically used for a library/object server with Tivoli Storage Manager on one workstation.

```
00 4 * * * /opt/ibm/ondemand/V10.1/bin/arsmaint -cdeimrsv -I archive
```

Scheduling application group maintenance on an object server

About this task

You can run the ARSMAINT program on an object server to maintain application group data in cache storage. See the *IBM Content Manager OnDemand for Multiplatforms: Administration Guide* for more information about the ARSMAINT program.

The instance must be started by the ARSOBJD program, otherwise the ARSMAINT program will fail.

The following is an example of a CRON record that automatically starts the ARSMAINT program every day at 4 am for the instance named `archive`. The ARSMAINT program will maintain application group data in cache storage, including copying report data to archive storage. This format of the command is typically used for an object server with Tivoli Storage Manager on some other workstation than the library server.

```
00 4 * * * /opt/ibm/ondemand/V10.1/bin/arsmaint -cmsv
```

Scheduling system table maintenance

About this task

You can run the ARSDB program to maintain the Content Manager OnDemand system tables on the library server. See the *IBM Content Manager OnDemand for Multiplatforms: Administration Guide* for more information about the ARSDB program.

The instance must be started by the ARSSOCKD program, otherwise the ARSDB program will fail.

The following is an example of a CRON record that automatically starts the ARSDB program to maintain the Content Manager OnDemand system tables for the instance named `archive`. The ARSDB program will run twice a month, on the 7th and 14th of each month, beginning at 5 am.

```
00 5 7,14 * * /opt/ibm/ondemand/V10.1/bin/arsdb -mv -I archive >> /tmp/
arsdb.log 2>&1
```

Scheduling the Content Manager OnDemand database backup

About this task

You can use the ARSDB program to create a backup image of the Content Manager OnDemand database. The ARSDB program supports table space backups and full database backups, online backups and offline backups, and the use of Tivoli Storage Manager to maintain the backup image files. See the *IBM Content Manager OnDemand for Multiplatforms: Administration Guide* for more information about the ARSDB program.

The following is an example of a CRON record that automatically starts the ARSDB program to create a full online backup image of the Content Manager OnDemand database for the instance named `archive` every day beginning at 5:30 am. The backup image is written to a tape in the device `/dev/rmt0`. A tape must be mounted in the device before the ARSDB program begins.

```
30 5 * * * /opt/ibm/ondemand/V10.1/bin/arsdb -v -z /dev/rmt0 -I archive >> /tmp/arsdb.log 2>&1
```

Next step on Linux™

After you have installed Content Manager OnDemand and related software on the system, configured the instance of Content Manager OnDemand, created the instance, and automated instance operations, you are now ready to verify the installation on Content Manager OnDemand.

Related tasks

Verifying the installation

After you have completed installation and configuration of the database manager, Content Manager OnDemand software, and Tivoli Storage Manager software, and have configured and initialized the system, perform the following tasks.

Chapter 4. Installing Content Manager OnDemand on Windows servers

This part of the IBM Content Manager OnDemand for Multiplatforms: Installation and Configuration Guide explains how to install and configure Content Manager OnDemand on a Windows server and how to install and configure related software to work with Content Manager OnDemand.

About this task

There are five basic phases to the installation:

- Preparing for the installation
- Installing and configuring Content Manager OnDemand and related software
- Verifying the installation
- Preparing the system for use
- Adding optional software

You will find checklists for each of these phases in “Checklist for installation on Windows” on page 92.

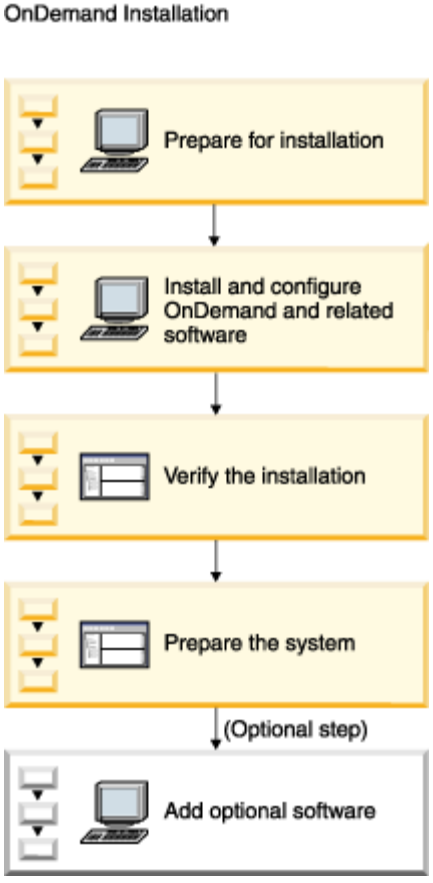

Figure 12: Installing Content Manager OnDemand on a Windows server

Checklist for installation on Windows

Review the pre-installation tasks before installing the product on Windows.

About this task

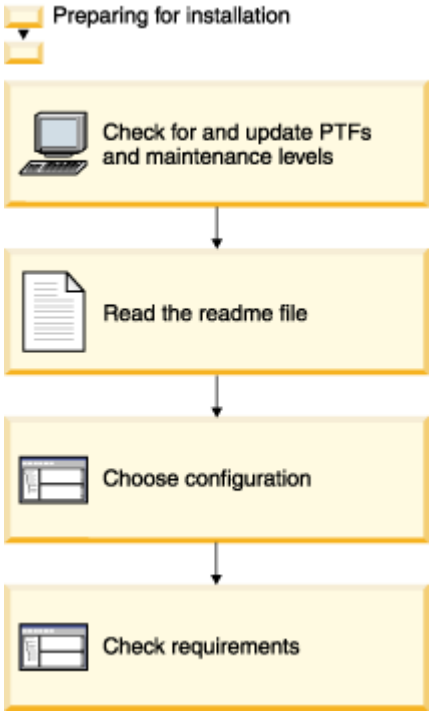

Figure 13: Pre-installation tasks

Procedure

Before beginning the installation, you should complete the following tasks:

1. Contact the IBM support center for the latest maintenance levels of DB2, Content Manager OnDemand, and optionally, Infoprint Manager (Infoprint) and Tivoli Storage Manager. If you are using Oracle instead of DB2, you should contact Oracle for information about the latest maintenance level of Oracle. If you are using SQL Server instead of DB2, you should contact Microsoft for information about the latest maintenance level of SQL Server.
2. Obtain a copy of the Content Manager OnDemand README file. Print and read the entire file before you begin.
3. Determine the type of Content Manager OnDemand system that you need to configure (see “Choosing a configuration” on page 2).
4. Check the Content Manager OnDemand prerequisites and verify the required and optional hardware and software products. See “Windows™ server requirements” on page 96 for information on specific server requirements.
5. Check the hardware and software requirements for all system components and features. See http://www.ibm.com/support/docview.wss?uid=swg27016455 for details.

Results

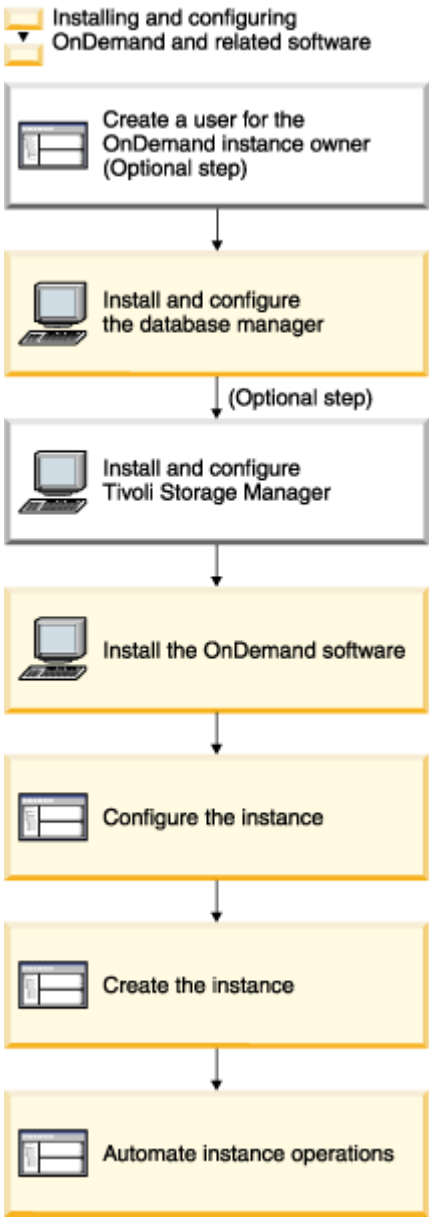

Figure 14: Installing Content Manager OnDemand and related software

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

1. Create a Content Manager OnDemand system administrator account on each workstation that is part of the Content Manager OnDemand system (see "Content Manager OnDemand system administrator account" on page 96).
2. Install and configure the database manager product on the library server (see "Installing the database manager on Windows" on page 98).
3. If you plan on using SSL for security, set up SSL on the Content Manager OnDemand server and client (see "SSL for Content Manager OnDemand" on page 101).
4. If you plan to maintain data in archive storage, install and configure Tivoli Storage Manager on the library server or on each object server that will be used to maintain data in archive storage (see "Installing and configuring Tivoli Storage Manager on Windows" on page 104).
5. Install the Content Manager OnDemand software on each workstation that is part of the Content Manager OnDemand system (see "Installing the Content Manager OnDemand software on Windows" on page 111).

6. Configure the Content Manager OnDemand software on each workstation that is part of the Content Manager OnDemand system (see “Configuring instances on Windows” on page 114). This step includes configuring the database manager (library server only), storage manager, services, and maintenance tasks and creating and initializing the database (library server only).

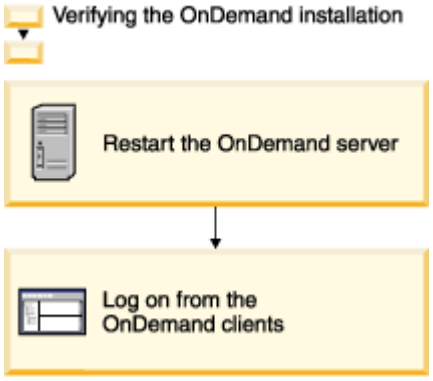

Figure 15: Verifying the installation

Verify the installation of Content Manager OnDemand (see “Verifying the installation” on page 137).

1. After installing and configuring each Content Manager OnDemand server, restart the system to reinitialize the operating system and start the services required by Content Manager OnDemand.
2. Share the Content Manager OnDemand client program folder so that other users on the network can install client software. You should use the share name `odclient` to share the `\Program Files\IBM \OnDemand Clients\V10.1\` folder.
3. Log on to the library server with a Content Manager OnDemand client program. To access the system, you must install at least one of the Content Manager OnDemand client programs on a PC running Microsoft Windows. See *IBM Content Manager OnDemand: Client Installation Guide* for installation information about the Content Manager OnDemand client. See *IBM Content Manager OnDemand for Multiplatforms: Administration Guide* for installation information about the administrative client.

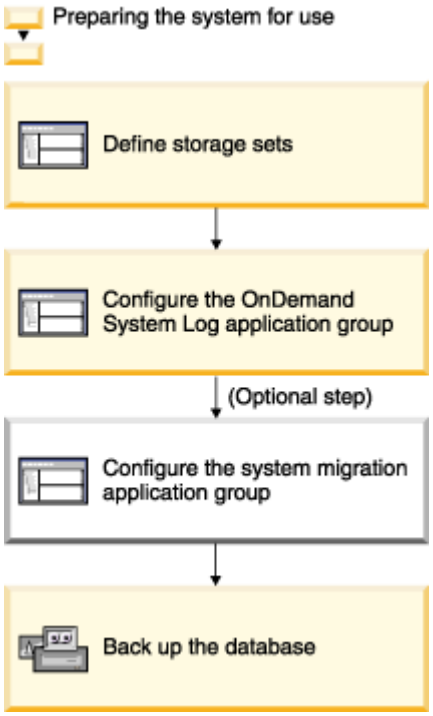

Figure 16: Preparing the system for use

Prepare the system for use:

1. Define storage sets (see "Define storage sets" on page 138). Before you define reports or load data into the system, you must define storage sets.
2. Configure the System Log application group (see "Configuring the System Log application group" on page 138). Before you define reports to the system, load data, or let users access the system, you should configure the System Log application group.
3. If you plan to migrate index data to archive storage, configure the System Migration application group (see "Configure the System Migration application group" on page 142).
4. Backup the database (see Chapter 7, "Backing up the Content Manager OnDemand database," on page 145). After configuring the system, you should create a full backup image of the Content Manager OnDemand database.

Installing and configuring optional software

You must complete specific tasks if you plan to use Download for the z/OS feature to transmit data from z/OS systems to servers, reprint documents by using the server print function, or if you need to customize and enhance the standard functionality within the product.

Procedure

To install and configure optional software:

1. If you plan to use Download for the z/OS feature (Download) to transmit data from z/OS systems to servers, then you must install and configure Download. Follow the instructions in *PSF for z/OS: Download for z/OS* to plan, install, configure, and verify the installation of the Download software. Then, configure Download on each server. Complete the following tasks:
   a) Obtain a copy of *PSF for z/OS: Download for z/OS*.
   b) Check the prerequisites and verify the z/OS and TCP/IP software levels for Download.
   c) Install and configure Download on the system.
   d) Configure the Content Manager OnDemand MVSD service on each server that will use Download to receive data sets.
2. If you plan to reprint documents by using the server print function, then you must install Infoprint on a workstation that belongs to the same network as the library server. Follow the instructions in the Infoprint documentation to plan, install, configure, and verify the installation of the Infoprint software. Then, configure the server print function on the library server. Complete the following tasks:
   a) Obtain a copy of the Infoprint documentation for your server.
   b) Install and configure Infoprint.
   c) Verify that all resources and fonts that your organization requires to reprint the reports that you plan to store in Content Manager OnDemand are installed on the Infoprint server.
   d) On the Infoprint server, define the print queues and devices that Infoprint uses to manage the Content Manager OnDemand server print function.
   e) Obtain the TCP/IP host name or TCP/IP address of the Infoprint server.
   f) On the library server, update the `ARSPRT.BAT` file with the TCP/IP host name or TCP/IP address of the Infoprint server and the fully qualified file name of the arslpr program. You can find the `ARSPRT.BAT` file and the `arslpr.exe` file in the `\Program Files\IBM\OnDemand\V10.1\bin` directory.
   g) Define a server printer on the Content Manager OnDemand library server with the administrative client.
3. If you need to customize and enhance the standard functions within the product, see the user exit documentation. A user exit is a point during processing that enables you to run a user-written program and return control of processing after your user-written program ends. Content Manager OnDemand provides the following user exit points:
   a) Download user exit.
   b) Report specifications archive definition user exit.

c) Retrieval preview user exit.

d) Security user exit.

e) System log user exit.

f) Table space creation user exit.

Windows™ server requirements

About this task

The exact hardware and software configuration that you need for Content Manager OnDemand to support your organization depends on the volume of data that you plan to maintain on the system, the number of concurrent users that the system must support, the backup and recovery requirements of your organization, and the performance levels that the system must meet. At a minimum, you need one processor for a standard Content Manager OnDemand library/object server.

For all Windows server requirements, see http://www.ibm.com/support/docview.wss?uid=swg27049168

Content Manager OnDemand system administrator account

Each library server must have a user account that will be used to install Content Manager OnDemand software products, administer the system, load data, and perform other Content Manager OnDemand functions. If you use DB2 to manage the database, the user name must meet the DB2 naming rules.

You should create the ODADMIN user account on each workstation on which you plan to install the Content Manager OnDemand server software. Assign the account a password. Specify the following characteristics:

- Enter a description, such as: Content Manager OnDemand system administrator account
- Add the account to the local group: Administrators
- All logon hours should be allowed
- Modify the local security policy settings to grant the following user rights to the new user:
  - Act as part of the operating system
  - Create a token object
  - Increase quotas
  - Log on as a service
  - Replace a process level token

Important if you are using a unified login: After you install and configure the system, remember to add the ODADMIN user to Content Manager OnDemand. Set the User Type to **System Administrator**. If you change the password in Windows, remember to change it in Content Manager OnDemand. Or, if you change the password in Content Manager OnDemand, remember to change it in Windows.

Unified login for user accounts

The unified login mode in Content Manager OnDemand works with the Windows authentication process to allow users to log on to Content Manager OnDemand using their Windows account user names.

This feature means that when an authorized user starts the Content Manager OnDemand client after logging on to Windows, the user does not have to enter a Content Manager OnDemand user ID and password. This is because Windows account user names are automatically associated with Content Manager OnDemand user IDs. The LAN Manager Security Support Provider service is used to associate the accounts and userIDs.

How a user is authenticated

You can authenticate a user with unified logins.

Procedure

A user is authenticated with unified login:

1. The user must first log on to Windows with a valid Windows account.
2. Start the Content Manager OnDemand client.

   The user's Windows account user name is routed from Windows to Content Manager OnDemand. The password is not routed to Content Manager OnDemand because it is not used to authenticate the user. The only requirement is that the Windows account user name exist in the Content Manager OnDemand User table.

   Content Manager OnDemand searches the User table for the Windows account user name:

   - If there is a match, then the user is authenticated and logged on to Content Manager OnDemand. A list of folders (or the user's default folder) is then displayed.
   - If there is no match, then Content Manager OnDemand displays the Logon dialog box. The user must enter a valid Content Manager OnDemand userid and password to proceed.

How to prepare for unified login

To use unified login mode, the client must be running Windows 7 or higher and the server must be running Windows 2008 R2.

About this task

All users that access Content Manager OnDemand must have valid Windows accounts. The Windows account user names of the users must be added to Content Manager OnDemand with the Content Manager OnDemand administrative client. You should create a Windows account for an administrative user that will run the Content Manager OnDemand services.

Content Manager OnDemand uses unified login when the user initially starts the client software. If the user logs off the server but does not exit the client software, then the user must use the Logon command the next time the user wants to log on to the server. The user must enter a valid Content Manager OnDemand userid and password to logon to the server.

In unified login mode, Content Manager OnDemand attempts to log on to the user's default server. The default server is established the first time a user logs on to Content Manager OnDemand. The default server is always the last server the user selected in the Logon dialog box, with one exception. The default server name can be fixed by specifying the /S startup parameter in the Properties of the shortcut used to start the client software.

A user logged on to Windows with an account user name that exists in the Content Manager OnDemand User table does not have to enter a Content Manager OnDemand user ID and password to run Content Manager OnDemand commands from the command line. However, Content Manager OnDemand validates the permissions of the userid, to verify that the user has the right to perform the requested action.

With the administrative client, you can log on to a server using a Content Manager OnDemand user ID that is different than the Windows account user name with the Logon As command. To access the Logon As command, point to the server and click the right mouse button. If the user is currently logged on to the server, the user must log off the server before selecting the Logon As command.

Installing the database manager on Windows

The Content Manager OnDemand library server maintains system information and user-defined index data in a relational database.

About this task

You can use DB2, Oracle, or Microsoft SQL Server as the database manager. For all products, see the product documentation for complete installation instructions. This section provides installation and configuration information specific to Content Manager OnDemand for DB2, Oracle, and Microsoft SQL Server.

Installing DB2®

You must install DB2 on the Content Manager OnDemand library server.

About this task

The DB2 Universal Database Enterprise Edition program DVDs or electronic images are provided with the Content Manager OnDemand program package. The DB2 technical information is available in HTML and PDF formats on separate DVDs or electronic images. The README file explains how to locate the information that you need. Follow the instructions in *IBM DB2 Universal Database Quick Beginnings for DB2 Servers* to plan, install, configure, and verify the installation of DB2.

Procedure

To install DB2 on the library server:

1. Log on with the Content Manager OnDemand system administrator account.
2. Insert the DB2 CD-ROM into the CD-ROM drive. The setup program automatically starts after you load the CD-ROM into the drive.
3. When prompted, select **Typical** as the installation type, to install all DB2 components required to support Content Manager OnDemand. You can take most default options (unless you have specific requirements of your own).
4. When prompted, enter the user name and password of the Content Manager OnDemand administrator account.
5. After you install the software, apply the latest fix pack for DB2.

   You can obtain the latest fix packs at http://www.ibm.com/support/docview.wss?uid=swg27007053. Print the README file. Follow the instructions in the README file to apply the service update. After installing a fix pack, you might need to update your database instances (for example, archive). See the DB2 README for details.

Installing Oracle

About this task

You must install Oracle on the Content Manager OnDemand library server.

What to do next

1. Create the database.

   You should create the Content Manager OnDemand database using Oracle utilities. The name that you specify for the database should match the value that you specify for the Content Manager OnDemand instance name. For example, ARCHIVE. See “Configuring instances on Windows” on page 114.
2. Change the Oracle parameters:

a. Connect to the database by using sqlplus. Connect as the Oracle user with dba privileges. For example:

```
sqlplus "/ as sysdba"
```

b. Change the Oracle parameters:

```
alter system set remote_os_authent = TRUE SCOPE = SPFILE;
alter system set os_authent_prefix = '' SCOPE = SPFILE;
```

c. Restart the database.

3. Create the userid of the Content Manager OnDemand instance owner in Oracle.

   All tables created by Content Manager OnDemand will be owned by the user that you create in this step. If you want to have a default Oracle table space for the user, then you should specify the table space when you create the user.

   Use the following format of the CREATE USER command:

```
CREATE USER userid IDENTIFIED EXTERNALLY ;
GRANT dba to userid ;
```

   Where `userid` is the Windows account user name that you will use to create the instance with the Content Manager OnDemand configurator program (see “Configuring instances on Windows” on page 114).

4. Integrate Content Manager OnDemand with the Oracle shared library.

   Verify that the Oracle program directory was added to the PATH during the software installation. Unless you specify otherwise, the Oracle program directory is `\oracle\ora92\bin`.

Installing SQL Server 2012

You typically install Microsoft SQL Server 2012 and Content Manager OnDemand on the same system. The database that is used by the Content Manager OnDemand instance is created locally and is on the local file system.

Before you begin

Although it is not preferred, it is possible to run the Content Manager OnDemand server against a remote SQL Server database. In this case, the database is on a different server than the one on which the Microsoft SQL Server is installed.

Microsoft SQL Server 2012 requires Microsoft .NET Framework 3.5.1 to be installed. On Windows 2008 R2 server, this step is not completed by the Microsoft SQL Server 2012 installation process. Complete this step by using the Windows 2008 R2, 2012 or 2012 R2 Server Manager outside the Microsoft SQL Server 2012 installation.

- In the Server Manager interface, select **Features** on the left pane and select **.NET Framework 3.5.1** from the list.
- Click **Add Features** on the right pane if .NET Framework 3.5.1 is not shown.
- Follow the screen directions to install the feature and restart, as instructed.

This limitation also applies to Microsoft SQL Server 2014, but does not apply to Microsoft SQL Server 2008 R2.

Procedure

To install Microsoft SQL Server on the library server:

1. Click `setup.exe` to start the SQL Server Installation Center.
2. Select **Installation** on the left pane.
3. Select **New SQL Server standalone installation** or add features to an existing installation.

Make sure that there are no failures on the Setup Support Rules panes.

4. On the Setup Role pane, select **SQL Server Feature Installation**.
5. On the **Feature Selection** screen, select **Database Engine Services**, **Client Tools Connectivity**, **Documentation Components**, and **Management Tools - Basic**.

