

EDMSuite OnDemand



Introduction and Planning Guide

Version 2.2

EDMSuite OnDemand



Introduction and Planning Guide

Version 2.2

Fifth Edition (June 1999)

This edition of *IBM EDMSuite OnDemand: Introduction and Planning 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 Software Solutions Representative.

IBM Software Solutions welcomes your comments. For your convenience, a form for reader's comments is provided at the back of this publication. You may send your comments by fax to 1-303-924-7522, by e-mail to Ondemand@us.ibm.com, or mail your comments to:

INFORMATION DEVELOPMENT
IBM SOFTWARE SOLUTIONS
DEPARTMENT MYGA BUILDING 001L
PO BOX 1900
BOULDER CO 80301-9191

Note: Visit our home page at <http://www.software.ibm.com/data/edmsuite>.

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.

© **Copyright International Business Machines Corporation 1996, 1999. All rights reserved.**
US Government Users Restricted Rights - Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

Figures	vii	OnDemand for the ADSM administrator	20
Tables	ix	Overview of OnDemand	20
Notices	xi	Operational considerations	21
Trademarks and Service Marks	xi	Chapter 3. Hardware and software requirements.	23
About this publication	xv	AIX server requirements	23
Who should use this publication	xv	HP-UX server requirements	24
How this publication is organized	xv	Solaris server requirements	24
Our use of typefaces	xvi	Windows NT server requirements	25
Product support	xvi	Optical and tape storage	26
OnDemand documentation	xvi	Download	27
Related documentation.	xvii	Enhanced indexing	27
Related IBM EDMSuite Products	xix	Server printing	27
ImagePlus VisualInfo	xix	OS/2 client.	27
ContentConnect	xix	Windows client	28
MQSeries Workflow (formerly FlowMark)	xx	Administrative clients	28
Domino Doc	xx	OS/2 client setup program	29
Windows client setup program		Administrator interface setup program	29
Part 1. Introducing OnDemand	1	Chapter 4. The role of the OnDemand system administrator	31
Chapter 1. About OnDemand	3	Part 2. Planning for OnDemand	33
Overview	3	Chapter 5. Preparing for OnDemand	35
The OnDemand system	4	Chapter 6. Planning for reports and documents	37
OnDemand concepts	6	Understanding end-user requirements.	37
Application	6	Input data formats supported by	
Application Group	6	OnDemand	37
Folder	6	Indexing data	39
OnDemand indexing methods	7	Indexing data with ACIF	40
OnDemand documents.	8	Document indexing	40
The OnDemand server environment	9	Report indexing	41
The OnDemand request manager	11	Indexing data with other programs.	42
The OnDemand database manager	12	Indexing reports using date fields	44
The OnDemand storage manager	12	Running ACIF on MVS and OS/390 Systems	44
Download	12	Running ACIF on OnDemand servers	45
Data indexing and loading	13	Chapter 7. Planning for applications	47
OnDemand management programs	15		
Chapter 2. OnDemand and ADSM	17		
ADSM for the OnDemand administrator	17		
Overview of ADSM	17		
Defining the storage configuration	19		

Overview	47	Setting up storage devices for Windows NT servers	91
The OnDemand folder	48	Recommended Storage devices for small to medium systems	91
The OnDemand application group	49	Recommended Storage devices for large systems	92
The OnDemand application	52	Using RAID storage with OnDemand	93
OnDemand users and groups	53	Overview	93
Understanding OnDemand permissions	53	RAID storage and the document cache	94
Overview	53	RAID storage and the OnDemand database	94
Setting folder permissions.	54	Configuring RAID storage	95
Setting application group permissions	55	The IBM 7133 Serial Disk System	95
About OnDemand names	56	The IBM 3527 SSA Storage Subsystem for PC Servers	95
OnDemand field types	58		
Chapter 8. Estimating storage requirements.	61	Chapter 10. Planning processor requirements and configurations.	97
Overview	61	AIX processor and memory requirements	97
Storage hierarchy	63	RS/6000 processor	97
Data compression	64	RS/6000 memory	97
Calculating magnetic storage requirements	65	HP-UX processor and memory requirements	97
Base system software	65	HP-UX processor	97
Download storage	65	HP-UX memory	98
Temporary space for indexing and loading data	66	Sun Solaris processor and memory requirements	98
Cache storage for reports	67	Solaris processor	98
OnDemand database storage.	67	Solaris memory	98
Database transaction log file storage	71	NT Processor and memory requirements	98
ADSM database and recovery log storage	74	NT processor	99
Temporary space for server print	74	System memory	99
Temporary space for importing index data	75	Server hardware configurations	99
Calculating archive media requirements	75	Small server configuration	100
Report storage space	75	Medium server configuration.	100
Database backup image file storage	76	Large server configuration	100
Database archived log file storage	78	Server software configurations	101
Migrated index storage space	79	Library/object server	101
Number of storage volumes and libraries	80	Library/object server with ADSM	102
Storage sizing examples	81	Library server and object server.	104
Storage for report containing logical items.	81	Library server and object server with ADSM	106
Storage for report containing transaction data	83	Library server with ADSM and object server with ADSM	108
Storage sizing worksheets.	84	UNIX license information	110
Storage sizing worksheet for reports containing logical items	85	NT License information	110
Storage sizing worksheet for reports containing transaction data	86		
Chapter 9. Planning a storage subsystem	89	Chapter 11. Assessing network and printer requirements	111
Overview	89	Network requirements	111
Setting up storage devices for a UNIX server	90		

Data download	111	Server information	124
OnDemand client access to OnDemand servers	112	OnDemand database	124
Print requirements	113	Database table spaces	124
Printing jobs to PSF/MVS.	114	Database backup.	125
Chapter 12. Planning for client programs	115	Using ADSM to manage backup images	126
OS/2 client.	115	Database logging.	126
Windows client	115	Using ADSM to manage DB2 archived log files	127
Administrative client	115	Database recovery	127
Installing and running client software on a network.	115	The ADSM database	128
Client start up parameters.	116	ADSM recovery log	129
Adobe software provided with OnDemand	116	ADSM volume history file	130
Chapter 13. Planning for programming interfaces	119	ADSM device configuration file	130
Client customization	119	ADSM database recovery	130
Server commands	119	Archived reports	131
Server logging	120	The document cache	132
System logging facility	120	Reports on archive media	132
System log user exit.	120	<hr/>	
Logon user exit	121	Part 3. Appendixes	135
Chapter 14. Planning for backup and recovery	123	Glossary	137
Overview	123	Index	159
Server software	123	Readers' Comments — We'd Like to Hear from You	165

Figures

1. OnDemand System Environment	5	14. Calculating Log File Storage	73
2. Folders, Application Groups, and Applications	7	15. Calculating ADSM Database Storage	74
3. Applications and Documents	9	16. Calculating Archive Media Storage Space	75
4. Data Indexing and Loading Overview	14	17. Calculating Database Archived Log File Storage	79
5. ADSM and OnDemand Environment	18	18. Calculating Migrated Index Storage	80
6. Document Indexing Method	41	19. Library/Object Server	101
7. Report Indexing Method	42	20. Library/Object Server with ADSM	102
8. Calculating Report Download Space	66	21. Library Server and Object Server	104
9. Calculating Temporary Work Space	66	22. Library Server and Object Server with ADSM	106
10. Calculating Cache Storage	67	23. Library Server with ADSM and Object Server with ADSM	108
11. Calculating database storage space	69		
12. Calculating Log File Storage	71		
13. Calculating Active Log File Storage	72		

Tables

1. Optical Libraries	26	18. Blank archive media storage requirements worksheet for reports containing logical items	86
2. Group Permissions	53	19. Blank report profile worksheet for reports containing transaction data	86
3. Database and Folder Field Types	58	20. Blank magnetic disk requirements worksheet for reports containing transaction data	87
4. Index Field Type and Index Size	68	21. Blank archive media storage requirements worksheet for reports containing transaction data.	87
5. Database calculation example for a report that contains logical items	69	22. The ars Volume Group	90
6. Database calculation example for a report that contains line data	70	23. Logical Volumes and File Systems	91
7. Database Columns	81	24. Storage Devices for Windows NT Servers	92
8. Example report profile for report containing logical items	81	25. Storage Devices for Windows NT Servers	92
9. Example management disk requirements for report containing logical items	82	26. RAID Implementations	93
10. Example optical storage requirements for report containing logical items	82	27. Small Server Configuration.	100
11. Database Columns	83	28. Medium Server Configuration	100
12. Example report profile for storage for report containing logical items	83	29. Large Server Configuration.	100
13. Example magnetic disk requirements for storage for report containing logical items	83	30. Software for Library/Object Server	101
14. Example archive media storage requirements for storage for report containing logical items	84	31. Software for Library/Object Server with ADSM	103
15. Database Columns Worksheet	84	32. Software for Library Server	105
16. Blank report profile storage sizing worksheet for reports containing logical items	85	33. Software for Object Server	105
17. Blank magnetic disk requirements worksheet for reports containing logical items	85	34. Software for Library Server	107
		35. Software for Object Server	107
		36. Software for Library Server with ADSM	109
		37. Software for Object Server with ADSM	109

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: SWS General Legal Counsel, IBM Corporation, Department TL3 Building 062, P.O. Box 12195, Research Triangle Park, NC 27709-2195 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 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
AIX
AIX/6000
AIXwindows
BookManager
BookMaster
CICS
CICS/ESA
DATABASE 2
DB2
DB2 Universal Database
eNetwork
IBM
Infoprint

IPDS
MVS
MVS/XA
MVS/ESA
OS/2
OS/390
Print Services Facility
RS/6000
S/370
S/390
SP2
SP
System/370
System/390
Ultrastar
VisualInfo
WIN-OS2

The following terms, used in this publication and other books in the 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

Cygnnet is a registered trademark of Cygnnet 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

Intel, MMX, and Pentium are trademarks or registered trademarks of Intel Corporation

Lotus is a trademark of Lotus Development Corporation

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

Solaris is a trademark of Sun Microsystems, Inc.

Tivoli is a trademark of Tivoli Systems, Inc.

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 contains information about EDMSuite OnDemand Version 2.2 (OnDemand). OnDemand requires a database management product, such as IBM's DB2 Universal Database (DB2). If you plan to archive data on optical or tape storage volumes, OnDemand requires the IBM ADSTAR Distributed Storage Manager (ADSM) product.

Note: PSF for AIX, PSF/MVS, PSF for OS/390, and PSF/2 are commonly referred to as PSF throughout this book. MVS Download and Download for OS/390 are commonly referred to as Download throughout this book.

Who should use this publication

This book is of primary interest to people planning for the installation of OnDemand and other people in an organization responsible for the hardware, software, network, physical location, business recovery, and application planning related to production information systems.

How this publication is organized

This publication is organized into two basic parts:

- “Part 1. Introducing OnDemand” on page 1 provides an overview of the OnDemand system and includes information that describes how OnDemand and its database and storage management subsystems, the operating system, and the network create an archival and retrieval environment that supports your business.
- “Part 2. Planning for OnDemand” on page 33 contains important information about how to implement an OnDemand system. This part of the publication includes details about defining reports¹ to OnDemand, indexing data, estimating storage, processor, memory, network, and printer requirements, and planning for OnDemand client programs.

The “Glossary” on page 137 contains terms and definitions that you might find helpful as you and others in your organization learn about OnDemand.

1. In OnDemand, the term *report* refers to any type of data you wish to store in the system. A 10,000-page general ledger generated by an OS/390 application and a two-page Lotus WordPro file are both reports to OnDemand.

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.
<i>Italic</i>	Used to emphasize concepts and terms.
Monospace	Indicates output of commands and programs in examples. Also used for information you are instructed to type.
UPPERCASE	Indicates parameter or command names and sometimes file and directory names.

Product support

The IBM Support Center maintains current information about OnDemand, including program temporary fixes (PTFs).

Before you install OnDemand, contact the IBM Support Center or your IBM software service representative to obtain the latest maintenance level of OnDemand.

If you encounter problems or errors running any of the OnDemand programs, you can call the IBM Support Center to obtain software problem and defect support. The phone number for the IBM Support Center is 1-800-237-5511. The OnDemand program number is 5622-622 for AIX, 5765-D60 for HP-UX, 5765-E20 for Solaris, and 5639-E12 for Windows NT. The OnDemand component ID is 5622-66200.

OnDemand documentation

The following publications contain information about OnDemand Version 2.2:

Introduction and Planning Guide, G544-5281

Installation and Configuration Guide for UNIX Servers, G544-5598

Installation and Configuration Guide for Windows NT Servers, G544-5526

Administrator's Reference, S544-5293

Indexing Reference, S544-5489

Getting Started with the Administrative Client, S544-5463

User's Guide, SC26-9810

OS/2 Client Customization Guide, S544-5465

OLE Control and Win32 Client Customization Guide and Reference, S544-5466

Related documentation

The following publications contain information about ADSM Version 3, DB2 Universal Database Version 5, and PSF:

ADSM Messages, GC35-0271

ADSM for AIX Version 3 Quick Start, GC35-0273

ADSM for AIX Version 3 Administrator's Guide, GC35-0274

ADSM for AIX Version 3 Administrator's Reference, GC35-0275

ADSM for Windows NT Version 3 Administrator's Guide, GC35-0292

ADSM for Windows NT Version 3 Administrator's Reference, GC35-0293

ADSM for Windows NT Version 3 Quick Start, GC35-0295

ADSM for HP-UX Version 3 Administrator's Reference, GC35-0321

ADSM for HP-UX Version 3 Administrator's Guide, GC35-0320

ADSM for HP-UX Version 3 Quick Start, GC35-0322

ADSM for Sun Solaris Version 3 Administrator's Guide, GC35-0328

ADSM for Sun Solaris Version 3 Administrator's Reference, GC35-0329

ADSM for Sun Solaris Version 3 Quick Start, GC35-0330

DB2 Universal Database Extended Enterprise Edition for UNIX Version 5 Quick Beginnings, S99H-8314

DB2 Universal Database Version 5 Administration Getting Started, S10J-8147

DB2 Universal Database for UNIX Version 5 Quick Beginnings, S10J-8148

DB2 Universal Database for Windows NT Version 5 Quick Beginnings, S10J-8149

DB2 Universal Database Version 5 Administration Guide, S10J-8157

DB2 Universal Database Version 5 Command Reference, S10J-8166

DB2 Universal Database Version 5 Message Reference, S10J-8168

PSF for AIX: Print Submission, S544-3878

PSF for AIX: Print Administration, S544-3817

PSF for MVS: MVS Download Guide, G544-5294

PSF for OS/390: Download for OS/390, S544-5624

Related IBM EDMSuite Products

OnDemand is one of the web-enabled products included in the IBM Enterprise Document Management Suite (EDMSuite), which is a portfolio of IBM software that includes imaging, computer output to laser disk (COLD) document management, and workflow. In addition to OnDemand, EDMSuite contains four more products:

- ImagePlus VisualInfo
- ContentConnect
- MQSeries Workflow
- Domino Doc

For more information on EDMSuite, visit the EDMSuite homepage at <http://www.software.ibm.com/data/edmsuite>.

ImagePlus VisualInfo

ImagePlus VisualInfo is a distributed and scalable client/server solution that enables the management and storage of a broad array of document types, such as document images, graphics, spreadsheets, text, audio, and video. ImagePlus consists of scalable, multiplatform offerings for organizations of all sizes. ImagePlus implementations range from simple store-and-retrieve applications to image and workflow-enabling, very high-volume, transaction-oriented business processes. Some ImagePlus capabilities allow you to:

- Develop a robust library to store documents
- Handle multiple file types in addition to images, such as word processing documents, audio and video clips, and spreadsheets
- Manage and migrate documents across magnetic, optical, and tape storage automatically

ContentConnect

ContentConnect is a client toolkit based on Java that provides single-query access to multiple repositories from Web browsers, Lotus Notes clients, and stand-alone clients. ContentConnect provides a search engine that supplies reliable and immediate access to advanced document imaging, workflow, COLD archiving, and collaborative document management components. ContentConnect presents its search results in an organized, seamless, easy-to-read manner. ContentConnect not only connects document archives

and processes to other components of EDMSuite, but also connects its users to the document management world outside EDMSuite.

MQSeries Workflow (formerly FlowMark)

MQSeries Workflow is a workflow engine designed for the client/server environment. MQSeries Workflow is dedicated to managing the flow of work, allowing companies to integrate the applications required to meet the needs of their business processes. MQSeries Workflow allows you to:

- Separate application logic from business process rules
- Tie legacy and client/server applications to business process steps
- Monitor processes to show where a specific piece of work is in the overall process
- Record live production process data for analysis by management
- Model, animate, and simulate your business processes rapidly
- Enable processes across the Internet and Intranets

Domino Doc

Domino.doc is a Web-based document and content management solution that allows organizations to capture, file, retrieve, and distribute content across the Internet using desktop applications, Web browsers, or any Lotus Notes client. Domino.doc allows you to:

- Create document profiles that enforce the capture of key document search terms, so documents can be searched and retrieved easily
- Create a common library to provide a knowledge base of documents for your enterprise
- Use version control to provide a history of previous changes
- Use check-in and check-out controls to ensure data integrity

Part 1. Introducing OnDemand

This section of the book provides an overview of the OnDemand system and contains information that can help you better understand how OnDemand works. This section describes how OnDemand manages reports and index data, includes important information about how OnDemand, the database manager, and ADSM work together to archive and retrieve documents², includes information about hardware and software requirements for OnDemand, and outlines the role of the OnDemand system administrator, including a list of tasks the administrator typically performs to manage an OnDemand system.

Topic	Page
About OnDemand	3
OnDemand and ADSM	17
OnDemand server requirements	23
The role of the OnDemand system administrator	31

2. In OnDemand, the term *document* means a part of a report, such as a statement, policy, or other logical grouping of pages.

Chapter 1. About OnDemand

Overview

OnDemand supports any organization that can benefit from hard copy or microfiche replacement and instant access to information. An OnDemand system can support small office environments and large enterprise installations with hundreds of system users. OnDemand can dramatically improve productivity and customer service in many businesses by providing fast access to information contained in archived reports.

OnDemand processes the output of application programs, extracts index fields, stores index information in a relational database, and archives reports on cache storage and in ADSM. OnDemand's storage management system supports archiving newly created and frequently accessed reports on high speed, magnetic storage volumes (cache storage) and can automatically migrate reports to other types of ADSM-managed storage volumes as files age.

OnDemand fully integrates the capabilities of Advanced Function Presentation (AFP), including management of resources, indexes, and annotations, and supports full fidelity reprint and FAX of reports to devices attached to a PC, LAN, and OnDemand server.

OnDemand provides administrators with management programs, including the OnDemand administrative client, that can help manage OnDemand servers, maintain OnDemand applications, archive reports, authorize users access to OnDemand servers, and back up critical files and databases.

OnDemand provides users the ability to view reports, print or FAX copies or pages of reports, and attach electronic notes to archived reports.

OnDemand provides several advantages over data archival and retrieval methods you may have used. For example, users can:

- Easily locate data without knowing the individual report file or cycle date.
- Retrieve the section of the report without processing the entire report.
- Print or FAX a replica of the original report.

OnDemand provides end-users with an information management tool that can increase their effectiveness when working with customers.

OnDemand does the following:

- Integrates data created by application programs into an online, electronic information archive and retrieval system.
- Provides access to millions of reports at near real-time speeds.
- Retrieves the data you need when you need it.
- Provides capabilities to view, annotate, print, and FAX reports.

These features mean that OnDemand can help users quickly retrieve the specific page of a report that they need to provide fast customer service.

The OnDemand system

An OnDemand system consists of client programs and server programs that communicate over a network running the TCP/IP communications protocol, database management programs that maintain index and server control data, and storage management programs maintain documents that provide support for various types of storage devices. Figure 1 on page 5 shows an example of the OnDemand system environment. In the example, OnDemand client programs running on PCs and terminals attached to the network communicate with the OnDemand servers. The OnDemand library server manages a database of information about the users of the system and the reports archived in OnDemand. An OnDemand object server, residing on a different server connected to the network, manages the reports that reside on magnetic and optical storage devices.




Figure 1. OnDemand System Environment

OnDemand client programs operate on personal computers running OS/2 or Windows. The client program is the user's gateway to reports archived in OnDemand. Using the client program, users construct queries and search for reports, retrieve items from OnDemand, view, print, and FAX copies or pages of reports, and attach electronic notes to report file pages.

OnDemand servers manage control information and index data, store and retrieve reports and resource group files, and process query requests from OnDemand client programs. reports can reside on magnetic, optical, and tape storage volumes. New reports can be archived in OnDemand every day. That way, OnDemand always retrieves the latest information generated by application programs.

OnDemand client programs and servers communicate over a computer network supported by TCP/IP. When a user submits a query, the client program sends a search request to the OnDemand library server. The library server returns the list of items that match the query to the end-user. When the

user selects an item for viewing, the client program retrieves a copy of the item from the object server where the item resides, opens a viewing window, and displays the item.

OnDemand concepts

The terms *application*, *application group*, and *folder* represent how OnDemand stores, manages, retrieves, views, and prints reports and index data. When defining a new report file or type of data to OnDemand, an administrator must define an application and an application group or identify an existing application group for the application. In order for end-users to retrieve reports associated with the application, an administrator must define a folder or update an existing folder that references the application group for the application.

Application

An application describes the physical characteristics of a report file to OnDemand. Typically you define an application for each program that produces output that will be archived in OnDemand. The application includes information about the format of the data, the orientation of data on the page, the paper size, the record length, and the code page of the data. You can also define the parameters that the indexing program uses to locate and extract index data and special instructions that OnDemand uses to prepare and load index data in the database and reports on storage volumes.

Application Group

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

Folder

A folder is the user's way to query and retrieve data archived in OnDemand. Folders provide users a convenient way to find related information archived in OnDemand, regardless of the source of the information or how it was prepared. A folder allows an administrator to set up a common query screen for several application groups that may use different indexing schemes, so

that a user can retrieve the data with a single query. For example, a folder called Student Information might contain transcripts, bills, and grades, which represents information stored in different application groups, defined in different applications, and created by different programs.

Figure 2 shows the relationship between folders, application groups, and applications. The picture depicts two folders. One folder references a single application group, the other folder references two application groups. Application groups can contain one or more applications.

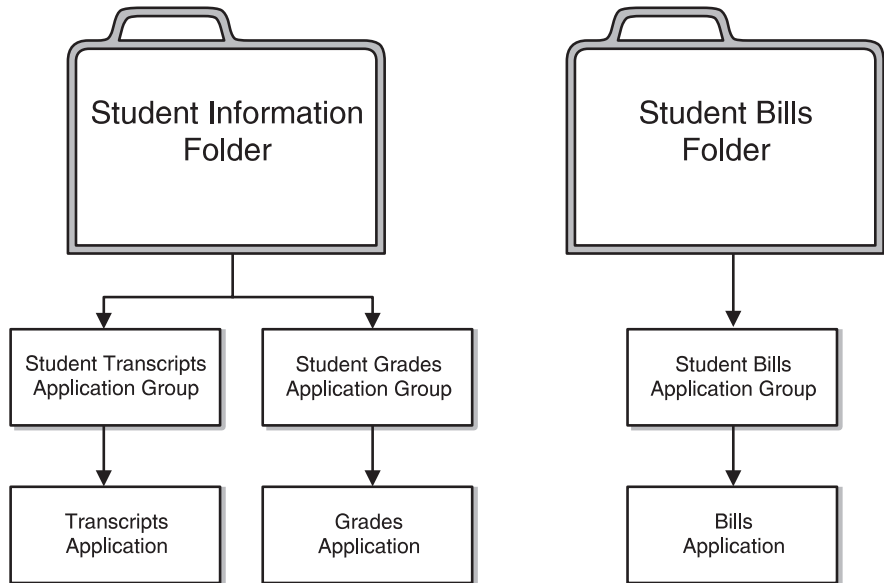


Figure 2. Folders, Application Groups, and Applications

OnDemand indexing methods

OnDemand supports two methods of indexing data:

- Document indexing is used for reports that contain logical items such as policies, and statements. Each of the items can be individually indexed on values such as account number, customer name, and balance. OnDemand supports up to 32 index values per item. With document indexing, the end-user does not necessarily have to be aware of reports or cycles to retrieve an item from OnDemand.
- Report indexing is used for reports that do not contain logical items, such as a transaction log. OnDemand stores the report as groups of pages and indexes each group. Users can retrieve report file information by page number. When reports include a sorted transaction value (for example,

invoice number), OnDemand can index the data on the transaction value. This is done by extracting the beginning and ending transaction values for each group of pages and storing the values in the database. This type of indexing lets users retrieve a specific transaction value directly.

OnDemand documents

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

Figure 3 on page 9 shows examples of OnDemand applications and documents. The picture depicts how an administrator might define OnDemand applications. Application One uses the document indexing method to segment the input report into individual customer statements. Users can retrieve statements by date and any combination of account number, customer name, balance, and transaction number. Application Two uses the report indexing method to segment the input report into groups of pages. Each group of pages is indexed using the first and last sorted transaction values that occur in the group. Users search the application by date and transaction number. OnDemand retrieves the 100 page segment that contains the value entered by the user.

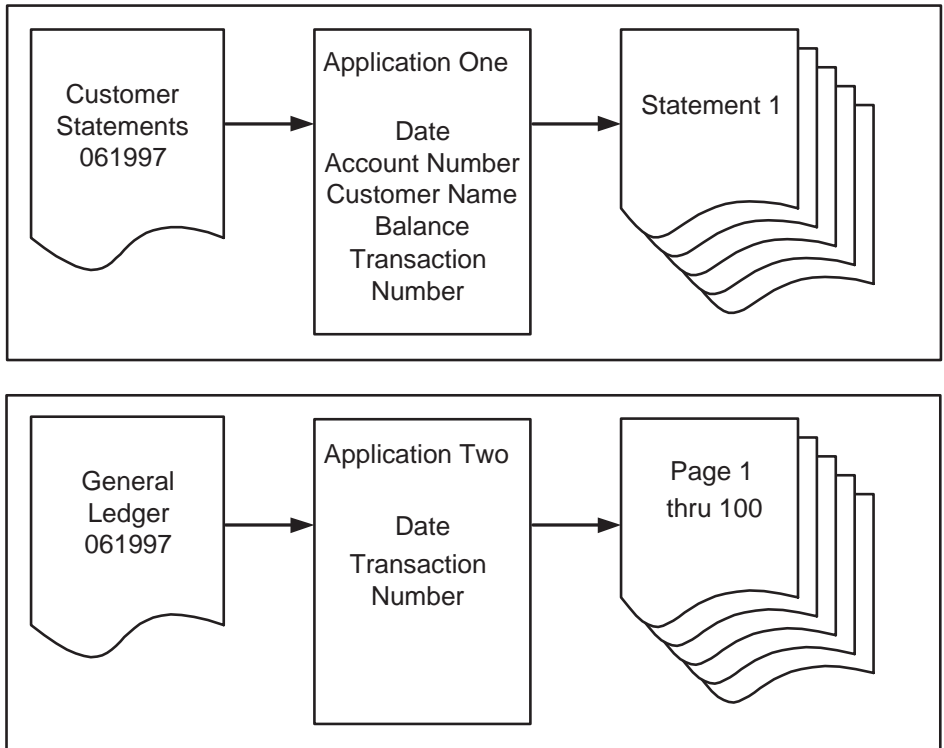


Figure 3. Applications and Documents

The OnDemand server environment

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

The library server maintains a central database about the files stored in OnDemand. The database manager provides the database engine and utilities to administer the database. The library server processes logins, queries, print requests, and database updates. The major functions that run on the library server are the request manager, the database manager, and the print manager.

An object server maintains files stored on cache storage volumes, and optionally, works with ADSM to manage files stored on archive media, such as optical and tape storage libraries. An object server processes data load,

retrieval, and expiration requests. The major functions that run on an object server are the storage manager, data loading and maintenance programs, and optionally, ADSM.

The basic OnDemand configuration is a library server and an object server on the same workstation or node. The single library/object server configuration supports the database functions and document cache on one workstation. You can add ADSM to the single library/object server configuration, to store files on archive media. You can also configure your OnDemand system with the library server on one workstation and one or more object servers on different workstations. This configuration is known as a distributed library/object server system. The distributed library/object server configuration supports caching of documents on different servers. You can add ADSM to one or more of the object servers to store files on archive media attached to different servers.

The OnDemand server environment contains several components:

- A *request manager* that provides client, network, and operating system services, security, and accounting. The request manager resides on the library server.
- A *database manager* that maintains indexes to the data archived in OnDemand. The database manager is a relational database management product, such as DB2 (which is included with EDMSuite OnDemand). The database manager resides on the library server.
- A *storage manager* that maintains cache storage and works with ADSM to manage files stored on archive media, such as optical and tape storage libraries. ADSM is an optional component and must be purchased separately. In addition, you can use ADSM to maintain archived log files and backup image files. This capability means that you do not have to manage these files on disk.
- Application and user *control data* that supports the applications, application groups, folders, storage sets, and printers that you define, determines who can access the system, and determines the folders a user can open and the application group data a user can query and retrieve. The database resides on the library server.
- A *download facility* that automatically transfers spool files to a server at high speed. The Download is part of PSF for MVS (or PSF for OS/390) and must be purchased separately.
- *Data indexing and conversion* programs. These programs create index data, collect required resources, and optionally convert line data reports to AFP files. OnDemand provides two indexing programs. The AFP Conversion and Indexing Facility (ACIF) can be used to index System/370 line data, ASCII data, and AFP files, collect resources required to view reports, and convert line data files to AFP data. The OnDemand Generic Indexer can be

used to create index data for other types of source data, such as PDF files and images. The indexing programs can run on any OnDemand server. In addition, ACIF can run on an MVS or OS/390 system.

- *Data loading* programs that can be set up to automatically store reports into application groups and update the database. The data loading programs can run on any OnDemand server.
- Archived reports and resources.
- A *server print* facility that allows users to reprint a large volume of documents at high speed. PSF is the OnDemand server print manager.
- OnDemand *management programs* to maintain the OnDemand database and cache storage volumes.
- A *system logging facility* that provides an administrator with tools to monitor server activity and respond to specific events as they occur. Your interface to the system logging facility is through a folder and a user exit.

The following topics provide additional information:

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

The OnDemand request manager

The request manager processes search requests from OnDemand client programs. When an end-user enters a query, the client program sends a request over the network to the request manager. The request manager works with the database manager to compile a list of items that match the query and returns the list to the client program. When the end-user selects an item for viewing, the request manager either retrieves the item from cache storage or sends a request to the storage manager to retrieve the item from a storage library. The storage manager retrieves data associated with the item. The OnDemand client program decompresses and displays the item.

OnDemand management programs include jobs that maintain the database and manage cache storage and automatically migrate files to optical and tape storage. These programs use the services of the request manager to manage index data, report file data, and resource files.

When a user logs on the system, OnDemand assigns a unique transaction number to that instance of the client program. All activity associated with that instance of the program contains the same transaction number. The request manager records messages generated by the OnDemand programs that service

end-user requests. For example, logon, query, print, and so forth. The messages contain the transaction number, userid, time stamp, and other information. You can open the System Log folder and review the messages. OnDemand also provides a system log user exit so that you can invoke a user-defined program to process messages, for example, send an alert to an administrator. The messages that the request manager creates can also be used to generate usage and billing reports.

The OnDemand database manager

OnDemand uses a database management product, such as IBM's DB2 UDB (provided with OnDemand), to maintain index data for the reports that you store in OnDemand. The database manager maintains the OnDemand database with tables that describe the applications, application groups, storage sets, folders, groups, users, and printers that you define and statistics to optimize the operation of OnDemand servers.

The OnDemand storage manager

The OnDemand storage manager maintains files on cache storage volumes and works with ADSM (ADSM) to manage files stored on archive media, such as optical or tape storage libraries. ADSM stores reports and resources on and retrieves files from archive media. ADSM maintains a database that contains the physical location of the files and resources. ADSM can automatically maintain backup copies of archived data. You decide which types of archive media that your OnDemand system must support, configure the storage devices on the system, and define the storage devices to ADSM.

