

OSF[®] DCE Application Development Guide — Directory Services

Release 1.2.2

Open Software Foundation
11 Cambridge Center
Cambridge, MA 02142

The information contained within this document is subject to change without notice.

OSF MAKES NO WARRANTY OF ANY KIND WITH REGARD TO THIS MATERIAL, INCLUDING BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE.

OSF shall not be liable for errors contained herein, or for any direct or indirect, incidental, special or consequential damages in connection with the furnishing, performance, or use of this material.

Copyright © 1995, 1996 Open Software Foundation, Inc.

This documentation and the software to which it relates are derived in part from materials supplied by the following:

Copyright © 1990, 1991, 1992, 1993, 1994, 1995, 1996 Digital Equipment Corporation

Copyright © 1990, 1991, 1992, 1993, 1994, 1995, 1996 Hewlett-Packard Company

Copyright © 1989, 1990, 1991, 1992, 1993, 1994, 1995, 1996 Transarc Corporation

Copyright © 1990, 1991 Siemens Nixdorf Informationssysteme AG

Copyright © 1990, 1991, 1992, 1993, 1994, 1995, 1996 International Business Machines

Copyright © 1988, 1989, 1995 Massachusetts Institute of Technology

Copyright © 1979, 1980, 1983, 1986, 1988, 1989, 1991, 1992, 1993, 1994 The Regents of the University of California

Copyright © 1995, 1996 Hitachi, Ltd.

All Rights Reserved

Printed in the U.S.A.

THIS DOCUMENT AND THE SOFTWARE DESCRIBED HEREIN ARE FURNISHED UNDER A LICENSE, AND MAY BE USED AND COPIED ONLY IN ACCORDANCE WITH THE TERMS OF SUCH LICENSE AND WITH THE INCLUSION OF THE ABOVE COPYRIGHT NOTICE. TITLE TO AND OWNERSHIP OF THE DOCUMENT AND SOFTWARE REMAIN WITH OSF OR ITS LICENSORS.

Open Software Foundation, OSF, the OSF logo, OSF/1, OSF/Motif, and Motif are registered trademarks of the Open Software Foundation, Inc.

X/Open is a registered trademark, and the X device is a trademark, of X/Open Company Limited.

The Open Group is a trademark of the Open Software Foundation, Inc. and X/Open Company Limited.

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

DEC, DIGITAL, and ULTRIX are registered trademarks of Digital Equipment Corporation.

DECstation 3100 and DECnet are trademarks of Digital Equipment Corporation.

HP, Hewlett-Packard, and LaserJet are trademarks of Hewlett-Packard Company.

Network Computing System and PasswdEtc are registered trademarks of Hewlett-Packard Company.

AFS, Episode, and Transarc are registered trademarks of the Transarc Corporation.

DFS is a trademark of the Transarc Corporation.

Episode is a registered trademark of the Transarc Corporation.

Ethernet is a registered trademark of Xerox Corporation.

AIX and RISC System/6000 are registered trademarks of International Business Machines Corporation.

IBM is a registered trademark of International Business Machines Corporation.

DIR-X is a trademark of Siemens Nixdorf Informationssysteme AG.

MX300i is a trademark of Siemens Nixdorf Informationssysteme AG.

NFS, Network File System, SunOS and Sun Microsystems are trademarks of Sun Microsystems, Inc.

PostScript is a trademark of Adobe Systems Incorporated.

Microsoft, MS-DOS, and Windows are registered trademarks of Microsoft Corp.

NetWare is a registered trademark of Novell, Inc.

FOR U.S. GOVERNMENT CUSTOMERS REGARDING THIS DOCUMENTATION AND THE ASSOCIATED SOFTWARE

These notices shall be marked on any reproduction of this data, in whole or in part.

NOTICE:Notwithstanding any other lease or license that may pertain to, or accompany the delivery of, this computer software, the rights of the Government regarding its use, reproduction and disclosure are as set forth in Section 52.227-19 of the FARS Computer Software-Restricted Rights clause.

RESTRICTED RIGHTS NOTICE:Use, duplication, or disclosure by the Government is subject to the restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 52.227-7013.

RESTRICTED RIGHTS LEGEND:Use, duplication or disclosure by the Government is subject to restrictions as set forth in paragraph (b)(3)(B) of the rights in Technical Data and Computer Software clause in DAR 7-104.9(a). This computer software is submitted with "restricted rights." Use, duplication or disclosure is subject to the restrictions as set forth in NASA FAR SUP 18-52.227-79 (April 1985) "Commercial Computer Software-Restricted Rights (April 1985)." If the contract contains the Clause at 18-52.227-74 "Rights in Data General" then the "Alternate III" clause applies.

US Government Users Restricted Rights - Use, duplication or disclosure restricted by GSA ADP Schedule Contract.

Unpublished - All rights reserved under the Copyright Laws of the United States.

This notice shall be marked on any reproduction of this data, in whole or in part.

Contents

Preface	xix
Audience	xix
Applicability	xix
Purpose	xx
Document Usage	xx
Related Documents	xx
Typographic and Keying Conventions	xxi
Problem Reporting	xxii
Pathnames of Directories and Files in DCE Documentation	xxii

Part 1. DCE Directory Service

Chapter 1. DCE Directory Service Overview	1-1
1.1 Introduction to This Guide	1-1
1.1.1 Use of This Guide	1-2
1.1.2 Directory Service Tools	1-3
1.2 Using the DCE Directory Service	1-3
1.3 DCE Directory Service Concepts	1-4
1.4 Structure of DCE Names	1-8
1.4.1 DCE Name Prefixes	1-9
1.4.2 Names of Cells	1-10
1.4.3 CDS Names	1-11
1.4.4 GDS Names	1-12
1.4.5 Junctions in DCE Names	1-12
1.4.6 Application Names.	1-13
1.5 The Federated DCE Namespace	1-13
1.5.1 The GDS Namespace	1-14

1.5.2	The CDS Namespace	1-16
1.5.3	Other Namespaces	1-18
1.6	Programming Interfaces to the DCE Directory Service . . .	1-18
1.6.1	The XDS Interface	1-18
1.6.2	The RPC Name Service Interface	1-18
1.6.3	Namespace Junction Interfaces	1-19

Part 2. CDS Application Programming

Chapter 2.	Programming in the CDS Namespace	2-1
2.1	Initial Cell Namespace Organization	2-1
2.1.1	The Cell Profile	2-4
2.1.2	The LAN Profile	2-4
2.1.3	The CDS Clearinghouse	2-4
2.1.4	The Hosts Directory	2-5
2.1.5	The Subsystems Directory	2-5
2.1.6	The /: DFS Alias	2-6
2.1.7	The DFS and DCE Security Service Junctions	2-6
2.2	Recommended Use of the CDS Namespace	2-7
2.2.1	Storing Data in CDS Entries	2-7
2.2.2	Access Control for CDS Entries	2-10
2.3	Valid Characters and Naming Rules for CDS	2-14
2.3.1	Metacharacters	2-17
2.3.2	Additional Rules	2-18
2.3.3	Maximum Name Sizes	2-20
2.4	Use of OIDs	2-24
Chapter 3.	XDS and the DCE Cell Namespace	3-1
3.1	Introduction to Accessing CDS with XDS	3-1
3.1.1	Using the Reference Material in This Chapter	3-2
3.1.2	What You Cannot Do with XDS	3-3
3.1.3	Registering A Nonlocal Cell	3-3
3.2	XDS Objects	3-3
3.2.1	Object Attributes	3-5
3.2.2	Interface Objects and Directory Objects	3-6
3.2.3	Directory Objects and Namespace Entries	3-9
3.2.4	Values That an Object Can Contain	3-11
3.2.5	Building a Name Object	3-11
3.2.6	A Complete Object	3-14
3.2.7	Class Hierarchy	3-15
3.2.8	Class Hierarchy and Object Structure	3-15

3.2.9	Public and Private Objects and XOM	3-16
3.2.10	XOM Objects and XDS Library Functions	3-17
3.3	Accessing CDS Using the XDS Step-by-Step Procedure	3-17
3.3.1	Reading and Writing Existing CDS Entry Attributes With XDS	3-17
3.3.2	Creating New CDS Entry Attributes	3-35
3.4	Object-Handling Techniques	3-39
3.4.1	Using XOM to Access CDS	3-39
3.4.2	Dynamic Creation of Objects	3-42
3.5	XDS/CDS Object Recipes	3-44
3.5.1	Input XDS/CDS Object Recipes	3-44
3.5.2	Input Object Classes for XDS/CDS Operations	3-45
3.6	Attribute and Data Type Translation.	3-60

Part 3. GDS Application Programming

Chapter 4.	GDS API: Concepts and Overview	4-1
4.1	Directory Service Interfaces	4-2
4.2	The X.500 Directory Information Model.	4-3
4.2.1	Directory Objects	4-3
4.2.2	Attribute Types	4-5
4.2.3	Object Identifiers	4-5
4.2.4	Object Entries	4-7
4.3	X.500 Naming Concepts	4-9
4.3.1	Distinguished Names	4-9
4.3.2	Relative Distinguished Names and Attribute Value Assertions	4-11
4.3.3	Multiple AVAs	4-11
4.3.4	Aliases	4-12
4.3.5	Name Verification	4-14
4.4	Schemas	4-14
4.4.1	The GDS Standard Schema	4-14
4.4.2	The Structure Rule Table	4-15
4.4.3	The Object Class Table	4-17
4.4.4	The Attribute Table	4-22
4.4.5	Defining Subclasses	4-24
4.5	Abstract Syntax Notation 1	4-25
4.5.1	ASN.1 Types	4-26
4.5.2	Basic Encoding Rules	4-27
4.6	GDS as a Distributed Service	4-28

4.6.1	The Directory Access Protocol	4-29
4.6.2	The Directory System Protocol	4-30
4.6.3	Referral	4-30
4.6.4	Chaining	4-31
4.6.5	The Directory User Agent Cache	4-33
4.6.6	GDS Configurations	4-39
4.6.7	GDS Security	4-40
4.6.8	GDS API Logging	4-41
Chapter 5. XOM Programming		5-1
5.1	OM Objects	5-2
5.1.1	OM Object Attributes	5-2
5.1.2	Object Identifiers	5-7
5.1.3	C Naming Conventions	5-8
5.1.4	Public Objects	5-10
5.1.5	Private Objects	5-22
5.1.6	Object Classes	5-23
5.2	Packages	5-31
5.2.1	The Directory Service Package	5-31
5.2.2	The Basic Directory Contents Package	5-32
5.2.3	The Strong Authentication Package	5-33
5.2.4	The GDS Package	5-33
5.2.5	The MHS Directory User Package	5-34
5.2.6	Package Closure	5-35
5.3	Workspaces	5-35
5.4	Storage Management	5-36
5.5	OM Syntaxes for Attribute Values	5-39
5.5.1	Enumerated Types	5-39
5.5.2	Object Types	5-41
5.5.3	Strings	5-41
5.5.4	Other Syntaxes	5-42
5.6	Service Interface Data Types	5-42
5.6.1	The OM_descriptor Data Type	5-43
5.6.2	Data Types for XDS API Function Calls	5-46
5.6.3	Data Types for XOM API Calls	5-47
5.7	OM Function Calls	5-48
5.7.1	Summary of OM Function Calls	5-48
5.7.2	Using the OM Function Calls	5-49
5.8	XOM API Header Files	5-56
5.8.1	XOM Type Definitions and Symbolic Constant Definitions	5-56

5.8.2	XOM API Macros	5–56
Chapter 6.	XDS Programming	6–1
6.1	XDS Interface Management Functions	6–2
6.1.1	The ds_initialize() Function Call	6–3
6.1.2	The ds_version() Function Call	6–4
6.1.3	The ds_shutdown() Function Call	6–6
6.2	Directory Connection Management Functions	6–6
6.2.1	A Directory Session	6–6
6.2.2	The ds_bind() Function Call	6–7
6.2.3	The ds_unbind() Function Call	6–9
6.2.4	Automatic Connection Management	6–9
6.3	XDS Interface Class Definitions	6–9
6.3.1	Example: The DS_C_FILTER Class	6–10
6.3.2	The DS_C_CONTEXT Parameter	6–10
6.4	Directory Class Definitions	6–11
6.5	The GDS Package	6–13
6.5.1	Authentication	6–13
6.5.2	Access Control	6–13
6.5.3	DUA Cache	6–15
6.5.4	Advanced Administration Operations	6–16
6.6	Directory Operation Functions	6–16
6.7	Directory Read Operations	6–17
6.7.1	Reading an Entry from the Directory	6–18
6.7.2	Step 1: Export Object Identifiers for Required Directory Classes and Attributes	6–18
6.7.3	Step 2: Declare Local Variables	6–19
6.7.4	Step 3: Build Public Objects	6–20
6.7.5	Step 4: Create an Entry-Information-Selection Parameter	6–21
6.7.6	Step 5: Perform the Read Operation	6–23
6.8	Directory Search Operations	6–26
6.8.1	Searching the Directory	6–27
6.8.2	Step 1: Export Object Identifiers	6–29
6.8.3	Step 2: Declare Local Variables	6–30
6.8.4	Step 3: Build Public Objects for the name Parameter to ds_search()	6–31
6.8.5	Step 4: Specify the Portion of the DIT To Be Searched	6–31
6.8.6	Step 5: Create a Filter	6–31

6.8.7	Step 6: Create an Entry-Information-Selection Parameter	6-34
6.8.8	Step 7: Perform the Search Operation	6-34
6.9	Directory Modify Operations	6-37
6.9.1	Modifying Directory Entries	6-37
6.9.2	Step 1: Export Object Identifiers for Required Directory Classes and Attributes	6-39
6.9.3	Step 2: Declare Local Variables	6-39
6.9.4	Step 3: Build Public Objects	6-40
6.9.5	Step 4: Create Descriptor Lists for Attributes	6-42
6.9.6	Step 5: Perform the Operations	6-44
6.10	Return Codes	6-47
Chapter 7.	Sample Application Programs	7-1
7.1	General Programming Guidelines	7-2
7.2	The example.c Program	7-2
7.2.1	The example.c Code	7-5
7.2.2	Error Handling.	7-12
7.3	The acl.c Program	7-14
7.3.1	The acl.c Code.	7-18
7.3.2	The acl.h Header File	7-35
7.3.3	The acl.h Code	7-36
7.4	The teldir.c Program	7-45
7.4.1	Predefined Static Public Objects	7-46
7.4.2	Partially Defined Static Public Objects	7-47
7.4.3	Dynamically Defined Public Objects	7-49
7.4.4	Main Program Procedural Steps	7-50
7.4.5	The teldir.c Code	7-51
Chapter 8.	Using Threads With The XDS/XOM API	8-1
8.1	Overview of Sample Threads Program	8-3
8.1.1	User Interface	8-4
8.1.2	Input File Format	8-5
8.1.3	Program Output	8-5
8.1.4	Prerequisites	8-6
8.2	Description of Sample Program	8-6
8.3	Detailed Description of Thread Specifics	8-8
8.4	The thradd.c Code	8-10
8.5	The thradd.h Header File	8-23

Chapter 9. XDS/XOM Convenience Routines	9-1
9.1 String Handling	9-2
9.1.1 Strings Representing GDS Attribute Information . . .	9-3
9.1.2 Strings Representing Structured GDS Attribute Information	9-3
9.1.3 Strings Representing a Structured GDS Attribute Value	9-6
9.1.4 Strings Representing a Distinguished Name . . .	9-7
9.1.5 Strings Representing Expressions	9-8
9.2 The acl2.c Program	9-10
9.2.1 The acl2.c Code	9-12
9.2.2 The acl2.h Header File	9-36
9.2.3 Example Strings	9-38

Part 4. XDS/XOM Supplementary Information

Chapter 10. XDS Interface Description	10-1
10.1 XDS Conformance to Standards	10-2
10.2 The XDS Functions	10-4
10.3 The XDS Negotiation Sequence	10-6
10.4 The session Parameter	10-6
10.5 The context Parameter	10-7
10.6 The XDS Function Arguments	10-8
10.6.1 Attribute and Attribute Value Assertion	10-8
10.6.2 The selection Parameter	10-10
10.6.3 The name Parameter	10-10
10.7 XDS Function Call Results	10-11
10.7.1 The invoke_id Parameter	10-11
10.7.2 The result Parameter	10-11
10.7.3 The DS_status Return Value	10-12
10.8 Synchronous Operations	10-12
10.9 Security and XDS	10-13
10.10 Other Features of the XDS Interface	10-13
10.10.1 Automatic Connection Management	10-14
10.10.2 Automatic Continuation and Referral Handling . . .	10-14
10.10.3 Abandoning Operations	10-15
Chapter 11. XDS Class Definitions	11-1

11.1	Introduction to OM Classes	11-1
11.2	XDS Errors	11-2
11.3	OM Class Hierarchy	11-4
11.4	DS_C_ABANDON_FAILED	11-6
11.5	DS_C_ACCESS_POINT	11-7
11.6	DS_C_ADDRESS	11-7
11.7	DS_C_ATTRIBUTE	11-8
11.8	DS_C_ATTRIBUTE_ERROR	11-9
11.9	DS_C_ATTRIBUTE_LIST	11-10
11.10	DS_C_ATTRIBUTE_PROBLEM	11-10
11.11	DS_C_AVA	11-11
11.12	DS_C_COMMON_RESULTS	11-12
11.13	DS_C_COMMUNICATIONS_ERROR	11-13
11.14	DS_C_COMPARE_RESULT	11-13
11.15	DS_C_CONTEXT	11-14
11.16	DS_C_CONTINUATION_REF	11-19
11.17	DS_C_DS_DN	11-20
11.18	DS_C_DS_RDN	11-21
11.19	DS_C_ENTRY_INFO	11-22
11.20	DS_C_ENTRY_INFO_SELECTION	11-23
11.21	DS_C_ENTRY_MOD	11-24
11.22	DS_C_ENTRY_MOD_LIST	11-25
11.23	DS_C_ERROR	11-25
11.24	DS_C_EXT	11-28
11.25	DS_C_FILTER	11-29
11.26	DS_C_FILTER_ITEM	11-31
11.27	DS_C_LIBRARY_ERROR	11-33
11.28	DS_C_LIST_INFO	11-34
11.29	DS_C_LIST_INFO_ITEM	11-35
11.30	DS_C_LIST_RESULT	11-36
11.31	DS_C_NAME	11-37
11.32	DS_C_NAME_ERROR	11-38

11.33	DS_C_OPERATION_PROGRESS	11-39
11.34	DS_C_PARTIAL_OUTCOME_QUAL	11-40
11.35	DS_C_PRESENTATION_ADDRESS	11-41
11.36	DS_C_READ_RESULT	11-42
11.37	DS_C_REFERRAL	11-43
11.38	DS_C_RELATIVE_NAME	11-43
11.39	DS_C_SEARCH_INFO	11-44
11.40	DS_C_SEARCH_RESULT	11-45
11.41	DS_C_SECURITY_ERROR	11-46
11.42	DS_C_SERVICE_ERROR	11-46
11.43	DS_C_SESSION	11-47
11.44	DS_C_SYSTEM_ERROR	11-49
11.45	DS_C_UPDATE_ERROR	11-50
Chapter 12.	Basic Directory Contents Package	12-1
12.1	Selected Attribute Types	12-2
12.2	Selected Object Classes	12-13
12.3	OM Class Hierarchy	12-15
12.4	DS_C_FACSIMILE_PHONE_NBR	12-16
12.5	DS_C_POSTAL_ADDRESS	12-17
12.6	DS_C_SEARCH_CRITERION	12-17
12.7	DS_C_SEARCH_GUIDE	12-19
12.8	DS_C_TELETEX_TERM_IDENT	12-20
12.9	DS_C_TELEX_NBR	12-21
Chapter 13.	Strong Authentication Package	13-1
13.1	SAP Attribute Types	13-2
13.2	SAP Object Classes	13-4
13.3	OM Class Hierarchy	13-5
13.4	DS_C_ALGORITHM_IDENT	13-5
13.5	DS_C_CERT	13-6
13.6	DS_C_CERT_LIST	13-8
13.7	DS_C_CERT_PAIR	13-9

13.8	DS_C_CERT_SUBLIST	13-9
13.9	DS_C_SIGNATURE	13-10
Chapter 14.	MHS Directory User Package	14-1
14.1	MDUP Attribute Types	14-2
14.2	MDUP Object Classes	14-6
14.3	MDUP OM Class Hierarchy	14-6
14.4	MH_C_OR_ADDRESS	14-7
14.5	MH_C_OR_NAME	14-27
14.6	DS_C_DL_SUBMIT_PERMS	14-28
Chapter 15.	GDS Package	15-1
15.1	GDSP Attribute Types	15-2
15.2	GDSP Object Classes	15-6
15.3	GDSP OM Class Hierarchy	15-7
15.4	DSX_C_GDS_ACL	15-8
15.5	DSX_C_GDS_ACL_ITEM	15-9
15.6	DSX_C_GDS_CONTEXT	15-10
15.7	DSX_C_GDS_SESSION	15-15
Chapter 16.	Distributed Management Environment Support	16-1
16.1	DME Attribute Types	16-2
16.2	DME Object Classes	16-3
Chapter 17.	Information Syntaxes	17-1
17.1	Syntax Templates	17-1
17.2	Syntaxes	17-2
17.3	Strings	17-3
17.4	Representation of String Values	17-5
17.5	Relationship to ASN.1 Simple Types	17-5
17.6	Relationship to ASN.1 Useful Types	17-6
17.7	Relationship to ASN.1 Character String Types	17-6
17.8	Relationship to ASN.1 Type Constructors	17-7

Chapter 18. XOM Service Interface	18-1
18.1 Standards Conformance	18-1
18.2 XOM Data Types	18-2
18.2.1 OM_boolean	18-4
18.2.2 OM_descriptor	18-5
18.2.3 OM_enumeration	18-7
18.2.4 OM_exclusions	18-8
18.2.5 OM_integer	18-8
18.2.6 OM_modification	18-9
18.2.7 OM_object	18-9
18.2.8 OM_object_identifier	18-9
18.2.9 OM_private_object	18-12
18.2.10 OM_public_object	18-12
18.2.11 OM_return_code	18-13
18.2.12 OM_string	18-14
18.2.13 OM_syntax	18-15
18.2.14 OM_type	18-17
18.2.15 OM_type_list	18-17
18.2.16 OM_value	18-18
18.2.17 OM_value_length	18-19
18.2.18 OM_value_position	18-19
18.2.19 OM_workspace	18-20
18.3 XOM Functions	18-20
18.4 XOM Return Codes	18-23
Chapter 19. Object Management Package	19-1
19.1 Class Hierarchy	19-1
19.2 Class Definitions	19-2
19.2.1 OM_C_ENCODING	19-2
19.2.2 OM_C_EXTERNAL	19-4
19.2.3 OM_C_OBJECT	19-5
Index	Index-1

List of Figures

Figure 1–1. A Federated DCE Namespace	1–14
Figure 1–2. GDS Namespace Entries and Directory Objects	1–15
Figure 2–1. The Cell Namespace After Configuration	2–3
Figure 2–2. A Possible Namespace Structure	2–9
Figure 2–3. Valid Characters in CDS, GDS, and DNS Names	2–16
Figure 3–1. One Object Descriptor	3–4
Figure 3–2. A Complete Object Represented	3–5
Figure 3–3. A Three-Layer Compound Object	3–6
Figure 3–4. Directory Objects and XDS Interface Objects	3–8
Figure 3–5. Directory Objects and Namespace Entries	3–10
Figure 3–6. The DS_C_READ_RESULT Object Structure	3–28
Figure 3–7. The DS_C_ENTRY_INFO Object Structure	3–31
Figure 3–8. The DS_C_ATTRIBUTE Object Structure	3–34
Figure 3–9. The DS_C_ATTRIBUTE_LIST Object	3–48
Figure 3–10. DS_C_DS_DN Object Attributes	3–51
Figure 3–11. The DS_C_ENTRY_MOD_LIST Object	3–55
Figure 3–12. The DS_C_ENTRY_INFO_SELECTION Object	3–58
Figure 4–1. XDS: Interface to GDS and CDS	4–3
Figure 4–2. The Structure of the DIB	4–4
Figure 4–3. Object Identifiers	4–6
Figure 4–4. A Directory Entry Describing Organizational Person	4–8
Figure 4–5. A Distinguished Name in a Directory Information Tree	4–10
Figure 4–6. An Alias in the Directory Information Tree	4–13
Figure 4–7. A Subtree Populated by Aliases	4–13
Figure 4–8. SRT DIT Structure for the GDS Standard Schema	4–17

Figure 4–9. A Partial Representation of the Object Class Table	4–20
Figure 4–10. The Relationship Between Schemas and the DIT	4–25
Figure 4–11. The Relationship Between the DSA and the DUA	4–29
Figure 4–12. An Example of a Referral	4–31
Figure 4–13. An Example of Chaining	4–32
Figure 4–14. GDS Components	4–34
Figure 5–1. The Internal Structure of an OM Object.	5–3
Figure 5–2. Mapping the Class Definition of DS_C_ENTRY_INFO_SELECTION	5–6
Figure 5–3. A Representation of a Public Object By Using a Descriptor List.	5–12
Figure 5–4. A Descriptor List for the Public Object: country	5–14
Figure 5–5. The Distinguished Name of “Peter Piper” in the DIT	5–15
Figure 5–6. Building a Distinguished Name	5–18
Figure 5–7. A Simplified View of the Structure of a Distinguished Name	5–19
Figure 5–8. Client-Generated and Service-Generated Objects	5–21
Figure 5–9. The OM Class DS_C_ENTRY_INFO_SELECTION.	5–24
Figure 5–10. Comparison of Two Classes With/Without an Abstract OM Class	5–26
Figure 5–11. Complete Description of Concrete OM Class DS_C_ATTRIBUTE	5–29
Figure 5–12. Data Type OM_descriptor_struct	5–44
Figure 5–13. Initializing Descriptors	5–45
Figure 5–14. An Object and a Subordinate Object	5–46
Figure 5–15. The Read Result	5–53
Figure 5–16. Extracting Information Using om_get()	5–54
Figure 6–1. Output from ds_read(): DS_C_READ_RESULT	6–25
Figure 6–2. Subtree for the acl.h Sample Program	6–28
Figure 6–3. OM Class DS_C_FILTER	6–33
Figure 6–4. OM Class DS_C_SEARCH_RESULT	6–36
Figure 6–5. A Sample Directory Tree	6–38
Figure 6–6. OM Class DS_C_LIST_RESULT	6–46
Figure 7–1. Entries With User Credentials Added to the Directory Tree	7–16
Figure 8–1. Issuing XDS/XOM Calls from Within Different Threads	8–2
Figure 8–2. Program Flow for the thradd Sample Program	8–7
Figure 18–1. OM_String Elements	18–15

List of Tables

Table 2–1. Metacharacters and Their Meaning	2–17
Table 2–2. Summary of CDS, GDS, and DNS Characteristics	2–18
Table 2–3. Maximum Sizes of Directory Service Names	2–21
Table 2–4. T61 Syntax	2–22
Table 2–5. Combinations of Diacritical Characters and Basic Letters	2–23
Table 3–1. Directory Service Functions With Their Required Input Objects	3–44
Table 3–2. CDS Attributes to OM Syntax Translation	3–60
Table 3–3. OM Syntax to CDS Data Types Translation	3–61
Table 3–4. CDS Data Types to OM Syntax Translation	3–62
Table 4–1. Object Identifiers for Selected Attribute Types	4–7
Table 4–2. Structure Rule Table Entries	4–15
Table 4–3. Object Class Table Entries	4–18
Table 4–4. Object Identifiers for Selected Directory Classes	4–21
Table 4–5. Attribute Table Entries	4–23
Table 4–6. Syntax for the Simple ASN.1 Types	4–27
Table 4–7. Cache Attributes: Read Cache First	4–36
Table 4–8. Cache Attributes: Read DSA First	4–36
Table 4–9. Cache Attributes: Read DSA Only	4–36
Table 4–10. Cache Attributes: DSX_USED_SA is OM_FALSE	4–37
Table 4–11. Cache Attributes: DSX_DUA_CACHE is OM_FALSE	4–37
Table 4–12. Cache Attributes: Error	4–37
Table 4–13. XDS_LOG Values	4–42
Table 5–1. C Naming Conventions for XDS	5–9
Table 5–2. C Naming Conventions for XOM	5–9
Table 5–3. Comparison of Private and Public Objects	5–22