   Make sure that there are no failures on the Installation Rules pane.
6. On the instance **Configuration** pane, select **Default** instance.
7. On the **Database Engine Configuration** pane, select **Mixed Mode (SQL Server authentication and Windows authentication)** and specify a password for the Microsoft SQL Server admin ID, `sa`.

   The Mixed Mode authentication is needed to create a new SQL Server user ID later to be used by the OnDemand instance owner.
8. Click **Add Current User** to also add the logged on Windows user ID to the Microsoft SQL Server and click **Next**.

   Make sure that there are no failures on the **Installation Configuration Rules** pane.
9. When the Microsoft SQL Server 2012 installation completes, make sure that there are no errors, close the SQL Server Installation Center, and restart the server, if needed.
10. Start the Microsoft SQL Server 2012 Management Studio and connect to Database Engine.
11. Click **Security** and expand **Logins** on the left pane.

    The Windows user ID used to install Microsoft SQL Server 2012 is listed but in the form of `<domainname>\user_name`. Create a SQL Server login without the host name or domain name.
12. Right-click **Logins** and select **New Login...**
13. Enter the name of an existing Windows user ID with administrator rights, such as `odadmin`.

    This name is the user ID that is used to create the OnDemand instance.
14. Select **SQL Server Authentication** and enter a password for the new Microsoft SQL Server user ID.
15. Accept the default database name master.
16. On the **Server Roles** page, make sure that the dbcreator box is checked.
17. Click **OK** to create the new login.

What to do next

Before you can create the Content Manager OnDemand Instance with Microsoft SQL Server 2012, you must verify that a correct version of the `regasm.exe` file exists. Content Manager OnDemand is released with an SMO (SQL Server Management Objects) DLL that needs to be registered by `regasm` at installation time. Content Manager OnDemand V9.0 uses a 64-bit `v2.0.50727 RegAsm.exe` file. Content Manager OnDemand V9.5 and V10.1 use a 64-bit `v4.0.30319 RegAsm.exe` file.

A `v2.0.50727 RegAsm.exe` file is normally available on the Windows Server 2008 R2 while a `v4.0.30319 RegAsm.exe` file is normally available on the Windows Server 2012 and Windows Server 2012 R2.

On systems without a `v2.0.50727 RegAsm.exe` file, you can install .NET Framework 3.5. This solution includes 2.0 and 3.0. If .NET Framework 3.5 is not installed, SQL 2008 R2 can install it automatically. Microsoft SQL Server 2012 and 2014, however, does not install .NET Framework 3.5 automatically. On Windows Server 2008 R2 and higher, you might need to use Server Manager to install it from the feature selection pane.

On systems without a `v4.0.30319 RegAsm.exe` file, you can install .NET Framework 4 or 4.5. If .NET Framework 4 is not installed, run the Windows Update to install it. To find out which version of the `regasm.exe` file is installed, run following commands from a Windows command prompt:

```
C:\>CD C:\Windows\Microsoft.NET

C:\Windows\Microsoft.NET>DIR regasm.exe /s/b
```

The sample output shows that both versions are installed.

```
C:\Windows\Microsoft.NET\Framework\v2.0.50727\RegAsm.exe
C:\Windows\Microsoft.NET\Framework\v4.0.30319\RegAsm.exe
C:\Windows\Microsoft.NET\Framework64\v2.0.50727\RegAsm.exe
C:\Windows\Microsoft.NET\Framework64\v4.0.30319\RegAsm.exe
```

If Content Manager OnDemand V10.1 is installed without a 64-bit `v4.0.30319 RegAsm.exe` file, the Content Manager OnDemand SMO DLL cannot be registered. Any attempt to run `arsdb` to create an SQL Server database might result in a system crash. Make sure that you have a 64-bit `v4.0.30319 RegAsm.exe` file before you install Content Manager OnDemand V10.1.

- If Content Manager OnDemand V10.1 is installed without a 64-bit `v4.0.30319 RegAsm.exe` file, install .NET Framework 4.5.1 manually or by using the Windows Update.
- After you verify that you have a 64-bit `v4.0.30319 RegAsm.exe` file, reinstall Content Manager OnDemand V10.1.
- Alternatively, you can apply the most recent Content Manager OnDemand V10.1 fix pack, such as 10.1.0.1, to register Content Manager OnDemand SMO during the installation.

To verify whether Content Manager OnDemand SMO is successfully registered, look for the following registry key:

```
HKEY_LOCAL_MACHINE\SOFTWARE\Classes\TypeLib\{ECC6E2B2-9901-4A26-
AEF0-085507A375EF}\v.r
```

where v.r can be 9.0, 9.5, or 10.1 for Content Manager OnDemand V9.0, V9.5, or V10.1, for example.

Microsoft SQL Server 2008 R2 can install .NET 3.5 automatically but does not install .NET Framework 4 automatically. For Content Manager OnDemand V10.1, .NET Framework 4 is required but not installed automatically by Microsoft SQL Server 2008 R2. For this reason, install Microsoft SQL Server 2012 for Content Manager OnDemand V10.1.

SSL for Content Manager OnDemand

Secure Sockets Layer (SSL) and its successor, Transport Layer Security (TLS), encrypts all transmissions between the Content Manager OnDemand servers.

Secure Sockets Layer (SSL) encrypts all transmissions between the Content Manager OnDemand servers and many of the supported clients (for example, ODWEK Java API, the Windows client, and ARSDOC command). The CICS client does not support SSL connections.

Before you begin setting up SSL on Content Manager OnDemand for Windows

Because of possible problems with system performance, create SSL connections only for communications requiring secure transmission. Consider adding additional processor resources on the Content Manager OnDemand server, client, or both to manage the increased overhead.

GSKit provides the GSKCapiCmd tool, which helps you create and manage digital certificates and key databases. The instructions in "Setting up SSL on the Content Manager OnDemand for Windows server" on page 102 provide examples of how to run the GSKCapCmd tool; however, to view the complete syntax and understand the behavior of this tool, see ftp://ftp.software.ibm.com/software/webserver/appserv/library/v80/GSK_CapiCmd_UserGuide.pdf.

Choose the scenario from the following list that fits your requirements, then follow the instructions for that scenario:

- Content Manager OnDemand server listens only on a non-SSL port. You cannot set up SSL for this situation. Continue to the next Content Manager OnDemand installation task.
- Content Manager OnDemand server listens only on a SSL port. You must do the following tasks:
  - Set up SSL on Content Manager OnDemand.

  - Install GSKit on all clients.
  - Configure the clients to support SSL.
- Content Manager OnDemand server listens on both a non-SSL port and a SSL port. You must do the following steps:
  - Set up SSL on Content Manager OnDemand.
  - Install GSKit on the clients connecting to the SSL port.
  - Configure those clients to support SSL.

Setting up SSL on the Content Manager OnDemand for Windows server

Procedure

To set up SSL on Content Manager OnDemand:

1. Create the key database and store it in the `config` subdirectory of Content Manager OnDemand server installation directory: `C:\Program Files\IBM\OnDemand Server\V10.1\config`.

   To create the key database, run commands similar to the following:

```
SET PATH=%PATH%;C:\Program Files\IBM\gsk8\bin;C:\Program Files\IBM\gsk8\lib64
gsk8capicmd_64 -keydb -create -db "ondemand.kdb" -pw "myKeyDBpasswd" -stash -
populate
```

   The following list describes why these parameters were chosen:

   -keydb -create -db "ondemand.kdb"

   Indicates that you want to create a key database called ondemand.kdb.

   -pw "myKeyDBpasswd" -stash

   Indicates that you want to create a stash file and store the password (`myKeyDBpasswd`) in that stash file. The GSKCapiCmd tool stores the stash file at the same path as the key database. You must remember this path because you must specify it in the ars.ini file. GSKCapiCmd creates the stash file with the same file name as the key database (`ondemand`), with the file extension of `.sth`. When Content Manager OnDemand starts, GSKit retrieves the password to the key database from this stash file.

   -populate

   Populates the key database with a set of predefined trusted certificate authority (CA) certificates. A trusted CA is a certificate authority root certificate is noted as trusted in the key database.

2. Create a digital certificate. You can create a self-signed certificate, which is useful for testing. When you are ready to move to a production environment, create a CA-signed digital certificate.

   To create a self-signed certificate, do the following steps:

   a. Create a self-signed certificate by using GSKCapiCmd. The following example creates a self-signed certificate with the label `myselfsigned`:

```
gsk8capicmd_64 -cert -create -db "ondemand.kdb" –pw "myKeyDBpasswd" -label
"myselfsigned" \
–dn
"CN=myhost.mycompany.com,O=myOrganization,OU=myOrganizationUnit,L=Boulder,ST=CO,C=US
"
```

   b. Extract the certificate to a file by using GSKCapiCmd. The following example extracts the certificate created into a file called ondemand.arm:

```
gsk8capicmd_64 -cert -extract -db "ondemand.kdb" -pw " myKeyDBpasswd " -label
"myselfsigned" \
-target "ondemand.arm" -format ascii
```

c. Distribute the file you created to all computers running clients that will establish SSL connections to your Content Manager OnDemand server.

To create a CA-signed digital certificate, do the following steps:

a. Create a Certificate Signing Request (CSR) by using GSKCapiCmd. You create a CSR for the following reasons:

- Create a new RSA private-public key pair and PKCS10 certificate request, which are stored in the key database in a file with the extension `.rdb`.
- Specify the name of the file, with the -fileoption, that you send to the CA.

The following example shows how to create a CSR that is stored in ondemand.kdb.

```
gsk8capicmd_64 -certreq -create -db "ondemand.kdb" -pw "myKeyDBpasswd" -label
"mycert" \
-dn
"CN=myhost.mycompany.com,O=myOrganization,OU=myOrganizationUnit,L=Boulder,ST=CO,C=US
" \
-file "mycertRequestNew"
```

b. Verify the contents of the CSR by using GSKCapiCmd. The following example shows how to display the contents of the CSR created:

```
gsk8capicmd_64 -certreq -details -db "ondemand.kdb" -pw "myKeyDBpasswd" –label
"mycert"
```

If you need to delete this CSR, run GSKCapiCmd similar to the following example:

```
gsk8capicmd_64 -certreq -delete -db "ondemand.kdb" -pw "myKeyDBpasswd" -label
"mycert"
```

c. Go to the web site of a well known CA (for example, Verisign) and follow their instructions for registering and obtaining a signed digital certificate. The instructions include paying the CA for their services and providing them with the file you specified with the -file option. In the following example and for the rest of these instructions, a trial version of a digital certificate is used.

d. The CA sends you an e-mail with the following information:

- The MyCertificate.arm file, your trial signed digital certificate.
- A link to download IntermediateCert.arm, the trial intermediate digital certificate.
- A link to download RootCert.arm, the root digital certificate.

Use a text editor (for example, notepad) to save each certificate into a file.

e. Add the trial root digital certificate to the key database. The following example adds RootCert.arm to ondemand.kdb:

```
gsk8capicmd_64 -cert -add -db "ondemand.kdb" -pw "myKeyDBpasswd" -label
"trialRootCACert" \
-file RootCert.arm -format ascii
```

f. Add the trial intermediate certificate to the key database. The following example adds IntermediateCert.arm to ondemand.kdb:

```
gsk8capicmd_64 -cert -add -db "ondemand.kdb" -pw "myKeyDBpasswd" –label
"trialIntermediateCACert" \
-file IntermediateCert.arm -format ascii
```

g. Receive your signed digital certificate to the key database. The following example receives MyCertificate.arm to ondemand.kdb:

```
gsk8capicmd_64 -cert -receive -file MyCertificate.arm -db "ondemand.kdb" –pw
“myKeyDBpasswd” \
–format ascii
```

h. Verify that all the certificates were stored in the key database by using GSKCapiCmd. The following example lists the certificates stored in ondemand.kdb:

```
gsk8capicmd_64 -cert -list all -db "ondemand.kdb" -pw "myKeyDBpasswd"
```

GSKCapCmd displays the following result:

```
Certificates found
* default, - personal, ! trusted
-! mycert
! trialIntermediateCACert
! trialRootCACert
-! myselfsigned
```

3. Configure the server:
   a) Start the Content Manager OnDemand OnDemand Configurator V10.1 program.
   b) Open the Server window.
   c) Select the **Enable** check box in the **SSL Port Number** group. Then, specify the port number you want to use for SSL communications.
   d) Click **Close**.
4. Restart the Content Manager OnDemand server. Because a trusted certificate authority provided the digital certificate, the Content Manager OnDemand server accepts the certificate.

   Both `ondemand.kdb` and `ondemand.sth` files need to be placed on the workstation where the Content Manager OnDemand clients are installed. Download both files to the `config` subdirectory under the client installation directory.

Installing and configuring Tivoli Storage Manager on Windows

This section explains how to set up Tivoli Storage Manager for Content Manager OnDemand on a Windows workstation.

About this task

Tivoli Storage Manager can be used with Content Manager OnDemand object servers to store report data on devices that are supported by Tivoli Storage Manager. Devices supported by Tivoli Storage Manager include optical libraries and tape media. The use of Tivoli Storage Manager is optional and is needed only if you want to provide long-term storage for your reports on devices other than the fixed disks attached to the object server. You can also use Tivoli Storage Manager facilities to maintain DB2 archived log files and backup image files.

You will need the *IBM Tivoli Storage Manager for Windows: Quick Start* publication to install and configure Tivoli Storage Manager. HTML and PDF versions of Tivoli Storage Manager publications, including the Quick Start Guide, are available at http://publib.boulder.ibm.com/tividd/td/tdprodlist.html.

Prerequisites

Content Manager OnDemand uses the Tivoli Storage Manager API client to store data into the Tivoli Storage Manager server.

Content Manager OnDemand supports Tivoli Storage Manager in the following configurations:

- Standard library/object server plus Tivoli Storage Manager on one workstation. Install the Server, Clients, API, Device Drivers, and Licenses on the library server workstation.
- Library server only (where Tivoli Storage Manager resides on some other workstation than the library server). Install the Clients and API on the library server workstation.
- Object server plus Tivoli Storage Manager on some other workstation than the library server. Install the Server, Clients, API, Device Drivers, and Licenses on the object server workstation.

The Tivoli Storage Manager server is managed and administered independently of Content Manager OnDemand. The Tivoli Storage Manager administrator must ensure that the following conditions are met:

- All the normal requirements for Tivoli Storage Manager storage are monitored and managed accordingly
- All required Tivoli Storage Manager policies, management classes, storage pools, and volumes are defined accordingly
- All required Tivoli Storage Manager storage pools and volumes are online
- All Tivoli Storage Manager storage pools and volumes have sufficient storage space to satisfy the needs of Content Manager OnDemand
- The Tivoli Storage Manager server is active when Content Manager OnDemand needs to read from or write to its storage repository

If your Tivoli Storage Manager configuration cannot support Content Manager OnDemand, system requests (that require Tivoli Storage Manager services) will fail. The Tivoli Storage Manager administrator should examine the system to ensure that it will support the storage and retrieval of data by Content Manager OnDemand.

Tivoli Storage Manager objects created during a typical installation

The objects defined to the Tivoli Storage Manager server after you install Tivoli Storage Manager, perform initial configuration, and update the configuration for Content Manager OnDemand.

The objects defined to the Tivoli Storage Manager server will depend on the number and types of devices that you configure on the system. The information in the table assumes that you will configure one automated library on the system (such as an IBM 3995-C64 optical library with two optical drives) and add one client node to hold Content Manager OnDemand data for seven years.

Table 7: Tivoli Storage Manager objects created during a typical installation

Object	**Name**
Automated Library	LB6.0.0.1
Drive 1	OP1.0.0.1
Drive 2	OP3.0.0.1
Storage Pool	OPTPOOL1
Device Class	OPTCLASS1
Client Node	OD7YRPRI
Policy Domain	DOM1
Policy Set	STANDARD
Management Class	STANDARD
Copy Group	STANDARD
Administrative Clients	ADMINODADMIN

Updating the configuration

Provides general guidance about how to configure Tivoli Storage Manager to maintain Content Manager OnDemand data in archive storage.

Before you begin, familiarize yourself with the following information available in the IBM Knowledge Center:

- For detailed information about the Tivoli Storage Manager commands, see the *IBM Tivoli Storage Manager for Windows: Administrator's Reference*. This guide is useful as your primary reference.
- If you encounter problems configuring Tivoli Storage Manager, see the Tivoli Storage Manager publications.

Complete these tasks to set up Tivoli Storage Manager for Content Manager OnDemand on a Windows workstation:

- Define server options
- Define client options
- Register a Content Manager OnDemand administrator
- Update the client node
- Set the expiration period for the activity log
- Update device classes
- Update storage pools
- Update policy information
- Prepare media

Configuring Tivoli Storage Manager to manage DB2 files

You can use Tivoli Storage Manager to maintain DB2 archived log files and backup image files.

About this task

This capability means that you do not have to manually maintain these files on disk. The tasks in this section are optional, and are only recommended for customers who need to use Tivoli Storage Manager facilities to backup and restore DB2 databases. For more information about using Tivoli Storage Manager to manage DB2 files, see *IBM DB2 Universal Database: Data Recovery and High Availability Guide and Reference*, SC09-4831.

Do the following tasks to configure Tivoli Storage Manager to maintain DB2 files:

- Define server options
- Define client options
- Define storage objects
- Register the client node
- Set the client node password
- Review space requirements
- Review backup considerations

Backing up Tivoli Storage Manager information

After you configure Tivoli Storage Manager, you should backup the Tivoli Storage Manager database and save Tivoli Storage Manager server files that contain important information.

The backup copy of the database can be used if you need to recover the database. (The backup copy should be saved until the next time that you create a full backup of the database.) The files contain important information that you must have if you need to recover the database.

You should backup the database and save server files whenever you make changes to the database. The database is modified when you store data in Tivoli Storage Manager and when you make changes to the Tivoli Storage Manager environment, such as adding devices and managing removable media operations.

Backing up the database
Before you backup the database, you must define the backup storage objects to Tivoli Storage Manager and label at least one tape storage volume. You can define one device class for full backups and a different device class for incremental backups.

About this task

For example, you can write full backups to a tape device and incremental backups to a disk device.

Procedure

At a minimum, you should define a tape backup device and its associated device class, library, and storage pool:

1. Start **Server Utilities** from the Tivoli Storage Manager program group.
2. Select **Device Configuration**.
3. Select the Device **Configuration** wizard.
4. Click **Start**.
5. Follow the instructions provided to add a manual tape device to Tivoli Storage Manager and define a device class, library, drive, and storage pool for the tape backup.

Results
Next, label a tape storage volume:

1. Place a blank, unlabeled tape in the drive.
2. Select **Media Labeling**.
3. Select the **Manual Media Labeling** wizard.
4. Click **Start**.
5. Follow the instructions provided to label one or more tape storage volumes.

After you have defined the device and storage objects to Tivoli Storage Manager and labeled at least one tape storage volume, you can backup the database. First, place a labeled tape storage volume in the drive. Then enter the following command from the Tivoli Storage Manager administrative command line interface:

```
backup db type=full devclass=dumptapedev
```

Replace the string `dumptapedev` with the name of the device class that you defined for tape backup. The backup command issues several messages, concluding with "Database dump process completed", after successfully creating the database backup.

Write down information about the database backup, such as the date and volume label, and store the backup copy of the database in a safe location, preferably offsite. (Keep the backup copy until you create another full backup copy of the database.)

Saving critical files
Save a copy of the files on removable media and store the copy in a safe location, preferably offsite. Save the copy until you create another backup copy of the files.

The following files contain important information that you must have if you need to recover the database:

- The server options file (`DSMSERV.OPT`)
- The volume history file (`VOLHIST.out`)
- The device configuration file (`DEVCNFG.out`)

- The Tivoli Storage Manager database and recovery log location file (`DSMSERV.DSK`)

Configure Content Manager OnDemand with DB2® to store logs, migrate database tables, and perform Tivoli Storage Manager backup copies
If you are using DB2 databases, you can configure Content Manager OnDemand to store archive logs, migrate database tables, and perform automatic back up files in Tivoli Storage Manager.

Set the environment variables used by the Tivoli Storage Manager API
The variables should be set as part of the startup environment for the DB2 database instance owner for the Content Manager OnDemand server.

For example in the environment settings for the Administrator user.

DSMI_DIR

Identifies the user-defined directory path where the API trusted agent dsmtca file is located.

The default value is defined in the CFG section of the registry under the instance. For example:

```
DSMI_DIR=c:\program files\tivoli\tsm\baclient
```

DSMI_CONFIG

Identifies the user-defined directory path to the `dsm.opt` file, which contains the Tivoli Storage Manager user options. This variable must contain a fully qualified path and file name.

The value should reference an option file specifically for DB2 so different Tivoli Storage Manager options might be specified:

```
DSMI_CONFIG=c:\program files\tivoli\tsm\baclient\dsm.db2.opt
```

DSMI_LOG

Identifies the user-defined directory path where the `dsierror.log` error file is created.

The default value is defined in the CFG section of the registry under the instance. For example:

```
DSMI_LOG=c:\temp
```

If any changes are made to these environment variables and either the database manager or the Content Manager OnDemand server is running, stop and restart the programs. For example:

1. Stop the Content Manager OnDemand server service using the Content Manager OnDemand configurator.
2. Stop the database manager using the DB2 service tool.
3. Start the database manager using the DB2 service tool.
4. Start the Content Manager OnDemand server service using the Content Manager OnDemand configurator.

Protecting data with the data retention protection (DRP) protocol

To avoid the accidental erasure or overwriting of critical data, Content Manager OnDemand supports the Tivoli Storage Manager APIs related to data retention.

Data retention protection (DRP)
Prohibits the explicit deletion of documents until their specified retention criterion is met. Although documents can no longer be explicitly deleted, they can still expire.

Important: DRP is permanent. After it is turned on, it cannot be turned off.

Event-based retention policy
Retention based on an external event other than the storage of data. For Content Manager OnDemand, the retention event is the call to delete the data. A load, unload, application group delete, or expiration of data triggers the retention event.

Restriction: Content Manager OnDemand does not support *deletion hold*, which is a feature that prevents stored data from being deleted until the hold is released.

If you decide to use these policies in Tivoli Storage Manager, then the following scenarios result:

Table 8: Scenarios of using data retention protection

	Creation-based object expiration policy	**Event-based retention object expiration policy**
Data retention protection off	Content Manager OnDemand issues a delete object command through the Tivoli Storage Manager API. Objects are deleted during the next inventory expiration. If a Content Manager OnDemand application group is being deleted, a `delete filespace` command is issued, and the object file space is immediately deleted with the file space.	Content Manager OnDemand issues an event trigger command through the Tivoli Storage Manager API. The status of the objects that are affected are changed from PENDING to STARTED, and the objects are expired by Tivoli Storage Manager based on their retention parameters. If the retention parameters are set to `NOLIMIT`, then the objects never expire. If a Content Manager OnDemand application group is being deleted, a `delete filespace` command is issued instead, and the objects are immediately deleted along with the file space.

Table 8: Scenarios of using data retention protection (continued)		
	Creation-based object expiration policy	**Event-based retention object expiration policy**
Data retention protection on	Content Manager OnDemand issues no commands to Tivoli Storage Manager. The objects are effectively orphaned by Content Manager OnDemand and are expired by Tivoli Storage Manager based on their retention parameters. If the retention parameters are set to `NOLIMIT`, then the objects never expire.	Content Manager OnDemand issues an event trigger command through the Tivoli Storage Manager API. The event status of the objects that are affected are changed from PENDING to STARTED and the objects will be expired by Tivoli Storage Manager based on their retention parameters. If the retention parameters are set to `NOLIMIT`, then the objects never expire. If aContent Manager OnDemand application group is being deleted, then a `delete filespace` cannot be used with DRP enabled, therefore, the operation is treated the same as if a delete were indicated. The status of all the affected objects is changed from PENDING to STARTED, and they will be expired by Tivoli Storage Manager based on their retention parameters. Because this leaves the file space entries in TSM, you must manually delete these entries when the file space is empty (even with DRP enabled).