In addition to managing reports on archive media, ADSM can also maintain DB2 archived log files and backup image files. This capability means that you do not have to maintain these files on disk. ADSM can assist you with automating database backup and archival on a regular schedule. When you use the `arsdb` command to create a database or table space backup image, you specify that you want ADSM to manage the image. After completing the backup image, the `arsdb` command copies the archived log files to ADSM-managed storage.

Download

Download is an optional feature of PSF/MVS and PSF for OS/390 that provides automatic, high-speed download of JES spool files to an OnDemand server. Download can be used to transfer reports created on MVS and OS/390 systems to the server, where OnDemand indexes the report file data and stores the report file and index data into an application group. Download operates as a JES Functional Subsystem Application (FSA) and can automatically route jobs based on a JES class or destination, reducing the need to modify JCL. Download uses TCP/IP protocols to stream data at high speed

over a LAN or channel connection from an MVS or OS/390 system to the OnDemand server. Please see *PSF for MVS: MVS Download Guide* and *PSF for OS/390: Download for OS/390* for more information about Download.

Data indexing and loading

Data that you store in OnDemand must be indexed. OnDemand supports several types of index data and indexing programs. For example, you can use the AFP Conversion and Indexing Facility (ACIF) to prepare index data for the reports that you want to archive. An administrator defines the index fields and other processing parameters ACIF uses to locate and extract index information from reports. OnDemand data loading programs read the index data generated by ACIF and load it into the OnDemand database. The data loading programs obtain other processing parameters from the OnDemand database, those used to segment, compress, and store report data in cache storage and on archive media. If you plan to index reports on an OnDemand server, you can define the parameters with the administrative client. The administrative client lets you to create ACIF indexing parameters by visually marking up sample report data. This easy to use *report wizard* runs under Windows NT, Windows 95, and Windows 98.

OnDemand also provides indexing programs that can be used to generate index data for Adobe PDF files and other types of source data, such as image files. See the *Indexing Reference* for details about the indexing programs provided with OnDemand.

Figure 4 on page 14 shows an overview of the data indexing and loading process.

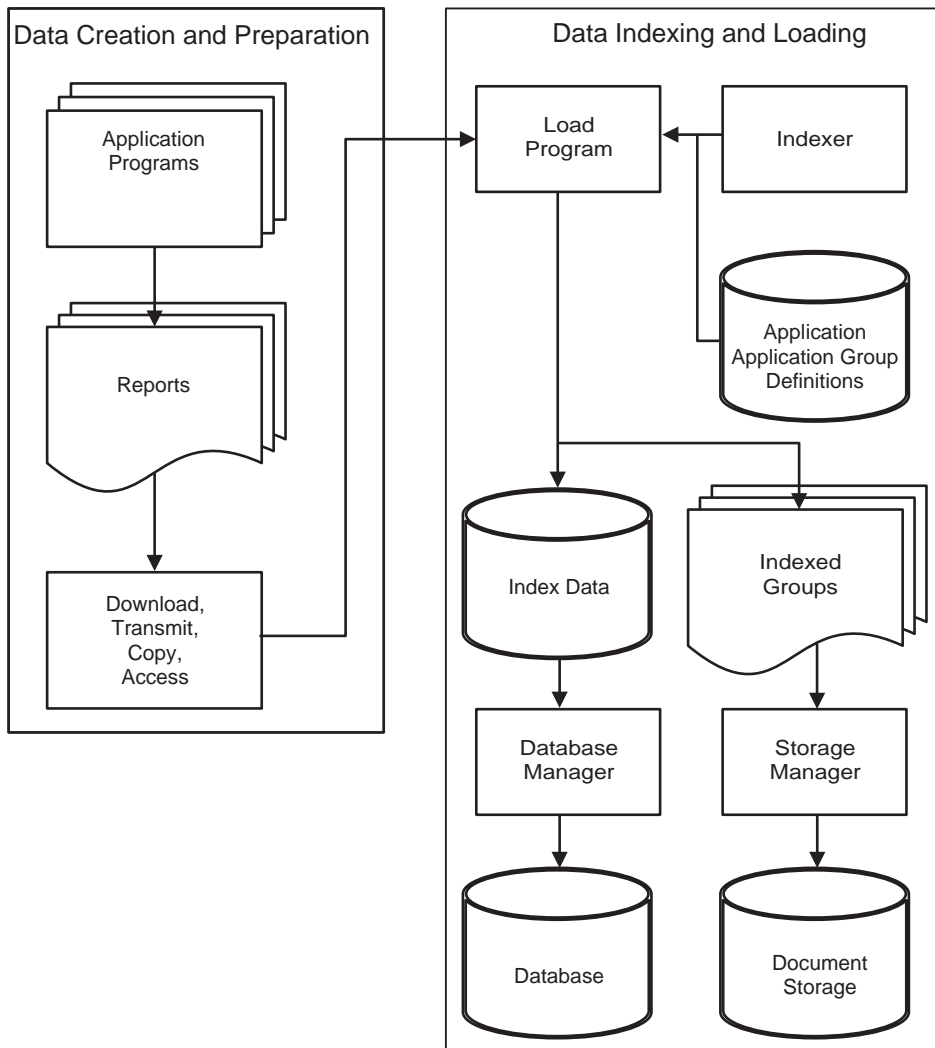


Figure 4. Data Indexing and Loading Overview

In the picture, reports created by application programs are processed by the main OnDemand data loading program. There are a number of methods you can use to provide the OnDemand data loading program access to the reports. For example, you can use MVS Download to transmit data from the JES spool to the OnDemand server. If the report needs to be indexed, the OnDemand data loading program calls the appropriate indexing program. The indexing program uses the application indexing parameters that you previously created with the administrator interface and stored in the database. The indexing

program can generate index data, collect resources, segment the report data into indexed groups. After indexing the report, the OnDemand data loading program processes the report data, resources, and index data using other application and application group parameters that you have defined. The database manager updates the OnDemand database with index data extracted from the report. Depending on the storage management definitions for the application group, the OnDemand data loading program can call the storage manager to segment and compress report data on cache storage volumes and ADSM to store copies of report data on archive media.

OnDemand management programs

OnDemand provides programs to manage databases and storage volumes and to optimize the performance of the database manager. An installer configures storage volumes and determines the processing parameters and the frequency the programs run during installation of OnDemand. When an administrator creates an application group, additional parameters are set that determine how and when OnDemand maintains index and report data. For example, an administrator defines data caching and migration values when creating an application group. OnDemand uses the information to schedule migration of report data from cache storage to archive media, deletion of report data from cache storage, migration of index data from the database to archive media, and expiration of index data from the database. OnDemand reclaims the database and cache storage space released by expired and migrated data. We recommend that you configure your OnDemand system to automatically start the management programs on a regular schedule, usually once per night or week.

ADSM independently deletes files from archive media when they reach their storage expiration date. An administrator defines ADSM management information for each ADSM domain that supports OnDemand data. The management information includes the storage volumes where ADSM stores data archived in OnDemand and data expiration information. It is mandatory that the same expiration information be specified to OnDemand and ADSM.

Chapter 2. OnDemand and ADSM

ADSM for the OnDemand administrator

Overview of ADSM

ADSM provides optical and tape storage management for the data that you archive in OnDemand. ADSM includes the following components:

- A server program that maintains a database of information about the devices and data managed by ADSM. The server program controls ADSM data storage media.
- An administrative client program that controls and monitors the ADSM server program activities and define storage management policies.
- A client API that is the OnDemand interface to ADSM. This API is part of the ADSM client.
- Device support modules which provide support for media storage libraries.

Based on the storage management policies that you define for OnDemand, ADSM stores data on and retrieves data from media volumes in ADSM storage, which is called storage pools.

Figure 5 on page 18 shows an example of ADSM supporting the OnDemand data archival and retrieval process.

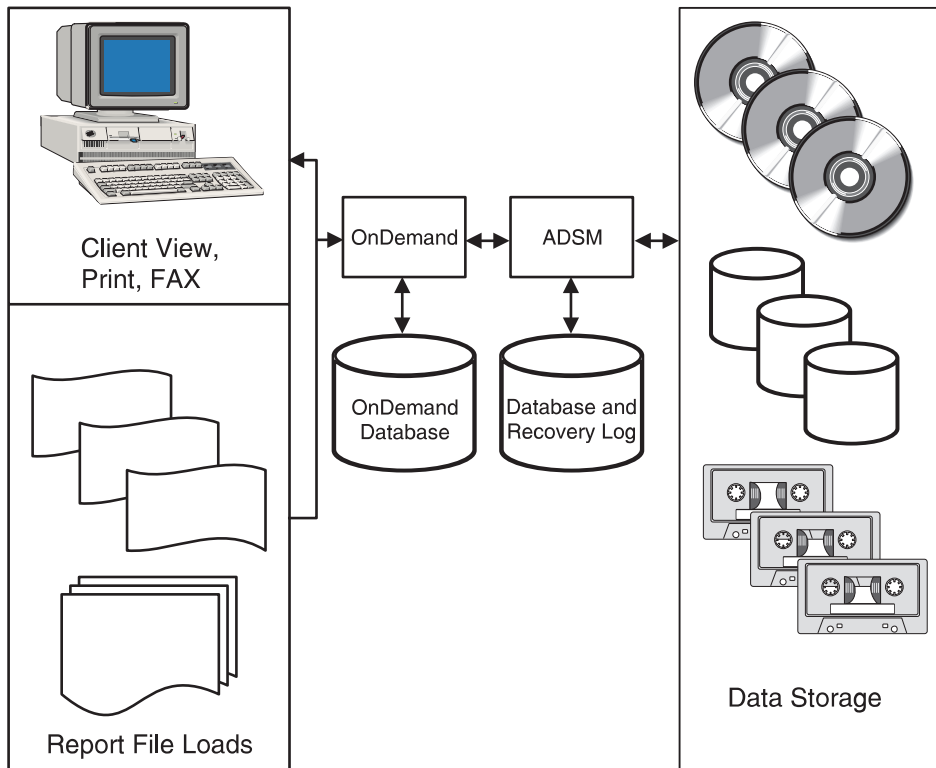


Figure 5. ADSM and OnDemand Environment

ADSM provides the administrator a graphical user interface and a command line interface to manage ADSM policies and storage. Examples of commands used to manage ADSM found in this book (and other books in the OnDemand library) use the command line interface. For UNIX servers, OnDemand provides the `ars_adsm` command to start and stop ADSM and create backup copies of the ADSM database.

The ADSM administrator defines storage management policies that carry out the requirements for archiving data stored in OnDemand. The administrator defines ADSM policy objects, which identify the library where data is stored, how long data is stored, and when data is deleted from storage volumes.

A library is a collection of drives used to read and write data on optical platters or tape cartridges. The ADSM administrator configures libraries and devices on the system and defines them to ADSM.

A device class represents a set of storage devices with similar storage characteristics. The ADSM administrator defines device classes for the drives

available to the ADSM server program. The administrator must specify a device class when defining a storage pool, which is a named collection of volumes for storing data archived in OnDemand.

OnDemand segments and compresses the report file data that you store in an application group into storage objects. ADSM places storage objects on groups of volumes called storage pools. Also, the data in these storage pools can be stored in a copy storage pool, providing a backup copy of report files stored in an application group.

For more information about ADSM and details about storage management policies and managing ADSM storage, please refer to the *ADSTAR Distributed Storage Manager: Administrator's Guide* for your server.

Defining the storage configuration

Your first task is to determine storage requirements for the report files that you plan to archive in OnDemand. You must collect physical space requirements, data retention information, media type specifications, and storage location requirements. For example:

- You determine that many of the report files generated on your system have the same basic space, media, and retention requirements. You can dedicate a high-capacity optical library to these applications. The IBM 3995 Model C68 Optical Library is an example of a high-capacity optical storage library. The 3995-C68 contains four drives that support optical platters with an unformatted capacity of up to 5.2 GB each. You can store up to 258 optical platters in the library at any one time, providing a total online (unformatted) capacity of 1.341 TB of data. Since OnDemand compresses data before storing it on storage volumes, you may obtain a much greater storage capacity. For example, if you achieve 5:1 compression on the files stored in a 3995-C68, you can store over six terabytes of data, online, in the library.
- You determine that several of the report files contain critical data and you must maintain a backup copy of the data in OnDemand. You can install, configure, and define a tape storage library where ADSM automatically stores a backup copy of the report file data.

After collecting the data storage requirements, you typically work with the ADSM administrator to configure the storage devices on the system and define the devices to ADSM. For example, when you define an optical library to ADSM, you specify the type of device, the number of drives in the library, and the capacity of a storage volume. The ADSM administrator defines a device class and a storage pool for each storage library. A storage pool contains a set of storage volumes that belong to the same library. ADSM keeps

track of the storage volumes that belong to each storage pool, including the utilization percentage and the date a storage volume was last written to or read from.

After defining the storage devices, the ADSM administrator defines the storage management policies, using the information that you collected about the report files that you plan to archive in OnDemand. Each storage node that you define to OnDemand must be registered as a node in an ADSM policy domain.

When you load data into an application group that requires optical or tape storage management, you must identify the storage set where OnDemand stores the report file data. A storage set can contain one or more storage nodes. A storage node identifies an OnDemand object server and ADSM node. ADSM stores the report file data in the optical library managed by the domain where the node is assigned.

OnDemand for the ADSM administrator

Overview of OnDemand

OnDemand manages the archival and retrieval of report files based on the storage characteristics of the data. The OnDemand administrator typically defines an OnDemand application for each report file to be archived in OnDemand. Applications with similar storage characteristics can be placed into a collection called an application group. The administrator specifies storage management characteristics for an application group, for example, the life of the data in OnDemand and the name of the storage set for the application group.

The storage set determines the type of storage support provided for an application group. One or more application groups can specify the same storage set. A storage set can contain one or more primary storage nodes. A storage node used to store files on optical or tape storage identifies an OnDemand object server and ADSM node.

OnDemand storage nodes that store files on optical or tape storage must be registered as nodes in ADSM.

The policy domain where you register the node must support the life of the data in OnDemand (the retention period in ADSM) and the type of media required by the application. A policy domain that supports OnDemand must include an active policy set that contains a management class and an archive copy group.

Operational considerations

OnDemand provides the `arsadsm` command for use with your UNIX server to start and stop ADSM and create backup copies of the ADSM database.

Each storage set that you define identifies the server and storage nodes where data written to the storage set can be stored. OnDemand can write data to one storage node at a time (OnDemand can read data from several storage nodes). If the storage set contains more than one storage node, you must identify the specific storage node where data is to be written. A system administrator can update the storage set to change the storage node where data is to be written and add storage nodes to the storage set.

OnDemand segments and compresses report file data into storage objects. ADSM places storage objects on the storage volumes it manages. OnDemand does not require ADSM to compress the storage objects. OnDemand extracts and decompresses a small portion of the storage object as required to support end-user retrieval requests.

The OnDemand administrator sets up data caching and migration values for each application group to determine if OnDemand copies data from cache to storage media and expires data on cache storage. ADSM independently expires data from the storage it manages. However, OnDemand and ADSM must specify the same expiration values so that data is removed from the system only when it is no longer required.

Chapter 3. Hardware and software requirements

The exact hardware and software configuration that you need to support your OnDemand for system depends on the volume of data you plan to manage, the number of concurrent users the system must support, back up and recovery requirements, and performance requirements. At a minimum, you need one workstation for a standard library/object server system.

AIX server requirements

The minimum configuration for an OnDemand for AIX server requires the following:

- AIX Version 4.2.1 or later
- IBM DB2 Universal Database Enterprise Edition Version 5.2 or later (included with OnDemand for AIX) for the library server
- RISC System/6000 Model F50
- 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 report 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.
- If you plan to use ADSM for backup and restore of your databases or to store reports on archive media, you must install ADSM Version 3.1.0.6 or later on at least one object server
- If you plan to use OnDemand server print or FAX, you must install PSF on a workstation that belongs to the same network as the OnDemand library server

HP-UX server requirements

The minimum configuration for an OnDemand for HP-UX server requires the following:

- HP-UX Version 10.20 or later with the following patches:
 - PHSS_15043, for OnDemand
 - PHSS_10556, for DB2 UDB V5.2
 - PHSS_10436, for DB2 UDB V5.2
 - PHSS_10113, for DB2 UDB V5.2
 - PHSS_10053, for DB2 UDB V5.2
- IBM DB2 Universal Database Enterprise Edition Version 5.2 or later (included with OnDemand for HP-UX) for the library server
- HP 9000 server
- 128 megabytes of memory
- Token Ring or Ethernet network adapter
- Magnetic and tape 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 report storage.
- A CD-ROM drive for program installation
- A 4mm or 8mm tape drive or automated tape library for data backup and recovery. A tape drive is recommended for a small system configuration. An automated tape library is recommended for medium and large system configurations. The *Introduction and Planning Guide* provides details about system configurations.
- If you plan to use ADSM for backup and restore of your databases or to store reports on archive media, you must install ADSM Version 3.1.0.6 or later on at least one object server
- If you plan to use OnDemand server print or FAX, you must install PSF on a workstation that belongs to the same network as the OnDemand library server

Solaris server requirements

The minimum configuration for an OnDemand for Sun Solaris server requires the following:

- Solaris Version 2.6 or later
- IBM DB2 Universal Database Enterprise Edition Version 5.2 or later (included with OnDemand for Sun Solaris) for the library server
- Sun Enterprise Ultra 10S server
- 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 report 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.
- If you plan to use ADSM for backup and restore of your databases or to store reports on archive media, you must install ADSM Version 3.1.0.6 or later on at least one object server
- If you plan to use OnDemand server print or FAX, you must install PSF on a workstation that belongs to the same network as the OnDemand library server

Windows NT server requirements

The minimum configuration for an OnDemand for Windows NT server requires the following:

- Microsoft Windows NT Version 4.0
- DB2 Universal Database Enterprise Edition Version 5.2 or later (included with OnDemand for NT) or Microsoft SQL Server 6.5
- An IBM compatible PC with an Intel 200 MHz Pentium Pro processor and 128 MB of RAM
- Token Ring or Ethernet connection
- A minimum of 8 GB of disk storage on at least two physical disks plus sufficient storage space for documents
- A CD-ROM drive for program installation (required)
- Standard TCP/IP support with Windows NT Server Version 4.0. To use IPX/SPX, the server must include NWLink in the protocol stack. To communicate with an OnDemand Windows NT server that uses IPX/SPX, clients must include an IPX/SPX-compatible network protocol in their protocol stack.
- 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.

- If you plan to use ADSM for backup and restore of your databases or to store reports on archive media, you must install ADSM Version 3.1.2.13 or later on at least one object server
- If you plan to use OnDemand server print or FAX, you must install PSF on a workstation that belongs to the same network as the OnDemand library server

Optical and tape storage

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

- ADSM Version 3.1.0.6 or later Server with the appropriate optical or tape device support module.
- ADSM Version 3.1.0.6 or later 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 IBM optical libraries supported by ADSM on AIX and Solaris servers. The IBM 3995 C-Series Optical Library family features the IBM industry-standard 5.25 inch, 5.2 GB (8x) Extended Multifunction Optical Drive. The 5.2 GB (8x) drives, when used with 5.2 GB optical cartridges, provide library capacities from 104 GB to 1.341 TB of online optical storage, depending on model configuration. The 3995 library supports magneto-optical (MO) rewritable, permanent write once/read many (WORM), and Continuous Composite WORM (CCW, also known as non-permanent WORM) optical media technologies.

Table 1. Optical Libraries

Device	Drives	Storage Cells	Uncompressed Capacity	SCSI IDs Required
IBM 3995-C60	1, 5.25 inch	20	104 GB	2
IBM 3995-C62	2, 5.25 inch	52	270 GB	3
IBM 3995-C64	2, 5.25 inch	104	540 GB	3
IBM 3995-C66	4 or 6, 5.25 inch	156	810 GB	5 or 7
IBM 3995-C68	4, 5.25 inch	258	1.341 TB	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. Visit the ADSM home page at <http://www.storage.ibm.com/software/adsm> for the latest information about supported devices.

Download

To enable seamless download from MVS and OS/390 systems to OnDemand servers, the following is required:

- IBM TCP/IP for MVS Version 3.2 or later
- PSF/MVS Version 2.2 or later with the MVS Download Feature or PSF Version 3.1 or later for OS/390 with Download for OS/390

Enhanced indexing

Enhanced indexing on MVS and OS/390 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 or PSF Version 3.1 or later for OS/390

You can index reports with either the standard ACIF (free with PSF) or an enhanced version of ACIF that can be ordered without charge by customers who are entitled to OnDemand.

Server printing

If you plan to use OnDemand server print or FAX, one of the following IBM program products is required:

- PSF for AIX Version 2.1 or later
- PSF/2 Version 2.00 or later

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 100 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, Windows 95, and Windows 98 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 100 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, Windows 95, and Windows 98, the standard TCP/IP support

Administrative clients

There are two administrative clients provided with OnDemand:

- **OnDemand Configurator.** The client used to configure and maintain Windows NT servers, services, and scheduled tasks. The configurator typically runs on an OnDemand Windows NT server.
- **Administrator Interface.** The client used to maintain users, groups, application groups, applications, folders, storage sets, and printers. The administrator interface runs under the Windows NT, Windows 95, and Windows 98 operating systems.

The administrative functions are Windows 32-bit applications and require the following hardware and software:

- Physical connection to the network, such as a Token Ring or Ethernet network adapter
- A minimum of 32 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 100 MB of free hard disk space
- The standard Windows TCP/IP support

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 27).

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 28).

Administrator interface setup program

The administrator interface setup program runs under Microsoft Windows NT 4.0, Windows 95, and Windows 98 and requires the same basic hardware and software as the administrative clients (refer to “Administrative clients” on page 28).

Chapter 4. The role of the OnDemand system administrator

The system administrator's job is to assume responsibility for and take care of the OnDemand environment. The OnDemand environment includes hardware, software, applications, and end-users.

- Hardware includes library and object server workstations, backup devices, archive media devices, client PCs, terminals, printers, and the networking equipment.
- Software includes the base operating system, prerequisite software, and client and server programs, configuration files and shell scripts.
- The system administrator defines OnDemand applications and decides how OnDemand manages data on the servers.
- The system administrator defines OnDemand groups and users to OnDemand and makes sure that the client software is installed and operating properly.

While the OnDemand system administrator is responsible for this collective environment from the viewpoint of OnDemand users, it is likely the system administrator is not the only person working on all these components.

Depending on the size of your organization, you might choose one person to administer OnDemand. If your organization is large, you might choose to divide administrator tasks among several people in your company. For example, an OnDemand system administrator could maintain OnDemand storage sets, printers, groups, and users, an OnDemand application administrator could maintain application groups and applications, an operating system administrator could apply base operating system upgrades and perform problem determination; and a service administrator could maintain records of system and network hardware and software and make equipment changes.

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

- Installing and maintaining OnDemand programs and prerequisite software.
- Defining and labeling storage volumes.
- Working with users to determine OnDemand application indexing and retrieval requirements.
- Managing OnDemand databases and defining applications.
- Managing and defining storage sets and storage nodes.
- Loading data into application groups.

- Defining OnDemand groups and users.
- Defining OnDemand printers.
- Answering end-user questions.
- Monitoring server activity and tuning system parameters.
- Managing backup and recovery of critical databases, files, and file systems.
- Mounting backup tapes each night.
- Solving server, network, and application problems.
- Installing equipment.

OnDemand provides an administrative client to allow administrators to maintain OnDemand objects through an easy-to-use, graphical user interface. The administrative client runs on a personal computer running Windows NT, Windows 95, or Windows 98. The administrative client allows administrators to define and maintain application groups, storage sets, storage nodes, folders, printers, applications, groups, and users. The administrative client includes features that allow administrators to process sample report data and create ACIF indexing parameters and logical views.

OnDemand provides a set of administrative commands to help an administrator to maintain the system. For example, OnDemand provides commands for loading and unloading data, maintaining the database and cache storage, and working with documents. Many of the administrative commands can be configured to run automatically, on a regular schedule.

Part 2. Planning for OnDemand

We provide this section of the book as an information and planning source for the OnDemand system administrator. Other readers may be technical and service support personnel, planners, database administrators, network administrators, and anyone else who has responsibility for making decisions about business systems, such as people responsible for physical site planning, production systems operations, and backup and recovery of business data.

This part of the book describes activities that we recommend the system administrator perform to plan for the installation of OnDemand and prepare for the production operation of OnDemand.

Topic	Page
Preparing for OnDemand	35
Planning for reports and documents	37
Planning for applications	47
Estimating storage requirements	61
Planning a storage subsystem	89
Planning processor requirements and server configurations	97
Assessing network and printer requirements	111
Planning for client programs	115
Planning for OnDemand programming interfaces	119
Planning for backup and recovery	123

Chapter 5. Preparing for OnDemand

This chapter of the book contains an outline that you may find helpful as you prepare your organization for the OnDemand environment and perform a pilot rollout of the system.

- Work with a single department or group of end-users. Send a memo to the end-users to explain how OnDemand will affect their daily work.
- Develop an end-user training course or contract with IBM to provide training for OnDemand.
- Establish a support plan for the end-users. The plan should include the names and phone numbers of persons to contact for assistance and a list of troubleshooting tips.
- Develop a set of evaluation and completion criteria that you can use to compare against the actual performance of the system.
- Choose a report or set of reports for an initial migration to OnDemand. Obtain hardcopy of the reports.
- Review the reports and determine the type of indexing required. Then select the fields from the reports for index, filter, and display fields.
- Review the selections with the end-users. Verify that the index, search, and display fields allow the end-users to retrieve the data that they need.
- Determine end-user viewing requirements.
- Identify the type of data contained in the report and determine how you will create index data.
- If you plan to use ACIF to index the report, identify resources used by the report. Resources are reusable objects found on pages of a report, such as overlays and page segments. Overlays contain constant data that is merged with variable report data during printing and viewing. Page segments are graphics and images that appear on pages of a report file, such as a company logo. Resources can be used by different applications. If you plan to index report files on an MVS or OS/390 system, resources can be gathered into a resource group file. If you plan to index report files on the server, you must either transmit the resource group file to the server or provide access to the file using the Network File System (NFS).
- If you plan to use ACIF to index the report, decide where to index reports: the MVS system, OS/390 system, or an OnDemand server. Determine how to transmit report file and index data from the MVS or OS/390 system to the server. We recommend that you use Download to transmit data from the JES spool to file systems on the server. See *PSF for MVS: MVS Download Guide* or *PSF for OS/390: Download for OS/390* for details about installing, configuring, and using Download.

- Configure optical and tape storage devices on the server. Define and configure the storage devices to ADSM. Create ADSM storage management policies to support OnDemand.
- Use the OnDemand administrative client to create the application group and applications required to support the report files.
- Use the administrative client to define the folders that end-users open to access data stored in the application group.
- Use the administrative client to define users and groups to OnDemand.
- Index the report files.
- Load the report files, resources, and index data into the application group.
- Begin end-user testing. Survey the end-users about initial testing and index, search, and display fields.
- Collect additional information from end-users, report suppliers, production scheduling, and capacity planning. For example:
 - Report file frequency.
 - Number of pages in a report file.
 - Number of indexed items, such as statements, in a report file.
 - Access frequency and patterns.
 - Length of time until the report file is out of date and the length of time that you want OnDemand to maintain the report file.
- Use the administrative client to update OnDemand with the information that you collect.
- Survey users about their satisfaction with OnDemand. Compare the performance of the system with the evaluation and completion criteria that you established. Prepare a list of issues to resolve.
- Update your company's vital records list to include the hardware and software required by the OnDemand system. Update your company's operations and recovery manuals with information required to operate, support, and backup the OnDemand system.

Chapter 6. Planning for reports and documents

This section of the book contains information that can help you plan for the reports that you will be storing in OnDemand. You can use the information to help determine the hardware configuration you need to support your OnDemand system. We list questions you might ask users of the reports, provide information about the types of data you can store in OnDemand, and provide information about indexing data.

Understanding end-user requirements

Planning for OnDemand requires that you understand how the system will be deployed, who will use the system and how they will use it, and other end-user requirements. Answers to these questions provide information that allows you to configure the proper OnDemand system, including the storage and network configuration, to support applications and end-users.

- Will you operate a single OnDemand server or a network of servers?
- What print data streams must the system support? What is the logical organization of the print data streams?
 - *Page* organization: a consistent stream of pages of transaction or ledger data.
 - Logical *groups* of information, such as statements or policies.
 - Data that may not have a consistent format, such as reference materials or product literature.
- Will OnDemand support short-term report file management, long-term archival storage, or both?
- What are the data volumes?
- How much time is available to load report files into OnDemand? Daily? Weekly?
- How long must OnDemand maintain report file data?
- How many concurrent, logged-on users do you anticipate?
- How many active users do you anticipate?
- What is the transaction rate of the active users?

Input data formats supported by OnDemand

OnDemand supports several types of report file data streams:

- AFP print data streams (AFP or MO:DCA-P), including line data mixed with AFP structured fields and line data formatted with a page definition.
- IBM System/370 line data with ANSI or machine carriage control characters.
- Unformatted ASCII data that is typically generated in the workstation environment.
- Adobe Portable Data Format (PDF) files. OnDemand provides 16- and 32-bit versions of Adobe PDF viewing software so that users can view PDF documents stored in OnDemand. The 16-bit version of the PDF viewing software can be run under Windows 3.1 or later. The 32-bit version of the PDF viewing software can be run under Windows NT, Windows 95, and Windows 98.
- Image files in the following formats:
 - BMP (Bitmap). A file that contains a bit-mapped graphic.
 - GIF (Graphic Interchange Format). A bit-mapped color graphics file format for IBM-compatible computers. GIF uses an efficient compression technique for high resolution graphics.
 - JFIF (JPEG Format Image File). A file that contains image data compressed using the JPEG (Joint Photographic Experts Group) standard.
 - PCX (Picture Exchange Format). A file that contains a graphic in the PCX file format, widely used by PC applications, such as the PC Paintbrush program.
 - TIFF (Tagged Image File Format). A bit-mapped graphics image format for scanned images with resolutions up to 300 DPI. TIFF simulates gray-scale shading.

OnDemand also supports user-defined file formats so that you can maintain other types of data in the system. For example, you can define an application for HTML documents. When you define the application you identify the file type of the data. The file type determines the program that is started when a user retrieves the file. For example, if the file type is HTM, Netscape Navigator could be started to view the file.

In the MVS or OS/390 environments, OnDemand allows applications that produce 1403 or 3211 data streams to take advantage of overlays, page segments, and typographic fonts. This is done using a page definition that specifies how data are mapped on the page. The definition allows text to be moved to different positions on the page, fonts to be changed, and conditional processing. When combined with a form definition, the page definition allows sophisticated pages to be produced by existing line data applications without changing the application that generates the data.

OnDemand can convert line data to AFP data before loading it into an application group. The resulting AFP data stream could add color or an

electronic form to line data, making presentation of the information more effective. However, archiving line data without conversion usually results in much higher compression ratios.

AFP supports graphics, presentation text, image, and bar code objects. Archiving AFP data streams allows full-fidelity viewing of presentation text and image objects. Viewing of graphics and bar-code objects is not currently supported. For example, end-users can retrieve and view customer statements that OnDemand presents using an electronic form, fonts, and images. The end-user views a copy of the statement that appears the same as the statement the customer received in the mail. AFP also supports navigation within a report file, using a table of contents.

When you archive report files that contain AFP data, you must store the resources in OnDemand. Resources include overlays, page segments, form definitions, and fonts. Resources must be resident on the processor where the data is indexed. If data will be indexed on the MVS system, the indexing program must gather the resources into a resource group so that the resource group can be transferred to the OnDemand server, where the data loading program processes the data. If data will be indexed on an OnDemand server, the resources must be resident on the OnDemand server (or be accessible from the OnDemand server) where the indexing program processes the data.

Indexing data

One of the main tasks that you perform with OnDemand is indexing report files. When you index a report file, OnDemand extracts data from the report file and stores it in the database, using the fields that you defined for the application group. When an end-user opens a folder, OnDemand displays a list of search fields, which represent the database fields. The end-user enters values in the search fields. OnDemand compares the search values with the values in the database fields and retrieves the items that match the query.