Table 5–4. Description of an OM Attribute By Using Syntax Enum(*)	5–40
Table 5–5. Description of an OM Attribute By Using Syntax Object(*)	5–41
Table 6–1. Representation of Values for Selected Attribute Types	6–12
Table 6–2. Mapping of XDS API Functions to the Abstract Services	6–17
Table 10–1. The XDS Interface Functions	10–5
Table 11–1. OM Attributes of DS_C_ACCESS_POINT	11–7
Table 11–2. OM Attributes of DS_C_ATTRIBUTE	11–8
Table 11–3. OM Attributes of DS_C_ATTRIBUTE_ERROR.	11–9
Table 11–4. OM Attribute of DS_C_ATTRIBUTE_LIST	11–10
Table 11–5. OM Attributes of DS_C_ATTRIBUTE_PROBLEM	11–11
Table 11–6. OM Attributes of DS_C_COMMON_RESULTS	11–12
Table 11–7. OM Attributes of DS_C_COMPARE_RESULT	11–13
Table 11–8. OM Attributes of DS_C_CONTEXT	11–14
Table 11–9. OM Attributes of DS_C_CONTINUATION_REF	11–19
Table 11–10. OM Attribute of DS_C_DS_DN	11–21
Table 11–11. OM Attribute of DS_C_DS_RDN	11–21
Table 11–12. OM Attributes of DS_C_ENTRY_INFO	11–22
Table 11–13. OM Attributes of DS_C_ENTRY_INFO_SELECTION	11–23
Table 11–14. OM Attribute of DS_C_ENTRY_MOD	11–24
Table 11–15. OM Attribute of DS_C_ENTRY_MOD_LIST	11–25
Table 11–16. OM Attribute of DS_C_ERROR	11–26
Table 11–17. OM Attributes of DS_C_EXT	11–29
Table 11–18. OM Attributes of DS_C_FILTER	11–30
Table 11–19. OM Attributes of DS_C_FILTER_ITEM	11–31
Table 11–20. OM Attributes of DS_C_LIST_INFO	11–35
Table 11–21. OM Attributes of DS_C_LIST_INFO_ITEM	11–36
Table 11–22. OM Attributes of DS_C_LIST_RESULT	11–37
Table 11–23. OM Attribute of DS_C_NAME_ERROR	11–38
Table 11–24. OM Attributes of DS_C_OPERATION_PROGRESS	11–39
Table 11–25. OM Attributes of a DS_C_PARTIAL_OUTCOME_QUAL	11–40
Table 11–26. OM Attributes of DS_C_PRESENTATION_ADDRESS	11–42
Table 11–27. OM Attribute of DS_C_READ_RESULT	11–43

Table 11–28. OM Attributes of DS_C_SEARCH_INFO	11–44
Table 11–29. OM Attributes of DS_C_SEARCH_RESULT	11–45
Table 11–30. OM Attributes of DS_C_SESSION	11–48
Table 12–1. Object Identifiers for Selected Attribute Types	12–4
Table 12–2. Representation of Values for Selected Attribute Types	12–5
Table 12–3. Object Identifiers for Selected Object Classes	12–14
Table 12–4. OM Attributes of DS_C_FACSIMILE_PHONE_NBR	12–16
Table 12–5. OM Attribute of DS_C_POSTAL_ADDRESS	12–17
Table 12–6. OM Attributes of DS_C_SEARCH_CRITERION	12–18
Table 12–7. OM Attributes of DS_C_SEARCH_GUIDE	12–19
Table 12–8. OM Attributes of DS_C_TELETEX_TERM_IDENT	12–20
Table 12–9. OM Attributes of DS_C_TELEX_NBR	12–21
Table 13–1. Object Identifiers for SAP Attribute Types	13–3
Table 13–2. Representation of Values for SAP Attribute Types	13–3
Table 13–3. Object Identifiers for SAP Object Classes	13–4
Table 13–4. OM Attributes of DS_C_ALGORITHM_IDENT	13–6
Table 13–5. OM Attributes of DS_C_CERT	13–7
Table 13–6. OM Attributes of DS_C_CERT_LIST	13–8
Table 13–7. OM Attributes of DS_C_CERT_PAIR	13–9
Table 13–8. OM Attributes of DS_C_CERT_SUBLIST	13–10
Table 13–9. OM Attributes of DS_C_SIGNATURE	13–11
Table 14–1. Object Identifiers for MDUP Attribute Types	14–2
Table 14–2. Representation of Values for MDUP Attribute Types	14–3
Table 14–3. Object Identifiers for MDUP Object Classes	14–6
Table 14–4. Attributes Specific to MH_C_OR_ADDRESS	14–7
Table 14–5. Forms of Originator/Recipient Address	14–22
Table 14–6. Attribute Specific to MH_C_OR_NAME	14–27
Table 14–7. OM Attributes of DS_C_DL_SUBMIT_PERMS	14–28
Table 15–1. Object Identifiers for GDSP Attribute Types	15–2
Table 15–2. Representation of Values for GDSP Attribute Types	15–4
Table 15–3. Object Identifier for GDSP Object Classes	15–7
Table 15–4. OM Attributes of DSX_C_GDS_ACL	15–8

Table 15–5. OM Attributes of DSX_C_GDS_ACL_ITEM	15–9
Table 15–6. OM Attributes of DSX_C_GDS_CONTEXT	15–10
Table 15–7. Default DSX_C_GDS_CONTEXT	15–14
Table 15–8. OM Attributes of DSX_C_GDS_SESSION	15–15
Table 15–9. Default DSX_C_GDS_SESSION	15–17
Table 16–1. Object Identifier for DME Attribute Type	16–2
Table 16–2. Representation of Values for DME Attribute Types	16–3
Table 16–3. Object Identifier for DME Object Class	16–4
Table 17–1. String Syntax Identifiers	17–4
Table 17–2. Syntax for ASN.1 Simple Types	17–5
Table 17–3. Syntaxes for ASN.1 Useful Types	17–6
Table 17–4. Syntaxes for ASN.1 Character String Types	17–6
Table 17–5. Syntaxes for ASN.1 Type Constructors	17–7
Table 18–1. XOM Service Interface Data Types	18–2
Table 18–2. Assigning Meanings to Values	18–17
Table 18–3. XOM Service Interface Functions	18–20
Table 19–1. Attributes Specific to OM_C_ENCODING	19–2
Table 19–2. Attributes Specific to OM_C_EXTERNAL	19–4
Table 19–3. Attribute Specific to OM_C_OBJECT	19–6

Preface

The *OSF DCE Application Development Guide* provides information about how to program the application programming interfaces (APIs) provided for each OSF® Distributed Computing Environment (DCE) component.

Audience

This guide is written for application programmers with UNIX operating system and C language experience who want to develop and write applications to run on DCE.

Applicability

This revision applies to the OSF® DCE Release 1.2.2 offering and related updates. See your software license for details.

Purpose

The purpose of this guide is to assist programmers in developing applications that use DCE. After reading this guide, you should be able to program the Application Programming Interfaces provided for each DCE component.

Document Usage

The *OSF DCE Application Development Guide* consists of three books, as follows:

- *OSF DCE Application Development Guide—Introduction and Style Guide*
- *OSF DCE Application Development Guide—Core Components*
 - Part 1. DCE Facilities
 - Part 2. DCE Threads
 - Part 3. DCE Remote Procedure Call
 - Part 4. DCE Distributed Time Service
 - Part 5. DCE Security Service
- *OSF DCE Application Development Guide—Directory Services*
 - Part 1. DCE Directory Service
 - Part 2. CDS Application Programming
 - Part 3. GDS Application Programming
 - Part 4. XDS/XOM Supplementary Information

Related Documents

For additional information about the Distributed Computing Environment, refer to the following documents:

- *Introduction to OSF DCE*
- *OSF DCE Command Reference*

- *OSF DCE Application Development Reference*
- *OSF DCE Administration Guide*
- *OSF DCE DFS Administration Guide and Reference*
- *OSF DCE GDS Administration Guide and Reference*
- *OSF DCE/File-Access Administration Guide and Reference*
- *OSF DCE/File-Access User's Guide*
- *OSF DCE Problem Determination Guide*
- *OSF DCE Testing Guide*
- *OSF DCE/File-Access FVT User's Guide*
- *Application Environment Specification/Distributed Computing*
- *OSF DCE Technical Supplement*
- *OSF DCE Release Notes*

Typographic and Keying Conventions

This guide uses the following typographic conventions:

Bold	Bold words or characters represent system elements that you must use literally, such as commands, options, and pathnames.
<i>Italic</i>	<i>Italic</i> words or characters represent variable values that you must supply. <i>Italic</i> type is also used to introduce a new DCE term.
Constant width	Examples and information that the system displays appear in constant width typeface.
[]	Brackets enclose optional items in format and syntax descriptions.
{ }	Braces enclose a list from which you must choose an item in format and syntax descriptions.
	A vertical bar separates items in a list of choices.
< >	Angle brackets enclose the name of a key on the keyboard.

... Horizontal ellipsis points indicate that you can repeat the preceding item one or more times.

This guide uses the following keying conventions:

<Ctrl-*x*> or **^*x***

The notation **<Ctrl-*x*>** or **^*x*** followed by the name of a key indicates a control character sequence. For example, **Ctrl-C** means that you hold down the control key while pressing **<C>**.

<Return>

The notation **<Return>** refers to the key on your terminal or workstation that is labeled with the word Return or Enter, or with a left arrow.

Problem Reporting

If you have any problems with the software or documentation, please contact your software vendor's customer service department.

Pathnames of Directories and Files in DCE Documentation

For a list of the pathnames for directories and files referred to in this guide, see the *OSF DCE Administration Guide—Introduction* and the *OSF DCE Testing Guide*.

Part 1

DCE Directory Service

Chapter 1

DCE Directory Service Overview

This chapter provides an overview of the DCE Directory Service for application programmers. The chapter begins with a description of this guide. It then introduces DCE Directory Service concepts, following which the structure of DCE names and the DCE namespace are described. The chapter then provides an overview of the programming interfaces used to access the DCE Directory Service.

1.1 Introduction to This Guide

This guide describes how application developers can access the DCE Directory Service. From the application programmer's perspective, the directory service has three main parts: the DCE Cell Directory Service (CDS), the DCE Global Directory Service (GDS), and the X/Open Directory Service (XDS) and X/Open OSI-Abstract-Data Manipulation (XOM) programming interfaces. This is reflected in the organization of the book, as follows:

- Part 1. DCE Directory Service
- Part 2. CDS Application Programming

- Part 3. GDS Application Programming
- Part 4. XDS/XOM Supplementary Information

Parts 2 and 3 contain conceptual material on CDS and GDS with descriptions of programming tasks, including the use of programming interfaces. Chapters in each of these parts (Chapter 3 of Part 2 and Chapter 7 of Part 3) contain annotated source code for sample applications.

Part 4 consists mostly of tables of values for the data structures used by the XDS and XOM application interfaces, which are the interfaces used to directly access the directory service. These chapters supplement the reference pages for the XDS and XOM function calls, which are located in the *OSF DCE Application Development Reference*.

1.1.1 Use of This Guide

Before reading this guide, you should read the *Introduction to OSF DCE*. It contains overviews, along with illustrations, of all the DCE components and of DCE as a whole. Many concepts and details are explained in the *Introduction to OSF DCE* that are necessary to a full understanding of what is described here. Next, read this chapter in its entirety.

Determine whether you will be programming primarily in the CDS namespace or the GDS namespace and read Part 2 or Part 3 accordingly. At this point, you are ready to begin programming and should proceed to Part 4. The main purpose of Part 4 is to provide a convenient location to look up the details of object values and structures needed when writing code.

If you do not find the information you need in either this guide or the *OSF DCE Application Development Reference*, see the *OSF DCE Administration Guide* and the *OSF DCE Command Reference*. For example, information about the CDS as a separate component is found in the *OSF DCE Administration Guide*. Although the DCE Security Service is documented in the *OSF DCE Application Development Guide*, some information of interest to programmers (such as adding new principals to the registry database) is also found in the *OSF DCE Administration Guide*.

1.1.2 Directory Service Tools

Both CDS and GDS have commands that allow system administrators to inspect and alter the contents of the directory. This can be useful when developing applications that access the DCE namespace.

For information on the CDS control program (**cdsep**), see the *OSF DCE Administration Guide—Core Components*. For information on the CDS browser (**cdsbrowser**), which is a utility based on Motif[®] that allows you to inspect the CDS namespace, see the *OSF DCE Administration Guide—Core Components*.

For information on the GDS system administration commands **gdssysadm**, **gdsdirinfo**, **gdsditadm**, and **gdscacheadm**, see the *OSF DCE GDS Administration Guide and Reference*.

1.2 Using the DCE Directory Service

The DCE Directory Service can be used in many ways. It is used by the DCE services themselves to support the DCE environment. For example, cells are registered in the global part of the directory service, enabling users from different cells to share information and resources.

The directory service is also useful to DCE applications. The client and server sides of an application can use it to find each other's locations. The directory service can also be used to store information that must be made available in a globally accessible, well-known place.

For example, one DCE application could be a print service consisting of a client side application that makes requests for jobs to be printed, and a server-side application that prints jobs on an available printer. The directory service could be used as a central place where the print clients could look up the location of a print server. It could also be used to store information about printers; for example, what type of jobs a printer can accept and whether it is currently up or down and lightly or heavily loaded.

In some ways, a directory service can be used in the same way that a file system has traditionally been used; that is, for containing globally accessible information in

a well-known place. An example is the use of configuration information stored in files in a UNIX `/etc` directory.

However, the directory service differs in important ways. It can be replicated so that information is available even if one server goes down. Replicas can be kept automatically up-to-date so that, unlike multiple copies of a file on different machines, the information in the replicas of the directory service can be kept current without manual intervention.

The directory service can also provide security for data that is kept in a globally accessible place. It supports access control lists (ACLs) that provide fine-grained control over who is able to read, modify, create, and perform other operations on its data.

As you learn about the directory service and how to access it, think about the ways in which your application can best take advantage of the services it provides.

1.3 DCE Directory Service Concepts

This section provides a description of DCE Directory Service concepts that are important to application developers. Concepts that are specific to GDS are covered in more detail in Part 3. The following concepts are intended to convey general definitions that are applicable to the directory service as a whole rather than specific to a particular directory service component. For more detailed definitions, see the glossary in the *Introduction to OSF DCE*.

- DCE namespace

The DCE namespace is the collection of names in a DCE environment. It can be made up of several domains, in which different types of servers own the names in different parts of the namespace. Typically, there are two high-level, or global, domains to a DCE namespace: the GDS namespace and the Domain Name System (DNS) namespace. At the next level is the CDS namespace, with names contained in the cell's CDS server. A DCE environment always contains a cell namespace, which is implemented by CDS. Parts of the DCE namespace may not be contained in any of the directory services; for example, the DFS (Directory File Service) namespace, also called the *filespace*, contains the names of files and

directories in DFS, and the security namespace contains principals and groups contained in the security server.

The term *DCE namespace* is used when referring to names, but not the information associated with them. For example, it would include the name of a printer in the directory service, but not its associated location attribute, and it would include the name of a DFS file, but not its contents.

- Cell namespace

All of the names found in a single DCE cell constitute the cell's namespace. This includes names managed by the cell's CDS server and security server, names in the cell's DFS if it has one, and any other names that reside within a particular cell.

- Hierarchy

The DCE namespace is organized into a hierarchy; that is, each name except the global root has a parent node and may itself have child nodes or leaves. The leaves are called *objects* or *entries*, and, in the CDS and DFS namespace, the nodes are called *directories*.

- Directory

The word *directory* has two meanings, which can be differentiated by their context. The first is the node of a hierarchy as mentioned in the previous definition. The second is a collection of objects managed by a directory service.

- Directory service

A directory service is software that manages names and their associated attributes. A directory service can store information, be queried about information, and be requested to change information. DCE contains two different directory services: CDS and GDS. It also interacts with a third directory service, DNS, which is not part of DCE.

- Junction

A junction is a point in the DCE namespace where two domains meet. For example, the point where the DFS entries are *mounted* into a CDS namespace is a junction. DCE also has junctions between the global directory services and CDS, and between CDS and the DCE Security Service.

- Object

The word *object* can have two meanings, depending on the context. Sometimes it means an entry in a directory service. Sometimes it means a real object that an entry in a directory service describes, such as a printer. In the context of XDS/XOM, the requested data is returned to the application in one or more *interface objects*, which are data structures that the application can manipulate.

- Entry

An entry is a unit of information in a directory service. It consists of a name and associated attributes. For example, an entry could consist of the name of a printer, its capabilities, and its network address.

- Class

In GDS, each entry has a class associated with it. The class determines what type of entry it is and what attributes may be associated with it.

- Link

A link is one type of object class. This type of object is a pointer to another object; it is similar to a soft link in a UNIX file system. A CDS link is similar to a GDS alias.

- Attribute

If an object is like a complex data structure, then its attributes are analogous to the separate member fields within that structure. Some of an object's attributes may be of significance only to the directory service that manages it. With attributes such as these, a directory service implements objects that contain various kinds of data about the directory itself, thus enabling the service to organize the entries into a meaningful structure. For example, directory objects can contain attributes whose values are other directory objects (called *child directories* or *subdirectories*) in the directory. Or link objects can contain attributes whose values are the names and internal identifiers of other directory entries, making a link object's entry name an alias of the other object to which its attributes indirectly refer.

- Type

Every attribute is characterized as being of a certain type. The attribute is used to hold a certain kind of data, such as a zip code or the name of a cat. Entries can also be classified by type; for entries, the term used is *class*.

— Value

An attribute can have one or more values.

- Object identifier

Directory attributes are uniquely identified by object identifiers (OIDs), which are administered by the International Organization for Standardization (ISO). In GDS, OIDs are also used to identify object classes. When it creates new attribute types, an application is responsible for tagging them with new, properly allocated OIDs (see your directory service administrator for OID assignments). In CDS, attribute types are identified by strings, which can be representations of OIDs.

- Name

A DCE name corresponds to an entry in some service participating in the DCE namespace, usually a directory service.

— Global name

A global name is a name that contains a path through one of the global namespaces (GDS or DNS).

— Local name

A local name is a name that uses the cell prefix `/:` to indicate the cell name and therefore does not have a specific path through a global namespace. The entry for a local name is always contained in the local cell.

- Access control list

Access to DCE namespace entries is determined by lists of entities that are attached through the DCE Security Service to both the entries and the objects when they are created. The lists, called *access control lists* (ACLs), specify the privileges that an entity or group of entities has for the entry the ACL is associated with. The security service provides servers with authenticated identification of every entity that contacts them; it is then the server's responsibility to check the ACL attached to the object that the potential client wants to access, and perform

or refuse to perform the requested operation on the basis of what it finds there. The ACLs are checked using security service library routines.

Objects in the GDS namespace have ACLs associated with them, but they are not security service ACLs.

- Replication

The DCE Directory Service can keep *replicas* (copies) of its data on different servers. This means that, if one server is unavailable, clients can still obtain information from another server.

- Caching

Both the CDS and GDS components of the directory service support caching of data on the client machine. When a client requests a piece of data from the directory service for the first time, the information must be obtained over the network from a server. However, the data can then be *cached* (stored) on the local machine, and subsequent requests for the same data can be satisfied more quickly by looking in the local cache instead of sending a request over the network. You need to be aware of caching because in some cases you will want to bypass the cache to ensure that the data you obtain is as up-to-date as possible.

1.4 Structure of DCE Names

The following subsections describe the structure of the names found in a DCE environment. DCE names can consist of several different parts, which reflect the federated nature of the DCE namespace. A DCE name has some combination of the following elements. They must occur in this order, but most elements are optional.

- Prefix
- GDS cell name or DNS cell name
- GDS name or CDS name
- Junction
- Application name

A DCE name can be represented by a string that is a readable description of a specific entry in the DCE namespace. The name is a string consisting of a series of elements separated by / (slash). The elements are read from left to right. Each consecutive element adds further specificity to the entry being described, until finally one arrives at the rightmost element, which is the simple name of the entry itself. Thus, in appearance, DCE names are similar to UNIX filenames.

In the discussion that follows, a DCE name *element* is the single piece of a name string enclosed between a consecutive pair of slashes. For example, consider the following string:

/.../C=US/O=OSF/OU=DCE/hosts/abc/self

In it, the following two substrings are both elements:

O=OSF

abc

The entire name contains (counting the ... element) a total of seven elements.

In GDS, an element is called a *relative distinguished name* (RDN) and the entire name is called a *distinguished name* (DN). In the preceding example, the attribute type **O** stands for the Organization type OID, which is 2.5.4.10.

1.4.1 DCE Name Prefixes

The leftmost element of any valid DCE name is a root prefix. The appearance and meaning of each is as follows:

/... This is the *global root*. It signifies that the immediately following elements form the name of a global namespace entry. Usually, the entry's contents consist of binding information for a DCE cell (more specifically, for some CDS server in the cell), and the name of the global entry is the name of the cell.

/: This is the *cell root*. It is an alias for the global prefix plus the name of the local cell; that is, the cell in which the prefix is being used. It

signifies that the immediately following elements taken together form the name of a cell namespace entry.

/: This is the *filespace root*. It is an alias for the global prefix, the name of the local cell, and the DFS junction.

DCE also supports a junction into the security service namespace, but there is no alias for this junction.

A prefix by itself is a valid DCE name. For example, you can list the contents of the **/:** directory to see the top-level entries in the CDS namespace, and you can use a file system command to list the contents of the **/:** directory to see the top-level entries in the filespace.

1.4.2 Names of Cells

After the global root prefix, the next section of a DCE name contains the name of the cell in which the object's name resides. The name of a cell can be expressed as either a GDS name or a DNS name, depending on which global directory service (GDS or DNS) the cell is registered in. The following subsections provide examples.

1.4.2.1 GDS Cell Names

GDS elements always consist of a substring in which an abbreviation or acronym in capital letters is followed by a = (equal sign), which is followed by a string value. As you will learn in more detail in Chapter 2, these substrings represent pairs of attribute types and attribute values.

For example, consider the following global DCE name:

/.../C=DE/O=SNI/OU=DCE/subsys/druecker/docs

In it, the *attribute=value* form of the leftmost elements after the **/...** indicates that the global part of the name is a GDS namespace entry, and that it ends after the **OU=DCE** element; therefore, the rest of the name is in the **/.../C=DE/O=SNI/OU=DCE** cell.

1.4.2.2 DNS Cell Names

If DNS is used as the global directory, a global name has a form like the following:

`/.../cs.univ.edu/subsys/printers/docs`

where the single element

cs.univ.edu

is the cell name; that is, the cell's name in the DNS namespace. The DNS name consists of up to four domain names (depending on the name assigned to the cell), separated by dots.

1.4.2.3 Discovering Your Local Cell's Name

A DCE cell consists of the machines that are configured into it; each DCE machine belongs to one DCE cell. Your local cell is therefore the cell to which the machine you are using belongs. Depending on the DCE name you are using, you can access your own cell or other (foreign) cells. If the name you are accessing is global, then its cell is explicitly named. If the name begins with the local cell prefix, then you are accessing a name within your local cell. You can find out what cell you are in by calling the **dce_cf_get_cell_name()** function.

Using the global directory services, applications can access resources and services in foreign cells; however, applications most frequently use resources from their local cell. If this is not the case, the cell boundaries may not have been well chosen.

1.4.3 CDS Names

After the cell name or cell root prefix, the next part of a DCE name is often a CDS name. For example, consider the following name:

`/.../C=DE/O=SNI/OU=DCE/subsys/druecker/docs`

The CDS part of this name is

/subsys/druecker/docs

Another example is the name

/.../cs.univ.edu/subsys/printers/docs

In this name, the CDS part is

/subsys/printers/docs

The following strings show equivalent names that use the cell root prefix, assuming that the name is used from within the **/.../C=DE/O=SNI/OU=DCE** and **/.../cs.univ.edu** cells, respectively:

/./subsys/druecker/docs

/./subsys/printers/docs

1.4.4 GDS Names

Some names fall entirely in the GDS namespace. These names are pure X.500 (and therefore GDS) names, since each element consists of a type and an attribute. The entries for these names are contained in a GDS server. The following is an example of a pure GDS name:

/.../C=US/L=Cambridge/CN=Kilroy

1.4.5 Junctions in DCE Names

Some junctions are implied in a DCE name; others can be seen. There is an implied junction between the global prefix and either GDS or DNS. It occurs after the global prefix. The junction between either of the global namespaces and the local cell namespace is also implied. It occurs after the cell name. The junction between the local cell namespace and either the DFS namespace or the security namespace is shown by the symbol **/fs** or **/sec**, respectively. The following are examples of visible junctions in DCE names:


```
./:/fs/usr/snowpaws  
/.../dce.osf.org/sec/principal/ziggy
```

1.4.6 Application Names

The part of a DCE name that occurs after a junction into a DCE application is the application name. DFS and security names are the currently supported examples; in the future, application programmers may also be able to create junctions in the namespace.

DFS names occur after the DFS junction. They are typeless and resemble UNIX file system names. After the global and CDS parts of a DFS name have been resolved by the appropriate directory services, the rest of the DFS name is handled within DFS. In the equivalent examples that follow, **/usr/snowpaws** is the DFS part of the DCE name:

```
/.../dce.osf.org/fs/usr/snowpaws  
./:/fs/usr/snowpaws  
/:/usr/snowpaws
```

Security names are similar to DFS names; first the parts of the name within the DCE Directory Service are resolved, then the rest of the name is handled by the security service. The entry is contained in the security registry database. Consider the following:

```
./:/sec/principal/ziggy
```


In this example, the security part of the DCE name is **/principal/ziggy**.

1.5 The Federated DCE Namespace

The DCE namespace is a single hierarchy of names, but the names can be contained in many different services. Because several services cooperate to make the DCE namespace, it is a federated namespace.

Figure 1-1 shows a typical DCE namespace and the different services in which names reside.

Figure 1-1. A Federated DCE Namespace



The following sections describe the different domains of the DCE namespace.

1.5.1 The GDS Namespace

This section provides a brief overview of the main characteristics of the GDS namespace regarded apart from the XDS interface used to access it. More detailed information about GDS and XDS can be found in Part 3 and Part 4, respectively.

In a GDS name such as

/.../C=US/O=OSF/OU=DCE

the **C=US** and **O=OSF** elements do not refer to directory entries that are fundamentally different from the one represented by **OU=DCE**, unlike in CDS or the UNIX file system.

Thus, in the name string

/C=US/O=OSF/OU=DCE

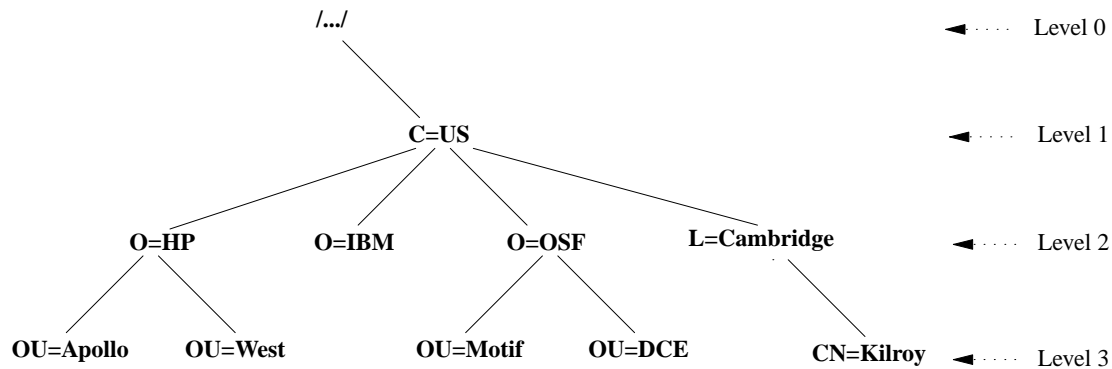
the element **C=US** refers to a one-level-down country entry whose value is **US**, then to a two-levels-down organization entry whose value is **OSF**, and then to a three-levels-down organization unit entry whose value is **DCE**. Concatenating these elements results in a valid path of entries from the directory root to the **DCE** entry. The entry itself is the namespace sign to a GDS directory object that contains binding information for the **/.../C=US/O=OSF/OU=DCE** cell.

1.5.1.1 An Example GDS Namespace

Figure 1-2 shows what a part of the DCE global namespace could look like. Levels in the tree of entries are numbered; the global root is at Level 0. The GDS structure rules as defined for DCE allow only country name entries at the next level under the root; organization name and locality name entries can exist at the level below a country name. An organizational unit name can be a child of an organizational name entry, and a common name can be a child of a locality name. The details of the GDS rules for the valid types and locations of entries in the directory tree can be found in the *OSF DCE Administration Guide*.

The object entry **/C=US/O=OSF/OU=DCE** belongs to the Organizational Unit class. One of the object's values is the CDS server binding information that is used to reach the cell from other DCE cells. The entire name is an attribute of the object that it refers to, as is the CDS server binding information that it contains.