Recommendations:

- Set up the application groups to expire by load.
- Define the Tivoli Storage Manager archive copy groups to be event-based, and retain data for 0 days.
- Run the Tivoli Storage Manager inventory expiration regularly to ensure that expired data is removed.

Configuring the Content Manager OnDemand server

After installing and configuring the Tivoli Storage Manager software, you need to configure the Content Manager OnDemand server software with information it needs to operate with Tivoli Storage Manager.

About this task

You can use the Content Manager OnDemand configurator program to configure the Content Manager OnDemand server software.

Configure each instance of Content Manager OnDemand that will use Tivoli Storage Manager. Do the following tasks:

- On the Storage page, select the **TSM** option. Then click **TSM Options**. Specify the location of the Tivoli Storage Manager program directory and the client options file that you create.
- If you plan to use Tivoli Storage Manager to maintain DB2 files, move to the Database page and select the **Use TSM** for **DB2 Files** option. Then **Click Advanced Options**. Specify the location of the client options file that you created.

Installing the Content Manager OnDemand software on Windows

You must install a copy of the Content Manager OnDemand software on each workstation or node that is part of the Content Manager OnDemand system.

Procedure

To install the Content Manager OnDemand software, follow these steps:

1. Log on with the Content Manager OnDemand system administrator account.
2. Insert the Content Manager OnDemand for Windows Servers DVD or unpack the electronic image to a temporary directory.
3. Choose **Run** from the **Start** menu. In the open box, type *x*`:\server\windows\odwin`, where *x* is the letter of the DVD drive, the hard drive, or the network drive where `odwin.exe` is located.
4. Read the Welcome screen and then click **Next**.
   The License Agreement window appears.
5. Select **Yes** to accept the license agreement. Click **Next**.
   The Destination window appears.
6. Accept the default directory name or enter the name of another directory. Click **Next**.
   The Start Copying Files window is displayed.
7. If you have a version of Content Manager OnDemand older than Version 8.5 installed, the installation program removes the previous version before installing the new version.
8. Click **Next**.
   The progress window appears. When the process completes, the Setup Complete windows appears.
9. If you want to view the README file now, click **Finish**. Otherwise, clear the View README checkbox and then click **Finish** to complete the installation and restart the computer.
10. After installing the software from the CD-ROM, apply the latest service update for Content Manager OnDemand. You can obtain the latest service update from IBM service at http://www.ibm.com/eserver/support/fixes/.

Installing optional Content Manager OnDemand software on Windows

Other software is available for installation in addition to Content Manager OnDemand software.

About this task

The command to install the Content Manager OnDemand PDF Indexing feature is:

```
odpdfwin.exe
```

or

```
odpdfwin.exe -i console
```

The command to install the IBM Content Manager OnDemand Distribution Facility feature is:

```
ododfwin.exe
```

or

```
ododfwin.exe -i console
```

The command to install the Content Manager OnDemand Full Text Search server feature is:

```
odftswin.exe
```

or

```
odftswin.exe -i console
```

To install the Content Manager OnDemand Enhanced Retention Management feature, see *Enhanced Retention Management Guide*.

Performing initial configuration

About this task

After you have installed Tivoli Storage Manager and restarted the system, Tivoli Storage Manager prompts you to configure the Tivoli Storage Manager server. The Initial Configuration wizard guides you through the process. You can find an overview of initial configuration and more information about configuring Tivoli Storage Manager in the *IBM Tivoli Storage Manager for Windows: Quick Start* publication. If you are not familiar with Tivoli Storage Manager functions and concepts, you should read the Tivoli Storage Manager Basics and the Configuring and Managing Server Storage chapters in *IBM Tivoli Storage Manager for Windows: Administrator's Guide* before you begin.

The initial configuration does not include all of the functions needed to support Content Manager OnDemand, but it does provide a system with the basic components that Content Manager OnDemand needs.

In general, initial configuration of Tivoli Storage Manager to support Content Manager OnDemand consists of:

- Defining the environment
- Configuring performance
- Configure services
- Registering licenses
- Adding devices
- Configuring volumes
- Adding client nodes

Defining the environment
The Environment Wizard asks you whether you want to configure a Tivoli Storage Manager stand-alone or network environment.

Most Content Manager OnDemand customers should select the Standalone option.

Configuring performance
The Performance Configuration Wizard helps you optimize the performance of the Tivoli Storage Manager server.

Most Content Manager OnDemand customers should select the Mostly Large Files option and the 2 – 49 Client Nodes option.

Configure services
The Service Configuration Wizard asks you if you want to start the Tivoli Storage Manager server and scheduler services and the Tivoli Storage Manager device driver automatically when the system is restarted.

Most Content Manager OnDemand customers should configure the Tivoli Storage Manager server service to start automatically and the Tivoli Storage Manager device driver to enable optical support at boot.

After you configure services, you should restart the system to enable optical support (so that Tivoli Storage Manager will recognize all of your SCSI-attached optical devices). After you restart the system, start the Tivoli Storage Manager Server Utilities and continue with initial configuration from the Device Configuration Wizard.

Registering licenses

You should register the device, client, and other licensed functions that you purchased. For current information about devices supported by Tivoli Storage Manager, contact the IBM support center.

About this task

When you install Tivoli Storage Manager, your system is licensed for the following base Tivoli Storage Manager support:

- An unlimited number of administrative clients
- One backup-archive client
- Enterprise administration functions
- Server-to-server virtual volume support

The License Wizard asks you to select the license options that you purchased. Select the **license options** and then click **Next**.

Most Content Manager OnDemand customers should select Advanced Device Support and enter the number of Tivoli Storage Manager client licenses that they purchased in the space provided.

Adding devices

The Device Configuration Wizard detects manual and automated devices that are attached to the server.

About this task

The Device Selection dialog shows devices that have not been defined to Tivoli Storage Manager in the left pane. The right pane lists devices that have been defined to Tivoli Storage Manager.

- To define a manual device to Tivoli Storage Manager and associate it with a manual library, move the device from the left pane to the right pane.
- To define an autochanger and its drives to Tivoli Storage Manager and associate them with an automated library, move the autochanger to the right and then drag and drop the drives on the autochanger.

  For an autochanger with multiple drives, you must associate the drives with the autochanger element number order as described in the autochanger documentation.

Configuring volumes

When you install Tivoli Storage Manager, the installation program creates a default 13 MB database volume (`db1.dsm`) and a 9 MB recovery log volume (`log1.dsm`).

The database size is determined by the amount of data that you plan to store on the server. You might need to increase the size of the recovery log, depending on the current utilization. The *IBM Content Manager OnDemand for Multiplatforms: Introduction and Planning Guide* provides formulas that you can use to estimate the database and recovery log sizes. You should start by increasing the database size by 256 MB and the recovery log size by 72 MB. As you load data on the server, you can monitor the utilization and increase or decrease the database and recovery log sizes accordingly.

Use the Volume Configuration wizards to increase the size of the Tivoli Storage Manager database volumes and recovery log volumes. The wizards format and define the additional volumes and place them where you want them.

Adding client nodes

The Client Node Configuration Wizard lets you register client nodes and specify policy information.

About this task

A client node links clients and their data with storage volumes and devices. Before Content Manager OnDemand can store data in Tivoli Storage Manager storage, you must register at least one client node. You must register at least one client node in each storage pool that will contain Content Manager OnDemand data.

To add a client node from the Client Node Configuration Wizard:

Procedure

1. Click **Add Node**.
2. In the **Add TSM Nodes** dialog, specify the following information:
   - Client node name.

     When you define a storage node to Content Manager OnDemand (with the Content Manager OnDemand administrative client), you specify the Tivoli Storage Managerclient node name.
   - Client node password. Enter the password in the Password and Verify Password fields.
   - Storage pool name. Select the name that Tivoli Storage Manager associated with one of the devices that you added.
3. Click **OK**.

Configuring instances on Windows

After installing software on the server, you need to configure Content Manager OnDemand to integrate the various software products and control information, building your specific Content Manager OnDemand operating environment.

About this task

In general, initial configuration of a Content Manager OnDemand system consists of:

- Defining an instance
- Specifying properties of the instance:
  - Server type and other options
  - Load information
  - Distribution information
  - Database manager options
  - Storage manager options
- Creating the instance
- Installing services
- Creating and initializing the database

After you complete the initial configuration of your system, you might need to perform advanced configuration, such as:

- Configuring services
- Configuring scheduled tasks
- Defining multiple instances on one workstation

Getting started

You configure servers by using the Content Manager OnDemand Configurator V10.1 program.

To begin, log on with the Content Manager OnDemand system administrator account (see “Content Manager OnDemand system administrator account” on page 96 for details).

Next, select **Start** > **Programs** > **IBM OnDemand Server V10.1** > **OnDemand Configurator V10.1**. The main Configurator window contains a menu bar, toolbar, navigator pane, list pane, and status bar.

The configurator provides online help to assist you with completing tasks. The online help contains information about the options, fields, and commands on the windows, dialog boxes, and property sheets that you see when using the configurator. To display online help, press F1 any time the configurator is

active in Windows. Help is available for dialog box commands and options. The main help topic for each dialog box usually contains information about the purpose of the dialog box and the commands and options that appear on the dialog box. To display an index of help topics, select Search from the Help menu. You can type search words to locate related topics in the help file. To learn about Windows help and for information about how to use Windows help, select Using Help from the Help menu.

System properties for defining instances

If your Content Manager OnDemand system consists of more than one workstation, you must define an instance for each server.

An instance owner is assigned during the process of creating the instance. By default, the user name that you use to log on to Windows is assigned as the instance owner. After an instance is created, only the creater-owner of the instance can update or delete the instance. When you want to create an instance, you should log on with the Content Manager OnDemand system administrator account.

After an instance is created, the following properties of the instance cannot be changed:

- Instance name
- Server type
- Database instance name
- Instance owner
- Database engine
- Location of database
- Size of the database
- Location of log files
- Size of log files
- Number of log files
- First cache file system named

Important: If you configure a separate object server, ensure that the port number of the object server matches the port number of the library server.

If you update an instance, you must stop and restart the Content Manager OnDemand library and object servers by using the configurator program or system services.

When defining an object server:

- You should use the instance name of the library server.
- You must identify the host name of the library server.
- You must specify the same language and code page as the library server.
- You don't specify database information.

When you create an instance, the configurator installs one or more services on the server. Not all of the services are set up to start automatically when the system is booted. Depending on your requirements, you might need to reconfigure one or more of the services before you begin system operation. For example, you might want to configure the Content Manager OnDemand MVSD service to start automatically on any server that will receive data from other systems using Download.

When you create an instance, the configurator creates scheduled tasks. Before you begin system operation, you must configure these tasks to use the correct runtime options for your system and enable them to run.

You can use the configurator to maintain servers locally or remotely. To identify drives, directories, and paths on a remote server, you must either enter the information in the space provided or use the Browse button to identify a shared folder on the remote server. To maintain a server remotely, the user must have sufficient authority on the remote server.

Defining an instance

You must define the Windows server after you install Content Manager OnDemand.

Procedure

1. From the **File** menu, select **New Local Instance**.
   The **Add a Server** dialog box appears.
2. Name the instance.

   For the first and only instance, you should accept the default provided (ARCHIVE). An instance name can be from one to eight character, and can contain the letters A through Z and the numbers 0 through 9.
3. Enter the temporary and print directories for the Content Manager OnDemand programs to use.
4. Click **Next** to continue and specify the properties of the instance.

   You can specify the following properties:

   - Server type and other options
   - Content Manager OnDemand Load program information
   - Content Manager OnDemand Distribution (ODF) program information
   - Database manager options
   - Storage manager options

Defining Server Type properties

You must specify the server properties to configure an instance.

To set the Server Type properties:

1. On the Server Type page, select Library and Object Server or Object Server Only.
2. If you selected Object Server Only, identify the Library Server Name and the Object Server Name.
3. If you selected Library and Object Server, click Advanced Options. If required, change the defaults provided for:

   - Number of Database Servers

     If you set the Number of Database Servers to a value other than `0` (zero) or `1` (one), you should update the license information for your database management product.
   - Content Federation Services (CFS-CMOD)
   - Full Text Index and Search

   Click **OK** to close the **Advanced Options** dialog box and return to the Server Type page.
4. Click **Communications**. If required, change the default Protocol and Port Number.
5. Click **OK** to close the **Communications** dialog box and return to the Server Type page.
6. Click **Next** to continue.

Database servers

When you configure the library server, you determine the number of processes that Content Manager OnDemand can start on the library server to support database requests. This provides a performance advantage by distributing the server workload over several processes, while balancing the impact on system resources.

In addition to database connections by Content Manager OnDemand client programs, the value that you specify must support the number of active Content Manager OnDemand commands and services such as `ARSLOAD`, `ARSDOC`, `ARSODF`, `ARSMAINT`, and `ARSADMIN`.

Each connection to the Content Manager OnDemand database requires a database agent. Content Manager OnDemand can start a database agent for each connection. However, each agent requires its own private memory and some portion of application shared memory. You configure the Number of Database Servers parameter to optimize the way that Content Manager OnDemand handles the database

load. For example, you can configure the server so that Content Manager OnDemand starts a fixed number of database agents, regardless of the number of concurrent database requests. While this might appear restrictive, database requests typically process very quickly. For example, ten database agents can handle many database requests, while balancing the impact on system resources.

You should set the Number of Database Servers parameter to support the peak number of concurrent database connections that you expect the library server to handle. A low value limits access to the database during periods of high database activity. A high value requires more system resources during periods of high database activity. The value that you choose also depends on the characteristics of the queries. For example, general queries typically keep a connection open longer than a more specific query.

Load properties
On the Load page, add directories used by Content Manager OnDemand programs and services such as temporary work space, print work space, and data directories. Also, add the user ID and password to use to run the load process.

Procedure

1. Type in the user ID and password that will run the load process.
2. You can add one or more directories to the list. After you create the instance, you can always add more directories to each list.
3. When finished, click **OK** to continue.

Distribution properties
On the Distribution page, specify the user ID and password to use to run the distribution process, add a directory used by Content Manager OnDemand distribution process for temporary work space, add a directory where Content Manager OnDemand stores the final output, the full path name of the Java program, and provide values for several sleep settings.

Procedure

1. Type in the user ID and password that will run the distribution process.
2. Define a temporary data directory.
3. Define a directory where Content Manager OnDemand stores the final output.
4. Type in, or use **Browse** to locate, the full path name of `java.exe`.
5. Enter values for the various sleep settings displayed in the Distribution page.
6. When finished, click **OK** to continue.

Database properties
The database properties are defined only for library servers.

Procedure

1. On the Database page, select the Database Engine.

   If you select Oracle, there are no other options that you can specify on the Database page. Click Next to continue configuring the instance.
2. If you plan to use Tivoli Storage Manager to maintain DB2 archived log files and backup image files, select the Use TSM for DB2 files option.
3. Click **Advanced Database Options**.
   The Advanced Options dialog box appears.
4. Configure the database options:
   - DB2:
     - Database Location

       Physically separating the database, the directories that contain the log files, and other system and application data improves performance and helps recovery.

- Primary Log File Path
- Archive Log Path (Or Archive Log TSM Option File, if you selected the Use TSM for DB2 files option on the Database page. The TSM for DB2 client options file is usually different from the client options file that is used to maintain application group data.)
- Log File Size
- Number of Primary Log Files

• SQL Server:
  - Database Path

    Physically separating the database, the directories that contain the log files, and other system and application data improves performance and helps recovery.
  - Database Size
  - Transaction Log Path
  - Log File Size
  - Number of Log Files

5. Click **OK** to close the Advanced Options dialog box and return to the Database page.

   If the Database Engine is DB2, optionally define DB2 File Systems. You can add one or more table space file systems to the list. After you create the instance, the first table space file system that you define holds control information and cannot be altered or removed. After you create the instance, you can always add more table space file systems to the list.

   Storing application group data in table space file systems is optional, but highly recommended. See your database management product documentation for more information about table spaces.
6. Click **Next** to continue.

Storage properties
You must configure the storage properties to configure an instance.

Procedure

1. On the Storage page, configure the storage manager. Select **Cache Only** or **Tivoli Storage Manager**.
2. If you selected Tivoli Storage Manager, click **Tivoli Storage Manager Options** to verify the location of the Tivoli Storage Manager program directory and client options file used to maintain application group data. The client options file used to maintain application group data is usually different from the client options file used to maintain DB2 files.
3. Define cache file systems.

   You can add one or more cache file systems to the list. After you create the instance, the first cache file system that you define holds control information and cannot be altered or removed. After you create the instance, you can always add more file systems to the list.

   Physically separating cache file systems and other system and application data improves performance and helps recovery.

Installing services

Use the Install Services dialog box to identify the Windows user account and password that Content Manager OnDemand uses to log on to its services.

About this task

By default, Content Manager OnDemand uses the same user account that you used to create the instance. You should use the Content Manager OnDemand system administrator account. If the user account was not assigned a password in Windows, you must select This User ID does not have a password. Otherwise, you must enter and verify the password in the spaces provided.

After you enter the information in the spaces provided, Content Manager OnDemand adds the instance, services, and tasks to the list.

Creating and initializing the DB2® and SQL Server
Before you can use your Content Manager OnDemand system, you must create the database and initialize the Content Manager OnDemand system tables on the library server.

In addition, if you plan to migrate index data from the database to archive storage, you must initialize the system migration facility. The configurator allows you to complete these steps using the Create Content Manager OnDemand Database dialog box.

When you are ready to proceed, click Create Database Now. The configurator creates and initializes the database. After completing these tasks, you can click View Log File to display messages generated during the creation and initialization process.

Creating and initializing the Oracle server
Oracle users must create the database using Oracle utilities.

About this task

After creating the database, you must create the Content Manager OnDemand system tables by using the ARSDB program.

Procedure

To create the database:

1. Create the Content Manager OnDemand system tables with the ARSDB program.

   The ARSDB program is installed in the `\Program Files\IBM\OnDemand\V10.1\bin` directory. Use the format: `arsdb -I` *`OnDemandInstanceName`* `-rtv`*`OnDemandInstanceName`* is the name of the Content Manager OnDemand instance.

   For example, ARCHIVE (see “Configuring instances on Windows” on page 114).

   Note: The Configurator will offer to create the database at the end of the instance creation if database engine is DB2 and SQL server is installed.
2. Initialize the Content Manager OnDemand System Log tables with the ARSSYSCR program.

   The ARSSYSCR program is installed in the `\Program Files\IBM\OnDemand\V10.1\bin` directory. Use the format: `arssyscr -I` *`OnDemandInstanceName`* `-l`

   Note: The Configurator may also create tables at the end of the database creation.
3. Initialize the Content Manager OnDemand system load tables with the ARSSYSCR program.

   The ARSSYSCR program is installed in the `\Program Files\IBM\OnDemand\V10.1\bin` directory. Use the command: `arssyscr -I` *`OnDemandInstanceName`* `-a`
4. Optional: If you plan to migrate index data from the database to archive storage, initialize the Content Manager OnDemand System Migration tables with the ARSSYSCR program.

   The ARSSYSCR program is installed in the `\Program Files\IBM\OnDemand\V10.1\bin` directory. Use the command: `arssyscr -I` *`OnDemandInstanceName`* `-m`

Creating the database
To create the database, click **Start** > **Programs** > **IBM OnDemand Server V10.1** > **OnDemand Command Window V10.1**.

About this task

Enter the following command at the prompt:

```
arsdb -cv
```

Initializing the system log

After you have successfully created the instance of Content Manager OnDemand, run the ARSSYSCR program to initialize the Content Manager OnDemand system logging facility for the instance:OnDemand creates the tables that support the system logging facility. This process can take several minutes.

Procedure

To initialize the system log:

1. Enter the following command in the Content Manager OnDemand command window prompt: `arssyscr -I archive -l` where `archive` is the name of the Content Manager OnDemand instance.
2. Press **Enter**.

   The ARSSYSCR program generates a series of messages. For example:

```
arssyscr: Updating ARSSERVR.ARSSYS
arssyscr: Adding to ARSSERVR.ARSAG with Storage Set Id = 0
arssyscr: Adding to ARSSERVR.ARSAGPERMS
arssyscr: Adding to ARSSERVR.ARSAGFLD
arssyscr: Adding to ARSSERVR.ARSAGFLDALIAS
arssyscr: Adding to ARSSERVR.ARSAG2FOL
arssyscr: Adding to ARSSERVR.ARSAPPUSR
arssyscr: Adding to ARSSERVR.ARSAPP
arssyscr: Adding to ARSSERVR.ARSFOL
arssyscr: Adding to ARSSERVR.ARSFOLPERMS
arssyscr: Adding to ARSSERVR.ARSFOLFLD
arssyscr: Adding to ARSSERVR.ARSFOLFLDUSR
arssyscr: Creation of System Log information was successful
```

Initializing system migration

The system migration facility is required only by customers who plan to migrate application group index data from the database to archive storage.

About this task

After you have successfully created the instance of Content Manager OnDemand, run the ARSSYSCR program to initialize the Content Manager OnDemand system migration facility for the instance. This process can take several minutes.

Procedure

1. Enter the following command at the Content Manager OnDemand window prompt: `arssyscr -I archive -a` where `archive` is the name of the Content Manager OnDemand instance.
2. Press **Enter**.

   The ARSSYSCR program generates a series of messages. For example:

```
arssyscr: Updating ARSSERVR.ARSSYS
arssyscr: Adding to ARSSERVR.ARSAG with Storage Set Id = 0
arssyscr: Adding to ARSSERVR.ARSAGPERMS
arssyscr: Adding to ARSSERVR.ARSAGFLD
arssyscr: Adding to ARSSERVR.ARSAGFLDALIAS
arssyscr: Adding to ARSSERVR.ARSAG2FOL
arssyscr: Adding to ARSSERVR.ARSAPPUSR
arssyscr: Adding to ARSSERVR.ARSAPP
arssyscr: Adding to ARSSERVR.ARSFOL
arssyscr: Adding to ARSSERVR.ARSFOLPERMS
arssyscr: Adding to ARSSERVR.ARSFOLFLD
arssyscr: Adding to ARSSERVR.ARSFOLFLDUSR
arssyscr: Creation of System Migration information was successful
```

Initializing the system load logging facility

Content Manager OnDemand provides a logging facility to enable tracking Content Manager OnDemand loading activity. When you enable load logging, Content Manager OnDemand stores the messages that are generated by OnDemand load programs in the system load log.

About this task

You use one of the Content Manager OnDemand client programs to search for and filter messages by load date, application group name, load ID, input file name, and other parameters.

Before you start Content Manager OnDemand for the first time, you must initialize the system load logging facility:

Procedure

1. Type the following command at the Content Manager OnDemand window prompt: `arssyscr -I archive -a`
2. Press the **Enter** key.
3. Content Manager OnDemand creates the tables that support the system load logging facility.

   This process may take several minutes.

   The ARSSYSCR program generates a series of messages.

