EDMSuite OnDemand for Windows NT
Installation and Configuration Guide
Version 2.2

G544-5526-03

April 30, 1998

EDMSuite OnDemand for Windows NT

Installation and Configuration Guide

Version 2.2

EDMSuite OnDemand for Windows NT

Installation and Configuration Guide

Version 2.2

G544-5526-03

Fourth Edition (May 1998)

This edition of IBM EDMSuite OnDemand for Windows NT: Installation and Configuration Guide applies to IBM EDMSuite OnDemand Version 2 Release 2 and to all subsequent releases of this product until otherwise indicated in new releases or technical newsletters.

The following paragraph does not apply to the United Kingdom or any country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS MANUAL ″AS IS″ WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions. Therefore, this may not apply to you.

IBM does not warrant that the contents of this publication or the accompanying source code examples, whether individually or as one or more groups, will meet your requirements or that the publication or the source code is error-free.

Requests for copies of this publication and for technical information about IBM products should be made to your IBM authorized Dealer, your IBM Marketing Representative, or your IBM Printing Systems Company Representative.

The IBM Printing Systems Company welcomes your comments. You may send your comments by fax to 1-800-524-1519, by e-mail to printpub@us.ibm.com, or mail your comments to:

INFORMATION DEVELOPMENT
IBM PRINTING SYSTEMS COMPANY
DEPARTMENT H7FE BUILDING 003G
PO BOX 1900
BOULDER CO 80301-9191

Note: Visit our home page at http://www.printers.ibm.com.

When you send information to IBM, you grant IBM a nonexclusive right to use or distribute the information in any way it believes appropriate without incurring any obligation to you.

However, the following copyright notice protects this documentation under the Copyright laws of the United States and other countries which prohibit such actions as, but not limited to, copying, distributing, modifying, and making derivative works.

Contents

Notices

References in this publication to products or services of IBM do not suggest or imply that IBM will make them available in all countries where IBM does business or that only products or services of IBM may be used. Noninfringing equivalents may be substituted, but the user must verify that such substitutes, unless expressly designated by IBM, work correctly. No license, expressed or implied, to patents or copyrights of IBM is granted by furnishing this document. You can send license inquiries, in writing, to the IBM Director of Licensing, IBM Corporation, 500 Columbus Avenue, Thornwood, NY 10594, USA.

Licensees of this program who wish to have information about it for the purpose of enabling: (1) the exchange of information between independently created programs and other programs (including this one) and (2) the mutual use of the information, which has been exchanged, should contact: IBM Corporation, Printing Systems Company Legal Department, Mail Drop 001W, Boulder, Colorado 80301 USA. Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

Trademarks and Service Marks

The following terms, used in this publication and other books in the IBM EDMSuite OnDemand library, are trademarks, registered trademarks, or service marks of the IBM Corporation in the United States and other countries:

ADSTAR
Advanced Function Printing
AFP
AFPDS
AIX
AIX/6000
AIXwindows
DATABASE 2
DB2
EDMSuite
IBM
MO:DCA-P
MVS

MVS/XA
MVS/ESA
OS/2
OS/390
PSF/2
PSF for AIX
PSF/MVS
RS/6000
SP
SP2
System/370
System/390
SNA
VisualInfo

The following terms, used in this publication and other books in the IBM EDMSuite OnDemand library, are trademarks of other companies as listed:

Acrobat, Acrobat Exchange, Adobe, Adobe Type Manager, ATM, Distiller, and PostScript are trademarks or registered trademarks of Adobe Systems Incorporated.

Cygnet is a registered trademark of Cygnet Systems, Inc.

Ethernet is a trademark of Xerox Corporation.

Exabyte is a trademark of Exabyte Corporation.

FaxBox is a trademark of DCE Corporation.

Hewlett-Packard, HP-UX, HP VUE, and PCL are trademarks of Hewlett-Packard Company.

iFOR is a registered trademark of Gradient Technologies, Inc.

Microsoft, TrueType, Windows, Windows 95, and Windows NT are registered trademarks of Microsoft Corporation.

MMX is a trademark and Pentium is a registered trademark of Intel Corporation.

Motif is a trademark of Open Software Foundation, Inc.

Netscape and Netscape Navigator are trademarks of Netscape Communications Corporation.

NFS is a trademark of SUN Microsystems Incorporated.

Novell is a trademark of Novell Inc.

PC/TCP is a trademark of FTP Software, Inc.

Post-it is a registered trademark of 3M.

UNIX is a registered trademark in the United States and other countries licensed exclusively through X/Open Company Limited.

WinFax Lite is a trademark of Delrina Corporation.

X Windows is a trademark of the Massachusetts Institute of Technology.

Portions of the OnDemand Windows client program contain licensed software from Adobe Systems Incorporated, © Adobe Systems Incorporated 1987-1997. All rights reserved.

Portions of the OnDemand Windows client program contain licensed software from Pixel Translations Incorporated, © Pixel Translations Incorporated 1990, 1997. All rights reserved.

About this publication

This book provides information about installing and configuring IBM EDMSuite OnDemand for Windows NT Version 2.2 (OnDemand) servers. You can use the information in this book when you initially set up an OnDemand server on a Windows NT system, for example, to configure devices and install and configure software. The information in this book can also be used as a reference for an operational OnDemand server, for example, to add devices, install and configure optional software, and change the database manager and storage manager configuration.

Who should use this publication

This book is of primary interest to people responsible for installing and configuring software products. This book can also be used by other people in an organization responsible for planning and maintaining hardware, software, and applications. This book provides the information an installer needs to install and configure an OnDemand system, such as installing OnDemand software and related software programs, configuring the services required to operate the server, and setting up maintenance tasks to run automatically on a regular schedule.

How this publication is organized

"Before you begin", beginning on page 1 contains important information you should review before you begin installing and configuring your OnDemand system.

"Installing software," beginning on page 21 provides instructions for installing OnDemand, DB2 UDB, SQL Server, and ADSM software.

"Configuring the system", beginning on page 31 describes how to configure OnDemand and related software programs to operate in your environment.

"Post installation", beginning on page 63 describes how to verify the installation of your OnDemand system and includes tasks we recommend you complete before you begin production use of the system.

Product support

The IBM support center maintains current information about EDMSuite OnDemand products.

Before you install OnDemand, contact the support center or your IBM software service representative to obtain the latest maintenance information about OnDemand.

If problems occur when you install or run any of the OnDemand programs, you can call the support center to obtain software problem and defect support. The phone number for the IBM support center is 1-800-237-5511. The OnDemand for Windows NT program ID is 5639-E12 and the component ID is 5622-66200.

Our use of typefaces

Throughout this book, words and phrases appear in **Bold**, *Italic*, and other fonts. The following explains our convention when using these fonts.

Bold	Used for paragraphs that call attention to especially relevant information about a topic or command.
Italic	Used to emphasize concepts and terms.
Monospace	Indicates output of commands and programs in examples. Also used for information you are instructed to type.

EDMSuite OnDemand publications

The following publications contain information about IBM EDMSuite OnDemand Version 2.2:

Introduction and Planning Guide, G544-5281

Installation and Configuration Guide for Windows NT, G544-5526

Installation Guide for OnDemand Clients, G544-5527

Administrator's Reference, S544-5293

Indexing Reference, S544-5489

Getting Started with the User Interface for OS/2, S544-5510

Getting Started with the User Interface for Windows 3.1, S544-5509

Getting Started with the User Interface for Windows NT and Windows 95, S544-5469

Getting Started with the Administrator Interface for Windows NT and Windows 95, S544-5463

Part 1. Before you begin

This section of the book contains important information you should review before you begin installing and configuring the OnDemand server software:

- Installation checklist
- Verifying the required and optional server hardware and software products
- Choosing a configuration
- Windows NT Server security basics
- Creating the OnDemand system administrator account

Chapter 1. Installation checklist

Setting up your OnDemand for Windows NT system typically requires completion of the following tasks:

__ 1. Contact the IBM Support Center for the latest maintenance levels of DB2 UDB, ADSM, and OnDemand for Windows NT.

__ 2. Obtain a copy of the OnDemand for Windows NT README file. Read the entire file before you begin.

__ 3. Check the OnDemand for Windows NT prerequisites (refer to "Chapter 2. Hardware and software requirements" on page 7) and verify the required and optional hardware and software products.

__ 4. Choose the type of server configuration that meets your needs (refer to "Chapter 3. Choosing a configuration" on page 13 for more information).

__ 5. Create the OnDemand system administrator account on each workstation that is part of the OnDemand system (refer to "Chapter 5. OnDemand system administrator account" on page 19).

__ 6. Install OnDemand on each workstation that is part of the OnDemand system (see "Chapter 6. Installing OnDemand" on page 23).

__ 7. Install DB2 UDB (see "Chapter 7. Installing DB2 UDB" on page 25) or SQL Server (see "Chapter 8. Installing SQL Server" on page 27) on the library server.

__ 8. If you plan to maintain data on archive media, install ADSM on each workstation that supports archive media (see "Chapter 9. Installing ADSM" on page 29).

__ 9. If you installed ADSM to maintain data on archive media, configure ADSM (see "Chapter 10. Configuring ADSM" on page 33). This step includes defining devices to ADSM, defining policies, registering client nodes, and preparing storage volumes.

__ 10. Configure OnDemand on each workstation that is part of the OnDemand system (see "Chapter 11. Configuring OnDemand servers" on page 51). This step includes configuring the database manager, storage manager, services, and maintenance tasks.

__ 11. Create and initialize the OnDemand database on the library server (see "Chapter 12. Creating and initializing the database" on page 61).

If you plan to reprint documents using the OnDemand server print facility, complete the following tasks:

__ 1. Obtain a copy of *PSF for AIX: Print Administration*.

__ 2. Install and configure PSF for AIX.

__ 3. Verify that all resources and fonts required to reprint documents stored in OnDemand are installed on the PSF for AIX workstation.

__ 4. Obtain the TCP/IP host name of the PSF for AIX workstation.

__ 5. On the library server, edit the arsprt.bat file and insert the host name of the PSF for AIX workstation. The arsprt.bat file is in the OnDemand for Windows NT program directory.

If you plan to transmit data from MVS systems to OnDemand servers using the MVS Download feature, complete the following tasks:

__ 1. Obtain a copy of the *PSF/MVS: MVS Download Guide*.

__ 2. Check the prerequisites and verify the MVS, PSF/MVS, and TCP/IP software levels for the MVS Download feature.

__ 3. Install and configure the MVS Download feature on the MVS system.

__ 4. Set up the OnDemand MVSD service on each workstation that is part of the OnDemand system and will use MVS Download to receive datasets from MVS systems (refer to "Chapter 11. Configuring OnDemand servers" on page 51).

Verify the installation of OnDemand (refer to "Chapter 13. Verifying the installation" on page 65).

__ 1. After installing and configuring an OnDemand server, **restart the system**. Windows NT reinitializes and starts the services required by OnDemand.

__ 2. Share the OnDemand client program folder so that other users on the network can install client software. We recommend that you use the share name `odclient` to share the `\Program Files\IBM\OnDemand for WinNT\client` folder.

__ 3. To access the server, you must install at least one of the OnDemand client programs. Refer to the *Installation Guide for OnDemand Clients* for details about installing OnDemand client software.

__ 4. Log on to the library server with an OnDemand client program.

Prepare the system for production use:

__ 1. Define storage sets (refer to "Chapter 14. Defining storage sets" on page 67). Before you define reports or load data into an application group, you must define storage sets. In addition, to support the system application groups, we recommend that you define at least one storage set that writes data to archive media.