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

Index information can be added to report files at the same time the application program generates the data or the output print data stream can be processed by an indexing program. The information and examples in the

sections that follow assume that you will use ACIF to index reports. ACIF is a post-processing program that tags key information within report files and builds the indexes using parameters that you create. You can index report files with ACIF on an OnDemand Server or an MVS system. Refer to “Indexing data with other programs” on page 42 for information about other methods you can use to index reports.

Indexing data with ACIF

ACIF is a powerful tool for indexing the print data streams of MVS system application programs. ACIF indexes report files based on the organization of the data in the report file. You can optionally convert the print data streams into AFP documents. ACIF processes three input sources:

- Index parameters that specify how MVS system line data should be indexed. You create the index parameters when you define an OnDemand application.
- AFP resources needed to view and print the data, if the data was created by an AFP application.
- A print data stream.

The output of ACIF is either a fully composed AFP data stream or the original line data input. ACIF automatically converts line data files to AFP files (you can specify no conversion), produces an index file that OnDemand uses to create DB2 index data, and optionally, collects resources into a resource group file. ACIF only produces a resource group file for AFP files. To create a resource group file, ACIF must have access to the resources required by the input data stream. OnDemand stores resources on storage volumes and retrieves the resources associated with a specific item when an end-user selects the item for viewing.

The report files that you index with ACIF generally fit into one of two categories:

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

Report files that do not contain logical items or sorted line data can usually be indexed with the report indexing method.

Document indexing

Document indexing supports report files made up of logical items or subsets of report files that can be uniquely identified by data such as an account number or customer name. OnDemand returns a list of the items that match a

user's query and transfers individual items to the OnDemand client program for viewing and printing. OnDemand supports up to 32 fields as indexes or filters for document-type data. The fields do not have to be sorted and can contain numeric or text information. The fields are stored in the database as indexes or filters. Figure 6 shows an example of a report file and document indexing.

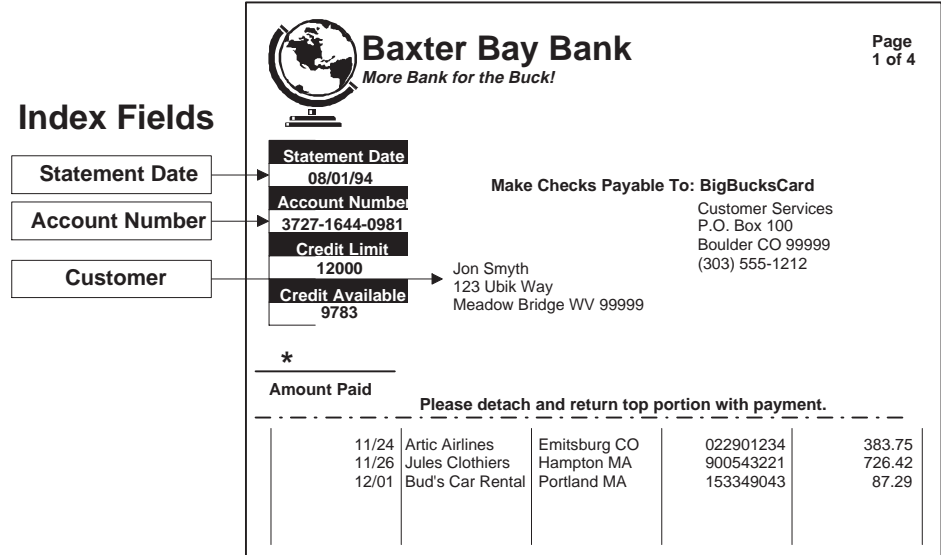


Figure 6. Document Indexing Method

Report indexing

Report indexing allows end-users to retrieve sorted report data and view the first occurrence of the value specified in the query. OnDemand divides the report data into groups of pages and stores the first and last index values contained in each group of pages in the database. When the end-user enters a query, OnDemand returns a list of the items that match the query. When the end-user selects an item for viewing, OnDemand performs a text search within the item for the value specified by the end-user. The OnDemand client program displays the first page that contains the value specified by the end-user. OnDemand uses a single, unique sorted index value for retrieval of report file data, for example, invoice number or transaction identifier. Figure 7 on page 42 shows an example of a report file and report indexing.

The diagram illustrates the report indexing method. It shows two pages of a 'Customer Invoice Report' for the date 05/31/94. Each page has a header with '1', 'Customer Invoice Report', 'Date: 05/31/94', and 'Page 1' or 'Page 2'. The data is organized into columns: Invoice Number, Customer, Date, Price, and Account Number. In both pages, a box labeled 'Sorted Invoice Numbers' has an arrow pointing to a list of invoice numbers that are sorted in ascending order. The first page lists invoice numbers from 00453051 to 00453060, and the second page lists them from 00453151 to 00453160.

Page 1		Date: 05/31/94		Page 1	
- Invoice Number	Customer	Date	Price	Account Number	
00453051	Barnshorn	3/16/94	22.50	0635588	
00453052	Smyth	3/16/94	30.10	0655388	
00453053	Ruben	3/16/94	88.60	0608458	
00453054	Dockert	3/16/94	21.80	0667588	
00453055	Doll	3/16/94	151.75	0603712	
00453056	Eckert	3/16/94	109.22	5063521	
00453057	Defino	3/16/94	320.90	5276586	
00453058	Malstrom	3/16/94	865.60	8669306	
00453059	Brubakan	3/16/94	12.14	8673066	
00453060	Little	3/16/94	82.81	6061665	

Page 2		Date: 05/31/94		Page 2	
- Invoice Number	Customer	Date	Price	Account Number	
00453151	Pratte	3/16/94	12.50	1063558	
00453152	Noble	3/16/94	320.11	1065538	
00453153	Harvey	3/16/94	8.60	1060845	
00453154	Edwards	3/16/94	71.30	1066758	
00453155	Davis	3/16/94	11.75	1060371	
00453156	Schmidt	3/16/94	19.23	6352183	
00453157	Pearl	3/16/94	120.49	6658662	
00453158	Lockhart	3/16/94	65.60	0084581	
00453159	Simpson	3/16/94	2.14	6753382	
00453160	Dean	3/16/94	72.31	6166568	

Figure 7. Report Indexing Method

Indexing data with other programs

ACIF is designed to index report files that contain line data with a consistent structure and format. Although you can use ACIF to index less structured report files and AFP files, there are some limitations to the indexing that ACIF can perform on these types of files. Alternatives you might consider when indexing AFP data include the following:

- Generate index data in the application program.
- Use AFP applications that automatically generate index data, for example, the IBM Document Composition Facility (DCF).
- Use the AFP API to generate indexed AFP data.
- Insert AFP records in the data stream that can be reliably extracted by ACIF.
- Create index data that supports the OnDemand generic index file format.
- Use the OnDemand PDF Indexer to index Adobe PDF files.

Generate index data in the application program

As an alternative to using ACIF to index report files, you can create index information in the application program that generates the report file. Some

application programs already provide support to add indexing information. However, you may find it necessary to modify your application program to add indexing functions.

Generate index data with AFP application programs

The IBM Document Composition Facility (DCF) is a product that can be used to create indexed AFP data. The primary function of DCF is to prepare and format documents for printing. Along with its many other features, DCF provides the ability to add both group-level and page-level indexing tags. DCF allows specific indexing information to be included in the output print data stream. You can process the output file created by DCF with ACIF to create the index file the OnDemand data loading program processes.

In addition, there are several popular third-party programs that can produce indexed AFP data.

Generate index data with the AFP Application Programming Interface

The AFP Application Programming Interface (AFP API) is a product that can be used to index print data. Using AFP API, a programmer who knows COBOL or PL/1 can format complex output without knowing the syntax and semantics of MO:DCA-P. Using AFP API, you can index AFP files with both group-level and page-level indexing tags, allowing more specific information to be included in the output file. Index information is added at the same time that the application program generates the print data. You then process the output file with an indexing program, such as ACIF, to create the index data that the OnDemand data loading program stores in the database.

Insert AFP records in the data stream

A common way of indexing unstructured, mixed-mode data is to add NOP structured fields to the data stream. ACIF can then be used to index the data stream to locate the NOP fields and extract index values.

Generate index data using the OnDemand generic index file format

OnDemand provides a generic index file format so that you can create index data for files that cannot be indexing using other methods, such as ACIF. For example, you can create a generic index file that contains index data that OnDemand stores in the database about a set of files in the TIFF image format. A generic index file contains index fields and values and offset and file information. OnDemand creates a row in the database for each index record contained in the file. Users can search the database using any combination of the index fields that you defined in the generic index file. The *Indexing Reference* describes the generic index file format.

Generate index data using the OnDemand PDF Indexer

The OnDemand PDF Indexer is a utility you can use to extract index data from or generate index data about Adobe PDF files. The index data can enhance your ability to store, retrieve, and view PDF documents with OnDemand. The PDF Indexer processes PDF input files. A PDF file is a distilled version of a PostScript file, adding structure and efficiency. A PDF file can be created by Acrobat Distiller or a special printer driver program called a PDFWriter. The PDF Indexer supports PDF Version 1.2 input and output files. The *Indexing Reference* provides details about the OnDemand PDF Indexer and shows examples about how to use it to process PDF input files.

Indexing reports using date fields

OnDemand requires that you specify a date field to index report files. OnDemand uses the date to determine one report's data from another, when querying the database. OnDemand also uses the date to determine when to remove reports from magnetic storage and how long to maintain report data in the application group.

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

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

Note: OnDemand is Year 2000 ready. When OnDemand is used in accordance with its associated documentation, it is capable of correctly processing, providing, or receiving date data within and between the twentieth and twenty-first centuries, provided all other products (for example, software, hardware, and firmware) used with the product properly exchange accurate date data with it.

Running ACIF on MVS and OS/390 Systems

Indexing report files on MVS and OS/390 systems can provide certain benefits in a production archival process, including:

- The potential to balance processing and resources across the MVS or OS/390 system and OnDemand servers. If you index data as a step in the jobs that generate reports for OnDemand, you can transmit report and index data directly to one or more OnDemand servers. The OnDemand servers can load the data into the appropriate application groups. Additional storage and processing cycles to do indexing on the OnDemand server would not be required.

- The indexing program can collect resources for reports into a resource group file. You can transmit the resource group file to the OnDemand server where the OnDemand data loading program processes the report or you can make the resource group file available to the data loading program using NFS. If you index reports on an OnDemand server, you must either transmit the individual resources to the OnDemand server where the indexing program processes the resources or make the resource library available to the indexing program using NFS.

Running ACIF on OnDemand servers

Indexing reports on an OnDemand server can provide certain benefits in a production archival process, including:

- You can schedule and run the indexing program as part of the data loading process. The *Installation and Configuration Guide* for your server describes how to configure the OnDemand data indexing and loading programs to run on a regular schedule, automatically processing input data.
- You can off-load part of the MVS and OS/390 processing load to an OnDemand server, by running the indexing program on the OnDemand server. If your OnDemand system consists of several OnDemand servers, you can distribute the indexing workload among the servers, increasing the resources available to process reports and archive data in OnDemand.

Chapter 7. Planning for applications

This chapter contains information that can help you plan for the application groups, applications, and folders you define to OnDemand.

Overview

When you install and configure OnDemand, you create and initialize a set of database tables that form the internal framework of the system. When you define report files to OnDemand, you add application-specific tables and control information to the system.

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

Users of an OnDemand system open a folder to query and access archived data. The folder is typically the only OnDemand object visible to users. A folder provides users the means to search for and retrieve data archived in OnDemand. Users open a folder and construct queries and display the report files that you stored in the application groups referenced by the folder. A folder can reference one or more application groups.

An application group represents the index and report file data that you archive in OnDemand. The OnDemand database contains tables of application group data. Records in an application group table contain index values extracted from report files and pointers to report files located on storage volumes. An application group can contain one or more applications that have the same storage characteristics and index fields.

An OnDemand application includes a description of the physical characteristics of a report file, such as the type of data contained in the report file and the record format of the file, instructions to the indexing and loading programs that process the report file data, and information that OnDemand uses to display and print report file pages. Typically, you define an application

for each type of report file that you plan to archive in OnDemand. You can group applications that have the same storage characteristics and index fields into an application group.

You assign a unique name to each object that you define to OnDemand: the database fields that you define for an application group; you assign applications to an application group by selecting the name of the application group; users open a folder by selecting the name of the folder from a list.

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

The OnDemand folder

A folder provides users the means to search for and display related report files archived in OnDemand. Users open folders, construct queries, and display and print report files stored in application groups. When you create a folder, you define the search and display fields that appear when an user opens the folder. You map the folder fields to database fields in the application groups referenced by the folder. The database fields contain index values extracted from report files. For example, the folder search field *Customer Account Number* could be mapped to the *acct#* application group database field. OnDemand creates database records that include index values for *acct#* when you load data into the application group. When the user enters a query, OnDemand retrieves records from the database if the values of the *acct#* database field match the value that the user typed in the *Customer Account Number* search field.

When you define a folder to OnDemand, you add one or more application groups to the folder, select index fields from the application group to appear as search and display fields when an user opens the folder, and specify properties of the search and display fields. For example, you can determine the layout of search fields on the screen and specify default values for search fields.

OnDemand maintains information about the name of the folder and its structure in the OnDemand database. For example, the database contains

information that describes the search and display fields the you defined and the database fields that you selected from application groups referenced by the folder.

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

The OnDemand application group

An application group is a collection of one or more applications that have the same index fields and storage characteristics. The application group is where OnDemand stores and manages the report files that you archive in OnDemand, including the index data for report files and the permissions for the groups and users authorized to access application group data.

When you define an application group, you specify the name and type of the database fields, determine whether a field is used to index or filter data, and specify other characteristics of the fields. OnDemand creates an application group table in the database with a column for each field that you define. Users search for report files using one or more of the fields. When you index and load data into an application group, OnDemand creates rows in the table for each indexed item in the report file. An indexed item can be a logical item, such as a policy or statement, or a group of pages, depending on how the report is organized and how you index the report.

- OnDemand supports up to thirty-two index and filter fields for each application group. Index fields allow fast access to a specific record using a key, but generally require a large amount of magnetic storage to implement and require longer to load data into the application group. OnDemand uses index fields to locate the records in the database that meet the search criteria entered by the user. The index record contains the physical location of an item on a storage volume.

- Filter fields are used to refine queries, retrieving only a subset of the records found with an index field. Filter fields are generally used with an index field to identify a specific item in the application group. Filter fields can also be used to display additional information in the document list, for example, an address.
- OnDemand requires a segment field for each application group that you define. OnDemand uses the segment field to organize and maintain application group data and to locate items that match a query. The segment field can be one of two types of date or date/time fields:
 - Report Date. The date that the application program created the report file. Typically the date found on pages of a report file.
 - Load Date. The date that you loaded the report file into the application group.

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

- Some report files are made up of logical groups of information, such as statements, invoices, and policies. These groups, or logical items, can contain one or more pages of information. OnDemand can index, store, and retrieve the logical items contained in a report file. Each logical item can be indexed on up to 32 values, for example, account number, customer name, and balance. OnDemand creates a row in the database for each logical item it finds in the report file.
- Other report files may be organized differently, and may not necessarily contain logical items. For example, a report file that contains thousands of pages of transaction or general ledger data. OnDemand can index, store, and retrieve information from these report files using index values, such as date and page number, and a sorted search key. OnDemand divides these types of report files into groups of pages and indexes each group of pages. While these report files may contain logical items, it probably would not be cost effective to index every item in the report file. That is, indexing every item in such a report file would result in thousands of index records added to the database each time a report file is loaded into the application group.

When you create an application group, you choose how OnDemand stores the index data for the report files that you load into the application group. OnDemand provides two database organization methods that determine how index records are loaded into the database and how users can query for report files:

- Multiple Loads per Database Table

Each time that you load a report file into the application group, OnDemand adds the index records to an existing database table. Index records for every report file loaded into the application group are stored in the same logical database table. OnDemand maintains the application group data so that, as far as the query manager is concerned, it appears to reside in one database table. OnDemand automatically segments application group data when it grows beyond a certain size. OnDemand maintains a segment table for the application group. A segment table provides faster query performance by limiting searches to a specific table of application group data, using a date value to construct the query. We recommend that you select this database organization when the users that search for data stored in the application group do not necessarily know or care what report file generated the information they are seeking.

- **Single Load per Database Table**

Each time that you load a report file into the application group, OnDemand stores the index records in a new database table. You can define a report field for the application group so that users can easily search for and retrieve a specific instance of a report file. We recommend that you select this database organization when the users that search for data stored in the application group typically query a specific report file for the item they need.

When you create an application group, you specify storage characteristics for the data, such as the length of time that OnDemand maintains data stored in the application group and the data caching and migration values. Storage characteristics also determine whether OnDemand stores a copy of report file data in optical or tape storage, whether OnDemand automatically creates a backup copy of report file data, and how OnDemand removes index and report file data when it is no longer needed.

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

- **Database expiration processing**

OnDemand can periodically delete index data for the report files that you store in an application group after you no longer require that OnDemand maintain the report file. Database expiration processing conserves magnetic storage space. The next time that you schedule the database index maintenance process, OnDemand can reclaim the space taken by expired index records.

- **Cache migration processing**

For application groups that support optical or tape storage, OnDemand can periodically migrate report files from cache to archive storage. Cache migration processing optimizes the use of magnetic storage, while providing excellent performance for short-term report archival. As a report

file ages, and in all likelihood, accesses becomes less frequent, OnDemand can automatically copy the report file to long-term storage. Cache migration frees magnetic storage space for newer versions of the report files that you load into the application group. You can also use cache migration and offload the storage of report file data to optical and tape storage volumes during periods of little or no other system activity.

- Cache expiration processing

OnDemand can delete report files from cache storage on a schedule that you establish. Cache expiration processing also optimizes the use of magnetic storage, while providing access to the latest versions of report files.

The OnDemand application

An OnDemand application describes the physical characteristics of a report file, processing instructions for the indexing and data loading programs, and information about how OnDemand displays and prints report file pages. You can establish default settings for viewing and printing report file pages at the OnDemand application level. For example, if you select a default printer for the application, whenever an user prints an item associated with the application, OnDemand sends the item to the selected printer. Typically you define an application for each different report file that you plan to archive in OnDemand.

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

The OnDemand application is where you specify information to the indexing and data loading programs, such as the technique that OnDemand uses to compress the report file, indexing parameters, and values that OnDemand uses to process index data before loading index records in the database. OnDemand uses the indexing parameters, options, and data values that you specify to locate index data in and extract index data from the report file.

You can set up one or more logical views of the report file. A logical view determines how OnDemand displays line data report files and governs other viewing characteristics. For example, you can set up a logical view so that when an user selects an item for viewing, the OnDemand client program automatically displays the item at 200 percent zoom and locks the report heading. That is, text and illustrations appear twice as large as the original and when the user scrolls from page to page of the report file, OnDemand keeps the lines that make up the heading in the same place on the screen.

OnDemand users and groups

Each user logs on to OnDemand with a userid. OnDemand authenticates userids and determines the usage and administrative authority available to the user based on the log on userid. An OnDemand userid does not necessarily have to identify an individual user. However, for accounting purposes, you probably want to assign an OnDemand userid to each person that uses the system.

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

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

Understanding OnDemand permissions

Overview

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

If you assign a user to more than one group, the user normally obtains the permissions of all of the groups. For example, using the group settings listed in Table 2:

Table 2. Group Permissions

Group	Folders
Accounting	Student Bills
Admissions	Student Transcripts

A user assigned to both groups can access the Student Bills and Student Transcripts folders.

However, there are exceptions to this rule. See *Setting OnDemand Permissions in the Getting Started with the Administrative Client* publication for details.

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

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

Setting folder permissions

You can set folder permissions at the folder, group, and user levels. Setting permissions at the folder level provides every OnDemand user and group with the permissions that you define. Setting permissions at the group level provides every user that you assign to the group with the permissions that you define. Setting permissions at the user level provides the user with the permissions that you define. In general, you set the least restrictive permissions at the folder level and the most restrictive permissions at the user level. That is, you allow every OnDemand user to open the folder. However, you assign administrator permission to a single user.

You can set four types of folder permissions:

- **Access.** Users can open the folder with OnDemand client programs and search for and retrieve data from the application groups referenced in the folder.
To search for and retrieve items, users must have access permission for the folder and access permission to one or more of the application groups referenced in the folder.
- **Fields.** Users can open the folder with OnDemand client programs and can modify the folder field information with the administrator interface. OnDemand maintains a set of folder fields for each user given fields permission for the folder.
- **Named Queries.** A named query is a set of search criteria, saved by name, that can be selected and restored into folder search fields. OnDemand supports two types of named queries: public, that is, a named query that is available to all users that can open the folder, and private, that is, a named

query available only to the user that created the named query. Users can be given authority to view, create, modify, and delete named queries.

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

Setting application group permissions

You can set application group permissions at the application group, group, and user levels. Setting permissions at the application group level provides every OnDemand user and group with the permissions that you define. Setting permissions at the group level provides every user that you add to the group with the permissions that you define. Setting permissions at the user level provides the user with the permissions that you define. In general, you set the least restrictive permissions at the application group level and the most restrictive permissions at the user level. That is, you allow every OnDemand user to access data stored in the application group. However, you assign administrator permission to a single user.

You can set six types of application group permissions:

- **Access.** Users can search for and retrieve data stored in the application group using OnDemand client programs.
- **Document.** Determines the types of document functions users can perform. The default document permissions are view, print, FAX, and copy.
- **Annotation.** Determines the types of annotation functions users can perform. The default annotation permissions are view and add.
- **Logical Views.** Logical views determine how OnDemand displays report file pages. Users can define their own logical views with OnDemand client programs.
- **Administrator.** An application group administrator can modify and delete the application group. An application group administrator can change user and group permissions, add and remove users and groups from the application group, change message logging options, update the storage management settings for the application group, and make changes to the application group field information.
- **Query restriction.** Limits access to application group data. You typically set up a query restriction to limit the data that a specific user or group of users can access.

About OnDemand names

When you administer OnDemand, you assign names to various objects.

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

When naming a user, the name that you specify:

- Can contain 1 to 20 characters (bytes).
- Cannot include the ' (apostrophe), * (asterisk), % (percentage) + (plus), _ (underscore), [(left bracket),] (right bracket), " (double quote), or blank character.
- Must be unique to the library server.
- By default, OnDemand converts lowercase letters in a user name to uppercase (laguarde is stored as LAGUARDE).

Note: If you define a logon user exit, you can determine the characteristics of OnDemand userids in your environment.

When naming a group, the name that you specify:

- Can contain 1 to 20 characters (bytes).
- Cannot include the: ' (apostrophe), * (asterisk), % (percentage) + (plus), _ (underscore), [(left bracket),] (right bracket), " (double quote), or blank character.
- Must be unique to OnDemand.
- Can be mixed case; however, OnDemand ignores the case (LaGuarde is the same as laguarde).

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

- Can contain 1 to 60 characters (bytes), including embedded blanks.
- Cannot include the: ' (apostrophe), % (percentage), _ (underscore) character, [(left bracket),] (right bracket), or " (double quote).
- Can be mixed case; however, OnDemand ignores the case (LaGuarde is the same as laguarde).
- An application name must be unique to the application group where you assign the application.
- An application group or folder name must be unique to OnDemand.

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

- Can contain 1 to 18 characters (bytes).

- Must begin with the letter A through Z.
- Can include A through Z, 0 through 9, @ (at sign), \$ (dollar), _ (underscore) character, and # (number sign).
- Can be mixed case; however, OnDemand doesn't create a unique name (rDate is the same as rdate).
- Must be unique to the application group.
- Cannot be any of the OnDemand reserved words:

annot	doc_name
comp_len	doc_off
comp_off	prt_nid
comp_type	resource
doc_len	sec_nid

- Cannot be any of the SQL reserved words, such as **date**. For a list of reserved words, see the SQL reference for your database management product.

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

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

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

- Can contain 1 to 60 characters (bytes), including embedded blanks.
- Cannot include the: ' (apostrophe), % (percentage), _ (underscore) [(left bracket),] (right bracket), or " (double quote) character.
- Can be mixed case.
- Must be unique to the folder.

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

- Can contain 1 to 60 characters (bytes).
- Can be mixed case; however, OnDemand ignores the case (LaGuarde is the same as laguarde).
- Must be unique to OnDemand.

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

- Can contain 1 to 60 characters (bytes).
- Can be mixed case; however, OnDemand ignores the case (LaGuarde is the same as laguarde).
- Must be unique to the storage set.

When naming a printer, the name that you specify:

- Can contain 1 to 60 characters (bytes).
- Can be mixed case; however, OnDemand ignores the case (LaGuarde is the same as laguarde).
- Must be unique to OnDemand.

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

- Must be a valid AIX printer queue name.

OnDemand field types

When you define an application group to OnDemand, you create a structure for a database table with the index and filter fields that you define. When you store a report file in the application group, OnDemand extracts index data from the report file, places the index data into the proper database fields, and inserts rows into the application group table. The database fields that you define for the application group can contain different types of data. When you define the database fields, you select a data type for each field. The field type tells OnDemand what kind of data can be stored in the field.

When you define a folder to OnDemand, you define search fields, where users enter values to construct queries, and display fields, that OnDemand uses to display items in the document list. A folder supports the same field data types as an application group, and three additional field types: *application group*, *segment*, and *text search*.

Table 3 contains a list of the field types supported by OnDemand.

Table 3. Database and Folder Field Types

Field Type	Field Size	Description
Small Integer	2 bytes	Contains whole numbers between $-32,767$ and $32,767$.
Integer	4 bytes	Contains whole numbers between -2147483648 and 2147483647 . Integer fields require more space to store in the database than small integer fields.
Decimal	8 bytes	Contains numbers between -10^{307} and 10^{308} with up to 15 significant digits. You can store money values in a decimal field, and use the precision attribute to format the decimal places.
String	1—254 bytes	Contains letters, numbers, special symbols, such as the % and #, and any other printable character. OnDemand supports fixed size strings and variable size strings. OnDemand requires an additional four bytes for a variable size string.

Table 3. Database and Folder Field Types (continued)

Field Type	Field Size	Description
Date	2 bytes	Contains a valid date from January 1, 1970 to December 31, 2069. If you need to index reports that contain dates that occur before January 1, 1970, you must define the date as a String field. OnDemand checks a date value to make sure it is valid.
Time	2 bytes	Contains times of day, stored in three-second increments, since midnight, and limited to 24 hours.
Date/Time	4 bytes	Contains both a date and time value. The date can be from January 1, 1970 to December 31, 2038. The time is stored in one-second increments.
Date/Time (TZ)	4 bytes	Contains both a date and time value. The date can be from January 1, 1970 to December 31, 2038. A Date/Time (TZ) field is exactly like a Date/Time field, but uses the time zone set by the time zone environment variable on the client PC. You must make sure that the time zone variable is set whenever the PC is restarted.
Application Group	N/A	For a search field, contains a list of application groups referenced by the folder. For a display field, lists the name of the application group where OnDemand found the item.
Segment	N/A	Contains a list of segment tables for an application group. Each item in the list represents a different segment of application group data.
Text Search	1-254	Used to search a document for a specified string before transferring the document to the client. Contains letters, numbers, special symbols, such as the % and #, and any other printable character.

When you create a folder that contains more than one application group, you can create an *application group* field. If enabled for query, users can select the name of the application group that OnDemand searches, rather than searching all of the application groups referenced by the folder (the default). If enabled for display, OnDemand displays the name of the application group next to each item that matches the query. Only one application group field can be defined for the folder.

You can define a *segment* search and display field for a folder. If enabled for query, the segment field provides users the ability to select a specific segment of application group data to search. OnDemand segments application group data by date. The OnDemand client programs list all of the segment tables in the application group. The user can select a segment of data to search. Only one segment field can be defined for the folder.

A *text search* field can be used to further refine the results of a query. If the document does not contain the text string specified in the text search field, it is not added to the document list. A typical use of a text search field is to provide users an additional search field without incurring database overhead. For example, assume a report is indexed on date and transaction number. A text search field would allow users to optionally enter a customer's name, phone number, or any other information contained in the document or documents that the user needs to retrieve (the information is not contained in the database).

Chapter 8. Estimating storage requirements

Overview

Estimating storage requirements for an OnDemand system begins with understanding and documenting end-user needs for storing and accessing data. “Chapter 6. Planning for reports and documents” on page 37 provides details about collecting end-user requirements.

Before you translate end-user needs into a storage subsystem that supports OnDemand, you must assess various operational and management issues. For example, OnDemand supports up to 32 index and filter fields for each report that you archive. Typically users do not require that many index and filter fields to locate data or a particular report. The number of index and filter fields that you define has a significant effect on the amount of magnetic disk space required for the OnDemand database. It is important to work with end-users and understand their data retrieval requirements. Define only the number of index and filter fields needed to balance end-user requirements with available magnetic disk space.

Caching reports can have a significant effect on the amount of magnetic storage that you configure on the system. OnDemand supports cache storage on magnetic storage volumes for high speed access to reports. Carefully review how end-users locate and retrieve information and other requirements for each report that you plan to store in OnDemand. For example, if most accesses occur in the first 90 days after reports are stored in OnDemand, you probably want to store reports in cache storage for that length of time. Choose a cache time frame for each report that meets end-user requirements and makes the best use of available magnetic disk space.

There are several areas that you must examine to determine the total amount of magnetic, optical, and tape storage required to support an OnDemand system. The following factors determine the amount of magnetic storage required:

- Storage space for required program products and system software, including the base operating system, OnDemand server software, the database manager, and ADSM.
- Storage space for configuration and control files.
- Storage space for the system logging facility.
- Temporary storage space for reports received from other systems. In general, you should allocate enough magnetic storage space to temporarily

hold the largest report to be loaded into OnDemand. In many companies, most occurrences of a report are similar in size. However, there may be times when a report is much larger than average. For example, the end of the month or the end of the quarter.

- Temporary storage space for indexing a report on the OnDemand server.
- Temporary storage space for loading a report on the OnDemand server.
- Cache storage. This may be zero, for reports that do not require cache storage. However, a very large amount of magnetic storage space may be required for reports that remain in cache storage for several months.

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

- Storage space for the OnDemand database, including control information and report index data. The amount of space that you need to allocate for the OnDemand database is a factor of the number of items contained in a report and the number of index fields and filters. For reports that contain sorted transaction data, OnDemand creates one index row for each 100 page segment ³ of the report.

For reports that contain logical items, such as statements, policies, and bills, OnDemand creates one index row for each logical item in the report. Typically the database space required for indexing line data is much less than that required for indexing reports made up of logical items. Index fields provide fast lookup, but require significantly more table space than filter fields.

The length of time that index data must remain in the database also affects the amount of storage space needed for the OnDemand database.

- Storage space for database log files. This includes space for primary and archived log files. If you use ADSM to manage DB2 archived log files, you will need additional primary log file storage space, but you can eliminate the archived log file storage space requirement (on disk).
- Storage space for the ADSM database and recovery log.
- Storage space for temporary data created by OnDemand programs that support server print and FAX.
- Storage space for temporary data created by OnDemand programs that import index data from archive media to the database.

The following factors determine the amount of optical storage required:

3. The examples and calculations that follow assume that OnDemand creates one indexed item for every 100 pages in a report that contains transaction data. This is a parameter that you can configure when you index a report with ACIF. The *Indexing Reference* provides details about using ACIF.

- The size and number of the reports that you archive in OnDemand and copy to optical storage.
- Backup copies of reports stored in an application group.
- The compression ratio achieved.
- Log file storage, if you use ADSM to maintain DB2 archived log files.
- Backup image file storage, if you use ADSM to maintain DB2 backup image files.

When you calculate optical storage requirements, you must determine the number of storage volumes and libraries required to support your system. A single optical library can support an enormous amount of data, depending on the compression ratio achieved. For example, an IBM 3995-C68 optical library can hold up to 1.341 TB (unformatted) of data online. If OnDemand can achieve 6:1 compression on files archived in the library, you can store over seven terabytes of data in the library.

You can replace full optical storage volumes as needed, if the availability requirements of your system allow you to do so. For example, you may decide to remove full storage volumes from a library one year after OnDemand copies report data to the storage volume. You can replace the full storage volumes with newly initialized storage volumes and continue to store new reports in OnDemand, with the latest versions of reports available in the library. However, if you must keep many years of data online in the library or you store massive amounts of data in your application groups, several optical libraries may be required.