   For example:

```
arssyscr: Updating ARSSERVR.ARSSYS
arssyscr: Adding to ARSSERVR.ARSAG with Storage Set Id = 0
arssyscr: Adding to ARSSERVR.ARSAGPERMS
arssyscr: Adding to ARSSERVR.ARSAGFLD
arssyscr: Adding to ARSSERVR.ARSAGFLDALIAS
arssyscr: Adding to ARSSERVR.ARSAG2FOL
arssyscr: Adding to ARSSERVR.ARSAPPUSR
arssyscr: Adding to ARSSERVR.ARSAPP
arssyscr: Adding to ARSSERVR.ARSFOL
arssyscr: Adding to ARSSERVR.ARSFOLPERMS
arssyscr: Adding to ARSSERVR.ARSFOLFLD
arssyscr: Adding to ARSSERVR.ARSFOLFLDUSR
arssyscr: Creation of System Load information was successful
```

Advanced configuring services

When you create an instance, the configurator installs one or more services on the server.

Not all services are set up to start automatically when the system is booted. Depending on your requirements, you might need to configure the services before you begin system operation. For example, you might want to configure the Content Manager OnDemand MVSD service to automatically start on any server that will receive data from other systems using Download.

Content Manager OnDemand provides the following services:

- Content Manager OnDemand LibSrvr or Content Manager OnDemand ObjSrvr, depending on the Server Type
- Content Manager OnDemand Load Data, one or more, depending on the number of object servers assigned to the instance
- Content Manager OnDemand MVSD, one or more, depending on the number of object servers assigned to the instance
- Content Manager OnDemand Scheduler

Important: The Content Manager OnDemand server service (LibSrvr or ObjSrvr) and the Scheduler service should always be running. If the server service is not running, the maintenance tasks will fail. If the Scheduler service is not running, scheduled tasks will not initiate.

You should verify the properties of the Content Manager OnDemand Load Data service if you plan to keep the Content Manager OnDemand data indexing and loading processes running at all times. You should verify the properties of the Content Manager OnDemand MVSD service if you plan to transmit data from other systems to the server using Download.

To verify the properties of a service:

1. Point to the service and right mouse click.
2. From the pop-up menu, select **Properties**.
3. On the Service page, verify the Startup Type. Refer to the online help for more information.
4. On the Directories page, assign directories to the service. To assign a directory to the service, select the directory in the Available Directories list and click Add. You must assign at least one directory to each service. You can assign up to ten directories to the MVSD service. You can add directories to the Available Directories list from the Directory page of the instance properties dialog box.
5. For the Load Data service, verify the properties on the Load Information page. Refer to the online help for information about data you can enter.
6. For the MVSD service, verify the properties on the Advanced page. See the online help for information about data you can enter.
7. Click **OK** to configure the service.

If more than one object server belongs to the instance, you need to configure the Content Manager OnDemand Load Data service and the Content Manager OnDemand MVSD service on each object server that requires the service.

Configuring scheduled tasks
When you create an instance, the configurator creates several database maintenance tasks.

Before you begin system operation, you must configure these tasks to use the correct runtime options for your system and enable them to run. However, if the database manager is SQL Server, you should use the Database Maintenance Plan wizard from Enterprise Manager to configure and schedule the tasks.

Content Manager OnDemand provides the ability for you to set up tasks to run automatically on a regular schedule:

- ApplGroup Data Maintenance, to maintain application group database tables and cache file systems
- System Table Maintenance, to maintain Content Manager OnDemand system tables
- Content Manager OnDemand Database Backup, to take backup images of the database

Some systems might need to schedule more than one instance of a database maintenance task. For example, a task can be set up to run with specific options every day. Another occurrence of the task can be needed to run once a week, with different options. To create more than one instance of a scheduled task, use the Duplicate command to copy the task. Then use the Properties command to configure the copy of the task.

Important: The Content Manager OnDemand server service (LibSrvr or ObjSrvr) and the Scheduler service should always be running. If the server service is not running, then the maintenance tasks will fail. If the Scheduler service is not running, then scheduled tasks will not initiate.

To configure and enable a scheduled task to run:

1. Click **Scheduled Tasks**.
2. Point to the task you want to configure and click the right mouse button.
3. From the pop-up menu, select **Properties**.
4. On the **Task** page, verify the Startup Path.
5. Click the **Enabled** check box.
6. On the Options page, verify the options used when the task is run. See the online help for information about data you can enter.

7. On the Schedule page, define the frequency and start time. See the online help for information about configuring a schedule.
8. Click **OK** to schedule the task.

If the instance contains more than one object server, you need to configure scheduled tasks on each object server.

Running more than one instance on a workstation
An instance name is unique to a library server. However, you can use the same instance name on more than one server.

For example, `server1` is a library server and `server2` is an object server; both servers can use the same instance name. As a matter of fact, you should use the same instance name to identify a library server and all of the object servers that communicate with the library server. However, you should not use the same instance name on more than one library server.

You can define one or more instances to run on the same workstation. For example, you can run one instance for production work and a different instance for development work. Each instance running on the same workstation must be configured to utilize a different TCP/IP port number and each distributed library and object server instance must operate on the same TCP/IP port number.

Each instance should be configured with different properties. For example:

- Instance name
- Database
- Database file systems
- Cache file systems

Configuring LDAP
To use the Content Manager OnDemand Configurator to set LDAP parameters, select **Start** > **Programs** > **IBM OnDemand Server V10.1** > **OnDemand Configurator V10.1**.

. The LDAP settings are available on the Server page for each instance. The LDAP parameters are explained in the Content Manager OnDemand Configurator contextual help.

Configuring the `ARSLDAP.INI` file
The ARS_LDAP_BIND_MESSAGES_FILE parameter enables Content Manager OnDemand to customize message text returned from an LDAP server that is used to alert users that their LDAP password is about to expire or their LDAP account is locked.

The messages displayed to users are contained in the file referenced by this parameter. To enable this user-configurable message functionality, create a file with the appropriate message strings, and set ARS_LDAP_BIND_MESSAGES_FILE to the full path of the file. The `ARSLDAP.INI` file is provided with example message strings that can be used by the ARS_LDAP_BIND_MESSAGES_FILE parameter.

The `ARSLDAP.INI` file contains the following three sections:

```
[BIND_MESSAGES]
PASSWORD_EXPIRED="C:\Program Files\IBM\OnDemand\V10.1\config
\password_expired.txt"
ACCOUNT_LOCKED="C:\Program Files\IBM\OnDemand\V10.1\config
\account_locked.txt"

[PASSWORD_EXPIRED]
TDS6="Password has expired"
AD="data 532"
UDEF1=
UDEF2=
UDEF3=

[ACCOUNT_LOCKED]
TDS6="Account is locked"
AD="data 775"
```

```
UDEF1=
UDEF2=
UDEF3=
```

The BIND_MESSAGES section specifies the path to the files containing the user-configurable message text that is displayed to users when their LDAP password is about to expire, or their LDAP account is locked. Generic files are supplied, and should be customized to reflect your actual Content Manager OnDemand environment.

An example message that would be displayed to a user:

```
Your LDAP password has expired and needs to be changed.
Log into <company intranet> for password setting instructions.
```

The entries in the PASSWORD_EXPIRED and ACCOUNT_LOCKED sections are for Tivoli Directory Server Version 6.x and Microsoft Active Directory (AD). These sections also contain three user-defined entries (UDEFx), allowing you to enter your own pattern strings for LDAP servers that are not directly supported.

The LDAP server may return additional information when the user's bind operation fails. When an error is returned from the LDAP server, Content Manager OnDemand reads the file referenced by the ARS_LDAP_BIND_MESSAGES_FILE parameter and searches under the two stanzas, [PASSWORD_EXPIRED] and [ACCOUNT_LOCKED], for user-defined text that matches the LDAP server error. If a match is found, Content Manager OnDemand will display the text found in the files defined under the [BIND_MESSAGES] stanza.

If the ARS_LDAP_BIND_MESSAGES_FILE parameter is not defined, has no file referenced, or the PASSWORD_EXPIRED or ACCOUNT_LOCKED files do not exist, the user will receive a default 'The server failed while attempting to logon' message.

Currently only two error conditions can be handled: PASSWORD_EXPIRED and ACCOUNT_LOCKED. The section titles for these two conditions cannot be changed, but you can change the pattern strings and message text presented to the user to define any two error conditions.

Next steps on Windows

After you have installed the Content Manager OnDemand and related software on the system, configured the instance of Content Manager OnDemand, created the instance, and automated instance operations, you are now ready to verify the installation on Content Manager OnDemand.

Related tasks

Verifying the installation

After you have completed installation and configuration of the database manager, Content Manager OnDemand software, and Tivoli Storage Manager software, and have configured and initialized the system, perform the following tasks.

Chapter 5. Configuring other external storage solutions

Content Manager OnDemand supports external cloud storage managers such as Amazon Simple Storage Service (S3), Apache Hadoop Distributed File System (HDFS), IBM Cloud Object Storage, 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, IBM Tivoli Storage Manager (TSM) has been the only option used by Content Manager OnDemand to maintain data stored in external storage. The addition of Amazon S3, Apache HDFS, IBM Cloud Object Storage, 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 TSM does. This means that data in Content Manager OnDemand can be stored in cache as well as stored in Amazon S3, Apache HDFS, IBM Cloud Object Storage, 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 the migration of the data can be scheduled at a later date.

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

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/

IBM Cloud Object Storage
https://www.ibm.com/cloud-computing/infrastructure/object-storage/

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 the migration of the data can be scheduled at a later date.

Configuring an Amazon S3 external storage manager

Content Manager OnDemand supports data storage in an Amazon Simple Storage Service (S3) repository. More information on Amazon S3 can be found at: https://aws.amazon.com/s3/

Updating the ARS.CFG file on AIX, Linux, Linux on System z, or Solaris servers

Perform these steps to configure Amazon S3 on an AIX, Linux, Linux on System z, or Solaris server.

1. Two new entries must be added to the ARS.CFG file.

   For AIX and Solaris servers:

   ```
   ARS_S3_CONFIG_FILE=/opt/IBM/ondemand/V10.1/config/ars.s3
   ARS_S3_CONFIG_DIR=/opt/IBM/ondemand/V10.1/config
   ```

   For Linux and Linux on System z servers:

   ```
   ARS_S3_CONFIG_FILE=/opt/ibm/ondemand/V10.1/config/ars.s3
   ARS_S3_CONFIG_DIR=/opt/ibm/ondemand/V10.1/config
   ```

   The ARS_S3_CONFIG_FILE entry specifies an existing Amazon S3 configuration file which the server uses by default.

   The ARS_S3_CONFIG_DIR entry specifies the directory in which any alternate configuration files are kept. This directory is used if additional Amazon S3 configuration files are defined. The names of these additional configuration files can be specified when defining storage nodes in Content Manager OnDemand. If no configuration file is specified in the storage node, the default configuration file is used.

   The configuration file name and directory path shown in the examples are the recommended values for these entries.
2. The ARS_STORAGE_MANAGER entry in the ARS.CFG file might also need to be changed. If you specify ARS_STORAGE_MANAGER=CACHE_ONLY, this disables all storage managers supported by Content Manager OnDemand.

   To configure the Content Manager OnDemand server to use Amazon S3 as a storage manager, the value must be set to one of the following:

   ARS_STORAGE_MANAGER=TSM
   This setting will enable all external storage managers supported by Content Manager OnDemand. The Content Manager OnDemand server requires additional software to utilize Tivoli Storage Manager (TSM) as a storage manager. If that additional software is not installed, the server will not start when the ARS_STORAGE_MANAGER value is set to TSM.

   ARS_STORAGE_MANAGER= NO_TSM
   This setting will enable all external storage managers supported by Content Manager OnDemand except Tivoli Storage Manager. This setting is used when the additional software to support Tivoli Storage Manager is not installed and Tivoli Storage Manager is not required as an external storage manager.

Updating an instance configuration on Windows servers

Perform these steps to configure Amazon S3 on a Windows server. Both steps use the OnDemand Configurator to create or update the configuration information.

1. Select Amazon S3 as an external storage manager.
2. Set the configuration entries:

   Configuration Directory
   The Configuration Directory specifies the directory in which any alternate configuration files are kept. This directory is used if additional Amazon S3 configuration files are defined. The names of these additional configuration files can be specified when defining storage nodes in Content Manager OnDemand. If no configuration file is specified in the storage node, the default configuration file is used. For example: C:\Program Files\IBM\OnDemand\V10.1\config

   Default Configuration File
   The Default Configuration File specifies an existing Amazon S3 configuration file which the server uses by default. For example: C:\Program Files\IBM\OnDemand\V10.1\config\ars.s3 A sample configuration file is included as part of the installation of Content Manager OnDemand.

Creating an Amazon S3 configuration file

An Amazon S3 configuration file for Content Manager OnDemand contains entries specific to your Amazon S3 implementation. You specify the location and name of the default configuration file in the ARS.CFG entry or via the OnDemand Configurator. Required entries must be specified. Optional entries are not required in the configuration file unless those values need to be changed.

The following list describes the entries that can be specified in an Amazon S3 configuration file.

ARS_S3_SERVER
Specifies the Amazon S3 server name. Do not include `http://` or `https://` in the name. This entry is required.

ARS_S3_REGION
Specifies the Amazon S3 region. This entry is required.

ARS_S3_USE_SSL
Indicates whether or not to use SSL in server communications. The possible values are:

- `0` - SSL will not be used
- `1` - SSL will be used

The default value is `1`. This entry is optional.

ARS_S3_CONNECT_TIMEOUT
Specifies the maximum number of seconds that Content Manager OnDemand waits for a response from the storage manager. The default is `60`. This entry is optional. **Warning:** Setting this value too low might cause connection failures.

ARS_S3_HLD
Specifies the high-level directory name. This attribute is available to group sets of Content Manager OnDemand data together which might be needed if sharing external storage among multiple Content Manager OnDemand servers. **Warning:** Once this value is set, it must not be changed. If it is changed, any data that is already stored will not be retrievable. There is no default value. This entry is optional.

As an example, for a URL such as `https://s3-us-west-2.amazonaws.com/`, the Amazon S3 configuration file contains:

```
ARS_S3_SERVER=s3.amazonaws.com
ARS_S3_REGION=us-west-2
```

Defining an Amazon S3 storage node with the Administrator client

You can define the settings for using the Amazon S3 access method on the **Add a Primary Node** dialog of the OnDemand Administrator client.

The Storage Node field becomes Bucket Name when the Access Method is set to Amazon S3. The bucket name must exist in your repository or access to Amazon S3 will fail.

The Logon and Password fields contain the Amazon S3 access key and password that Content Manager OnDemand needs to access your Amazon S3 repository.

The Access Method radio button is set to `Amazon S3`. For Content Manager OnDemand servers running on all platforms except Windows, the Configuration File Name defaults to the value specified by the `ARS_S3_CONFIG_FILE` parameter in the ARS.CFG file if no value is entered. Otherwise, Content Manager OnDemand looks for the configuration file in the directory defined by the `ARS_S3_CONFIG_DIR` parameter specified in the ARS.CFG file. For Content Manager OnDemand servers running on Windows, the server uses the Configuration File Name field and the Configuration Directory field that are specified in the OnDemand Configurator instead of using the ARS.CFG file parameters.

Configuring an Apache HDFS external storage manager

Content Manager OnDemand supports data storage in an Apache Hadoop Distributed File System (HDFS).

The Apache® Hadoop® project develops a variety of open-source software for reliable, scalable, distributed computing. The project includes Apache HDFS, which is a distributed file system that provides high-throughput access to application data. More information on Apache HDFS can be found at: https://hadoop.apache.org/

Updating the ARS.CFG file on AIX, Linux, Linux on System z, or Solaris servers

Perform these steps to configure Apache HDFS on an AIX, Linux, Linux on System z, or Solaris server.

1. Two new entries must be added to the ARS.CFG file.

   For AIX and Solaris servers:

   ```
   ARS_HDFS_CONFIG_FILE=/opt/IBM/ondemand/V10.1/config/ars.hdfs
   ARS_HDFS_CONFIG_DIR=/opt/IBM/ondemand/V10.1/config
   ```

   For Linux and Linux on System z servers:

   ```
   ARS_HDFS_CONFIG_FILE=/opt/ibm/ondemand/V10.1/config/ars.hdfs
   ARS_HDFS_CONFIG_DIR=/opt/ibm/ondemand/V10.1/config
   ```

   The ARS_HDFS_CONFIG_FILE entry specifies an existing Apache HDFS configuration file which the server uses by default.

   The ARS_HDFS_CONFIG_DIR entry specifies the directory in which any alternate configuration files are kept. This directory is used if additional Apache HDFS configuration files are defined. The names of these additional configuration files can be specified when defining storage nodes in Content Manager OnDemand. If no configuration file is specified in the storage node, the default configuration file is used.

   The configuration file name and directory path shown in the examples are the recommended values for these entries.
2. The ARS_STORAGE_MANAGER entry in the ARS.CFG file might also need to be changed. If you specify ARS_STORAGE_MANAGER=CACHE_ONLY, this disables all storage managers supported by Content Manager OnDemand.

   To configure the Content Manager OnDemand server to use Apache HDFS as a storage manager, the value must be set to one of the following:

   ARS_STORAGE_MANAGER=TSM
   This setting will enable all external storage managers supported by Content Manager OnDemand. The Content Manager OnDemand server requires additional software to utilize Tivoli Storage Manager (TSM) as a storage manager. If that additional software is not installed, the server will not start when the ARS_STORAGE_MANAGER value is set to TSM.

   ARS_STORAGE_MANAGER= NO_TSM
   This setting will enable all external storage managers supported by Content Manager OnDemand except Tivoli Storage Manager. This setting is used when the additional software to support Tivoli Storage Manager is not installed and Tivoli Storage Manager is not required as an external storage manager.

Updating an instance configuration on Windows servers

Perform these steps to configure Apache HDFS on a Windows server. Both steps use the OnDemand Configurator to create or update the configuration information.

1. Select Apache HDFS as an external storage manager.
2. Set the configuration entries:

   Configuration Directory
   The Configuration Directory specifies the directory in which any alternate configuration files are kept. This directory is used if additional Apache HDFS configuration files are defined. The names of these additional configuration files can be specified when defining storage nodes in Content

Manager OnDemand. If no configuration file is specified in the storage node, the default configuration file is used. For example: `C:\Program Files\IBM\OnDemand\V10.1\config`

Default Configuration File
The Default Configuration File specifies an existing Apache HDFS configuration file which the server uses by default. For example: `C:\Program Files\IBM\OnDemand\V10.1\config\ars.hdfs` A sample configuration file is included as part of the installation of Content Manager OnDemand.

Creating an Apache HDFS configuration file

An Apache HDFS configuration file for Content Manager OnDemand contains entries specific to your Apache HDFS implementation. You specify the location and name of the default configuration file in the ARS.CFG entry or via the OnDemand Configurator. Required entries must be specified. Optional entries are not required in the configuration file unless those values need to be changed.

The following list describes the entries that can be specified in an Apache HDFS configuration file.

ARS_HDFS_SERVER
Specifies the Apache HDFS server name. Do not include `http://` or `https://` in the name. This entry is required.

ARS_HDFS_PORT
Specifies the Apache HDFS server port number. This entry is optional if using a standard port. Content Manager OnDemand assumes port 80 for HTTP or port 443 for HTTPS communications.

ARS_HDFS_TLD
Specifies the Apache HDFS top-level directory name. This is any additional path information after the server name and port in the URL. This entry is optional.

ARS_HDFS_USE_SSL
Indicates whether or not to use SSL in server communications. The possible values are:

- `0` - SSL will not be used
- `1` - SSL will be used

The default value is `0`. This entry is optional.

ARS_HDFS_AUTH_TYPE
Specifies the user authentication type. The possible values are:

- `NONE` - Open system
- `KNOX` - Access and authenticate through Apache Knox

The default value is `NONE`. This entry is optional.

ARS_HDFS_CONNECT_ TIMEOUT
Specifies the maximum number of seconds that Content Manager OnDemand waits for a response from the storage manager. The default is `60`. This entry is optional. **Warning:** Setting this value too low might cause connection failures.

ARS_HDFS_FILE_PERMS
Specifies the permissions for new files. The default is `440`. This entry is optional.

ARS_HDFS_HLD
Specifies the high-level directory name. This attribute is available to group sets of Content Manager OnDemand data together which might be needed if sharing external storage among multiple Content Manager OnDemand servers. **Warning:** Once this value is set, it must not be changed. If it is changed, any data that is already stored will not be retrievable. There is no default value. This entry is optional.

As an example, for a URL such as `http://hdfs.example.com/webhdfs/v1`, the Apache HDFS configuration file contains:

```
ARS_HDFS_SERVER=hdfs.example.com
ARS_HDFS_TLD=/webhdfs/v1
```

Defining an Apache HDFS storage node with the Administrator client

You can define the settings for using the Apache HDFS access method on the **Add a Primary Node** dialog of the OnDemand Administrator client.

The Storage Node field is not used for communication with the Apache HDFS server and can be set to any name you choose.

The Logon field is the user name from the Apache HDFS system which Content Manager OnDemand uses to store and retrieve data. A password might not be required for open Apache HDFS systems, so this field is optional.

The Access Method radio button is set to Apache HDFS. For Content Manager OnDemand servers running on all platforms except Windows, the Configuration File Name defaults to the value specified by the ARS_HDFS_CONFIG_FILE parameter in the ARS.CFG file if no value is entered. Otherwise, Content Manager OnDemand looks for the configuration file in the directory defined by the ARS_HDFS_CONFIG_DIR parameter specified in the ARS.CFG file. For Content Manager OnDemand servers running on Windows, the server uses the Configuration File Name field and the Configuration Directory field that are specified in the OnDemand Configurator instead of using the ARS.CFG file parameters.

Configuring an IBM Cloud Object Storage external storage manager

Content Manager OnDemand supports data storage in an IBM Cloud Object Storage repository. More information on IBM Cloud Object Storage can be found at: https://www.ibm.com/cloud-computing/infrastructure/object-storage/

Updating the ARS.CFG file on AIX, Linux, Linux on System z, or Solaris servers

Perform these steps to configure IBM Cloud Object Storage on an AIX, Linux, Linux on System z, or Solaris server.

1. Two new entries must be added to the ARS.CFG file.

   For AIX and Solaris servers:

   ```
   ARS_ICOS_CONFIG_FILE=/opt/IBM/ondemand/V10.1/config/ars.icos
   ARS_ICOS_CONFIG_DIR=/opt/IBM/ondemand/V10.1/config
   ```

   For Linux and Linux on System z servers:

   ```
   ARS_ICOS_CONFIG_FILE=/opt/ibm/ondemand/V10.1/config/ars.icos
   ARS_ICOS_CONFIG_DIR=/opt/ibm/ondemand/V10.1/config
   ```

   The ARS_ICOS_CONFIG_FILE entry specifies an existing IBM Cloud Object Storage configuration file which the server uses by default.

   The ARS_ICOS_CONFIG_DIR entry specifies the directory in which any alternate configuration files are kept. This directory is used if additional IBM Cloud Object Storage configuration files are defined. The names of these additional configuration files can be specified when defining storage nodes in Content Manager OnDemand. If no configuration file is specified in the storage node, the default configuration file is used.

   The configuration file name and directory path shown in the examples are the recommended values for these entries.

2. The ARS_STORAGE_MANAGER entry in the ARS.CFG file might also need to be changed. If you specify ARS_STORAGE_MANAGER=CACHE_ONLY, this disables all storage managers supported by Content Manager OnDemand.