__ 2. Configure the System Log application group (refer to "Chapter 15. Configuring the System Log application group" on page 69). Before you

define OnDemand objects, load data, or let users access the system, update the System Log application group so that OnDemand maintains data written to the system logging facility for the life of your OnDemand system. Also, if you define table space file systems to OnDemand, update the System Log application group to store index data in table spaces.

__ 3. Configure the System Migration application group (refer to "Chapter 16. Configuring the System Migration application group" on page 73). If you plan to maintain index data on archive media, you must assign the System Migration application group to a storage set that writes data to archive media. Also, if you define table space file systems to OnDemand, update the System Migration application group to store index data in table spaces.

__ 4. Backup the database (see "Chapter 17. Creating a back up copy of the database" on page 75). After configuring the system, we recommend that you create a full backup image of the database.

Chapter 2. Hardware and software requirements

Server hardware requirements

The exact hardware configuration that you need to support your OnDemand for Windows NT system depends on the volume of data you plan to manage and the number of concurrent users the system must support. The PC processor family supports a wide range of configurations, with the high-end capable of supporting one gigabyte of memory, 256 gigabytes of DASD, and over two terabytes of optical storage from a single processor.

The minimum configuration for an OnDemand for Windows NT server requires the following:
- A 200 MHz Pentium Pro or 233 MHz Pentium II or faster PC
- A minimum of 128 megabytes of memory.
- Token Ring or Ethernet network adapter.
- Magnetic and optical storage for the database, report data, and temporary work space. A minimum of eight gigabytes of space and two separate disks are required, in addition to sufficient disk space for document storage.
- A CD-ROM drive for program installation.
- A 4mm or 8mm tape drive, automated tape library, or optical library for data backup and recovery. A tape drive is recommended for a small system configuration. An automated tape library or dedicated optical library is recommended for medium and large system configurations. The *Introduction and Planning Guide* provides details about system configurations.

Server software requirements

Version and release levels are the minimum levels. Subsequent versions and releases are acceptable:
- Microsoft Windows NT Server Version 4.0.
- IBM DB2 Universal Database Enterprise Edition Version 5 (included with OnDemand for Windows NT) or Microsoft SQL Server Version 6.5.

Optical and tape storage

The recommended option for OnDemand for Windows NT to support tape and optical storage management, including archiving, requires:

7

- ADSM for Windows NT Version 3 Server with the appropriate optical or tape device support module.
- ADSM for Windows NT Administrative Client with a minimum of a 10 concurrent-user ADSM license purchased. More users may be required depending on the number of storage pools defined to support OnDemand.

If you plan to use ADSM to maintain files on optical and tape storage volumes, consult the documentation provided with ADSM for a list of the devices supported by the ADSM server. Table 1 lists the optical libraries manufactured by IBM and supported by ADSM.

Table 1. Optical Libraries

Device	Drives	Storage Cells	Uncompressed Capacity	SCSI IDs Required
IBM 3995-C60	1, 5.25 inch	20	52 GB	2
IBM 3995-C62	2, 5.25 inch	52	135 GB	3
IBM 3995-C64	2, 5.25 inch	104	270 GB	3
IBM 3995-C66	4 or 6, 5.25 inch	156	405 GB	5 or 7
IBM 3995-C68	4, 5.25 inch	258	671 GB	5

ADSM also supports optical libraries from Hewlett-Packard, ATG Cygnet, Kodak, DISC, and Phillips/LMS, tape libraries from IBM and StorageTek, and many other storage devices.

PSF/MVS: MVS Download feature

To enable seamless download from MVS systems to the server, the following is required:
- IBM TCP/IP for MVS Version 3.1 or later.
- PSF/MVS Version 2.2 or later with the MVS Download feature.

Enhanced indexing for MVS

Enhanced indexing for MVS systems requires:
- One of the following operating systems:
  - OS/390 Version 2.4 or later.
  - MVS/ESA SP JES2 Version 4.2.2.
  - MVS/ESA SP JES3 Version 4.2.2.
  - MVS/ESA SP JES2 Version 5.1.

- – MVS/ESA SP JES3 Version 5.1.
- PSF/MVS Version 2.2 or later.

MVS host indexing can be processed with either the standard ACIF (free with PSF/MVS) or an enhanced version of ACIF that can be ordered without charge by customers who are entitled to OnDemand for Windows NT.

OS/2 client

The OnDemand OS/2 client runs under IBM OS/2 Version 3.0 or later and IBM OS/2 Warp Connect and requires the following hardware and software:
- Physical connection to the network, such as a Token Ring or Ethernet network adapter.
- A minimum of 16 megabytes of memory.
- An 80386 or faster processor; an 80486 processor is recommended when viewing AFP or image documents.
- A super-VGA display and adapter with a minimum resolution of 800x600.
- A minimum of 50 megabytes of free hard disk space.
- For OS/2 Version 3.0 or later, the IBM OS/2 TCP/IP Version 2.0 Base kit is required.
- For OS/2 Warp Connect, the standard Warp Connect TCP/IP support is required.

Windows client

The OnDemand Windows client runs under Microsoft Windows 3.1 or later, Windows NT 4.0, and Windows 95 and requires the following hardware and software:
- Physical connection to the network, such as a Token Ring or Ethernet network adapter.
- A minimum of 12 megabytes of memory.
- An 80386 or faster processor; an 80486 processor is recommended when viewing AFP or image documents.
- A super-VGA display and adapter with a minimum resolution of 800x600.
- A minimum of 50 megabytes of free hard disk space.
- For Windows 3.1 or later, a Windows-compatible TCP/IP socket program, such as the IBM TCP/IP for DOS Version 2.1.1 or later.

   Note: Other TCP/IP packages compatible with the Windows sockets API may also be supported. Contact your IBM representative for the latest information about other TCP/IP packages.

- For Windows NT 4.0 and Windows 95, the standard TCP/IP support is required.

Administrative clients

There are two administrative clients provided with OnDemand:

- OnDemand Configurator. The client used to configure the server, maintain services and maintenance tasks, and schedule server programs. The configurator runs on an OnDemand server.
- Administrator Interface. The client used to maintain users, groups, application groups, applications, folders, storage sets, and printers. The administrator interface can run on any Windows NT or Windows 95 system.

The administrative functions run under Microsoft Windows NT 4.0 and Windows 95 and require the following hardware and software:

- Physical connection to the network, such as a Token Ring or Ethernet network adapter.
- A minimum of 16 megabytes of memory.
- An 80386 or faster processor; an 80486 processor is recommended.
- A super-VGA display and adapter with a minimum resolution of 800x600.
- A minimum of 18 MB of free hard disk space.
- The standard Windows NT or Windows 95 TCP/IP support is required.

OS/2 client setup program

The OnDemand OS/2 client setup program runs under IBM OS/2 Version 3.0 or later and IBM OS/2 Warp Connect and requires the same basic hardware and software as the OS/2 client (refer to "OS/2 client" on page 9), except as noted below:

- Up to 50 MB of free hard disk space, depending on the setup options selected.

Windows client setup program

The OnDemand Windows client setup program runs under Microsoft Windows 3.1 or later and requires the same basic hardware and software as the Windows client (refer to "Windows client" on page 9), except as noted below:

- Up to 50 MB of free hard disk space, depending on the setup options selected.

Administrator interface setup program

The administrator interface setup program runs under Microsoft Windows NT 4.0 and Windows 95 and requires the same basic hardware and software as the administrative clients (refer to "Administrative clients" on page 10), except as noted below:

- Up to 50 MB of free hard disk space, depending on the setup options selected.

Chapter 3. Choosing a configuration

Before you install your OnDemand system, you need to learn about the different configurations and components. The *Introduction and Planning Guide* provides details about the different configurations and the components required on each type of server. You should also read the README file on the OnDemand CD-ROM. Then you need to choose the components that you want to install and prepare to install those components by setting up your environment. Setting up your environment consists of two tasks: (1) planning for sufficient performance and capacity and (2) installing the platform-specific prerequisites.

OnDemand supports several installation configurations. However, most customers will install one of two configurations:
- Standard library/object server system. A single workstation that includes a complete OnDemand system and performs both the library server and object server functions.
- Distributed library/object server system. Multiple workstations make up the OnDemand system. One workstation performs the library server role and one or more workstations provide object server functions.

Note: In a production environment, we recommend that you dedicate each workstation to OnDemand work and processes. That is, you should not run other applications on a workstation with an OnDemand library or object server. For example, in a Windows NT domain, you should not install OnDemand software on a PDC, BDC, or other application server.

Library server

The 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 and Windows NT Server and can use IBM DB2 Universal Database or Microsoft SQL Server to manage the library contents. An 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 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 accomodate 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 OnDemand system.

Installation considerations for the library server

When you create an OnDemand instance[1], a single language is defined. The default locale is `English` and the default code page is `1252`. If you need to change these defaults, you must do so before you create the instance. Refer to "OnDemand National Language Support," on page 77 for more information.

Object server

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 OnDemand administrative programs. The object server supports attachment of disk and ADSM-managed storage devices. An OnDemand system can have many object servers distributed across networks to provide convenient user access. Object servers run on AIX and Windows NT Server.

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

---

1. An OnDemand instance is a logical OnDemand server environment, with one library server (and database) and one or more object servers. Each OnDemand instance (the servers and the database) must run in a single code page.

- Prerequisite software and OnDemand program files
- Staging areas for data download and indexing
- Temporary spaces for loading and printing
- Document cache
- Long-term storage of documents

The *Introduction and Planning Guide* provides details, examples, and worksheets to help you estimate the amount of storage needed to support your environment.

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

ADSM considerations for object servers

OnDemand uses ADSM to maintain object storage on media other than fixed disks. To store the primary copy of a document, the object server writes to ADSM using an ADSM archive copy group. Because there is no need for multiple generations of the object server's objects, we recommend that you set the maximum number of backup copies to 1 (one). You can choose to maintain more than one copy of data in a storage pool. However, you must define the backup copy to ADSM 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 4. Windows NT security basics

You should take the following precautions on the workstations that you plan to use as OnDemand library or object servers:

- Disable the Guest account
- Rename the Administrator account and assign the account a unique password
- Prohibit local and network logons, except for users in the Administrators group
- Limit access to the Administrators group

You should also secure resources on your system to prevent authorized local or network access:

- When a volume is formatted with NTFS, the special group Everyone is automatically given Full Control permission to the volume. Folders and files created on the volume inherit this default permission. If you make no changes to the volume's security settings, any user that can log on locally to the server can add, change, and delete files and folders on the volume. Since, by default, all volumes are created in this manner, it means that the volumes that contain operating system and application files, databases and log files, and documents and other types of application data may be exposed to unintentional or malicious damage or tampering. We strongly encourage you to secure all volumes. At a minimum, you should provide the Administrators group Full Control access to each volume and change the permission of the Everyone group to Read Only. Apply these changes to all existing folders and files on each volume.
- When a volume is formatted with FAT, there is no way to secure it from users who can log on locally. Shared folder permissions can be used to secure a FAT volume when users connect to the server over the network. Since the default shared folder permission is Full Control for the Everyone group, you should change the shared folder permissions before you share volumes that contain operating system and application files, databases and log files, and documents and other types of application data.

Chapter 5. OnDemand system administrator account

You need to create a user account that will be used to install OnDemand software products, administer the system, load data, and perform other OnDemand functions. The Username must meet DB2's naming rules.

Create the **odadmin** user account with User Manager. Assign the account a password. Include the following characteristics:

- Under Description, enter: **OnDemand administrator account**
- Add the account to the local group **Administrators** on the computer where OnDemand is to be installed.
- All logon hours are allowed.
- Grant the **Act as part of the operating system** advanced user rights policy.

Part 2. Installing software

This section of the book contains instructions for installing OnDemand software on Windows NT servers. The required software includes the OnDemand product files and a relational database management product, such as IBM DB2 Universal Database. This section also describes how to install optional software, such as ADSM, which is required if you plan to maintain data on archive media.