Storage hierarchy

There are several different storage management strategies that can be implemented with OnDemand and ADSM.

ADSM is a hierarchical storage management system that can manage storage pools of disk devices, optical devices, and tape devices. ADSM allows data to be migrated from one storage pool to another using criteria defined by the administrator. In most installations, OnDemand will not use the hierarchical storage management capabilities of ADSM, because of the time it takes to migrate data from one storage medium to another.

OnDemand provides a high-speed disk cache outside ADSM. In general, OnDemand should only use ADSM for managing storage pools of optical or tape devices. When OnDemand loads a report, it can simultaneously write the file to cache and ADSM-managed media storage. OnDemand also offers the option to store reports only in cache storage and then later migrate files to ADSM-managed optical or tape storage. We highly recommend that you select

the option where OnDemand writes files to cache and ADSM at the same time. This approach eliminates the need to backup cache storage, because files written to cache are already backed up on optical or tape storage. Writing reports to cache storage and ADSM-managed storage eliminates the need to migrate large amounts of report data.

OnDemand deletes reports from cache storage on a schedule that you determine. For example, you can specify that OnDemand deletes reports after they have been stored in cache storage for 90 days. The amount of time to cache data on disk can be specified at the application group level.

The management and life of data stored in ADSM by OnDemand is specified in ADSM, independently of OnDemand. Once data is stored into ADSM, you can let ADSM manage it completely separate from OnDemand.

Data compression

OnDemand compresses reports using a variety of different data compression algorithms before storing data on magnetic cache and on optical storage volumes. The compression ratio achieved has a significant impact on the amount of space required for cache, optical, and tape storage.

Compression ratios vary widely depending on the type and format of the data and cannot always be estimated accurately by examining the data. You can expect to achieve between 2:1 and 15:1 compression for AFP documents and up to 30:1 compression for line data reports. Compression for AFP documents is based on the output data file produced by ACIF, and not the input file, which could have been line data. When ACIF formats line data with a page definition, it may increase the size of the data by adding AFP controls for positioning text.

To properly size an application, we recommend that you measure the compression ratio achieved with sample report data. This can be done by using the `arsadmin` command.

- For reports that contain logical items, such as statements and policies, use the following format of the command:

```
arsadmin compress -l 200000 -s report_file -o output_file
```

Where `report_file` is the input file that you want to measure and `output_file` is the compressed output.

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

- For reports that contain line data and include a sorted transaction value, such as a general ledger, first determine the size of an indexed group of pages, for example, 100 pages. Then extract a group of pages from a larger report and run the arsadmin command. Use the following format of the arsadmin command:

```
arsadmin compress -s group_pages -o output_file
```

Where `group_pages` is a file that contains a representative group of pages from a larger report and `output_file` is the compressed output.

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

See the *Administrator's Reference* for more information about the arsadmin command, options, and parameters.

Calculating magnetic storage requirements

Base system software

Most OnDemand servers require 2 GB of magnetic disk storage for the base operating system software, swap space, temporary work space, the database manager, ADSM, and OnDemand software and configuration and control files.

The base system software is a per server requirement, not per application.

Download storage

OnDemand requires temporary storage space for downloading data from another system, such as an OS/390 system. Many customers download report data during the day when application programs generate the reports, but load data into OnDemand at night or during other periods of little or no other system activity. This approach requires enough disk space to hold all of the data generated in one day (or to hold all of the data before you begin loading data). To minimize the impact to users of the system, we recommend that you dedicate one or more magnetic storage volumes to data download storage.

Use the following calculation to determine data download space requirements:

$$\begin{array}{l} \text{Report} \\ \text{Download} = \text{Total data for} * 1.20 \\ \text{space} \quad \quad \text{largest cycle} \end{array}$$

Figure 8. Calculating Report Download Space

For example, if you download 400 MB of data in a single day, then the download space required on the OnDemand server is:

$$\begin{array}{l} \text{Report} \\ \text{Download} = 400 \text{ MB} * 1.20 = 480 \text{ MB} \\ \text{space} \end{array}$$

Temporary space for indexing and loading data

OnDemand requires temporary storage space when indexing and loading data on the OnDemand server. When preparing to store reports into an application group, OnDemand compresses and segments report data into storage objects and places the storage objects in cache, optical, and tape storage, depending on the storage management requirements of the application group. The temporary space required is a factor of the largest single report that you plan to store in OnDemand and whether the data is indexed on an MVS or OS/390 system or the OnDemand server. A single day's data typically includes many different files.

Use the following calculations to determine the amount of temporary space required to support data loading:

- If the report is indexed on the MVS or OS/390 system, no temporary space is required.
- If the report is indexed on the OnDemand server:

$$\begin{array}{l} \text{Temporary} = \text{Largest report} * 1.5 \\ \text{space} \quad \quad \text{file size} \end{array}$$

Figure 9. Calculating Temporary Work Space

For example, if the largest report is 400 MB and the report is indexed on the OnDemand server, the temporary space required to support data indexing and loading is:

$$\begin{array}{l} \text{Temporary} = 400 \text{ MB} * 1.5 = 600 \text{ MB} \\ \text{space} \end{array}$$

Cache storage for reports

Cache storage requirements vary based on customer requirements. You typically store reports in cache storage for a short period of time, to provide the fastest retrieval time for the most frequently used documents. As documents age, and use of the documents declines, you can migrate them to optical or tape storage volumes. You would typically cache reports when many users access the system concurrently. Because OnDemand may require between six and sixty seconds to mount an optical or tape storage volume, it is not possible to support a high transaction rate with optical or tape libraries.

Another use of cache storage is for reports that have a short life, for example, one week or one month, and do not require long-term archival. You can store these reports in cache storage and OnDemand will automatically delete them when they reach their expiration date.

OnDemand compresses reports before storing them in cache storage.

Use the following calculation to determine the amount of magnetic storage space required to support cache storage:

$$\begin{aligned} \text{Cache Storage} &= \text{Size of Data per week} \\ &\quad * \text{Number of Weeks to cache} \\ &\quad * \text{Data Compression ratio} \\ &\quad * 1.1 \end{aligned}$$

Figure 10. Calculating Cache Storage

For example, if you plan to store 2 GB of report data per week, reports must be maintained in cache storage for 12 weeks, and the compression ratio is 3:1 (0.33), the cache storage space required can be calculated as follows:

$$\begin{array}{l} \text{Cache} \\ \text{Storage} \end{array} = 2 \text{ GB} * 12 * .33 * 1.1 = 8.71 \text{ GB}$$

OnDemand database storage

OnDemand stores the index data used to locate reports in database tables. For reports that contain logical items, such as statements and invoices, OnDemand creates one database row for every item in the report. For reports that contain sorted transaction data, OnDemand creates one database row for every 100 pages of report data.

A database row contains a fixed amount of information that OnDemand uses to maintain reports (approximately 40 bytes) and any additional index and

filter fields that you define. Index fields, which provide fast lookup, require significantly more storage space than filter fields and require more time to load into OnDemand.

There are four major factors that determine the amount of magnetic disk space required for the OnDemand database:

1. The number of index and filter fields.
2. The size of the index and filter fields.
3. The number of indexed items per month.
4. The number of months that OnDemand maintains the index data on magnetic storage.

Table 4 lists the types of index fields supported by OnDemand and the number of bytes required to store a value in each type of index field.

Table 4. Index Field Type and Index Size

Field Type	Field Size
Small Integer	2 bytes
Integer	4 bytes
Decimal	8 bytes
String (fixed)	n bytes; where n is the fixed size of a string
String (variable)	$n+4$ bytes; where n is the maximum size of a string
Date	2 bytes
Time	2 bytes
Date/Time	4 bytes
Date/Time (TZ)	4 bytes

Calculating the size of the database

You can use the following calculations to determine the space required in the OnDemand database to store a report. This calculation can be used for reports that contain logical items or reports that contain a sorted transaction value.

$$\begin{aligned}
 \text{TableSize} &= (\text{Field 1 size} + \text{Field 2 size} + \text{Field n size}) \\
 \text{IndexSize} &= ((\text{Index 1 size} + 8) + (\text{Index 2 Size} + 8) \\
 &\quad + (\text{Index n size} + 8)) \\
 \text{DatabaseSize} &= ((\text{TableSize} + 40) * 1.5) + (\text{IndexSize} * 3.2)) \\
 &\quad * \text{Number of indexed items per month} \\
 &\quad * \text{Number of months to keep index online}
 \end{aligned}$$

Figure 11. Calculating database storage space

- Field n size is the size of a database field. For example, acct_num, a string field, fixed length, 12 bytes.
- Index n size is the size of a database field for which you want OnDemand to build an index. For example, report_date, a date field, 2 bytes. the OnDemand database requires eight bytes for each index that you define.
- OnDemand adds approximately 40 bytes of control information to each indexed item.
- When the report contains logical items, the number of indexed items per month is the number of statements or policies.
- When the report contains a sorted transaction value, the number of indexed items per month is the number of report pages per month divided by the size of an indexed group of pages, for example, 100 pages. You determine the size of an indexed group of pages when you index a report with ACIF.

Examples

1. The following example illustrates how to calculate the database storage space required for a report that contains logical items, such as statements. In the example, we plan to index one million items per month and keep the index data on magnetic storage for 24 months. Table 5 lists the pertinent information about the database fields.

Table 5. Database calculation example for a report that contains logical items

Field Name	Field Type	Field Size	Index or Filter
Report Date	Date	2 bytes	Index
Account Number	Fixed String	12 bytes	Index
Invoice Balance	Decimal	8 bytes	Filter
Customer Name	Variable String	20+4 bytes	Filter

$$\text{TableSize} = (2 + 12 + 8 + (20 + 4)) = 46$$

$$\text{IndexSize} = ((2 + 8) + (12 + 8)) = 30$$

$$\begin{aligned} \text{DatabaseSize} &= ((46 + 40) * 1.5) + (30 * 3.2) \\ &\quad * 1,000,000 \\ &\quad * 24 = 5.4 \text{ GB} \end{aligned}$$

OnDemand requires approximately 5.4 GB of magnetic disk space to store 24 months of report index data in the database.

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

Reports that contain line data with a sorted transaction value use a fixed type of indexing, where each database row contains the beginning value, the ending value, and the beginning page number of the group of pages. OnDemand maintains the beginning and ending values as indexes and the page number as a filter. The main parameters for the calculation are the length, in bytes, of the sorted transaction value, the number of report pages created per month, the size of a group of indexed pages, and the number of months that OnDemand maintains the index data.

In the example, we plan to index one million pages per month, in groups of 100 pages, and store the index on magnetic storage for 24 months.

Table 6 lists the pertinent information about the database fields.

Table 6. Database calculation example for a report that contains line data

Field Name	Field Type	Field Size	Index or Filter
Report Date	Date	2 bytes	Index
Begin Transaction Value	Fixed String	10 bytes	Index
End Transaction Value	Fixed String	10 bytes	Index
Page Number	Integer	4 bytes	Filter

$$\text{TableSize} = (2 + 10 + 10 + 4) = 26$$

$$\text{IndexSize} = ((2 + 8) + (10 + 8) + (10 + 8)) = 46$$

$$\begin{aligned} \text{DatabaseSize} &= ((26 + 40) * 1.5) + (46 * 3.2) \\ &\quad * (1,000,000/100) \\ &\quad * 24 = 59.088 \text{ MB} \end{aligned}$$

OnDemand requires 59.088 MB of magnetic disk to store 24 months of report index data in the database.

Database transaction log file storage

When you load a report into OnDemand, the database manager records changes to the database in a log file. If you are using DB2, when a log file fills, DB2 closes the full file and opens a new log file. When all changes to the database have been made, the DB2 closes the last log file. If you are using SQL Server, when a log file fills, SQL Server goes to the beginning of the log file and begins recording again. For DB2, after the load process disconnects from the database, OnDemand copies the closed log files to the archive log file directory. OnDemand deletes the archived log files after you create a full back up copy of the OnDemand database with the `arsdb` command.

Total log file storage space is a factor of the number and size of the log files and the length of time between full backups of the database.

When you install and configure OnDemand, you set parameters that determine the number and size of the log files and the location of the active (or primary) and archived log files.

Report load log space

$$\begin{aligned} \text{LoadLogSpace} = & (((\text{TableSize} + 40) * 1.5) + (\text{IndexSize} * 3.2)) \\ & * \text{the number of Indexed Items} \\ & * 4 \end{aligned}$$

Figure 12. Calculating Log File Storage

- The `TableSize` and `IndexSize` parameters were calculated in “OnDemand database storage” on page 67.
- OnDemand adds approximately 40 bytes of control information to each indexed item.
- The number of Indexed Items is the number of logical items or indexed groups of pages contained in the report. The number of indexed items depends on the organization of data in the report and how you index the report.

The following example illustrates the log storage required for a report, where the `TableSize` is 46, the `IndexSize` is 30, and the number of items added to the database is fifty thousand.

$$\begin{aligned} \text{LoadLogSpace} = & (((46 + 40) * 1.5) + (30 * 3.2)) \\ & * 50,000 \\ & * 4 = 45 \text{ MB} \end{aligned}$$

For the sample report, OnDemand requires approximately 45 MB of log file storage to add fifty thousand items to the database.

Active log file storage

The storage space required for the active log files is a factor of the largest report load that OnDemand must process, the maximum number of report loads that OnDemand processes at any one time, and buffer space. In addition, if you use ADSM to manage archived log files, we recommend that you triple the amount of space you calculate for the active log file storage.

- Determine the storage space required for the largest report load, using the calculation in “Report load log space” on page 71.
- Determine the maximum number of report loads that OnDemand must process at any one time. Determine the storage space required for each report load, using the calculation in “Report load log space” on page 71. Total the values.
- Double the sum of the two values. The result is the storage space required for the active log files.

It is critical that you allocate adequate free space for the active log file directory. Report loads will fail if the database manager does not have enough space to allocate log files.

Databases use several parameters to determine the amount of space available for active log files.

$$\text{ActiveLogSpace} = ((\text{logprimary} + \text{logsecond}) * (\text{logfilsiz} + 2) * 4096) + 8192)$$

Figure 13. Calculating Active Log File Storage

The default database configuration values provided by OnDemand include the following:

- The `logprimary` is 40
- The `logsecond` is 2
- The `logfilsiz` is 1000

$$\begin{aligned} \text{ActiveLogSpace} &= (((40 + 2) \\ &\quad * (1000 + 2) * 4096) + 8192) \\ &= 172 \text{ MB} \end{aligned}$$

By default, OnDemand allocates approximately 172 MB of active log file space (assuming you do not use ADSM to maintain the archived log files).

You should adjust the values of the `logprimary` and `logfilesiz` database parameters so they reflect your calculation for active log file storage. For example, suppose the log space required for the largest report load is 20 MB and you plan to run a maximum of four concurrent report loads, each

requiring 10 MB of log file storage. You must allocate a total of 120 MB for active log file storage. You can do so by changing the value of the logprimary database parameter to 28.

Archive log file storage with DB2

If you are using DB2, the storage space required for archived log files is a factor of the number and size of the log files created between full backups of the database.

Note: If you use ADSM to manage DB2 archived log files, do not allocate disk space for archive log file storage on disk. Allocate optical storage space for the archived log files (refer to “Database archived log file storage” on page 78).

We strongly encourage you to backup the database on a regular schedule. For example, after each report load or once a week. When you backup the database with the arsdm command, OnDemand automatically removes log files from the archived log file directory, releasing the space taken by files no longer needed. Taking regular backups can also reduce the time required to rebuild the database, in the event you need to do so.

It is critical that you allocate adequate free space for the archived log file directory. If the database manager does not have enough space to copy log files to the archived log file directory, it leaves the files in the active log file directory.

As a guideline, we recommend that you allocate two times the space required for the active log files. However, you must allocate enough space to hold all of the log files created between full backups of the database.

$$\text{ArchiveLogSpace} = (2 * \text{ActiveLogSpace})$$

Figure 14. Calculating Log File Storage

The following example illustrates the archive log space required, when the space allocated for the active log files is 172 MB.

$$\text{ArchiveLogSpace} = (2 * 172 \text{ MB}) = 344 \text{ MB}$$

In the example, OnDemand requires approximately 344 MB of archive log file storage.

ADSM database and recovery log storage

ADSM maintains a database of information on magnetic storage about the devices that provide optical and tape storage management for OnDemand and the storage objects maintained on optical and tape storage. OnDemand uses the information to retrieve reports from optical and tape storage volumes. By default, OnDemand stores application group data in 10 MB storage objects.

You can use the following calculation to determine the size of the ADSM database:

$$\begin{aligned} \text{ADSMDatabase} &= (\text{Data per Month} / \text{Object Size}) \\ &\quad * 700 \\ &\quad * \text{Life of data in months} \end{aligned}$$

Figure 15. Calculating ADSM Database Storage

For example, if you plan to store 8 GB of data per month into an application group, the size of a storage object is 10 MB, and ADSM must maintain the data for seven years (84 months), plan to allocate approximately 47 MB of magnetic storage space for the ADSM database. The following example shows how to calculate the size of the ADSM database, given those parameters:

$$\begin{aligned} \text{ADSMDatabase} &= (8,000,000,000 / 10,000,000) \\ &\quad * 700 \\ &\quad * 84 = 47.04 \text{ MB} \end{aligned}$$

If you require that OnDemand maintain a backup copy of data stored in an application group, double the space allocation for the ADSM database. If you plan to mirror the ADSM database, double the allocation for the ADSM database. If you plan both a backup copy of data and mirroring the database, quadruple the allocation for the ADSM database.

Temporary space for server print

OnDemand requires temporary work space to process server print requests. You must allocate enough disk space to support the maximum number of concurrent print requests that the server must manage. We strongly encourage you to define a dedicated file system where OnDemand can store temporary print files. We recommend that at least 500 MB of free space be available in this file system at all times. If your storage configuration permits, we recommend that you allocate 1 GB of free space to this file system.

Temporary space for importing index data

Note: If you do not plan to migrate index data from the database to archive media, you do not need to allocate temporary space for importing index data.

OnDemand requires temporary work space to import index data from archive media to the database. You must allocate enough disk space to support the maximum number of concurrent import requests that the server must manage. We strongly encourage you to define a dedicated file system where OnDemand can store temporary data created by the programs that import index data. The amount of space you allocate to this file system is based on the size of your database tables and the number of tables that you must import to satisfy a query for migrated data. We recommend at least 500 MB of free space be available in this file system at all times. If your storage configuration permits or the size of your database tables dictates, you may need to allocate 1 GB or more free space to this file system. For example, based on the estimate we made for the sample report statements and making some assumptions about how the data is stored in OnDemand, we need approximately 500 MB of space to store one database table. If we need to import two database tables to satisfy a query, the import program requires at least 1 GB of temporary disk space.

Calculating archive media requirements

Report storage space

When you estimate the amount of archive media space required to store a report, you must consider the size of the report, the compression ratio achieved, and the length of time that ADSM must maintain the report. Archive media can be optical storage or magnetic tape. You can use the following calculation to determine archive media space for reports:

$$\begin{aligned} \text{ArchiveSpace} &= (\text{Data per month} * \text{life of data in months}) \\ &\quad * \text{compression ratio} \\ &\quad * 1.1 \end{aligned}$$

Figure 16. Calculating Archive Media Storage Space

For example, if you plan to store 8 GB of report data per month, ADSM must maintain the data for seven years, and OnDemand can achieve a compression ratio of 3:1 (0.33), you would require approximately 244 GB of optical storage space.

$$\begin{aligned} \text{OpticalSpace} &= (8 \text{ GB} * 84) \\ &\quad * 0.33 \\ &\quad * 1.1 = 244 \text{ GB} \end{aligned}$$

Backup report storage space

You can use OnDemand to maintain a backup copy of reports stored on archived media. You typically maintain two copies of reports that are critical to the operation of your company or difficult or impossible to recreate.

OnDemand provides two methods for maintaining a backup copy of reports:

- Configure a copy storage pool in ADSM. With this method, ADSM manages a backup copy of files stored in a primary storage pool independently and transparently to OnDemand. ADSM retrieves the backup copy if the primary copy becomes damaged, lost, or unusable.
- Configure a secondary storage node in OnDemand. With this method, OnDemand stores two copies of the report, one copy in the primary storage node and one copy in the secondary storage node. However, if the copy in the primary storage node becomes damaged, lost, or unusable, OnDemand does not automatically retrieve the copy from the secondary storage node.

If you need to maintain two copies of a report on archive media, we encourage you to configure a copy storage pool in ADSM.

If you require that OnDemand maintain a backup copy of a report, double the report storage space you calculated in “Report storage space” on page 75.

Database backup image file storage

Note: If you do not plan to use ADSM to maintain DB2 database backup image files, do not allocate space for the backup image files on archive media.

The storage pool where ADSM maintains DB2 database backup image files must contain enough storage to hold the files needed to recover your database. The number of backup image files ADSM maintains depends on the type of database backups taken and how often you take backup images. The storage required to hold the files also depends on the size of the database and table spaces. In addition, if you plan to migrate application group index data to table spaces, ADSM must maintain a backup image for each migrated table. Finally, ADSM can maintain multiple copies of each backup image. For example, for added protection, you may want ADSM to maintain two copies of each backup image.

OnDemand supports full database backups and incremental table space backups. To recover a database using incremental table space backups, you must create and maintain at least one full database backup image (taken before any changes are made to the database and prior to the first incremental table space backup).

You must configure ADSM to maintain a backup image as long as it is needed. For example, if you plan to create a full backup image of the database every week, we recommend that you configure ADSM to maintain two versions of the backup image. We recommend that you configure ADSM to maintain two copies of each version (full or table space backup). If you need to recover the database, you would always start with the latest version of the backup image. The latest version should be no more than one week old. If, for some reason, either copy of the latest version could not be used, you could use the prior version, which should be no more than two weeks old.

You may need to regularly initialize and load scratch storage volumes into the storage library where ADSM maintains the database backup image files. If ADSM determines there is not enough space available in the storage pool, it can request a mount for a scratch storage volume. However, the backup command cannot complete until the mount request is satisfied. If you operate in an unattended environment, this could have an adverse affect on system availability, especially when running an offline backup.

Note: Depending on the variables that we listed above, you may need to maintain hundreds of backup image files in ADSM-managed storage. We strongly encourage you to contact ADSM and database specialists to help plan your storage requirements. The following storage calculations may not accurately estimate the amount of storage needed to hold backup image files to support your system.

Full database backup calculation

Use the following calculation to estimate the optical storage space required to maintain full backup images of the database. The calculation uses the maximum size of the database, to allocate enough storage space to hold the largest backup image file required to recover the database.

$$\text{DB2BackupImageSpace} = (\text{MaxDBSize} * \text{compression ratio}) \\ * \text{CopiesMaintained} \\ * \text{VersionsMaintained}$$

For example, if the maximum size of the database is 5.4 GB and we want ADSM to maintain two versions of the backup image and two copies of each version, then the optical storage required to hold the backup image files is:

$$\begin{aligned} \text{DB2BackupImageSpace} &= (5.4 \text{ GB} * .33) \\ &\quad * 2 \\ &\quad * 2 = 7.2 \text{ GB} \end{aligned}$$

In the example, OnDemand requires approximately 7.2 GB of optical storage to hold the backup image files.

Table space backup

Use the following calculation to estimate the optical storage space required to maintain backup images of a table space. The calculation uses the maximum size of the table space, to allocate enough storage space to hold the largest backup image file required to recover the table space.

$$\begin{aligned} \text{TSBackupImageSpace} &= (\text{MaxTSSize} * \text{compression ratio}) \\ &\quad * \text{CopiesMaintained} \\ &\quad * \text{VersionsMaintained} \end{aligned}$$

For example, if the maximum size of the table space is 560 MB and we want ADSM to maintain two versions of the backup image and two copies of each version, then the optical storage required to hold the backup image files is:

$$\begin{aligned} \text{TSBackupImageSpace} &= (560 \text{ MB} * .33) \\ &\quad * 2 \\ &\quad * 2 = 740 \text{ MB} \end{aligned}$$

In the example, OnDemand requires approximately 740 MB of optical storage to hold the backup image files.

Database archived log file storage

Note: If you do not plan to use ADSM to maintain DB2 archived log files, do not allocate space for archive log files on archive media. Allocate disk storage space for the archived log files (refer to “Archive log file storage with DB2” on page 73).

The storage pool where ADSM maintains DB2 archived log files must contain enough storage to hold the files needed to recover your database. There are many variables you need to consider:

- For user data, how often you load reports
- For system data, how often you update system tables, such as users, groups, applications, application groups, and so forth. In addition, the System Log table space may be updated every time someone logs on or off the system, data is stored, queried, retrieved, and printed, and so forth.
- Whether you store user data in table spaces
- The size of the database or table spaces
- The frequency and type of backups taken

- How long you need to keep archived log files

If you regularly take full backup images of the database, we recommend that you allocate two times the space required for the active log files. However, you must allocate enough space to hold all of the archived log files created between full backups of the database. After a full database backup image is created, archived log files created prior to the backup are no longer needed and can be deleted. The following calculation can be used to estimate the amount of archived log file storage for a database:

$$\text{DB2ArchiveLogSpace} = (2 * \text{ActiveLogSpace}) \\ * \text{compression ratio}$$

Figure 17. Calculating Database Archived Log File Storage

The following example illustrates the archive log space required, when the space allocated for the active log files is 516 MB.

$$\text{DB2ArchiveLogSpace} = (2 * 516 \text{ MB}) \\ * .33 = 340 \text{ MB}$$

In the example, OnDemand requires approximately 340 MB of archive log file storage.

If you do not take full backup images of the database, we recommend that you keep archived log files indefinitely. Accordingly, you must carefully estimate the archive media storage requirement for the logs. For example, a single log file requires approximately 1.3 MB of (uncompressed) storage space. Depending on the variables that we listed above, you may need to maintain hundreds of log files in ADSM-managed storage. We strongly encourage you to contact ADSM and database specialists to help plan your storage requirements.

You may need to regularly initialize and load scratch storage volumes into the storage library where ADSM maintains the database archived log files. If ADSM determines there is not enough space available in the storage pool, it can request a mount for a scratch storage volume. However, the backup command cannot complete until the mount request is satisfied. If you operate in an unattended environment, this could have an adverse affect on system availability, especially when running an offline backup.

Migrated index storage space

After a period of time, OnDemand can migrate tables of application group index data from the database to archive media storage, such as magnetic tape or optical storage. This should be done only when the data is no longer

queried or queries to the data are very infrequent. For example, most queries occur in the first 24 months after a report is stored in OnDemand. After that time, queries are very infrequent and the index data for the report can be migrated from the database to archive media. Migration of index data is optional; you can choose to migrate index data for all, some, or none of the application groups on your system. In addition, you can determine the length of time before OnDemand migrates index data from the database to archive media.

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

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

Figure 18. Calculating Migrated Index Storage

For example, if the database size is 202 MB per month, OnDemand must maintain the data for 84 months, and data remains in the database for 24 months before being migrated, then the archive media storage required to hold the migrated index data is:

$$\begin{aligned} \text{ArchiveMediaDBSpace} &= (202 \text{ MB} * .33) \\ &* (84 - 24) = 4 \text{ GB} \end{aligned}$$

Number of storage volumes and libraries

Estimating the amount of archive media storage required to hold the reports you plan to store in OnDemand determines the number of storage volumes needed to maintain the reports.

In the previous example (refer to “Report storage space” on page 75), 244 GB of optical storage is required to hold the compressed report storage objects. Assuming the formatted capacity of a 5.25 inch storage volume is about 5.2 GB, approximate 48 storage volumes would be required to hold the data. An IBM 3995 Model C68 optical library supports 258 online storage volumes. Therefore, a single IBM 3995-C68 optical library can hold the entire seven years of data online.

You may need additional storage volumes and libraries to support other operational requirements. For example, if you use ADSM to manage DB2 backup image files and archived log files, we recommend that you dedicate a library for that purpose. If you need to maintain a backup copy of reports stored on archive media, we recommend that you configure a backup copy group in ADSM and store the backup copy in a different library than the primary copy.

It is possible to reduce the number of storage libraries by removing storage volumes from a library and placing them in offline storage. For example, you may find that you can remove a storage volume from a library one year after the last write or read access to the storage volume. ADSM provides commands that you can use to determine when a storage volume was last written to or read from and to dismount a storage volume from a library. However, before report data can be retrieved from an offline storage volume, an operator must locate the storage volume and mount it in the library.

You can also reduce the number of storage libraries by storing different types of reports in the same library. However, ADSM uses one policy domain to maintain all data stored in a library. The policy domain determines the length of time that ADSM maintains data in the library.

Storage sizing examples

The following examples show the estimated storage requirements for two reports:

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

The worksheet examples are separated into three parts:

- Report Profile
- Magnetic Disk Requirements (in bytes)
- Archive Media Storage Requirements (in bytes)

Storage for report containing logical items

Table 7. Database Columns

Column Number	Name	Index or Filter	Bytes
1	Report Date	Index	2
2	Account Number	Index	12
3	Balance	Filter	8
4	Customer Name	Filter	24

Table 8. Example report profile for report containing logical items

Report Profile	
Report Characteristic	Report Estimate
Volume of data per month (bytes)	8,000,000,000
Average statement size (bytes)	8,000
Number of statements per month	1,000,000

Table 8. Example report profile for report containing logical items (continued)

Number of cycles per month	20
Largest cycle data size (bytes)	400,000,000
Largest single report file size (bytes)	400,000,000
Largest cycle - number of statements	50,000
Number DB columns from Table 7	4
Life of data (days)	2555
Number of days to cache data	90
Number of days to keep index on magnetic	730
Compression percentage (ratio)	0.33 (3:1)
ACIF on MVS, OS/390, or OnDemand server	OnDemand Server

Table 9. Example management disk requirements for report containing logical items

Magnetic Disk Requirements (in bytes)	
Storage Component	Storage Requirement
Base system software	2,000,000,000
Data download storage	480,000,000
Temp space for data index/load	600,000,000
Cache storage	8,712,000,000
OnDemand database	5,400,000,000
Database log files	516,000,000
ADSM database	47,040,000
Temp space for server print	500,000,000
Temp space for import data	500,000,000
Total Magnetic Storage Required	18,755,040,000 (bytes)

Table 10. Example optical storage requirements for report containing logical items

Archive Media Storage Requirements (in bytes)	
Report data	244,000,000,000
Migrated index data	4,000,000,000
Database log file storage	0
Database backup image file storage	0
Total Archive Media Storage Required	248,000,000,000 (bytes)

Storage for report containing transaction data

Table 11. Database Columns

Column Number	Name	Index or Filter	Bytes
1	Report Date	Index	2
2	Beginning Invoice Number	Index	10
3	Ending Invoice Number	Index	10
4	Page Number	Filter	2

Table 12. Example report profile for storage for report containing logical items

Report Profile	
Report Characteristic	Report Estimate
Volume of data per month (bytes)	5,000,000,000
Average page size (bytes)	5,000
Number of pages per month	1,000,000
Number of cycles per month	20
Largest Cycle - data size (bytes)	200,000,000
Largest Cycle - number of pages	50,000
Largest single report file size (bytes)	200,000,000
Group of indexed pages	100
Number DB columns from Table 11	4
Life of data (days)	730
Number of days to cache data	0
Number of days to keep index on magnetic	730
Compression percentage (ratio)	0.25 (4:1)
ACIF on MVS, OS/390, or OnDemand Server	OnDemand Server

Table 13. Example magnetic disk requirements for storage for report containing logical items

Magnetic Disk Requirements (in bytes)	
Storage Component	Storage Requirement
Base System Software	2,000,000,000
Data download storage	240,000,000
Temp space for data index/load	300,000,000
Cache storage	0
OnDemand database	58,368,000
Database log files	18,000,000

Table 13. Example magnetic disk requirements for storage for report containing logical items (continued)

ADSM database	5,250,000
Temp space for server print	500,000,000
Temp space for import data	0
Total Magnetic Storage Required	3,121,618,000

Table 14. Example archive media storage requirements for storage for report containing logical items

Archive Media Storage Requirements (in bytes)	
Report data	20,350,000,000
Migrated index data	0
Database log file storage	0
Database backup image file storage	0
Total Archive Media Storage Required	20,350,000,000 (bytes)

Storage sizing worksheets

Make a copy of Table 15 and the worksheet on page 85 (reports that contain logical items) or page 86 (reports that contain sorted transaction data) for each report that you want to store in OnDemand. Complete the table and worksheet to calculate the storage requirements for the report. Refer to “Calculating magnetic storage requirements” on page 65 and “Calculating archive media requirements” on page 75 for formulas used to calculate storage requirements.