Figure 1-2. GDS Namespace Entries and Directory Objects



1.5.1.2 The GDS Schema

The schema defines the shape and format of entries in the GDS directory. It contains four types of rules, which describe the following:

- The legal hierarchy of entries. What entries may be subordinate to other entries. These rules are what prevents, for example, countries from being subordinate to states.
- The allowable object classes, the mandatory and optional attributes of entries, and which attributes are the naming attributes.
- The allowable attribute types, associating a unique OID and an attribute syntax with each attribute type.
- The syntaxes of attributes that describe what attribute values look like, such as strings, numbers, or OIDs.

By installing the proper schema, an entry of any particular object class can have the two attributes needed to identify a cell. See the *OSF DCE Administration Guide* for a full description of how to set up a cell entry by using either GDS or DNS.

1.5.2 The CDS Namespace

The CDS namespace is the part of the DCE namespace that resides in the local cell's CDS. DCE itself is made up of components that, like the applications that use them, are distributed client/server applications. These components rely on CDS to make themselves available as services to DCE applications. This requires that the structure of the cell namespace be stable, known, and have parts that are not alterable by casual users or applications.

1.5.2.1 The CDS Schema

The cell namespace's hierarchy model is different from the GDS model, and the CDS rules do not enforce any particular model; CDS allows entries containing any kind of data to be created anywhere in the namespace. Thus, CDS offers a free-form namespace in which entries and directories can be organized as desired, and in which any entry or directory can contain any attributes. The CDS administrator

can create additional directories, and applications can add name entries as needed; applications *cannot* create CDS directory entries. Because of this, and because the cell namespace is so important to the operation of the cell, application developers and system administrators have more responsibility in planning and regulating their use of it.

The cell namespace has a structure similar to that of a UNIX file system. The CDS namespace is a tree of entries that grows from the root downward. The name entries are organized under directory entries, which can themselves be subentries of other directories. The cell root (represented by the prefix `/.`) can be thought of as the location you get when you dereference the cell's global name. New directories and new entries within the directories can be added anywhere in the tree, subject to the restrictions mentioned previously.

1.5.2.2 CDS Entries and CDS Attributes

There are three different kinds of CDS entries that are of significance to application programmers, as follows:

- Object
- Soft link
- Directory

The object entries are the most primitive form. These are where data is stored. Directory entries contain other entries (that is, can have children) just like UNIX file system directories. Soft link entries are essentially alias names for other directory or object entries. Only object entries can be created by applications; soft links and directories have to be created and manipulated with the **cdscp** command.

Thus, any CDS entry is defined as a directory, a soft link, or an object entry by the presence of a certain combination of attributes belonging to that kind of entry. You can use the **cdscp** command to get a display of all the attributes of any CDS entry.

The term *attribute* as applied to namespace entry objects has roughly the same meaning in CDS and GDS. The main difference is that CDS does not restrict or control the combinations of attributes attached to entries written in its namespace.

1.5.3 Other Namespaces

For information about names contained in the DFS namespace (the filespace) and the security namespace, refer to the chapters on those components in this guide.

1.6 Programming Interfaces to the DCE Directory Service

The following two subsections describe two programming interfaces for accessing the DCE Directory Service.

1.6.1 The XDS Interface

The main programming interface to all services within the directory service is XDS/XOM, as defined by X/Open. The calls correspond to the X.500 service requests, including Read, List (enumerate children), Search, Add Entry, Modify Entry, Modify RDN, and Remove Entry. XDS uses XOM to define and manipulate data structures (called *objects*) used as the parameters to these calls, and used to describe the directory entries manipulated by the calls. XOM is extremely flexible, but also somewhat complex. The interfaces are used in different ways, depending on which underlying directory service is being addressed. For example, CDS entries are typeless, but GDS entries are typed. This difference is reflected in the use of the interface.

1.6.2 The RPC Name Service Interface

The DCE Remote Procedure Call (RPC) facility supports an interface to the directory service that is specific to RPC and is layered on top of directory service interfaces; it is called the Name Service Independent (NSI) interface. NSI can manipulate three object classes — entries, groups, and profiles — which were created to store RPC binding information. NSI data is stored in CDS. Programming using this interface is discussed in the *OSF DCE Application Development Guide—Core Components* and *OSF DCE Application Development Guide—Introduction and Style Guide* volumes.

1.6.3 Namespace Junction Interfaces

For information about programming interfaces to names that occur in namespace junctions, see the documentation for that component.

Part 2

CDS Application Programming

Part 2 describes DCE Directory Service application programming in the CDS namespace. Chapter 2 describes the contents of the CDS namespace, where applications should put their data, and what the valid CDS characters and names are. Chapter 3 describes how to use the XDS programming interface to access data in the CDS namespace.

Chapter 2

Programming in the CDS Namespace

This chapter provides information about writing applications that use the XDS/XOM interface to access the portion of the DCE namespace contained in CDS.

The XDS/XOM interface provides generalized access to CDS. However, if you only need to use CDS to store information related to RPC (for example, storing the location of a server so that clients can find it), you should use the NSI interface of DCE RPC. NSI implements RPC-specific use of the namespace. For information on using RPC NSI, see the *OSF DCE Application Development Guide—Core Components*.

For information on the details of accessing the CDS namespace through the XDS/XOM interface, see Chapter 3.

2.1 Initial Cell Namespace Organization

The following subsections describe the organization of a cell's namespace after it has initially been configured. (For more information on configuring a cell, see the *OSF DCE Administration Guide*.)

Every DCE cell is set up at configuration with the basic namespace structure necessary for the other DCE components to be able to find each other and to be accessible to applications. The vital parts of the namespace are protected from being accessed by unauthorized entities by ACLs that are attached to the entries and directories.

Figure 2-1 shows what the cell namespace looks like after a cell has been configured and before any additional directories or entries have been added to it by system administrators or applications. In the figure, ovals represent directories, rectangles represent simple entries, circles represent soft links, and triangles represent namespace junctions.

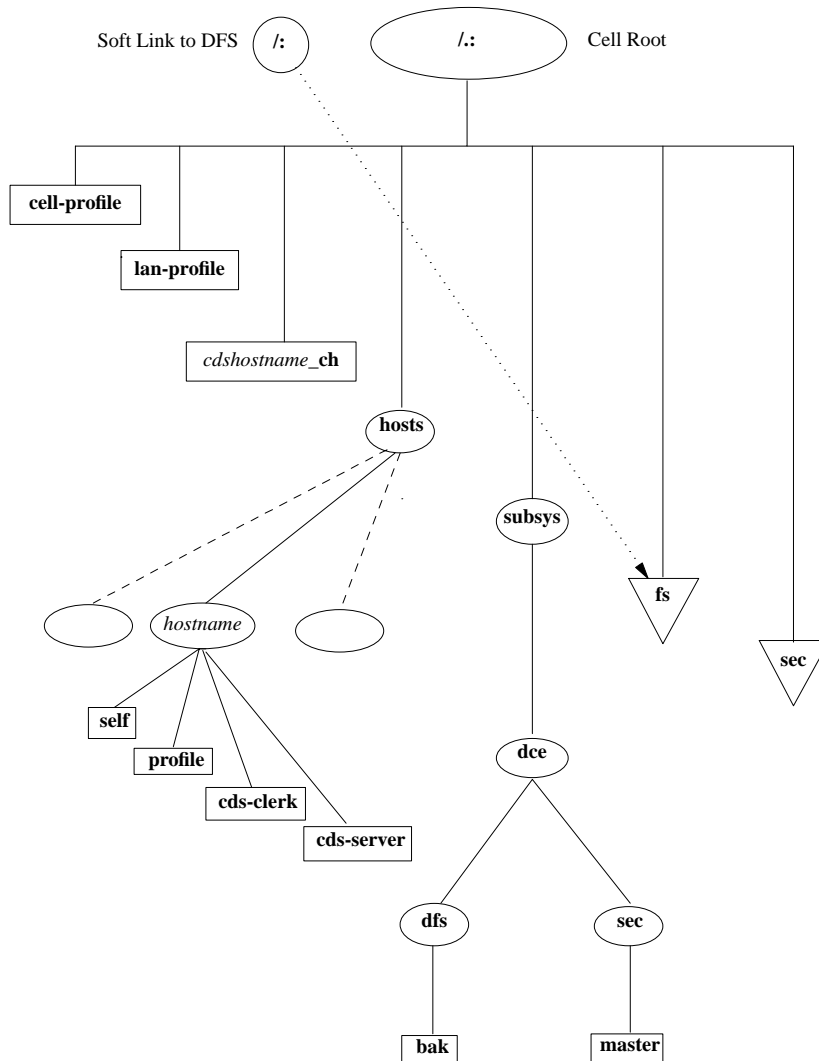
All of the simple entries shown in the figure are created for use with RPC NSI routines; that is, they all contain server-binding information and exist to enable clients to find servers. These are referred to as *RPC entries*.

Note that only the name entries (those in boxes) and junction entries (those in triangles) are RPC entries. The directories (entries indicated by ovals) are normal CDS directories.

Some of the namespace entries in the figure are intended to be used (if desired) directly by applications; namely, **./cell-profile**, **./lan-profile**, and, through the **./** soft link alias, **./fs**. The **self** and **profile** name entries under **hosts** also fall into this category. Others, such as those under **./subsys/dce**, are for the internal use of the DCE components themselves.

Each of the entries is explained in detail in the following subsections. See the *OSF DCE Administration Guide* for detailed information on the contents of the initial DCE cell namespace.

Figure 2–1. The Cell Namespace After Configuration



2.1.1 The Cell Profile

The `./cell-profile` entry is an RPC profile entry that contains the default list of namespace entries to be searched by clients trying to bind to certain basic services. An RPC profile is a class of namespace entry used by the RPC NSI routines. When a client imports bindings from such an entry, it imports, through the profile, from an ordered list of RPC entries containing appropriate bindings. The list of entries is keyed by their interface universal unique identifiers (UUIDs) so that only bindings to servers offering the interface sought by the client are returned. The entries listed in the profile exist independently in the namespace, and can be separately accessed in the normal way. The profile is simply a way of organizing clients' searches.

The main purpose of **cell-profile** is as a path of last resort for prospective clients. All other profile entries in the cell namespace are required to have the **cell-profile** entry in *their* entry lists so that, if a client exhausts a particular profile's list of entries, it tries those in **cell-profile**.

2.1.2 The LAN Profile

The `./lan-profile` entry is a local area network (LAN)-oriented default list of services' namespace entries that is used when servers' relative positions in the network topography are of importance to their prospective clients.

2.1.3 The CDS Clearinghouse

The `./cdshostname_ch` entry is the namespace entry for *cdshostname*'s clearinghouse, where *cdshostname* is the name of the host machine on which a CDS server is installed.

A *clearinghouse* is the database managed by a CDS server; it is where CDS directory replicas are physically stored. For more information about clearinghouses, see the *OSF DCE Administration Guide*. All clearinghouse namespace entries reside at the cell root, and there must be at least one in a DCE cell. The first clearinghouse's name must be in the form shown in Figure 2-1, but additional clearinghouses can be named as desired.

2.1.4 The Hosts Directory

The `./:/hosts` entry is a directory containing entries for all of the host machines in the cell. Each host has a separate directory under **hosts**; its directory has the same name as the host. Four entries are created in each host's directory:

- **self**

This entry contains bindings to the host's DCE daemon (**dced**), which is responsible for, among other things, dynamically resolving the partial bindings that it receives in incoming RPCs from clients attempting to reach servers resident on this host.

- **profile**

This entry is the default profile entry for the host. This profile contains in its list of entries at least the `./:/cell-profile` entry described in Section 2.1.1.

- **cds-clerk**

This entry contains bindings to the host's resident CDS clerk.

- **cds-server**

This entry contains bindings to a CDS server.

2.1.5 The Subsystems Directory

The `./:/subsys` entry is the directory for subsystems. Subdirectories below **subsys** are used to hold entries that contain location-independent information about services, particularly RPC binding information for servers.

The **dce** directory is created below `./:/subsys` at configuration. This directory contains directories for the DCE Security Service and Directory File Service (DFS) components. The functional difference between these two directories and the **fs** and **sec** junctions described in Section 2.1.7 is that the latter two entries are the access points for the components' special databases, whereas the directories under **subsys/dce** contain the services' binding information.

Subsystems that are added to DCE should place their system names in directories created beneath the `./:subsys` directory. Companies adding subsystems should conform to the convention of creating a unique directory below **subsys** by using their trademark as a directory name. Use these directories for storage of location-independent information about services. You should store server entries, groups and profiles for the entire cell in the directories below **subsys**. For example, International Air Freight-supplied subsystems should be placed in `./:subsys/IAF`.

2.1.6 The `./:` DFS Alias

The entry `./:` is created and set up as a soft link to the `./:fs` entry, which is the DFS database junction. The name `./:` is equivalent to `./:fs`. Note, however, that the name `./:` is well-known, whereas the name `./:fs` is not, so using `./:` makes an application more portable. A CDS soft link entry is an alias to some other CDS entry. A soft link is created through the **cdscp** command. The procedure is described in the *OSF DCE Administration Guide*.

2.1.7 The DFS and DCE Security Service Junctions

The `./:fs` entry is the DFS junction entry. This is the entry for a server that manages the DFS file location database.

The `./:sec` entry is the DCE Security Service junction entry. This is the entry for a server that manages the security service database (also called the *registry database*).

The `./:fs` and `./:sec` root entries in Figure 2-1 are junctions maintained by DCE components. The `./:sec` junction is the security service's namespace of principal identities and related information. The DFS's fileset location servers are reached through the `./:fs` entry, making `./:fs` effectively the entry point into the cell's distributed file system.

Note that `./:sec` and `./:fs` are both actually RPC group entries; the junctions are implemented by the servers whose entries are members of the group entries. (See the *OSF DCE Administration Guide* for further details on the security service and DFS junctions.)

2.2 Recommended Use of the CDS Namespace

CDS data is maintained in a loosely consistent manner. This means that, when the writeable copy of a replicated name is updated, the read-only copies may not be updated for some period of time, and applications reading from those nonsynchronized copies can receive stale data. This is in contrast to distributed databases, which use multiphase commit protocols that prevent readers from accessing potentially stale or inconsistent data while the writes are being propagated to all copies of the data. It is possible to specifically request data from the master copy, which is guaranteed to be up-to-date, but replication advantages are then lost. This should only be done when it is important to obtain current data.

2.2.1 Storing Data in CDS Entries

Some CDS entries may contain information that is immediately useful or meaningful to applications. Other entries may contain RPC information that enables application clients to reach application servers; that is, binding handles for servers, which are stored and retrieved using the RPC NSI routines. In either case, the entry's name should be a meaningful identification label for the information that the entry contains. This is because the namespace entry names are the main clue that users and applications have to the available set of resources in the DCE cell. Using the CDS namespace to store and retrieve binding information for distributed applications is the function of DCE RPC NSI.

In general, applications can store data into CDS object entry attributes in any XDS-expressible form they wish. Refer to Tables 3-3 and 3-4 in Chapter 3 for XDS-to-CDS data type translations. If you add new attributes to the `/opt/dcelocal/etc/cds_attributes` file, together with a meaningful CDS syntax (that is, a data type identifier) and name, then the attribute is displayed by `cdscp show` commands when executed on objects containing instances of that attribute.

There are three main questions to consider when using CDS to store data through application calls to XDS:

1. **Where in the CDS namespace should the new entries be placed?**

You are free to create new directories as long as you do not disturb the namespace's configured structure. Keep in mind that CDS directories must be created with the **cdscp** command; they cannot be created by applications.

Only two root-level directories are created at configuration: **hosts** and **subsys**. Applications should not add entries under the **hosts** tree; the host's default profile should instead be set up by a system administrator. The **subsys** directory is intended to be populated by directories (for example, **./subsys/dce**) in which the servers and other components of independent vendors' distributed products are accessed. Thus, the typical cell should usually have a series of root-level CDS directories that represent a reasonable division of categories.

One obvious division could be between entries intended for RPC use (that is, namespace entries that contain bindings for distributed applications), and entries that contain data of other kinds. On the other hand, it may be very useful to add supplementary data attributes to RPC entries in which various housekeeping or administrative data can be held. In this way, for example, performance data for printers can be associated with the print servers' name entries. You can either add new attributes to the server entries themselves, where, for example, the following is the name of a server entry that receives the new attributes:

./applications/printers/pr1

Or you can change the subtree structure so that new *entries* are added to hold the data, the server bindings are still held in separate wholly RPC entries, and each group of entries is located under a directory named for the printer:

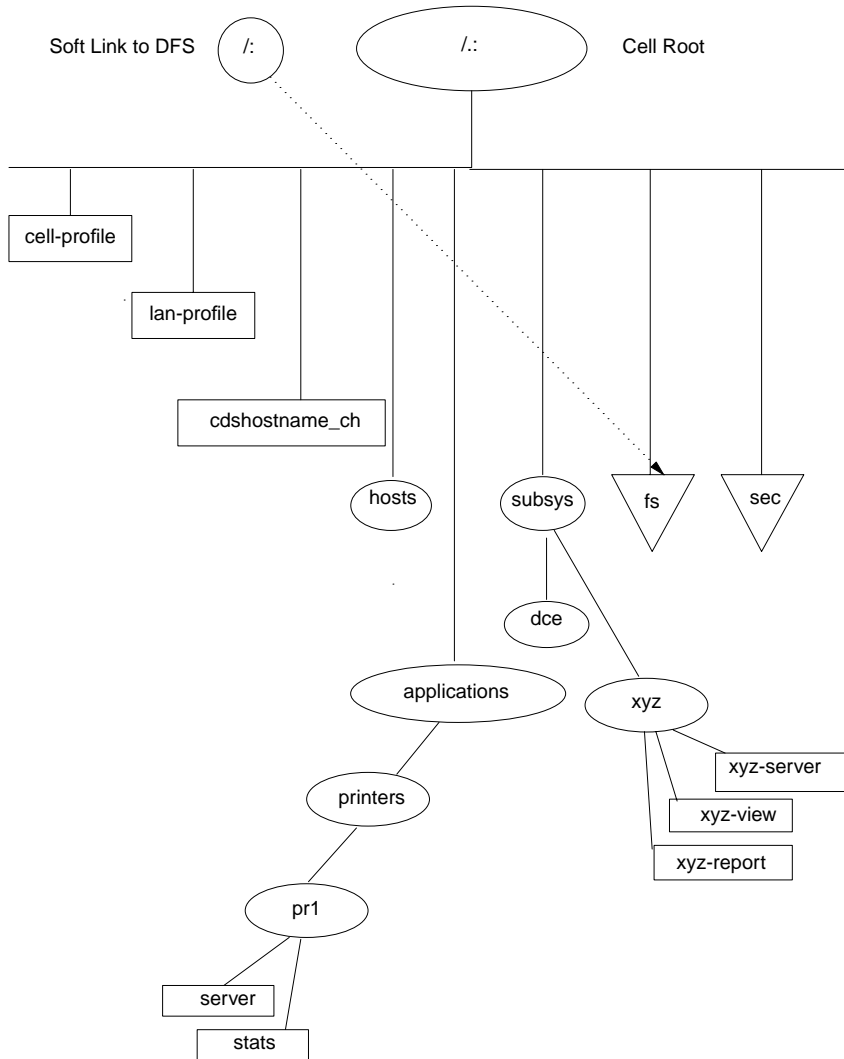
./applications/printers/pr1 — *directory*
./applications/printers/pr1/server — *server bindings*
./applications/printers/pr1/stats — *extra data*

In general, the same principals of logic and order that apply to the organization of a file system apply to the organization of a namespace. For example, server entries should *not* be created directly at the namespace root because this is the place for default profiles, clearinghouse entries, and directories.

Figure 2-2 illustrates some of the preceding suggestions, added to the initial configuration namespace structure shown in Figure 2-1. In Figure 2-2, the vendor of the **xyz** subsystem has set up an **xyz** directory under **./subsys** in which the system's servers are exported. This cell also has an **./applications** directory

in which the **printers** directory contains separate directories for each installed printer available on the system; the directory for **pr1** is illustrated in the figure. In the **pr1** directory, **server** is an RPC entry containing exported binding handles, and **stats** is an entry created and maintained through the XDS interface.

Figure 2–2. A Possible Namespace Structure



2. How should the entries be constructed?

Because CDS allows you to add as many attributes as you wish to an object entry, it is up to you to impose some restraint in doing this. In view of the XDS overhead involved in reading and writing single CDS attributes, it makes sense to combine multiple related attributes under single entries (that is, in the same directory object) where they can be read and written in single calls to **ds_read()** or **ds_modify_entry()**. This way, for example, you only have to create one interface input object (to pass to **ds_read()**) to read all the attributes, which you can do with one call to **ds_read()**. You can then separate out the returned subobjects that you are interested in and ignore the rest. Chapter 3 contains detailed discussions of XDS programming techniques.

In any case, you should define object types for use in applications so that namespace access operations can be standardized and kept efficient. A CDS object type consists of a specific set of attributes that belong to an object of that type, with no other attributes allowed. Note again that CDS, unlike GDS, does not force you to do things this way. You could theoretically have hundreds of CDS object entries, each of which would contain a different combination of attributes.

3. Should a directory or an entry be created?

When you consider adding information to the namespace, you can choose between creating a new directory, possibly with entries in it, or creating simply one or more entries. When making your decision, take into consideration the following:

- a. Directories cannot be created using XDS; they must be created using administrative commands. Directories are more expensive; they take up more space and take more time to access. However, they can contain entries and can therefore be used to organize information in the namespace.
- b. Entries can be created using XDS and they are cheaper to create and use than directories. However, they must be created in existing directories, and cannot themselves contain other entries.

2.2.2 Access Control for CDS Entries

Each object in the CDS namespace is automatically equipped with a mechanism by which access to it can be regulated by the object's owner or by another authority.

For each object, the mechanism is implemented by a separate list of the entities that can access the object in some way; for example, to read it, write to it, delete it, and so on. Associated with each entity in this list is a string that specifies which operations are allowed for that entity on the object. The object's list is automatically checked by CDS whenever any kind of access is attempted on that object by any entity. If the entity can be found in the object's list, and if the kind of access the entity intends is found among its permissions, then the operation is allowed to proceed by CDS; otherwise, it is not allowed.

DCE permission lists are called *access control lists* (ACLs). ACLs are one of the features of the DCE Security Service used by CDS. ACLs are used to test the entities' (that is, the principals') authorization to do things to the objects they propose to do them to. The authorization mechanism for all CDS objects is handled by CDS itself. All that users of the CDS namespace have to do is make sure that ACLs on the CDS objects that they create are set up with the appropriate permissions.

2.2.2.1 Creation of ACLs

Whenever you create a new entry in the CDS namespace, an ACL is created for it implicitly, and its initial list of entries and their permission sets are determined by the ACL templates associated with the CDS directory in which you create the entry.

Each CDS directory has the following two ACL templates associated with it:

- Initial Container

This template is used to generate the initial ACL for any directories created within the directory.

- Initial Object

This template is used to generate ACLs for entries created within the directory.

Like other CDS objects, each CDS directory also has its own ACL, generated from the parent directory's Initial Container template when the child directory is created. The Initial Container template also serves as a template for the child directories' own Initial Container templates.

2.2.2.2 Manipulating ACLs

There are two ways to manipulate ACLs: either through the **acl_edit** command (see the **acl_edit(8sec)** reference page) or through the DCE ACL application interface (see the **sec_acl_*(3sec)** reference pages).

2.2.2.3 Initializing ACLs

After creating a CDS directory by using the **cdscp** command, your first step is usually to run the **acl_edit** command to set up the new directory's ACLs the way you want them. (The new directory will have inherited its ACLs and its templates from the directory in which it was created, as explained in Section 2.2.2.1.) You may want to modify not only the directory's own ACLs, but also its two templates. To edit the latter, you can specify the **-ic** option (for the Initial Container template) or the **-io** option (for the Initial Object template); otherwise, you will edit the object ACL.

You can modify a directory's ACL templates from an application, assuming that you have control permission for the object, with the same combination of **sec_acl_lookup()** and **sec_acl_replace()** calls as for the object ACL. An option to these routines lets you specify which of the three possible ACLs on a directory object you want the call applied to. The ACLs themselves are in identical format.

The **-e** (entry) option to **acl_edit** can be used to make sure that you get the ACL for the specified namespace entry object, and not the ACL (if any) for the object that is *referenced* by the entry. This distinction has to be made clear to **acl_edit** because it finds the object (and hence the ACL) in question by looking it up in the namespace and binding to its ACL manager. Essentially, the **-e** option tells **acl_edit** whether it should bind to the CDS ACL manager (if the entry ACL is wanted), or to the manager responsible for the referenced object's ACL. This latter manager would be a part of the server application whose binding information the entry contained.

An example of such an ambiguous name would be a CDS clearinghouse entry, such as the *cdshostname_ch* entry discussed previously. With the **-e** option, you would edit the ACL on the namespace entry, as follows:

```
acl_edit -e ./cdshostname_ch
```

Without the **-e** option, you would edit the ACL on the clearinghouse itself, which you presumably do *not* want to do.

Similarly, there is a *bind_to_entry* parameter by which the caller of **sec_acl_bind()** can indicate whether the entry object's ACL or the ACL to which the entry refers is desired.

2.2.2.4 Namespace ACLs at Cell Configuration

The ACLs attached to the CDS namespace at configuration are described in *OSF DCE Administration Guide*. The following ACL permissions are defined for CDS objects. The single letter in parentheses for each item represents the DCE notation for that permission. These single letters are identical to the untokenized forms returned by **sec_acl_get_printstring()**.

- read (**r**)

This permission allows a principal to look up an object entry and view its attribute values.

- write (**w**)

This permission allows a principal to change an object's modifiable attributes, except for its ACLs.

- insert (**i**)

This permission allows a principal to create new entries in a CDS directory. It is used with directory entries only.

- delete (**d**)

This permission allows a principal to delete a name entry from the namespace.

- test (**t**)

This permission allows a principal to test whether an attribute of an object has a particular value, but does not permit it actually to see any of the attribute values (in other words, read permission for the object is not granted). The test permission allows an application to verify a particular CDS attribute's value without reading it.

- control (**c**)

This permission allows a principal to modify the entries in the object's ACL. The control permission is automatically granted to the creator of a CDS object.

- administer (**a**)

This permission allows a principal to issue **cdsep** commands that control the replication of directories. It is used with directory entries only.

Detailed instructions on the mechanics of setting up ACLs on CDS objects can be found in the *OSF DCE Administration Guide*.

For CDS directories, read and test permissions are sufficient to allow ordinary principals to access the directory and to read and test entries therein. Principals who you want to be able to add entries in a CDS directory should have insert permission for that directory. Entries created by the RPC NSI routines (for example, when a server exports bindings for the first time) are automatically set up with the correct permissions. However, if you are creating new CDS directories for RPC use, you should be sure to grant prospective user principals insert permission to the directory so that servers can create entries when they export their bindings. A general list of the permissions required for the various RPC NSI operations can be found in the **rpc_intro(3rpc)** and **rpc_ns_*(3rpc)** (RPC NSI) reference pages.

Note that CDS names do not behave the same way as file system names. A principal does not need to have access to an entire entry name path in order to have access to an entry at the end of that path. For example, a principal can be granted read access to the following entry:

./:/applications/utilities/pr2

and yet not have read access to the **utilities** directory itself.

2.3 Valid Characters and Naming Rules for CDS

The following subsections discuss the valid character sets for DCE Directory Service names as used by CDS interfaces. They also explain some characters that have

special meaning and describe some restrictions and rules regarding case matching, syntax, and size limits.

The use of names in DCE often involves more than one directory service. For example, CDS interacts with either GDS or DNS to find names outside the local cell.

Figure 2-3 details the valid characters in CDS names, and the valid characters in GDS and DNS names as used by CDS interfaces.

Note: Because CDS, GDS, and DNS all have their own valid character sets and syntax rules, the best way to avoid problems is to keep names short and simple, consisting of a minimal set of characters common to all three services. The recommended set is the letters A to Z, a to z, and the digits 0 to 9. In addition to making directory service interoperations easier, use of this subset decreases the probability that users in a heterogeneous hardware and software environment will encounter problems creating and using names.

Although spaces are valid in both CDS and GDS names, a CDS simple name containing a space must be enclosed in "" (double quotes) when you enter it through the CDS control program. Additional interface-specific rules are documented in the modules where they apply.