   To configure the Content Manager OnDemand server to use IBM Cloud Object Storage as a storage manager, the value must be set to one of the following:

ARS_STORAGE_MANAGER=TSM
This setting will enable all external storage managers supported by Content Manager OnDemand. The Content Manager OnDemand server requires additional software to utilize Tivoli Storage Manager (TSM) as a storage manager. If that additional software is not installed, the server will not start when the ARS_STORAGE_MANAGER value is set to TSM.

ARS_STORAGE_MANAGER= NO_TSM
This setting will enable all external storage managers supported by Content Manager OnDemand except Tivoli Storage Manager. This setting is used when the additional software to support Tivoli Storage Manager is not installed and Tivoli Storage Manager is not required as an external storage manager.

Updating an instance configuration on Windows servers

Perform these steps to configure IBM Cloud Object Storage on a Windows server. Both steps use the OnDemand Configurator to create or update the configuration information.

1. Select `IBM Cloud Object Storage` as an external storage manager.
2. Set the configuration entries:

Configuration Directory
The Configuration Directory specifies the directory in which any alternate configuration files are kept. This directory is used if additional IBM Cloud Object Storage configuration files are defined. The names of these additional configuration files can be specified when defining storage nodes in Content Manager OnDemand. If no configuration file is specified in the storage node, the default configuration file is used. For example: `C:\Program Files\IBM\OnDemand\V10.1\config`

Default Configuration File
The Default Configuration File specifies an existing IBM Cloud Object Storage configuration file which the server uses by default. For example: `C:\Program Files\IBM\OnDemand \V10.1\config\ars.icos` A sample configuration file is included as part of the installation of Content Manager OnDemand.

Creating an IBM Cloud Object Storage configuration file

An IBM Cloud Object Storage configuration file for Content Manager OnDemand contains entries specific to your IBM Cloud Object Storage implementation. You specify the location and name of the default configuration file in the ARS.CFG entry or via the OnDemand Configurator. Required entries must be specified. Optional entries are not required in the configuration file unless those values need to be changed.

The following list describes the entries that can be specified in an IBM Cloud Object Storage configuration file.

ARS_ICOS_SERVER
Specifies the IBM Cloud Object Storage server name. Do not include `http://` or `https://` in the name. This entry is required.

ARS_ICOS_USE_SSL
Indicates whether or not to use SSL in server communications. The possible values are:

- `0` - SSL will not be used
- `1` - SSL will be used

The default value is 0. This entry is optional.

ARS_ICOS_CONNECT_ TIMEOUT
Specifies the maximum number of seconds that Content Manager OnDemand waits for a response from the storage manager. The default is 60. This entry is optional. **Warning:** Setting this value too low might cause connection failures.

ARS_ICOS_HLD
Specifies the high-level directory name. This attribute is available to group sets of Content Manager OnDemand data together which might be needed if sharing external storage among multiple Content

Manager OnDemand servers. **Warning:** Once this value is set, it must not be changed. If it is changed, any data that is already stored will not be retrievable. There is no default value. This entry is optional.

As an example, for a URL such as `http://sample.cleversafe.com/`, the IBM Cloud Object Storage configuration file contains:

```
ARS_ICOS_SERVER=sample.cleversafe.com
```

Defining an IBM Cloud Object Storage storage node with the Administrator client

You can define the settings for using the IBM Cloud Object Storage access method on the **Add a Primary Node** dialog of the OnDemand Administrator client.

The Storage Node field becomes Vault Name when the Access Method is set to IBM Cloud Object Storage. The vault name must exist in your repository or access to IBM Cloud Object Storage will fail.

The Logon and Password fields contain the IBM Cloud Object Storage logon and password that Content Manager OnDemand needs to access your IBM Cloud Object Storage repository.

The Access Method radio button is set to `IBM Cloud Object Storage`. For Content Manager OnDemand servers running on all platforms except Windows, the Configuration File Name defaults to the value specified by the `ARS_ICOS_CONFIG_FILE` parameter in the ARS.CFG file if no value is entered. Otherwise, Content Manager OnDemand looks for the configuration file in the directory defined by the `ARS_ICOS_CONFIG_DIR` parameter specified in the ARS.CFG file. For Content Manager OnDemand servers running on Windows, the server uses the Configuration File Name field and the Configuration Directory field that are specified in the OnDemand Configurator instead of using the ARS.CFG file parameters.

Configuring an OpenStack Swift external storage manager

Content Manager OnDemand supports data storage in an OpenStack Swift repository. OpenStack Swift is a highly available, distributed, eventually consistent object/blob store. You can use OpenStack Swift to store lots of data efficiently, safely, and inexpensively. More information on OpenStack Swift can be found at: http://docs.openstack.org/developer/swift/

Updating the ARS.CFG file on AIX, Linux, Linux on System z, or Solaris servers

Perform these steps to configure OpenStack Swift on an AIX, Linux, Linux on System z, or Solaris server.

1. Two new entries must be added to the ARS.CFG file.

   For AIX and Solaris servers:

   ```
   ARS_SWIFT_CONFIG_FILE=/opt/IBM/ondemand/V10.1/config/ars.swift
   ARS_SWIFT_CONFIG_DIR=/opt/IBM/ondemand/V10.1/config
   ```

   For Linux and Linux on System z servers:

   ```
   ARS_SWIFT_CONFIG_FILE=/opt/ibm/ondemand/V10.1/config/ars.swift
   ARS_SWIFT_CONFIG_DIR=/opt/ibm/ondemand/V10.1/config
   ```

   The `ARS_SWIFT_CONFIG_FILE` entry specifies an existing Swift configuration file which the server uses by default.

   The `ARS_SWIFT_CONFIG_DIR` entry specifies the directory in which any alternate configuration files are kept. This directory is used if additional Swift configuration files are defined. The names of these additional configuration files can be specified when defining storage nodes in Content Manager OnDemand. If no configuration file is specified in the storage node, the default configuration file is used.

The configuration file name and directory path shown in the examples are the recommended values for these entries.

2. The ARS_STORAGE_MANAGER entry in the ARS.CFG file might also need to be changed. If you specify ARS_STORAGE_MANAGER=CACHE_ONLY, this disables all storage managers supported by Content Manager OnDemand.

   To configure the Content Manager OnDemand server to use Swift as a storage manager, the value must be set to one of the following:

   ARS_STORAGE_MANAGER=TSM
   This setting will enable all external storage managers supported by Content Manager OnDemand. The Content Manager OnDemand server requires additional software to utilize Tivoli Storage Manager (TSM) as a storage manager. If that additional software is not installed, the server will not start when the ARS_STORAGE_MANAGER value is set to TSM.

   ARS_STORAGE_MANAGER= NO_TSM
   This setting will enable all external storage managers supported by Content Manager OnDemand except Tivoli Storage Manager. This setting is used when the additional software to support Tivoli Storage Manager is not installed and Tivoli Storage Manager is not required as an external storage manager.

Updating an instance configuration on Windows servers

Perform these steps to configure OpenStack Swift on a Windows server. Both steps use the OnDemand Configurator to create or update the configuration information.

1. Select `OpenStack Swift` as an external storage manager.
2. Set the configuration entries:

   Configuration Directory
   The Configuration Directory specifies the directory in which any alternate configuration files are kept. This directory is used if additional OpenStack Swift configuration files are defined. The names of these additional configuration files can be specified when defining storage nodes in Content Manager OnDemand. If no configuration file is specified in the storage node, the default configuration file is used. For example: `C:\Program Files\IBM\OnDemand\V10.1\config`

   Default Configuration File
   The Default Configuration File specifies an existing OpenStack Swift configuration file which the server uses by default. For example: `C:\Program Files\IBM\OnDemand\V10.1\config \ars.swift` A sample configuration file is included as part of the installation of Content Manager OnDemand.

Creating an OpenStack Swift configuration file

An OpenStack Swift configuration file for Content Manager OnDemand contains entries specific to your Swift implementation. You specify the location and name of the default configuration file in the ARS.CFG entry or via the OnDemand Configurator. Required entries must be specified. Optional entries are not required in the configuration file unless those values need to be changed.

The following list describes the entries that can be specified in a Swift configuration file.

ARS_SWIFT_SERVER
Specifies the OpenStack Swift server name. Do not include `http://` or `https://` in the name. This entry is required.

ARS_SWIFT_PORT
Specifies the OpenStack Swift server port number. This entry is optional if using a standard port. Content Manager OnDemand assumes port 80 for HTTP or port 443 for HTTPS communications.

ARS_SWIFT_TLD
Specifies the OpenStack Swift top-level directory name. This contains any additional path information after the server name and port in the URL. This entry is optional but usually necessary.

ARS_SWIFT_USE_SSL
Indicates whether or not to use SSL in server communications. The possible values are:

- `0` - SSL will not be used
- `1` - SSL will be used

The default value is 0. This entry is optional.

ARS_SWIFT_CONNECT_ TIMEOUT
Specifies the maximum number of seconds that Content Manager OnDemand waits for a response from the storage manager. The default is `60`. This entry is optional. **Warning:** Setting this value too low might cause connection failures.

ARS_SWIFT_HLD
Specifies a high-level directory name. This attribute is available to group sets of Content Manager OnDemand data together which might be needed if sharing external storage among multiple Content Manager OnDemand servers. **Warning:** Once this value is set, it must not be changed. If it is changed, any data that is already stored will not be retrievable. There is no default value. This entry is optional.

As an example, for a URL such as `https://swift.example.com:8088/v1/account/`, the Swift configuration file contains:

```
ARS_SWIFT_SERVER=swift.example.com
ARS_SWIFT_PORT=8088
ARS_SWIFT_TLD=/v1/account
ARS_SWIFT_USE_SSL=1
```

Defining an OpenStack Swift storage node with the Administrator client

You can define the settings for using the OpenStack Swift access method on the **Add a Primary Node** dialog of the OnDemand Administrator client.

The Storage Node field becomes Container Name when the Access Method is set to `OpenStack Swift`. The Container Name field is used with the OpenStack Swift server and determines the storage hierarchy for objects stored to this node. Containers are created if they do not already exist.

The Logon field is the OpenStack Swift user name which is used to store and retrieve data from the OpenStack Swift system. The password is also required.

The Access Method radio button is set to `OpenStack Swift`. For Content Manager OnDemand servers running on all platforms except Windows, the Configuration File Name defaults to the value specified by the `ARS_SWIFT_CONFIG_FILE` parameter in the ARS.CFG file if no value is entered. Otherwise, Content Manager OnDemand looks for the configuration file in the directory defined by the `ARS_SWIFT_CONFIG_DIR` parameter specified in the ARS.CFG file. For Content Manager OnDemand servers running on Windows, the server uses the Configuration File Name field and the Configuration Directory field that are specified in the OnDemand Configurator instead of using the ARS.CFG file parameters.

Using a file system for external storage

Content Manager OnDemand supports data storage in a file system repository. The file system must be locally accessible to the Content Manager OnDemand library or object server.

Updating the ARS.CFG file on AIX, Linux, Linux on System z, or Solaris servers

Perform these steps to configure Content Manager OnDemand to use a file system as external storage on an AIX, Linux, Linux on System z, or Solaris server.

1. Two new entries must be added to the ARS.CFG file.

For AIX and Solaris servers:

```
ARS_FILESYSTEM_CONFIG_FILE=/opt/IBM/ondemand/V10.1/config/ars.fs
ARS_FILESYSTEM_CONFIG_DIR=/opt/IBM/ondemand/V10.1/config
```

For Linux and Linux on System z servers:

```
ARS_FILESYSTEM_CONFIG_FILE=/opt/ibm/ondemand/V10.1/config/ars.fs
ARS_FILESYSTEM_CONFIG_DIR=/opt/ibm/ondemand/V10.1/config
```

The ARS_FILESYSTEM_CONFIG_FILE entry specifies an existing file system configuration file which the server uses by default.

The ARS_FILESYSTEM_CONFIG_DIR entry specifies the directory in which any alternate configuration files are kept. This directory is used if additional file system configuration files are defined. The names of these additional configuration files can be specified when defining storage nodes in Content Manager OnDemand. If no configuration file is specified in the storage node, the default configuration file is used.

The configuration file name and directory path shown in the examples are the recommended values for these entries.

2. The ARS_STORAGE_MANAGER entry in the ARS.CFG file might also need to be changed. If you specify ARS_STORAGE_MANAGER=CACHE_ONLY, this disables all storage managers supported by Content Manager OnDemand.

   To configure the Content Manager OnDemand server to use a file system as external storage, the value must be set to one of the following:

   ARS_STORAGE_MANAGER=TSM
   This setting will enable all external storage managers supported by Content Manager OnDemand. The Content Manager OnDemand server requires additional software to utilize Tivoli Storage Manager (TSM) as a storage manager. If that additional software is not installed, the server will not start when the ARS_STORAGE_MANAGER value is set to TSM.

   ARS_STORAGE_MANAGER= NO_TSM
   This setting will enable all external storage managers supported by Content Manager OnDemand except Tivoli Storage Manager. This setting is used when the additional software to support Tivoli Storage Manager is not installed and Tivoli Storage Manager is not required as an external storage manager.

Updating an instance configuration on Windows servers

Perform these steps to configure Content Manager OnDemand to use a file system as external storage on a Windows server. Both steps use the OnDemand Configurator to create or update the configuration information.

1. On the Storage tab, select Use File System under the External Storage Manager heading.
2. Set the configuration entries:

   Configuration Directory
   The Configuration Directory specifies the directory in which any alternate configuration files are kept. This directory is used if additional file system configuration files are defined. The names of these additional configuration files can be specified when defining storage nodes in Content Manager OnDemand. If no configuration file is specified in the storage node, the default configuration file is used. For example: C:\Program Files\IBM\OnDemand\V10.1\config

   Default Configuration File
   The Default Configuration File specifies an existing file system configuration file which the server uses by default. For example: C:\Program Files\IBM\OnDemand\V10.1\config\ars.fs

Creating a file system configuration file

A file system configuration file for Content Manager OnDemand contains entries specific to your file system implementation. You specify the location and name of the default configuration file in the ARS.CFG entry or via the OnDemand Configurator. Required entries must be specified. Optional entries are not required in the configuration file unless those values need to be changed.

The following list describes the entries that can be specified in a file system configuration file.

ARS_FILESYSTEM_NAME
Specifies the name of the file system for Content Manager OnDemand to use as a storage location. On a Windows server, this is a local directory, drive path, or UNC name. For all platforms, the file system permissions and ownership must be set to allow Content Manager OnDemand to read, write, and delete data to and from this location. This entry is required.

ARS_FILESYSTEM_HLD
Specifies the high-level directory name. This attribute is available to group sets of Content Manager OnDemand data together which might be needed if sharing this external storage among multiple Content Manager OnDemand servers. **Warning:** Once this value is set, it must not be changed. If it is changed, any data that is already stored may not be retrievable. There is no default value. This entry is optional.

As an example, on a Windows server, to use drive X, use the drive path:

```
ARS_FILESYSTEM_NAME=X:\
```

Defining a file system storage node with the Administrator client

You can define the settings for using the file system access method on the **Add a Primary Node** dialog of the OnDemand Administrator client.

The Storage Node field can be set to any name you choose. It is only used internally by Content Manager OnDemand.

The Logon and Password fields are not used.

The Access Method radio button is set to `Use a file system`. For Content Manager OnDemand servers running on all platforms except Windows, the Configuration File Name defaults to the value specified by the `ARS_FILESYSTEM_CONFIG_FILE` parameter in the ARS.CFG file if no value is entered. Otherwise, Content Manager OnDemand looks for the configuration file in the directory defined by the `ARS_FILESYSTEM_CONFIG_DIR` parameter specified in the ARS.CFG file. For Content Manager OnDemand servers running on Windows, the server uses the Configuration File Name field and the Configuration Directory field that are specified in the OnDemand Configurator instead of using the ARS.CFG file parameters.

Chapter 6. Preparing the system for use

Verify the installation of Content Manager OnDemand on all supported platforms.

Other important tasks described in this section include:

- Creating storage sets. You must add storage sets to the system before you can create application groups or assign the system-defined application groups to a storage set. Depending on the storage management characteristics of the reports that you plan to store on the system, you might need to add more than one storage set.
- Configuring the System Log application group. You should assign the System Log application group to a storage set that specifies a client node in storage managed by Tivoli Storage Manager so that the system can maintain a permanent copy of the system log data. You should also store system log index data in table spaces.
- Configuring the System Load application group. If you plan to use the system load logging facility provided by Content Manager OnDemand, then you must create a storage set that specifies a client node in storage managed by Tivoli Storage Manager. After you add the storage set to the system, you can assign the System Load application group to the storage set.
- Configuring the System Migration application group. If you plan to migrate index data from the database to archive storage, then you must create a storage set that specifies a client node in storage managed by Tivoli Storage Manager. After you add the storage set to the system, you can assign the System Migration application group to the storage set. You should also store system migration index data in table spaces.
- Creating a backup copy of the database. After installing and configuring the system, you should create a backup copy of the Content Manager OnDemand database. If you configured the system to use Tivoli Storage Manager, you should also backup the Tivoli Storage Manager database at this time.

Verifying the installation

After you have completed installation and configuration of the database manager, Content Manager OnDemand software, and Tivoli Storage Manager software, and have configured and initialized the system, perform the following tasks.

Procedure

1. Shut down and restart the system. Reinitializing the operating system starts the Content Manager OnDemand services.
2. If you have not already done so, install at least one of the Content Manager OnDemand client programs on a PC. See the *IBM Content Manager OnDemand: Client Installation Guide* for installation information.
3. Check the hardware and software requirements for all system components and features.
4. Start the Content Manager OnDemand client program. Content Manager OnDemand displays the **Logon to Server** dialog box.
5. Click **Update Servers**. Content Manager OnDemand displays the Update Servers dialog box.
6. Add the name of the Content Manager OnDemand library server.
7. Click **Close** to return to the **Logon to Server** dialog box.
8. Select the name of the server that you added in the **Update Servers** dialog box, if it is not already selected.
9. Type a Content Manager OnDemand user ID and password in the fields provided. (The first time that you log on to the system, you must specify the built-in Content Manager OnDemand userid, `admin`. Initially, there is no password. However, you will be prompted to enter and verify a password.)
10. Open and search the System Log folder.

Results

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

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

If the client program issues a message indicating a problem, follow the instructions in the message window. For more information about the messages, see *IBM Content Manager OnDemand: Messages and Codes*. If the problem persists, contact the IBM support center for help with resolving the problem.

Related reference

Next steps on AIX®
After you have installed the Content Manager OnDemand and related software on the system, configured the instance of Content Manager OnDemand, created the instance, and automated instance operations, you are now ready to verify the installation on Content Manager OnDemand.

Next step on Linux™
After you have installed Content Manager OnDemand and related software on the system, configured the instance of Content Manager OnDemand, created the instance, and automated instance operations, you are now ready to verify the installation on Content Manager OnDemand.

Next steps on Windows
After you have installed the Content Manager OnDemand and related software on the system, configured the instance of Content Manager OnDemand, created the instance, and automated instance operations, you are now ready to verify the installation on Content Manager OnDemand.

Define storage sets

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

About this task

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

A storage set must contain at least one primary storage node. A primary storage node can use cache storage (the default) or specify a client node in storage managed by Tivoli Storage Manager (or both). The administrative client online help provides details about defining storage sets and storage nodes.

If you plan to migrate index data to archive storage, you must assign the System Migration application group to a storage set that specifies a client node in storage managed by Tivoli Storage Manager. Also, you should assign the System Log and System Load application group to a storage set that specifies a client node in storage managed by Tivoli Storage Manager so that the system can maintain a permanent copy of the data that is written to the system log and system load log. The following topics provide additional details:

- “Configuring the System Log application group” on page 138
- “Configuring the System Load application group” on page 140
- “Configure the System Migration application group” on page 142

Configuring the System Log application group

When you install and configure Content Manager OnDemand, you initialize the system log.

The system log comprises the System Log application group, a set of system log applications, and the System Log folder. The System Log application group contains the storage management information that

Content Manager OnDemand uses to maintain the data written to the system log. When you initialize the system, the application group is not assigned to a storage set. Because the application group is not assigned to a storage set, the system does not maintain a permanent copy of the system log data.

Before you begin defining reports to Content Manager OnDemand, loading data on the system, or allowing users to access the system, you should configure the System Log application group to maintain a permanent copy of the data that is written to the system log. You can do this by first defining a storage set that specifies a client node in storage managed by Tivoli Storage Manager and then by updating the System Log application group and assigning it to the storage set.

If your system does not use Tivoli Storage Manager, you should assign the System Log application group to a cache-only storage set and change the length of time that Content Manager OnDemand maintains the system log data to the maximum permitted value.

If you define table space file systems to Content Manager OnDemand, you should store the system log data in table spaces.

Maintaining system log data in archive storage

You should create a storage set that specifies a client node in storage managed by Tivoli Storage Manager. You must add at least one primary storage node to the storage set.

About this task

The primary storage node must identify a client node in Tivoli Storage Manager that maintains data indefinitely. The Logon Name and Password in the primary storage node must be identical to the client node and password in Tivoli Storage Manager.

When you register the client node in Tivoli Storage Manager, you must specify the name of the Tivoli Storage Manager policy domain that maintains data on the required media for the required length of time. If you do not specify the name of a domain, Tivoli Storage Manager assigns the client node to the default domain. Individual chapters for each platform provide details about defining storage devices to Tivoli Storage Manager and policy domains to support Content Manager OnDemand and describe how to register a client node in Tivoli Storage Manager.

After you create the storage set, you must update the System Log application group and assign it to the storage set. After assigning the application group to the storage set and restarting the system, the system automatically maintains a copy of the system log data in archive storage.

Complete the following steps to assign the System Log application group to a storage set. The administrative client online help provides information about the options on the Storage Management page.

Procedure

1. Start the administrative client.
2. Log on to the server with a userid that has system administrator authority. (The built-in userid `admin` has system administrator authority.)
3. Click **Application Groups**.
4. Point to the **System Log application group** and right-mouse click.
5. From the pop-up menu, select **Update** to open the **Update an Application Group** window.
6. Click the **Storage Management** tab.
7. In the Storage Set Name list, select the name of the storage set. The storage set that you select should specify a client node in a Tivoli Storage Manager policy domain that maintains data indefinitely.
8. Click **Advanced** to open the Advanced Storage Management dialog box.
9. Select **Next Cache Migration** under Migrate Data from Cache. This causes Content Manager OnDemand to copy the system log data to archive storage the next time that the `ARSMAINT` command runs.

Maintaining system log data in cache storage

If your system does not use Tivoli Storage Manager, you should assign the System Log application group to a cache-only storage set and change the length of time that Content Manager OnDemand maintains the data to the maximum permitted value.

About this task

Ensure that Content Manager OnDemand does not delete the data from cache storage for a very long time.

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

Procedure

Complete the following steps to configure the System Log application group:

1. Start the administrative client.
2. Log on to the server with a userid that has system administrator authority. (The built-in userid `admin` has system administrator authority.)
3. Click **Application Groups**.
4. Point to the **System Log** application group and right-mouse click.
5. From the pop-up menu, select **Update** to open the **Update an Application Group** window.
6. Click the **Storage Management** tab.
7. In the Storage Set Name list, select the name of the storage set.

   The storage set named `Cache Only - Library Server` is a cache-only storage set created on the library server when you initialized the system.
8. Replace the contents of the Cache Data for ___ Days field with 99999.

   This value causes Content Manager OnDemand to maintain data for approximately 273 years.