The installation process for each product is designed to be efficient and intuitive. The installation program guides you through the process, prompting you for the necessary information. In most cases, we recommend that you accept the default values suggested by the installation program.

Note: The typical OnDemand system uses ADSM to manage archive media. If you plan to use another product to manage archive media, you must obtain installation instructions from the product vendor and detail information about how to configure the product for use with OnDemand.

Chapter 6. Installing OnDemand

1. Log on with the OnDemand system administrator account (refer to "Chapter 5. OnDemand system administrator account" on page 19 for details).

2. Insert the OnDemand for Windows NT CD-ROM into the CD-ROM drive. The setup program automatically starts after you load the CD-ROM into the drive.

3. Read the Welcome screen and then click Next. The License Agreement window appears.

4. Select Yes to accept the license agreement. Click Next. The Destination window appears.

5. We recommend that you accept the default directory name. Click Next. The Select Components window appears.

6. Select the type of server you want to install:

   - Typical - Library server and object server on the same workstation
   - Compact - Object server only
   - Custom - You choose the components to install. After selecting components, click Next to continue.

   After making your choice, the Select Program Folder window appears.

7. We recommend that you accept the default folder name. Click Next. The Start Copying Files window appears.

8. Click Next. The progress window appears.

9. When the process completes, the Setup Complete windows appears.

10. If you want to view the README file now, click Finish. Otherwise, clear the View README checkbox and then click Finish.

11. Click Finish to complete the installation and restart the computer.

Chapter 7. Installing DB2 UDB

OnDemand stores system information and user-defined index data in a relational database. OnDemand for Windows NT supports the IBM DB2 UDB for Windows NT and Microsoft SQL Server relational database management products.

This topic provides an overview of installing DB2 UDB. If you need more information about installing DB2 or you encounter problems installing DB2, please refer to the *DB2 Universal Database for Windows NT Version 5 Quick Beginnings* publication.

Important: After installing the base DB2 software you must immediately apply the latest DB2 FixPak. See "Installing DB2 FixPak" on page 26 for more information.

1. Log on with the OnDemand system administrator account (refer to "Chapter 5. OnDemand system administrator account" on page 19 for details).

2. Insert the DB2 UDB for Windows NT CD-ROM into the CD-ROM drive. The setup program automatically starts after you load the CD-ROM into the drive.

3. Read the Welcome screen and then click Next. The Select Products window appears.

4. Select DB2 Universal Database Enterprise Edition. Click Next. The Setup Options window appears.

5. Select the setup option. We recommend that you select the Typical option. After you select the setup option button, the Destination Directory window appears.

6. We recommend that you accept the default directory name. Click Next. The User Name and Password window appears.

7. Enter the username and password of the OnDemand administrator account. Click Next. The Start Copying Files window appears.

8. Click Next. The progress window appears.

9. When the process completes, the Setup Complete windows appears.

10. Click Finish to complete the installation and restart the computer.

When the system restarts, log on with the OnDemand system administrator account. The DB2 First Steps tool executes automatically. With the First Steps tools, create the sample database and run a simple query to verify the installation of DB2.

Installing DB2 FixPak

After installing the base DB2 software, you must immediately apply the latest DB2 FixPak. The FixPak is required for proper operation of DB2 with OnDemand. The current FixPak as of this publication date is WR09014. However, FixPaks are cummulative and you may apply the latest FixPak available.

You can obtain the FixPak from IBM Service on the World Wide Web. The URL is:

    `ftp://service.software.ibm.com/ps/products/db2/fixes`

- From the `fixes` directory, go to the directory for your language. For example, `english-us`
- Go to the `db2ntv5` directory
- Go to the `%L9014` directory, where `%L` is the locale, for example, `us`
- Print the `readme.1st` and `readme.txt` files
- Get the `%L9014.zip` file

Follow the instructions in the readme files to apply the service.

Chapter 8. Installing SQL Server

OnDemand stores system information and user-defined index data in a relational database. OnDemand for Windows NT supports the IBM DB2 UDB for Windows NT and Microsoft SQL Server relational database management products.

This topic provides an overview of installing SQL Server. If you need more information about installing SQL Server or you encounter problems installing SQL Server, please refer to the *Microsoft SQL Server Setup* publication.

When you install SQL Server, you must log on to Windows NT with an account that belongs to the Administrators group.

1. Create a SQL Executive user account. Assign the account a password. Include the following characteristics:
   - The user cannot change the password.
   - The password never expires.
   - All logon hours are allowed.
   - Grant the **Log On As A Service** advanced user rights policy.
   - Add the account to the local group **Administators** on the computer where SQL Server is installed.
2. Insert the SQL Server CD-ROM into the CD-ROM drive.
3. Run the Setup program from the \I386 directory.
4. The options dialog box appears with the Install SQL Server and Utilities option selected. Click Continue to begin the installation process.
5. The setup program prompts you to type your name, company, and product ID. Click Next. The Destination Directory window appears.
6. We recommend that you accept the default directory name.
7. After accepting or changing the installation drive and directory, you must provide configuration information about:
   - Licensing mode
   - MASTER device creation
   - SQL Server Books Online
   - Installation options, including character set, sort order, additional network support, and auto-start options
   - SQL Executive log on account

   Follow the on-screen instructions for each option. For additional instructions, consult the online Help.

8. After you verify all installation options, Click Next. The progress window appears.
9. When the process completes, the Setup Complete windows appears.
10. Click Finish to complete the installation and restart the computer.

When the system restarts, log on with an account that belongs to the Administrators group. Verify the installation by completing the following steps:

- Verify that the SQL Server services are started
- Connect to the server and run a simple query.

Chapter 9. Installing ADSM

OnDemand uses ADSM to manage data stored on archive media.

This topic provides an overview of the installation process. If you need more information about installing ADSM or you encounter problems installing ADSM, please refer to the *ADSM for Windows NT Version 3 Quick Start* publication.

1. Log on with the OnDemand system administrator account (refer to "Chapter 5. OnDemand system administrator account" on page 19 for details).

2. Insert the ADSM for Windows NT CD-ROM into the CD-ROM drive. The setup program automatically starts after you load the CD-ROM into the drive.

3. Read the Welcome screen and then click Next. The Setup Options window appears.

4. Select the setup option. We recommend that you select the Typical option. After you click the setup option button, the Client Location window appears.

5. Click Next to install the client and administrator in separate directories. The Language window appears.

6. Select the language. Click Next. The Destination Directory window appears.

7. We recommend that you accept the default directory name. Click Next. The Select Program Folder window appears.

8. We recommend that you accept the default folder name. Click Next. The Start Copying Files window appears.

9. Click Next. The progress window appears.

10. When the process completes, the Setup Complete windows appears.

11. Click Finish to complete the software installation and **restart the computer later**.

12. Click Yes to run the server utilities. Verify the server options. Then exit the server utilities.

13. In Control Panel, Services, change the ADSM service Startup Type to Automatic.

14. Shutdown and restart the system.

Part 3. Configuring the system

This section of the book provides a guide to configuring OnDemand servers, creating the OnDemand database, and configuring services to support OnDemand system operations.

Configuration tasks include the following:

- Configure ADSM, if you plan to use ADSM to maintain data on archive media.
- Configure OnDemand servers: defining servers, creating instances, and configuring the database manager and the storage manager.
- Create and initialize the OnDemand database.

Note: The typical OnDemand system uses ADSM to manage archive media. If you plan to use another product to manage archive media, you must obtain detail information from the product vendor about how to configure the product for use with OnDemand.

Chapter 10. Configuring ADSM

This chapter provides general information and examples about how to configure ADSM Version 3 to maintain data on archive media. ADSM can maintain reports that you store in OnDemand and migrated index data. If you do not plan to use ADSM to support any of these functions, skip this chapter and continue with "Chapter 11. Configuring OnDemand servers" on page 51.

Before you begin, we recommend that you familiarize yourself with Configuring and Managing Server Storage in the *ADSM for Windows NT Version 3 Administrator's Guide.* In addition, the *ADSM for Windows NT Version 3 Administrator's Reference* provides detailed information about all of the commands used in this chapter and should be your primary reference when you work with ADSM. Refer to the ADSM publications if you encounter problems configuring ADSM or if the examples that follow do not provide the information you need to implement your server storage devices, policies, and operations.

In general, initial configuration of ADSM to support OnDemand consists of:

- Verifying the server and client options.
- Registering ADSM licenses.
- Registering the OnDemand ADSM administrator.
- Configuring devices.
- Defining devices and storage pools.
- Defining policy domains.
- Preparing storage pool volumes.
- Backing up the ADSM database and saving critical files.

After you complete the initial configuration of ADSM, you may need to perform advanced configuration of ADSM:

- Implementing a copy storage pool.

Verify server and client options

The dsmserv.opt file contains options for the ADSM server. You can verify and set server options by starting ADSM Server Utilities from the ADSM program group. Then open Server Options and following the instructions provided.

The following example shows server options that enable TCP/IP as the communications protocol and set the port address of the server, specify a

maximum of 255 simultaneous client sessions, establish values for the database and recovery log buffer pools, expire data once a day, and specify the names of volume and device configuration files:

```
COMMmethod TCPIP
TCPPort 1500
TCPWindowsize 0
MAXSessions 255
COMMTimeout 60
IDLETimeout 60
BUFPoolsize 2048
LOGPoolsize 512
EXPInterval 24
VOLUMEHistory volhist.out
DEVCONFig devcnfg.out
```

Note: Verify the language settings for the server. Language settings, such as the locale and display format, can be found on the Language page in Server Options.

The dsm.opt file contains options for the ADSM client. You can verify and set client options by starting Backup Client Options from the ADSM program group. Remove the comment character (*) from column one of the lines you want the server to read. To change the value of a parameter, replace the default value in the file.

The following example shows a client options file that sets the fully qualified TCP/IP host name of the server, enables TCP/IP as the communications protocol, sets the TCP/IP port address of the server, and turns off compression:

```
TCPServeraddress    fully.qualified.tcpip.host.name
COMMmethod          TCPip
TCPPort             1500
COMPRESSION OFF
```

Register licenses

When you install ADSM for Windows NT, your system is licensed for the base ADSM support, which provides the following:

*   An unlimited number of administrative clients
*   One backup-archive client
*   Certain types of removable media devices

License certificate files are included with your server package in the server installation directory. Each license certificate file licenses the server for a single option or feature. For current information about devices supported by ADSM, contact the IBM Support Center.

To register ADSM licenses:

1. Start ADSM Server from the ADSM program group.
2. To register any ADSM licenses beyond the base support, issue the REGISTER LICENSE command to read in a license certificate file. For example, to license an additional ten clients, the TCP/IP communication method, and all of the supported devices, issue the following commands:

```
register license file(10client.lic)
register license file(network.lic)
register license file(advdev.lic)
```

Register OnDemand administrator

We recommend that you register the OnDemand administrator account as an ADSM administrator and grant the account the system authority level. For details about authority levels, refer to the *ADSM for Windows NT Version 3 Administrator's Guide*.

Enter the following commands at the ADSM Server command line interface:

```
register admin odadmin ondemand
grant authority odadmin classes=system
set passexp 9999
set actlogretention 365
```

Replace the string ondemand with the password that you assigned to the OnDemand administrator account.

When complete, stop (halt) the ADSM Server command line interface.

Configure devices

This topic provides an overview of configuring devices. When you configure a device, you may install a controller card, install drives in an automated library, attach the device to a controller card, and identify device names. The *ADSM for Windows NT Version 3 Administrator's Guide* provides details about how to configure devices. In addition, you'll need documentation provided by the manufacturer of each device you plan to attach to the server. The manufacturer's documentation explains how to install drives in an automated library, attach the device to the system, and configure it for use.