Figure 2–3. Valid Characters in CDS, GDS, and DNS Names

SP	0	@	P	'	p
!	1	A	Q	a	q
"	2	B	R	b	r
#	3	C	S	c	s
\$	4	D	T	d	t
%	5	E	U	e	u
&	6	F	V	f	v
'	7	G	W	g	w
(8	H	X	h	x
)	9	I	Y	i	y
*	:	J	Z	j	z
+	;	K	[k	{
,	<	L	\	l	
-	=	M]	m	}
.	>	N	^	n	~
/	?	O	_	o	

Key: ☐ Valid in CDS, GDS, and DNS names
☐ Valid only in CDS and GDS names
☐ Valid only in CDS names

2.3.1 Metacharacters

Certain characters have special meaning to the directory services; these are known as *metacharacters*. Table 2-1 lists and explains the CDS, GDS, and DNS metacharacters.

Table 2–1. Metacharacters and Their Meaning

Directory Service	Character	Meaning
CDS	/	Separates elements of a name (simple names).
	*	When used in the rightmost simple name of a name entered in a cdscp show or list command, acts as a wildcard, matching zero or more characters.
	?	When used in the rightmost simple name of a name entered in a cdscp show or list command, acts as a wildcard, matching exactly one character.
	\	Used where necessary in front of * (asterisk) or ? (question mark) to escape the character (indicates that the following character is not a metacharacter).
GDS	/	Separates RDNs.
	,	Separates multiple attribute type/value pairs (attribute value assertions) within an RDN.
	=	Separates an attribute type and value in an attribute value assertion.
	\	Used in front of / (slash), , (comma), or = (equal sign) to escape the character (indicates that the following character is not a metacharacter).
DNS	.	Separates elements of a name.

Some metacharacters are not permitted as normal characters within a name. For example, a \ (backslash) cannot be used as anything but an escape character in GDS. You can use other metacharacters as normal characters in a name, provided that you escape them with the backslash metacharacter.

2.3.2 Additional Rules

Table 2-2 summarizes major points to remember about CDS, GDS, and DNS character sets, metacharacters, restrictions, case-matching rules, internal storage of data, and ordering of elements in a name. For additional details, see the documentation for each technology.

Table 2-2. Summary of CDS, GDS, and DNS Characteristics

Characteristic	CDS	GDS	DNS
Character Set	a to z, A to Z, 0 to 9 plus space and special characters shown in Figure 2-3	a to z, A to Z, 0 to 9 plus . : , ' + - = () ? / and space	a to z, A to Z, 0 to 9 plus . -
Metacharacters	/ * ? \	/ , = \	.

Characteristic	CDS	GDS	DNS
Restrictions	<p>Simple names cannot contain slashes.</p> <p>The first simple name following the global cell name (or <code>/:</code> prefix) cannot contain an equal sign.</p> <p>When entering a name as part of a cdscp show or list command, you must use a backslash to escape any asterisk or question mark character in the rightmost simple name. Otherwise, the character is interpreted as a wildcard.</p>	<p>Relative distinguished names cannot begin or end with a slash.</p> <p>Attribute types must begin with an alphabetic character, can contain only alphanumerics, and cannot contain spaces.</p> <p>An alternate method of specifying attribute types is by object identifier, a sequence of digits separated by <code>.</code> (dots).</p> <p>You must use backslash to escape a slash, a comma, and an equal sign when using them as anything other than metacharacters.</p> <p>Multiple consecutive unescaped occurrences of slashes, commas, equal signs and backslashes are not allowed.</p> <p>Each attribute value assertion contains exactly one unescaped equal sign.</p>	<p>The first character must be alphabetic.</p> <p>The first and last characters cannot be <code>.</code> (dot) or <code>-</code> (dash).</p> <p>Cell names in DNS must contain at least one dot; they must be more than one level deep.</p>

Characteristic	CDS	GDS	DNS
Case-Matching Rules	Case exact	Attribute types are matched case insensitive. The case-matching rule for an attribute value can be case exact or case insensitive, depending on the rule defined for its type at the DSA.	Case insensitive
Internal Representation	Case exact	Depends on the case-matching rule defined at DSA. If the rule says case insensitive, alphabetic characters are converted to all lowercase characters. Spaces are removed regardless of the case-matching rule.	Alphabetic characters are converted to all lowercase characters.
Ordering of Name Elements	Big endian (left to right from root to lower-level names).	Big endian (left to right from root to lower-level names).	Little endian (right to left from root to lower-level names).

2.3.3 Maximum Name Sizes

Table 2-3 lists the maximum sizes for directory service names. Note that the limits are implementation specific, not architectural.

Table 2–3. Maximum Sizes of Directory Service Names

Name Type	Maximum Characters (Size)
CDS simple name (character string between two slashes)	254
CDS full name (including global or local prefix, cell name, and slashes separating simple names)	1023
GDS relative distinguished name	64
GDS distinguished name	1024
DNS relative name (character string between two dots)	64
DNS fully qualified name (sum of all relative names)	255

2.3.3.1 Valid Characters for GDS Attributes

This section describes the valid character sets for the GDS attributes.

The values of the country attributes are restricted to the ISO 3166 Alpha-2 code representation of country names. (For more information, see the *OSF DCE Administration Guide*.)

The character set for all other naming attributes is the T61 graphical character set. It is described in the next section.

2.3.3.2 T61 Syntax

Table 2-4 shows the T61 graphical character set.

Note: The 1) entry in the table indicates that it is not recommended that you use the codes in Column 2 Row 3, and Column 2 Row 4. Instead, use the appropriate code in Column A.

Table 2–4. T61 Syntax

	0	1	2	3	4	5	6	7	8	9	A	B	C	D	E	F
0			SP	0	@	P		p							Ω	K
1			!	1	A	Q	a	q			i	±	`		Æ	æ
2			”	2	B	R	b	r			¢	²	´			đ
3			1)	3	C	S	c	s			¶	³	^		a	—
4			1)	4	D	T	d	t			\$	x	~		H	h
5			%	5	E	U	e	u				μ	—			
6			&	6	F	V	f	v			#	⊥			ı	ij
7			'	7	G	W	g	w			§	•			L•	l•
8			(8	H	X	h	x				÷				
9)	9	I	Y	i	y								
A			*	:	J	Z	j	z							Œ	œ
B			+	;	K	[k				<<	>>				ß
C			,	<	L		l					¼	—			
D			—	=	M]	m					½	”		T	t
E				>	N		n					¾			η	η
F			/	?	0	—	o								'n	

The administration interface supports only characters smaller than 0x7e for names. The XDS application programming interface (API) supports the full T61 range as indicated in the preceding table.

Some T61 alphabetical characters have a 2-byte representation. For example, a lowercase letter **a** with an acute accent is represented by 0xc2 (the code for an acute accent) followed by 0x61 (the code for a lowercase **a**).

Only certain combinations of diacritical characters and basic letters are valid. They are shown in Table 2-5.

Table 2–5. Combinations of Diacritical Characters and Basic Letters

Name	Repr.	Code	Valid Basic Letters Following
grave accent	`	0xc1	a, A, e, E, i, I, o, O, u, U
acute accent	´	0xc2	a, A, c, C, e, E, g, i, I, l, L, n, N, o, O, r, R, s, S, u, U, y, Y, z, Z
circumflex accent	^	0xc3	a, A, c, C, e, E, g, G, h, H, i, I, j, J, o, O, s, S, u, U, w, W, y, Y
tilde	~	0xc4	a, A, i, I, n, N, o, O, u, U
macron	¯	0xc5	a, A, e, E, i, I, o, O, u, U
breve	˘	0xc6	a, A, g, G, u, U
dot above	·	0xc7	c, C, e, E, g, G, I, z, Z
umlaut	¨	0xc8	a, A, e, E, i, I, o, O, u, U, y, Y
ring	°	0xca	a, A, u, U
cedilla	¸	0xcb	c, C, G, k, K, l, L, n, N, r, R, s, S, t, T
double accent	¨	0xcd	o, O, u, U
ogonek	˛	0xce	a, A, e, E, i, I, u, U
caron	ˇ	0xcf	c, C, d, D, e, E, l, L, n, N, r, R, s, S, t, T, z, Z

The nonspacing underline (code 0xcc) must be followed by a Latin alphabetical character; that is, a basic letter (a to z or A to Z), or a valid diacritical combination.

2.4 Use of OIDs

OIDs are not seen by applications that restrict themselves to using only the RPC NSI routines (`rpc_ns_...()`), but these identifiers are important for applications that use the XDS interface to read entries directly or to create new attributes for use with namespace entries.

RPC makes use of only four different entry attributes in various application-specified or administrator-specified combinations. CDS, however, contains definitions for many more than these, which can be added by applications to RPC entries through the XDS interface. Attributes that already exist are already properly identified so applications that use these attributes do not have to concern themselves with the OIDs, except to the extent of making sure that they handle them properly.

Unlike UUIDs, OIDs are not generated by command or function call. They originate from ISO, which allocates them in hierarchically organized blocks to recipients. Each recipient, typically some organization, is then responsible for ensuring that the OIDs it receives are used uniquely.

For example, the following OID block was allocated to OSF by ISO:

1.3.22

OSF can therefore generate, for example, the following OID and allocate it to identify some DCE directory object:

1.3.22.1.1.4

(The OID **1.3.22.1.1.4** identifies the RPC profile entry object attribute.) OSF is responsible for making sure that **1.3.22.1.1.4** is not used to identify any other attribute. Thus, as long as all OIDs are generated only from within each owner's properly obtained block, and as long as each block owner makes sure that the OIDs generated within its block are properly used, each OID will always be a universally valid identifier for its associated value.

OIDs are encoded and internally represented as strings of hexadecimal digits, and comparisons of OIDs have to be performed as hexadecimal string comparisons (not as comparisons on NULL-terminated strings since OIDs can have NULL bytes as part of their value).

When applications have occasion to handle OIDs, they do so directly because the numbers do not change and should not be reused. However, for users' convenience, CDS also maintains a file (called **cds_attributes**, found in **/opt/dcelocal/etc**) that lists string equivalents for all the OIDs in use in a cell in entries like the following one:

1.3.22.1.1.4 RPC_Profile byte

This allows users to see **RPC_Profile** in output, rather than the meaningless string **1.3.22.1.1.4**. Further details about the **cds_attributes** file and OIDs can be found in the *OSF DCE Administration Guide*.

In summary, the procedure you should follow to create new attributes on CDS entries consists of three steps:

1. Request and receive from your locally designated authority the OIDs for the attributes you intend to create.
2. Update the **cds_attributes** file with the new attributes' OIDs and labels if you want your application to be able to use string name representations for OIDs in output.
3. Using XDS, write the routines to create, add, and access the attributes.

Your cell administrator should be able to provide you with a name and an OID. The *name* is a guaranteed-unique series of values for a global directory entry name. If the directory is GDS, the name is a series of type/value pairs, such as

C=US O=OSF

The cell administrator can also obtain an OID block. From this OID space, the administrator can assign you the OIDs you need for your application.

Note that there is no need for new OIDs in connection with cell names. The OIDs for Country Name and Organization Name are part of the X.500 standard implemented in GDS; only the values associated with the OIDs (the values of the objects) change from entry name to entry name. Instead, being able to generate new OIDs gives you the ability to invent and add new details to the directory itself. For example, you can create new kinds of CDS entry attributes by generating new OIDs to identify them. The same thing can be done to GDS, although the procedure is more complicated because it involves altering the directory schema.

Chapter 3

XDS and the DCE Cell Namespace

This chapter describes the use of the XDS programming interface when accessing the CDS namespace. The first section provides an introduction to using XDS in the CDS namespace. Section 3.2 describes XDS objects and how they are used to access CDS data. Section 3.3 provides a step-by-step procedure for writing an XDS program to access CDS. Section 3.4 provides examples of using the XOM interface to manipulate objects. Section 3.5 provides details of the structure of XDS/CDS objects. Finally, Section 3.6 provides translation tables between XDS and CDS for attributes and data types.

3.1 Introduction to Accessing CDS with XDS

Outside of the DCE cells and their separate namespaces is the global namespace in which the cell names themselves are entered, and where all intercell references are resolved. Two directory services participate in the global namespace. The first is the X.500-compliant GDS supplied with DCE. The second is DNS, with which DCE interacts, but is not a part of DCE.

The global and cell directory services are accessed implicitly by RPC applications using the NSI interface. GDS and CDS can also be accessed explicitly by using the XDS interface. With XDS, application programmers can gain access to GDS, a powerful general-purpose distributed database service, which can be used for many other things besides intercell binding. XDS can also be used to access the *cell* namespace directly, as this chapter describes.

An XDS application looks very different from the RPC-based DCE applications. This is partly because there is no dependency in XDS on the DCE RPC interface, although you can use both interfaces in the same application. Also, XDS is a generalized directory interface, oriented more toward performing large database operations than toward fine-tuning the contents of RPC entries. Nevertheless, XDS can be used as a general access mechanism on the CDS namespace.

3.1.1 Using the Reference Material in This Chapter

Complete descriptions of all the XDS and XOM functions used in CDS operations can be found in the *OSF DCE Application Development Reference*, which you should have beside you as you read through the examples in this chapter. In particular, refer to that manual for information about XDS error objects, which are not discussed in this chapter.

Complete descriptions for all objects required as *input* parameters by XDS functions when accessing a CDS namespace can be found in Section 3.5. Abbreviated definitions for these same objects can be found with all the others in Part 4. XOM functions are general-purpose utility routines that operate on objects of any class, and take the rest of their input in conventional form.

Slightly less detailed descriptions of the *output* objects you can expect to receive when accessing CDS through XDS are also given in Section 3.5. You do not have to construct objects of these classes yourself; you just have to know their general structure so that you can disassemble them using XOM routines.

No information is given in this chapter about the **DS_status** error objects that are returned by unsuccessful XDS functions; a description of all the subclasses of **DS_status** requires a chapter in itself. Code for a rudimentary general-purpose **DS_status**-handling routine can be found in the **teldir.c** XDS sample program in Chapter 7 of this guide.

3.1.2 What You Cannot Do with XDS

XDS allows you to perform general operations on CDS entry attributes, something which you cannot do through the DCE RPC NSI interface. However, there are certain things you cannot do to cell directory entries even through XDS:

- You cannot create or modify directory entries; the **ds_modify_rdn()** function does not work in a CDS namespace. These operations must be performed through the CDS control program (**cdscp**). For more information, see the *OSF DCE Command Reference*.
- You cannot perform XDS searches on the cell namespace; the XDS function **ds_search()** does not work. This is mainly because the CDS has no concept of a hierarchy of entry attributes, such as the X.500 schema. The **ds_compare()** function, however, does work.

3.1.3 Registering A Nonlocal Cell

If you are planning to use XDS to access the cell namespace in a one-cell environment (that is, your cell does not need to communicate with other DCE cells), you do not need to set up a cell entry in GDS for your cell because the XDS functions simply call the appropriate statically linked CDS routines to access the cell namespace. In these circumstances, XDS always tries to access the local cell when given an untyped (non-X.500) name.

For XDS to be able to access any nonlocal cell namespace, that cell must be registered (that is, have an entry) in the global namespace.

For information on setting up your cell name, see the *OSF DCE Administration Guide*.

3.2 XDS Objects

The XDS interface differs from the other DCE component interfaces in that it is *object oriented*. The following subsections explain two things: first, what object-oriented programming means in terms of using XDS; and second, how to use XDS to access CDS.

Imagine a generalized data structure that always has the same data *type*, and yet can contain any kind of data, and any amount of it. Functions could pass these structures back and forth in the same way all the time, and yet they could use the same structures for any kind of data they wanted to store or transfer. Such a data structure, if it existed, would be a true *object*. Programming language constructs allow interfaces to pretend that they use objects, although the realities of implementation are not usually so simple.

XDS is such an interface. For the most part, XDS functions neither accept nor return values in any form but as objects. The objects themselves are indeed always the same data type; namely, pointers to arrays of *object descriptor* (C **struct**) elements. Contained within these **OM_descriptor** element structures are unions that can actually accommodate all the different kinds of values an object can be called on to hold. In order to allow the interface to make sense of the unions, each **OM_descriptor** also contains a **syntax** field, which indicates the data type of that descriptor's union. There is also a record of what the descriptor's value actually is; for example, whether it is a name, a number, an address, a list, and so on. This information is held in the descriptor's **type** field.

These **OM_descriptor** elements, which are referred to as *object descriptors* or *descriptors*, are the basic building blocks of *all* XDS objects; every actual XDS object reduces to arrays of them. Each descriptor contains three items:

- A **type** field, which identifies the descriptor's value
- A **syntax** field, which indicates the data type of the **value** field
- The **value** field, which is a union

Figure 3-1 illustrates one such object descriptor.

Figure 3–1. One Object Descriptor

type:OM_CLASS
syntax:OID string
value:DS_C_DS_DN

Note that, from an abstract point of view, **syntax** is just an implementation detail. The scheme really consists only of a type/value pair. The **type** gives an identity to the object (something like CDS entry attribute, CDS entry name, or DUA access


point), and the **value** is some data associated with that identity, just as a variable has a name that gives meaning to the value it holds, and the value itself.

In order to make the representation of objects as logical and as flexible as possible, these two logical components of every object, **type** and **value**, are themselves each represented by separate object descriptors. Thus, the first element of every complete object descriptor array is a descriptor whose **type** field identifies its **value** field as containing the name of the kind (or *class*) of this object, and the **syntax** field indicates how that name **value** should be read. Next is usually one (or more, if the object is multivalued) object descriptor whose **type** field identifies its **value** field as containing some value appropriate for this class of object. Finally, every complete object descriptor array ends with a descriptor element that is identified by its fields as being a NULL-terminating element.

Thus, a minimum-size descriptor array consists of just two elements: the first contains its class identity, and the second is a NULL (it is legitimate for objects not to have values). When an object does have a value, it is held in the **value** field of a descriptor whose **type** field communicates the value's meaning.

Figure 3-2 illustrates the arrangement of a complete object descriptor array.

Figure 3-2. A Complete Object Represented



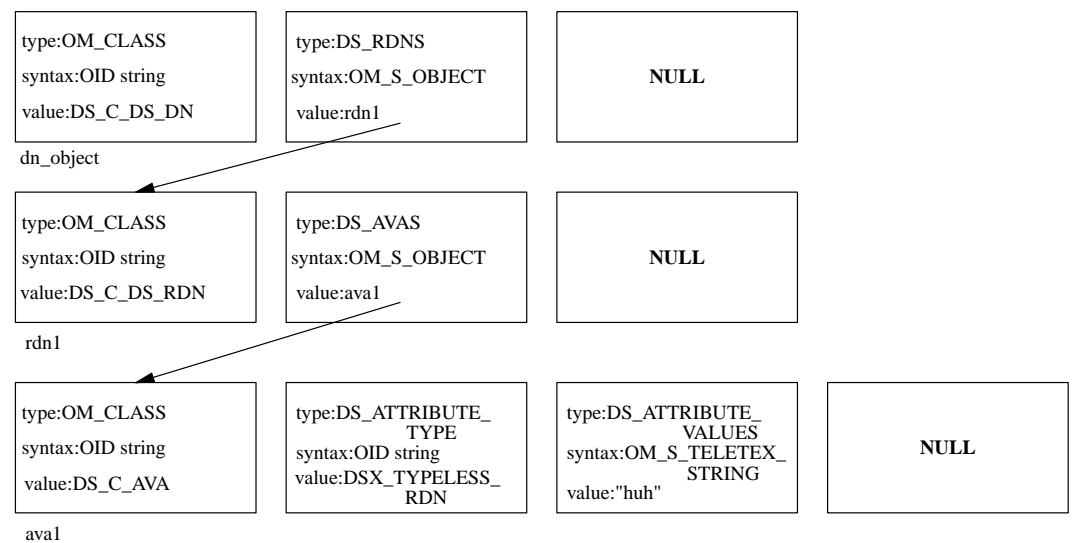
3.2.1 Object Attributes

The generic term for any object value is *attribute*. In this sense, an object is nothing but a collection of attributes, and every object descriptor describes one attribute. The first attribute's value identifies the object's class, and this determines all the other attributes the object is supposed to have. One or more other attributes follow, which contain the object's working values. The NULL object descriptor at the end is an implementation detail, and is not a part of the object.

Note that, depending on the attribute it represents, a descriptor’s **value** field can contain a pointer to another array of object descriptors. In other words, an object’s value can be another object.

Figure 3-3 shows a three-layer compound object: the top-level superobject, **dn_object**, contains the subobject **rdn1**, which in turn contains the subobject **ava1**.

Figure 3–3. A Three-Layer Compound Object



3.2.2 Interface Objects and Directory Objects

GDS is composed of directory objects that reflect the X.500 design. The XDS interface also works with objects. However, there is a big difference between directory and XDS objects. Programmers do not work directly with the directory objects; they are composed of attributes that make up the directory service’s implementation of entries.

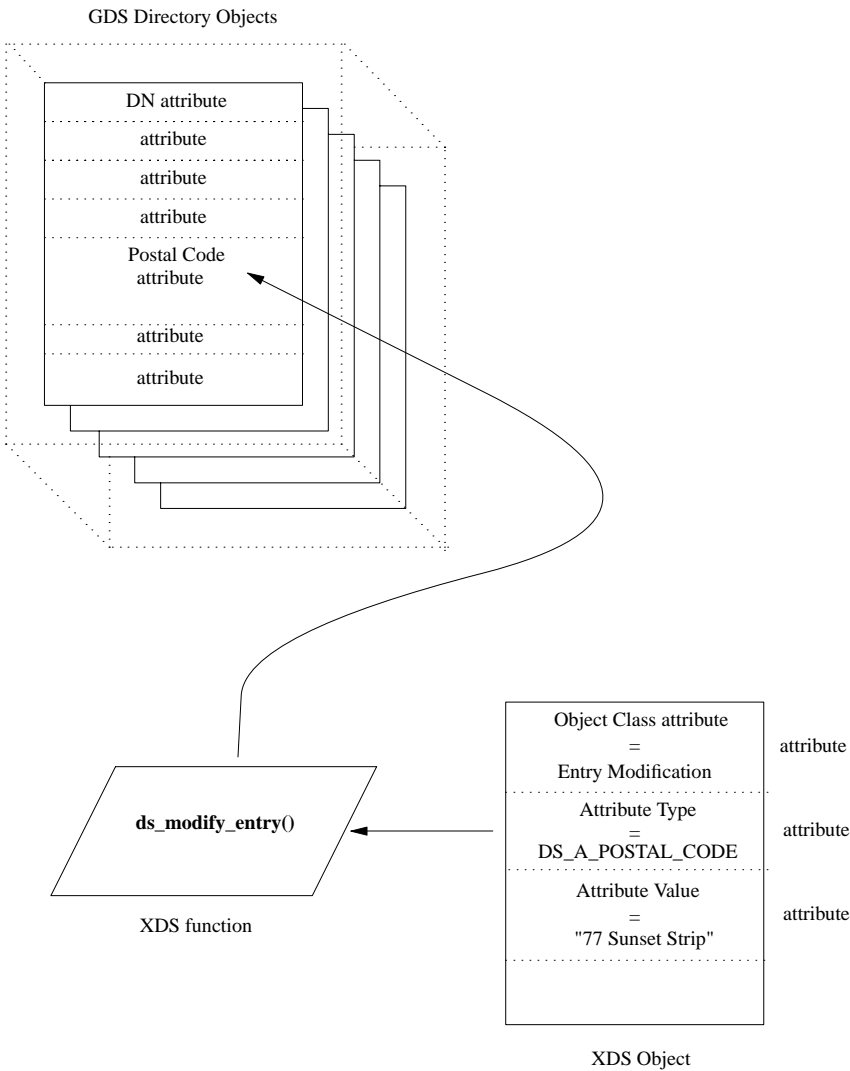
Programmers work with XDS objects. XDS objects have explicit data representations that can be directly manipulated with programming language operators. Some of these techniques are described in this chapter; others can be found in Chapter 7.

XDS and GDS terminology sometimes suggests that XDS objects are somehow direct representations of the directory objects to which they communicate information. This is not the case, however. You never directly see or manipulate the directory objects; the XDS interface objects are used only to pass parameters to the XDS calls, which in turn request GDS (or CDS) to perform operations on the directory objects. The XDS objects are therefore somewhat arbitrary structures defined by the interface.

Figure 3-4 illustrates the relationship between XDS (also called *interface*) objects and directory objects. The figure shows an application passing several properly initialized XDS objects to some XDS function; it then takes some action, which affects the attribute contents of certain directory objects. The application never works with the directory objects; it works with the XDS interface objects.

A side effect of the existence of a separate XDS interface and GDS or CDS directory objects is the existence of attributes for both kinds of objects as well. Because the purpose of XDS objects is to feed data into and extract data from directory objects, programmers work with XDS objects whose attributes have *directory* object attributes as their values. You should keep in mind the distinction between directory objects and interface objects.

Figure 3–4. Directory Objects and XDS Interface Objects

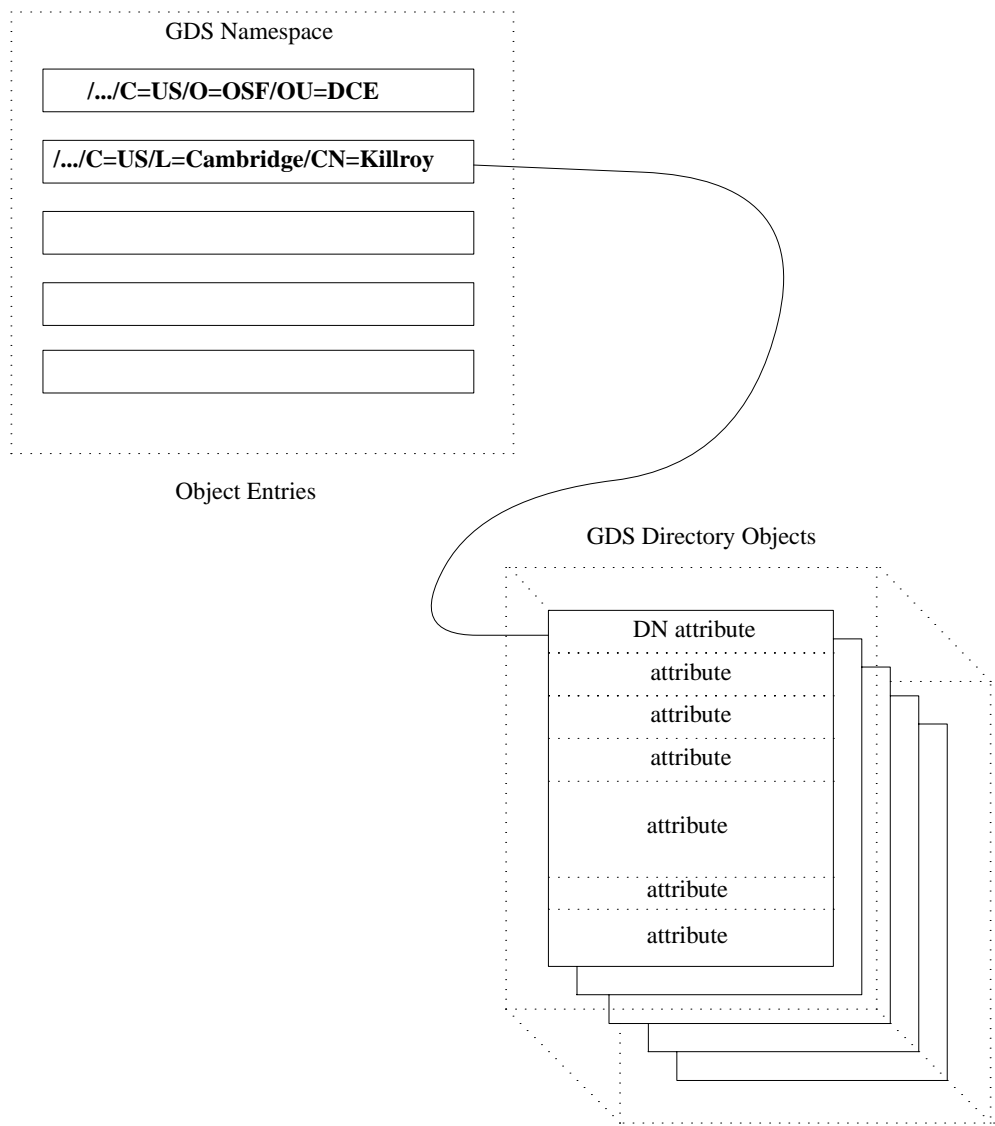


3.2.3 Directory Objects and Namespace Entries

The GDS namespace is a hierarchical collection of entries. The name of each of these entries is an attribute of a directory object. The object is accessed through XDS by stating its name attribute.

Figure 3-5 shows the relationship of entry names in the GDS namespace to the GDS directory objects to which they refer.

Figure 3–5. Directory Objects and Namespace Entries



3.2.4 Values That an Object Can Contain

There are many different classes of objects defined for the XDS interface; still more are defined by the X.500 standard for general directory use. But only a small number of classes are needed for XDS/CDS operations, and only those classes are discussed in this chapter. Information about other classes can be found in Part 4 of this guide.