Storing system log data in table spaces

If you define table space file systems to Content Manager OnDemand, you should configure the System Log application group to store index data in table spaces.

Procedure

To update the System Log application group:

1. Start the administrative client.
2. Log on to the server with a userid that has system administrator authority. (The built-in userid `admin` has system administrator authority.)
3. Click **Application Groups**.
4. Point to the System Log application group and right-mouse click.
5. From the pop-up menu, select **Update** to open the **Update an Application Group** window.
6. On the General page, select **Advanced** to open the **Database Information** dialog box.
7. Under Create Tablespace Type, select **SMS**.

Configuring the System Load application group

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

The system load logging facility comprises the System Load application group, a system load application, and the System Load folder. The System Load application group contains the storage management information that Content Manager OnDemand uses to maintain the data written to the system load logging

facility. When you initialize the system, the application group is not assigned to a storage set. Because the application group is not assigned to a storage set, the system does not maintain a permanent copy of the system load data.

Before you begin defining reports to Content Manager OnDemand, loading data on the system, or allowing users to access the system, you should configure the System Load application group to maintain a permanent copy of the data that is written to the system load logging facility. You can do this by first defining a storage set that specifies a client node in storage managed by Tivoli Storage Manager and then by updating the System Load application group and assigning it to the storage set.

If your system does not use Tivoli Storage Manager, you should assign the System Load application group to a cache-only storage set and change the length of time that Content Manager OnDemand maintains the system load data to the maximum permitted value.

If you define table space file systems to Content Manager OnDemand, you should store the system load data in table spaces.

Maintaining system load data in archive storage

You should create a storage set that specifies a client node in storage managed by Tivoli Storage Manager. You must add at least one primary storage node to the storage set.

About this task

The primary storage node must identify a client node in Tivoli Storage Manager that maintains data indefinitely. The Logon Name and Password in the primary storage node must be identical to the client node and password in Tivoli Storage Manager.

When you register the client node in Tivoli Storage Manager, you must specify the name of the Tivoli Storage Manager policy domain that maintains data on the required media for the required length of time. If you do not specify the name of a domain, Tivoli Storage Manager assigns the client node to the default domain. Individual chapters for each platform provide details about defining storage devices to Tivoli Storage Manager and policy domains to support Content Manager OnDemand and describe how to register a client node in Tivoli Storage Manager.

After you create the storage set, you must update the System Load application group and assign it to the storage set. After assigning the application group to the storage set and restarting the system, the system automatically maintains a copy of the system load data in archive storage.

Procedure

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

1. Start the administrative client.
2. Log on to the server with a userid that has system administrator authority. (The built-in userid `admin` has system administrator authority.)
3. Click **Application Groups**.
4. Point to the System Load application group and right mouse click.
5. From the pop-up menu, select **Update** to open the **Update an Application Group** window.
6. Click the **Storage Management** tab.
7. In the Storage Set Name list, select the name of the storage set.

   The storage set that you select should specify a client node in a Tivoli Storage Manager policy domain that maintains data indefinitely.
8. Click **Advanced** to open the **Advanced Storage Management** dialog box.
9. Select **Next Cache Migration** under Migrate Data from Cache.

   This causes Content Manager OnDemand to copy the system load data to archive storage the next time that the ARSMAINT command runs.

Maintaining system load data in cache storage

If your system does not use Tivoli Storage Manager, you should assign the System Load application group to a cache-only storage set and change the length of time that Content Manager OnDemand maintains the data to the maximum permitted value.

About this task

Doing so ensures that Content Manager OnDemand does not delete the data from cache storage for a very long time.

Procedure

Complete the following steps to configure the System Load application group:

1. Start the administrative client.
2. Log on to the server with a userid that has system administrator authority. (The built-in userid `admin` has system administrator authority.)
3. Click **Application Groups**.
4. Point to the System Load application group and right mouse click.
5. From the pop-up menu, select **Update** to open the **Update an Application Group** window.
6. Click the **Storage Management** tab.
7. In the Storage Set Name list, select the name of the storage set.

   The storage set named `Cache Only - Library Server` is a cache-only storage set created on the library server when you initialized the system.
8. Replace the contents of the Cache Data for ___ Days field with 99999.

   This value causes Content Manager OnDemand to maintain data for approximately 273 years.

Storing system load data in table spaces

If you define table space file systems to Content Manager OnDemand, you should configure the System Load application group to store index data in table spaces.

Procedure

To update the System Load application group:

1. Start the administrative client.
2. Log on to the server with a userid that has system administrator authority. (The built-in userid `admin` has system administrator authority.)
3. Click **Application Groups**.
4. Point to the System Load application group and right mouse click.
5. From the pop-up menu, select Update to open the Update an Application Group window.
6. On the General page, select **Advanced** to open the **Database Information** dialog box.
7. Under Create Tablespace Type, select **SMS**.

Configure the System Migration application group

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

About this task

This process optimizes database storage space while allowing you to maintain index data for a very long time. You typically migrate index data after users no longer need to access the reports, but for legal or other requirements, you still need to maintain the data for some number of years or months. Content Manager OnDemand uses the storage management settings in application groups to determine whether

or not to migrate index data to archive storage. All migrated data is managed through the System Migration application group.

When you install and configure Content Manager OnDemand, you can optionally initialize the system migration function. The system migration function comprises the System Migration application group, a system migration application, and the System Migration folder. The System Migration application group contains the storage management information that Content Manager OnDemand uses to maintain index data migrated to archive storage. Until you assign the application group to a storage set that specifies a client node in storage managed by Tivoli Storage Manager, Content Manager OnDemand cannot migrate index data from the database to archive storage.

If you define table space file systems to Content Manager OnDemand, you should store system migration data in table spaces.

Assigning the System Migration application group to a storage set

If you need the system to maintain index data in archive storage, you must assign the System Migration application group to a storage set that identifies a client node in Tivoli Storage Manager.

About this task

You must register the client node in a Tivoli Storage Manager policy domain that maintains data indefinitely. When you define the storage set, the Logon Name and Password of the primary storage node must be identical to the client node and password in Tivoli Storage Manager. Individual chapters for each platform provide details about defining storage devices and policy domains to support Content Manager OnDemand and show how to register a client node in Tivoli Storage Manager.

After you define the storage set, you must update the System Migration application group and assign it to the storage set. After assigning the application group to the storage set and restarting the system, the system automatically migrates index data to archive storage, whenever migration processing (the ARSMAINT program) runs.

Procedure

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

1. Start the administrative client.
2. Log on to the server with a userid that has system administrator authority. (The built-in userid `admin` has system administrator authority.)
3. Click **Application Groups**.
4. Point to the System Migration application group and right mouse click.
5. From the pop-up menu, select **Update** to open the **Update an Application Group** window.
6. Click the **Storage Management** tab.
7. In the Storage Set Name list, select the name of the storage set.

   The storage set that you select should identify a client node in a Tivoli Storage Manager policy domain that maintains data indefinitely.

Storing system migration data in table spaces

If you define table space file systems to Content Manager OnDemand, you should configure the System Migration application group to store data in table spaces.

Procedure

To update the system migration application group:

1. Start the administrative client.
2. Log on to the server with a userid that has system administrator authority. (The built-in userid `admin` has system administrator authority.)
3. Click **Application Groups**.

4. Point to the System Migration application group and right mouse click.
5. From the pop-up menu, select **Update** to open the **Update an Application Group** window.
6. On the General page, select **Advanced** to open the **Database Information** dialog box.
7. Under Create Tablespace Type, select **SMS**.

Chapter 7. Backing up the Content Manager OnDemand database

To complete the installation and configuration process, you should run the ARSDB program to create a full, offline backup image of the Content Manager OnDemand database on removable media.

About this task

A full backup image of the database is required to rebuild the database, in the event that you need to do so. You cannot rebuild the database unless you have a full database backup (and any table space backups and log files generated since the last full database backup). See the *IBM Content Manager OnDemand for Multiplatforms: Administration Guide* for more information about the ARSDB program, including the backup options.

Procedure

Complete the following steps to create a full, offline backup of the database with the ARSDB program:

1. Log on with the proper authority: the `root` user on a UNIX server, or an administrator on a Windows server.
2. Make sure there are no other users logged on to the library server.
3. Make sure there are no other applications connected to the database.
4. If you plan to backup the database to tape, place a blank, formatted tape storage volume in the tape drive.
5. On a UNIX server, change to the Content Manager OnDemand program directory. On a Windows server, select **Start** > **Programs** > **IBM OnDemand Server V10.1** > **Command Window V10.1 for Windows**.
6. Enter the backup command.
   For example:

   ```
   arsdb -v -y <device>
   ```

   Replace the string `<device>` with the name of the output device
7. Write down the information about the database backup, including the date and time that the backup was taken and the label of the storage volume. Keep the backup copy in a safe location, preferably offsite. Save the backup copy at least until the next time that you create another full backup image of the database.

Chapter 8. Silently installing Content Manager OnDemand

The installation programs for Content Manager OnDemand for Multiplatforms are based on InstallAnywhere.

About this task

You cannot install Content Manager OnDemand 10.1 silently without a response to accept the license agreement.

To generate a response file, you can use the `-r` command line switch that is followed by the response file name. Run the installation program normally until it finishes recording the responses. For example:

```
./odaix.bin -r /tmp/my_response_file
```

To use the generated response in a silent installation, you can use the `-f` command line switch followed by the response file name. For example:

```
./odaix.bin -f /tmp/my_response_file -i silent
```

Chapter 9. Uninstalling Content Manager OnDemand

The Content Manager OnDemand Version 10.1 installation detects previous versions of Content Manager OnDemand on the same system.

About this task

To uninstall Content Manager OnDemand on AIX, use the following command:

```
/opt/IBM/ondemand/V10.1/_uninst1010/uninstallod
```

or

```
/opt/IBM/ondemand/V10.1/_uninst1010/uninstallod -i console
```

To uninstall Content Manager OnDemand on Linux or Linux on zSeries, use the following command:

```
/opt/ibm/ondemand/V10.1/_uninst1010/uninstallod
```

or

```
/opt/ibm/ondemand/V10.1/_uninst1010/uninstallod -i console
```

You can also use the Java command to uninstall Content Manager OnDemand.

Important: Use the Java command only if you cannot use the above commands because the system cannot find a suitable JVM.

On AIX, enter this command:

```
java -jar /opt/IBM/ondemand/V10.1/_uninst1010/uninstallod.jar
```

or

```
java -jar /opt/IBM/ondemand/V10.1/_uninst1010/uninstallod.jar -i console
```

On Linux or Linux on zSeries, enter this command:

```
java -jar /opt/ibm/ondemand/V10.1/_uninst1010/uninstallod.jar
```

or

```
java -jar /opt/ibm/ondemand/V10.1/_uninst1010/uninstallod.jar -i console
```

To uninstall the Content Manager OnDemand Enhanced Retention Management feature on AIX, Linux, or Linux on zSeries, use the following command:

```
<CMOD server directory>/_uninst1010erm/uninstalloderm
```

or

```
<CMOD server directory>/_uninst1010erm/uninstalloderm -console
```

To uninstall the Content Manager OnDemand PDF Indexer feature on AIX, Linux, or Linux on zSeries, use the following command:

```
 <CMOD server directory>/_uninst1010pdf/uninstallodpdf
```

or

```
<CMOD server directory>/_uninst1010pdf/uninstallodpdf -console
```

To uninstall the Content Manager OnDemand Distribution Facility feature on AIX, Linux, or Linux on zSeries, use the following command:

```
 <CMOD server directory>/_uninst1010odf/uninstallododf
```

or

```
<CMOD server directory>/_uninst1010odf/uninstallododf -console
```

On Windows, use the **Program and Features** from the **Control Panel** to uninstall any or all of the following features:

- IBM OnDemand V10.1

  Remember: Uninstalling IBM OnDemand V10.1 will also uninstall the following features:

  - IBM OnDemand Enhanced Retention Management Feature V10.1
  - IBM OnDemand PDF Indexing Feature V10.1
  - IBM OnDemand Distribution Feature V10.1
- IBM OnDemand Full Text Search Server Feature V10.1

Chapter 10. User exit programming

User exits provided by Content Manager OnDemand are specific points in the program where an experienced programmer can specify processing routines to enhance or replace the default Content Manager OnDemand functions.

For example, the security user exit provides a point on the library server where you can identify and authenticate users that log on to the system. Programmers require a working knowledge of the tools needed to develop a user exit program. The following list identifies the main skills and tools that are needed:

- Skills
  - C and C++ programming
  - Operating system programming
  - Experience with relational database technology
  - Knowledge of compiling and linking programs in the C, C++, and operating system environment
  - DB2 UDB, Oracle, or SQL Server (if writing your own SQL code)
- Tools
  - IDE
  - C or C++ compiler

.A makefile is provided for use with the following compilers:

- **AIX:** IBM XL C/C++ Compiler, Version 11.01.0000.0004
- **Windows:** Visual Studio Version 10.0.40219.1 SP1 Rel
- **Sun:** Sun C 5.11 SunOS_sparc 2010/08/13
- **Linux/Intel:** gcc version 4.1.2 20070115 (SUSE Linux)
- **Linux on System z®:** gcc version 4.1.2 20070115 (SUSE Linux)

C/C++ compilers used outside of this list must provide equivalent compile options as documented in the makefile.

Download user exit

When processing files transmitted by Download, unless you specify otherwise, the ARSLOAD program uses a part of the name of the file that is saved on the server to identify the application group to load.

If the ARSLOAD program cannot determine the correct application group to load from the file name, then the load will fail. If the application group to load contains more than one application, then you must identify the application to load. Otherwise, the load will fail.

Content Manager OnDemand provides ways for you to identify the correct application group and application to load. For example, you can run the ARSLOAD program with the `-A` and `-G` parameters to specify the parts of the file name that identify the application group and application. However, if the file name does not contain information that can be used to identify the application group and application, then you must use some other method to determine the correct application group and application before the load process can proceed.

Using Download

You can use Download to transmit reports from z/OS systems to Content Manager OnDemand servers.

If you use Download to automate the data loading process, then you may need to provide a user-written program to process the files transmitted by Download before they can be processed by the ARSLOAD

program. You must provide a user-written program if the file name does not contain information that can be used to identify the application group to load.

For example, suppose that you use a report distribution system to place the output of your application programs on the spool data set. Download selects the output and transmits the data to a Content Manager OnDemand server. However, after the output has been processed by the report distribution system, the resulting file name on the spool data set can no longer be associated with the application program that generated the output. Therefore, the ARSLOAD program cannot use the file name to determine the application group to load. (And because the file name does not contain information that can be use used to identify the application group and application to load, you cannot run the ARSLOAD program with the -A and -G parameters.) You must find some other way to identify the application group and application to load.

Download provides a user exit (APSUX15) that allows you to provide additional job information to Content Manager OnDemand. Download includes the additional job information in the data stream that is transmitted from the spool data set to the server. See *PSF for z/OS: Download for z/OS* for detailed information about Download, the user exit, and the additional job information that can be included in the data that is transmitted to the server.

On the Content Manager OnDemand server, the ARSJESD program provides the -x parameter so that you can run a user-written program to process the additional job information after Download successfully transmits a file to the server. See the ARSJESD program reference in *IBM Content Manager OnDemand for Multiplatforms: Administration Guide* for information about using the -x parameter.

Invoking the Download user exit

When Download selects output data from the spool data set for transmission to a Content Manager OnDemand server, it invokes the APSUX15 user exit program.

The user exit program concatenates a string of additional job information to the print parameters that Download transmits to the server. Upon completion, the user exit program passes the location of the string and the string length to Download, which transmits the output data set with associated JCL and the additional job information to the server.

The ARSJESD program receives the data sets into file systems on the server. If you start the ARSJESD program with the -x parameter, the ARSJESD program invokes the specified user-written program. The program specified with the -x parameter can be any user-written program.

For example, you could provide a user-written program that parses the additional job information transmitted by Download and the APSUX15 user exit program. The user-written program could extract the value of the WRITER parameter to identify the application to load. Using this value, the user-written program could then query the database to determine the name of the application group to which the application belongs. The user-written program could then run the ARSLOAD program with the -a parameter to identify the application to load and the -g parameter to identify the application group to load. (The user-written program could also rename the input file and then run the ARSLOAD program with the -A parameter to specify the part of the new file name that identifies the application and the -G parameter to specify the part of the new file name that identifies the application group.)

By using the Download user exit program, the -x parameter with the ARSJESD program, and a user-written program, you can configure the system so that each file that Download transmits to the server is automatically processed and loaded into the correct application group and application.

ARSJESD processing

The ARSJESD program is the component of Download for the z/OS feature that runs on the workstation. The -x parameter of the ARSJESD program may be used to specify the name of a user-written program to process additional job information sent by PSF through the APSUX15 user exit.

If the ARSJESD program was invoked with the -x parameter, it calls the specified user-written program. The ARSJESD program passes the file name and the additional job information to the user-written program.

The additional job information is installation dependent. See *IBM Print Services Facility™ for z/OS: Download for z/OS*, S550-0429-01 for details about the APSUX15 user exit and the content, format, and

purpose of the additional job information. The processing done by the user-written program is also installation dependent. See your Infoprint Manager or PSF information for information about processing the additional job information with a user-written program.

Sample user exit program

In Content Manager OnDemand, it is possible to use the ARSJESD exit point in order to customize and enhance the standard functionality within the product. A user exit is a point during processing that enables you to run a user-written program and return control of processing after your user-written program ends.

The ARSJESD exit point is a server exit. Content Manager OnDemand provides data at the ARSJESD exit that serves as input to the user-written program. The data consists of the file name and the additional job information. Using this exit, it is possible to do functions such as parse the additional job information that is sent by PSF and rename the input file by using one of the PSF parameters.

The program invoked at the ARSJESD exit point is defined by specifying the –x *userProgram* parameter when starting the ARSJESD program, where *userProgram* is the name or full path name of the user-written program.

The following example demonstrates a common use for the ARSJESD exit point. The example processes the file name and the additional job information. The example renames the input file by using the value of the WRITER parameter that was sent by PSF. For the sake of simplicity, the sample is not demonstrated across all of the supported platforms. IBM recognizes that the scripting languages between platforms do vary, but the principles which are described here are uniform across all supported platforms; only the syntax differs. The sample is provided for an AIX system.

The example program is provided on an as-is basis. The licensee of the Content Manager OnDemand product is free to copy, revise modify, and make derivative works of this program sample as they see fit.

```
#!/bin/ksh
#
#  COPYRIGHT:  5697-G34 (C) COPYRIGHT IBM CORPORATION 2003
#              ALL RIGHTS RESERVED
#              LICENSED MATERIALS - PROPERTY OF IBM
#
#              US Government Users Restricted Rights - Use, duplication or
#              disclosure restricted by GSA ADP Schedule Contract with IBM Corp.
#
#  MODULE:    writer_rename
#
#  DESCRIPTION:
#    This Korn shell script is designed as a exit routine to be executed by
#    the ARSJESD program after a file has been downloaded from z/OS.
#    This script parses and prints the additional job options that were sent
#    by PSF and renames the file name by using the value of the WRITER parameter.
#
#
# set -x
#
ME=$(basename $0)
PRODUCT="OnDemand"
PRODUCT_DIR=/opt/IBM/ondemand/V10.1
ARS_BIN_DIR=${PRODUCT_DIR}/bin
AWK=/bin/awk
DATE=/bin/date

status_msg ()
{

  print "\n$(${DATE}) -- ${ME} -- \nERROR: Unable to determine WRITER field"
  print "\t\t$1"

  return

}

parse_objects ()
{
   while getopts o: i
   do
      print ${OPTARG} | ${AWK} -F [=] '{print $1}' | read option
      case $option in
         pa) value=${OPTARG#*=}
             while [ "$value" != "$oldvalue" ]
             do
                oldvalue=$value
                arg=${value%%=*}
                value=${value#*=}
                case $arg in
```

```
                    segmentid   ) SEGID=${value%%,*} ;;
                    forms       ) FORM=${value%%,*} ;;
                    class       ) CLASS=${value%%,*} ;;
                    destination ) DEST=${value%%,*} ;;
                    writer      ) WRITER=${value%%,*} ;;
                 esac
                 value=${value#*,}
              done ;;
          *) print ${OPTARG} | ${AWK} -F [=] '{print $1, $2}'| read arg value
             case $arg in
                fileformat ) FORMAT="$value" ;;
                datat      ) DATAT="$value" ;;
                chars      ) CHARS="$value" ;;
                cc         ) CC="$value" ;;
                cctype     ) CCTYPE="$value" ;;
                pagedef    ) PAGEDEF="$value" ;;
                prmode     ) PRMODE="$value" ;;
                trc        ) TRC="$value" ;;
                cop        ) COPIES="$value" ;;
                datac      ) DATAC="$value" ;;
                f          ) FORMDEF="$value" ;;
                outbin     ) OUTBIN="$value" ;;
                jobn       ) JOBNAME="$value" ;;
                us         ) USER="$value" ;;
                no         ) MVSNAME="$value" ;;
                pr         ) TEXT="$value" ;;
                address1   ) ADDRESS1="$value" ;;
                address2   ) ADDRESS2="$value" ;;
                address3   ) ADDRESS3="$value" ;;
                address4   ) ADDRESS4="$value" ;;
                bu         ) BLDG="$value" ;;
                de         ) DEPT="$value" ;;
                na         ) NAME="$value" ;;
                ro         ) RO="$value" ;;
                ti         ) TITLE="$value" ;;
             esac ;;

       esac

   done
}

FILENAME="$1"
shift
print "\n$(${DATE}) -- ${ME} -- \nStart writer_move for file ${FILENAME}\n"
print "File Name = ${FILENAME}"

print ${FILENAME} | awk -F [.] '{print $1, $2, $3, $4, $5, $6, $7}' | read MVS JOBNAME DATA SET FORMS YYDDD
HHMMM EXT

PARMS=$(print $@ | tr -d '\"' )
print "Parameters = ${PARMS}"

parse_objects ${PARMS}

APPGRP=${WRITER}

if [ "${WRITER}" != "" ]
then
   status_msg "WRITER field specified."
   NEWFILENAME="${MVS}.${JOBNAME}.${WRITER}.${FORMS}.${YYDDD}.${HHMMM}.${EXT}"
else
   status_msg "WRITER field not specified."
   NEWFILENAME="${MVS}.${JOBNAME}.${DATASET}.${FORMS}.${YYDDD}.${HHMMM}.NOWTR"
fi

print "File renamed to ${NEWFILENAME}"

mv ${FILENAME} ${NEWFILENAME}
```

Report specifications archive definition exit

The Report Specifications Archive Definition exit allows an installation to 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.

Interface exit components

The Report Specifications Archive Definition exit interface consists of the ARSUUPDT program, which is provided in source form in the `exits` directory.

ARSUUPDT is a DLL module written in the C programming language. It is provided in both source and executable forms, with the source being provided mainly in understanding how the exit is driven.

The Report Specifications Archive Definition exit is implemented by a single DLL - ARSUUPDT. The sample shipped with the product initializes the data structure and calls the exit driver.

ARSUUPDT DLL

When the ARSLOAD program loads a document, it makes two calls of the DLL.

The first (called Names) allows the exit to modify specifications for:

- Application Group name
- Application name
- Object Server Name
- Storage Node Name

The second call (Parameters) allows the exit to modify the following parameters:

- Indexer Parameters
- Viewer parameters:
  - Logical Record length
  - Record Format
  - Control character type

The ARSUUPDT DLL must reside in the `/opt/IBM/ondemand/V10.1/bin/exits` directory. To call the exit, you must specify the `-E` parameter when you run the ARSLOAD program.