Table 15. Database Columns Worksheet

Column Number	Name	Index or Filter	Bytes
1			
2			
3			
4			
5			
6			
7			
8			
9			
10			
11			

Table 15. Database Columns Worksheet (continued)

Column Number	Name	Index or Filter	Bytes
12			
13			
14			
15			
16			

Storage sizing worksheet for reports containing logical items

Table 16. Blank report profile storage sizing worksheet for reports containing logical items

Report Profile	
Report Characteristic	Report Estimate
Volume of data per month (bytes)	
Average item size (bytes)	
Number of items per month	
Number of cycles per month	
Largest cycle data size (bytes)	
Largest single report file size (bytes)	
Largest cycle - number of items	
Number DB columns from Table 15	
Life of data in days	
Number of days to cache data	
Number of days to keep index on magnetic	
Compression ratio	
ACIF on MVS, OS/390, or OnDemand Server	

Table 17. Blank magnetic disk requirements worksheet for reports containing logical items

Magnetic Disk Requirements (in bytes)	
Storage Component	Storage Requirement
Base system software	2,000,000,000
Note: The base system software requirement of 2 GB is per server	
Data download storage	
Temp space for data index/load	
Magnetic cache storage	
OnDemand database	

Table 17. Blank magnetic disk requirements worksheet for reports containing logical items (continued)

Database log files	
ADSM database	
Temp space for server print	
Temp space for import data	
Total Magnetic Storage Required	

Table 18. Blank archive media storage requirements worksheet for reports containing logical items

Archive Media Storage Requirements (in bytes)	
Storage Component	Storage Requirement
Report data	
Migrated index data	
Database log file storage	
Database backup image file storage	
Total Archive Media Storage Required	

Storage sizing worksheet for reports containing transaction data

Table 19. Blank report profile worksheet for reports containing transaction data

Report Profile	
Report Characteristic	Report Estimate
Volume of data per month (bytes)	
Average page size (bytes)	
Number of pages per month	
Number of cycles per month	
Size of largest cycle (bytes)	
Size of largest cycle in pages	
Size of largest single report file (bytes)	
Pages in a group	
Number DB columns from Table 15	
Life of data in days	
Number of days to cache data	
Number of days to keep index on magnetic	
Compression ratio	
ACIF on MVS, OS/390, or OnDemand Server	

Table 20. Blank magnetic disk requirements worksheet for reports containing transaction data

Magnetic Disk Requirements (in bytes)	
Storage Component	Storage Requirement
Base system software	2,000,000,000
Note: The base system software requirement of 2 GB is per server	
Data download storage	
Temp space for data index/load	
Cache storage	
OnDemand database	
Database log files	
ADSM database	
Temp space for server print	
Temp space for import data	
Total Magnetic Storage Required	

Table 21. Blank archive media storage requirements worksheet for reports containing transaction data

Archive Media Storage Requirements (in bytes)	
Storage Component	Storage Requirement
Report data	
Migrated index data	
Database log file storage	
Database backup image file storage	
Total Archive Media Storage Required	

Chapter 9. Planning a storage subsystem

Overview

Before you define a report to OnDemand and archive data on the server, it is important to estimate the data storage requirements of your applications. “Chapter 8. Estimating storage requirements” on page 61 contains information that can help you estimate storage requirements for the reports that you plan to store in OnDemand.

After you estimate the storage requirements for your applications, you can configure the magnetic storage devices installed on the OnDemand Server that support the storage required by OnDemand programs, data download and indexing, databases, log files, and cache storage.

Depending on the type of reports that you define and the number and size of the files that you plan to archive in OnDemand, you may need to configure the magnetic storage devices into volume groups and logical volumes and assign logical volumes to specific physical volumes. Configuring storage devices in this way allows you to extend volume groups and logical volumes as your storage needs grow and configure volume groups for high availability and maximum performance.

If you are planning for UNIX servers, see your operating system documentation for details about managing devices and defining logical volumes and file systems. If you are not familiar with logical volumes and file systems, we recommend that you review information in your operating system documentation before you continue.

If you are planning for NT servers, the *Microsoft Windows NT Server: Concepts and Planning (Version 4.0)* and *Microsoft Windows NT Server: Resource Guide* publications provide important information about file systems and describe how to manage devices and file systems. If you are not familiar with file systems, we recommend that you review the material in your NT documentation before you continue.

Setting up storage devices for a UNIX server

The examples that follow provide a suggested configuration for magnetic storage devices to hold software programs and control files, data downloaded from an MVS or OS/390 system, index data on the OnDemand server, storage for the OnDemand (DB2) database and log files, storage for the ADSM database, and cache storage for reports and resources.

We recommend that you organize storage volumes into volume groups. A volume group can support up to 32 physical volumes of varying sizes and types. You can define up to 255 volume groups per system. Regardless of the number of physical devices configured on the OnDemand server or the capacity of the devices, we recommend that you organize the available storage as described in the sections that follow. We strongly encourage you to adopt the suggested convention for naming the volume groups, logical volumes, file systems, directories, and files.

If you are configuring an OnDemand system with large magnetic storage requirements, we suggest that you organize available physical devices into five volume groups, supporting the following components:

- Software programs, control files, resources, and temporary storage.
- Data transmitted from other systems, data indexing, and data loading.
- OnDemand database and relational database log files on the library server.
- ADSM database and recovery log.
- Cache storage.

If you are configuring an OnDemand system with low to medium magnetic storage requirements that supports few users, we recommend that you organize available physical devices into two volume groups, supporting the following components:

- The *root* volume group for software programs, control files, resources, and temporary work space.
- The *ars* volume group for data transmitted from other systems, data indexing, data loading, the OnDemand database, DB2 log files, the ADSM database and recovery log, and cache storage.

Table 22 shows an example of configuring the *ars* volume group with the recommended names to support a small to medium OnDemand system.

Table 22. The ars Volume Group

Volume Group	Logical Volume	File System	Physical Volumes
arsvg	aciflv1	/arsacif/acif1	any
arsvg	cachelv1	/arscache/cache1	any

Table 22. The ars Volume Group (continued)

Volume Group	Logical Volume	File System	Physical Volumes
arsvg	dblv	/arsdb	hdisk1
arsvg	primloglv	/arsdb_primarylog	hdisk2
arsvg	archloglv	/arsdb_archivelog	hdisk3
arsvg	dsmdblv	none	hdisk3
arsvg	dsmloglv	none	hdisk2

Table 23 shows an example of configuring volume groups, logical volumes, and file systems with the recommended names to support a large OnDemand system.

Table 23. Logical Volumes and File Systems

Volume Group	Logical Volume	File System	Physical Volumes
acifvg	aciflv1	/arsacif/acif1	hdisk1
acifvg	aciflv2	/arsacif/acif2	hdisk2
cachevg	cachelv1	/arscache/cache1	hdisk3, hdisk4
cachevg	cachelv2	/arscache/cache2	hdisk5, hdisk6
cachevg	cachelv3	/arscache/cache3	hdisk7
db2vg	dblv	/arsdb	hdisk8, hdisk9
db2vg	primloglv	/arsdb_primarylog	hdisk10
db2vg	archloglv	/arsdb_archivelog	hdisk11
rootvg	arstmplv	/arstmp	hdisk12
adsmvg	dsmdblv	none	hdisk13
adsmvg	dsmloglv	none	hdisk14

Setting up storage devices for Windows NT servers

The examples that follow provide a suggested configuration for magnetic storage devices that support software programs and control files, downloading data from an MVS or OS/390 system, indexing data on the OnDemand server, storage for the database and log files, and cache storage for reports and resources.

Recommended Storage devices for small to medium systems

For OnDemand systems with small to medium magnetic storage requirements, we recommend that you use at least four 4.5 GB physical drives and organize the physical drives as shown in Table 24 on page 92. (You may

need more or fewer drives, depending on how much data you need to maintain in cache storage.)

Table 24. Storage Devices for Windows NT Servers

Physical Disk	Partition	File System(s)	Purpose
Disk1	C:	\WINNT	Windows NT system files
Disk1	D:	\Program Files	OnDemand for Windows NT and other applications
Disk2	E:	\arsdb, \arsdbpri, \arsdbarc	OnDemand database and log files
Disk3	F:	\arscache1	Cache storage
Disk4	G:	\arsacif1, \arstmp	Data indexing and loading; temporary storage

Recommended Storage devices for large systems

For OnDemand systems with large magnetic storage requirements, we recommend that you use at least ten 4.5 GB physical drives and organize the physical drives as shown in Table 25. (You may need more or fewer drives, depending on the size of the database and how much data you need to maintain in cache storage.)

Table 25. Storage Devices for Windows NT Servers

Physical Disk	Partition	File System(s)	Purpose
Disk1	C:	\WINNT	Windows NT system files
Disk1	D:	\Program Files	OnDemand for Windows NT and other applications
Disk2	E:	\arsdb	OnDemand database
Disk3	F:	\arsdbpri, \arsdbarc	Database log files
Disk4	G:	\arscache1	Cache storage
Disk5	H:	\arscache2	Cache storage
Disk6	I:	\arscache3	Cache storage
Disk7	J:	\arscache4	Cache storage
Disk8	K:	\arsacif1	Data download and indexing

Table 25. Storage Devices for Windows NT Servers (continued)

Physical Disk	Partition	File System(s)	Purpose
Disk9	L:	\arsload1	Data loading
Disk10	M:	\arstmp	Temporary storage

Using RAID storage with OnDemand

Overview

This chapter provides information about using a RAID storage subsystem with OnDemand.

RAID stands for Redundant Array of Inexpensive Disks and provides a method of classifying the different methods of using multiple disks to increase availability. With RAID, multiple physical disks appear to the OnDemand server as one logical disk. RAID carries out the concept of data striping by spreading data over multiple disks; a single file is segmented and stored on multiple disks. RAID carries out the concept of data mirroring by duplicating data from one disk to a second disk; a single file is stored twice, on two different disks. A failed disk still allows users to access data on the array, and a replacement disk or online spare can be recreated while the array is in use. Table 26 provides an overview of RAID implementations.

Table 26. RAID Implementations

RAID Level	Description	Protection	Performance
RAID 0	Data striping on multiple disk drives.	Poor; single disk failure.	Best; read and write requests can be met by any disk.
RAID 1	Disk mirroring.	Good; any disk can fail and data is still accessible.	Good; read request can be met by any disk.
RAID 3	Disk striping with parity disk, using interleaved bytes.	Good; if any disk fails, data can be accessed by using information from other disks and parity disk.	Good for large data transfers.
RAID 4	Disk striping with parity disk, using interleaved sectors.	Good; if any disk fails, data can be accessed by using information from other disks and parity disk.	Good for large data transfers.
RAID 5	Disk striping with distributed parity data.	Good; if any disk fails, data can be accessed by using information from other disks and parity information.	Good for small block sizes.

Table 26. RAID Implementations (continued)

RAID Level	Description	Protection	Performance
RAID 5 Orthogonal	Disk striping with distributed parity data, using dual controllers.	Best; if any disk fails, data can be accessed by using information from other disks and parity information, with additional protection from any single disk controller failure.	Good for small block sizes; improved performance because of use of dual controllers to read and write data.

Important: For most OnDemand systems, we recommend that you use the RAID 1 or RAID 5 implementations.

RAID storage and the document cache

A typical use of a RAID storage subsystem in the OnDemand environment is to support the document cache. Orthogonal RAID 5 (redundant disk controllers) provides excellent protection from a single disk or controller failure. Disk striping with distributed parity data allows the OnDemand server to remain available if a single disk fails. Redundant disk controllers provide excellent availability, enabling users to continue to access reports if a controller fails. We strongly recommend RAID storage for your object servers that do not use ADSM to maintain copies of reports on archive media. In addition, if you do not use ADSM to maintain copies of reports on archive media, we recommend that you use ADSM to backup your cache file systems. Without a copy of reports on archive media or an up-to-date backup of cache file systems, your system is exposed to loss of data that may be difficult or impossible to recreate.

RAID storage and the OnDemand database

It is possible to implement the entire OnDemand storage subsystem (as depicted in Table 23 on page 91) using one or more RAID storage subsystems. However, depending on the hardware and RAID implementation, the server may not achieve the same level of database performance as a library server using non-arrayed disk storage. That is, when loading reports or when many concurrent users query the database, you may not experience the same level of performance on a server using arrayed storage as you would on a server with non-arrayed storage. However, the availability benefits provided by RAID storage subsystems typically outweigh any performance degradation.

Configuring RAID storage

This topic provides information about configuring an IBM 7135 RAIDiant Array, a fault tolerant disk solution for the IBM RS/6000 family of workstations. For information about configuring other types of RAID devices, see the product information.

The IBM 7135 RAIDiant Array Model 210 can be configured with six banks of five disks and up to 135 GB of storage capacity. With RAID 5, the drives are configured into six logical units of five drives. Each logical unit appears as a physical volume to the operating system. You can assign one or more of the physical volumes to a volume group, for example, the *cachevg* volume group. You can create logical volumes, for example, the *cachelv1* logical volume, and file systems, such as the */arscache/cache1* file system. A fully populated 7135-210 can be configured into two cache file systems by assigning the disks to two volume groups, creating one logical volume for each volume group, and creating one file system for each logical volume.

The IBM 7133 Serial Disk System

The IBM 7133 Serial Disk System Advanced Models D40 and T40 provide high performance and enhanced availability for I/O intensive applications on UNIX and Windows NT servers. The 7133 Advanced Models use Serial Storage Architecture (SSA), the highest performing IBM disk subsystem, supporting up to 3,000 I/O operations per second per adapter and a data transfer rate of up to 160 MB per second. SSA includes a loop design so that a single adapter failure will not cause loss of access to data. If there is a disk failure, the hot swappable disks can be removed without loss of communication between the adapter and the other disks on the loop.

Each 7133 Advanced Model can be configured with up to 16 disks and 291 GB of storage. The 7133 Advanced Models can be populated with 18.2 GB, 9.1 GB, and 4.5 GB IBM Ultrastar disk drives with auto-docking capabilities. Drive capacities can be intermixed, providing the flexibility to build storage environments ranging in capacity from gigabytes to terabytes.

The 7133 Advanced Models can be attached to IBM RS/6000, HP, and Sun workstations and selected Intel-based servers. The 7133 Advanced Models support multi-host attachment, with one or more servers sharing the disk subsystem, so that if one server fails, the others can continue to process data.

The IBM 3527 SSA Storage Subsystem for PC Servers

The IBM 3527 SSA Storage Subsystem for PC Servers uses Serial Storage Architecture technology to deliver outstanding performance, highly scalable

capacity, and enhanced availability. The 3527 SSA Storage Subsystem can sustain data rates as high as 35 MB per second in RAID arrays, to support the most demanding storage applications.

The 3527 subsystem mini-tower design features five optional hot-swap drive bays that can hold any combination of 4.5 GB or 9.1 GB SSA disk drives. This results in a total storage capacity of up to 45.5 GB per enclosure. As many as 48 drives can be attached to each of two loops on a single adapter. This provides a maximum of 98 drives with 873 GB of storage on PC Server Models 320, 325, 330, and 520. For larger storage applications, PC Server Models 704 and 720 support two adapters per server, for a maximum of 192 drives totaling 1.75 TB of storage.

The 3527 subsystem uses SSA's bi-directional loop architecture to ensure that data access is never interrupted by a broken cable or failure of a single disk drive. All drives are self-configuring with optional hot-swap capability, for each service or replacement without shutting down the server. The SSA RAID Adapter for PC Servers helps protect against loss of data access using RAID 1 or 5. The system is quickly returned to normal operation once the failed drive has been replaced or data path restored.

Chapter 10. Planning processor requirements and configurations

AIX processor and memory requirements

You can use the following general guidelines to help determine the type of RS/6000 required to support your OnDemand system.

RS/6000 processor

The typical OnDemand production archive and retrieval system requires a RS/6000 Model J40 server.

For an OnDemand production system that supports a large number of concurrently logged on users, a large amount of magnetic storage for the OnDemand database and cache storage, and optical storage for documents, one or more RS/6000 Enterprise Server Model R50 systems or a Scalable POWERparallel (SP) large scale server may be required.

A system used to test OnDemand or an OnDemand production system that supports a small number of concurrently logged on users and little magnetic or optical storage can usually be supported by a RS/6000 Workgroup Server Model F40.

RS/6000 memory

You should plan to configure each RS/6000 workstation with at least 128 MB of memory to support the basic OnDemand programs. You should plan to add 2 MB of memory for each concurrently logged on OnDemand user.

HP-UX processor and memory requirements

You can use the following general guidelines to help determine the type of HP-UX system required to support your OnDemand system.

HP-UX processor

The typical OnDemand production archive and retrieval system requires an HP-UX 9000 server.

For an OnDemand production system that supports a large number of concurrently logged on users, a large amount of magnetic storage for the

OnDemand database and cache storage, and optical storage for documents, one or more HP-UX 9000 Series 800 Business servers may be required.

A system used to test OnDemand or an OnDemand production system that supports a small number of concurrently logged on users and little magnetic storage can usually be supported by an HP-UX 9000 server.

HP-UX memory

You should plan to configure each HP-UX workstation with at least 128 MB of memory to support the basic OnDemand programs. You should plan to add 2 MB of memory for each concurrently logged on OnDemand user.

Sun Solaris processor and memory requirements

You can use the following general guidelines to help determine the type of Sun Solaris system required to support your OnDemand system.

Solaris processor

The typical OnDemand production archive and retrieval system requires a Sun Enterprise 3500 server.

For an OnDemand production system that supports a large number of concurrently logged on users, a large amount of magnetic storage for the OnDemand database and cache storage, and optical storage for documents, one or more Sun Enterprise 5500 servers or a Sun Enterprise 10000 server may be required.

A system used to test OnDemand or an OnDemand production system that supports a small number of concurrently logged on users and little magnetic storage can usually be supported by a Sun Enterprise Ultra 10S server.

Solaris memory

You should plan to configure each Solaris workstation with at least 128 MB of memory to support the basic OnDemand programs. You should plan to add 2 MB of memory for each concurrently logged on OnDemand user.

NT Processor and memory requirements

You can use the following general guidelines to help determine the type of NT processor required to support your OnDemand system.

NT processor

The typical OnDemand production archive and retrieval system requires an IBM compatible PC with multiple Intel Pentium II (266 MHz) processors.

For an OnDemand production system that supports a large number of concurrently logged on users, a large amount of magnetic storage for the OnDemand database and cache storage, and optical storage for documents, two or more IBM compatible PCs with multiple Intel Pentium III (450 MHz) processors may be required.

A system used to test OnDemand or an OnDemand production system that supports a small number of concurrently logged on users and little magnetic or optical storage can usually be supported by an IBM compatible PC with an Intel Pentium Pro (200 MHz) processor.

System memory

You should plan to configure each Windows NT server with at least 128 MB of memory to support the basic OnDemand programs. You should plan to add 2 MB of memory for each concurrently logged on OnDemand user. You should also consider using ECC (Error Correction Code) memory for increased system reliability.

Server hardware configurations

You can use the following general guidelines to help determine the server hardware required to support your OnDemand system. The guidelines are based on the maximum number of users logged on to the server at any one time and the average number of pages stored in OnDemand each month. The configurations include enough magnetic storage to cache reports for three months and provide a recommended backup device.

When you plan your OnDemand server configuration, the business, legal, and operational requirements of your organization may result in a system configuration much different than those in this book. For example, if you need to maintain 24 by 7 availability of systems and data, you will likely require more than one physical server and significantly more archive media storage, devices, and controllers. If you plan to use ADSM to manage relational database log files and backup image files, we recommend that you plan for an automated tape library or a dedicated optical library. In addition, the configurations in this book **do not** include archive media storage for mirroring of databases or maintaining a backup copy of reports. Besides understanding your organization's business and operational needs, we recommend that you use the information provided in "Chapter 8. Estimating storage requirements" on page 61

on page 61 and “Chapter 14. Planning for backup and recovery” on page 123 to determine the hardware needed to support your OnDemand system.

Small server configuration

The following OnDemand server hardware can generally support up to 15 concurrently logged on users. The amount of data archived on the system in one year, at a rate of 200,000 pages a month, requires approximately 6 GB of archive media storage. The configuration includes sufficient cache storage to cache report files for three months.

Table 27. Small Server Configuration

Memory	Cache Storage	Archive Media Storage	Backup Device
128 MB	8 GB	One ATL or IBM 3995-C60	One 8 mm tape drive

Medium server configuration

The following OnDemand server hardware can generally support up to 100 concurrently logged on users. The amount of data archived on the system in one year, at a rate of 4,000,000 pages a month, requires approximately 120 GB of archive media storage. The configuration includes sufficient cache storage to cache report files for three months.

Table 28. Medium Server Configuration

Memory	Cache Storage	Archive Media Storage	Backup Device
256 MB	54 GB	One ATL or IBM 3995-C64	Second ATL or optical library

Large server configuration

The following OnDemand server hardware can generally support up to 500 concurrently logged on users. The amount of data archived on the system in one year, at a rate of 20,000,000 pages a month, requires approximately 600 GB of optical storage space. The configuration includes sufficient magnetic storage to cache report files for three months.

Table 29. Large Server Configuration

Memory	Cache Storage	Archive Media Storage	Backup Device
1024 MB	234 GB	One ATL or IBM 3995-C68	Second ATL or optical library

Server software configurations

Library/object server

Figure 19 shows the standard OnDemand library/object server environment.

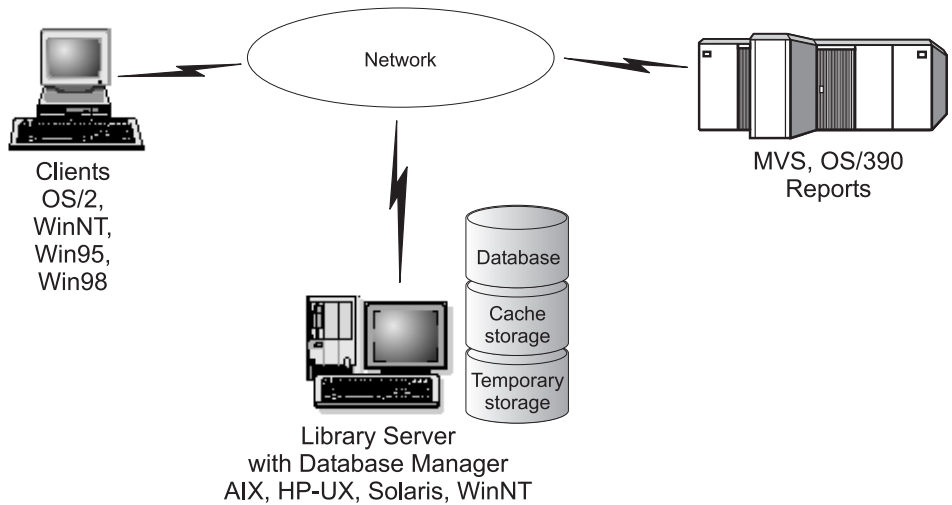


Figure 19. Library/Object Server

The standard library/object server includes a storage manager to maintain reports on cache storage volumes and programs required to index reports and load index data into the database. You can stage reports on temporary storage volumes for the data indexing and loading programs. This environment is ideal for those who do not require backup copies of files on archive media and need to run OnDemand on a single system.

Table 30 lists the software requirements for the standard library/object server.

Table 30. Software for Library/Object Server

Function	Req/Opt	Notes
OnDemand base	Required	Base OnDemand functions, such as enhanced ACIF. Manage files on cache storage volumes.
Database manager	Required	Database engine and administration.
OnDemand client	Required	OS/2, Windows client programs.
Download	Optional	Transmit files from MVS and OS/390.
PSF	Optional	Server print and FAX.

Library/object server with ADSM

Figure 20 shows the library/object server with ADSM environment.




Figure 20. Library/Object Server with ADSM

ADSM is a companion product of OnDemand that manages files stored on archive media, such as optical and tape storage volumes. The OnDemand library server includes a storage manager to maintain reports on cache storage volumes and includes programs required to index reports and load index data into the database. You can stage reports on temporary storage volumes for the data indexing and loading programs. This environment is ideal for those who require backup copies of files on archive media attached to the library server.

Table 31 on page 103 lists the software requirements for an OnDemand library server with ADSM.

Table 31. Software for Library/Object Server with ADSM

Function	Req/Opt	Notes
OnDemand base	Required	Base OnDemand functions, such as enhanced ACIF. Manage files on cache storage volumes. Interface to ADSM to store files on and retrieve files from archive media.
Database manager	Required	Database engine and administration.
OnDemand client	Required	OS/2, Windows client programs.
ADSM	Required	Manages files on optical and tape storage volumes.
Download	Optional	Transmit files from MVS and OS/390.
PSF	Optional	Server print and FAX.

Library server and object server

Figure 21 shows the distributed OnDemand library server and object server environment.




Figure 21. Library Server and Object Server

OnDemand supports storing data on multiple servers. In this environment, end-users submit queries to the library server; OnDemand retrieves items from the server where the data is stored. The OnDemand object server includes a storage manager to maintain reports on cache storage volumes attached to the object server. The OnDemand library server includes a storage manager to maintain reports on cache storage volumes attached to the library server and programs required to index reports and load index data into the database. You can stage reports on temporary storage volumes for the data indexing and loading programs. This environment is ideal for those who need to cache files on more than one server. The servers can reside on nodes in one physical machine, such as an SP processor, or separate systems in different physical locations.

Table 32 lists the software requirements for the workstation where you install the OnDemand library server. Table 33 lists the software requirements for the workstation where you install the OnDemand object server.

Table 32. Software for Library Server

Function	Req/Opt	Notes
OnDemand base	Required	Base OnDemand functions, such as enhanced ACIF. Maintain files on cache storage volumes attached to the library server.
Database manager	Required	Database engine and administration.
OnDemand client	Required	OS/2, Windows client programs.
Download	Optional	Transmit files from MVS and OS/390.
PSF	Optional	Server print and FAX.

Table 33. Software for Object Server

Function	Req/Opt	Notes
OnDemand base	Required	Maintain files on cache storage volumes attached to the object server.

Library server and object server with ADSM

Figure 22 shows the distributed library server and object server with ADSM environment.

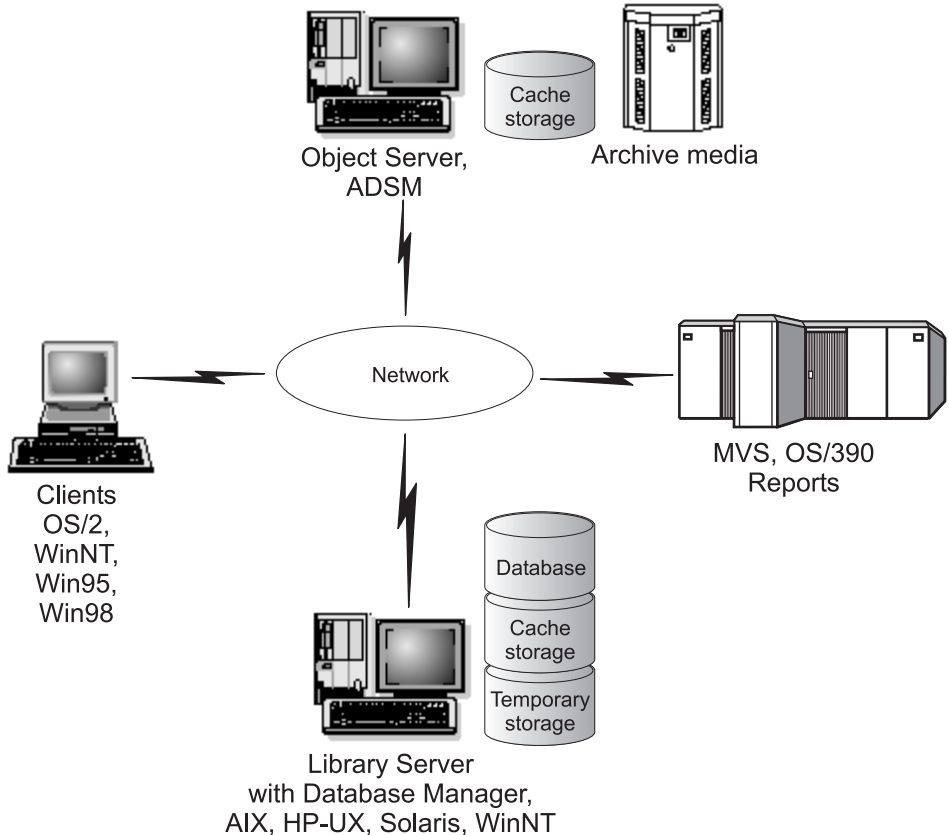


Figure 22. Library Server and Object Server with ADSM

OnDemand supports storing data on multiple servers. In this environment, end-users submit queries to the library server; OnDemand retrieves items from the server where the data is stored. ADSM can be installed on the object server to manage optical and tape storage devices attached to the object server. ADSM is a companion product of OnDemand that manages files stored on archive media, such as optical and tape storage volumes. The OnDemand object server includes a storage manager to maintain reports on cache storage volumes attached to the object server. The OnDemand library server includes a storage manager to maintain reports on cache storage volumes attached to the library server and programs required to index reports and load index data into the database. You can stage reports on temporary storage volumes for the

data indexing and loading programs. This environment is ideal for those who need to cache reports on different servers and require backup copies of files on archive media attached to a different server than the OnDemand library server. The servers can reside on nodes in one physical machine, such as an SP processor, or separate systems in different physical locations.

Table 34 lists the software requirements for the workstation where you install the OnDemand library server. Table 35 lists the software requirements for the workstation where you install the OnDemand object server with ADSM.

Table 34. Software for Library Server

Function	Req/Opt	Notes
OnDemand base	Required	Base OnDemand functions, such as enhanced ACIF. Maintain files on cache storage volumes attached to the library server.
Database manager	Required	Database engine; database administration.
OnDemand client	Required	OS/2, Windows client programs.
Download	Optional	Transmit files from MVS and OS/390.
PSF	Optional	Server print and FAX.

Table 35. Software for Object Server

Function	Req/Opt	Notes
OnDemand base	Required	Maintain files on cache storage volumes attached to the object server. Interface to ADSM to store files on and retrieve files from archive media.
ADSM	Required	Manage files on optical and tape storage volumes attached to the object server.

Library server with ADSM and object server with ADSM

Figure 23 shows a distributed library server and object server environment, with ADSM on both servers.