After you physically attach and define storage devices to the server, you must define them to ADSM so that ADSM can use and manage the devices and their media.

For the following example, an IBM 3995-C68 library containing four drives is attached to the server. You can use this example as a guide when configuring other automated devices. The example presents the procedure with a minimum of customization. Refer to the *ADSM for Windows NT Version 3 Administrator's Guide* if you want to do more.

Note: In most installations, a customer engineer (CE) installs drives in an automated library and sets the SCSI IDs for each drive. It is imperative that the customer engineer provide you with a completed device configuration worksheet. The worksheet must contain the SCSI ID and element number assigned to each drive.

1. Install the SCSI adapter card in the server, if it is not already installed.
2. Determine the SCSI IDs available on the SCSI adapter. Find one unused SCSI ID for each drive and one unused SCSI ID for the library controller (also called autochanger).

   Finding unused SCSI IDs: Your SCSI adapter card may have come with a utility that can help you find out the SCSI IDs that are in use. You can also use the device detector utility included with ADSM to determine the available SCSI IDs. From the ADSM program folder, open the ADSM Server Utilities, then open Diagnostics.

   Note: In some automated libraries, the drives and the autochanger share a single SCSI ID, but have different LUNs. For these libraries only a single SCSI ID is required. Check the documentation for the device.

3. Follow the manufacturer's instructions to set the SCSI ID for each device to one of the unused SCSI IDs that you found. Usually this means setting switches on the back of the device.

   Attention: Each device connected in a chain to a single SCSI bus must be set to a unique SCSI ID. If each device does not have a unique SCSI ID, you may have serious system problems.

4. Follow the manufacturer's instructions to attach the device to the server system hardware.

   Attention:
   a. To prevent damage to the hardware, power off your system before attaching a device.
   b. You must attach a terminator to the last device in the chain of devices connected on one SCSI adapter card. Detailed instructions should be in the documentation provided with your hardware.

5. Make sure that you have started the ADSMSCSI device driver, changed the startup from manual to boot, and enabled optical support. Use the ADSM device utilities to start the device driver and modify its properties.
6. Determine the name for each device and record the names on a worksheet. Table 2 shows an example.

Table 2. Example Worksheet: IBM 3995-C68

Device	SCSI ID	Device Name
Optical Drive 1 (element 1)	0	op0,0,0,0
Optical Drive 2 (element 2)	1	op1,0,0,0
Optical Drive 3 (element 3)	2	op2,0,0,0
Optical Drive 4 (element 4)	3	op3,0,0,0
Autochanger (element 6)	6	lb6,0,0,0

The device name is required to define the device to ADSM. For an optical drive, the device name at a minimum has the form **op**x, where x is the SCSI ID for the device on the SCSI adapter card. For a library controller, the device name at a minimum has the form **lb**x, where x is the SCSI ID for the device on the SCSI adapter card. The element address is a number that indicates the physical location of a drive within an automated library.

You can use the device detector utility included with ADSM to determine the device name. From the ADSM program folder, open the ADSM Server Utilities, then open Diagnostics.

Use the Device Configuration Worksheet for your library to determine the element addresses. Device Configuration Worksheets are provided in the *ADSM for Windows NT Version 3 Administrator's Guide* for all devices supported by ADSM.

Define devices

In general, the following steps are required before ADSM can use a newly configured library:
1. Define the devices to ADSM. You must define the storage objects that represent the physical device and media, such as the library, drives, device class, and storage pool.
2. Define an ADSM policy that links data with media for the new library.
3. Register client nodes to the policy. This step links clients and their data with storage volumes and devices.
4. Prepare volumes for use by the device. At a minimum, you must label volumes for the device. For automated libraries, you must also add the volumes to the device's volume inventory by checking in the volumes.

This topic explains how to define devices to ADSM. Later in this chapter we'll explain how to define a policy, register client nodes, and prepare storage volumes.

In the following example, we'll define an IBM 3995-C68 library containing two drives to ADSM. You can use this example as a guide when defining other automated devices. The example presents the procedure with a minimum of customization. If you want to do more, refer to the *ADSM for Windows NT Version 3 Administrator's Guide.*

1. Define the library to ADSM. The library type is `SCSI` because the library is a SCSI-attached automated library. Enter the following command at the ADSM Admin Command Line interface:

   ```
   def library archlib0 libtype=scsi device=lb6,0,0,0
   ```

   The string `archlib0` is an arbitrary name that you assign to the library. Make sure that the name is not already defined to ADSM. The `device` parameter gives the device name for the library. The name was determined in step **6** on page 37.

2. Define the drives to ADSM. Enter the following commands:

   ```
   def drive archlib0 optdrv0 device=op0,0,0,0 element=1
   def drive archlib0 optdrv1 device=op1,0,0,0 element=2
   def drive archlib0 optdrv2 device=op2,0,0,0 element=3
   def drive archlib0 optdrv3 device=op3,0,0,0 element=4
   ```

   The drives belong to the `archlib0` library (which you defined in step 1). The strings `optdrv0`, `optdrv1`, `optdrv2`, and `optdrv3` are arbitrary names you assign to the drives. Make sure the names are not already defined to ADSM. The `device` parameter gives the device name for the drive. The name was determined in step **6** on page 37. The element parameter identifies the physical location of the drive. The element address is required if there is more than one drive in an automated library. The element address was determined in step **6** on page 37.

3. Classify drives according to type and recording format by defining ADSM device classes. The following example shows how to define a device class for optical devices that use the 2600 recording format. The library contains four, 5.25 inch optical drives with an estimated unformatted capacity of 2600 MB each. ADSM keeps idle optical platters in the drives for ten minutes before dismounting them. Enter the following command:

   ```
   def devclass odlib0 devtype=optical format=2600 -
                mountlimit=4 library=archlib0 -
                estcapacity=2600M mountretention=10
   ```

   The string `odlib0` is an arbitrary name that you assign to the optical device class. Make sure the name is not already defined to ADSM. The string `archlib0` is the library that you defined in step 1.

38

If the optical library you are defining to ADSM uses 650 MB drives, change the value of the `format` and `estcapacity` parameters. If the optical library you are defining has fewer or more than four drives, change the value of the `mountlimit` parameter accordingly. If the optical library you are defining to ADSM uses WORM optical platters, change the value of the `devtype` parameter to `WORM`.

You can set the `mountretention` parameter to a larger value, if you are configuring a storage subsystem where the most frequently accessed data resides on the same storage volumes or if you expect a low volume of retrieval activity to data stored in the library. The `mountretention` parameter governs how long ADSM retains idle storage volumes before dismounting them. For example, if you set the value to 360, ADSM retains idle storage volumes for six hours before dismounting them.

4. Create a storage pool to use the devices in the device class you just defined. For example, define a storage pool named `odnt1` associated with the device class `odlib0`. Enter the following command:

```
def stgpool odnt1 odlib0 -
            desc='Optical storage for OnDemand' -
            collocate=no reclaim=60 -
            maxscratch=1 reusedelay=30
```

The string `odnt1` is an arbitrary name that you assign to the storage pool. Make sure that the name is not already defined to ADSM. The string `odlib0` is the device class that you created in step 3 on page 38.

We recommend that you specify `collocate=no`. Refer to the ADSM documentation for more information.

The `reclaim` parameter specifies when to initiate reclamation of fragmented space on volumes in the storage pool, based on the percentage of reclaimable space on a volume. Reclamation makes the fragmented space usable again by moving any remaining active files from one volume to another volume, thus making the original volume available for reuse. Refer to the ADSM documentation for details.

The `maxscratch` parameter can be used to enable ADSM to automatically mount new scratch volumes when needed, without intervention from an ADSM administrator. Setting the value to `1` means that ADSM can mount a new scratch volume once. Use caution when setting this parameter to a larger value. Refer to the ADSM documentation for more information.

The `reusedelay` parameter determines when a storage volume can be rewritten or returned to the scratch pool. That is, the number of days that must elapse after all of the files on a storage volume have been deleted

before ADSM can write over the volume or return the volume to the scratch pool. Refer to the ADSM documentation for details.

5. To check what you have defined, enter the following commands:

```
query library
query drive
query devclass
query stgpool
```

Define policy domains

An ADSM policy domain links data with media for the new library. A policy domain supports a single storage pool, which in turn supports a single library.

Note: When you create an application group, you specify the length of time that OnDemand maintains data stored in the application group. If you plan to store application group data in ADSM, you must make sure that the policy domain where you store the data maintains the data for the required length of time. The retention period that you specify for the domain must be equal to or greater than the life of the application group data.

For the following example, we'll define a policy domain to maintain data for a period of seven years (2557 days). You can use this example as a guide when defining other policy domains. The example presents the procedure with a minimum of customization. If you want to do more, refer to the *ADSM for Windows NT Version 3 Administrator's Guide*.

1. Define the domain. Enter the following command at the ADSM Admin Command Line interface:

```
define domain OD7YRPD desc='OnDemand Policy Domain'
```

Replace the string OD7YRPD with the name of an ADSM domain that you plan to use to manage storage for OnDemand. Make sure that the name is not already defined to ADSM.

2. Define the policy set:

```
define policyset OD7YRPD OD7YRPS desc='OnDemand Policy Set'
```

Replace the string OD7YRPD with the name of the domain that you defined in step 1. Replace the string OD7YRPS with the name of the policy set. Make sure that the name is not already defined to ADSM.

3. Define the management class:

```
define mgmtclass OD7YRPD OD7YRPS OD7YRMG -
                desc='OnDemand Management Class'
```

Replace the string OD7YRPD with the name of the policy domain that you defined in step 1 on page 40. Replace the string OD7YRPS with the name of the policy set that you defined in step 2 on page 40. Replace the string OD7YRMG with the name of the management class. Make sure that the name is not already defined to ADSM.

4. Define an archive copy group. The copy group identifies the name of the storage pool where ADSM stores the data and the length of time that ADSM maintains the data:

```
define copygroup OD7YRPD OD7YRPS OD7YRMG standard -
                  type=archive dest=odnt1 retver=2557
```

Replace the string OD7YRPD with the name of the policy domain that you defined in step 1 on page 40. Replace the string OD7YRPS with the name of the policy set that you defined in step 2 on page 40. Replace the string OD7YRMG with the name of the management class that you defined in step 3 on page 40. Replace the string odnt1 with the name of the storage pool that you defined in step 4 on page 39. Replace the string 2557 with the length of time (in days) that you want ADSM to maintain data in the storage pool.

5. To check what you have defined, enter the following commands:

```
query domain
query policyset
query mgmtclass
query copygroup
```

6. After you define your policy sets and the management classes they contain, you must assign a default management class for each policy set. To assign the OD7YMG management class as the default management class for the OD7YRPS policy set in the OD7YRPD policy domain, enter:

```
assign defmgmtclass OD7YRPD OD7YRPS OD7YRMG
```

Replace the string OD7YRPD with the name of the policy domain that you defined in step 1 on page 40. Replace the string OD7YRPS with the name of the policy set that you defined in step 2 on page 40. Replace the string OD7YRMG with the name of the management class that you defined in step 3 on page 40.

7. After you define a default management class for the policy set, validate and activate the policy set. Enter the following commands:

```
validate policyset OD7YRPD OD7YRPS
activate policyset OD7YRPD OD7YRPS
```

Replace the string OD7YRPD with the name of the policy domain that you defined in step 1 on page 40. Replace the string OD7YRPS with the name of the policy set that you defined in step 2 on page 40.

Register client nodes

A client node links clients and their data with storage volumes and devices. Before OnDemand can store data in ADSM storage, you must register at least one client node. You must register at least one client node in each policy domain that will contain OnDemand data.