C language ARSUUPDT

The Format example shows the C language ARSUUPDT.

C language ARSUUPDT sample:

```
#pragma export(UPDTEXIT)
#include "arscsxit.h"

int UPDTEXIT( ArsCSXitUpdtExit  updt )
```

Function field

The updt->Function field should be interrogated to determine the which type of call is being made.

Names

ARCCSXIT_PROCESS_NAMES

Parameters

ARCCSXIT_PROCESS_PARMS

Do not attempt to use parameters that are not specified as being valid for a given call.

Table 9: Names call parameters

Parameter	Meaning
updt->pFileName	Address of the null delimited file name.

Table 9: Names call parameters (continued)	
Parameter	**Meaning**
updt->ApplGrpName	Null delimited application group name. This is the application group name that ARSLOAD will attempt to use if no action is performed by the Names call.
updt->ApplName	Null delimited application name. This is the application name that ARSLOAD will attempt to use if no action is performed by the Names call.
updt->ObjServer	Null delimited object server name. This is the object server that ARSLOAD will attempt to use if no action is performed by the Names call. If no object server is explicitly specified, ARSLOAD will use the object server specified by the storage node that is designated for loading in the storage set assigned to the application group.
updt->StorageNode	Null delimited storage node name. This is the storage node that ARSLOAD will attempt to use if no action is performed by the Names call. If no node is explicitly specified, ARSLOAD will use the primary storage node specified by the storage node that is designated for loading in the storage set assigned to the application group.
updt->ArsCSXitUpdtExit-pJES	Pointer to the JES information. If this pointer is not null, the object being loaded is being read from the JES SPOOL. The JES information contains a DDNAME that is currently allocated to the SPOOL file and a pointer to the JES SSS2 SSOB extension. If null, the file being processed is either an MVS™ data set or an HFS file.

Table 10: Parameters call parameters	
Parameter	**Meaning**
updt->pFileName	Address of the null delimited file name.
updt->ApplGrpName	Null delimited application group name. This is the application group name the document will be stored under.
updt->ApplName	Null delimited application name. This is the application name the document will be stored under.
updt->ObjServer	Null delimited object server name. This is the object server that will be stored in. If the object server is X'00', ARSLOAD will use the object server specified by the storage node that is designated for loading in the storage set assigned to the application group.

Table 10: Parameters call parameters (continued)	
Parameter	**Meaning**
updt->StorageNode	Null delimited storage node name. This is the storage node that ARSLOAD will be stored in. If the node is X'00', ARSLOAD will use the primary storage node that is specified by the storage node that is designated for loading in the storage set that is assigned to the application group.
updt->ArsCSXitUpdtExit-pJES	Pointer to the JES information. If this pointer is not null, the object being loaded is being read from the JES SPOOL. The JES information contains a DDNAME that is currently allocated to the SPOOL file and a pointer to the JES SSS2 SSOB extension. If null, the file being processed is either an MVS data set or an HFS file.
updt->IndexerParms	The indexer parameter that will be passed to the indexer. The indexers parameters are a series of records separated by newline characters (X'15'). If altered by the exit, these parameters will be stored back into indexer parameters that are associated with the application.
updt->CCType	Carriage control type. If altered by the exit, this value will be stored in the view information that is associated with the application. Changing this value after reports have been loaded may cause previously loaded documents to display incorrectly.
updt->LRECL	For fixed record format the length of each line. This has no meaning for variable format. If altered by the exit, this value will be stored in the view information that is associated with the application. Changing this value after reports have been loaded may cause previously loaded documents to display incorrectly.
updt->RECFM	The record format the document is stored in OnDemand. If altered by the exit, this value will be stored in the view information that is associated with the application. Changing this value after reports have been loaded may cause previously loaded documents to display incorrectly.
updt->update_appl	Indicates that the application is to be updated. If zero, any changes made by the exit will not be reflected in the application definitions. If non-zero, the application will be updated with the new values.
updt->Delim	For documents stored as stream, this is the NULL terminated string that contains the string that is used to determine record boundaries. For example, records that use EBCDIC newline characters as record delimiters would specify X'15'.

Valid values for record format are described in this table.

Table 11: Record formats	
Format	**Value**
Fixed	ARCCSXIT_DOC_FORMAT_FIXED
Variable	ARCCSXIT_DOC_FORMAT_VARIABLE
Stream	ARCCSXIT_DOC_FORMAT_STREAM

Valid values for carriage control are described in this table.

Table 12: Carriage control	
Format	**Value**
ANSI	ARCCSXIT_CC_ANSI
Machine	ARCCSXIT_CC_MACHINE
None	ARCCSXIT_CC_NONE

1. Carriage control, lrecl, recfm, and record delimiter are only valid for documents stored as Line data.
2. Documents stored as variable in OnDemand are stored with a two-byte length prefix followed by the data for the record. The length does not include the two-byte prefix.

Returned values
In addition to updating the ArsCSXitPrepExit as appropriate, the DLL should set a return code of zero to indicate success, and a non-zero to indicate failure.

Retrieval preview user exit

The retrieval preview user exit allows for the modification of document data prior to the data being retrieved from the server. The retrieval preview user exit allows you to run a user-written program to process documents that belong to a specified application.

The user-written program is activated by selecting the Use Preview Exit option on the Miscellaneous Options page of an application. When the option is selected, the user-written program will be called any time that a request is made to retrieve a document. Any information that is specified in the Parameters field will be passed to the user-written program.

The retrieval preview user exit allows an installation to process document data before the document is presented to the client. The retrieval 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 ("white out") 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.

Programming considerations

The retrieval preview user exit is not called for all document retrievals. In particular, the user exit is not called for functions that use the so-called Bulk Retrieval method of retrieving documents or for server printing.

For example, running the ARSDOC GET function without specifying the `-n` parameter performs a bulk retrieval, and documents retrieved will not be presented to the client preview exit.

If a request is made to retrieve a large object document, care should be taken to make certain that the retrieval preview user exit does not remove any pages from the document. The large object segment size and page navigation information are based on the number of pages that existed when the document was loaded on the server. Unexpected results may occur if this information is changed. The retrieval preview user exit may be enabled for all data types, except for None.

The user-written program can be used to remove sensitive information from a document or to perform other types of data manipulation. However, because the user-written program is not used during server reprinting or bulk retrievals, the restricted data may still be accessible by the user.

When modifying the data, the format and type of the data must not be changed; only the content may be changed. For example:

- If the format of the data is EBCDIC data with a fixed length of 133 bytes, the format must not be changed to something different, such as ASCII data delimited by the line feed character (X'0A')
- If the data type is AFP, the document may not be converted to some other type of data, such as PDF

When the modified data is viewed by the Windows client, the format of the data and the data type that is defined in the application on the View Information page will be used to display the data. If the format or data type has changed, the document will not view properly.

The retrieval user exit point may be enabled for more than one application. However, all applications must be processed by the same user-written program (only one user-written program is supported). The system passes the name of the application that is associated with the document to the user-written program. The user-written program can perform processing based on the application or it can perform the same processing for all documents regardless of the application.

A sample user exit program arsuprep.c and header file arscsxit.h are provided by IBM. The files are located in the `/opt/IBM/ondemand/V10.1/exits` directory on AIX, in the `/opt/ibm/ondemand/V10.1/exits` directory on Linux, and in the `\Program Files\IBM\OnDemand\V10.1\exits` directory on Windows.

Return values

If the exit wants a different file presented to the user, it should set `prep->OutFileName` to the name of the file. This file must be formatted to agree with the specification in ArcCSXitAppl. For example, the file cannot be in variable format if the ArcCSXitAppl indicates fixed.

This file is deleted when its contents are updated to the user. The exit should set a return code of 0 (zero).

Security user exit

Content Manager OnDemand provides a user exit that allows you to implement your own user exit program to identify and authenticate users that log on to the system.

You can use the security user exit to authenticate a user's password by some means other than the way that is built in to Content Manager OnDemand. For example, you may want to deny access to the system after three incorrect logon attempts are made by a user; you may want to enforce some sort of password uniqueness; and so forth. You can also use the security user exit to allow users that are not already in the Content Manager OnDemand user database to access the system.

The security user exit allows you to augment the security related processing of the following activities or events:

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

When driven for these activities, a user-written exit routine (or set of exit routines) can interact with some other security system to determine if the given activity is to be allowed or disallowed.

The security user exit runs the ARSUSEC program when a user attempts to logon to the system. A sample C program is provided in the EXITS directory. To implement your own security user exit program, you should add your specific code to the sample provided (for example, you could call another program from the ARSUSEC program). See the ARSCSXIT.H file for information about functions, parameters, and return codes. You then compile the ARSUSEC program (a Makefile is provided) and move or copy the executable program to the BIN directory. Then restart the library server to begin using the security user exit program.

Important: When you implement your own security user exit program, you bypass the logon verification processing that is built into the base Content Manager OnDemand product. IBM advises caution when you bypass the Content Manager OnDemand user and password restrictions. The security of the system could easily be subverted by malicious or defective code. Only use code that you trust.

Sample security user exit program

The following example demonstrates how to use the security user exit. The example program prevents users from changing their passwords. For the sake of simplicity, the sample is not demonstrated across all of the supported platforms.

IBM recognizes that the scripting languages between platforms do vary, but the principles which are described here are uniform across all supported platforms; only the syntax differs. The sample is provided for an AIX system.

To customize the password functions of the system, IBM recommends that you enable the Content Manager OnDemand user security exit and provide a user exit program to handle these activities. This will allow you to determine if a given activity is allowed or disallowed. For example, you could customize the system so that only system administrators and user administrators can change passwords; you could specify a list of users that may not change their passwords; and so on.

These hints and tips apply to the security user exit:

- Once enabled, the user exit program processes all logons to the system. This includes logons from the Windows (end-user) client, the administrative client, client programs such as ARSDOC and ARSLOAD, and the ODWEK Java APIs.
- The user security exit is enabled by setting the SRVR_FLAGS_SECURITY_EXIT parameter in the ARS.INI file to 1 (one).
- The name of the security user exit program is arsusec. See the arsusec.c file in the EXITS directory for a sample source file. See the arscsxit.h file in the EXITS directory for the function and structure declarations for the security user exit program.
- There are several approaches to the main logic of the user exit program, depending on the exact requirements. For example:

  Permit a specific user ID to change passwords. For example, you could allow the built-in user ID (admin) to change passwords. Some customers may need to permit a list of users to change passwords (such as system administrators and user administrators).

  Maintain a list of users that cannot change their passwords. When the actual user ID matches one of the user IDs in the "User Cannot Change Password" list, set the return code to ARCCSXIT_SECURITY_RC_FAILED or ARCCSXIT_SECURITY_RC_PERMS.
- Set the Maximum Password Age parameter to the value that best matches the main logic of the user exit program (permit / deny). The Maximum Password Age parameter is set on the System Parameters dialog box, which is accessed by using the administrative client.

  Set the Maximum Password Age parameter to Never Expires so that users are not prompted to change their passwords. If you are restricting the change password function to a very limited number of users, then this is probably the best overall setting because most users will never automatically be prompted to change their passwords. However, if your organization mandates that passwords must be changed periodically, then the administrator (user or users that can change passwords) must be prompted outside of the system when the password interval has elapsed and it is time to change passwords.

Set the Maximum Password Age parameter to Expires in n Days so that users will be prompted to change their passwords periodically. For the (hopefully) small number of users that cannot change their passwords, any attempt to do so will fail; the standard password expiration policy for the system handles all other users.

The example program is provided on an as-is basis. The licensee of the Content Manager OnDemand product is free to copy, revise modify, and make derivative works of this program sample as they see fit.

```
#define _ARSUSEC_C

/*****************************************************************/
/*                                                               */
/* MODULE NAME: ARSUSEC.C                                        */
/*                                                               */
/*                                                               */
/* SYNOPSIS:  OnDemand Security Exit                             */
/*                                                               */
/*                                                               */
/* DESCRIPTION:  This module contains the SECURITY function, which   */
/*               is used to validate userids and passwords       */
/*                                                               */
/* COPYRIGHT:                                                    */
/*  5697-G34 (C) COPYRIGHT IBM CORPORATION 2001.                 */
/*  All Rights Reserved                                          */
/*  Licensed Materials - Property of IBM                         */
/*                                                               */
/*  US Government Users Restricted Rights - Use, duplication or      */
/*  disclosure restricted by GSA ADP Schedule Contract with IBM Corp.*/
/*                                                               */
/* NOTE: This program sample is provided on an as-is basis.      */
/*       The licensee of the OnDemand product is free to copy,   */
/*       revise modify, and make derivative works of this program    */
/*       sample as they see fit.                                 */
/*                                                               */
/*****************************************************************/

#include "arscsxit.h"
#include <stdio.h>
#include <string.h>

ArcCSXitSecurityRC
ARSCSXIT_EXPORT
ARSCSXIT_API
SECURITY( char *act_userid,
          char *cur_userid,
          char *cur_passwd,
          char *new_userid,
          char *new_passwd,
          ArcCSXitSecurityAction action,
          char *msg,
          char *clnt_id,
          char *instance,
          char *passthru_text,
          ArcU32 passthru_size,
          ArcByte *passthru_buf       )
{
   ArcCSXitSecurityRC rc;

   msg[0] = '\0';

   if      ( action == ARCCSXIT_SECURITY_USER_LOGIN )
     {
     rc = ARCCSXIT_SECURITY_RC_OKAY_BUT_VALIDATE_IN_OD;
     }
```

```
   else if ( action == ARCCSXIT_SECURITY_USER_UPDATE )
     {
     /* action = change password */

     if ( !strcmp(act_userid,"ADMIN") )
       {
       /* actual userid is ADMIN, who may change passwords */

       rc = ARCCSXIT_SECURITY_RC_OKAY;
       }
     else
       {
       /* actual userid is NOT ADMIN */

       if ( !strcmp(cur_passwd,new_passwd) &&
     *cur_passwd != NULL )
  {
  /* new password must match current password and
     current password must not be blank,
     required for initial login/change password */

  rc = ARCCSXIT_SECURITY_RC_OKAY;
  }
       else
  {
         /* all other users may not change their own password */

         strcpy(msg,"You do not have permission to change your password.");
         rc = ARCCSXIT_SECURITY_RC_PERMS;
         }
       }
     }
   else
     {
     /* all other actions */

     rc = ARCCSXIT_SECURITY_RC_OKAY;
     }

   return( rc );
}
```

System log user exit

Content Manager OnDemand generates messages about the various actions that occur on the system.

For example, when a user logs on the system, Content Manager OnDemand generates a message that contains the date and time, the type of action, the userid, and other information. Unless you specify otherwise, certain messages are automatically saved in the system logging facility. You can configure the system to save other messages in the system logging facility.

In addition to saving messages in the system log, OnDemand sends the messages to the system log user exit. The system log user exit is a point at which you may run a user-written program to process the messages and return control of processing after your user-written program ends. The user-written program can process the messages in any way that you want. For example, the user-written program could send alerts to administrators, compile statistics, or generate accounting information.

The standard system log user exit program is named ARSLOG. The system log user exit program on UNIX servers is named `arslog`; the system log user exit program on Windows servers is named ARSLOG.BAT. The program resides in the Content Manager OnDemand executable directory (`bin`). The ARSLOG program supplied by IBM does not perform any functions. However, you can replace the ARSLOG program

supplied by IBM with a user-written program that does specific functions, such as checking the message number and issuing alerts to administrators.

Important: The name of the user-written program must be ARSLOG. You can change the function of the ARSLOG program, but you may not specify a different program name.

You must do the following to configure the system to send the messages to the system log user exit:

- Enable Content Manager OnDemand to generate system messages and specify the types of messages generated by selecting the appropriate options in the System Parameters dialog box.
- Enable Content Manager OnDemand to generate application group messages by selecting the appropriate option in the System Parameters dialog box.
- Specifying the types of application group messages generated by selecting options on the Message Logging page in application groups.
- Enable Content Manager OnDemand to send messages to a user-defined program by selecting the appropriate options in the System Parameters dialog box.

After you have completed these steps, Content Manager OnDemand automatically saves the messages in the system log and sends the messages to the system log user exit.

The messages that Content Manager OnDemand sends to the system log user exit contain the parameters listed .

Table 13: System log user exit message parameters

Parameter	Purpose	Size (in bytes)	Example
$1	Content Manager OnDemand instance	10	archive
$2	Time stamp	20	08/13/95 14:24:31
$3	Log record identifier	10	57049
$4	Content Manager OnDemand user ID	128	ADMIN
$5	User's accounting information	60	Z76-001J/999999
$6	Severity: **1** Alert **2** Error **3** Warning **4** Information **5** Debug	1	3
$7	Message number	5	31
$8	Message text	255	Failed Login: odaixlib1 7.52.365.12

Table 13: System log user exit message parameters (continued)			
Parameter	**Purpose**	**Size (in bytes)**	**Example**
$9	Document file	variable	The file is stored in the directory that is specified by the ARS_TMP parameter in the ARS.CFG file. The file is deleted immediately after the exit program returns control to Content Manager OnDemand.

If you create your own ARSLOG program, remember that:

- The ARSLOG program and any programs that it might call run under the UID of the `root` user on a UNIX server, or with Administrator authority on a Windows server. These are privileged accounts with unrestricted access to all files and commands on the system.
- You should specify the full path name of all files used by the ARSLOG program and all programs invoked by the ARSLOG program. For example, you should specify `/bin/mail`, and not simply `mail`.

Content Manager OnDemand programs are coded in the C language. However, the ARSLOG program can call any executable file. You are responsible for developing a user-written program. You must validate the quality and performance of the user-written program and any other programs that it calls. The actual mechanism for taking action based on the messages provided by Content Manager OnDemand is dependent on the software that you are using on the system.

The online help for the administrative client provides information about enabling Content Manager OnDemand to generate messages and send them to the ARSLOG program. The online help also provides information about how to select the application group messages that Content Manager OnDemand generates and sends to the ARSLOG program.

The System Log table contains one row for each message that Content Manager OnDemand generates.

Table 14: System Log table			
Column Name	**Data Type**	**Size (in bytes)**	**Description**
time_stamp	Date/Time(TZ)	4	The time stamp of the log record. See the ARSDATE program reference in *IBM Content Manager OnDemand for Multiplatforms: Administration Guide* for information about date formats.
userid	VARCHAR	20	The userid of the user that generated the log record.
severity	CHAR	1	The severity of the log record. **1** Alert **2** Error **3** Warning **4** Information **5** Debugging
msg_num	SMALLINT	2	The message number of the log record.
msg_text	VARCHAR	254	The message text of the log record.

Table 14: System Log table (continued)

Column Name	Data Type	Size (in bytes)	Description
appl_id	CHAR	1	Determines whether Content Manager OnDemand overhead information is valid. **A** Not applicable. The overhead information does not apply to the log record. However, the overhead information can be useful for other purposes. For example, a log record created when a document is retrieved contains overhead information about the document. **N** No. The overhead information does not contain useful information. **Y** The overhead information contains information about the document belonging to this particular log record.
log_id	INTEGER	4	The identifier for the Content Manager OnDemand client connection.
account	VARCHAR	60	The user's accounting information.
doc_name	VARCHAR	11	The name of the object.
doc_off	INTEGER	4	The offset of the document within the compressed object.
doc_len	INTEGER	4	The length of the document within the compressed object. A `0` (zero) means all of the data.
comp_off	INTEGER	4	The offset of the document within the compressed object.
comp_len	INTEGER	4	The length of the document within the compressed object. A `0` (zero) means all of the data.
annot	CHAR	1	Determines whether annotations exist for the document. Applies only if the annotation flag is set (YES) for the application group.
comp_type	CHAR	1	The method used to compress document data.
resource	INTEGER	4	The resource identifier for the document.
pri_nid	SMALLINT	2	The primary storage node identifier.
sec_nid	SMALLINT	2	The secondary storage node identifier.

Sample ARSLOG user exit script for UNIX™

You can use the sample ARSLOG user exit script for UNIX that is provided by IBM.

Sample ARSLOG user exit script for UNIX:

```
# $1 - OnDemand Instance Name
# $2 - Time Stamp
# $3 - Log Identifier
# $4 - Userid
# $5 - Account
```

```
# $6 - Severity
# $7 - Message Number
# $8 - Message Text
# $9 - Document File
#

echo "$@" >> ${ARS_TMP}/syslog.log

if [ -n "$9" ];then
   if [ -f "$9" ];then
      print "Copy log doc $9\n" >> ${ARS_TMP}/syslog.log
      cp $9 /tmp/syslogdocs/$(basename $9).doc 2>> ${ARS_TMP}/syslog.log
   else
      print "$9 does not exist\n" >> ${ARS_TMP}/syslog.log
   fi
fi

exit 0
```

Sample ARSLOG user exit batch file for Windows™

You can use the sample ARSLOG user exit batch file for Windows that is provided by IBM.

Sample ARSLOG user exit batch file for Windows:

```
REM %1 - OnDemand Instance Name
REM %2 - Time Stamp
REM %3 - Log Identifier
REM %4 - Userid
REM %5 - Account
REM %6 - Severity
REM %7 - Message Number
REM %8 - Message Text
REM %9 - Document File
REM

ECHO %1 %2 %3 %4 %5 %6 %7 %8 %9>%ARS_TMP%\System.log

REM make sure the %ARS_TMP%\Syslog directory exists
IF EXIST %9 COPY %9 %ARS_TMP%\Syslog

EXIT
```

Table space creation user exit

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

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

Interface exit components

The table space creation exit consists of the ARSUTBL sample program, which contains a sample table space creation exit in C source form.

The table space creation is implemented by a DLL which in this document is called arsutbl. However, the DLL can have any name and can reside in any directory. The sample ARSUTBL places the DLL in `/opt/IBM/ondemand/V10.1/bin/exits`.

The following statement must exist in the ARS.CFG file that is associated with the instance so that the arsutbl DLL can be invoked:

```
ARS_DB_TABLESPACE_USEREXIT=absoulte path name
```

For the sample arsutbl, you would specify the following statement in the ARS.CFG file:

```
ARS_DB_TABLESPACE_USEREXIT=/opt/IBM/ondemand/V10.1/bin/exits/arsutbl
```

C language arsutbl

The Format example shows the C language arsutbl.

C language arsutbl sample:

```
#pragma export(TBLSPCRT)
#include "arscsxit.h"
int TBLSPCRT( ArcCSXitApplGroup appl_grp,
              char tblsp_name,
              char table_name,
              char idx_name,
              char sql,
              int  action,
              int  created
            )
```

General description

On entry, action should be interrogated to determine the type of action being performed.

The actions are:

Table space creation
1

Table creation
2

Index creation
3

Final call
4

Do not attempt to use parameters that are not specified as being valid for a given action. Do not attempt to access storage beyond the terminating X'00'. The exit should not modify any input parameters except the SQL string if supplied for the action.

Index Creation is called once for each index that is being created on the table.

Table 15: Tablespace Create

Parameter	Meaning
appl_grp	Contains information related to the application group for which the table space is being created. This includes the application group name, the application group identifier, and the internal application group name.
tblsp_name	Null delimited table space name being created.

Table 16: Table Create	
Parameter	**Meaning**
appl_grp	Contains information related to the application group for which the table space is being created. This includes the application group name, the application group identifier, and the internal application group name.
tblsp_name	Null delimited table space name being created.
table_name	Null delimited table name being created.
sql	Null delimited SQL that will be used to create the table. The installation can alter this, however the resultant string plus the tailing X'00' must not exceed 16384 bytes.