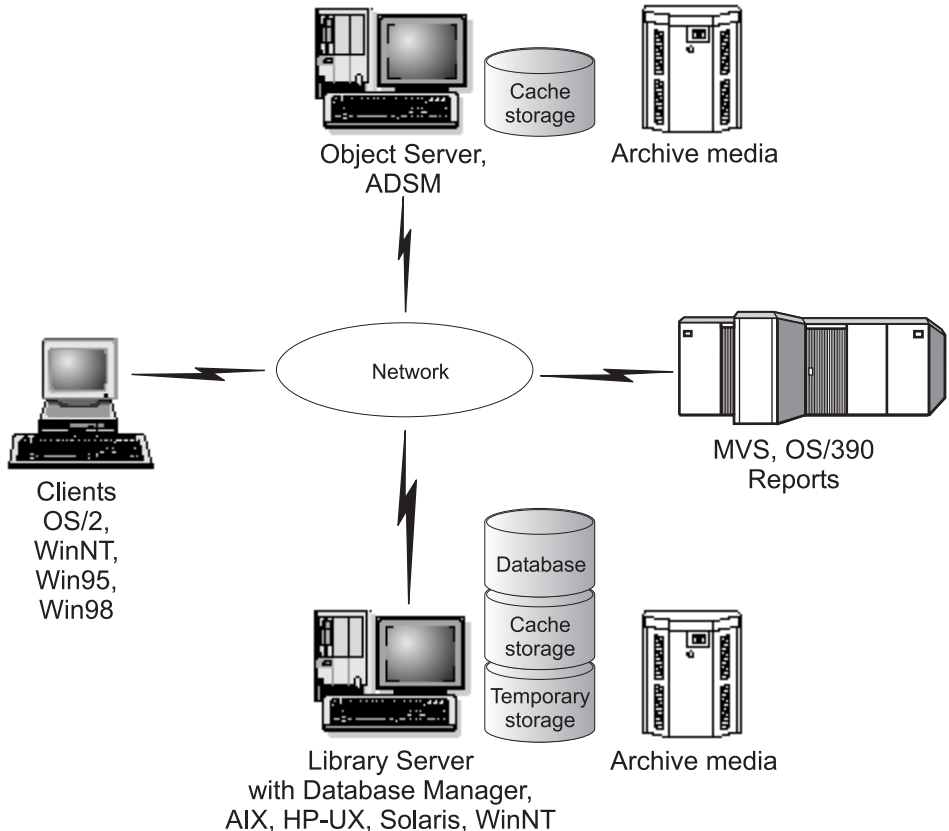


Figure 23. Library Server with ADSM and Object Server with ADSM

OnDemand object servers support storing data on multiple servers. In this environment, end-users submit queries to the library server; OnDemand retrieves items from the server where the data is stored. ADSM is a companion product of OnDemand that manages files stored on archive media, such as optical and tape storage volumes. OnDemand programs on the object server maintain reports on cache storage volumes attached to the object server. The OnDemand library server includes programs that maintain reports on cache storage volumes attached to the library server and index reports and load index data into the database. You can stage reports on temporary storage volumes for the data indexing and loading programs. This environment is ideal for those who need to cache files on more than one server and require backup copies of files on archive media attached to more than one server. The

servers can reside on nodes in one physical machine, such as an SP processor, or separate systems in different physical locations.

Table 36 lists the software requirements for the workstation where you install the OnDemand library server with ADSM. Table 37 lists the software requirements for the workstation where you install the OnDemand object server with ADSM.

Table 36. Software for Library Server with ADSM

Function	Req/Opt	Notes
OnDemand base	Required	Base OnDemand functions, such as enhanced ACIF. Maintain files on cache storage volumes attached to the library server. Interface to ADSM to store files on and retrieve files from archive media.
Database manager	Required	Database engine; database administration.
OnDemand client	Required	OS/2, Windows client programs.
ADSM	Required	Manage files on optical and tape storage volumes attached to the library server.
Download	Optional	Transmit files from MVS and OS/390.
PSF	Optional	Server print and FAX.

Table 37. Software for Object Server with ADSM

Function	Req/Opt	Notes
OnDemand base	Required	Maintain files on cache storage volumes attached to the object server. Interface to ADSM to store files on and retrieve files from archive media.
ADSM	Required	Manage files on optical and tape storage volumes attached to the object server.

UNIX license information

The license information to use OnDemand is available with the product.

The OnDemand base product allows one library server, one object server, and one concurrent user.

The number of concurrent users is defined as the maximum number of client workstations connected to the library server at a given point in time. If you need to support more than one user, you must acquire entitlements for additional users. Entitlements for additional users do not require Product License Keys. It is up to the purchaser of the product to ensure compliance with the license agreement with regards to the number of licensed users. The Product License Key that comes with the base product unlocks the product for use and does not restrict the number of concurrent users.

If you need to support more than one library server, you must purchase an additional base product license for each additional library server.

If you need to support more than one object server, you must acquire entitlements for additional object servers.

NT License information

The license information to use OnDemand is available with the product.

The OnDemand base product allows one library server, one object server, and one concurrent user.

The number of concurrent users is defined as the maximum number of client workstations connected to the library server at a given point in time. If you need to support more than one user, you must acquire entitlements for additional users. It is up to the purchaser of the product to ensure compliance with the license agreement with regards to the number of licensed users.

If you need to support more than one library server, you must purchase an additional base product license for each additional library server.

If you need to support more than one object server, you must acquire entitlements for additional object servers.

Chapter 11. Assessing network and printer requirements

Network requirements

The network requirements of an OnDemand system can be grouped into two categories:

- Requirements for downloading reports from other systems to OnDemand servers.
- Requirements for client programs accessing OnDemand servers.

Data download

There are three methods to transfer reports from an MVS or OS/390 system to an OnDemand server:

- SNA protocols
- TCP/IP protocols
- Download

There are advantages and disadvantages to each approach.

SNA protocols

You can use third-party RJE packages to provide a seamless download process from the JES spool to queues on an OnDemand server. This option works well for low volume applications or line of business applications with less critical availability requirements.

TCP/IP

TCP/IP provides a wide variety of options for data download. TCP/IP is a powerful and flexible industry-standard way of connecting multiple systems. There are many ways you can customize TCP/IP to fit the specific needs of your system.

TCP/IP provides very fast data download. Up to one gigabyte of data can be transferred per hour over a 16 megabyte Token Ring gateway (IBM 3172 controller). Even faster speeds can be obtained by directly attaching the host system and the OnDemand server over an ESCON channel.

Several TCP/IP protocols can be used to move print data between the host system and an OnDemand server:

- FTP

FTP permits data transfer between dissimilar host systems. FTP is designed to be used by applications, creating and removing datasets and directories and transferring multiple files in a single request. FTP provides security by passing user and account passwords between hosts.

- LPR/LPD

LPR is the TCP/IP queueing facility. LPR places print data streams on print queues by copying files into spooling datasets. LPD is the printer daemon that controls remote printing between print servers and other hosts on the network. With OnDemand, LPD can be used to transfer files from a host or LAN system to an OnDemand server.

- NFS

NFS allows host systems on the TCP/IP network to share datasets and disks. The shared datasets and disks appear to each remote host system as just another local dataset or disk.

Download

Download is an optional feature that provide a seamless, high-speed download for report files from the JES spool to the OnDemand server, using TCP/IP communications protocols. After storing the report file in a file system on the OnDemand server, OnDemand can automatically create index data and load the report file into an application group. The download programs use a JES class or printer destination to automatically transmit report files to OnDemand. See *PSF for MVS: MVS Download Guide* or *PSF for OS/390: Download for OS/390* for information about installing, configuring, and operating Download on your system.

OnDemand client access to OnDemand servers

OnDemand client programs must use TCP/IP sockets protocols to communicate with an OnDemand server. TCP/IP sockets support a wide range of hardware configurations, including Ethernet, Token Ring, and WAN and can even be used over telephone lines with modems using the Serial Line Internet Protocol (SLIP) or Point-to-Point Protocol (PPP) connections.

If you plan to run the OnDemand OS/2 client program under IBM OS/2 3.0 or later, you must install the IBM TCP/IP Version 2.0 for OS/2 Base kit. If you plan to run the OnDemand OS/2 client program under IBM OS/2 WARP Connect, install the standard OS/2 WARP Connect TCP/IP support.

If you plan to run the OnDemand Windows client program under Windows 3.1 or later, TCP/IP sockets support is part of IBM TCP/IP for DOS and other compatible TCP/IP software packages. Contact your IBM representative for up-to-date information about supported TCP/IP packages. If you plan to run

the OnDemand Windows client program under Windows NT, Windows 95, or Windows 98, install the standard TCP/IP support provided with the operating system.

Print requirements

The OnDemand client program allows users to print directly from a document or to print selected items from the document list. There are three ways to print from the OnDemand client program:

- Print to a local printer. OnDemand can reprint all document formats (AFP, line data, and image file) under OS/2 and Windows. This approach is most commonly used to print to PostScript or HP PCL printers that are already defined and used by other applications on the user's workstation.

When printing AFP documents to local printers, the fidelity of the output can vary. The best fidelity is achieved when printing documents formatted with the Core Interchange Fonts. If an exact reprint of the document is required, then you must use one of the other printing options, which requires PSF.

- Print to devices attached to the network and managed by Print Services Facility. Some user's workstations will have a local printer port, such as LPT1, redirected to a Print Services Facility print queue. This method can allow reprinting of high volumes of documents at high speed, depending on the capabilities of the network printer, because the printing is offloaded to a print server.
- Print through the OnDemand server print facility. This is the highest performance print option, because the documents are not retrieved to the user's workstation before printing. It is also designed to allow many documents to be selected for reprint from the document list. When the OnDemand server print command is chosen, the client program sends a print request to the OnDemand server. The server then sorts the documents by optical platter or tape cartridge before retrieving them. After the documents are retrieved, a print job is submitted to a TCP/IP print queue, which is then processed by Print Services Facility. When an AFP document is printed, the resource group that was captured at the time the document was archived is put inline, to ensure that the document prints with the correct resources.

An OnDemand server printer is an interface between the user and a server print device. OnDemand supports two types of server print devices: a FAX machine and a physical printer. A server print device can be physically connected to the workstation where the library server is running or attached to another workstation in the network. Server print devices are managed by Print Services Facility. A server printer names a PSF print queue. In addition to defining a server printer with the administrative client, you must also identify the host name or IP address of the server

where PSF is running. The *Installation and Configuration Guide* for your server describes how to configure OnDemand server printing.

Printing jobs to PSF/MVS

PSF/2 and PSF for AIX provide facilities for uploading a print job to MVS for printing through PSF/MVS. This is desirable for many companies that want bill or statement reprints to go through the same process as the original bill. We recommend that you use PSF for AIX for this purpose, since it provides higher performance when uploading data to MVS than PSF/2.

Chapter 12. Planning for client programs

OS/2 client

The OnDemand OS/2 client software runs under IBM OS/2 Version 3.0 or later and IBM OS/2 WARP Connect and requires TCP/IP support. Please see “Chapter 3. Hardware and software requirements” on page 23 for a complete list of the hardware and software requirements for the OnDemand OS/2 client program.

Windows client

The OnDemand Windows client software runs under Microsoft Windows 3.1 or later, Windows 95, Windows 98, and Windows NT 4.0 and requires TCP/IP or IPX/SPX support. Please see “Chapter 3. Hardware and software requirements” on page 23 for a complete list of the hardware and software requirements for the OnDemand Windows client program.

Administrative client

The OnDemand administrative client runs under Microsoft Windows NT 4.0, Windows 95, and Windows 98 and requires TCP/IP or IPX/SPX support. Please see “Chapter 3. Hardware and software requirements” on page 23 for a complete list of the hardware and software requirements for the OnDemand administrative client.

Installing and running client software on a network

A personal computer attached to the network can share a single copy of the OnDemand client software. In general, you install client software on a network file server to save hard disk space on PCs and make it easier to upgrade software later. You typically install OnDemand client software on a PC for better performance and to reduce network traffic.

OnDemand provides a network installation procedure to copy OnDemand control files from a network file server to a user's PC (no program files are copied to the user's PC). As part of a network installation, the setup program

builds a directory structure on the user's PC and creates folders and icons on the desktop. The OnDemand program icons point to a copy of the software installed on the network file server.

When the user runs an OnDemand client program from a network file server, the operating system loads the OnDemand programs from the network file server into memory on the user's PC. When the user selects items for viewing, OnDemand allocates temporary work space on the user's PC for documents and resources.

Before you install a client program using the network installation option, you must install the client software on a network file server. See the *User's Guide* for an overview of installation options and details about the network installation procedure.

Client start up parameters

OnDemand provides command line parameters you can use to customize the operation of the client program. For example, you can automate the logon process so that the user does not have to specify a server, userid, or password. The *OS/2 Customization Guide* provides information about customizing the OS/2 client. The *OLE Control and 32-Bit GUI Customization Guide and Reference* provides information about customizing the client under Windows 95, Windows 98, and Windows NT. The books can be found on the OnDemand web site:

<http://www.software.ibm.com/data/ondemand/library.html>

Adobe software provided with OnDemand

OnDemand provides 16- and 32-bit versions of Adobe PDF viewing software so that users can view PDF documents stored in OnDemand. The 16-bit version of the PDF viewing software can be run under Windows 3.1 or later. The 32-bit version of the PDF viewing software can be run under Windows NT 4.0, Windows 95, and Windows 98. The *User's Guide* provides information about installing the Adobe PDF viewing software on a PC.

The Adobe Acrobat Exchange software license agreement provided with OnDemand allows you to use Acrobat Exchange as part of the OnDemand Windows client to view PDF documents stored in OnDemand. If you wish to view PDF documents outside of OnDemand, or to use Acrobat Exchange for other purposes, you must purchase a full Acrobat Exchange license from Adobe.

The OnDemand client CD-ROM includes versions of Adobe Type Manager Version (ATM) for each of the supported client operating systems. ATM Version 4.0 is provided for Windows NT 4.0, Windows 95, and Windows 98. You should install ATM if you plan to view AFP documents that use the IBM Core Outline Fonts or Sonoran Metric Outline Fonts (provided with OnDemand).

Chapter 13. Planning for programming interfaces

OnDemand provides several kinds of programming interfaces that can be used to customize OnDemand clients and perform additional server functions.

Client customization

OnDemand provides information about Object Linking and Embedding (OLE) control and how to customize OS/2 and Windows clients by specifying command line parameters, by invoking and manipulating OnDemand from other OS/2 and Windows 32-bit applications with the Dynamic Data Exchange (DDE) interface, or by creating a Product Information File (PIF).

The *OS/2 Customization Guide* provides information about customizing the OS/2 client. The *OLE Control and 32-Bit GUI Customization Guide and Reference* provides information about customizing the Windows 95, Windows 98, and Windows NT clients. The books can be found on the OnDemand web site:

<http://www.software.ibm.com/data/ondemand/library.html>

Server commands

OnDemand provides two server commands that can be used to automate processes. The commands are `arsadm` and `arsdoc`.

The `arsadm` command is a multi-purpose program for maintaining users, groups, and printers. You can use the `arsadm` command to add, delete, and update users, groups, and printers. You can run the `arsadm` command from the command line or a user-defined program.

The `arsdoc` command is a multi-purpose document processing program. You can use the `arsdoc` command to query the library server and generate a list of items that match a query; retrieve documents from the system; add, delete, and update documents; and send documents to the server print facility. You can run the `arsdoc` command from the command line or a user-defined program.

The *Administrator's Reference* provides details about these commands.

Server logging

System logging facility

OnDemand provides the system logging facility to help an administrator track activity and monitor the system. OnDemand can record messages generated by various client and server programs that support system events, user events, and application group events. For example, you can configure the system to record a message in the system log every time a user logs on the system. (For that matter, you can configure the system to record a message in the system log every time an unsuccessful log on attempt occurs.) When you use the administrative client to add and update the database, OnDemand records the information in the system log. You can use one of the OnDemand client programs to search for and view messages by time stamp, severity, message number, and userid.

System log user exit

OnDemand provides a configurable exit that can be used to process messages written to the OnDemand system log. A common use of the system log user exit is to look for error conditions or certain messages and take the appropriate action, such as notifying an administrator or operator or running some other program. The system log user exit runs the `arslog` command (UNIX servers) or `arslog.bat` file (Windows NT servers) after writing a record to the system log. However, the `arslog` command and the `arslog.bat` file provided with OnDemand do not perform any functions. You must replace them with programs that perform user-defined functions. For example, you would create an `arslog` command that checks the message number and severity of each message written to the system log and, when appropriate, sends an alert to Tivoli system management software.

OnDemand passes eight parameters to the system log user exit: the name of the OnDemand instance, time stamp, log record identifier, userid, the user's accounting information, message severity, message number, and message text. The information that appears in the accounting information part of the message can be set for each user with the `add/update` a user command. You can customize the message text by selecting the application group fields (and values) to include in the message. You can further configure OnDemand to provide specific information to the system log user exit by setting system and application group parameters with the administrative client.

See the *Installation and Configuration Guide* for your server for more information about the system log user exit.

Logon user exit

OnDemand provides a user exit that allows you to implement your own user exit program to identify and authenticate users that logon to the system.

The OnDemand request manager can optionally call a user exit program to establish a user's identity, using alternative methods of authenticating the user name and password.

When you implement your own logon user exit program, you bypass the logon verification processing built into the base OnDemand product. We advise caution when you bypass the OnDemand user and password restrictions. The security of the system could easily be subverted by malicious or defective code. Only use code that you trust.

Please contact the IBM support center for more information about implementing a logon user exit.

Chapter 14. Planning for backup and recovery

Overview

This chapter describes backup and recovery for OnDemand and provides recommendations about methods and procedures that an administrator can use to ensure the following critical OnDemand components can be recovered when needed:

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

OnDemand provides support for storing index data in DB2 table spaces and incremental backup of table spaces. DB2 table spaces enhance the management of index data and provide improved performance, especially for database backups. An incremental table space backup completes much quicker than a full database backup, providing you with increased flexibility in scheduling report loads. Incremental backup images also require less storage space than full database backups.

You can use ADSM to manage DB2 backup image files and archived log files. This capability means that you do not have to manage these files on disk.

Server software

If a media failure or some other unforeseen event occurs, you may be required to restore OnDemand software programs, unmodified OnDemand commands, and ADSM and database software. You can use the original product media to accomplish this task.

It is important that you store the original product media in a safe location. We recommend that you register OnDemand as part of your business recovery plan and store the original product media with other vital information systems records.

Server information

When you installed and configured OnDemand, you specified information that customized OnDemand to operate in your environment. In UNIX, this information is stored in various control files. In Windows NT, this information is stored in the Registry. We recommend that you backup the control files or Registry immediately after you have verified the installation of OnDemand. In addition, if you periodically make changes to the OnDemand server information, we recommend that you backup the control files or Registry on a regular basis, perhaps once a day.

Also, if you periodically make changes to or create new system, OnDemand, ADSM, and database control files or Registry, you may find it helpful to backup these files on a regular basis, perhaps once a week. Your system administrator probably schedules regular backup copies of your file systems. Check with the system administrator to make sure that server, OnDemand, ADSM, and database control files required for OnDemand operation are included in this backup.

You can use several different commands to make backup copies of files and file systems and to schedule regular backups of selected file systems, directories, and files. For example in UNIX, you can schedule the tar and mksysb commands with the cron facility. In Windows NT, you can use the NTBACKUP, REGBACK, and RDISK commands to backup the Registry and schedule the backup with the AT command. In addition, ADSM can be used to backup files, including the Registry, maintain the backups, and assist with recovery. ADSM is available on all of the server operating systems supported by OnDemand. See the operating system and device publications for your server for details about backup and restore concepts and commands.

OnDemand database

Database table spaces

Database table space support provides enhanced flexibility and improved performance for your application group data. For example, after you store a report in OnDemand, you can create a backup image of the table that changed during the load process, rather than creating a backup image of the entire database. You can also create an incremental backup image of the database, which contains only those tables that changed since the last backup image. Because the backup image only contains the changes made to the database, the backup process typically runs much faster than a full backup.

OnDemand creates one table space for each segment of application group data. After OnDemand closes the segment and you back up the table space, you do not need to back up the table space again, unless it is recovered or restored.

When you use the incremental table space backup capability, we recommend that you backup the OnDemand database after each report file load. If your schedule does not permit you to run the backup command after each load, we recommend that you backup the database once a day (assuming that you load multiple reports each day). While incremental backup images can be used to recover the database, we recommend that you periodically create a full backup image of the database. A full backup image of the database is the quickest way to recover the database in the event that you need to do so. However, if your OnDemand database is very large, and it cannot be backed up in a reasonable amount of time or requires a prohibitive number of storage volumes to hold, you may find that maintaining full backup images of the database is not possible.

The *Installation and Configuration Guide* for your server provides details about how to configure the system to support table spaces.

Database backup

OnDemand provides the `arsdb` command so that you can create backup images of the OnDemand database. The `arsdb` command supports table space and full database backups:

- OnDemand provides support for incremental table space backups and full database backups.
- An online backup can be taken when other applications or processes are connected to the database. That is, other applications and processes can continue to read or modify data while the backup is in process.
- During an offline backup, only the backup task is connected to the database. Before starting an offline backup, we recommend that you stop the OnDemand system to make sure that no other applications or processes are connected to the database.
- When you back up the database with the `arsdb` command, OnDemand removes the log files from the archived log file directory, releasing the space taken by files that are no longer needed. However, if you use ADSM to manage DB2 log files, the policy domain determines when archived log files are removed.

If your production schedule allows, we strongly encourage you to create offline backups on a regular schedule, perhaps once a week. Regularly scheduled offline backups can reduce the time required to rebuild table spaces or the database, if you need to do so. We recommend that you write offline

backup images to removable media or ADSM-managed storage. Keep backup images in a safe place, until the next time that you create an offline backup image of the table space or database.

If your schedule does not provide time to take offline backups (that is, your system must always remain available to users), you should take online backups on a regular schedule. The *Installation and Configuration Guide for UNIX Servers* shows how to use the cron facility to create online backups of the database with the `arsdb` command automatically on a regular schedule. The *Installation and Configuration Guide for Windows NT Servers* shows how to use the OnDemand NT configurator to create online backups of the database automatically on a regular schedule.

The *Administrator's Reference* provides details about the `arsdb` command, parameters, and options.

See the administration guide for your database product for details about backing up a database.

Using ADSM to manage backup images

You can use ADSM to manage DB2 backup image files. This eliminates the need for you to manage DB2 backup image files on disk. When you use the `arsdb` command to create table space backup images or backup the database, you can specify that you want ADSM to manage the database or table space backup images.

Before you can use ADSM to manage DB2 backup image files, you must define an ADSM storage hierarchy to manage the files. The storage hierarchy includes definitions that identify the type of media and storage devices that ADSM stores files on, the length of time that ADSM maintains files, and the number of backup copies that ADSM maintains. Before you schedule a backup command, you must make sure that ADSM storage volumes exist with sufficient free space to hold the backup image.

The *Installation and Configuration Guide* for your server provides details about how to configure the system to use ADSM to manage backup image files. See the *Administrator's Reference* for details about how to use the `arsdb` command to backup table spaces to ADSM-managed storage.

Database logging

The database manager uses transaction logging to record changes to the OnDemand database. The information in the log file is used to recover from corruption of data in the database. Logging ensures that no data is lost. By combining the information in the log files with a backup copy of the database, the OnDemand database can be recovered to any point in time.

Using ADSM to manage DB2 archived log files

You can use ADSM to manage DB2 archived log files. This eliminates the need for you to manage the log files on disk.

Before you can use ADSM to manage DB2 archived log files, you must define an ADSM storage hierarchy to manage the files. The storage hierarchy includes definitions that identify the type of media and storage devices that ADSM stores files on, the length of time that ADSM maintains files, and the number of backup copies that ADSM maintains. Before DB2 creates archived log files, you must make sure that ADSM storage volumes exist with sufficient free space to hold the files.

The *Installation and Configuration Guide* for your server provides details about how to configure the system to use ADSM to manage DB2 log files.

Database recovery

There are two types of database recovery. The first type recovers from failures that occur while update transactions are taking place. The log helps correct this type of failure by allowing the transactions received before the failure to either be reapplied to the database or to be *rolled-out*. Rolling-out transactions is a way to return the database to the state it was in before the transaction that caused the failure.

The second type of recovery deals with corruption of the OnDemand database and is usually caused by media failure. The combination of log files and a backup copy of the database can be used to recreate an image of the OnDemand database at a particular point in time.

If a catastrophic failure occurs, the system administrator will need to intervene to recover the database. Recovery from catastrophic failure starts with restoration of the latest full backup copy of the database. Next, the system administrator reapplies the transactions recorded in the log files. These steps will recreate a mirror image of the OnDemand database before the catastrophic failure.

The OnDemand database and database log files should reside on different physical volumes. The database backup image should be written to removable media. Unless multiple disk and tape volumes are damaged or lost **at the same time**, there is no possibility of losing the information contained in the OnDemand database.

The ADSM database

ADSM maintains a database that contains information about the devices and files it manages. When you load a report file into an application group that requires archive media support, ADSM updates the database and stores a copy of the file on a storage volume. When you define devices and register nodes, ADSM updates the database. When ADSM maintains storage volumes, it updates the database with status information about files and storage volumes. The ADSM database is critical to proper operation of ADSM in storing objects on and retrieving objects from the optical and tape storage volumes it manages.

Mirroring the ADSM database is strongly encouraged. When you mirror the database, ADSM replicates the database onto different physical storage. ADSM automatically keeps track of and refreshes both copies of the database. When you configure physical storage so that ADSM can mirror the database on different physical devices and adapters, you can provide protection for the database because of a failure of a single device. With mirroring, ADSM can continue operation without interruption if a database volume fails by using a mirrored copy of the failed volume. Mirroring requires additional storage space for the mirrored volumes. The *ADSM Administrator's Guide* provides details about mirroring the ADSM database.

To protect the information in the database, and ensure that it can be restored if a disaster occurs, you must periodically create a backup copy of the database. You can recover the database to its most current state or to a specific point in time with the backup copy.

- A full backup image of the database should be taken after you install and configure ADSM with OnDemand. In addition, we recommend that you create a full backup of the database on a regular schedule, such as once a week. A full backup copy of the database should be written to removable media.
- An incremental backup image of the database records changes that occurred since the last backup of the database (full or incremental). You can create a maximum of 32 incremental backups between full backups of the database. If you write incremental backup images of the database to disk, make sure that the disk is on a different controller than any of the database or recovery log volumes.

There are several factors to consider when you decide the type and frequency of backups.

- A full backup takes longer to run than an incremental backup.
- Recovery time is faster with a full backup. Incremental backups increase the time it takes to recover the database because a full backup must be loaded first, followed by some or all of the incremental backups.

- A full backup is required under specific conditions. You must create a full backup after installing and customizing ADSM with OnDemand. You can run up to 32 incremental backups between each full backup.

We recommend that you back up the database after you load report files into OnDemand and following ADSM maintenance of storage volumes (expiration and reclamation). Under typical conditions, you should back up the database each day.

ADSM includes a central scheduling component that allows the automatic processing of administrative commands, such as database backup. 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 DATABASE command, and activating the schedule. The *ADSM Administrator's Guide* provides details about scheduling operations.

For UNIX Servers, OnDemand provides the `ars_adsm` command to create a full backup of the database. If you invoke the data loading program from a shell script, you can add a step to the script that runs the `ars_adsm` command after loading reports into OnDemand. The *Installation and Configuration Guide for UNIX Servers* describes how to start the `ars_adsm` command and create a full backup image of the database automatically on a regular schedule. The *Administrator's Reference* provides details about the `ars_adsm` command, parameters, and options.

ADSM recovery log

The recovery log is critical to the operation of the ADSM server. If the recovery log is unusable, the entire server is unavailable. With the recovery log available, and a restored backup image of the database, you can recover the database to its most current state.

To ensure fast recovery time and high availability of the database, you should always mirror the recovery log. Mirroring the recovery log requires much less space than mirroring the database. If you do not mirror the recovery log, you should allocate the recovery log on a disk other than the one on which the database resides. Please refer to the *ADSM Administrator's Guide* for information about mirroring the ADSM recovery log.

When a database backup is completed, recovery log records preceding the backup are deleted, freeing up recovery log storage for reuse. Taking frequent database backups reduces recovery log storage requirements, and reduces the time required to recover the database.

ADSM volume history file

Volume history information is vital for recovery of a lost or damaged database. The volume history file contains information that ADSM needs about the volumes to use for database backups. The volume history file contains information that you need to know about storage pool volumes to audit after a recovery.

ADSM cannot obtain volume history information from the database during a restore of the database. Therefore, you should store at least one backup copy of the volume history file on a disk other than the one on which the database resides.

The *ADSM Administrator's Guide* provides details about establishing volume history backup files.

ADSM device configuration file

When you define, update, or delete a device class, drive, or library, ADSM updates the database and makes an entry in the device configuration file. To restore the database, ADSM requires a definition for the device class from which backup data is to be read. This definition is in the device configuration file.

When the database is being restored, no definitions can be read from the database. Therefore, you should have at least one backup copy of the device configuration information on a disk other than the one on which the database resides.

The *ADSM Administrator's Guide* provides details about establishing device configuration backup files.

ADSM database recovery

Recovering using mirrored copies of the database

If a database volume fails because of media failure and you have enabled mirroring, you can recover the database by using mirrored copies of the database. After fixing the failing device, you can allocate space for the new mirrored copy and define the volume to ADSM. After you define the volume to ADSM, the server synchronizes the volume with the database.

Recovering using backup copies of the database

ADSM provides commands to recover the database, should a catastrophic failure occur. These commands restore the database from the latest available

full backup copy, apply all incremental backups that apply, and use the recovery log to apply any changes made to the database since the last backup was created.

If you restore the database to its most current state, ADSM automatically synchronizes the database and storage volumes.

If you restore the database to a specific point in time, you must audit all storage pool volumes to check for and resolve any inconsistencies between database information and storage pool volumes. Depending on the number of storage pool volumes and the amount of activity that occurred after the database backup that you restored, the audit may require a significant amount of time.

To perform a database recovery, you should have the following information, preferably stored at a different location:

- Back up volumes of the database
- Copy storage pool volumes
- Server options file
- Volume history file
- Device configuration file
- Output from ADSM commands that provide details of the database and recovery log setup

The *ADSM Administrator's Guide* provides details about recovering data.

Archived reports

OnDemand can store copies of reports and resources in the document cache and on archive media.

- The primary purpose of the document cache is short-term, high-speed storage and retrieval. The document cache consists of magnetic storage volumes maintained by OnDemand on one or more object servers.
- The primary purpose of archive media is long-term storage and retrieval. Reports on archive media can also be used as backup copies, in the event cache storage becomes corrupted or unavailable. Archive media consists of optical or tape storage volumes managed by ADSM.

Most customers copy reports to cache storage and archive media at the same time, when they load a report into the system.

OnDemand can retrieve a copy of a report from archive media after the report has been deleted from cache storage or if the copy on cache storage is unavailable. However, you must properly configure your system to support

multiple copies of reports. You must install and configure ADSM, add archive media devices to your system, and configure OnDemand to use archive media. You configure OnDemand to use archive media by defining storage nodes in ADSM, assigning application groups to the appropriate storage sets, and setting application group data migration and caching information.

Note: If you do not copy reports to archive media, we recommend that you use ADSM to protect the cache file systems. ADSM can provide a range of storage management services. You can use ADSM to make backup copies of cache file systems. If a cache storage device fails, you can use ADSM to restore all or part of the cache file system. However, without a copy of reports on archive media, users cannot retrieve reports until cache storage is restored.

The document cache

The document cache is the primary, short-term storage location for reports.

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

The document cache can be protected by maintaining it on RAID storage subsystems. RAID storage can provide excellent availability, allowing users to access reports even if a disk or controller fails. “Chapter 9. Planning a storage subsystem” on page 89 provides information about using RAID storage with OnDemand. However, RAID storage is not fail safe. There may be situations when, because of multiple disk or controller failures, users cannot access reports. We encourage you to use ADSM to maintain a backup or secondary copy of reports.

Reports on archive media

The primary storage node identifies the object server and ADSM node where OnDemand stores the primary copy of a report on archive media. OnDemand retrieves the primary copy of the report from archive media after the report has been deleted from cache storage. Customers with critical, high availability requirements may need to create a backup or secondary copy of the report on archive media. The backup or secondary copy can be used if the primary copy becomes corrupted or unavailable.

There are two methods that you can use to create a backup or secondary copy of reports:

- Define a *copy storage pool* to ADSM. With this method, ADSM manages a backup copy of files that are stored in a primary storage pool

independently and transparently to OnDemand. The backup copy is stored in a copy storage pool that can be used to restore the original files if they become damaged, lost, or unusable. 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. Copy storage pools require additional space in the ADSM database. A copy storage pool must reside on the object server where the primary storage pool resides.

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 Administrator's Guide* provides details about defining and managing a copy storage pool, and storage pool backup and recovery, and scheduling operations.

- Assign the primary storage node to a secondary storage node. With this method, ADSM maintains a secondary copy of files that are stored in a primary storage node. The secondary copy can be used if the primary copy becomes damaged, lost, or unusable. However, administrator intervention is required before OnDemand can use the secondary copy. For this reason, we discourage the use of secondary storage nodes.