Note: When you define a storage node to OnDemand, you identify an ADSM client node and specify the client node password.

In the following example, we'll register the client node PRI7YR and assign it to the OD7YPD policy domain (that we created in step 1 on page 40). The node password is password. You can use this example as a guide when registering other client nodes. The example presents the procedure with a minimum of customization. If you want to do more, refer to the *ADSM for Windows NT Version 3 Administrator's Guide.* Enter the following command at the ADSM Admin Command Line interface:

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

Prepare storage pool volumes

ADSM allows you to use and reuse removable media to store data. You need to perform some steps to prepare removable media for initial use by ADSM. This topic provides general information and examples showing how to label storage pool volumes and check them into an automated library. For details about these tasks and important information about managing removable media operations, please refer to the *ADSM for Windows NT Version 3 Administrator's Guide.*

To prepare a storage volume:
1. Label the volume. Any volumes associated with the optical or tape device types must be labeled before the server can use them.
2. For storage pools in automated libraries, check the volume into the library.

You can use the LABEL LIBVOL command to label and check in a volume in one operation. Provide parameters that specify:
• The name of the library
• The name of the storage volume
• Whether to overwrite a label on the volume. By default, the label command does not overwrite an existing label on a volume.
• Whether to search the library for volumes for labeling. The label command can search all of the storage slots in the library for volumes and try to label each one that it finds.

- Whether to read media labels and prompt for volume names or read the bar-code label of each cartridge. If the library has a bar-code reader, the label command can use the reader to obtain volume names. Otherwise, the label command will prompt you for volume names.
- Whether to check in the volume

The following example shows how to label a new storage volume for use in the library named `archlib0`. You will manually insert the new volume into the library, and you want the volume to be placed in a storage slot inside the library after the label is written. You know that the new volume does not contain any data, so it is acceptable to overwrite an existing volume label. Replace the string `<volume label>` with the label you want ADSM to write on the storage volume. You want to use all of the library's drives for the label operation. Enter the following command at the ADSM Admin Command Line interface:

```
label libvolume archlib0 <volume label> overwrite=yes checkin=scratch
```

Wait for the LABEL LIBVOL operation to complete. Because the label and check in commands involve device access, it may take a long time to complete.

When you label a volume, you must supply a label for the volume. Labels can be from 1 to 6 alphanumeric characters in length. You can use the characters A through Z and 0 through 9. The label cannot include embedded blank characters. ADSM converts lowercase letters in a volume label to uppercase.

When you have completed labeling and checking in all the storage volumes, enter the following command to list the storage volumes in the `archlib0` library:

```
query libvol archlib0
```

Back up the database and save critical files

After you configure ADSM, we recommend that you backup the database and save server files that contain important information. The backup copy of the database can be used if you need to recover the database. (And 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.

Note: We recommend that you backup the database and save server files whenever you make changes to the database. The database is modified when you store data in ADSM and when you make changes to the ADSM environment, such as defining new devices and managing removable media operations.

Backing up the database

Before you backup the database, you must define the backup storage objects to ADSM 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. For example, you can write full backups to a tape device and incremental backups to a disk device. At a minimum, we recommend that you define a tape backup device, device class, library, and storage pool:

1. Start the ADSM Server Utilities.
2. Select Wizards.
3. Click the Device Configuration tab.
4. Choose define manual tape devices to ADSM.
5. Follow the instructions provided to define a device class, library, drive, and storage pool for the tape backup.

Next, label a tape storage volume:

1. Place a blank, unlabeled tape in the drive.
2. Click the Volume Labeling tab.
3. Choose label media on manual devices.
4. Follow the instructions provided to label one or more tape storage volumes.

After you have defined the storage objects to ADSM and labeled a tape storage volume, you can backup the database. First, place a labeled tape storage volume in the drive. Then enter the following command at the ADSM Admin 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.

Record 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

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)

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.

Implementing a copy storage pool

Overview

This chapter of the book describes how you can use ADSM to manage a backup copy of data stored on archive media. This capability means that OnDemand can automatically retrieve a backup copy of a report file from archive media, when needed, without administrator intervention.

Maintaining a backup copy of reports can increase the availability of your system. For example, assume you need to have backup copies of the files stored in the ODNT1 primary storage pool. All of the files stored in the ODNT1 storage pool are critical to the operation of the company. You can define a copy storage pool to contain the backup copies. The copy storage pool allows ADSM to maintain multiple copies of files on archive media, thus reducing the potential for data integrity loss because of media failure. If the primary file is not available or becomes corrupted, ADSM accesses and uses the duplicate file from the copy storage pool.

After you define the copy storage pool, you must copy all of the files that existed before you defined the copy storage pool. You can then schedule daily incremental backups of the files in the primary storage pool.

To use ADSM to independently manage more than one copy of a file on archive media, you must complete the following tasks:
- Identify the primary storage pool that contains the data that you need to back up.
- Identify a library and storage volumes for the copy storage pool.
- Define the library and drives to ADSM, including a device class for the library. Refer to topic "Defining the storage library" on page 47 for details.
- Define a copy storage pool. Topic "Defining the copy storage pool" on page 48 provides details.
- Copy the files from the primary storage pool to the copy storage pool. Topic "Copying existing files" on page 48 provides details.

- Define a schedule to incrementally backup the primary storage pool to the copy storage pool. Topic "Scheduling the backup copy" on page 49 provides details.

Refer to the *ADSM for Windows NT Version 3 Administrator's Guide* for details about configuring and managing server storage and protecting and recovering your data.

About storage pools

OnDemand stores copies of reports in cache storage, for short-term, high-speed access, and on archive media, to support long-term storage. ADSM maintains report data stored on archive media. When you load reports into an application group that identifies an ADSM client node, ADSM stores the data in a primary storage pool. When OnDemand requests a file from archive media, ADSM usually retrieves the requested file from a primary storage pool.

A copy storage pool stores copies of data that resides in a primary storage pool. With this method, ADSM maintains a backup copy of files that are stored in a primary storage pool independently and transparently to OnDemand. The backup copy can be used to restore the original file if it become damaged, lost, or unusable.

When you implement a copy storage pool, remember the following points:
- Copy storage pools require additional space in the ADSM database.
- The copy storage pool can be assigned to the same library as the primary storage pool. However, you would typically assign the copy storage pool to a different library.
- You can copy data from one or more primary storage pools to the same copy storage pool.
- A copy storage pool must be defined on the object server where the primary storage pool resides.

  To backup files to a storage node that resides on a different object server, you must define a secondary storage node to OnDemand and ADSM. Use the administrator interface to define a secondary storage node.
- A primary storage pool file must always exist for a copy storage pool file to exist. When a primary storage pool file is deleted, so is the copy of the file in the copy storage pool.

Copies in a copy storage pool are produced by using the BACKUP STGPOOL command. The copies are made incrementally while the server is operational. Central scheduling allows for producing copies at convenient times for your installation.

Defining the storage library

When you add an optical or tape library to the system, you must define it to ADSM. When you define a library to ADSM, you define a device class for the library and define the library and the drives contained in the library. You also define a storage pool for the collection of storage volumes that belong to the library.

The following example shows how to define an optical library, device class, and storage pool as the destination for files backed up from a primary storage pool. You can use this example as a guide when defining other copy storage pools. The example presents the procedure with a minimum of customization. If you want to do more, refer to the *ADSM for Windows NT Version 3 Administrator's Guide*.

Define the library

The following example shows how to define a SCSI-attached library:

```
def library archlib9 libtype=scsi device=lb6,0,0,0
```

The string `archlib9` is an arbitrary name that you assign to the library. Make sure that the name is not already defined to ADSM. The `device` parameter gives the device name for the library. "Configure devices" on page 35 describes how to determine device names.

Define the drives

The following example shows how to define the drives in a library that contains four drives, such as an IBM 3995-C68 optical storage library:

```
def drive archlib9 optdrv4 device=op0,0,0,0 element=1
def drive archlib9 optdrv5 device=op1,0,0,0 element=2
def drive archlib9 optdrv6 device=op2,0,0,0 element=3
def drive archlib9 optdrv7 device=op3,0,0,0 element=4
```

The string `archlib9` is the library that you defined in "Define the library". The strings `optdrv3`, `optdrv4`, `optdrv5`, and `optdrv6` are arbitrary names you assign to the optical disk drives. Make sure the names are not already defined to ADSM. The `device` parameter gives the device name for the drive. The `element` parameter identifies the physical location of the drive. "Configure devices" on page 35 describes how to determine device names and element addresses.

Define the device class

The following example shows how to define a device class for optical devices that use the 2600 recording format. The library contains four drives. ADSM keeps idle optical platters in the drives for ten minutes before dismounting them:

```
def devclass odlib9 devtype=optical format=2600    -
              mountlimit=4 library=archlib9          -
              estcapacity=2600M mountretention=10
```

The string `odlib9` is an arbitrary name that you assign to the optical device class. Make sure the name is not already defined to ADSM. The string `archlib9` is the library that you defined in "Define the library" on page 47.

Defining the copy storage pool

The following example shows how to define a copy storage pool to ADSM. All of the storage volumes in the copy storage pool are located on the same object server as the primary storage pool:

```
define stgpool ODNT1-COPY odlib9                    -
               desc='Copy storage pool for ODNT1'   -
               collocate=no reclaim=60               -
               maxscratch=1 reusedelay=30            -
               pooltype=copy
```

The string `ODNT1-COPY` is an arbitrary name that you assign to the storage pool. Make sure that the name is not already defined to ADSM. The string `odlib9` is the device class that you defined in "Define the device class".

Preparing copy storage pool volumes

You must define the same number of volumes in the copy storage pool as were already defined in the primary storage pool. If you copy data from more than one primary storage pool to the same copy storage pool, you must define the total number of volumes in all of the primary storage pools.

To prepare a storage volume, you must label the storage volume and check it into the library. Refer to "Prepare storage pool volumes" on page 42 for details.

Copying existing files

Before ADSM schedules the first incremental backup of the files in the primary storage pool, you must copy the files that existed in the primary

storage pool before you created the copy storage pool. If no data exists in the primary storage pool, skip this step and continue with "Scheduling the backup copy".

The following example shows how to copy the files from the primary storage pool to the copy storage pool:

```
backup stgpool ODNT1 ODNT1-COPY
```

Replace the string ODNT1 with the name of the primary storage pool. Replace the string ODNT1-COPY with the name of the copy storage pool.

If you copy data from more than one primary storage pool to the same copy storage pool, issue the backup storage pool command for each primary storage pool.

Scheduling the backup copy

ADSM includes a central scheduling component that allows the automatic processing of administrative commands, such as copying data from a primary storage pool to a copy storage pool. Each administrative command is called an event. Each scheduled event is tracked by the server and recorded in the database. You set up an administrative command schedule by defining schedule parameters, such as the start day, date, and time, specifying the command to be executed, such as the BACKUP STGPOOL command, and activating the schedule. The *ADSM for Windows NT Version 3 Administrator's Guide* provides details about scheduling operations.

The following example shows how to define a schedule for an incremental backup of the primary storage pool to the copy storage pool every day beginning at 11 p.m:

```
define schedule backup_ODNT1 type=administrative     -
               cmd='backup stgpool ODNT1 ODNT1-COPY'  -
               active=yes startime=23:00 period=1
```

Replace the string ODNT1 with the name of the primary storage pool. Replace the string ODNT1-COPY with the name of the copy storage pool.

If you copy data from more than one primary storage pool to the same copy storage pool, define a backup schedule for each primary storage pool.