Table 17: Index Create	
Parameter	**Meaning**
appl_grp	Contains information related to the application group for which the table space is being created. This includes the application group name, the application group identifier, and the internal application group name.
tblsp_name	Null delimited table space name being created.
table_name	Null delimited table name being created.
idx	Null delimited index name being created.
sql	Null delimited SQL that will be used to create the table. The installation can alter this, however the resultant string plus the tailing X'00' must not exceed 16384 bytes.

Returned values

The *created return value should be set.

Table 18: Return codes	
Action	**Meaning**
0	OnDemand should create the table space, table, or index.
non-zero	The exit has created the table space, table or index.

The exit should set a return code of 0 (zero).

Chapter 11. National Language Support

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

Unless otherwise indicated, information applies to all supported operating systems.

Conversion between different code pages

A code page maps each character from a character set, such as the Latin alphabet, to a numeric representation.

Each code page is identified by a numeric identifier. For example, code page 850 represents the character A as hexadecimal 41.

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

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

When does character conversion occur?

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

Character conversion can occur

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

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

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

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

  Character data conversion takes place on the server using Unicode code page mapping tables. Content Manager OnDemand converts the characters from the code page used by ACIF to Unicode and then from Unicode to the code page of the instance. For example, index data generated by ACIF and encoded in code page 500 (ISO EBCDIC) can be stored in an instance that has character data encoded in code page 819.

Character conversion will not occur

Documents stored in Content Manager OnDemand.

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

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

Character mapping

For double-byte character set (DBCS) AFP data and DBCS and single-byte character set (SBCS) line data, the Content Manager OnDemand client automatically converts characters in a document from the code page of the server to the code page of the client using ICU converters.

This method of character mapping works with the Windows client and supports DBCS (for AFP and line data) and SBCS (for line data) languages, including most DBCS User Defined Character (UDC) mappings. The ICU converters automatically map the user-defined area of a code page to the standard user-defined area of the corresponding ICU table. If the code page contains UDC mappings outside the standard user-defined area, you can create and use your own ICU converter. (Otherwise, the viewing program will not be able to display the characters correctly.) Use the Character Data Representation Architecture (CDRA) utility to create your own ICU converter. For AFP data, change the IconvLocalePath parameter in the Preferences section of the FLDPORT2.INI file to use your own ICU converter.

For more information on ICU, see http://www.icu-project.org/.

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

How does Content Manager OnDemand determine code page values?

The client code page is determined from the operating environment when the connection to the instance is made.

For example, the Windows client derives the code page from the locale as specified in the Regional Settings under Control Panel. The instance code page is derived from the value specified at the time the instance is created. The instance is in one and only one code page.

The code page of index data generated by ACIF is determined by the value of the CPGID parameter. When index data is stored in an instance, it is converted from the code page used by ACIF to the code page of the instance. (When index data is retrieved from an instance, it is converted from the code page of the instance to the code page of the client.)

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

Creating application groups

An application group is a container that holds report data.

You store reports and the index data used to retrieve and maintain them in an application group. You define database fields for each application group. The database fields represent categories of information in a report. When you load a report into an application group, you store index information about the report in the database.

When you define database fields, you specify attributes of the fields. Attributes include the field name, type, and length. For character data, the field length must specify the number of bytes required to hold the field data in the database. For multi-byte languages, character string conversion between code pages might result in either an increase or decrease in the length of the string when data is loaded into the database. For example, the client might require two bytes to display a double-byte character and the server might require three bytes to store the character in the euc code page of the database. In the example, the string length must be increased so that the database field is large enough to hold the converted string. The maximum length of a string field in Content Manager OnDemand is 254 bytes. Verify the length of each database (string) field you define:

- If you use the Report Wizard to generate application groups, the Report Wizard converts strings you select to the code page of the database and displays the number of characters required to hold the string in the database. You can accept the value generated by the Report Wizard or replace it with another value.
- If you use the Add an Application Group command to add application groups, you must calculate the number of bytes required to hold the field in the database and enter the value on the Field Information page.

Create applications

You typically create a Content Manager OnDemand application for each type of report or source of data that you plan to store in Content Manager OnDemand. When you create an application, you specify attributes of the application. The attributes include:

- The data type of the report as it is stored in Content Manager OnDemand (for example, AFP). The data type determines the viewing program used to display pages of the report.
- The program used to index the report. If you use one of the indexing programs provided with Content Manager OnDemand, the application typically includes the parameters that the indexing program uses to process the report and generate the index data.
- Logical views of report data. Logical views provide different ways to view pages of a line data report.

You can create an application by using the Report Wizard or by using the Add an Application command. You can create indexing information by entering parameters and values directly into the application, specifying the name of a parameter file that contains the information, or using the Graphical Indexer to generate indexing information. You can create logical views by entering values directly into the application or using the sample data window to generate the logical view information.

Data Type

The Data Type of the application identifies the format of the data as it is stored in Content Manager OnDemand and the viewer that the client calls to display documents stored in the application.

If you plan to store line data in Content Manager OnDemand or create indexing parameters with the graphical indexer (using a line data source file), you must set the Data Type of the application to Line. When you set the Data Type to Line:

- Verify the code page of the data. The code page of the data is typically the code page of the operating system where the data was created. In Content Manager OnDemand, the default code page for line data is 500 (ISO EBCDIC).
- If the line data contains shift-in and shift-out (SOSI) codes, indicate how ACIF handles them. Shift-in and shift-out codes indicate when the code points in a record change from single byte to double byte and double byte to single byte. Select from SOSI1, SOSI2, and SOSI3.

  MBCS Considerations: Multi-byte character set (MBCS) code pages that contain shift-in and shift-out codes are supported in line data but not AFP. Examples of MBCS code pages are 939 (Japan) and 933 (Korea).

For all other types of data, the code page is encapsulated in the data. For AFP data, it is possible that characters are encoded in more than one code page. The AFP viewer uses mapping files to display single-

and double-byte data in the proper code page. You might need to map AFP fonts a document was created with to outline fonts on the PC to properly display some characters. The *Client Customization Guide* provides details about mapping AFP fonts.

Indexing with ACIF

If you use ACIF to index the input data, the application indexing parameters determine how ACIF indexes the data and the code page of the index data generated by ACIF.

The CPGID (Code Page Global Identifier) parameter identifies the code page of the index data generated by ACIF. The CPGID should be the same as the code page of the source data. You must code ACIF trigger and index string values in the code page of the source data. See the *IBM Content Manager OnDemand Indexing Reference* for more information about ACIF, including examples that show how to code trigger and index string values for EBCDIC data.

Important: If you plan to use ACIF to generate index data in a multi-byte language, then you must allow ACIF to create the Map Coded Font Format 2 (MCF2) structured fields with coded fonts. To do so, you must set the MCF2REF parameter to CF (MCF2REF=CF).

Indexing with the Generic indexer

If you use the Generic Indexer to index the input data, the default code page is 819 (Latin 1, code set ISO 8859-1).

If you need to generate index data in some other code page, you must specify the code page by using the CODEPAGE: parameter in the parameter file that is used by the Generic Indexer. See the *IBM Content Manager OnDemand Indexing Reference* for more information.

Running Content Manager OnDemand programs

Certain Content Manager OnDemand programs accept input data (parameters and values) from a parameter file. When you work with a multi-byte language database, the data in the parameter file must be encoded in the euc code page of the database.

When you need to create a parameter file, you should log on at a workstation that is running under the same code page as the database. Do not create the parameter file on a workstation and then use the FTP program to send it to a server. Do not log on to a server with the TELNET program. (However, you can use the TELNET program to run commands from the prompt, so long as you enter all of the parameters and values at the command prompt.)

Troubleshooting incorrect NLS characters

Check the code page of the database. When Content Manager OnDemand is installed, the default code page is 1208. The 1208 is an example of an integer identifier known as a Coded Character Set Identifier (CCSID).

If your characters display incorrectly, try the following troubleshooting tips:

- The 819 CCSID maps to the UTF-8 code set. If the default code page needs to be changed , it must be changed before the Content Manager OnDemand database instance is created.
- Check to make sure you are using the appropriate language version of the Content Manager OnDemand Windows or administrative client, during the thick client install options are provided to install the appropriate language version.
- Check to see if custom template HTML files specify a language encoding. Character conversion could be taking place based on the encoding that is specified.
- You can adjust the encoding of your browser.

Chapter 12. SSL, certificates, certificate authorities, and public-key cryptography

Review this information to understand the technology involved in creating an SSL connection between a server and a client.

During an SSL handshake, a client and server securely exchange digital signatures and encryption keys by using a public-key algorithm (usually RSA). The client and server establish a secure connection with this identity and key information. After the client and server establish a secure session, they transmit the data to each other, encrypting it with a symmetric algorithm, such as AES.

The client and server do the following steps during the SSL handshake:

1. The client requests an SSL connection and includes a list of its supported cipher suites in that request.
2. The server responds by selecting a cipher suite from the list.
3. The server sends its digital certificate to the client.
4. The client authenticates the server certificate by checking with the trusted certificate authority that issued the server certificate or by checking its own key database.
5. The client and server securely negotiate a session key and a message authentication code (MAC).
6. The client and server securely exchange information using the key and the selected MAC.

The optional authentication of the client is not supported during the SSL handshake.

Overview of the SSL handshake

During an SSL handshake, a client and server securely exchange digital signatures and encryption keys by using a public-key algorithm (usually RSA). The client and server establish a secure connection with this identity and key information. After the client and server establish a secure session, they transmit the data to each other, encrypting it with a symmetric algorithm, such as AES.

The client and server do the following steps during the SSL handshake:

1. The client requests an SSL connection and includes a list of its supported cipher suites in that request.
2. The server responds by selecting a cipher suite from the list.
3. The server sends its digital certificate to the client.
4. The client authenticates the server certificate by checking with the trusted certificate authority that issued the server certificate or by checking its own key database.
5. The client and server securely negotiate a session key and a message authentication code (MAC).
6. The client and server securely exchange information using the key and the selected MAC.

Content Manager OnDemand does not support the (optional) authentication of the client during the SSL handshake.

Digital certificates and certificate authorities

Trusted parties, called certificate authorities (CA), issue digital certificates to verify the identity of an entity, such as a client or a server.

The digital certificate serves the following purposes:

- Verify the identity of the owner.
- Make the public key of the owner available.

The certificate authority issues the certificate with an expiration date, after which the certificate is no longer guaranteed by the certificate authority.

To obtain a digital certificate, you send a request to the CA of your choice; for example, Verisign or RSA. The request includes your distinguished name, your public key, and your signature. A distinguished name (DN) is a unique identifier for each user or host for which you are applying for a certificate. The CA checks your signature using your public key and performs some level of verification of your identity. (The verification process varies between CAs). After verification, the CA sends you a signed digital certificate that contains your distinguished name, your public key, the distinguished name of the CA, and the signature of the CA. You store this signed certificate in your key database.

When you send this certificate to a receiver, the receiver does the following steps to verify your identity:

1. Uses your public key that comes with the certificate to check your digital signature.
2. Verifies that the CA that issued your certificate is legitimate and trustworthy. To do this, the receiver needs the public key of the CA. The receiver might already hold an assured copy of the public key of the CA in their key database, but if not, the receiver must acquire an additional digital certificate to obtain the public key of the CA. This certificate might in turn depend on the digital certificate of another CA; there might be a hierarchy of certificates issued by multiple CAs, each depending on the validity of the next. Eventually, however, the receiver needs the public key of the root CA. The root CA is the CA at the top of the hierarchy. To trust the validity of the digital certificate of the root CA, the public-key user must receive that digital certificate in a secure manner, such as through a download from an authenticated server, or with preloaded software received from a reliable source, or on a securely delivered diskette.

Many applications that send a digital certificate also send all of the CA digital certificates necessary to verify the hierarchy of certificates up to the root CA certificate. For a digital certificate to be entirely trustworthy, the owner of the digital certificate must protect the private key, for example, by encrypting it on their computer's hard drive. If their private key has been compromised, an imposter could misuse their digital certificate.

You can use self-signed digital certificates for testing purposes. A self-signed digital certificate contains your distinguished name, your public key, and your signature.

Public-key cryptography

SSL uses public-key algorithms to exchange encryption key information and digital certificate information for authentication. Public-key cryptography (also known as asymmetric cryptography) uses two different encryption keys: a public key to encrypt data and an associated private key to decrypt it.

Conversely, symmetric key cryptography uses just one key, which is shared by all parties involved in the secure communication. This secret key is used both to encrypt and decrypt information. The key must be safely distributed to, and stored by, all parties, which is difficult to guarantee. With public-key cryptography, the public key is not secret, but the messages it encrypts can only be decrypted by using it's associated private key. The private key must be securely stored, for example, in your key database, or encrypted on your computer's hard drive.

Public-key algorithms alone do not guarantee secure communication; you also need to verify the identity of whoever is communicating with you. To perform this authentication, SSL uses digital certificates. When you send your digital certificate to someone, the certificate provides them with your public key. You have used your private key to digitally sign your certificate and so the receiver of the communication can use your public key to verify your signature. The validity of the digital certificate itself is guaranteed by the certificate authority (CA) that issued it.

Supported cipher suites

During an SSL handshake, the client and server negotiate which cipher suite to use to exchange data. A cipher suite is a set of algorithms that are used to provide authentication, encryption, and data integrity.

Content Manager OnDemand runs GSKit in FIPS mode to provide SSL support. GSKit supports the following cipher suites:

- TLS_RSA_WITH_AES_256_CBC_SHA
- TLS_RSA_WITH_AES_128_CBC_SHA
- TLS_RSA_WITH_3DES_EDE_CBC_SHA

The name of each cipher suite specifies the algorithms that it uses for authentication, encryption, and data integrity checking. For example, the cipher suite TLS_RSA_WITH_AES_256_CBC_SHA uses RSA for authentication; AES 256-bit and CBC for encryption algorithms; and SHA-1 for the hash function for data integrity.

You cannot prioritize which cipher suite is selected.

Default GSKit trusted root certificates

Below is a list of trusted root certificates.

- Thawte Personal Premium CA
- Thawte Personal Freemail CA
- Thawte Personal Basic CA
- Thawte Premium Server CA
- Thawte Server CA
- VeriSign Class 3 Secure Server CA
- VeriSign International Server CA - Class 3
- VeriSign Class 4 Public Primary Certification Authority - G3
- VeriSign Class 3 Public Primary Certification Authority - G5
- VeriSign Class 3 Public Primary Certification Authority - G3
- VeriSign Class 2 Public Primary Certification Authority - G3
- VeriSign Class 1 Public Primary Certification Authority - G3
- VeriSign Class 4 Public Primary Certification Authority - G2
- VeriSign Class 3 Public Primary Certification Authority - G2
- VeriSign Class 2 Public Primary Certification Authority - G2
- VeriSign Class 1 Public Primary Certification Authority - G2
- VeriSign Class 3 Public Primary Certification Authority
- VeriSign Class 2 Public Primary Certification Authority
- VeriSign Class 1 Public Primary Certification Authority
- Entrust.net Global Secure Server Certification Authority
- Entrust.net Global Client Certification Authority
- Entrust.net Client Certification Authority
- Entrust.net Certification Authority (2048)
- Entrust.net Secure Server Certification Authority

Setting up SSL on the Content Manager OnDemand client

You can configure SSL on the Content Manager OnDemand client.

Do the following tasks:

1. If you obtained and installed a CA-signed digital certificate on your Content Manager OnDemand server (as describe in the following topics):
   - “Setting up SSL on Content Manager OnDemand for AIX” on page 18
   - Setting up SSL on the Content Manager OnDemand for Solaris server
   - “Setting up SSL on the Content Manager OnDemand for Linux™ server” on page 61
   - “Setting up SSL on the Content Manager OnDemand for Windows server” on page 102
2. If you created and installed a self-signed digital certificate on your Content Manager OnDemand server, add the self-signed digital certificate to the key database of the client by using GSKCapiCmd. The following command adds the digital certificate from the file `ondemand.arm` into the key database `ondemand.kdb`:

```
gsk8capicmd_64 -cert -add -db "ondemand.kdb" -pw "myKeyDBpasswd" -label
"dbselfsigned" \
-file "ondemand.arm" -format ascii
```

3. Configure the Content Manager OnDemand client to connect to the SSL port.

Chapter 13. Creating Content Manager OnDemand system tables into user-defined table spaces

When you create the Content Manager OnDemand tables or indexes, the ARSDB command can build the tables and indexes in the database default table space or in table spaces that you create.

About this task

When you run the ARSDB command, Content Manager OnDemand validates the existence of the table space of the Content Manager OnDemand system tables. If the table spaces do not exist, the ARSDB command creates the Content Manager OnDemand system tables and indexes into the current database default table space.

If you want the ARSDB command to build the Content Manager OnDemand system tables and indexes into tables spaces that you create (user-defined table spaces), do the following tasks before you run the ARSDB command:

Procedure

1. Create the table spaces by using the database utility tools. If you choose the Content Manager OnDemand default table space names then you can skip the next step.
2. If you want to create your own table space names, modify the Content Manager OnDemand configuration file (AIX or Linux) or properties of an instance (Windows). Specify the names of the tables spaces that you created by adding the parameters listed.

Parameters to specify names for system tables and table spaces

The following list describes the parameters you can use to specify a name for each Content Manager OnDemand system table space.

ARS_ARSAG_TABLESPACE
ARS_ARSAG2FOL_TABLESPACE
ARS_ARSAGFLD_TABLESPACE
ARS_ARSAGFLDALIAS_TABLESPACE
ARS_ARSAGINDEX_TABLESPACE
ARS_ARSAGPERMS_TABLESPACE
ARS_ARSANN_TABLESPACE
ARS_ARSAPP_TABLESPACE
ARS_ARSAPPUSR_TABLESPACE
ARS_ARSCAB_TABLESPACE
ARS_ARSCAB2FOL_TABLESPACE
ARS_ARSCABNAMES_TABLESPACE
ARS_ARSCABPERMS_TABLESPACE
ARS_ARSCFSODWORK_TABLESPACE
ARS_ARSFOL_TABLESPACE
ARS_ARSFOLFLD_TABLESPACE
ARS_ARSFOLFLDUSR_TABLESPACE
ARS_ARSFOLNAMES_TABLESPACE
ARS_ARSFOLPERMS_TABLESPACE
ARS_ARSFTIWORK_TABLESPACE
ARS_ARSGROUP_TABLESPACE
ARS_ARSHOLD_TABLESPACE
ARS_ARSHOLDMAP_TABLESPACE
ARS_ARSHOLDNAMES_TABLESPACE

ARS_ARSHOLDPERMS_TABLESPACE
ARS_ARSHOLDWORK_TABLESPACE
ARS_ARSLOAD_TABLESPACE
ARS_ARSLOADWORK_TABLESPACE
ARS_ARSNAMEQ_TABLESPACE
ARS_ARSNODE_TABLESPACE
ARS_ARSPRT_TABLESPACE
ARS_ARSPRTOPTS_TABLESPACE
ARS_ARSPRTUSR_TABLESPACE
ARS_ARSRES_TABLESPACE
ARS_ARSSEG_TABLESPACE
ARS_ARSSET_TABLESPACE
ARS_ARSSYS_TABLESPACE
ARS_ARSUSER_TABLESPACE
ARS_ARSUSRGRP_TABLESPACE
ARS_ARSUSRGRPID_TABLESPACE

The following list describes the parameters you can use to specify a name for each Content Manager OnDemand Distribution Facility table space:

ARS_ARSDFBDT_TABLESPACE
ARS_ARSDFCRT_TABLESPACE
ARS_ARSDFDCT_TABLESPACE
ARS_ARSDFDRT_TABLESPACE
ARS_ARSDFDST_TABLESPACE
ARS_ARSDFEML_TABLESPACE
ARS_ARSDFLIS_TABLESPACE
ARS_ARSDFPPT_TABLESPACE
ARS_ARSDFUOT_TABLESPACE

Content Manager OnDemand system tables and the default table space names

Each table name has a coexisting table space name.

Table 19: OnDemand system tables and the default names for the table spaces.

Table name	Table space names
ARSAG	ARSAGT
ARSAG2FOL	ARSAG2FT
ARSAGFLD	ARSAGFLT
ARSAGFLDALIAS	ARSAGFAT
ARSAGINDEX	ARSAGIDT
ARSAGPERMS	ARSAGPET
ARSANN	ARSANNT
ARSAPP	ARSAPPT
ARSAPPUSR	ARSAPPUT
ARSCAB	ARSCABT
ARSCAB2FOL	ARSCABFT
ARSCABPERMS	ARSCABPT

Table 19: OnDemand system tables and the default names for the table spaces. (continued)	
Table name	**Table space names**
ARSCABNAMES	ARSCABNT
ARSCFSODWORK	ARSCFSWT
ARSFOL	ARSFOLT
ARSFOLFLD	ARSFOLFT
ARSFOLFLDUSR	ARSFOLUT
ARSFOLNAMES	ARSFOLNT
ARSFOLPERMS	ARSFOLPT
ARSFTIWORK	ARSFTIWT
ARSGROUP	ARSGROUT
ARSHOLD	ARSHLDT
ARSHOLDMAP	ARSHLDMT
ARSHOLDNAMES	ARSHLDNT
ARSHOLDPERMS	ARSHLDPT
ARSHOLDWORK	ARSHLDWT
ARSLOAD	ARSLOADT
ARSLOADWORK	ARSLDWKT
ARSNAMEQ	ARSNAMET
ARSNODE	ARSNODET
ARSPRT	ARSPRTT
ARSPRTOPTS	ARSPRTOT
ARSPRTUSR	ARSPRUST
ARSRES	ARSREST
ARSSEG	ARSSEGT
ARSSET	ARSSETT
ARSSYS	ARSSYST
ARSUSER	ARSUSERT
ARSUSRGRP	ARSUSRGT
ARSUSRGRPID	ARSUSGIT

Table 20: ODF system tables and the default names for the table spaces.	
Table name	**Table space names**
ARSDFBDT	ARSDFBDT
ARSDFCRT	ARSDFCRT
ARSDFDCT	ARSDFDCT
ARSDFDRT	ARSDFDRT

Table 20: ODF system tables and the default names for the table spaces. (continued)	
Table name	**Table space names**
ARSDFDST	ARSDFDST
ARSDFEML	ARSDFEML
ARSDFLIS	ARSDFLIS
ARSDFPPT	ARSDFPPT
ARSDFUOT	ARSDFUOT

Accessibility information for Content Manager OnDemand

For complete information about accessibility features that are supported by this product, see your *Content Manager OnDemand Administration Guide*.

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-IBMproducts. 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.

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" www.ibm.com/legal/copytrade.shtml.

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

Intel, Intel logo, Intel Inside, Intel Inside logo, Intel Centrino, Intel Centrino logo, Celeron, Intel Xeon, Intel SpeedStep, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

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

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

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

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

Portions of the OnDemand Windows client program contain licensed software from Pixel Translations Incorporated, © Pixel Translations Incorporated 1990, 2003. All rights reserved.

All rights reserved.

FairCom Corp. and other company, product or service names might be trademarks or service marks of others.

Privacy policy considerations

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 http://www.ibm.com/privacy and IBM’s Online Privacy Statement at http://

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 http://www.ibm.com/software/info/product-privacy.

Index

A

B

C

D

E

F

G

H

I

L

M

N

O

P

R

S

T

U

V

W

GC19-3342-02