Part 3. Appendixes

Glossary

This glossary includes definitions from the following sources:

- Definitions reprinted from the *American National Dictionary for Information Processing Systems*, copyright 1982 by the Computer Business Equipment Manufacturers Association (CBEMA), are identified by the symbol (A) following the definition. Copies can be purchased from the American Standards Institute, 1430 Broadway, New York, New York 10018.
 - Definitions reprinted from a published section of the International Organization for Standardization's (ISO) *Vocabulary—Information Processing* or from a published section of the ISO *Vocabulary—Office Machines* are identified by the symbol (I) following the definition. Because many ISO definitions are also reproduced in the *American National Dictionary for Information Processing Systems* ISO definitions may also be identified by the symbol (A).
 - Definitions reprinted from working documents, draft proposals, or draft international standards of ISO Technical Committee 97, Subcommittee 1 (Vocabulary) are identified by the symbol (T) following the definition, indicating that final agreement has not yet been reached among its participating members.
 - Definitions are also reprinted from the *CCITT Eighth Plenary Assembly Red Book, Terms and Definitions* and working documents published by the International Telegraph and Telephone Consultative Committee of the International Telecommunication Union, Geneva, 1985.
- Definitions that are specific to IBM products are so labeled, for example, “In OnDemand,” or “In MVS.”

The following cross references are used in this glossary:

Contrast with. This refers to a term that has an opposed or substantively different meaning.

Synonym for. This indicates that the term has the same meaning as a preferred term, which is defined in its proper place in the dictionary.

Synonymous with. This is a backward reference from a defined term to all other terms that have the same meaning.

See. This refers the reader to multiple-word terms that have the same last word.

See also. This refers the reader to terms that have a related, but not synonymous, meaning.

A

access. To obtain data from or to put data in storage.

ACIF. See Advanced Function Presentation Conversion and Indexing Facility.

Acrobat. The Adobe viewer for PDF files. Acrobat is similar to the IBM AFP Workbench, that is, a stand-alone viewer. Acrobat also supports a robust set of APIs. It is through these APIs that Acrobat is integrated with the OnDemand client program.

active log file. The subset of files consisting of primary log files and secondary log files that are currently needed by the database manager for rollbacks and recovery.

active policy set. In ADSM, the policy set within the policy domain that contains the most

recently activated policy currently in use by all client nodes that have been assigned to that policy domain. See Policy Set.

adapter. A part that electrically or physically connects a device to a computer or to another device.

addressable point. Any point in a presentation surface that can be identified by a coordinate from the coordinate system of the presentation medium. See also Pel.

administrative client. (1) In ADSM, the program that allows administrators to control and monitor the server through administrator commands. (2) In OnDemand, the program that provides administrators with functions to manage OnDemand groups, users, printers, applications, application groups, storage sets, and folders.

ADSM. See ADSTAR Distributed Storage Manager.

ADSTAR Distributed Storage Manager. A program that provides storage management for archived files.

Advanced Function Presentation (AFP). A set of licensed programs that use the all-points-addressable concept to print data on a wide variety of printers or display data on a variety of display devices. AFP also includes creating, formatting, archiving, viewing, retrieving, and distributing information.

Advanced Function Presentation Application Programming Interface (AFP API). An AFP program shipped with PSF/MVS 2.1.1 and PSF/VM 2.1.1 that creates the AFP data stream from the COBOL and PL/1 high-level programming languages.

Advanced Function Presentation Conversion and Indexing Facility. An AFP program shipped with OnDemand that you can use to convert a print file into a MO:DCA-P document, to retrieve resources used by the document, and to index the file for later retrieval and viewing.

Advanced Function Presentation data stream (AFP data stream). A presentation data stream that is processed in the AFP environment. MO:DCA-P is the strategic AFP interchange data stream. IPDS is the strategic AFP printer data stream.

AFP. See Advanced Function Presentation.

AFP API. See Advanced Function Presentation Application Programming Interface.

AFPDS. A term formerly used to identify the composed page, MO:DCA-P-based data stream interchanged in AFP environments.

AIX. (1) Advanced Interactive Executive. (2) IBM's version of the UNIX operating system.

AIX Acrobat Libraries. A subset of the Acrobat Libraries ported to AIX for use by OnDemand.

all-points-addressable (APA). The capability to address, reference, and position data elements at any addressable position in a presentation space or on a physical medium. An example of all points addressability is the positioning of text, graphics, and images at any addressable point on the physical medium. See also Picture Element.

all-points-addressable mode. Synonym for Page Mode.

alphabetic character. A letter or other symbol, excluding digits, used in a language. Usually the uppercase and lowercase letters A through Z plus other special symbols (such as \$ and _) allowed by a particular language. See also Alphanumeric Character.

alphanumeric character. Consisting of letters, numbers, and often other symbols, such as punctuation marks and mathematical symbols. See also Alphabetic Character.

alphanumeric string. A sequence of characters consisting solely of the letters a through z and the numerals 0 through 9.

American National Standards Institute (ANSI). An organization for the purpose of establishing voluntary industry standards.

anchor point. The point in a document that signals to ACIF the beginning of a group of pages, after which it adds indexing structured fields to delineate this group.

ANSI. American National Standards Institute.

ANSI carriage control character. A character that specifies that a write, space, or skip operation should be performed before printing the line containing the carriage control. ANSI carriage control characters are encoded in ASCII or EBCDIC.

APA. All points addressable.

API. Application Program Interface.

application. In OnDemand, a definition of the attributes of a report, such as the data format and the compressing method. An application is defined for each output print data stream to be stored in OnDemand.

application group. One or more OnDemand applications with common indexing and storage management attributes, for example, invoice number and life of the data in OnDemand.

Application Program Interface (API). A formally defined programming language interface that is between a program and the user of a program.

archive copy group. In ADSM, a policy object containing attributes that control the generation, destination, and expiration of archive files. An archive copy group belongs to a management class.

archive log file. The subject of files consisting of primary log files and secondary log files that are no longer needed for normal database processing.

archive media. Devices and volumes used to store long-term copies of reports. For example, an optical storage library is one type of archive media supported by OnDemand.

ASCII (American Standard Code for Information Interchange). The standard code, using a coded character set consisting of 7-bit coded characters (8-bits including parity check), that is used for information interchange among data processing systems, data communication systems, and associated equipment. The ASCII set consists of control characters and graphic characters. (A)

attachment. A device or feature attached to a processing unit, including required adapters. Contrast with Adapter.

authentication. The process of checking a user's password before allowing the user access to resources or the server.

authorize. (1) To grant to a user the right to communicate with or make use of a computer system or display station. (2) To give a user either complete or restricted access to an object, resource, or function.

B

backend. In the AIX operating system, the program that sends output to a particular device. Synonymous with Backend Program.

backend program. Synonym for Backend.

bitmap. A file that contains a bit-mapped graphic.

BMP. Bitmap.

byte. The amount of storage required to represent 1 character; a byte is 8 bits.

C

cache. Short-term, magnetic storage. OnDemand loads the most recent and frequently used versions of reports and documents in the cache.

carriage control character. The first character of an output record (line) that is to be printed; it determines how many lines should be skipped before the next line is printed.

case-sensitive. Able to distinguish between uppercase and lowercase letters.

CCITT. Consultative Committee on International Telegraphy and Telephone.

CD-ROM. Compact disc read-only memory.

channel. A device connecting the processor to input and output devices.

channel adapter. A communication controller hardware unit used to attach the controller to a System/370 data channel.

channel-attached. (1) Pertaining to devices attached to a controlling unit by cables, rather than by telecommunication lines. (2) Synonymous with Local.

character. A letter, digit, or other symbol representing, organizing, or controlling data.

character rotation. The alignment of a character with respect to its character baseline, measured in degrees in a clockwise direction. Examples are 0°, 90°, 180°, and 270°. Zero-degree character rotation exists when a character is in its customary alignment with the baseline.

character set. A group of characters used for a specific reason; for example, the set of characters a printer can print or a keyboard can support.

click. To press the left mouse button while pointing to an object such as a command button or a toolbar button.

client. (1) In a distributed file system environment, a system that is dependent on a server to provide it with programs or access to programs. (2) A personal computer connected to a network running OnDemand software that can log on and query the library server, retrieve documents from OnDemand, and view and print documents.

client domain. The set of optical drives and storage volumes used by ADSM to store report files and resources belonging to an application group.

client node. An application group that has been registered to the ADSM server.

COBOL. Common business-oriented language. A high-level programming language, based on English, that is used primarily for business applications.

code page. An ordered set of up to 256 predefined display symbols. The first 32 code points of each code page are reserved for control codes and are the same for all code pages, leaving up to 224 distinct display symbols per page.

Code Page Global Identifier (CPGID). A unique code page identifier that can be expressed as either a two-byte binary or a five-digit decimal value.

code point. A character within a code page.

coded font. An AFP font that associates a code page and a font character set.

command. A request to perform an operation or run a program. When parameters values, flags, or other operands are associated with a command, the resulting character string is a single command.

command line. The area of the screen where commands are displayed as they are typed.

communication method. The method used by OnDemand and ADSM for exchanging information.

communication protocol. A set of defined interfaces that allow computers to communicate with each other.

compact disc read-only memory (CD-ROM). High capacity read-only memory in the form of an optically read compact disk.

composed page. In Advanced Function Presentation, a page that can be printed only on an all-points-addressable output medium. It may contain composed text and raster images.

composed-text data file. A file containing text data and text control information that dictates the format, placement, and appearance of the data to be printed.

compression. A technique for removing strings of duplicate characters, gaps, empty fields, and trailing blanks before transmitting data.

concatenate. (1) To link together. (2) To join two character strings.

concatenated field. Two or more fields from a physical file record format that have been combined to make one field in a logical file record format.

conditional processing. A page definition function that allows input data records to partially control their own formatting.

configuration. The process of describing to a system the devices, optional features, and program products that have been installed so that these features can be used. Contrast with Customization.

configuration file. A file that specifies the characteristics of a system or subsystem; for example, the operating system queueing system.

configure. To describe to a system the devices, optional features, and licensed programs installed on a system.

console. The main operating system display station.

constant. A data item with a value that does not change during the running of a program. Contrast with Variable.

Consultative Committee on International Telegraphy and Telephone (CCITT). A United Nations Specialized Standards group whose membership includes common carriers concerned with devising and proposing recommendations for international telecommunications representing alphabets, graphics, control information, and other fundamental information interchange issues.

control character. A character that is not a graphic character such as a letter, number, or punctuation mark. Such characters are called control characters because they frequently act to control a peripheral device.

controller. A device that coordinates and controls the operation of one or more input/output devices, such as workstations, and synchronizes the operation of the system as a whole.

conversion. In programming languages, the transformation between values that represent the same data item but belong to different data types.

copies. See Copy Group.

copy group. In ADSM, a policy object that contains attributes that control the generation, destination, and expiration of backup and archive files. There are two kinds of copy groups: backup and archive. Copy groups belong to management classes.

copy storage pool. A named collection of storage volumes that contains copies of files that reside in primary storage pools. Copy storage pools are used to back up the data stored in primary storage pools.

CPGID. Code Page Global Identifier

customization. The process of describing optional changes to defaults of a software program that is already installed on the system and configured so that it can be used. Contrast with Configuration.

customize. To describe the system, the devices, programs, users, and user defaults for a particular data processing system or network. Contrast with Configure.

D

daemon. In UNIX, a process begun by the root user or by the root shell that can be stopped only by the root user. Daemon processes generally provide services that must be available at all

times, such as sending data to the printer. A daemon runs continuously, looking for work to do, performing that work, and waiting for more work. A daemon does not have a controlling terminal associated with it. The OnDemand data download program (arsjesd) is an example of a daemon.

database. (1) The collection of information about all objects managed by OnDemand, including documents, groups, users, printers, application groups, storage sets, applications, and folders. (2) The collection of information about all objects managed by ADSM, including policy management objects, administrators, and client nodes.

Database Managed Space (DMS). A type of DB2 table space. A DSM table space is managed by the database manager.

data set. Synonym for File.

data stream. A continuous stream of data elements being transmitted, or intended for transmission, in character or binary-digit form using a defined format.

data transfer. The movement, or copying, of data from one location and the storage of the data at another location.

data type. The type, format, or classification of a data object.

DCF. Document Composition Facility.

decimal. Pertaining to a system of numbers to the base 10. The decimal digits range from 0 through 9.

decompression. A function that expands data to the length that preceded data compression. See also Compression.

default. A value, attribute, or option that is assumed when no alternative is specified by the user.

default directory. The directory name supplied by the operating system if none is specified.

default printer. A printer that accepts all the printed output from a display station assigned to it.

default value. A value stored in the system that is used when no other value is specified. See also Default.

desktop printer. In this publication, an IBM LaserPrinter 4019 or 4029, or compatible printer.

device class. A named group of ADSM storage devices. Each device class has a unique name and represents a device type of disk, tape, or optical disk.

device driver. A program that operates a specific device, such as a printer, disk drive, or display.

device type. A type of ADSM storage device. Each device class must be categorized with one of the following device types: disk, tape, or optical disk.

device-independent. Pertaining to a function that can be accomplished without regard for the characteristics of particular types of devices.

dialog box. An application window on the display that requests information from the user.

directory. (1) A type of file containing the names and controlling information for other files or directories. (2) A listing of related files arranged in a useful hierarchy.

disk operating system (DOS). An operating system for computer systems that use disks and diskettes for auxiliary storage of programs and data.

Distiller. A batch utility that converts PostScript files to PDF files. The distiller runs under AIX, HP-UX, Solaris, and Windows NT.

DMS. Database Managed Space.

document. (1) In OnDemand, a logical section of a larger file, such as an individual invoice within a report file of thousands of invoices. A document can also represent a 100 page segment

of a report. (2) A file containing an AFP data stream document. An AFP data stream document is bounded by Begin Document and End Document structured fields and can be created using a text formatter such as Document Composition Facility (DCF).

Document Composition Facility. An IBM licensed program used to prepare printed documents.

domain. See Policy Domain or Client Domain.

DOS. Disk operating system.

double-click. To rapidly press the left mouse button twice while pointing to an object.

download. To transfer data from one computer for use on another one. Typically, users download from a larger computer to a diskette or fixed disk on a smaller computer or from a system unit to an adapter.

drag. To hold down the left mouse button while moving the mouse.

driver. The end of a stream closest to an external interface. The principal functions of the driver are handling any associated device, and transforming data and information between the external device and stream.

E

EBCDIC. Extended Binary-Coded Decimal Interchange Code. This is the default type of data encoding in an MVS environment. Contrast with ASCII.

enqueue. To place items in a queue.

enter. (1) An instruction to type specific information using the keyboard. (2) A keyboard key that, when pressed, confirms or initiates the selected command.

environment variable. A variable that is included in the current software environment and is therefore available to any called program that requests it.

error condition. The state that results from an attempt to run instructions in a computer program that are not valid or that operate on data that is not valid.

error log. A file in a product or system where error information is stored for later access.

error log entry. In AIX, a record in the system error log describing a hardware or software failure and containing failure data captured at the time of the failure.

error message. An indication that an error has been detected. (A)

error recovery. The process of correcting or bypassing the effects of a fault to restore a computer system to a prescribed condition. (T)

error type. Identifies whether an error log entry is for a permanent failure, temporary failure, performance degradation, impending loss of availability, or undetermined failure.

Ethernet. A 10-megabit baseband local area network using CSMA/CD (carrier sense multiple access with collision detection). The network allows multiple stations to access the medium at will without prior coordination, avoids contention by using carrier sense and deference, and resolves contention by using collision detection and transmission.

exit program. A user-written program that is given control during operation of a system function.

exit routine. A routine that receives control when a specified event occurs, such as an error.

expiration. The process of deleting index data and reports based on application configuration information. The OnDemand database and storage managers perform expiration processing to remove reports and documents that are no longer needed from storage volumes and reclaim the space.

Extended Binary-Coded Decimal Interchange Code (EBCDIC). A coded character set consisting of eight-bit coded characters.

external library resource (member). Objects that can be used by other program products while running print jobs; for example, coded fonts, code pages, font character sets, form definitions, page definitions, and page segments. Synonym for Resource Object.

external object. Synonym for Resource Object.

F

FCB. Forms control buffer.

field. A specified area in a record used for a particular type of data; for example, a group of characters that represent a customer's name.

file. (1) A named set of records stored or processed as a unit. (T) (2) The major unit of data storage and retrieval. A file consists of a collection of data in one of several prescribed arrangements and described by control information to which the operating system has access.

file system. The collection of files and file management structures on a physical or logical mass storage device, such as a diskette or a minidisk.

file transfer. In remote communications, the transfer of a file or files from one system to another over a communications link.

File Transfer Protocol (FTP). In TCP/IP, the protocol that makes it possible to transfer data among hosts and to use foreign hosts indirectly.

fixed disk. A flat, circular, nonremovable plate with a magnetizable surface layer on which data can be stored by magnetic recording. A rigid magnetic disk.

fixed-disk drive. The mechanism used to read and write information on a fixed disk.

folder. In OnDemand, the end-user view of data archived in one or more application groups. Folders provide end-users a convenient way to

find related information, regardless of the source of the information or how the data was prepared.

font. (1) A family of characters of a given size and style, for example 9-point Helvetica. (2) A set of characters in a particular style. See Raster Font.

font character set. Part of an AFP font that contains the raster patterns, identifiers, and descriptions of characters. Often synonymous with Character Set. See also Coded Font.

form definition (FORMDEF). A form definition is a resource used by OnDemand. A form definition specifies the number of copies to be printed, whether the sheet should be printed on both sides, the position of a page of data on the sheet, text suppression, and overlays to be used (if any). Synonymous with FORMDEF.

FORMDEF. See Form Definition.

FSA. Functional SubSystem Application. A collection of programs residing in the FSS address space that control a device.

FSI. Functional SubSystem Interface. An MVS or OS/390 interface that allows communication between JES and a FSS and FSS applications. Download uses an FSI to communicate with the operating system and JES to process spool datasets created by application programs.

FSS. Functional SubSystem. An MVS or OS/390 subsystem comprised of programs residing in the same address space that provide JES-related functions. For example, a print programs that extend the scope of JES processing could be defined as a FSS.

FTP. File Transfer Protocol.

G

GB. Gigabyte.

GIF. Graphic Interchange Format.

gigabyte. A unit of memory or space measurement equal to approximately one billion bytes. One gigabyte equals 1,000 megabytes.

graphic. A symbol produced by a process such as handwriting, drawing, or printing. (I) (A)

graphic character. A character that can be displayed or printed.

Graphical User Interface. A type of user interface that takes advantage of a high-resolution monitor, including some combination of graphics, the use of pointing devices, menu bars, overlapping windows, and icons.

graphics. A type of data created from such fundamental drawing units such as lines, curves, polygons, and so forth.

Graphic Interchange Format (GIF). A bit-mapped color graphics file format for IBM and IBM-compatible computers. GIF employs an efficient compression technique for high resolution graphics.

group. (1) A named collection of sequential pages that form a logical subset of a document. (2) A named collection of users assigned a specific role or belonging to a specific department.

GUI. Graphical user interface.

H

hardware. The physical equipment of computing and computer-directed activities. The physical components of a computer system. Contrast with Software.

help. One or more files of information that describe how to use application software or how to perform a system function.

hex. See Hexadecimal.

hexadecimal (hex). Pertaining to a system of numbers in the base sixteen; hexadecimal digits range from 0 (zero) through 9 (nine) and A (ten) through F (fifteen).

host. (1) The primary or controlling computer in the communications network. (2) See Host System.

host-based computer. (1) In a computer network a computer that provides end users with services such as computation and data bases and that usually performs network control functions. (T) (2) The primary or controlling computer in a multiple-computer installation.

host system. (1) The controlling or highest level system in a data communication configuration, for example, an OS/390 system is the host system for the terminals connected to it. (2) In TCP/IP, a computer that is a peer system in a network.

I

icon. A 32 by 32 pixel bitmap used by the windows manager to represent an application or other window.

image. (1) An electronic representation of a picture produced by means of sensing light, sound, electron radiation, or other emanations coming from the picture or reflected by the picture. An image can also be generated directly by software without reference to an existing picture. (2) An electronic representation of an original document recorded by a scanning device.

index. (1) A process of segmenting a print file into uniquely identifiable groups of pages (a named collection of sequential pages) for later retrieval. (2) A process of matching reference points within a file and creating structured field tags within the MO:DCA-P document and the separate index object file.

index object file. An index-information file created by the OnDemand Conversion and Indexing Facility that contains the Index Element (IEL) structured fields, which identify the location of tagged groups in the AFP file. The indexing tags are contained in the Tagged Logical Element (TLE) structured fields.

indexing. (1) A process of segmenting a print file into uniquely identifiable groups of pages (a named collection of sequential pages) for later retrieval. (2) In ACIF, a process of matching reference points within a file and creating structured field tags within the MO:DCA-P document and the separate index object file.

indexing with data values. Adding indexing tags to a MO:DCA-P document using data that is already in the document and that is consistently located in the same place in each group of pages.

indexing with literal values. Adding indexing tags to a MO:DCA-P document by assigning literal values as indexing tags, because the document is not organized such that common data is located consistently throughout the document.

informational message. (1) A message that provides information to the end-user or system administrator but does not require a response. (2) A message that is not the result of an error condition.

input file. A file opened in order to allow records to be read.

install. (1) To add a program, program option, or software program to the system in a manner such that it may be executed and will interact properly with all affected programs in the system. (2) To connect a piece of hardware to the processor.

intelligent printer data stream (IPDS). An all-points-addressable data stream that allows users to position text, images, and graphics at any defined point on a printed page.

interface. Hardware, software, or both, that links systems, programs, or devices.

Internet. A wide area network connecting thousands of disparate networks in industry, education, government, and research. The Internet network uses TCP/IP as the protocol for transmitting information.

Internet Protocol (IP). In TCP/IP, a protocol that routes data from its source to its destination in an Internet environment.

IP. Internet Protocol.

IPDS. Intelligent printer data stream.

J

job. One or more related procedures or programs grouped into a procedure, identified by appropriate job control statements.

job queue. A list of jobs waiting to be processed by the system.

Joint Photographic Experts Group (JPEG). An image compression standard developed to handle larger images with many colors. JPEG uses a lossy algorithm, which means there is some loss of detail when saving and viewing images in this format. However, JPEG files can offer as much as 35% improvement in file size and compression.

JPEG. See Joint Photographic Experts Group.

K

kernel. The part of an operating system that performs basic functions such as allocating hardware resources.

kernel extension. A program that modifies parts of the kernel that can be customized to provide additional services and calls. See Kernel.

K-byte. See Kilobyte.

keyword. Part of a command operand that consists of a specific character string.

kilobyte (K-byte). 1024 bytes in decimal notation when referring to memory capacity; in all other cases, it is defined as 1000.

L

LAN. Local area network.

LAN server. A data station that provides services to other data stations on a local area network; for example, file server, print server, mail server. (T)

laser printer. A nonimpact printer that creates, by means of a laser beam directed on a photosensitive surface, a latent image which is then made visible by toner and transferred and fixed on paper. (T)

Lempel Ziv Welsh (LZW). A data compression algorithm. OnDemand uses the 16-bit version of LZW to compress data.

library. System storage for generated form definitions and page definitions.

library resource (member). A named collection of records or statements in a library.

library resource name. A name by which an object may be called from a library by Advanced Function Printing as part of a print job. Includes the 2-character prefix for the type of object, such as P1 for page definitions, F1 for form definitions, or O1 for overlays (also known as *resource name*).

library server. In OnDemand, the node in the TCP/IP network that contains the core OnDemand code and the OnDemand database.

licensed program. A separately priced program and its associated materials that bear a copyright and are offered to customers under the terms and conditions of a licensing agreement.

line data. Data prepared for printing on a line printer, such as an IBM 3800 Model 1 Printing Subsystem. Line data is usually characterized by carriage-control characters and table reference characters.

line-data print file. A file that consists of line data, optionally supplemented by a limited set of structured fields.

line printer. A device that prints a line of characters as a unit. (I) (A) Contrast with Page Printer.

literal. (1) A symbol or a quantity in a source program that is itself data, rather than a reference to data. (2) A character string whose value is given by the characters themselves; for example, the numeric literal 7 has the value 7, and the character literal CHARACTERS has the value CHARACTERS.

loading. The logical process of archiving reports in OnDemand. During the loading process, OnDemand programs process reports, create index data, and store report data and resources on cache storage volumes and archive media.

local. Pertaining to a device accessed directly without use of a telecommunication line.

local area network (LAN). (1) A computer network located on a user's premises within a limited geographical area. Communication within a local area network is not subject to external regulations; however, communication across the LAN boundary may be subject to some form of regulation. (2) A network in which a set of devices is connected to one another for communication and that can be connected to a larger network. See also Token-Ring Network.

logical volume. The combined space from all volumes defined to either the ADSM database or recovery log. The database resides on one logical volume and the recovery log resides on a different logical volume.

log file. A fixed-length file used to record changes to a database.

LPD. Line Printer Daemon. In TCP/IP, the command responsible for sending data from the spooling directory to a printer.

LPR. Line Printer Requestor. In TCP/IP, a client command that allows the local host to submit a file to be printed on a remote print server.

LZW. See Lempel Ziv Welsh.

M

M byte. Megabyte.

MB. Megabyte.

machine carriage control character. A character that specifies that a write, space, or skip operation should be performed either immediately or after printing the line containing the carriage control.

mainframe. A large computer, particularly one to which other computers can be connected so that they can share facilities the mainframe provides. The term usually refers to hardware only.

management class. In ADSM, a policy object that is a named collection of copy groups. A management class can contain one backup copy group, one archive copy group, a backup and archive copy group, or zero copy groups. Users can bind each file to a management class to specify how the server should manage backup versions or archive copies of files. See Copy Group.

mapping. (1) A list that establishes a correspondence between items in two groups. (2) The process of linking database fields in an application group to folder search and display fields.

megabyte (MB). When used with hard drive, diskette, or removable media storage capacity, 1,000,000 bytes. When referring to system memory capacity, 1,048,576 bytes.

memory. Program-addressable memory from which instructions and other data can be loaded directly into registers for subsequent running or processing. Memory is sometimes referred to as “storage”.

menu bar. The area at the top of a window that contains choices that give a user access to actions available in that window.

message. Information from the system that informs the user of a condition that may affect further processing of a current program.

migration. (1) The process of moving data from one computer system to another without converting the data. (2) The process of moving report files, resources, and index data from cache storage to long-term (optical or tape) storage.

mirroring. In ADSM, a feature that protects against data loss with the database or recovery log by writing the same data to multiple disks at the same time. Mirroring supports up to three exact copies of each database or recovery log.

Mixed Object Document Content Architecture - Presentation (MO:DCA-P). (1) A strategic, architected, device-independent data stream for interchanging documents. (2) A printing data stream that is a subset of the Advanced Function Presentation data stream.

MO:DCA-P. Mixed Object: Document Content Architecture for Presentation.

mount. To make a file system accessible.

mouse. A hand-held locator that a user operates by moving it on a flat surface. It allows the user to select objects and scroll the display screen by pressing buttons.

Multiple Virtual Storage (MVS). Consists of MVS/System Product Version 1 and the MVS/370 Data Facility Product operating on a System/370 processor.

MVS. Multiple virtual storage; an IBM operating system.

N

network. A collection of data processing products that are connected by communication lines for information exchange between locations.

Network File System (NFS). A protocol developed by Sun Microsystems that uses Internet Protocol to allow a set of cooperating computers to access each other's file system as if they were local.

NFS. Network File System.

node. A workstation that operates as an OnDemand library server or object server and is connected to a TCP/IP network.

notes. Comments, clarifications, and reminders that can be attached to an OnDemand document.

non-IPDS printer. In this publication, a printer that is not channel-attached and which does not accept the Intelligent Printer Data Stream.

numeric. Pertaining to any of the digits 0 through 9.

O

object. (1) A collection of structured fields. The first structured field provides a begin-object function and the last structured field provides an end-object function. The object may contain one or more other structured fields whose content consists of one or more data elements of a particular data type. An object may be assigned a name, which may be used to reference the object. Examples of objects are text, graphics, and image objects. (2) A resource or a sequence of structured fields contained within a larger entity, such as a page segment or a composed page. (3) A collection of data referred to by a single name.

object server. A storage manager node connected to a TCP/IP network that provides cache management and, optionally, support for files archived on optical and tape storage volumes.

offset. The number of measuring units from an arbitrary starting point in a record, area, or control block to some other point.

online. Being controlled directly by or directly communicating with the computer.

operating environment. (1) The physical environment; for example, temperature, humidity, and layout. (2) All of the basic functions and the user programs that can be executed by a store controller to enable the devices in the system to perform specific operations. (3) The collection of store controller data, user programs, lists, tables, control blocks,

and files that reside in a subsystem store controller and control its operation.

operating system. Software that controls the running of programs and that also can provide such services as resource allocation, scheduling, input and output control, and data management.

optical library. A disk storage device that houses optical disk drives and optical disks, and contains a mechanism for moving optical disks between a storage area and optical disk drives.

optimize. To improve the speed of a program or to reduce the use of storage during processing.

outline fonts. (1) Fonts whose graphic character shapes are defined as mathematical equations rather than by raster patterns. (2) Fonts created in the format described in *Adobe Type 1 Font Format*, a publication available from Adobe Systems, Inc. Synonymous with Type 1 fonts.

overlay. A collection of predefined, constant data such as lines, shading, text, boxes, or logos, that is electronically composed and stored as an AFP resource file that can be merged with variable data on a page while printing or viewing.

P

page. (1) A collection of data that can be printed on one side of a sheet of paper or a form. (2) The boundary for determining the limits of printing. See also Logical Page and Physical Page. (3) Part of an AFP document bracketed by a pair of Begin Page and End Page structured fields.

page definition. A resource used by OnDemand that defines the rules of transforming line data into composed pages and text controls.

page printer. A device that prints one page as a unit. (I) (A) Contrast with Line Printer.

page segment. In Advanced Function Presentation, a resource that can contain text and images and can be positioned on any addressable point on a page or an electronic overlay.

PAGEDEF. Page definition

parallel device. A device that can perform two or more concurrent activities. Contrast with Serial Device.

parameter. (1) Information that the user supplies to a panel, command, or function. (2) In the AIX operating system, a keyword-value pair.

partitioned data set. A data set in direct access storage that is divided into partitions, called members, each of which can contain a program, part of a program, or data.

path. In a network, any route between any two nodes.

path name. A name that specifies the location of a directory within a file system. Path names are used to locate and reference directories and their contents.

PC. Personal Computer.

PCL. Printer control language.

PCX. Picture Exchange Format.

PDF. Portable Document Format.

permissions. Codes that determine how the file can be used by any users who work on the system.

personal computer. A microcomputer primarily intended for stand-alone use by an individual. (T)

Picture Exchange Format (PCX). A file that contains a graphic in the PCX graphics file format, which was originally developed for the PC Paintbrush program, but is now widely used by other programs.

piobe. The printer input/output back end program used by AIX for printing tasks.

pipe. To direct the data so that the output from one process becomes the input to another process. The standard output of one command can be connected to the standard input of

another with the pipe operator (`|`). Two commands connected in this way constitute a pipeline.

point. (1) To move the mouse pointer to a specific object. (2) A unit of typesetting measure equal to 0.01384 inch (0.35054 mm), or about 1/72 of an inch. There are 12 points per pica.

point size. The height of a font in points. See also Point.

policy domain. A policy object that contains policy sets, management classes, and copy groups that is used by a group of client nodes. See Policy Set, Management Class, Copy Group, and Client Node.

policy set. A policy object that contains a group of management class definitions that exist for a policy domain. At any one time, there can be many policy sets within a policy domain but only one policy set can be active. See Management Class and Active Policy Set.

port. (1) A part of the system unit or remote controller to which cables for external devices (display stations, terminals, or printers) are attached. The port is an access point for data entry or exit. (2) A specific communications end point within a host. A port is identified by a port number.

Portable Document Format. A distilled version of PostScript data that adds structure and efficiency. PDF data has the same imaging model as PostScript but does not have its programmability. PDF also provides direct access to pages and allows hypertext links, bookmarks, and other navigational aids required for viewing. The text in a PDF file is usually compressed using LZW methods. The images in a PDF file are usually compressed using CCITT or JPEG methods.