Chapter 11. Configuring OnDemand servers

Overview

After installing software on the server, you need to configure OnDemand to integrate the various software products and control information, building your specific OnDemand operating environment. In general, initial configuration of an OnDemand Windows NT system consists of:

- Defining the server or servers
- Defining an instance on each server
- Specifying properties of the instance:
  - Server type and other options
  - NLS
  - Directories for OnDemand programs to use
  - Database manager options
  - Storage manager options
- Creating the instance
- Installing services

After you complete the initial configuration of your system, you may need to perform advanced configuration, such as:

- Configuring services
- Configuring scheduled tasks
- Managing multiple servers from one workstation
- Defining multiple instances on one server

Getting started

You configure servers using the OnDemand Configurator program.

To begin, select the OnDemand for WinNT program group and start OnDemand Configurator. The Configurator window contains a menu bar, toolbar, navigator pane, list pane, and status bar. Initially, the navigator pane contains no servers and the list pane contains no details. In just a minute, you'll add your first server.

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 NT. 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 NT help and for information about how to use Windows NT help, select Using Help from the Help menu.

Important things to remember

If your OnDemand system consists of more than one workstation, you must define each server and create an instance for each server.

After an instance is created, only the creater-owner of the instance can modify the instance. Other users can view the properties of the instance.

After an instance is created, the following properties of the instance cannot be changed:
- Instance name
- Server type
- Language and code page
- Database instance name
- Instance owner
- Database engine
- Location of database
- Size of the database (SQL Server)
- Location of primary log files (DB2)
- Size of log files
- Number of primary log files (DB2)
- First file system named (DBFS and Cache)

If you update an instance, you must stop and restart the system.

When defining a object server:
- We recommend that you 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 may need to reconfigure one or more of the services before you begin system operation. For example, you may want to configure the OnDemand MVSD service to start automatically on any server that will receive data from MVS systems using MVS Download.

When you create an instance, the configurator creates two 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.

Defining a server

1. From the File Menu, select New Server. The Add a Server dialog box appears.
2. Enter the information in the spaces provided.
3. When complete, click OK to define the server.

   The configurator adds an entry for the server in the navigator pane. A server contains three items: Instances, Services, and Scheduled Tasks. For initial configuration of a server, you must create an instance and configure scheduled tasks and enable them to run. If you plan to run the OnDemand data indexing and loading processes as a service or you plan to transmit data from MVS systems to the server using the MVS Download feature, you should verify the configuration of these services.

Defining an instance

1. Click Instances. The Instance dialog box appears.
2. Name the instance. For the first or only instance, we recommend that you use the default provided (ARCHIVE).
3. Click Next to continue and specify the properties of the instance:
   - Server type and other options
   - NLS parameters
   - Directories for OnDemand programs to use
   - Database manager options (library server only)
   - Storage manager options

Specifying properties of an instance

Server Type

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

     Important: If you set the Number of Database Servers to a value other than 0 (zero) or 1 (one), you must update the license information for your database management product. For example, with DB2, use the Nodelock Administration tool to update the number of concurrent licences (equal to or greater than the Number of Database Servers).

   - Number of OnDemand Licenses

   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.

Language

Note: An instance can run in one and only one code page.
1. On the Language page, verify the Language and Code Page.
2. To change the defaults provided, select a different Language. The configurator automatically selects the correct code page for the language you select.
3. Click Next to continue.

Directory

Note: Physically separating temporary work space, print work space, data index and load, and MVS download directories and other system and application data improves performance.
1. On the Directory page, define directories used by various OnDemand programs and utilities:
   - Temporary work space directory

- Print work space directory
- Data index and load directory
- MVS download directory

You must define at least one temporary work space, print work space, and data index and load directory. An MVS download directory is required if you plan to transmit data from MVS systems to the server using the MVS Download feature.

2. You can add one or more directories to each list. After you create the instance, you can always add more directories to each list.

3. When finished, click Next to continue.

Database

Note: The database properties are defined only for library servers.

Physically separating the database, the directories that contain the log files, and other system and application data improves performance and helps recovery.

1. On the Database page, select the Database Engine.
2. Click Advanced Database Options. The Advanced Options dialog box appears.
3. Configure the database options for DB2:
   - Database Location
   - Primary Log File Path
   - Archive Log Path
   - Log File Size
   - Number of Primary Log Files

   or SQL Server:
   - Database Path
   - Database Size
   - Transaction Log Path
   - Log File Size
4. Click OK to close the Advanced Options dialog box and return to the Database page.
5. If the Database Engine is DB2, optionally define DB2 File Systems. You can add one or more DMS and SMS file systems to the list. After you create the instance, the first DB2 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.
6. Click Next to continue.

Storage

> **Note:** Physically separating cache file systems and other system and application data improves performance and helps recovery.

1. On the Storage page, configure the storage manager. Select Cache Only or ADSM.
2. If you selected ADSM, click ADSM Options to verify the location of the ADSM program directory and client options file.
3. Define cache file systems. You can add one or more cache file system 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.

Creating an instance

1. Verify that you have completed all of the required options.
2. Verify options that cannot be altered after you create the instance.
3. On the Storage page, click Finish. The configurator displays the Install Services dialog box.

Installing services

Use the Install Services dialog box to identify the Windows NT user account and password that OnDemand uses to log on to its services. By default, OnDemand uses the current logon account. We recommend that you use the OnDemand system administrator account. Refer to "Chapter 5. OnDemand system administrator account" on page 19 for more information.

> **Note:** If the user account was not assigned a password in User Manager, you must select the This User ID does not have a password check box. Otherwise, you must enter and verify the password in the spaces provided.

After you enter the information in the spaces provided, OnDemand adds the instance, services, and tasks to the list. You can proceed to "Advanced configuration" on page 57 and learn about configuring services and scheduled tasks.

Advanced configuration

Configuring services

Note: 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 may need to configure the services before you begin system operation. For example, you may want to configure the OnDemand MVSD service to automatically start on any server that will receive data from MVS systems using MVS Download.

OnDemand provides the following services:

- OnDemand LibSrvr or OnDemand ObjSrvr, depending on the Server Type
- OnDemand Load Data, one or more, depending on the number of object servers assigned to the instance
- OnDemand MVSD, one or more, depending on the number of object servers assigned to the instance
- OnDemand Scheduler

You should verify the basic and advanced properties of the OnDemand Load Data service, if you plan to keep the OnDemand data indexing and loading processes running at all times. You should verify the basic properties of the OnDemand MVSD service, if you plan to transmit data from MVS systems to the server using the MVS Download feature.

To verify the properties of a service:

1. Point to the service and click the right mouse button.
2. From the pop-up menu, select Properties. The basic service properties dialog box appears.
3. Verify the Startup Type and other basic service properties. Refer to the online help for information about data you can enter.
4. For the Load Data Service, verify the advanced service properties. Refer to the online help for information about data you can enter.
5. Click OK to configure the service.

Note: If more than one object server belongs to the instance, you need to configure the OnDemand Load Data service and the OnDemand MVSD service on each object server that requires the service.

Configuring scheduled tasks

Note: When you create an instance, the configurator creates two 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, we recommend that you use the Database Maintenance Plan wizard from SQL Enterprise Manager to configure and schedule the tasks.

OnDemand provides the ability for you to set up tasks to run automatically on a regular schedule:

- ApplGroup Data Maintenance, maintains application group database tables and cache file systems
- System Table Maintenance, maintains OnDemand system tables

Note: Some systems may need to schedule more than one instance of a database maintenance task. For example, a task may be set up to run with specific options every day. Another occurrence of the task may 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.

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. The Properties dialog box appears.
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.

Note: If the instance contains more than one object server, you need to configure scheduled tasks on each object server.

Managing multiple servers

You can use the OnDemand Configurator to manage OnDemand servers locally or remotely. However, to remotely manage servers, you must establish a common Windows NT user account and password. That is, you must define the same account name and password on all of the workstations. Since we recommend that you create an OnDemand system administrator account when you install and configure OnDemand, this should not be a problem. Just

remember to log on to Windows NT using the OnDemand system administrator account whenever you want to maintain OnDemand servers 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. A resource on a remote server must be shared. Otherwise, it will not appear in the Browse list.

Defining multiple instances

Note: An instance name is unique to a server. You can use the 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, we recommend that you identify a library server and all of the object servers assigned to the library server with the same instance name. However, we recommend that you do not use the same instance name on more than one library server.

You can define one or more instances on the same OnDemand server. Each instance can be configured with different properties. For example:
- You can define instances that use different databases and database file systems
- You can define instances that use different cache file systems
- You can define instances that run different versions of OnDemand software
- You can define instances that maintain data in different ways, on different devices, and with different maintenance attributes
- You can define one instance that uses DB2 as the database manager and another instance that uses SQL Server as the database manager

You typically define multiple instances on one server when you want to test changes to your system, such as a new release of OnDemand, DB2, or ADSM software or a new user-defined program or when you want to generate performance data or debug problems.

Chapter 12. Creating and initializing the database

Creating the database

To create the database, select the OnDemand program group and open
OnDemand Command Line. Then enter the following command:

```
arsdb -cv
```

Initializing the system log

After successfully creating the database, initialize the system logging facility.
Enter the following command:

```
arssyscr -l
```

Initializing system migration

Note: You must initialize the system migration facility if you plan to migrate
application group index data from the database to archive media.

To initialize the system migration facility enter the following command:

```
arssyscr -m
```

Part 4. Post installation

This section of the book describes how to verify the installation of the OnDemand server.

Other important tasks described in this section include:

- Defining storage sets

  You must define storage sets before you can define you own application groups or assign the system defined application groups to a storage set. Depending on the storage management characteristics of the application groups on your system, you may need to define more than one storage set.

- Configuring the system log application group

  We recommend that you define a storage set that writes files to archive media and then update the storage management settings in the System Log application group so that OnDemand copies data written to the system logging facility to archive media. We also recommend that you store system log index data in table spaces.

- Configuring the system migration application group

  If you need to migrate index data from the database to archive media, you must define a storage set that writes files to archive media and then update the storage management settings in the System Migration application group to identify the storage set where OnDemand migrates the data. We also recommend that you store system migration index data in table spaces.

- Create a backup copy of the database

  After installing and configuring OnDemand, we recommend that you create a backup copy of the database.

Chapter 13. Verifying the installation

1. Complete all of the required steps in Part 2. Installing software beginning on page 21.

2. Complete all of the required steps in Part 3. Configuring the system beginning on page 31.

3. Shutdown and restart Windows NT. Reinitializing Windows NT starts the OnDemand services.

4. If you have not already done so, share the OnDemand client program folder so that other users on the network can install client software. We recommend that you use the share name `odclient` to share the `\Program Files\IBM\OnDemand for WinNT\client` folder.

5. If you have not already done so, install at least one of the OnDemand client programs on a PC.

6. Start the OnDemand client program. OnDemand displays the Logon to a Server dialog box.

7. Click Update Servers. OnDemand displays the Update Servers dialog box.

8. Add the name of the OnDemand library server. Click Help for information about the fields and options.

9. Click Close to return to the Logon to a Server dialog box.

10. Select the name of the server that you added in the Update Servers dialog box, if it is not already selected.

11. Type an OnDemand userid and password in the fields provided.

12. Press Enter.

13. Open and search the System Log folder.

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

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

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

Chapter 14. Defining storage sets

You must define storage sets before you define reports to OnDemand or load data into application groups. If you plan to migrate index data to archive media, you must assign the System Migration application group to a storage set that writes data to archive media. Also, we recommend that you assign the System Log application group to a storage set that writes data to archive media.

You can define storage sets that write data to cache storage or archive media (or typically, both). The storage management requirements of the application groups that you add to the system determine the types of media you need and how you configure storage sets.

A storage set must contain at least one primary storage node. A primary storage node can write data to cache storage or, if you installed ADSM, archive media. If the primary storage node writes data to archive media and you need to maintain a backup copy of the data, you can define a secondary storage node or implement a copy storage pool in ADSM. The administrator interface online help provides details about defining storage sets and storage nodes. "Implementing a copy storage pool" on page 45 describes how to implement a copy storage pool in ADSM.