The class that an object belongs to determines what sort of information the object can contain. Each object class consists of a list of attributes that objects must have. For example, you would expect an object in the directory entry name class to be required to have an attribute to hold the entry name string. However, it is not sufficient to simply place a string like the following into an object descriptor:

```
/.../C=US/O=OSF/OU=DCE/hosts/tamburlaine/self
```

A full directory entry name such as the preceding one is called in XDS a *distinguished name* (DN), meaning that the entry name is fully qualified (distinct) from root to entry name. To properly represent the entry name in an object, you must look up the definition of the XDS distinguished name object class and build an object that has the set of attributes that the definition prescribes.

3.2.5 Building a Name Object

Complete definitions for all the object classes required as input for XDS functions can be found in Section 3.5. Among them is the class for distinguished name objects, called **DS_C_DS_DN**. There you will learn that this class of object has two attributes: its class attribute, which identifies it as a **DS_C_DS_DN** object, and a second attribute, which occurs multiple times in the object. Each instance of this attribute contains as its value one piece of the full name; for example, the directory name **hosts**.

The **DS_C_DS_DN** attribute that holds the entry name piece, or relative distinguished name (RDN), is defined by the class rules to hold, not a string, but another object of the RDN class (**DS_C_DS_RDN**).

Thus, a static declaration of the descriptor array representing the **DS_C_DS_DN** object would look like the following:

The definition for **DS_C_DS_RDN** can be found in Section 3.5.2.6. This class object is defined, like **DS_C_DS_DN**, to have only one attribute (with the exception of the **OM_Object** attribute, which is mandatory for all objects). The one attribute, **DS_AVAS**, holds the value of one relative name. The syntax of this value is **OM_S_OBJECT**, meaning that **DS_AVAS**'s value is a pointer to yet another object descriptor array:

```
static OM_descriptor    Country_RDN[] = {

    OM_OID_DESC(OM_CLASS, DS_C_DS_RDN),

    {DS_AVAS, OM_S_OBJECT, {0, Country_Value}},

    OM_NULL_DESCRIPTOR
};
```

Note that there should also be five other similar declarations, one for each of the other **DS_C_DS_RDN** objects held in **DS_C_DS_DN**.

The declarations have the same meanings as they did in the previous example. **Country_Value** is the name of the descriptor array that represents the object of class **DS_C_AVA**, which we are now about to look up.

The rules for the **DS_C_AVA** class can be found in this chapter just after **DS_C_DS_RDN**. They tell us that **DS_C_AVA** objects have two attributes aside from the omnipresent **OM_Object**; namely:

- **DS_ATTRIBUTE_VALUES**

This attribute holds the object's value.

- **DS_ATTRIBUTE_TYPE**

This attribute gives the meaning of the object's value.

In this instance, the meaning of the string **US** is that it is a country name. There is a particular directory service attribute value for this; it is identified by an OID that is associated with the label **DS_A_COUNTRY_NAME** (the OIDs held in objects are represented in string form). Accordingly, we make that OID the value of **DS_ATTRIBUTE_TYPE**, and we make the name string itself the value of **DS_ATTRIBUTE_VALUES**, as shown.

```
static OM_descriptor    Country_Value[] = {

    OM_OID_DESC(OM_CLASS, DS_C_AVA),

    OM_OID_DESC(DS_ATTRIBUTE_TYPE, DS_A_COUNTRY_NAME),

    {DS_ATTRIBUTE_VALUES, OM_S_PRINTABLE_STRING, OM_STRING("US")},
    /*                      ^^^^^^^^^^^^^^^^^^ */
    /*                      Macro to properly */
    /*      fill the "value" union with the NULL-terminated string. */

    OM_NULL_DESCRIPTOR

};
```

There are also five other **DS_C_AVA** declarations, one for each of the five other separate name piece objects referred to in the **DS_C_DS_RDN** superobjects.

3.2.6 A Complete Object

The previous sections described how an object is created: you look up the rules for the object class you require, and then add the attributes called for in the definition. Whenever some attribute is defined to have an object as its value, you have to look up the class rules for the new object and declare a further descriptor array for it. In this way, you continue working down through layers of subobjects until you reach an object class that contains no subobjects as values; at that point, you are finished.

Normally, you do not statically declare objects in real applications. The steps outlined in the preceding text are given as a method for determining what an object looks like. Once you have done that, you can then write routines to create the objects dynamically. An example of how to do this can be found in the **teldir.c** example application in Chapter 7 of this guide.

The process of object building is somewhat easier than it sounds. There are only five different object classes needed for input to XDS functions when accessing CDS, and only one of those, the **DS_C_DS_DN** class, has more than one level of subobjects. The rules for all five of these classes can be found in Part 4 of this guide. In order to use the GDS references, you should know a few things about class hierarchy.

3.2.7 Class Hierarchy

Object classes are hierarchically organized so that some classes may be located above some classes in the hierarchy and below others in the hierarchy. In any such system of subordinate classes, each next lower class inherits all the attributes prescribed for the class immediately above it, plus whatever attributes are defined peculiarly for it alone. If the hierarchy continues further down, cumulative collection of attributes continues to accumulate. If there were a class for every letter of the alphabet, starting at the highest level with A and continuing down to the lowest level with Z, and if each succeeding letter was a subclass of its predecessor, the Z class would possess all the attributes of all the other letters, as well as its own, while the A class would possess only the A class attributes.

XDS/XOM classes are seldom nested more than two or at most three layers. All inherited attributes are explicitly listed in the object descriptions that follow, so you do not have to worry about class hierarchies here. However, the complete descriptions of XDS/XOM objects in Part 4 of this guide rely on statements of class inheritance to fill out their attribute lists for the different classes. Refer to Part 4 for information about the classes of objects that can be returned by XDS calls in order to be able to handle those returned objects.

3.2.8 Class Hierarchy and Object Structure

Note that *class hierarchy* is different from *object structure*. Object structure is the layering of object arrays that was previously described in the **DS_C_DS_DN** declaration in Section 3.2.5. It occurs when one object contains another object as the value of one or more of its attributes.

This is what is meant by recursive objects: one object can point to another object as one of its attribute values. The layering of subobjects below superobjects in this way is described repeatedly in Section 3.5.

The only practical significance of class hierarchy is in determining all the attributes a certain object class must have. Once you have done this, you may find that a certain attribute requires as its value some other object. The result is a compound object, but this is completely determined by the attributes for the particular class you are looking at.

3.2.9 Public and Private Objects and XOM

In Section 3.2.5, you saw how a multilevel XDS object can be statically declared in C code. Now imagine that you have written an application that contains such a static **DS_C_DS_DN** object declaration. From the point of view of your application, that object is nothing but a series of arrays, and you can manipulate them with all the normal programming operators, just as you can any other data type. Nevertheless, the object is syntactically perfectly acceptable to any XDS (or XOM) function that is prepared to receive a **DS_C_DS_DN** object.

Objects are also created by the XDS functions themselves; this is the way they usually return information to callers. However, there is a difference between objects generated by the XDS interface and objects that are explicitly declared by the application: you cannot access the former, *private*, objects in the direct way that you can the latter, *public*, objects.

These two kinds of objects are the same as far as their classes and attributes are concerned. The only difference between them is in the way they are accessed. The public objects that an application explicitly creates or declares in its own memory area are just as accessible as any of the other data storage it uses. However, private objects are created and held in the XDS interface's own system memory. Applications get handles to private objects, and, in order to access the private objects' contents, they have to pass the handles to object management functions. The object management (XOM) functions make up a sort of all-purpose companion interface to XDS. Whereas XDS functions typically require some specific class object as input, the XOM functions accept objects of any class and perform useful operations on them.

If a private object needs to be manipulated, one of the XOM functions, **om_get()**, can be called to make a public copy of the private object. Then, calling the XOM **om_create()** function allows applications to generate private objects manipulable by **om_get()**. The main significance of private as opposed to public objects is that they do not have to be explicitly operated on; instead, you can access them cleanly through the XOM interface and let it do most of the work. You still have to know something about the objects' logical representation, however, to use XOM.

Except for a few more details, which will be mentioned as needed, this is practically all there is to XDS object representation.

3.2.10 XOM Objects and XDS Library Functions

To call an XDS library function, do the following:

1. Decide what input parameters you must supply to the function.
2. Bundle up these parameters in objects (that is, arrays of object descriptors) in an XDS format.

Almost all data returned to you by an XDS function is enclosed in objects, which you must parse to recover the information that you want. This task is made almost automatic by a library function supplied with the companion X/Open OSI-Abstract-Data Manipulation (XOM) interface.

With XDS, the programmer has to perform a lot of call parameter management, but in other respects the interface is easy to use. The XDS functions' dependence on objects makes them easy to call, once you have the objects themselves correctly set up.

3.3 Accessing CDS Using the XDS Step-by-Step Procedure

You now know all that you need to know to work with a cell namespace through XDS. The following subsections provide a walk-through of the steps of some typical XDS/CDS operations. They describe what is involved in using XDS to access existing CDS attributes. They then describe how you can create and access new CDS entry attributes.

3.3.1 Reading and Writing Existing CDS Entry Attributes With XDS

Suppose that you want to use XDS to read some information from the following CDS entry:

```
/.../C=US/O=OSF/OU=DCE/hosts/tamburlaine/self
```

As explained in the *OSF DCE Administration Guide*, the `./:/hosts/hostname/self` entry, which is created at the time of cell configuration, contains binding information for the machine *hostname*. Since this is a simple RPC NSI entry, there is not very much in the entry that is interesting to read, but this entry is used as an example anyway as a simple demonstration.

Following are the header inclusions and general data declarations:

```
#include <stdio.h>
#include <stdlib.h>
#include <string.h>
#include <xom.h>
#include <xds.h>
#include <xdsbdcp.h>
#include <xds cds.h>
```

Note that the **xom.h** and **xds.h** header files must be included in the order shown in the preceding example. Also note that the **xds cds.h** header file is brought in for the sake of **DSX_TYPELESS_RDN**. This file is where the CDS-significant OIDs are defined. The **xdsbdcp.h** file contains information necessary to the Basic Directory Contents Package, which is the basic version of the XDS interface you can use in this program.

The XDS/XOM interface defines numerous object identifier string constants, which are used to identify the many object classes, parts, and pieces (among other things) that it needs to know about. In order to make sure that these OID constants do not collide with any other constants, the interface refers to them with the string **OMP_O_** prefixed to the user-visible form; for example, **DS_C_DS_DN** becomes **OMP_O_DS_C_DS_DN** internally. In order to make application instances consistent with the internal form, use **OM_EXPORT** to import *all* XDS-defined or XOM-defined OID constants used in your application.

```
OM_EXPORT( DS_A_COUNTRY_NAME )
OM_EXPORT( DS_A_OBJECT_CLASS )
OM_EXPORT( DS_A_ORG_UNIT_NAME )
OM_EXPORT( DS_A_ORG_NAME )

OM_EXPORT( DS_C_ATTRIBUTE )
OM_EXPORT( DS_C_ATTRIBUTE_LIST )
OM_EXPORT( DS_C_AVA )
OM_EXPORT( DS_C_DS_DN )
```

```
OM_EXPORT( DS_C_DS_RDN )
OM_EXPORT( DS_C_ENTRY_INFO_SELECTION )
OM_EXPORT( DSX_TYPELESS_RDN )

/* ... Special OID for an untyped (that is, nonX.500)      */
/* relative distinguished name. Defined in xdscds.h header. */
```

A further important effect of **OM_EXPORT** is that it builds an **OM_string** structure to hold the exported OID hexadecimal string. As explained in the previous chapter, OIDs are not numeric values, but strings. Comparisons and similar operations on OIDs must access them as strings. Once an OID has been exported, you can access it by using its declared name. For example, the hexadecimal string representation of **DS_C_ATTRIBUTE** is contained in **DS_C_ATTRIBUTE.elements**, and the length of this string is contained in **DS_C_ATTRIBUTE.length**.

3.3.1.1 Significance of Typed and Untyped Entry Names

Next are the static declarations for the lowest layer of objects that make up the global name (distinguished name) of the CDS directory entry you want to read. These lowest-level objects contain the string values for each part of the name. Remember that the first three parts of the name (excluding the global prefix **/.../**, which is not represented) constitute the cell name, as follows:

```
/C=US/O=OSF/OU=DCE/
```

In this example, assume that GDS is being used as the cell's global directory service, so the cell name is represented in X.500 format, and each part of it is typed in the object representation; for example, **DS_A_COUNTRY_NAME** is the **DS_ATTRIBUTE_TYPE** in the **Country_String_Object**. If you were using DNS, the cell name might be something like the following:

```
osf.org.dce
```

In this case, the entire string **osf.org.dce** would be held in a single object whose **DS_ATTRIBUTE_TYPE** would be **DSX_TYPELESS_RDN**.

DSX_TYPELESS_RDN is a special type that marks a name piece as not residing in an X.500 namespace. If the object resides under a typed X.500 name, as is the case

in the declared object structures, then it serves as a delimiter for the end of the cell name GDS looks up, and the beginning of the name that is passed to a CDS server in that cell, assuming that the cell has access to GDS; if not, such a name cannot be resolved. In the following name, the untyped portion is at the beginning:

`/.../osf.org.dce/hosts/zenocrate/self`

In this case, the name is passed immediately by XDS via the local CDS (and the GDA) to DNS for resolution of the cell name. Thus, the typing of entry names determines which directory service a global directory entry name is sent to for resolution.

3.3.1.2 Static Declarations

The following are the static declarations you need:

```
/* **** */
/* Here are the objects that contain the string values for each */
/* part of the CDS entry's global name... */

static OM_descriptor Country_String_Object[] = {
    OM_OID_DESC(OM_CLASS, DS_C_AVA),
    OM_OID_DESC(DS_ATTRIBUTE_TYPE, DS_A_COUNTRY_NAME),
    {DS_ATTRIBUTE_VALUES, OM_S_PRINTABLE_STRING, OM_STRING("US")},
    OM_NULL_DESCRIPTOR
};

static OM_descriptor Organization_String_Object[] = {
    OM_OID_DESC(OM_CLASS, DS_C_AVA),
    OM_OID_DESC(DS_ATTRIBUTE_TYPE, DS_A_ORG_NAME),
    {DS_ATTRIBUTE_VALUES, OM_S_PRINTABLE_STRING, OM_STRING("OSF")},
    OM_NULL_DESCRIPTOR
};

static OM_descriptor Org_Unit_String_Object[] = {
    OM_OID_DESC(OM_CLASS, DS_C_AVA),
    OM_OID_DESC(DS_ATTRIBUTE_TYPE, DS_A_ORG_UNIT_NAME),
    {DS_ATTRIBUTE_VALUES, OM_S_PRINTABLE_STRING, OM_STRING("DCE")},
    OM_NULL_DESCRIPTOR
};
```



```

};

static OM_descriptor    Hosts_Dir_String_Object[] = {
    OM_OID_DESC(OM_CLASS, DS_C_AVA),
    OM_OID_DESC(DS_ATTRIBUTE_TYPE, DSX_TYPELESS_RDN),
    {DS_ATTRIBUTE_VALUES, OM_S_PRINTABLE_STRING, OM_STRING("hosts")},
    OM_NULL_DESCRIPTOR
};

static OM_descriptor    Tamburlaine_Dir_String_Object[] = {
    OM_OID_DESC(OM_CLASS, DS_C_AVA),
    OM_OID_DESC(DS_ATTRIBUTE_TYPE, DSX_TYPELESS_RDN),
    {DS_ATTRIBUTE_VALUES, OM_S_PRINTABLE_STRING, OM_STRING("tamburlaine")},
    OM_NULL_DESCRIPTOR
};

static OM_descriptor    Self_Entry_String_Object[] = {
    OM_OID_DESC(OM_CLASS, DS_C_AVA),
    OM_OID_DESC(DS_ATTRIBUTE_TYPE, DSX_TYPELESS_RDN),
    {DS_ATTRIBUTE_VALUES, OM_S_PRINTABLE_STRING, OM_STRING("self")},
    OM_NULL_DESCRIPTOR
};

```

The string objects are contained by a next-higher level of objects that identify the strings as being pieces (RDNs) of a fully qualified directory entry name (DN). Thus, the **Country_RDN** object contains **Country_String_Object** as the value of its **DS_AVAS** attribute; **Organization_RDN** contains **Organization_String_Object**, and so on.

```

/*****
/* Here are the "relative distinguished name" objects.

static OM_descriptor    Country_RDN[] = {
    OM_OID_DESC(OM_CLASS, DS_C_DS_RDN),
    {DS_AVAS, OM_S_OBJECT, {0, Country_String_Object}},
    OM_NULL_DESCRIPTOR
};

static OM_descriptor    Organization_RDN[] = {

```

```
OM_OID_DESC(OM_CLASS, DS_C_DS_RDN),
{DS_AVAS, OM_S_OBJECT, {0, Organization_String_Object}},
OM_NULL_DESCRIPTOR
};

static OM_descriptor    Org_Unit_RDN[] = {
    OM_OID_DESC(OM_CLASS, DS_C_DS_RDN),
    {DS_AVAS, OM_S_OBJECT, {0, Org_Unit_String_Object}},
    OM_NULL_DESCRIPTOR
};

static OM_descriptor    Hosts_Dir_RDN[] = {
    OM_OID_DESC(OM_CLASS, DS_C_DS_RDN),
    {DS_AVAS, OM_S_OBJECT, {0, Hosts_Dir_String_Object}},
    OM_NULL_DESCRIPTOR
};

static OM_descriptor    Tamburlaine_Dir_RDN[] = {
    OM_OID_DESC(OM_CLASS, DS_C_DS_RDN),
    {DS_AVAS, OM_S_OBJECT, {0, Tamburlaine_Dir_String_Object}},
    OM_NULL_DESCRIPTOR
};

static OM_descriptor    Self_Entry_RDN[] = {
    OM_OID_DESC(OM_CLASS, DS_C_DS_RDN),
    {DS_AVAS, OM_S_OBJECT, {0, Self_Entry_String_Object}},
    OM_NULL_DESCRIPTOR
};
```

At the highest level, all the subobjects are gathered together in the DN object named **Full_Entry_Name_Object**.

```
/*
/*****/

static OM_descriptor    Full_Entry_Name_Object[] = {
    OM_OID_DESC(OM_CLASS, DS_C_DS_DN),
    {DS_RDNS, OM_S_OBJECT, {0, Country_RDN}},
    {DS_RDNS, OM_S_OBJECT, {0, Organization_RDN}},
    {DS_RDNS, OM_S_OBJECT, {0, Org_Unit_RDN}},
    {DS_RDNS, OM_S_OBJECT, {0, Hosts_Dir_RDN}},
    OM_NULL_DESCRIPTOR
};
```

```
    {DS_RDNS, OM_S_OBJECT, {0, Tamburlaine_Dir_RDN}},  
    {DS_RDNS, OM_S_OBJECT, {0, Self_Entry_RDN}},  
    OM_NULL_DESCRIPTOR  
};
```

3.3.1.3 Other Necessary Objects for **ds_read()**

The **ds_read()** procedure takes requests in the form of a **DS_C_ENTRY_INFO_SELECTION** class object. However, if you refer to the recipe for this object class in Section 3.5, you will find that it is much simpler than the name object; it contains no subobjects, and its declaration is straightforward.

The value of the **DS_ALL_ATTRIBUTES** attribute specifies that all attributes be read from the CDS entry, which is specified in the **Full_Entry_Name_Object** variable.

Note that the term *attribute* is used slightly differently in CDS and XDS contexts. In XDS, attributes describe the values that can be held by various object classes; they can be thought of as *object fields*. In CDS, attributes describe the values that can be associated with a directory entry. The following code fragment shows the definition of a **DS_C_ENTRY_INFO_SELECTION** object:

```
static OM_descriptor    Entry_Info_Select_Object[] = {  
  
    OM_OID_DESC(OM_CLASS, DS_C_ENTRY_INFO_SELECTION),  
    {DS_ALL_ATTRIBUTES, OM_S_BOOLEAN, OM_TRUE},  
    {DS_INFO_TYPE, OM_S_ENUMERATION, DS_TYPES_AND_VALUES},  
    OM_NULL_DESCRIPTOR  
};
```

3.3.1.4 Miscellaneous Declarations

The following are declarations for miscellaneous variables:

```
OM_workspace            xdsWorkspace;  
    /* ...will contain handle to our "workspace" */
```

```
DS_feature featureList[] = {

    { OM_STRING(OMP_O_DS_BASIC_DIR_CONTENTS_PKG), OM_TRUE },
    { 0 }
};

/* ...list of service "packages" we will want from XDS */

OM_private_object      session;
/* ...will contain handle to a bound-to directory session */

DS_status              dsStatus;
/* ...status return from XDS calls */

OM_return_code         omStatus;
/* ...status return from XOM calls */

OM_sint               dummy;
/* ...for unsupported ds_read() argument */

OM_private_object      readResultObject;
/* ...to receive entry information read from CDS by "ds_read()" */

OM_type I_want_entry_object[] = {DS_ENTRY, OM_NO_MORE_TYPES};
OM_type I_want_attribute_list[] = {DS_ATTRIBUTES, OM_NO_MORE_TYPES};
OM_type I_want_attribute_value[] = {DS_ATTRIBUTE_VALUES, \
                                   OM_NO_MORE_TYPES};
/* ...arrays to pass to "om_get()" to extract subobjects */
/* from the result object returned by "ds_read()" */

OM_value_position number_of_descriptors;
/* ...to hold number of attribute descriptors returned */
/* by "om_get()" */

OM_public_object entry;
/* ...to hold public object returned by "om_get()" */
```

3.3.1.5 The Main Program

This section describes the main program. Three calls usually precede any use of XDS.

First, **ds_initialize()** is called to set up a *workspace*. A workspace is a memory area in which XDS can generate objects that will be used to pass information to the application. If the call is successful, it returns a handle that must be saved for the **ds_shutdown()** call. If the call is unsuccessful, it returns NULL, but this example does not check for errors.

```
xdsWorkspace = ds_initialize();
```

If GDS is being used as the global directory service, the service packages are specified next. Packages consist of groups of objects, together with the associated supporting interface functionality, designed to be used for some specific end. For example, to access the (X.500) Global Directory, specify **DSX_GDS_PKG**. This example uses the basic XDS service so **DS_BASIC_DIR_CONTENTS_PKG** is specified. The *featureList* parameter to **ds_version()** is an array, not an object, since packages are not being handled yet:

```
dsStatus = ds_version(featureList, xdsWorkspace);
```

Note that, if you are *not* using GDS as your global directory service (in other words, if you are using XDS by itself), then you should *not* call **ds_version()**.

From this point on, status is returned by XDS functions via a **DS_status** variable. **DS_status** is a handle to a private object, whose value is **DS_SUCCESS** (that is, NULL) if the call was successful. If something went wrong, the information in the (possibly complex) private error object has to be analyzed through calls to **om_get()**, which is one of the general-purpose object management functions that belongs to XDS's companion interface XOM. Usage of **om_get()** is demonstrated later on in this program, but return status is not checked in this example.

The third necessary call is to **ds_bind()**. This call brings up the directory service, which binds to a Directory System Agent (DSA), the GDS server, through a Directory User Agent (DUA), the GDS client. The **DS_DEFAULT_SESSION** parameter calls for a default session. The alternative is to build and fill out your own **DS_C_SESSION** object, specifying such things as DSA addresses, and pass that. The default is used in this example:

```
dsStatus = ds_bind(DS_DEFAULT_SESSION, xdsWorkspace, &session);
```

3.3.1.6 Reading a CDS Attribute

At this point, you can read a set of object attributes from the cell namespace entry. Call **ds_read()** with the two objects that specify the entry to be read and the specific entry attribute you want:

```
dsStatus = ds_read(session, DS_DEFAULT_CONTEXT, Full_Entry_Name_Object,  
                  Entry_Info_Select_Object, &readResultObject, &dummy);
```

The **DS_DEFAULT_CONTEXT** parameter could be substituted with a **DS_C_CONTEXT** object, which would typically be reused during a series of related XDS calls. This object specifies and records how GDS should perform the operation, how much progress has been made in resolving a name, and so on.

If the call succeeds, the private object **readResultObject** contains a series of **DS_C_ATTRIBUTE** subobjects, one for each attribute read from the cell name entry. A complete recipe for the **DS_C_READ_RESULT** object can be found in Chapter 11, but the following is a skeletal outline of the object's structure:

```
DS_C_READ_RESULT  
    DS_ENTRY: object(DS_C_ENTRY_INFO)  
    DS_ALIAS_DEREFERENCED: OM_S_BOOLEAN  
    DS_PERFORMER: object(DS_C_NAME)  
  
    DS_C_ENTRY_INFO  
        DS_FROM_ENTRY: OM_S_BOOLEAN  
        DS_OBJECT_NAME: object(DS_C_NAME)  
        DS_ATTRIBUTES: one or more object(DS_C_ATTRIBUTE)  
  
    DS_C_NAME == DS_C_DS_DN  
        DS_RDNS: object(DS_C_DS_RDN)  
  
    DS_C_DS_RDN  
        DS_AVAS: object(DS_C_AVA)
```

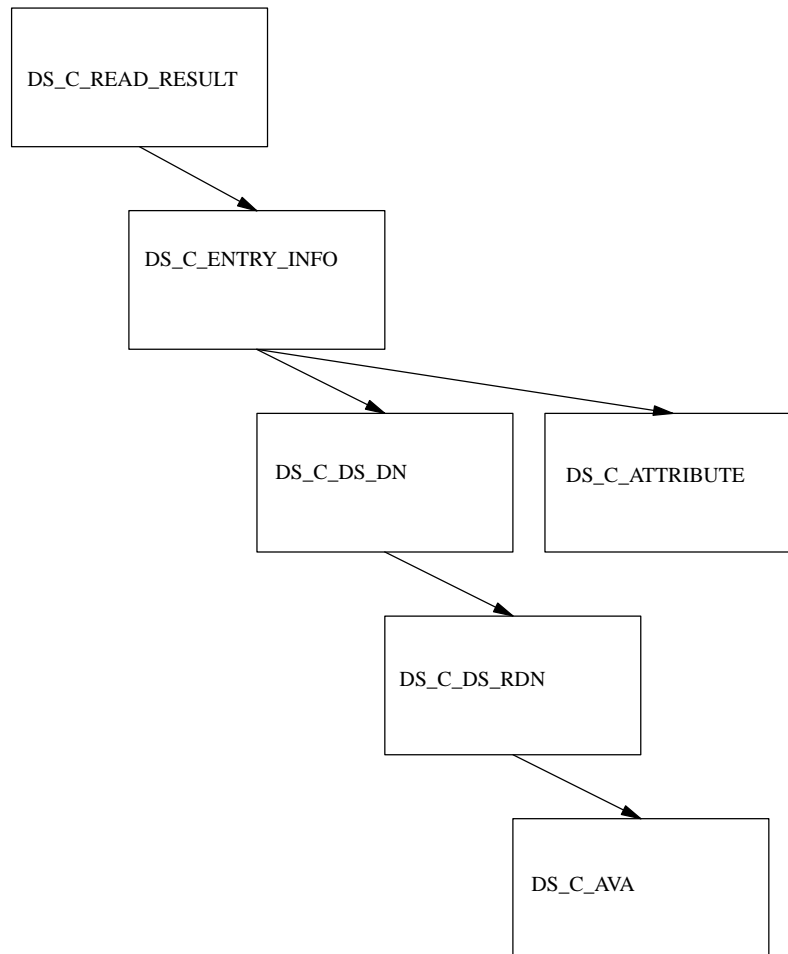
```
DS_C_AVA
    DS_ATTRIBUTE_TYPE: OID string
    DS_ATTRIBUTE_VALUES: anything

DS_C_ATTRIBUTE ---one for each attribute read
    DS_ATTRIBUTE_TYPE: OID string
    DS_ATTRIBUTE_VALUES: anything

DS_C_ATTRIBUTE
    DS_ATTRIBUTE_TYPE: OID string
    DS_ATTRIBUTE_VALUES: anything
```

Figure 3-6 illustrates the general object structure of a **DS_C_READ_RESULT**, showing only the object-valued attributes, and only one **DS_C_ATTRIBUTE** subobject.

Figure 3–6. The DS_C_READ_RESULT Object Structure



3.3.1.7 Handling the Result Object

The next goal is to extract the instances of the **DS_C_ATTRIBUTE** subclass, one for each attribute read, from the returned object. The first step is to make a public copy of **readResultObject**, which is a *private* object, and therefore does not allow

access to the object descriptors themselves. Using the XOM **om_get()** function, you can make a public copy of **readResultObject**, and at the same time specify that only the relevant parts of it be preserved in the copy. Then, with a couple of calls to **om_get()**, you can reduce the object to manageable size, leaving a superobject whose immediate subobjects are fairly easily accessed.