PostScript. Adobe's page description language used for printing. PostScript is a very flexible programming language and imaging model but is not as structured as AFP. PostScript cannot be parsed to determine page boundaries, it must be interpreted. Because of this limitation, PostScript

is not practical for archiving and viewing. Adobe created PDF for archiving and viewing.

press. To touch a specific key on the keyboard.

primary log file. A set of one or more log files used to record changes to a database. Storage for these files is allocated in advance.

primary storage pool. A named collection of storage volumes that ADSM uses to store archive copies of files.

print file. A file that a user wants to print.

print job. A series of print files scheduled for printing. At print submission time, the user can request one or more files to be printed; therefore, a print job consists of one or more print files.

print queue. A file containing a list of the names of files waiting to be printed.

Print Services Facility (PSF). A sophisticated IBM print subsystem that drives IPDS page printers. PSF is supported under MVS, VSE, VM, OS/2, AIX, and is a standard part of the operating system under OS/400. PSF manages printer resources such as fonts, images, electronic forms, form definitions, and page definitions, and provides error recovery for print jobs.

When printing line data, PSF supports external formatting using page definitions and form definitions. This external formatting extends page printer functions such as electronic forms and use of typographic fonts without any change to applications that generate the data.

Print Services Facility/2 (PSF/2). PSF/2 is an OS/2-based print server that drives IPDS page printers, as well as IBM PPDS and HP-PCL compatible printers. PSF/2 manages printer resources and provides error recovery for print jobs. PSF/2 supports distributed printing of AFP print jobs from PSF for AIX, PSF/MVS, PSF/VSE, PSF/VM, and OS/400. PSF/2 also supports printing from a wide range of workstation applications, including Microsoft Windows and OS/2 Presentation Manager, as well as the ASCII, PostScript, and AFP data streams.

Print Services Facility for AIX (PSF for AIX). An IBM licensed program that produces printer commands from the data sent to it and it runs on the AIX/6000 operating system.

print spooler. The print spooler directs the printing of data from different applications. It temporarily stores information in separate files until they are printed.

Printer Control Language (PCL). The data stream used by Hewlett-Packard LaserJet II and III and other compatible printers.

process. An activity within the system that is started, such as a command, a shell program, or another process.

profile. (1) A file containing customized settings for a system or user. (2) Data describing the significant features of a user, program, or device.

program level. The version, release, modification, and fix levels of a program.

prompt. A displayed symbol or message that requests information or operator action.

protocol. A set of semantic and syntactic rules that determines the behavior of functional units in achieving communication.

PSF. Print Services Facility.

PSF/2. Print Services Facility/2.

PSF for AIX. Print Services Facility for AIX.

PTF. Program temporary fix.

Q

qdaemon. The daemon process that maintains a list of outstanding jobs and sends them to the specified device at the appropriate time.

qualified name. (1) A data name explicitly accompanied by a specification of the class to which it belongs in a specified classification system. (I) (A) (2) A name that has been made unique by the addition of one or more qualifiers.

queue. (1) A line or list formed by items waiting to be processed. (2) To form or arrange in a queue.

queue device. A logical device defining characteristics of a physical device attached to a queue.

R

radio button. Round option buttons grouped in dialog boxes; one is preselected. Like a radio in an automobile, select only one button (“station”) at a time.

RAM. Random access memory. Specifically, the memory used for system memory. Sometimes this memory is referred to as main storage.

raster. In Advanced Function Presentation, an on/off pattern of electrostatic images produced by the laser print head under control of the character generator.

raster font. A font in which the characters are defined directly by the raster bit map. See Font. Contrast with Outline Font.

raster graphics. Computer graphics in which a display image is composed of an array of pixels arranged in rows and columns.

read access. In computer security, permission to read information.

record. (1) In programming languages, an aggregate that consists of data objects, possibly with different attributes, that usually have identifiers attached to them. (2) A set of data treated as a unit. (3) A collection of fields treated as a unit.

recovery log. In ADSM, a log of updates that are about to be written to the database. The log can be used to recover from system and media failures.

recovery procedure. (1) An action performed by the operator when an error message appears on the display screen. This action usually permits the program to run the next job. (2) The method

of returning the system to the point where a major system error occurred and running the recent critical jobs again.

register. Define a client node to ADSM.

remote. Pertaining to a system or device that is accessed through a communications line. Contrast with Local.

remote print. Issuing print jobs to one machine (client) to print on another machine (server) on a network.

remote system. A system that is connected to your system through a communication line.

report file. A print data stream produced by an application program that can contain hundreds or thousands of pages of related information. Some report files can be segmented into single and multiple page objects called documents.

resolution. (1) In computer graphics, a measure of the sharpness of an image, expressed as the number of lines and columns on the display screen. (2) The number of pels per unit of linear measure.

resource. A collection of printing instructions, and sometimes data to be printed, that consists entirely of structured fields. A resource can be stored as a member of a directory and can be called for by the Print Services Facility when needed. The different resources are: coded font, character set, code page, page segment, overlay, and form definition.

resource directory. A place in which resource files are stored.

resource management. The function that protects serially accessed resources from concurrent access by computing tasks.

retention. The amount of time, in days, that archived files will be retained in ADSM before they are deleted.

retry. To try the operation that caused the device error message again.

return code. (1) A value that is returned to a program to indicate the results of an operation issued by that program. (2) A code used to influence the running of succeeding instructions.

root. The user name for the system user with the most authority.

root file system. The file system that contains all the default installation and program directories in the system.

root user. In UNIX environments, an expert user who can log in and execute restricted commands, shut down the system, and edit or delete protected files.

root volume group. The volume group, identified with a single / (forward slash) that contains all the directories in the root file system.

rotation. (1) The alignment of a character with respect to its character baseline, measured in degrees in a clockwise rotation. Examples are 0°, 90°, 180°, and 270°. Zero-degree character rotation exists when a character is in its customary alignment with the baseline. Synonymous with Character Rotation. (2) The number of degrees a character is turned relative to the page coordinates. (3) The orientation of the characters of a font with respect to the baseline.

routing. The assignment of the path by which a message will reach its destination.

S

secondary log file. A set of one or more log files used to record changes to a database. Storage for these files is allocated as needed when the primary log fills up.

segment. (1) A collection of composed text and images, prepared before formatting and included in a document when it is printed. See Page Segment. (2) The resource that contains the structured-field definition of a page segment. (3) A 100 page portion of a report file. OnDemand divides report files into segments to provide enhanced performance and maintenance.

select. To pick a menu command or other object with a single click of the mouse.

serial device. A device that performs functions sequentially, such as a serial printer that prints one byte at a time. Contrast with Parallel Device.

server. (1) On a network, the computer that contains the data or provides the facilities to be accessed by other computers on the network. (2) A program that handles protocol, queuing, routing, and other tasks necessary for data transfer between devices in a computer system. (3) A workstation connected to a TCP/IP network that runs the OnDemand programs that store, retrieve, and maintain report files. OnDemand supports two types of servers: a library server and an object server.

server options file. The ADSM file that specifies processing options for communication methods, tape handling, pool sizes, language, and date, time, and number formats.

shell. In UNIX environments, a software interface between a user and the operating system of a computer. Shell programs interpret commands and user interactions on devices such as keyboards and pointing devices and communicate them to the operating system.

skip-to-channel control. A line printer control appearing in line data. Allows space to be left between print lines. Compatible with page printers when the data is formatted by page definitions.

SMIT. System Management Interface Tool.

SMS. System Managed Space.

software. Programs, procedures, rules, and any associated documentation pertaining to the operating of a system. Contrast with Hardware.

spool file. (1) A disk file containing output that has been saved for later printing. (2) Files used in the transmission of data among devices.

spooling (simultaneous peripheral operation online). Performing a peripheral operation such as printing while the computer is busy with other work.

spooling subsystem. A synonym for the queuing system that pertains to its use for queuing print jobs.

stand-alone workstation. A workstation that can perform tasks without being connected to other resources such as servers or host systems.

standard input. The primary source of data going into a command. Standard input comes from the keyboard unless redirection or piping is used, in which case standard input can be from a file or the output from another command.

standard output. The primary destination of data coming from a command. Standard output goes to the display unless redirection or piping is used, in which case standard output can be to a file or another command.

status. (1) The current condition or state of a program or device. For example, the status of a printer. (2) The condition of the hardware or software, usually represented in a status code.

storage. (1) The location of saved information. (2) In contrast to memory, the saving of information on physical devices such as disk or tape.

storage device. A functional unit for storing and retrieving data.

storage hierarchy. A logical ordering of storage devices. Generally, the ordering is based on the speed and capacity of the devices.

storage node. A named object that identifies an ADSM domain on an OnDemand object server. There are two types of storage nodes: a primary storage node and a secondary storage node. Application group data is archived in a primary storage node. A secondary storage node contains a backup copy of data stored in a primary node.

storage object. A portion of a storage volume managed as a single entity. A storage object can contain many report file segments.

storage pool. A named collection of storage volumes that is the destination for archived files.

storage pool volume. A volume that has been assigned to a storage pool to store archived files.

storage set. A named collection of storage nodes with the same storage management characteristics. For example, the life of the data in OnDemand.

storage volume. A disk volume that has been assigned to store reports and documents on the OnDemand server.

string. A series or set of alphabetic or numeric characters. A string can be composed of letters, numbers, and special characters.

structure. A variable that contains an ordered group of data objects. Unlike an array, the data objects within a structure can have varied data types.

structured field. (1) A self-identifying, variable-length, bounded record that can have a content portion that provides control information, data, or both. (2) A mechanism that permits variable length data to be encoded for transmission in the data stream. See Field.

subdirectory. In the file system hierarchy, a directory contained within another directory.

subroutine. (1) A sequenced set of statements or coded instructions that can be used in one or more computer programs and at one or more points in a computer program. (2) A routine that can be part of another routine.

syntax. The grammatical rules for constructing a command, statement, or program.

syntax diagram. A diagram for a command that displays how to enter the command on the command line.

system console. A console, usually equipped with a keyboard and display screen, that is used by an operator to control and communicate with a system. Synonymous with Console.

system customization. Specifying the devices, programs, and users for a particular data processing system. See also Configuration.

system integrity. In computer security, the quality of a system that can perform its intended function in an unimpaired manner, free from deliberate or inadvertent unauthorized manipulation of the system.

System Managed Space (SMS). A type of DB2 table space. An SMS table space is managed by the filesystem manager.

system management. The tasks involved in maintaining the system in good working order and modifying the system to meet changing requirements.

System Management Interface Tool (SMIT). In the AIX operating system, a series of panels that allow you to perform system functions without directly issuing any commands.

system memory. Synonymous with Main Storage, but used in hardware to refer to semiconductor memory (modules).

system prompt. Synonym for command line. The system prompt is the symbol that appears at the command line of an operating system. The system prompt indicates that the operating system is ready for the user to enter a command.

T

table. A named collection of data consisting of rows and columns.

table reference character (TRC). (1) Usually, the second byte on a line in the user's data. This byte contains a value (0–126) that is used to select a font to be used to print that line. (2) In the 3800 Printing Subsystem, a numeric character (0, 1, 2, or 3) corresponding to the order in which the character arrangement table names have been

specified with the **CHARS** keyword. It is used for selection of a character arrangement table during printing.

table space. An abstraction of a collection of containers into which database objects are stored. A table space provides a level of indirection between a database and the tables stored within the database. A table space:

- Has space on media storage devices assigned to it.
- Has tables created within it.

tag. (1) A type of structured field used for indexing in an AFP document. Tags associate an index attribute-value pair with a specific page or group of pages in a document. (2) In text formatting markup language, a name for a type of document element that is entered in the source document to identify it.

Tagged Image File Format (TIFF). A bit-mapped graphics format for scanned images with resolutions of up to 300 dpi. TIFF simulates gray scale shading.

TB. Terabyte.

TCP. Transmission Control Protocol.

TCP/IP. Transmission Control Protocol/Internet Protocol.

terabyte. A unit of memory or space measurement capacity equal to approximately one trillion bytes. One terabyte is equal to 1,000 gigabytes, or one million megabytes.

text. (1) A type of data consisting of a set of linguistic characters (letters, numbers, and symbols) and formatting controls. (2) In word processing, information intended for human viewing that is presented in a two-dimensional form, such as data printed on paper or displayed on a screen.

throughput. A measure of the amount of work performed by a computer system over a period of time, for example, the number of jobs per day. (I)

TIFF. Tagged Image File Format.

token name. An eight-byte name that can be given to all data stream objects.

token-ring network. A ring network that allows unidirectional data transmission between data stations, by a token passing procedure, such that the transmitted data return to the transmitting station. (T)

toolbar. The region directly beneath the menu bar of the main window in OnDemand client programs that support a graphical user interface.

toolbar button. A small bitmap on the toolbar that represents a command in OnDemand client programs that support a graphical user interface. Click a toolbar button to quickly access a command.

trigger. Data values that ACIF searches for in a print data stream, to delineate the beginning of a new group of pages. The first trigger is then the anchor point ACIF uses to locate index values.

transfer. To send data to one place and to receive data at another place.

transform. To change the form of data according to specified rules without significantly changing the meaning of the data. (I) (A)

Transmission Control Protocol (TCP). A communications protocol used in Internet and in any network that follows the U.S. Department of Defense standards for inter-network protocol. TCP provides a host-to-host protocol between hosts in packet-switched communications networks and in interconnected systems of such networks. It assumes that the Internet protocol is the underlying protocol.

Transmission Control Protocol/Internet Protocol (TCP/IP). A set of communications protocols that support peer-to-peer connectivity functions for both local and wide area networks.

TRC. Table reference character.

type. To enter specific information using the keyboard, typing characters exactly as given.

U

unformatted print data. Data that is not formatted for printing. A page definition can contain controls that map unformatted print data to its output format.

UNIX operating system. An operating system developed by Bell Laboratories that features multiprogramming in a multi-user environment. The UNIX operating system was originally developed for use on minicomputers but has been adapted for mainframes and microcomputers.

upload. To transfer data from one computer to another. Typically, users upload from a small computer to a large one.

user. A person authorized to logon to an OnDemand server.

user exit. (1) A point in an IBM-supplied program at which a user exit routine may be given control. (2) A programming service provided by an IBM software product that may be requested during the execution of an application program for the service of transferring control back to the application program upon the later occurrence of a user-specified event.

user interface. The hardware, software, or both that implements a user interface, allowing the user to interact with and perform operations on a system, program, or device. Examples are a keyboard, mouse, command language, or windowing subsystem.

V

value. (1) A set of characters or a quantity associated with a parameter or name. (2) A quantity assigned to a constant, variable, parameter, or symbol.

variable. (1) A name used to represent a data item whose value can change while the program is running. (2) In programming languages, a language object that can take different values at

different times. (3) A quantity that can assume any of a given set of values.

version number. The version level of a program, which is an indicator of the hardware and basic operating system upon which the program operates. The version, release, modification, and fix levels together comprise the program level or version of a program.

virtual printer. A view of a printer that refers only to the high-level data stream, such as ASCII or PostScript, that the printer understands. It does not include any information about how the printer hardware is attached to the host computer or the protocol used for transferring data to and from the printer.

volume. The basic unit of storage for a database, log file, or a storage pool. A volume can be an LVM logical volume, a standard file system file, a tape cartridge, or an optical platter. Each volume is identified by a unique volume identifier.

W

wildcard. Search characters that represent other letters, numbers, or special characters. In OnDemand, the %(percentage) and the _(underscore) are wildcard characters.

window. A part of a display screen with visible boundaries in which information is presented.

workstation. A terminal or microcomputer, usually one that is connected to a mainframe or to a network, at which a user can perform applications.

write access. In computer security, permission to write to an object.

writer. A JES function that processes print output.

Index

Numerics

- 2000 ready, Year 44
- 3527 SSA storage subsystem 95
- 7133 serial disk system 95
- 7135 RAIDiant array 95

A

- about OnDemand 3
- access to servers 112
- ACIF
 - MVS systems 44
 - OnDemand servers 45
 - OS/390 systems 44
 - running 44, 45
- Acrobat Exchange 116
- administrative client
 - planning 115
- administrative clients
 - hardware and software requirements 28
- administrator interface setup program
 - hardware and software requirements 29
- Adobe Acrobat Exchange 116
- Adobe PDF viewing software 116
- Adobe Type Manager 116
- ADSM
 - database backup and recovery 128
 - database recovery 130
 - database storage requirements 74
 - device configuration file 130
 - recovering the database 130
 - recovery log 129
 - recovery log storage requirements 74
 - report backup and recovery 131
 - storage manager overview 12
 - volume history file 130
- ADSM and OnDemand Environment 18
- AFP Application Programming Interface
 - generate index data with 43
- AFP conversion and indexing facility (ACIF)
 - overview 13

- AFP print data
 - evaluating 37
- application
 - overview 6
 - properties 52
- application group
 - overview 6
 - properties 49
- application group database organization 50
- application group field 59
- application group field types 58
- application group indexing methods 50
- application planning overview 47
- archive media libraries,
 - estimating 80
- archive media storage requirements
 - calculating 75
- archive media storage volumes,
 - estimating 80
- arsacif
 - data indexing planning and requirements 37
- assessing network and printer requirements 111
- ATM 116

B

- backup and recovery
 - ADSM database 128, 130
 - ADSM device configuration file 130
 - ADSM recovery log 129
 - ADSM volume history file 130
 - control files 123
 - database 124
 - OnDemand 123
 - OnDemand database 124
 - overview 123
 - product software 123
 - recovery log 129
 - reports and resources 131
- base system software storage requirements 65
- Bitmap (BMP) 38
- BMP (Bitmap) 38

C

- cache expiration processing 52
- cache migration processing 51
- cache storage requirements 67
- Calculating ADSM Database Storage 74
- Calculating Archive Log File Storage 73
- calculating archive media requirements 75
- Calculating Archive Media Storage Space 75
- Calculating Cache Storage 67
- calculating database archived log file storage 79
- Calculating Database Storage Space 69
- Calculating Database Transaction Log File Storage 71
- Calculating Log File Storage 72
- calculating magnetic storage requirements 65
- Calculating Migrated Index Storage 80
- Calculating Report Download Space 66
- Calculating Temporary Work Space 66
- client access to servers 112
- client print requirements 113
- client/server environment
 - system overview 4
- commands 119
- compliant
 - Year 2000 44
- compression 64
- concepts
 - application 6
 - application group 6
 - client/server system overview 4
 - document 8
 - folder 6
 - indexing methods 7
- configuration
 - 3527 SSA storage subsystem 95
 - 7133 serial disk system 95
 - 7135 RAIDiant array 95

- configuration (*continued*)
 - OnDemand server
 - large system
 - configuration 100
 - medium system
 - configuration 100
 - small system
 - configuration 100
 - RAID storage 93
 - RAIDiant array 93
- Configuring Logical Volumes and File Systems 91
- Configuring the ars Volume Group 90
- conventions xvi
 - highlighting xvi
 - OnDemand names 56

D

- data
 - compression 64
 - download
 - overview 12
 - download, planning for 111
 - evaluating indexing requirements 39
 - indexing 39
 - overview 13
- database
 - backup 125
 - backup and recovery 124
 - logging 126
 - recovery 127
 - storage requirements 67
 - transaction log file storage requirements 71
- Database and Folder Field Types 58
- database backup and recovery 128
- Database Calculation Example 69, 70
- database expiration processing 51
- database field names 56
- database field types 58
- database management product
 - DB2 12
 - overview 12
- database organization
 - multiple loads per database table 50
 - single load per database table 51
- date fields
 - database field type 58
 - folder field type 58
 - indexing 44
- disk system 95
- document
 - in OnDemand 8
- Document Indexing Method 41
 - download
 - data from MVS and OS/390, overview 12
 - Download for OS/390
 - about 12
 - overview 12
 - planning for 112
 - software requirements 27
 - download storage requirements 65

E

- end-user requirements 37
- environment variables
 - time zone (TZ) 59
 - TZ (time zone) 59
- estimating
 - memory requirements 97, 98
 - processor requirements 97, 98
 - storage requirements 61
 - storage subsystem 89
- estimating storage requirements
 - example of 81, 83
- evaluating image file formats 37
- evaluating input data formats 37
- evaluating print data formats 37
- examples
 - estimating storage requirements 81, 83
- Exchange 116
- expiration processing 51

F

- field names 56
- field types, overview 58
- file formats supported by OnDemand
 - PDF (Portable Data Format) 38
 - Portable Data Format (PDF) 38
 - user-defined 38
- folder
 - application group field 59
 - overview 6, 48
 - segment field 59
 - text search field 60
- folder field names 56
- folder field types 58
- folders, application groups, and applications
 - overview 7

G

- generate index data in the application program 42
- generate index data using the OnDemand generic index file format 43
- generate index data using the OnDemand PDF Indexer 44
- generate index data with AFP application programs 43
- generate index data with the AFP Application Programming Interface 43
- generic index file
 - generating index data with 43
- GIF (Graphic Interchange Format) 38
- Graphic Interchange Format (GIF) 38
- groups
 - overview 53

H

- hardware requirements
 - administrative clients 28
 - administrator interface setup program 29
 - memory 97, 98, 99
 - OS/2 client 27
 - OS/2 client setup program 29
 - processor 97, 98, 99
 - Windows client 28
 - Windows client setup program 29
- highlighting conventions xvi
- HP-UX
 - memory requirements 97
 - planning for 97
 - processor requirements 97

I

- image file formats 37
- image file formats supported by OnDemand
 - Bitmap (BMP) 38
 - BMP (Bitmap) 38
 - GIF (Graphic Interchange Format) 38
 - Graphic Interchange Format (GIF) 38
 - JFIF (JPEG Format Image File) 38
 - JPEG Format Image File (JFIF) 38

- image file formats supported by OnDemand (*continued*)
 - PCX (Picture Exchange Format) 38
 - Picture Exchange Format (PCX) 38
 - Tagged Image File Format (TIFF) 38
 - TIFF (Tagged Image File Format) 38
- image files
 - evaluating 37
- imported index data storage requirements 75
- index
 - migration storage space 79
- index data storage requirements 75
- Index Field Type and Index Size 68
- indexing
 - ACIF 44, 45
 - alternatives to ACIF 42
 - data 39
 - evaluating requirements for 39
 - methods 7
 - MVS systems 44
 - OnDemand servers 45
 - OS/390 systems 44
 - other indexing programs 42
 - overview 13
 - programs 42
 - using date fields 44
- indexing data 43, 44
 - generate index data in the application program 42
 - insert AFP records in the data stream 43
 - with AFP application programs 43
- indexing methods 50
- input data formats 37
- insert AFP records in the data stream 43
- interfaces
 - commands 119
 - logging 120
 - login user exit 121
 - logon user exit 121
 - OnDemand server
 - commands 119
 - OnDemand server logging 120
 - OnDemand server logon 121
 - server commands 119
 - server logging 120
 - system log 120

- introduction
 - about OnDemand 3
 - data indexing 39
 - hardware requirements 23
 - indexing data 39
 - input data formats 37
 - print data formats 37
 - role of the system
 - administrator 31
 - software requirements 23
 - to OnDemand 1

J

- JFIF (JPEG Format Image File) 38
- JPEG Format Image File (JFIF) 38

L

- large OnDemand system
 - storage devices with Windows NT 92
- Large Server Configuration 100
- Library/Object Server 101
- Library/Object Server with ADSM 102
- Library Server and Object Server 104, 106
- Library Server with ADSM and Object Server with ADSM 108
- license for Adobe software 116
- line data report files
 - evaluating 37
- logging 120
- login user exit 121
- logon user exit 121

M

- management programs
 - overview 15
- medium OnDemand system
 - storage devices with Windows NT 91
- Medium Server Configuration 100
- memory requirements 97, 98, 99
- migrated index storage space requirements 79
- multiple loads per database table 50
- MVS Download
 - about 12
 - overview 12
 - software requirements 27
- MVS download feature
 - planning for 112
- MVS indexing
 - software requirements 27

N

- names, in OnDemand 56
- network
 - client software 115
- network requirements, planning for 111
- NT
 - memory requirements 98
 - planning for 98
 - processor requirements 98
- NT license information 110

O

- OnDemand 90
 - applications
 - overview of documents 8
 - client/server environment 4
 - concepts 6
 - conventions 56
 - data download
 - overview 12
 - database management product
 - overview 12
 - database storage requirements 67
 - document
 - overview 8
 - evaluating image file formats 37
 - evaluating indexing requirements 39
 - evaluating input data formats 37
 - evaluating print data formats 37
 - field types 58
 - generic index file
 - generating index data with 43
 - groups overview 53
 - hardware requirements 23
 - image file formats 37
 - indexing data 39
 - indexing methods 7
 - indexing overview 13
 - input data formats 37
 - introduction 1
 - management programs 15
 - names 56
 - object names 56
 - overview of servers 9
 - PDF Indexer
 - generating index data with 44
 - permissions 53
 - planning 33
 - preparing for 35

- OnDemand 90 *(continued)*
 - print data formats 37
 - request manager overview 11
 - servers, overview 9
 - software requirements 23
 - storage manager overview 12
 - system overview 4
 - understanding end-user requirements 37
 - users overview 53
 - OnDemand and ADSM 17
 - OnDemand Applications and Documents 9
 - OnDemand folders, application groups, and applications 7
 - OnDemand server commands 119
 - OnDemand server logging 120
 - OnDemand server logon 121
 - OnDemand system environment 5
 - Optical Libraries 26
 - OS/2 client
 - hardware and software requirements 27
 - planning 115
 - setup program
 - hardware and software requirements 29
 - OS/390 indexing
 - software requirements 27
 - overview of a folder 48
 - overview of ADSM
 - for the OnDemand administrator 17
 - Overview of Data Indexing and Loading 14
 - overview of OnDemand
 - for the ADSM administrator 20
- P**
- PCX (Picture Exchange Format) 38
 - PDF (Portable Data Format) 38
 - PDF Indexer
 - generating index data with 44
 - PDF viewing software 116
 - permissions, overview 53
 - Picture Exchange Format (PCX) 38
 - planning
 - administrative client 115
 - applications 47
 - backup and recovery 123
 - client programs 115
 - client software on a network 115
 - documents 37
 - end-user requirements 37
 - planning *(continued)*
 - for OnDemand 33
 - interfaces 119
 - memory requirements 97, 98
 - networks 111
 - NT license information 110
 - OnDemand applications 47
 - OS/2 client 115
 - overview 35
 - preparing for OnDemand 35
 - printers 111
 - processor requirements 97, 98
 - programming interfaces 119
 - RAID storage 93
 - recovery 123
 - reports 37
 - storage requirements 61
 - storage sizing worksheets 84
 - storage subsystem 89
 - UNIX license information 110
 - Windows client 115
 - planning applications, overview 47
 - planning for NT based server
 - memory requirements 98
 - processor requirements 98
 - server configurations 98
 - Portable Data Format (PDF) 38
 - preparing for OnDemand 35
 - print data formats 37
 - print requirements 113
 - print storage space requirements 74
 - printing
 - software requirements 27
 - processor requirements 97, 98, 99
 - programming interfaces
 - commands 119
 - logging 120
 - login user exit 121
 - logon user exit 121
 - OnDemand server commands 119
 - OnDemand server logging 120
 - OnDemand server logon 121
 - server commands 119
 - server logging 120
 - system log 120
 - RAID storage 93
 - RAIDiant array 93, 95
 - recovering the ADSM database 130
 - report backup and recovery 131
 - report download storage requirements 65
 - Report Indexing Method 42
 - report storage space requirements 75
 - request manager
 - overview 11
 - requirements
 - ADSM database storage 74
 - ADSM recovery log storage 74
 - archive media 75
 - base system software 65
 - cache storage 67
 - data download 111
 - data loading storage space 66
 - database, transaction log file storage 71
 - database storage 67, 74
 - download storage 65
 - end-users 37
 - estimating storage 61
 - hardware 23, 97, 98, 99
 - HP-UX memory 98
 - HP-UX processor 97
 - log file storage 71
 - magnetic storage 65
 - memory requirements 97, 98, 99
 - migrated index storage on optical 79
 - network 111
 - NT memory 99
 - NT processor 99
 - OnDemand database storage 67
 - overview of storage 61
 - printing 113
 - processor requirements 97, 98, 99
 - recover log storage 74
 - report cache 67
 - report download storage 65
 - report storage on optical 75
 - RS/6000 memory 97
 - RS/6000 processor 97
 - software 23
 - Sun Solaris memory 98
 - Sun Solaris processor 98
 - temporary space 66
 - temporary space for importing index data 75
 - temporary space for server print 74
 - transaction log file storage 71
 - role of the system administrator 31
 - RS/6000
 - memory requirements 97
 - planning for 97
 - processor requirements 97

S

- segment field 59
- serial disk system 95
- server
 - access to 112
 - client access to 112
 - hardware requirements 23, 97, 98, 99
 - memory requirements 97, 98, 99
 - NT license information 110
 - planning client access 112
 - print requirements 113
 - processor requirements 97, 98, 99
 - software requirements 23
 - UNIX license information 110
- server commands 119
- server logging 120
- server planning
 - memory requirements 97
 - processor requirements 97
 - server configurations 97
- server print
 - software requirements 27
- servers
 - environment 9
 - overview 9
- setting the time zone environment variable 59
- single load per database table 51
- small OnDemand system
 - storage devices with Windows NT 91
- Small Server Configuration 100
 - software requirements
 - administrative clients 28
 - administrator interface setup program 29
 - Download for OS/390 27
 - indexing 27
 - MVS Download 27
 - MVS indexing 27
 - OS/2 client 27
 - OS/2 client setup program 29
 - OS/390 indexing 27
 - printing 27
 - server print 27
 - Windows client 28
 - Windows client setup program 29
- software storage requirements 65
- storage
 - 3527 SSA storage subsystem 95
 - 7133 serial disk system 95
 - 7135 RAIDiant array 95

- storage (*continued*)
 - planning for RAID
 - subsystems 93
 - RAID 93
- storage devices with Windows NT
 - large OnDemand system 92
 - medium OnDemand system 91
 - setting up 91
 - small OnDemand system 91
- storage requirements
 - ADSM database storage 74
 - ADSM recovery log storage 74
 - cache storage 67
 - database 67, 74
 - database, transaction log file 71
 - download space 65
 - estimating 61
 - loading data 66
 - log file 71
 - OnDemand database 67
 - planning for 61
 - recovery log 74
 - report download space 65
 - temporary space 66
 - transaction log file 71
- storage requirements overview
 - estimating storage 61
 - overview of storage 61
- storage sizing examples 81
- storage sizing worksheets 84
- storage subsystem 95
 - estimating 89
 - planning 89
- Sun Solaris
 - memory requirements 98
 - planning for 98
 - processor requirements 98
- system administrator
 - roles and responsibilities 31
- system log 120
- system software storage requirements 65

T

- Tagged Image File Format (TIFF) 38
- TCP/IP
 - data download, planning for 111
- temporary space requirements 66
- text search field 60
- TIFF (Tagged Image File Format) 38
- time zone environment variable 59
- transmitting files
 - Download for OS/390 112

- transmitting files (*continued*)
 - MVS download feature 112
- TZ environment variable 59

U

- understanding end-user requirements 37
- UNIX license information 110
- user-defined data format 38
- user-defined file format 38
- user exit
 - login 121
 - logon 121
- users
 - overview 53
- using RAID storage 93

V

- volume groups 90

W

- Windows client
 - hardware and software requirements 28, 29
 - installing and running on a network 115
 - planning 115
- worksheets
 - storage sizing 84

Y

- Year 2000 ready 44

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

EDMSuite OnDemand
Introduction and Planning Guide
Version 2.2

Publication No. G544-5281-04

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

	Very Satisfied	Satisfied	Neutral	Dissatisfied	Very Dissatisfied
Overall satisfaction	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

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

	Very Satisfied	Satisfied	Neutral	Dissatisfied	Very Dissatisfied
Accurate	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Complete	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Easy to find	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Easy to understand	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Well organized	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Applicable to your tasks	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

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.



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 Software Solutions
Department MYGA Building 001L
PO Box 1900
Boulder, CO 80301-9817



Fold and Tape

Please do not staple

Fold and Tape



Part Number: CT7KUNA



Printed in the United States of America
on recycled paper containing 10%
recovered post-consumer fiber.

G544-5281-04



CT7KUNA



Spine information:



EDMSuite OnDemand

Introduction and Planning Guide

Version 2.2