You must assign each system application group to a storage set to ensure proper operation of the system. You should assign these application groups to storage sets that write data to archive media and maintain the data for the life of your OnDemand system. Refer to "Chapter 15. Configuring the System Log application group" on page 69 and "Chapter 16. Configuring the System Migration application group" on page 73 for details.

When you define an application group, you must assign it to a storage set. Use the administrator interface to define an application group and assign it to a storage set.

Chapter 15. Configuring the System Log application group

Before you begin defining applications to OnDemand, loading data into application groups, or allowing end-users to access the system, we recommend that you configure the System Log application group to permanently store data written to the system logging facility on archive media. You can do this by first defining a storage set that writes data to archive media and then updating the System Log application group and assigning it to the storage set. If your system does not support archive media, you should assign the System Log application group to a cache-only storage set and change the length of time that OnDemand maintains system log data in cache storage to the maximum permitted value. In addition, if you define table space file systems to OnDemand, we recommend that you store system log data in table spaces.

When you install and configure OnDemand, you initialize the system logging facility. The system logging facility contains the System Log application group, a set of system log applications, and the System Log folder. The System Log application group contains the storage management information that OnDemand uses to maintain the data written to the system logging facility. When you initialize the system, the application group is not assigned to a storage set and data is maintained in cache storage for seven years. Because the application group is not assigned to a storage set, a permanent copy of the system log data is not maintained.

Maintaining system log data on archive media

We recommend that you create a storage set that writes system log data to archive media. You must add at least one primary storage node to the storage set. The primary storage node must identify a client node in ADSM that maintains data indefinitely. You may want to implement a copy storage pool in ADSM, so that a backup copy of the data is automatically maintained. The administrator interface online help provides information about defining storage sets.

When you define the primary storage node, you must provide the name of an ADSM client node. The client node must be registered in an ADSM policy domain that maintains data on archive media. The policy domain must also maintain the data indefinitely. The primary storage node name and password must be identical to the client node and password in ADSM. When you register the client node in ADSM, you must specify the name of the ADSM 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, ADSM assigns the client node to the default domain. "Chapter 10. Configuring ADSM" on page 33 provides details about defining storage devices to ADSM and policy domains to support OnDemand and shows how to register a client node in ADSM.

After you define 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, OnDemand automatically copies data written to the system log to archive media. Complete the following steps to assign the System Log application group to a storage set:

1. Start the administrator interface.
2. Log on to the server with the ADMIN user. The ADMIN user is a system administrator user created when you initialize the system.
3. Click Application Groups.
4. Point to the System Log application group and click the right mouse button.
5. From the pop-up menu, select Update. OnDemand displays 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 write data to archive media and maintain the data indefinitely.
8. Replace the contents of the Length of Time before copying Cache to Archive Media field with a 0 (zero). This causes OnDemand to migrate system log records to archive media when the arsmaint command runs.
9. Click OK. OnDemand updates the database.

The administrator interface online help provides information about assigning an application group to a storage set.

Maintaining system log data in cache storage

If your OnDemand system does not support archive media, you should assign the System Log application group to a cache-only storage set and change the length of time that OnDemand maintains system log data in cache storage to the maximum permitted value. Doing so ensures that OnDemand does not expire system log data from cache storage for a very long time. Complete the following steps to assign the System Log application group to a storage set and change the length of time that OnDemand maintains system log data in cache storage:

1. Start the administrator interface.

2. Log on to the server with the ADMIN user. The ADMIN user is a system administrator user created when you initialize the system.
3. Click Application Groups.
4. Point to the System Log application group and click the right mouse button.
5. From the pop-up menu, select Update. OnDemand displays 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 Length of Time to Cache Data on Magnetic field with 99999. This value causes OnDemand to maintain system log data in cache storage for approximately 273 years.
9. Click OK. OnDemand updates the database.

The administrator interface online help provides information about the settings in the storage management category.

Storing system log data in table spaces

If you define table space file systems to OnDemand, we recommend that you configure the system log application group to store index data in table spaces.

To update the system log application group:
1. Start the administrator interface.
2. Log on to the server with the ADMIN user. The ADMIN user is a system administrator user created when you initialize the system.
3. Click Application Groups.
4. Point to the System Log application group and click the right mouse button.
5. From the pop-up menu, select Update. OnDemand displays the Update an Application Group window.
6. On the General page, select the type of table space. We recommend that you store the system log index data in SMS table spaces.
7. Click OK. OnDemand updates the database.

Chapter 16. Configuring the System Migration application group

If you plan to migrate index data from the database to archive media, you must assign the System Migration application group to a storage set that identifies a client node in an ADSM domain. Migration is the process by which OnDemand moves index data from the database to archive media. This process optimizes database storage space while allowing you to maintain index data for a very long time. You typically migrate index data after users no longer need to access the reports, but, for legal or other requirements, you still need to maintain the data for some number of years or months. OnDemand uses the storage management settings for application groups that you define to maintain migrated index data on archive media. When you define an application group, you determine if OnDemand migrates index data and how long OnDemand maintains data stored in the application group. In addition, if you define table space file systems to OnDemand, we recommend that you store system migration data in table spaces.

When you install and configure OnDemand, you initialize the system migration facility. The system migration facility contains the System Migration application group, a set of system migration applications, and the System Migration folder. The System Migration application group contains the storage management information that OnDemand uses to maintain index data migrated to archive media. Until you assign the application group to a storage set, OnDemand cannot migrate index data from magnetic storage to archive media.

Assigning the system migration application group to a storage set

If you need to store index data on archive media, you must assign the System Migration application group to a storage set that identifies a client node in ADSM. You must register the client node in an ADSM policy domain that maintains data indefinitely. You may want to implement a copy storage pool in ADSM, so that a backup copy of the migrated data is automatically maintained. When you define the storage set, the primary storage node name and password must be identical to the client node and password in ADSM. The administrator interface online help provides information about defining storage sets and primary storage nodes. "Chapter 10. Configuring ADSM" on page 33 provides details about defining storage devices to ADSM and policy domains to support OnDemand and shows how to register a client node in ADSM.

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, OnDemand automatically migrates index data to archive media, whenever the arsmaint command runs. Complete the following steps to assign the System Migration application group to a storage set:

1. Start the administrator interface.
2. Log on to the server with the ADMIN user. The ADMIN user is a system administrator user created when you initialize the system.
3. Click Application Groups.
4. Point to the System Migration application group and click the right mouse button.
5. From the pop-up menu, select Update. OnDemand displays 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 write data to archive media and maintain the data indefinitely.
8. Click OK. OnDemand updates the database.

The administrator interface online help provides information about assigning an application group to a storage set.

Storing system migration data in table spaces

If you define table space file systems to OnDemand, we recommend that you configure the system migration application group to store index data in table spaces.

To update the system migration application group:

1. Start the administrator interface.
2. Log on to the server with the ADMIN user. The ADMIN user is a system administrator user created when you initialize the system.
3. Click Application Groups.
4. Point to the System Migration application group and click the right mouse button.
5. From the pop-up menu, select Update.
6. On the General page, select the type of table space. We recommend that you store the system log index data in SMS table spaces.
7. Click OK to update the application group and return to the administrator window.

Chapter 17. Creating a back up copy of the database

To complete the installation and configuration process, we recommend that you create a full, offline backup image of the OnDemand database on removable media. 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).

Complete the following steps to create a full, offline backup of the database.

1. Make sure there are no other users logged on to the library server.
2. Make sure there are no other applications connected to the database.
3. If you plan to backup the database to tape, place a blank, formatted tape storage volume in the tape drive.
4. Select the OnDemand program group and open OnDemand Command Line.
5. Enter the backup command at the prompt. For example:

   ```
   arsdb -v -y <device>
   ```

   Replace the string <device> with the name of the output device.
6. Record 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. Save the backup copy until you create another full backup image of the database.

75

Appendix A. OnDemand national language support

Overview

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

Conversion between different code pages

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

Ideally, for optimal performance, OnDemand clients and applications should always use the same code page as the OnDemand instance[2]. However, this is not always practical or possible. OnDemand provides support for character conversion that allows clients, applications, and instances to use different code pages. This means that, while an OnDemand instance must run in a single code page, clients that access the instance can operate in any code page and reports that you store in OnDemand can contain characters encoded in any code page.

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

When does character conversion occur?

Character conversion can occur in the following situations:

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

  Character data conversion takes place on the server using Unicode code page mapping tables. Unicode code page mapping tables exist for all single- and double-byte languages. For example, a Windows client operating in the Latin 1 code page 1252 can access an OnDemand instance that has character data encoded in the Latin 1 code page 819 (code set ISO 8819-1).

---

2. An OnDemand instance is a logical OnDemand server environment, with one library server (and database) and one or more object servers. Each OnDemand instance (the servers and the database) must run in a single code page.

　　　　　　　　　　　　　　　　　　　　　77

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

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

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

  Character data conversion takes place on the server using Unicode code page mapping tables. 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 for:

- Documents stored in OnDemand.

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

- Documents retrieved from OnDemand.

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

  Note: Although OnDemand sends documents from the server to the client without converting them, the client viewing program may perform character conversion to map fonts a document was created with to fonts available at the client.

How does OnDemand determine code page values?

The client code page is determined from the operating environment when the connection to the instance is made. For example:

- In OS/2, the primary and secondary code pages are specified in the CONFIG.SYS file. You can use the chcp command to display and change code pages within a given session.
- In Windows, the code page is derived from the country ID, as specified in the International section of the WIN.INI file (16-bit client) or the locale, as specified in the Regional Settings under Control Panel (32-bit client).

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). For all other types of documents, the code page is derived from the data. The server never performs character conversion on documents.

Configuring your system for NLS

When you install OnDemand, the default locale is `English` and the default code page is `1252`. If you need to change the defaults, you must do so **before** you create the instance. After you create the instance, you cannot change the values unless you delete the instance, drop the database, and then recreate the instance and database.

The following topics provide additional information about configuring your system for NLS:

- Language settings
- Creating the database
- Mapping AFP fonts
- Creating applications

Language settings

OnDemand obtains information about the instance from values you supplied when you configured the system. For example, when you create the database,

OnDemand determines the locale and code page of the database by using the values you set when you configured the system.

"Chapter 11. Configuring OnDemand servers" on page 51 provides details about setting the locale and code page.

Creating the OnDemand database

We recommend that you use the arsdb program to create the database. The arsdb program is copied to the OnDemand program directory when you install OnDemand. Before you create the database you must verify, and if necessary, modify the database environment variables used by the arsdb program. For NLS, the key environment variables are the language and code page. "Chapter 11. Configuring OnDemand servers" on page 51 provides details about setting the language and code page.

Note: You must set the correct language and code page **before** you create the database. If you create the database using the incorrect language and code page, you must drop and then reinitialize the database. If you drop the database, you must reload any data that you previously loaded into the system and recreate any OnDemand objects you defined, such as users, folders, and application groups.

Mapping AFP fonts

You may need to map the AFP fonts your documents were created with to fonts that can be displayed at the client. Mapping AFP fonts is done with font definition files supplied with OnDemand. The files are stored in the FONT subdirectory under the directory in which OnDemand was installed (by default, \ARS, \ARSOS2, or \ARS32, depending on the client).

Note: OnDemand provides a set of font definition files for documents that contain characters in each of the supported client languages. Font definition files are provided for the following DBCS languages:
- Simplified Chinese
- Traditional Chinese
- Japanese
- Korean

When a client is running in a DBCS environment, OnDemand automatically uses the DBCS font definition files to map the fonts used in a document to the AFP fonts supplied with OnDemand.