The **om_get()** function takes as its third input parameter an **OM_type_list**, which is an array of **OM_type**. Possible parameters are **DS_ENTRY**, **DS_ATTRIBUTES**, **DS_ATTRIBUTE_VALUES**, and anything that can legitimately appear in an object descriptor's **type** field. The types specified in this parameter are interpreted according to the options specified in the preceding parameter. For example, the relevant attribute from the read result is **DS_ENTRY**. It contains the **DS_C_ENTRY_INFO** object, which in turn contains the **DS_C_ATTRIBUTE** objects. The **DS_C_ATTRIBUTE** objects contain the data read from the cell directory name entry. Therefore, you should specify the **OM_EXCLUDE_ALL_BUT_THESE_TYPES** option, which has the effect of excluding everything but the contents of the object's **DS_ENTRY** type attribute.

The **OM_EXCLUDE_SUBOBJECTS** option is also ORed into the parameter. Why would you not preserve the subobjects of **DS_C_ENTRY_INFO**? Because **om_get()** works only on private, not on public, objects. If you were to use **om_get()** on the entire object substructure, you would not be able to continue getting the subobjects, and instead you would have to follow the object pointers down to the **DS_C_ATTRIBUTE**. However, when **om_get()** excludes subobjects from a copy, it does not really leave them out; it merely leaves the subobjects private, with a handle to the private objects where pointers would have been. This allows you to continue to call **om_get()** as long as there are more subobjects.

The following is the first call:

```
/* The DS_C_READ_RESULT object that ds_read() returns has */
/* one subobject, DS_C_ENTRY_INFO; it in turn has two sub- */
/* objects, that is a DS_C_NAME which holds the object's */
/* distinguished name (which we don't care about here), */
/* and a DS_C_ATTRIBUTE which contains the attribute info */
/* we read; that one we want. So we climb down to it ... */
/* This om_get() will "return" the entry-info object ... */

omStatus = om_get(readResultObject,
                  OM_EXCLUDE_ALL_BUT_THESE_TYPES +
```

```
OM_EXCLUDE_SUBOBJECTS,  
I_want_entry_object,  
OM_FALSE,  
OM_ALL_VALUES,  
OM_ALL_VALUES,  
&entry,  
&number_of_descriptors);
```

The **number_of_descriptors** parameter contains the number of attribute descriptors returned in the public copy, not in any excluded subobjects.

If an XOM function is successful, it returns an **OM_SUCCESS** code. Unsuccessful calls to XOM functions do not return error objects, but rather return simple error codes. The interface assumes that, if the XOM function does not accept your object, then you will not be able to get much information from any further objects. The return status is not checked in this example.

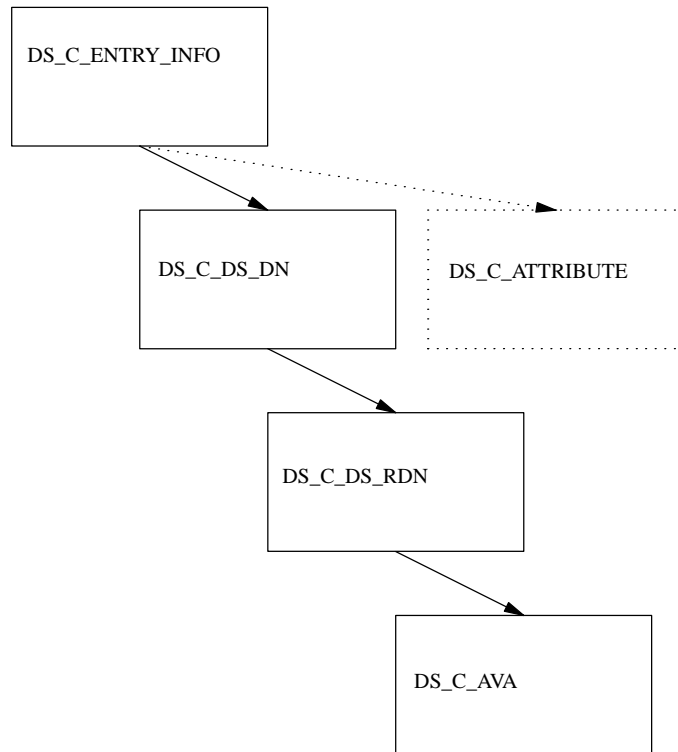
The return parameter **entry** should now contain a pointer to the **DS_C_ENTRY_INFO** object with the following immediate structure. (The number of instances of **DS_ATTRIBUTES** depends on the number of attributes read from the entry.)

```
DS_C_ENTRY_INFO  
DS_FROM_ENTRY: OM_S_BOOLEAN  
DS_OBJECT_NAME: object(DS_C_NAME)  
DS_ATTRIBUTES: object(DS_C_ATTRIBUTE)  
DS_C_ATTRIBUTE  
DS_ATTRIBUTE_TYPE: OID string  
DS_ATTRIBUTE_VALUES: anything  
  
DS_ATTRIBUTES: object(DS_C_ATTRIBUTE)  
                  object(DS_C_ATTRIBUTE)  
  
DS_C_ATTRIBUTE  
DS_ATTRIBUTE_TYPE: OID string  
DS_ATTRIBUTE_VALUES: anything
```

The italics indicate private subobjects. Figure 3-7 shows the **DS_C_ENTRY_INFO** object. Only one instance of a **DS_C_ATTRIBUTE** subobject is shown in the figure; usually there are several such subobjects, all at the same level, each containing information about one of the attributes read from the entry. These

subobjects are represented in **DS_C_ENTRY_INFO** as a series of descriptors of type **DS_ATTRIBUTES**, each of which has as its value a separate **DS_C_ATTRIBUTE** subobject.

Figure 3–7. The DS_C_ENTRY_INFO Object Structure



Now extract the separate attribute values of the entry that was read. These were returned as separate object values of **DS_ATTRIBUTES**; each one has an object class of **DS_C_ATTRIBUTE**. To return any one of these subobjects, a second call to **om_get()** is necessary, as follows:

```

/* The second om_get() returns one selected subobject */
/* from the DS_C_ENTRY_INFO subobject we just got. The */
/* contents of "entry" as we enter this call is the */
/* private subobject which is the value of DS_ATTRIBUTES. */

```

```
/* If we were to make the following call with the          */
/* OM_EXCLUDE_SUBOBJECTS and without the                  */
/* OM_EXCLUDE_ALL_BUT_THESE_VALUES flags, we would get    */
/* back an object consisting of six private subobjects,    */
/* one for each of the attributes returned. Note the      */
/* values for initial and limiting position: "2"          */
/* specifies that we want only the third DS_C_ATTRIBUTE   */
/* subobject to be gotten (the subobjects are numbered   */
/* from 0, not from 1), and the "3" specifies that we want */
/* no more than that--in other words, the limiting value  */
/* must always be one more than the initial value if the  */
/* latter is to have any effect.                          */
/* OM_EXCLUDE_ALL_BUT_THESE_VALUES is likewise required */
/* for the initial and limiting values to have any       */
/* effect ...                                             */

omStatus = om_get(entry->value.object.object,
                  OM_EXCLUDE_ALL_BUT_THESE_TYPES
                  + OM_EXCLUDE_SUBOBJECTS
                  + OM_EXCLUDE_ALL_BUT_THESE_VALUES,
                  I_want_attribute_list,
                  OM_FALSE,
                  ((OM_value_position) 2),
                  ((OM_value_position) 3),
                  &entry,
                  &number_of_descriptors);
```

Note the value that is passed as the first parameter. Since **om_get()** does not work on public objects, pass it the handle of the private subobject explicitly. To do this you have to know the arrangement of the descriptor's value union, which is defined in **xom.h**.

3.3.1.8 Representation of Object Values

The following is the layout of the **object** field in a descriptor's **value** union:

```
typedef struct {
    OM_uint32      padding;
```

```
    OM_object      object;  
} OM_padded_object;
```

The following is the layout of the **value** union itself:

```
typedef union OM_value_union {  
    OM_string      string;  
    OM_boolean     boolean;  
    OM_enumeration enumeration;  
    OM_integer     integer;  
    OM_padded_object object;  
} OM_value;
```

The following is the layout of the descriptor itself:

```
typedef struct OM_descriptor_struct {  
    OM_type          type;  
    OM_syntax        syntax;  
    union OM_value_union value;  
} OM_descriptor;
```

Thus, if **entry** is a pointer to the **DS_C_ENTRY_INFO** object, then the private handle to the **DS_C_ATTRIBUTE** object you want next is the following:

entry->value.object.object

3.3.1.9 Extracting an Attribute Value

The last call yielded one separate **DS_C_ATTRIBUTE** subsubobject from the original returned result object:

```
DS_C_ATTRIBUTE  
DS_ATTRIBUTE_TYPE: OID string  
DS_ATTRIBUTE_VALUES: anything
```

Figure 3-8 illustrates what is left.

Figure 3–8. The DS_C_ATTRIBUTE Object Structure



A final call to **om_get()** returns the single object descriptor that contains the actual value of the single attribute you selected from the returned object:

```
omStatus = om_get(entry->value.object.object,
                  OM_EXCLUDE_ALL_BUT_THESE_TYPES,
                  I_want_attribute_value,
                  OM_FALSE,
                  OM_ALL_VALUES,
                  OM_ALL_VALUES,
                  &entry,
                  &number_of_descriptors);
```

At this point, the value of **entry** is the base address of an object descriptor whose **entry->type** is **DS_ATTRIBUTE_VALUES**. Depending on the value found in **entry->syntax**, the value of the attribute can be read from **entry->value.string**, **entry->value.integer**, **entry->value.boolean**, or **entry->value.enumeration**.

For example, suppose the value of **entry->syntax** is **OM_S_OCTET_STRING**. The attribute value, represented as an octet string (*not* terminated by a NULL), is found in **entry->value.string.elements**; its length is found in **entry->value.string.length**.

You can check any attribute value against the value you get from the **cdscp** command by entering the following:

```
cdscp show object ./:/hosts/tamburlaine/self
```

For further information on **cdscp**, see the *OSF DCE Command Reference*.

Note that you can always call **om_get()** to get the *entire* returned object from an XDS call. This yields a full structure of object descriptors that you can manipulate like any other data structure. To do this with the **ds_read()** return object would have required the following call:

```
/* make a public copy of ENTIRE object... */
```

```
omStatus = om_get(readResultObject,
                  OM_NO_EXCLUSIONS,
                  ((OM_type_list) 0),
                  OM_FALSE,
                  ((OM_value_position) 0),
                  ((OM_value_position) 0),
                  &entry,
                  &number_of_descriptors);
```

At the end of every XDS session, you need to unbind from GDS and then deallocate the XDS and XOM structures and other storage. You must also explicitly deallocate any service-generated objects, whether public or private, with calls to **om_delete()**, as follows:

```
/* delete service-generated public or private objects... */

omStatus = om_delete(readResultObject);
omStatus = om_delete(entry);

/* unbind from the GDS... */
dsStatus = ds_unbind(session);

/* close down the workspace... */
dsStatus = ds_shutdown(xdsWorkspace);

exit();
```

3.3.2 Creating New CDS Entry Attributes

The following subsections provide the procedure and some code examples for creating new CDS entry attributes.

3.3.2.1 Procedure for Creating New Attributes

To create new attributes of your own on cell namespace entries, you must do the following:

1. Allocate a new ISO OID for the new attribute. For information on how to do this, see Chapter 2 of this guide and the *OSF DCE Administration Guide*.
2. Enter the new attribute's name and OID in the file `./opt/dcelocal/etc/cds_attributes`. This text file contains OID-to-readable string mappings that are used, for example, by `cdscp` when it displays CDS entry attributes. Each entry also gives a syntax for reading the information in the entry itself. This should be congruent with the format of the data you intend to write in the attribute. For more information about the `cds_attributes` file, see the *OSF DCE Administration Guide*.
3. In the `xdscds.h` header file, define an appropriate OID string constant to represent the new attribute.

For example, the following shows the `xdscds.h` definition for the CDS `CDS_Class` attribute:

```
#define OMP_O_DSX_A_CDS_Class    "\x2B\x16\x01\x03\x0F"
```

Note the XDS internal form of the name. This is what `DSX_A_CDS_Class` looks like when it has been exported using `OM_EXPORT` in an application, as all OIDs must be. Thus, if you wanted to create a CDS attribute called `CDS_Brave_New_Attrib`, you would obtain an OID from your administrator and add the following line to `xdscds.h`:

```
#define OMP_O_DSX_A_CDS_Brave_New_Attrib "your_OID"
```

4. In an application, call the XDS `ds_modify_entry()` routine to add the attribute to the cell namespace entry of your choice.

3.3.2.2 Coding Examples

In the following code fragments, a set of declarations similar to those in the previous examples is assumed.

The `ds_modify_entry()` function, which is called to add new attributes to an entry or to write new values into existing attributes, requires a `DS_C_ENTRY_MOD_LIST` input object whose contents specify the attributes and values to be written to the

entry. The name, as always, is specified in a **DS_C_DS_DN** object. The following is a static declaration of such a list, which consists of two attributes:

```
static OM_descriptor    Entry_Modification_Object_1[] = {
    OM_OID_DESC(OM_CLASS, DS_C_ENTRY_MOD),
    OM_OID_DESC(DS_ATTRIBUTE_TYPE, DSX_A_CDS_Brave_New_Attrib),
    {DS_ATTRIBUTE_VALUES, OM_S_PRINTABLE_STRING,
    OM_STRING("O brave new attribute")},
    {DS_MOD_TYPE, OM_S_ENUMERATION, DS_ADD_ATTRIBUTE},
    OM_NULL_DESCRIPTOR
};

static OM_descriptor    Entry_Modification_Object_2[] = {
    OM_OID_DESC(OM_CLASS, DS_C_ENTRY_MOD),
    OM_OID_DESC(DS_ATTRIBUTE_TYPE, DSX_A_CDS_Class),
    {DS_ATTRIBUTE_VALUES, OM_S_PRINTABLE_STRING, \
    OM_STRING("Miscellaneous")},
    {DS_MOD_TYPE, OM_S_ENUMERATION, DS_ADD_ATTRIBUTE},
    OM_NULL_DESCRIPTOR
};

static OM_descriptor    Entry_Modification_List_Object[] = {
    OM_OID_DESC(OM_CLASS, DS_C_ENTRY_MOD_LIST),
    {DS_CHANGES, OM_S_OBJECT, {0, Entry_Modification_Object_1}},
    {DS_CHANGES, OM_S_OBJECT, {0, Entry_Modification_Object_2}},
    OM_NULL_DESCRIPTOR
};
```

A full description of this object can be found in Section 3.5. There could be any number of additional attribute changes in the list; this would mean additional **DS_C_ENTRY_MOD** objects declared, and an additional **DS_CHANGES** descriptor declared and initialized in the **DS_C_ENTRY_MOD_LIST** object.

With the **DS_C_ENTRY_MOD_LIST** class object having been declared as shown previously, the following code fragment illustrates how to call XDS to write a new attribute value (actually two new values since two attributes are contained in the list object). Note that any of the attributes may be new, although the entry itself must already exist.

```
dsStatus = ds_modify_entry(session,      /* Directory session */
                           /* from "ds_bind()" */
                           DS_DEFAULT_CONTEXT, /* Usual directory context */
                           Full_Entry_Name_Object, /* Entry name object */
                           Entry_Modification_List_Object, /* Entry Modification */
                           /* object */
                           &dummy);      /* Unsupported argument */
```

If the entire entry is new, you must call **ds_add_entry()**. This function requires an input object of class **DS_C_ATTRIBUTE_LIST**, whose contents specify the attributes (and values) to be attached to the new entry. Following is the static declaration for an attribute list that contains three attributes:

```
static OM_descriptor    Class_Attribute_Object[] = {
    OM_OID_DESC(OM_CLASS, DS_C_ATTRIBUTE),
    OM_OID_DESC(DS_ATTRIBUTE_TYPE, DSX_A_CDS_Class),
    {DS_ATTRIBUTE_VALUES, OM_S_PRINTABLE_STRING, OM_STRING("Printer")},
    OM_NULL_DESCRIPTOR
};

static OM_descriptor    ClassVersion_Attribute_Object[] = {
    OM_OID_DESC(OM_CLASS, DS_C_ATTRIBUTE),
    OM_OID_DESC(DS_ATTRIBUTE_TYPE, DSX_A_CDS_ClassVersion),
    {DS_ATTRIBUTE_VALUES, OM_S_PRINTABLE_STRING, OM_STRING("1.0")},
    OM_NULL_DESCRIPTOR
};

static OM_descriptor    My_Own_Attribute_Object[] = {
    OM_OID_DESC(OM_CLASS, DS_C_ATTRIBUTE),
    OM_OID_DESC(DS_ATTRIBUTE_TYPE, DSX_A_CDS_My_OwnAttribute),
    {DS_ATTRIBUTE_VALUES, OM_S_PRINTABLE_STRING, OM_STRING("zorro")},
    OM_NULL_DESCRIPTOR
};

static OM_descriptor    Attribute_List_Object[] = {
    OM_OID_DESC(OM_CLASS, DS_C_ATTRIBUTE_LIST),
    {DS_ATTRIBUTES, OM_S_OBJECT, {0, Class_Attribute_Object}},
    {DS_ATTRIBUTES, OM_S_OBJECT, {0, ClassVersion_Attribute_Object}},
    {DS_ATTRIBUTES, OM_S_OBJECT, {0, My_Own_Attribute_Object}},
    OM_NULL_DESCRIPTOR
};
```

```
};
```

The **ds_add_entry()** function also requires a **DS_C_DS_DN** class object containing the new entry's full name, for example:

```
./.../osf.org.dce/subsys/doc/my_book
```

where every member of the name exists except for the last one, **my_book**. Assuming that **Full_Entry_Name_Object** is a **DS_C_DS_DN** object, the following code shows what the call would look like:

```
dsStatus = ds_add_entry(session,      /* Directory session    */
                        /* from "ds_bind()"                */
                        DS_DEFAULT_CONTEXT, /* Usual directory context */
                        Full_Entry_Name_Object, /* Name of new entry      */
                        Attribute_List_Object, /* Attributes to be       */
                        /* attached to new entry, with values */
                        &dummy);          /* Unsupported argument   */
```

3.4 Object-Handling Techniques

The following subsections describe the use of XOM and discuss dynamic object creation.

3.4.1 Using XOM to Access CDS

The following code fragments demonstrate an alternative way to set up the entry modification object for a **ds_modify_entry()** call, mainly for the sake of showing how the **om_put()** and **om_write()** functions are used.

The following technique is used to initialize the modification object:

1. The **om_create()** function is called to generate a private object of a specified class.

2. The **om_put()** function is called to copy statically declared attributes into a declared private object.
3. The **om_write()** function is called to write the value string, which is to be assigned to the attribute, into the private object.
4. The **om_get()** function is called to make the private object public.
5. The object is now public, and its address is inserted into the **DS_C_ENTRY_MOD_LIST** object's **DS_CHANGES** attribute.

The following new declarations are necessary:

```
OM_private_object newAttributeMod_priv;
    /* ...handle to a private object to "om_put()" to */

OM_public_object newAttributeMod_pub;
    /* ...to hold public object from "om_get()" */

OM_type types_to_include[] = {DS_ATTRIBUTE_TYPE, DS_ATTRIBUTE_VALUES,
                             DS_MOD_TYPE, OM_NO_MORE_TYPES};
    /* ...that is, all attribute values of the Entry Modification */
    /* object. For "om_put()" and "om_get()" */

char *my_string = "O brave new attribute";
    /* ...value I want to write into attribute */

OM_value_position number_of_descriptors;
    /* ...to hold value returned by "om_get()" */
```

First, use XOM to generate a private object of the desired class:

```
omStatus = om_create(DS_C_ENTRY_MOD, /* Class of object */
                    OM_TRUE, /* Initialize attributes per defaults */
                    xdsWorkspace, /* Our workspace handle */
                    &newAttributeMod_priv); /* Created object handle */
```

Next, copy the public object's attributes into the private object:

```

omStatus = om_put(newAttributeMod_priv, /* Private object to copy */
                  /* attributes into */
                  OM_REPLACE_ALL, /* Which attributes to replace in */
                  /* destination object */
                  Entry_Modification_Object, /* Source object to copy */
                  /* attributes from */
                  types_to_include, /* List of attribute types we */
                  /* want copied */
                  0, 0); /* Start-stop index for multivalued */
                  /* attributes; ignored with OM_REPLACE_ALL */

```

Since **om_put()** ignores the class of the source object (the object from which attributes are being copied), it is not necessary to declare class descriptors for the source objects. In other words, the static declarations could have omitted the **OM_CLASS** initializations if this technique were being used, for example:

```

static OM_descriptor    Entry_Modification_Object_2[] = {

/*      OM_OID_DESC(OM_CLASS, DS_C_ENTRY_MOD),          */
/*      Not needed for "om_put()" ...                    */

    OM_OID_DESC(DS_ATTRIBUTE_TYPE, DSX_A_CDS_Class),
    {DS_ATTRIBUTE_VALUES, OM_S_PRINTABLE_STRING, \
      OM_STRING("Miscellaneous")},
    {DS_MOD_TYPE, OM_S_ENUMERATION, DS_ADD_ATTRIBUTE},
    OM_NULL_DESCRIPTOR
};

```

The **OM_CLASS** was already properly initialized by **om_create()**.

Next, write the attribute value string into the private object:

```

omStatus = om_write(newAttributeMod_priv, /* Private object to write to */
                    DS_ATTRIBUTE_VALUES, /* Attribute type whose value */
                    /* we're writing */
                    0, /* Descriptor index if attribute is multivalued */
                    OM_S_PRINTABLE_STRING, /* Syntax of value */
                    0, /* Offset in source string to write from */
                    my_string); /* Source string to write from */

```

Now make the whole thing public again:

```
omStatus = om_get(newAttributeMod_priv, /* Private object to get      */
                  0, /* Get everything                                */
                  types_to_include, /* All attribute types            */
                  0, /* Unsupported argument                        */
                  0, 0, /* Start-stop descriptor index for multival- */
                      /* ued attributes; ignored in this case      */
                  &newAttributeMod_pub, /* Pointer to returned copy    */
                  &number_of_descriptors); /* Number of attribute        */
                                           /* descriptors returned        */
```

Finally, insert the address of the subobject into its superobject:

```
Entry_Modification_List_Object[1].value.object.object = \
    newAttributeMod_pub;
```

3.4.2 Dynamic Creation of Objects

Objects can be completely dynamically allocated and initialized; however, you have to implement the routines to do this yourself. The examples in this section are code fragments; for complete examples, see Chapter 7.

Initialization of object structures can be automated by declaring macros or functions to do this. For example, the following macro initializes one object descriptor with a full set of appropriate values:

```
/* Put a C-style (NULL-terminated) string into an object and */
/* set all the other descriptor fields to requested values    */
#define FILL_OMD_STRING( desc, index, typ, syntax, val ) \
    desc[index].type = typ; \
    desc[index].syntax = syntax; \
    desc[index].value.string.length = \
        (OM_element_position)strlen(val); \
    desc[index].value.string.elements = val;
```

When generating objects, use **malloc()** to allocate space for the number of objects desired, and then use macros (or functions) such as the preceding one to initialize the descriptors. The following code fragment shows how this can be done for the top-level object of a **DS_C_DS_DN** object, such as the one described near the beginning of this chapter. Recall that **DS_C_DS_DN** has a separate **DS_RDNS** descriptor for each name piece in the full name.

```
/* Calculate number of "DS_RDNS" attributes there should be ... */
numberOfPieces = number_of_name_pieces;

/* Allocate space for that many descriptors, plus one for the
/* object class at the front, and a NULL descriptor at the back */

Name_Object = (OM_object)malloc((numberOfPieces + 2) \
    * sizeof(OM_descriptor));
if(Name_Object == NULL)                /* "malloc()" failed */
    return OM_MEMORY_INSUFFICIENT;

/* Initialize it as a DS_C_DS_DN object by placing that class
/* identifier in the first position... */

FILL_OMD_XOM_STRING(Name_Object, 0, OM_CLASS,
    OM_S_OBJECT_IDENTIFIER_STRING, DS_C_DS_DN)
```

Note that all these steps would have to be repeated for each of the **DS_C_DS_RDN** objects required as attribute values of the **DS_C_DS_DN**. Then a tier of **DS_C_AVA** objects would have to be created in the same way, since each of the **DS_C_DS_RDNs** requires one of them as *its* attribute value.

You could now use **om_create()** and **om_put()** to generate a private copy of this object, if so desired.

The application is responsible for managing the memory it allocates for such dynamic object creation.

3.5 XDS/CDS Object Recipes

The following subsections contain shorthand for object classes. For example, if you look at the reference pages for the **ds_***() functions, you will see that an object of class **DS_C_NAME** is required to hold entry names you want to pass to the call, *not* **DS_C_DS_DN** as is stated in this chapter. However, **DS_C_NAME** is in fact an abstract class with only one subclass **DS_C_DS_DN** so, in this chapter, **DS_C_DS_DN** is used.

3.5.1 Input XDS/CDS Object Recipes

In general, the objects you work with in an XDS/CDS application fall into two categories:

- Objects you have to supply as *input parameters* to XDS functions
- Objects returned to you as *output* by XDS functions

This section describes only the first category, since you have to construct these input objects yourself.

Table 3-1 shows XDS functions and the objects given to them as input parameters.

Only items significant to CDS are listed in the table. **DS_C_SESSION** and **DS_C_CONTEXT** are ignored. **DS_C_SESSION** is returned by **ds_bind()**, which usually receives the **DS_DEFAULT_SESSION** constant as input. **DS_C_CONTEXT** is usually substituted by the **DS_DEFAULT_CONTEXT** constant.

Note: **DS_C_NAME** is an abstract class that has the single subclass **DS_C_DS_DN**. Therefore, **DS_C_NAME** is practically the same thing as **DS_C_DS_DN**.

Table 3–1. Directory Service Functions With Their Required Input Objects

Function	Input Object
ds_add_entry()	DS_C_NAME
	DS_C_ATTRIBUTE_LIST

Function	Input Object
ds_bind()	None
ds_compare()	DS_C_NAME DS_C_AVA
ds_initialize()	None
ds_list()	DS_C_NAME
ds_modify_entry()	DS_C_NAME DS_C_ENTRY_MOD_LIST
ds_read()	DS_C_NAME DS_C_ENTRY_INFO_SELECTION
ds_remove_entry()	DS_C_NAME
ds_shutdown()	None
ds_unbind()	None
ds_version()	None

3.5.2 Input Object Classes for XDS/CDS Operations

The following subsections contain information about all the object types required as input to any of the XDS functions that can be used to access CDS. In order to use these functions successfully, you must be able to construct and modify the objects that the functions expect as their input parameters. XDS functions require most of their input parameters to be wrapped in a nested series of data structures that represent objects, and these functions deliver their output returns to callers in the same object form.

Objects that are returned to you by the interface are not difficult to manipulate because the **om_get()** function allows you to go through them and retrieve only the value parts you are interested in, and discard the parts of data structures you are not interested in. However, any objects you are required to supply as *input* to an XDS or XOM function are another matter: you must build and initialize these object structures yourself.

The basics of object building have already been explained earlier in this chapter. Each object described in the following subsections is accompanied by a static declaration in C of a very simple instance of that object class. The objects in an application are usually built dynamically (this technique was demonstrated earlier in this chapter). The static declarations that follow give a simple example of what the objects look like.

An object's properties, such as what sort of values it can hold, how many of them it can hold, and so on, are determined by the *class* the object belongs to. Each class consists of one or more *attributes* that an object can have. The attributes hold whatever values the object contains. Thus, the objects are data structures that all look the same (and can be handled in the same way) from the outside, but whose specific data fields are determined by the class each object belongs to. At the abstract level, objects consist of attributes, just as structures consist of fields.

3.5.2.1 XDS/CDS Object Types

Following is a list of all the object types that are described in the following subsections. Most of these objects are object structures; that is, compounds consisting of superobjects that contain subobjects as some of their values. These subobjects may in turn contain other objects, and so on. Subobjects are indicated by indentation. A **DS_C_DS_DN** object contains at least one **DS_C_DS_RDN** object, and each **DS_C_DS_RDN** contains one **DS_C_AVA** object. Note that subobjects can, and often do, exist by themselves, depending on what object class is called for by a given function. This list contains all the possible kinds of objects that can be required as input for any XDS/CDS operation.