You need to map AFP fonts if your documents contain fonts that are not defined to OnDemand, if your documents contain modified IBM AFP fonts or

if your documents contain AFP fonts that you created. Otherwise, the viewing program may not be able to display the characters properly.

The *Installation Guide for OnDemand Clients* provides details about mapping AFP fonts.

Creating applications

You typically create an OnDemand application for each type of report or source of data that you plan to store in OnDemand. The application identifies the format of the data as it is stored in OnDemand (for example, AFP). The format also determines the viewer that the client calls to display documents stored in the application. The application also identifies the program used to index the data. If you use ACIF to index the data, the application includes the parameters ACIF uses to index the data.

The Data Type of the application identifies the format of the data as it is stored in OnDemand and the viewer that the client calls to display documents stored in the application. If you plan to store line data in OnDemand or create indexing parameters with the graphical indexer (using a line data source file), you must set the Data Type of the application to Line. When you set the Data Type to Line:

- Verify the code page of the data. The code page of the data is typically the code page of the operating system where the data was created. In OnDemand, the default code page for line data is 500 (ISO EBCDIC).
- If the line data contains shift-in and shift-out codes, indicate how OnDemand client programs handle 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.

For all other types of data, the code page is encapsulated in the data. For AFP data, it is possible that characters are encoded in more than one code page. The AFP viewer uses mapping files to display single- and double-byte data in the proper code page. You may need to map fonts a document was created with to AFP fonts supplied with OnDemand to properly display some characters. Refer to the *Installation Guide for OnDemand Clients* for details about how to map AFP fonts.

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 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. Refer to the *Indexing Reference* for more

information about ACIF, including examples that show how to code trigger and index string values for EBCDIC data.

Appendix B. Configuring the system log user exit

You can configure OnDemand to record information, warning, and error messages in the system logging facility. OnDemand can record messages about system activity, such as when users log on and log off the system, and application group activity, such as when clients query and retrieve data. In addition, you can configure OnDemand to send the messages to the system log user exit program. You can use the system log user exit program to generate alerts, compile statistics, and generate accounting information.

Configuring the system log user exit program requires the following:
- Enable OnDemand to record system messages and determine the types of system messages recorded in the system log by selecting the appropriate options in the System Parameters dialog box.
- Enable OnDemand to record application group messages by selecting the appropriate option in the System Parameters dialog box.
- Determine the types of application group messages OnDemand records in the system log by setting Message Logging options when you create or update an application group.
- Enable OnDemand to send messages to the system log user exit by selecting the appropriate options in the System Parameters dialog box.

After you've completed these steps, OnDemand automatically records messages in the system log and sends the messages to the system log user exit.

OnDemand sends a message to the system log user exit by calling the arslog.bat file. By default, the arslog.bat file does not perform any function. However, you can replace the file we supply with one that does specific functions, such as check the message number and issue an alert.

A message sent to the arslog.bat file contains the fields listed in Table 3.

Table 3. System Log User Exit Program Variables

Variable	Purpose	Example
$1	OnDemand instance	archive
$2	Time stamp	08/13/95 14:24:31
$3	Log record identifier	57049
$4	OnDemand userid	ADMIN
$5	User's accounting information	Z76-001J/999999

Table 3. System Log User Exit Program Variables (continued)

Variable	Purpose	Example
$6	Severity: **1** Alert **2** Error **3** Warning **4** Information **5** Debug	3
$7	Message number	31
$8	Message text	Failed Login: odaixlib1 7.52.365.12

If you create your own arslog.bat file, remember that:

- The file and any programs called by the file run with Administrator authority. An Administrator is a privileged account with unrestricted access to all files and commands.
- You should specify the full path name of all files and programs invoked from the file. For example, specify `c:\winnt\system32\backup.exe`, not `backup`.

OnDemand programs are coded in the C language. However, the arslog.bat file can call any executable file. It is your responsibility to develop the program and validate the quality and performance of the arslog.bat file and any programs it calls. The actual mechanism for taking action based on the messages provided by OnDemand is dependent on the software that you are using on the system.

The administrator interface online help provides information about how to enable OnDemand to write messages to the system logging facility and send messages to the system log user exit program. The administrator interface online help also provides information about how to choose the application group messages that OnDemand writes to the system log and sends to the system log user exit program.

The System Log table contains one row for each record written to the system logging facility. Table 4 on page 85 lists the fields, data types, and descriptions of the fields that make up a system log record.

Table 4. System Log Table

Column Name	Data Type	Size	Description
time_stamp	Date/Time(TZ)	4	The time stamp of the log record in the OnDemand internal date format. Refer to the arsdate command for details.
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.
appl_id	CHAR	1	Determines whether OnDemand overhead information is valid. **A** Not applicable. The overhead information does not apply to the log record. However, the overhead information may 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 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.

Table 4. System Log Table (continued)

Column Name	Data Type	Size	Description
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 compressed offset of the document within the compressed object.
comp_len	INTEGER	4	The compressed 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.

Appendix C. About database servers

When you configure the library server, you determine the number of processes that 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 OnDemand end-user client programs, the value that you set must support:

- Active database manager tools and utilities.
- The number of active arsdb processes.
- The number of active arsload processes (two for each).

Description

Each connection to the OnDemand database requires a database agent. 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 to optimize the way that OnDemand handles the database load. For example, you can configure the server so that OnDemand starts a fixed number of database agents, regardless of the number of concurrent database requests. While this may appear restrictive, database requests typically process very quickly. For example, ten database agents can handle a heavy database request load, while balancing the impact on system resources.

The Number of Database Servers parameter supports high performance database access for OnDemand servers with the following characteristics:

- Many concurrent database requests. For example, a system with 20 concurrent users.
- One or more load processes active.
- One or more database maintenance processes active.
- Database administration using database manager tools and utilities.

We recommend that you 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

87

depends on the characteristics of the queries. For example, general queries typically keep a connection open longer than a more specific query.

Notes

1. The default setting of 4 (four) is recommended for servers that support a small to medium number of concurrent connections to the database. OnDemand starts four threads for all client connections and four threads to handle all database connections. Use caution when you change the default to a larger value.

2. For DB2, the size of the DB2 MAXAPPLS database parameter must exceed the size specified in the Number of Database Servers parameter:

   The MAXAPPLS parameter determines the maximum number of simultaneous database connections. The default value is 20 (twenty). In general, we recommend that you set this parameter to three times the value specified in the Number of Database Servers parameter.

   If you set the Number of Database Servers parameter to zero and you expect the server to handle more than twenty concurrent database connections, you must increase the value of the MAXAPPLS parameter.

   Refer to the DB2 documentation for details about modifying the MAXAPPLS parameter.

3. For SQL Server, the User Connections setting determines how many simultaneous connections are available for clients. The default value is 20 (twenty). If you set the Number of Database Servers parameter to zero and you expect the server to handle more than twenty concurrent database connections, you must increase the value of the User Connections setting. Refer to the SQL Server documentation for details about modifying the User Connections setting.

4. If the database manager is SQL Server, you need a client access license for each concurrent connection to the database. If the database manager is DB2, your OnDemand license allows an unlimited number of connections to the database.

Appendix D. About table space file systems

Overview

When you create the OnDemand database, the database manager allocates space for the database using configuration information that you provide. For example, the Database Location determines the base file system where the database manager stores system tables and allocates the USERSPACE1 table space.

Before you begin loading reports in OnDemand, you must decide whether to store index data in the base file system or table space file systems. If you decide to store index data in table space file systems, you must define the file systems to OnDemand.

Regardless of where you decide to store index data, you must make sure that sufficient free space is available in the file systems to hold the index data created during load processing. If you need to allocate additional free space, you can increase the size of a file system or add table space file systems.

- You must use the Control Center tool to increase the size of a database file system. The Control Center tool is a DB2 UDB graphical user interface to database objects. Detailed information is provided through the online help facility within the Control Center.
- Adding table space file systems allows OnDemand to improve database performance and provides flexible backup and recovery options. Table spaces allow OnDemand to locate tables of application group index data on different file systems. You can choose to backup table spaces, rather than the entire database. If a media failure occurs and you need to recover the database, you may be able to restore the affected table space, rather than the entire database.

Before you increase the size of a file system or add table space file systems, you must prepare the physical storage volumes for use.

About table spaces

There are two types of table spaces:
- System Managed Space (SMS) Table Space, where the operating system controls the storage space.
- Database Managed Space (DMS) Table Space, where the database manager controls the storage space.

89

In general, a DMS table space provides better performance than a SMS table space. However, OnDemand allocates all of the disk space required for a DMS table space at the time that you define the application group. The size of a DMS table space is a factor of the maximum number of rows permitted for the table and the size of a row. A database row is calculated using the lengths of the application group fields and adding overhead required by OnDemand. Space is allocated across all defined DMS table space file systems.

Note: We do not recommend using DMS table space file systems because all of the space for a DMS table space must be allocated when the table space is created. We estimate the amount of space required for a DMS table space, using the maximum number of rows and the lengths of the fields defined for the application group. Because we estimate the amount of space needed, we may allocate space that is never used. It is also possible that we may not allocate enough space for a table space, which could result in data loading errors.

Space is allocated for a SMS table space when it is required. The operating system increases the size of a SMS table space, as required, during a load process. The operating system allocates space in the SMS table space file system with the most free space.

When you create an application group, you choose the type of table space used to store application group data (SMS, DMS, or None):
- We recommend that you use SMS table spaces.
- If your organization defines few application groups, you can use DMS table spaces.
- If you choose None, application group data is stored in the USERSPACE1 table space.

When defining table space file systems, consider the following:
- We recommend that table space file systems contain only OnDemand application group data.
- Do not define SMS and DMS table spaces in the same file system.
- We recommend that you allocate equal amounts of disk space to each SMS file system.

For more information about table spaces, please refer to the *DB2 Administration Guide.*

Index

Readers' Comments — We'd Like to Hear from You

EDMSuite OnDemand for Windows NT
Installation and Configuration Guide
Version 2.2

Publication No. G544-5526-03

Overall, how satisfied are you with the information in this book?

	Very Satisfied	Satisfied	Neutral	Dissatisfied	Very Dissatisfied
Overall satisfaction	☐	☐	☐	☐	☐

How satisfied are you that the information in this book is:

	Very Satisfied	Satisfied	Neutral	Dissatisfied	Very Dissatisfied
Accurate	☐	☐	☐	☐	☐
Complete	☐	☐	☐	☐	☐
Easy to find	☐	☐	☐	☐	☐
Easy to understand	☐	☐	☐	☐	☐
Well organized	☐	☐	☐	☐	☐
Applicable to your tasks	☐	☐	☐	☐	☐

Please tell us how we can improve this book:

Thank you for your responses. May we contact you?   ☐ Yes   ☐ No

When you send comments to IBM, you grant IBM a nonexclusive right to use or distribute your comments in any way it believes appropriate without incurring any obligation to you.

Name

Address

Company or Organization

Phone No.

Readers' Comments — We'd Like to Hear from You

G544-5526-03

IBM ®

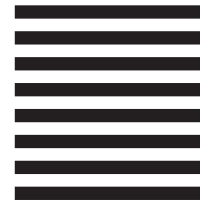Fold and Tape   **Please do not staple**   Fold and Tape

NO POSTAGE
NECESSARY
IF MAILED IN THE
UNITED STATES

BUSINESS REPLY MAIL

FIRST-CLASS MAIL   PERMIT NO. 40   ARMONK, NEW YORK

POSTAGE WILL BE PAID BY ADDRESSEE

IBM Corporation
Information Development
IBM Printing Systems Company
Department H7FE Building 003G
PO Box 1900
Boulder CO 80301-9817

Fold and Tape   **Please do not staple**   Fold and Tape

G544-5526-03

IBM ®