- **DS_C_ATTRIBUTE_LIST**
 - **DS_C_ATTRIBUTE**
- **DS_C_DS_DN**
 - **DS_C_DS_RDN**
 - DS_C_AVA**
- **DS_C_ENTRY_MOD_LIST**
 - **DS_C_ENTRY_MOD**
- **DS_C_ENTRY_INFO_SELECTION**

In each section, information is provided for the described object's attributes. All its attributes are listed.

The illustrations in the following sections can be compared to the same object classes' tabular definitions later in this guide.

3.5.2.2 The **DS_C_ATTRIBUTE_LIST** Object

A **DS_C_ATTRIBUTE_LIST** class object is required as input to **ds_add_entry()**. The object contains a list of the directory attributes you want associated with the entry that is to be added.

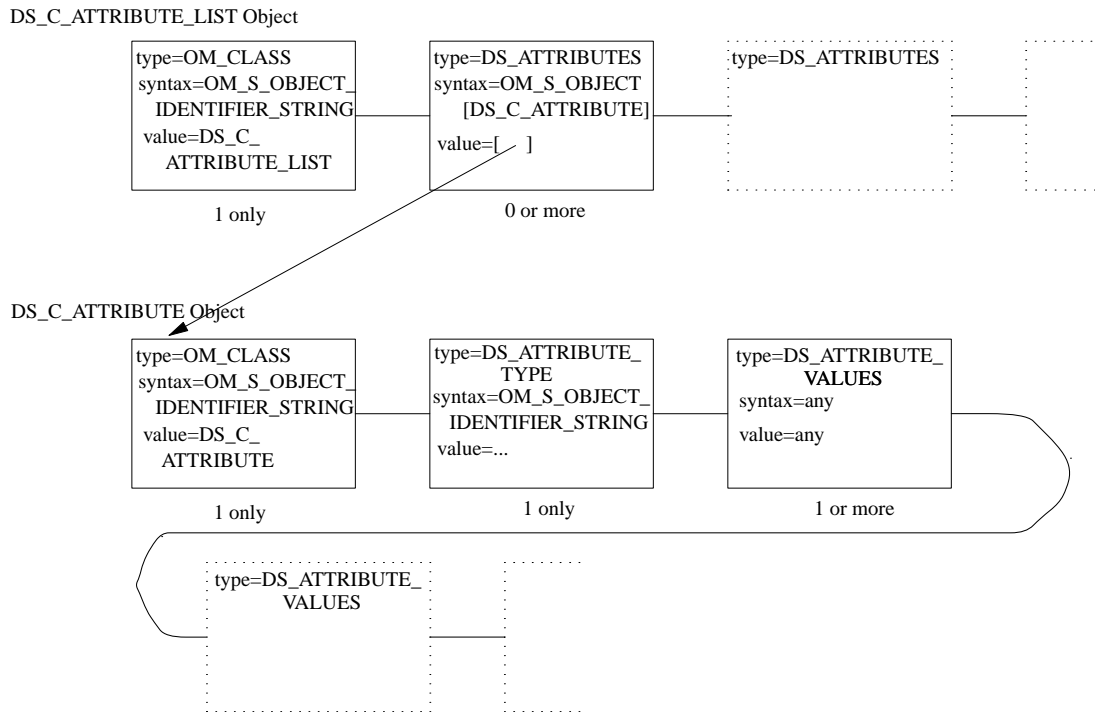
Its general structure is as follows:

- Attribute List class type attribute
- Zero or more Attribute objects:
 - Attribute class type attribute
 - Attribute Type attribute
 - Zero or more Attribute Value(s)

Thus, a **DS_C_ATTRIBUTE_LIST** object containing one attribute consists of two object descriptor arrays because each additional attribute in the list requires an additional descriptor array to represent it. The subobject arrays' names (that is, addresses) are the contents of the value fields in the **DS_ATTRIBUTES** object descriptors.

Figure 3-9 shows the attributes of the **DS_C_ATTRIBUTE_LIST** object.

Figure 3–9. The DS_C_ATTRIBUTE_LIST Object



- **OM_CLASS**

The value of this attribute is an OID string that identifies the object's class; its value is always **DS_C_ATTRIBUTE_LIST**.

- **DS_ATTRIBUTES**

This is an attribute whose value is another object of class **DS_C_ATTRIBUTE** (see Section 3.5.2.3). The attribute is defined by a separate array of object descriptors whose base address is the value of the **DS_ATTRIBUTES** attribute. Note that there can be any number of instances of this attribute and, therefore, any number of subobjects.

3.5.2.3 The **DS_C_ATTRIBUTE** Object

An object of this class can be an attribute of a **DS_C_ATTRIBUTE_LIST** object (see Section 3.5.2.2).

- **OM_CLASS**

The value of this attribute is an OID string that identifies the object's class; its value is always **DS_C_ATTRIBUTE**.

- **DS_ATTRIBUTE_TYPE**

The value of this attribute, which is an OID string, identifies the directory attribute whose value is contained in this object.

- **DS_ATTRIBUTE_VALUES**

These are the actual values for the directory attribute represented by this **DS_C_ATTRIBUTE** object. Both the value syntax and the number of values depend on what directory attribute this is; that is, they depend on the value of **DS_ATTRIBUTE_VALUE**.

3.5.2.4 Example Definition of a **DS_C_ATTRIBUTE_LIST** Object

The following code fragment is a definition of a **DS_C_ATTRIBUTE_LIST** object.

```
static OM_descriptor    Single_Attribute_Object[] = {
    OM_OID_DESC(OM_CLASS, DS_C_ATTRIBUTE),
    OM_OID_DESC(DS_ATTRIBUTE_TYPE, DSX_A_CDS_Class),
    {DS_ATTRIBUTE_VALUES, OM_S_PRINTABLE_STRING, \
     OM_STRING("Printer")},
    OM_NULL_DESCRIPTOR
};
static OM_descriptor    Attribute_List_Object[] = {
    OM_OID_DESC(OM_CLASS, DS_C_ATTRIBUTE_LIST),
    {DS_ATTRIBUTES, OM_S_OBJECT, {0, Single_Attribute_Object}},
    OM_NULL_DESCRIPTOR
};
```

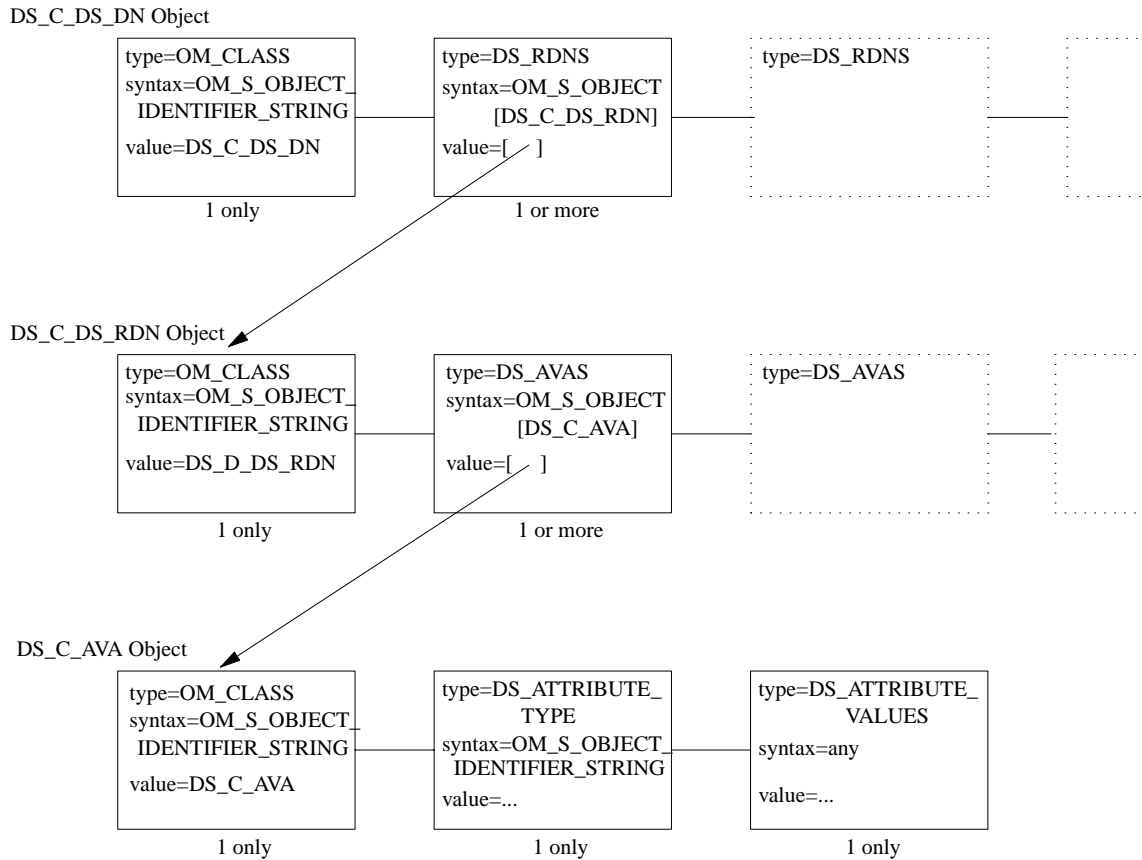
3.5.2.5 The **DS_C_DS_DN** Object

DS_C_DS_DN class objects are used to hold the full names of directory entries (distinguished names). You need an object of this class to pass directory entry names to the following XDS functions:

- **ds_add_entry()**
- **ds_compare()**
- **ds_list()**
- **ds_modify_entry()**
- **ds_read()**
- **ds_remove_entry()**

Figure 3-10 shows the attributes of a **DS_C_DS_DN** object.

Figure 3–10. DS_C_DS_DN Object Attributes



- **OM_CLASS**

The value of this attribute is an OID string that identifies the object's class; its value is **DS_C_DS_DN**.

- **DS_RDNS**

This is an attribute whose value is another object of class **DS_C_DS_RDN** (see Section 3.5.2.6). The **DS_C_DS_RDN** object is defined by a separate array of object descriptors whose base address is the value of the **DS_RDNS** attribute.

There are as many **DS_RDNS** attributes in a **DS_C_DS_DN** object as there are separate name components in the full directory entry name. For example, suppose you wanted to represent the following CDS entry name:

/.../C=US/O=OSF/OU=DCE/hosts/brazil/self

This would require a total of six instances of the **DS_RDNS** attribute in the **DS_C_DS_DN** object. The **/.../** (global root prefix) is not represented. This means that another six object descriptor arrays are required to hold the RDN objects, as well as six object descriptors in the present object, one to hold (as the value of a **DS_RDNS** attribute) a pointer to each array.

Note that the order of these **DS_RDNS** attributes is significant; that is, the first **DS_RDNS** should contain as its value a pointer to the array representing the **C=US** part of the name; the next **DS_RDNS** should contain as its value a pointer to the array representing the **O=OSF** part, and so on. The root part of the name is not represented at all.

3.5.2.6 The **DS_C_DS_RDN** Object

DS_C_DS_RDN class objects are required as values for the **DS_RDNS** attributes of **DS_C_DS_DN** objects. (For an illustration of its structure, see Figure 3-10.) **RDN** refers to the X.500 term *RDN* that is used to signify a part of a full entry name. Separate objects of this class are not usually required as input to XDS functions.

The standard permits multiple AVAs in an RDN, but the DCE Directory and XDS API restrict an RDN to one AVA.

- **OM_CLASS**

The value of this attribute is an OID string that identifies the object's class; its value is always **DS_C_DS_RDN**.

- **DS_AVAS**

This is an attribute whose value is yet another object of class **DS_C_AVA** (see Section 3.5.2.7). The **DS_C_AVA** object is defined by a separate array of object descriptors whose base address is the value of the **DS_AVAS** attribute.

Note that there can only be one instance of this attribute in the **DS_C_RDN** object. The object descriptor array describing this object always consists of three object descriptor structures: the first describes the object's class, the second describes the **DS_AVAS** attribute, and the third descriptor is the terminating NULL.

3.5.2.7 The **DS_C_AVA** Object

The **DS_C_AVA** class object is used to hold an actual value. The value is usually in the form of one of the many different XOM string types. (For an illustration of its structure, see Figure 3-10.)

In calls to **ds_compare()**, an object of this type is required to hold the type and value of the attribute that you want compared with those in the entry you specify. It holds the type and value in a separate **DS_C_DS_DN** object.

DS_C_AVA is also included here because it is a required subsubobject of **DS_C_DS_DN** itself. **DS_C_AVA** is the subobject in which the name part's actual literal value is held.

- **OM_CLASS**

The value of this attribute is an OID string that identifies the object's class; its value is always **DS_C_AVA**.

- **DS_ATTRIBUTE_TYPE**

The value of this attribute, which is an OID string, identifies the directory attribute whose value is contained in this object.

- **DS_ATTRIBUTE_VALUES**

This is the literal value of what is represented by this **DS_C_AVA** object.

If the **DS_C_AVA** object is a subobject of **DS_C_DS_RDN** (and therefore also of **DS_C_DS_DN**), then the value is a string representing the part of the directory entry name represented by this object. For example, if the **DS_C_DS_RDN** object contains the **O=OSF** part of an entry name, then the string **OSF** is the value of the **DS_ATTRIBUTE_VALUES** attribute, and **DS_A_COUNTRY_NAME** is the value of the **DS_ATTRIBUTE_TYPE** attribute.

On the other hand, if **DS_C_AVA** contains an entry attribute type and value to be passed to **ds_compare()**, then **DS_ATTRIBUTE_TYPE** identifies the type of the attribute, and **DS_ATTRIBUTE_VALUES** contains a value, which is appropriate for the attribute type, to be compared with the entry value.

For example, suppose you wanted to compare a certain value with a CDS entry's **CDS_Class** attribute's value. The identifiers for all the valid CDS entry attributes are found in the file `./opt/dcelocal/etc/cds_attributes`. The value of **DS_ATTRIBUTE_TYPE** would be **CDS_Class**, which is the label of an object identifier string, and **DS_ATTRIBUTE_VALUES** would contain some desired value, in the correct syntax for **CDS_Class**. The syntax also is found in the `cds_attributes` file; for **CDS_Class** it is **byte**; that is, a character string.

3.5.2.8 Example Definition of a **DS_C_DS_DN** Object

The following code fragment shows an example definition for a **DS_C_DS_DN** object.

```
static OM_descriptor    Entry_String_Object[] = {
    OM_OID_DESC(OM_CLASS, DS_C_AVA),
    OM_OID_DESC(DS_ATTRIBUTE_TYPE, DSX_TYPELESS_RDN),
    {DS_ATTRIBUTE_VALUES, OM_S_PRINTABLE_STRING, \
     OM_STRING("brazil")},
    OM_NULL_DESCRIPTOR
};

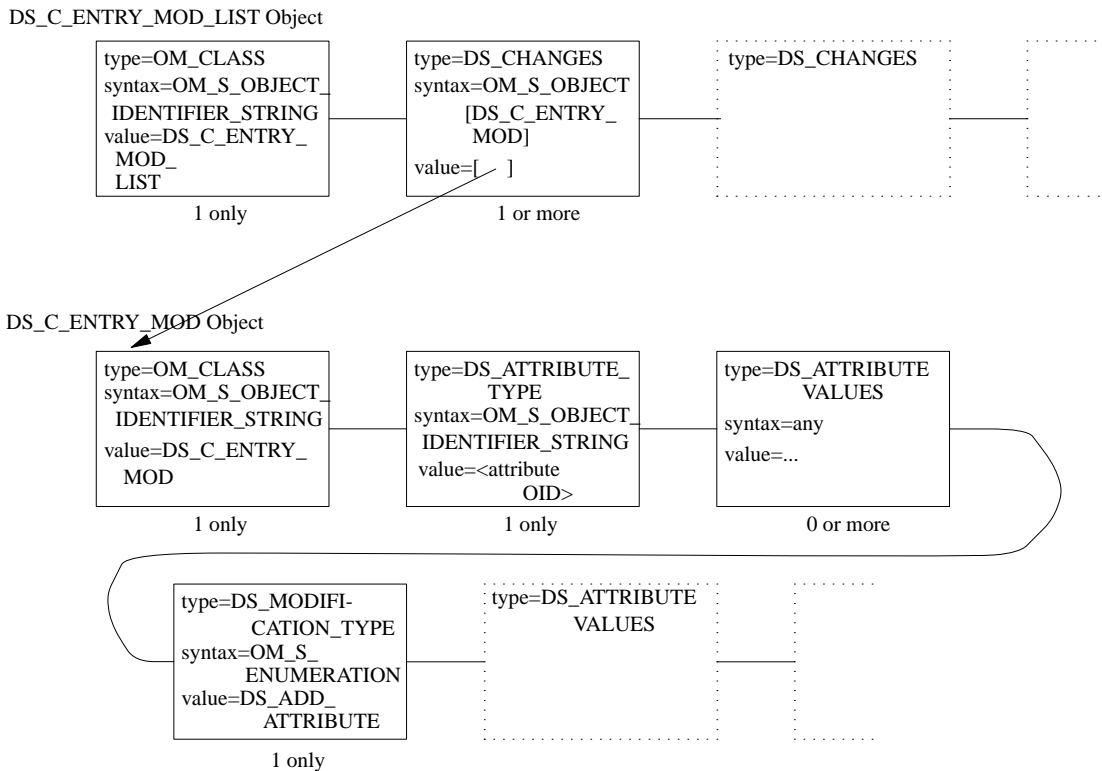
static OM_descriptor    Entry_Part_Object[] = {
    OM_OID_DESC(OM_CLASS, DS_C_DS_RDN),
    {DS_AVAS, OM_S_OBJECT, {0, Entry_String_Object}},
    OM_NULL_DESCRIPTOR
};

static OM_descriptor    Entry_Name_Object[] = {
    OM_OID_DESC(OM_CLASS, DS_C_DS_DN),
    {DS_RDNS, OM_S_OBJECT, {0, Entry_Part_Object}},
    OM_NULL_DESCRIPTOR
};
```

3.5.2.9 The DS_C_ENTRY_MOD_LIST Object

DS_C_ENTRY_MOD_LIST class objects, which contain a list of changes to be made to some directory entry, must be passed to **ds_modify_entry()**. **DS_C_ENTRY_MOD_LIST** objects have the attributes shown in Figure 3-11.

Figure 3–11. The DS_C_ENTRY_MOD_LIST Object



- **OM_CLASS**

The value of this attribute is an OID string that identifies the object's class; its value is always **DS_C_ENTRY_MOD_LIST**.

- **DS_CHANGES**

This is an attribute whose value is another object of class **DS_C_ENTRY_MOD** (see Section 3.5.2.10). The **DS_C_ENTRY_MOD** object is defined by a separate array of object descriptors whose base address is the value of the **DS_CHANGES** attribute.

Note that there can be one or more instances of this attribute in the object, which is why it is called **_LIST**. Each attribute contains one separate entry modification. To learn how the modification itself is specified, see Section 3.5.2.10. The order of multiple instances of this attribute is significant because, if more than one modification is specified, the modifications are performed by **ds_modify_entry()** in the order in which the **DS_CHANGES** attributes appear in the **DS_C_ENTRY_MOD_LIST** object.

3.5.2.10 The **DS_C_ENTRY_MOD** Object

The **DS_C_ENTRY_MOD** class object holds the information associated with a directory entry modification. (For an illustration of its structure, see Figure 3-11.) Each **DS_C_ENTRY_MOD** object describes one modification. To create a list of modifications suitable to be passed to a **ds_modify_entry()** call, describe each modification in a separate **DS_C_ENTRY_MOD** object, and then insert these objects as multiple instances of the **DS_CHANGES** attribute in a **DS_C_ENTRY_MOD_LIST** object (see Section 3.5.2.9).

- **OM_CLASS**

The value of this attribute is an OID string that identifies the object's class; its value is always **DS_C_ENTRY_MOD**.

- **DS_ATTRIBUTE_TYPE**

The value of this attribute, which is an OID string, identifies the directory attribute whose modification is described in this object.

- **DS_ATTRIBUTE_VALUES**

These are the values required for the entry modification; their type and number depend on both the entry type and the modification requested.

- **DS_MOD_TYPE**

The value of this attribute identifies the kind of modification requested. It can be one of the following:

— **DS_ADD_ATTRIBUTE**

The attribute specified by **DS_ATTRIBUTE_TYPE** is not currently in the entry. It should be added, along with the value(s) specified by **DS_ATTRIBUTE_VALUES**, to the entry. The entry itself is specified in a separate **DS_C_DS_DN** object, which is also passed to **ds_modify_entry()**.

— **DS_ADD_VALUES**

The specified attribute is currently in the entry. The value(s) specified by **DS_ATTRIBUTE_VALUES** should be added to it.

— **DS_REMOVE_ATTRIBUTE**

The specified attribute is currently in the entry and should be deleted from the entry. Any values specified by **DS_ATTRIBUTE_VALUES** are ignored.

— **DS_REMOVE_VALUES**

The specified attribute is currently in the entry. One or more values, specified by **DS_ATTRIBUTE_VALUES**, should be removed from it.

3.5.2.11 Example Definition of a **DS_C_ENTRY_MOD_LIST** Object

The following code fragment is an example definition of a **DS_C_ENTRY_MOD_LIST** object.

```
OM_string my_uuid;

static OM_descriptor Entry_Mod_Object[] = {
    OM_OID_DESC(OM_CLASS, DS_C_ENTRY_MOD),
    OM_OID_DESC(DS_ATTRIBUTE_TYPE, DSX_UUID),
    {DS_ATTRIBUTE_VALUES, OM_S_OCTET_STRING, my_uuid},
    {DS_MOD_TYPE, OM_S_ENUMERATION, DS_ADD_ATTRIBUTE},
    OM_NULL_DESCRIPTOR
};
```

```
static OM_descriptor    Entry_Mod_List_Object[] = {
    OM_OID_DESC(OM_CLASS, DS_C_ENTRY_MOD_LIST),
    {DS_CHANGES, OM_S_OBJECT, {0, Entry_Mod_Object}},
    OM_NULL_DESCRIPTOR
};
```

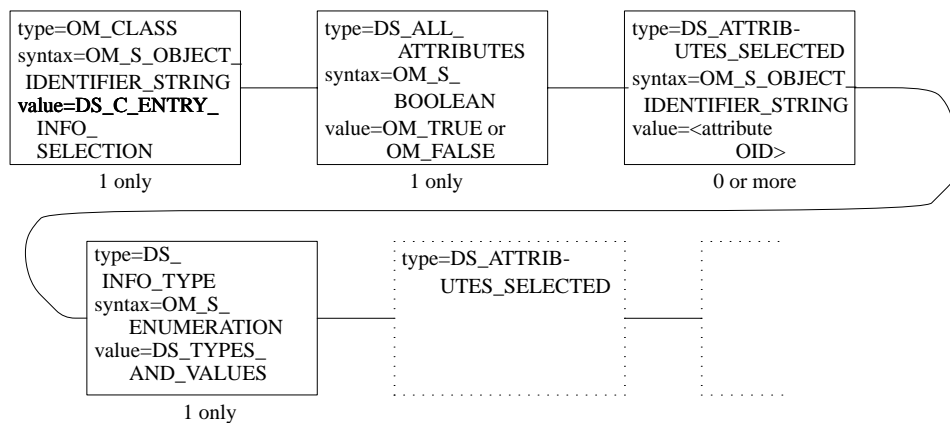
3.5.2.12 The **DS_C_ENTRY_INFO_SELECTION** Object

When you call **ds_read()** to read one or more attributes from a CDS entry, you specify in the **DS_C_ENTRY_INFO_SELECTION** object the entry attributes you want to read.

The **DS_C_ENTRY_INFO_SELECTION** object contains the attributes shown in Figure 3-12.

Figure 3-12. The **DS_C_ENTRY_INFO_SELECTION** Object

DS_C_ENTRY_INFO_SELECTION Object



Note that this object class has no subobjects.

- **OM_CLASS**

The value of this attribute is an OID string that identifies the object's class; its value is always **DS_C_ENTRY_INFO_SELECTION**.

- **DS_ALL_ATTRIBUTES**

This attribute is a simple Boolean option whose value indicates whether all the entry's attributes are to be read, or only some of them. Its possible values are as follows:

- **OM_TRUE**, meaning that all attributes in the directory entry should be read. Any values specified by the **DS_ATTRIBUTES_SELECTED** attribute are ignored.
- **OM_FALSE**, meaning that only some of the entry attributes should be read; namely, those specified by the **DS_ATTRIBUTES_SELECTED** attribute.

- **DS_ATTRIBUTES_SELECTED**

The value of this attribute, which is an OID string, identifies the entry attribute to be read. Note that this attribute's value has meaning only if the value of **DS_ALL_ATTRIBUTES** is **OM_FALSE**; if it is **OM_TRUE**, the value of **DS_ATTRIBUTES_SELECTED** is ignored.

Note also that there are multiple instances of this attribute if more than one attribute, but not all of them, is to be selected for reading. Each separate instance of **DS_ATTRIBUTES_SELECTED** has as its value an OID string that identifies one directory entry attribute to be read. If **DS_ATTRIBUTES_SELECTED** is present but does not have a value, **ds_read()** reads the entry but does not return any attribute data; this technique can be used to verify the existence of a directory entry.

- **DS_INFO_TYPE**

The value of this attribute specifies what information is to be read from each attribute specified by **DS_ATTRIBUTES_SELECTED**. The two possible values are as follows:

- **DS_TYPES_ONLY**, meaning that only the attribute types of the selected attributes should be read.
- **DS_TYPES_AND_VALUES**, meaning that both the attribute types and the attribute values of the selected attributes should be read.

3.5.2.13 Example Definition of a **DS_C_ENTRY_INFO_SELECTION** Object

The following code fragment provides an example definition of a **DS_C_ENTRY_INFO_SELECTION** object.

```
static OM_descriptor    Entry_Info_Select_Object[] = {
    OM_OID_DESC(OM_CLASS, DS_C_ENTRY_INFO_SELECTION),
    OM_OID_DESC(DS_ATTRIBUTES_SELECTED, DSX_A_CDS_Class),
    {DS_ALL_ATTRIBUTES, OM_S_BOOLEAN, OM_FALSE},
    {DS_INFO_TYPE, OM_S_ENUMERATION, DS_TYPES_AND_VALUES},
    OM_NULL_DESCRIPTOR
};
```

3.6 Attribute and Data Type Translation

This section provides translations between CDS and XDS for attributes and data types. Table 3-2 provides the OM syntax for CDS attributes. Table 3-3 provides the OM syntax for CDS data types. Table 3-4 defines the mapping of CDS data types to OM syntaxes.

Table 3–2. CDS Attributes to OM Syntax Translation

CDS Attribute	OM Syntax
CDS_CTS	OM_S_OCTET_STRING
CDS_UTS	OM_S_OCTET_STRING
CDS_Class	OM_S_OCTET_STRING
CDS_ClassVersion	OM_S_INTEGER
CDS_ObjectUID	OM_S_OCTET_STRING
CDS_AllUpTo	OM_S_OCTET_STRING
CDS_Convergence	OM_S_INTEGER
CDS_InCHName	OM_S_INTEGER

CDS Attribute	OM Syntax
CDS_DirectoryVersion	OM_S_INTEGER
CDS_UpgradeTo	OM_S_INTEGER
CDS_LinkTimeout	OM_S_INTEGER
CDS_Towers	OM_S_OCTET_STRING

Table 3–3. OM Syntax to CDS Data Types Translation

OM Syntax	CDS Data Type
OM_S_TELETEX_STRING	cds_char
OM_S_OBJECT_IDENTIFIER_STRING	cds_byte
OM_S_OCTET_STRING	cds_byte
OM_S_PRINTABLE_STRING	cds_char
OM_S_NUMERIC_STRING	cds_char
OM_S_BOOLEAN	cds_long
OM_S_INTEGER	cds_long
OM_S_UTC_TIME_STRING	cds_char
OM_S_ENCODING_STRING	cds_byte

Table 3–4. CDS Data Types to OM Syntax Translation

CDS Data Type	OM Syntax
cds_none	OM_S_NULL
cds_long	OM_S_INTEGER
cds_short	OM_S_INTEGER
cds_small	OM_S_INTEGER
cds_uuid	OM_S_OCTET_STRING
cds_Timestamp	OM_S_OCTET_STRING
cds_Version	OM_S_PRINTABLE_STRING
cds_char	OM_S_TELETEX_STRING
cds_byte	OM_S_OCTET_STRING

Part 3

GDS Application Programming

Part 3 is an overview of programming GDS using XDS.

Chapter 4 discusses GDS concepts and gives an overview of GDS programming. Chapter 5 describes XOM programming, and Chapter 6 describes XDS programming. Chapter 7 contains programming examples. Chapter 8 describes how to use threads with XDS and XOM, and Chapter 9 describes the XDS and XOM convenience routines.

Chapter 4

GDS API: Concepts and Overview

The Global Directory Service (GDS) is a distributed, replicated directory service. It is distributed because information is stored in different places in the network. Requests for information may be routed by GDS to directory servers throughout the network. It is replicated because information can be stored in more than one location for easier and more efficient access by its users.

GDS is based on the CCITT X.500/ISO 9594 (1988) international standard. The aim of this standard, also referred to as the OSI Directory standard, is to provide a global directory that supports network users and applications with information required for communications. The directory plays a significant role in allowing the interconnection of information processing systems from different manufacturers, under different managements, of different levels of complexity, and of different ages.

GDS is the DCE implementation of the OSI Directory standard. Together with the Cell Directory Service (CDS), it provides its users with a centralized place to store information required for communications, which can be retrieved from anywhere in a distributed system. GDS maintains information describing objects such as people, organizations, applications, distribution lists, network hardware, and other distributed services dispersed over a large geographical area.

CDS stores names and attributes of resources located in a DCE cell. A DCE cell consists of various combinations of DCE machines connected by a network. Each DCE cell contains its own cell directory server, which provides access to local resource information. CDS is optimized for local information access by its users. For a more detailed description of cells and their resource services, see the *Introduction to OSF DCE*.

GDS serves as a general-purpose information repository. It provides information about resources outside a DCE cell. It ties together the various cells by helping to find remote cells. A detailed discussion of the DCE namespace and its various servers and their interaction is provided in Chapter 1.

4.1 Directory Service Interfaces


X/Open Directory Service (XDS) and X/Open OSI-Abstract-Data Manipulation (XOM) are application programming interfaces. XOM and XDS application interfaces are based on X/Open standards specifications. Together these interfaces provide the application programmer with a library of functions with which to develop applications that access the directory service.

The XOM application programming interface (XOM API) is an interface for creating, deleting, and accessing information objects. The XOM API defines an object-oriented information model. Objects belong to classes and have attributes associated with them. The XOM API also defines basic data types, such as Boolean, string, object, and so on. The representation of these objects are transparent to the programmer. Objects can only be manipulated through the XOM interface, not directly.

DCE programmers use the XDS API to make directory service calls. In DCE, the XDS API directs the calls it receives to either GDS or CDS by examining the names of the information objects to be looked up as shown in Figure 4-1. It uses the XOM interface for defining and handling information objects. These objects are passed as parameters and return values to the XDS routines. The XDS API contains functions for managing connections with a directory server: reading, comparing, adding, removing, modifying, listing, and searching for directory entries. The GDS package provides additional information objects that provide for security and cache management when using GDS.

GDS supports additional functions, called *convenience functions*, at the XDS/XOM API. These functions, described in Chapter 9, provide GDS programmers with a toolkit to allow more efficient production of XDS/XOM based applications.

Figure 4–1. XDS: Interface to GDS and CDS



4.2 The X.500 Directory Information Model

This section describes the directory information model of X.500, which GDS is based on. A directory is a collection of information about some part of the world. The most familiar type of directory is the list of names and numbers that make up a city telephone directory. A name is provided with some information about the named object, such as an address and telephone number. The ISO and CCITT standards define a *directory information model* that defines the abstract structure of directory information, services, and protocols for a computer network environment, such as DCE.

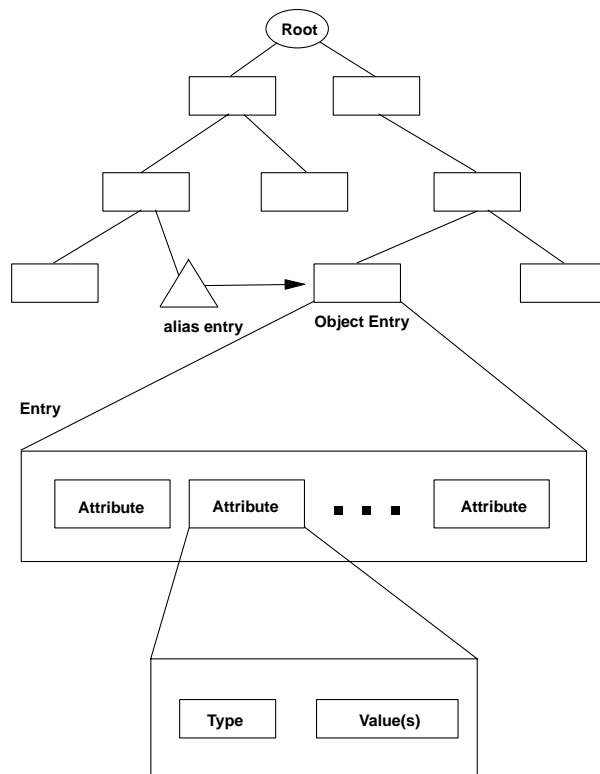
4.2.1 Directory Objects

The directory contains information about objects. The standard defines an object very broadly as “anything in some ‘world,’ generally the world of telecommunications and information processing or some part thereof, which is identifiable (can be named).” Some examples of objects include people, corporations, and application processes.

Each object known to the directory is represented by an entry. The set of all entries is called the Directory Information Base (DIB), which is a hierarchical tree. Each entry consists of a set of attributes representing specific information about the object. Each attribute, in turn, has a type and one or more values of that type. Attributes with more than one value are referred to as *multivalued* or *recurring* attributes.

Figure 4-2 shows the structure of the DIB.

Figure 4-2. The Structure of the DIB



The attributes that constitute a single entry can be of various types. For example, an entry for a person may contain that person's name, address, and phone number. If the person has a second telephone number, the attribute of type telephone number may have two values, one for each telephone number.

Object entries are composed of mandatory and optional attributes. Mandatory and optional attributes are discussed in Section 4.4.3.

4.2.2 Attribute Types


All attributes in a particular entry must be of different attribute types. Each attribute type is assigned a unique object identifier value. The directory standard assigns object identifiers for several commonly used attribute types, including surname, country name, telephone number, and presentation address. Other international standards may define additional attribute types. For example, the X.400 Message Handling standard defines mail-specific attributes like O/R address. It is expected that various national and private organizations will also define attribute types of their own. The CDS attributes (defined in the **xdscds.h** header file) and the GDS package attributes (defined in the **xdsgds.h** header file) are examples of additional attribute definitions.

4.2.3 Object Identifiers

Objects in a network environment, such as DCE, require unique names to distinguish them from one another. To provide these names, object identifiers are allocated by an administrative organization, such as a standards body. An object identifier is a hierarchical sequence of numbers uniquely identifying an object. Associated with each object identifier is a character string to make it easier to document.

The possible values of object identifiers are defined in a tree. Part of this tree is shown in Figure 4-3. It begins with three numbered branches coming from the root: branch 0 (assigned to CCITT), branch 1 (assigned to ISO), and branch 2 (a joint ISO-CCITT branch). Below each of these branches are other numbered branches assigned to various standards such as the directory service (**ds(5)**) and electronic mail service (**mhs-motis(6)**) with each ending in a named object. Thus, the name of any of these objects is a series of integers describing a path down this tree to the leaf node.

Figure 4–3. Object Identifiers



The object identifier associated with the XDS package is defined as follows:

**{iso(1) identified-organization(3) icd-ecma(12) member-company(2)
dec(1011) xopen(28) dsp(0)}**

All object classes and object attributes in the directory service package have these numbered branches associated with them. The classes and attributes, in turn, have their own unique numbers. These object identifiers are defined in header files included as part of the XDS and XOM API software. For example, the attribute type **Common-Name** is identified by the object identifier 2.5.4.3.

Table 4-1 contains a sample list of object identifiers for selected attributes. The complete list is provided in Chapter 12.

Table 4-1. Object Identifiers for Selected Attribute Types

Attribute Type	Object Identifier
Aliased-Object-Name	2.5.4.1
Business-Category	2.5.4.15
Common-Name	2.5.4.3
Country-Name	2.5.4.6
Description	2.5.4.13

Note: The object identifiers in Table 4-1 stem from the root **{joint-iso-ccitt(2) ds(5) attributeType(4)}**.

4.2.4 Object Entries

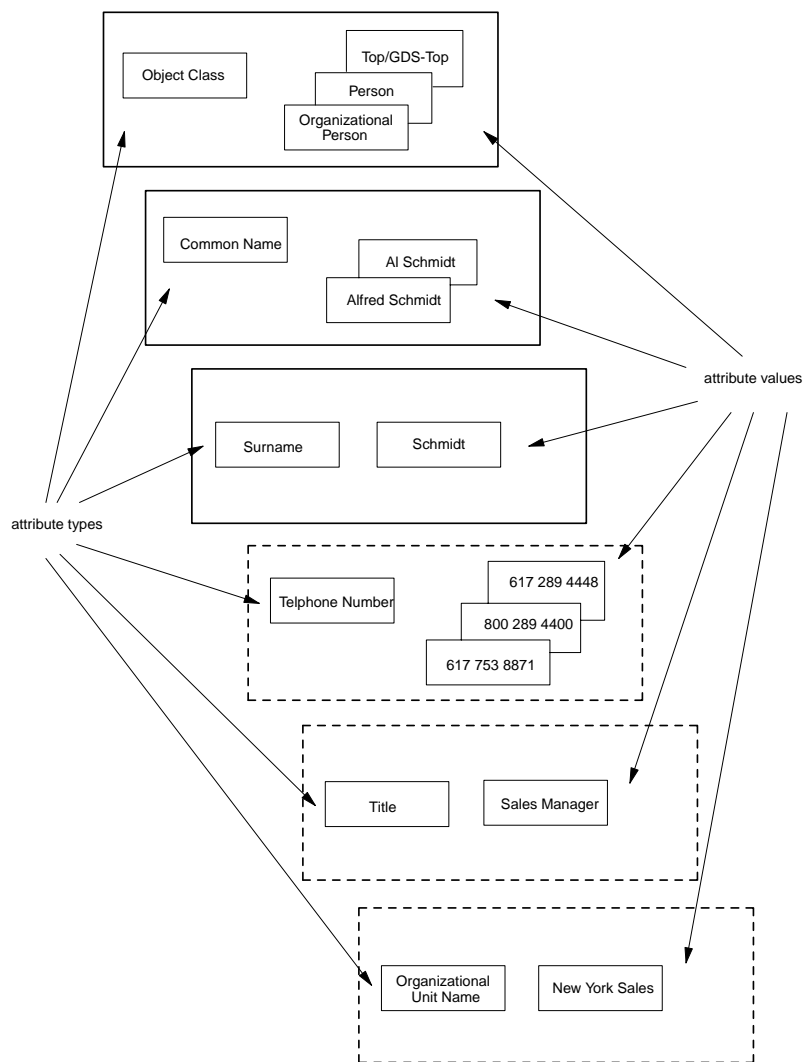
Entries are grouped into generic object classes based on the type of object they represent. Examples of object classes are **Country**, **Organizational-Person**, and **Application-Entity**. All entries contain a special attribute, the object class attribute, indicating to which object class (or classes) they belong.

Entries that model a certain object and contain information about the object in terms of attributes are called *object entries*. The directory contains a second type of entry, which is a pointer to an object entry, called an *alias entry*. Alias entries are discussed in Section 4.3.4.

In summary (as shown previously in Figure 4-2), the DIB is made up of entries, each of which contains information about objects. Entries consist of attributes; each attribute has a type and one or more values.

Section 4.3 describes how objects are organized in the DIB. Figure 4-4 shows an example of an entry describing **Organizational-Person**.

Figure 4-4. A Directory Entry Describing Organizational Person



4.3 X.500 Naming Concepts

Large amounts of information need to be organized in some way to make efficient retrieval possible and ensure that names are unique. Information in the DIB is organized into a hierarchical structure known as the Directory Information Tree (DIT). The structure and naming of the nodes in the DIT are specified by registration authorities for a standardized set of X.500 names and by implementors of the directory service (such as OSF) for implementation-specific names. The DIT hierarchy is described by a schema. Schemas are described in more detail in Section 4.4.

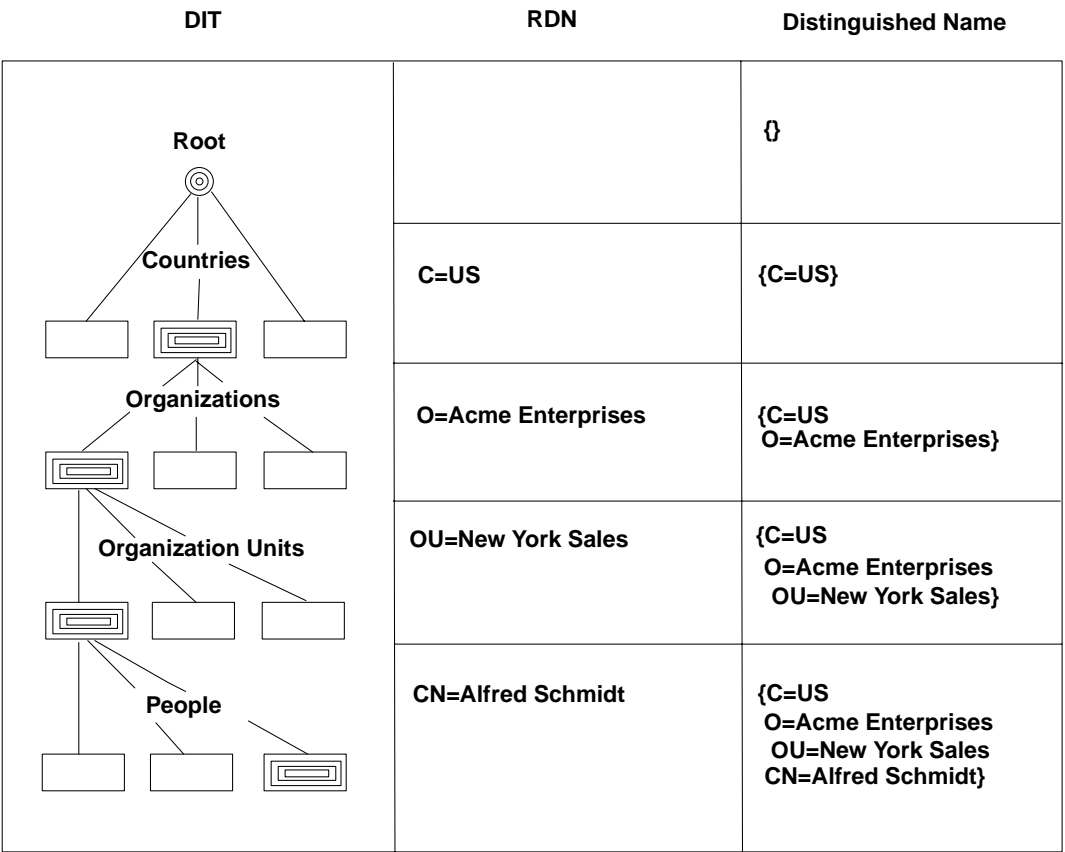
Although the X.500 standard does not mandate a specific schema, it does make general recommendations. For example, countries and organizations should be named close to the root of the DIT; people, applications, and devices should be named further down in the hierarchy. GDS supplies a default schema that complies with these recommendations.

4.3.1 Distinguished Names

A hierarchical path exists from the root of the DIT to any entry in the DIB. To access information stored in an entry, a name that uniquely describes that entry must be given. An RDN distinguishes an entry from other entries with the same superior node in the DIT. A sequence of RDNs, starting from the root of the tree, can identify a unique path down the tree, and thus a unique entry. This sequence of RDNs, each of which identifies a particular entry, is the distinguished name of that entry. Each entry in the DIB can be referred to by its distinguished name.

Figure 4-5 shows an example of a distinguished name. The shaded boxes in the DIT represent the entries that are named in the column labeled RDN. The schema dictates that countries are named directly below the root, followed by organizations, organization units, and people.

Figure 4-5. A Distinguished Name in a Directory Information Tree



Every entry in the DIB has a distinguished name, not just the leaf nodes. For example, the entry for the organization, Acme Enterprises (shown in Figure 4-5) is represented by the shaded box in the **Organizations** subtree. Its distinguished name is the concatenation of the distinguished name of the previous entry above with its relative distinguished name. The entry for **People**, Alfred Schmidt, is represented by the shaded box in the **People** subtree.

4.3.2 Relative Distinguished Names and Attribute Value Assertions

Each entry has a unique *relative distinguished name* (RDN), which distinguishes it from all other entries with a particular immediate superior in the DIT.

An RDN consists of one or more assertions of the type and value of an attribute. A pair consisting of an attribute type and a value of that type is known as an *attribute value assertion* (AVA). All attribute types in an RDN must be different. The attribute value of an attribute in an RDN's AVA is called the *distinguished value* of that attribute, as opposed to the other possible values of that attribute.

The assertion is TRUE if the entry contains an attribute of the specified type, and if one of that attribute's values matches the AVA's distinguished attribute value. An entry commonly has an RDN that consists of a single AVA. In some cases, however, more than one AVA may be required to distinguish an entry. (Multiple AVAs are discussed in Section 4.3.3.)

The entry shown in Figure 4-4 contains the RDN **Common-Name = Alfred Schmidt**. The attribute consists of two values: **Alfred Schmidt** and **Al Schmidt**. The AVA **Common-Name = Alfred Schmidt** contains the value **Alfred Schmidt**, which has been designated as the distinguished value in the AVA.

4.3.3 Multiple AVAs

Frequently, as shown in the previous section, an entry contains a single distinguished value; therefore, the RDN consists of a single AVA. However, under certain circumstances, additional values (and hence multiple AVAs) can be used.

Figure 4-4 shows the contents of an entry describing **Organizational-Person**. The RDN of an **Organizational-Person** entry is usually composed of a single AVA, such as the **Common-Name** attribute type with a distinguished value (in Figure 4-5, the AVA **CN = Alfred Schmidt**). Depending on the schema, the RDN of an **Organizational-Person** entry may contain more than one AVA. For example, the RDN in Figure 4-5 could contain the AVAs **CN = Alfred Schmidt** and **OU = New York Sales**, with Alfred Schmidt and New York Sales as distinguished values.

In summary:

- A DIT consists of a collection of distinguished names.
- Distinguished names result from a concatenation of RDNs.
- RDNs consist of an unordered collection of attribute type and value pairs (AVAs).

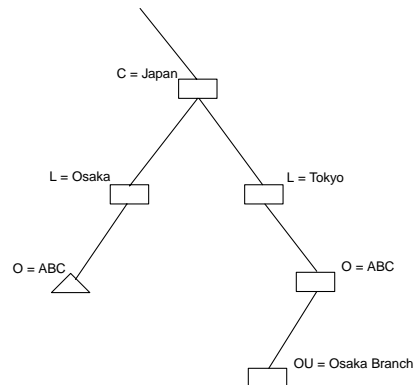
4.3.4 Aliases

An alternative name or alias is supported in the DIT by the use of special pointer entries called *alias entries*. Alias entries do not contain any other attributes beyond their distinguished attributes, the object class attribute, and the aliased object name attribute; that is, the distinguished name of the aliased object entry. Furthermore, an alias entry has no subordinate entries, making it, by definition, a leaf entry of the DIT as shown in Figure 4-6. Alias entries point to object entries and provide the basis for alternative names for the corresponding objects.

Aliases are used to do such things as provide more user-friendly names, direct the search for a particular entry, reduce the scope of a search, provide for common alternate abbreviations and spellings, or provide continuity after a name change.


Figure 4-6 demonstrates how an alias name provides continuity after a name change. The ABC company's branch office located originally in Osaka has moved to Tokyo. To make the transition easier for directory service users and to guarantee that a search based on the old information finds its target, an alias for **O=ABC** has been added to the directory beneath **L=Osaka**. This alias entry points to the object entry **O=ABC**. A search for ABC under **L=Osaka** in the DIT finds the entry **/C=Japan/L=Tokyo/O=ABC**.

Figure 4–6. An Alias in the Directory Information Tree



Another use of alias entries is as an alternative to *filtering*; that is, by using assertions about particular attributes to search through the DIT. Although this approach does not require any special information to be set up in the DIT, it can be expensive to search where there is a large population of entries and attributes. An alternative approach is to set up special subtrees whose naming structures are designed for “Yellow Pages” type searching. Figure 4-7 shows an example of such a subtree populated by alias entries only. In reality, the entries within these subtrees can be a mixture of object and alias entries, as long as there exists only one object entry for each object stored in the directory.

Figure 4–7. A Subtree Populated by Aliases



An object with an entry in the DIT can have zero or more aliases. Several alias entries can point to the same object entry. An alias entry can point to an object that

is not a leaf entry. Only object entries can have aliases. Thus, aliases of aliases are not permitted.

4.3.5 Name Verification

A directory user identifies an entry by supplying an ordered set of RDNs (each of which consists of an unordered set of AVAs) that form a purported name. The purported name is mapped onto the desired entry by the process of name verification, which performs a distributed tree walk through the DIT. When a purported name is a valid name, a distinguished name exists with the same number of RDNs and matching AVAs within the RDNs.

4.4 Schemas

The structure of directory information is governed by a set of rules called a *schema*. Schemas specify rules for the following:

- The structure of the DIT
- The contents of entries in terms of attributes
- The syntax of attribute values and rules for comparing and matching them

4.4.1 The GDS Standard Schema

When the DCE software package is shipped to a customer, it includes a default or *standard* schema for GDS. This is the GDS proprietary interpretation of the X.500 schema.

Each attribute in the schema is assigned a unique object identifier and the syntax of its value. In addition, the schema specifies the mechanism by which attributes of this type are compared with one another. Each entry in the DIT belongs to an object class governed by the schema. Object class definitions can be used to derive subclasses, supporting the inheritance and refinement of the attribute types defined for the superclass.

Included with the GDS standard schema are the following tables that define the structure of the directory:

- Structure rule table (SRT)
- Object class table (OCT)
- Attribute table (AT)

4.4.2 The Structure Rule Table

The SRT specifies the relationship of object classes in the structure of the directory. The SRT supplied with the GDS standard schema contains the entries shown in Table 4-2.

Table 4–2. Structure Rule Table Entries

Rule Number	Superior Rule Number	Acronym of Naming Attribute	Acronym of Structural Object Class
1	0	CN	SCH
2	0	C	C
3	2	O	ORG
4	3	OU	OU
5	4	CN	ORP
6	4	CN, OU	ORP
7	4	CN	ORR
8	4	CN	MDL
9	4	CN	APP
10	9	CN	APE
11	9	CN	DSA
12	9	CN	MMS
13	9	CN	MTA

Rule Number	Superior Rule Number	Acronym of Naming Attribute	Acronym of Structural Object Class
14	9	CN	MUA
15	2	L	LOC
16	15	CN	REP
17	15	CN, STA	REP

The SRT determines how the object classes are laid out in the DIT by assigning rule numbers to each object class. An object class' superior rule number specifies the object class directly above it in the DIT.

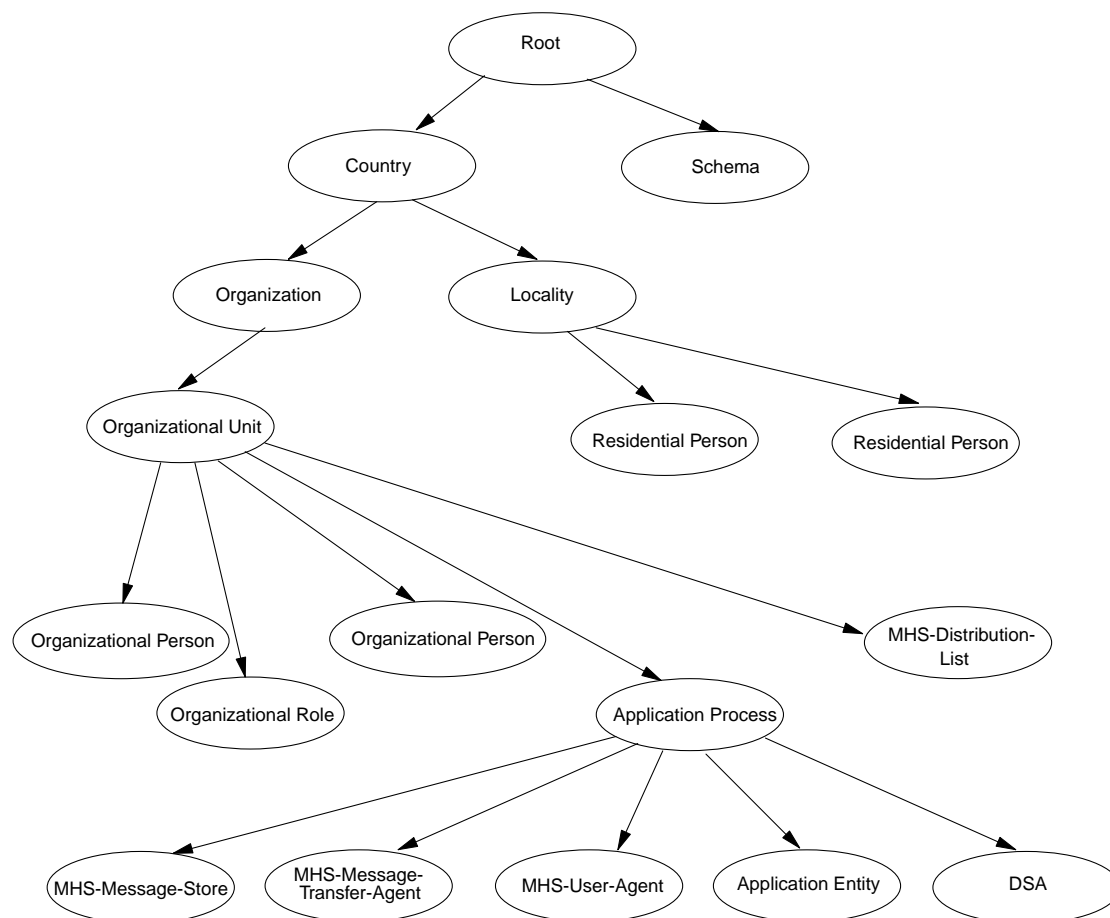
For example, the object class **Organization** (abbreviated with the acronym **ORG** in the SRT) has a superior rule number of 2, indicating that it is located in the DIT beneath the object class **Country (C)**, which has a rule number of 2. **Organization Unit (OU)** is located beneath **Organization** because it has a superior rule number of 3 and so forth.

The SRT only contains structured object classes; that is, classes that form branches in the DIT. Other object classes, such as abstract and alias classes, are not included.

The SRT specifies the attribute(s) used to name entries belonging to each object class. These attributes, called *naming attributes*, are used to define the RDN and therefore the distinguished name of directory entries.

Figure 4-8 shows the structure of the DIT as defined by the SRT of the GDS standard schema.

Figure 4–8. SRT DIT Structure for the GDS Standard Schema



4.4.3 The Object Class Table

The object classes that make up the GDS standard schema are defined in the OCT. Table 4-3 contains a partial listing of the OCT (refer to the *OSF DCE GDS Administration Guide and Reference* for a complete listing of the OCT for the GDS

standard schema). Each column in Table 4-3 contains information about an object class entry in the schema.

Table 4–3. Object Class Table Entries

Object Class							
Acronym	Name	Kind	Super-class	OID	File No.	Mandatory Attributes	Optional Attributes
TOP	Top	Abstract	None	85.6.0	-1	OCL	None
ALI	Alias	Alias	TOP	85.6.1	-1	AON	None
C	Country	Structural	GTP	85.6.2	1	C	DSC SG CDC CDR
LOC	Locality	Structural	GTP	85.6.3	4	None	DSC L SPN STA SEA SG CDC CDR
ORG	Organization	Structural	GTP	85.6.4	1	O	DSC L SPN STA PDO PA PC POB FTN IIN TN TTI TXN X1A PDM DI RA SEA UP BC SG CDC CDR

Note: The object identifiers in Table 4-3 stem from the root {**joint-iso-ccitt(2) ds(5) objectClass(6)**}.


Column 4, Superclass acronyms, provides the class from which an object class inherits its attributes. Using the information in Column 4, it is possible to derive a graphical representation of the inheritance properties of object classes in the DIT as shown in Figure 4-9.

In the figure, the object class **Top** is the root of the tree, with **Alias** and **GDS-Top** as the main branches. **Top** contains the attribute type object class, which is inherited by all the other object classes.

Do not confuse the information in the OCT with that presented in the SRT. There is no direct relationship between the relative location of branches and leaves in the DIT structure and the inheritance properties of classes with their superclasses and subclasses. For example, when a directory service request is made by a directory user, such as a read operation, the SRT is used by the directory service to indicate its position in the DIT. The directory service uses the information defined in the SRT for tree traversal so that the requested object can be located in the directory. Figure 4-8 shows the object class **Organization** located beneath **Country** in the DIT.

On the other hand, the OCT defines, among other things, the attributes of an object class along with its inherited attributes from its superclass. The superclass, in turn, inherits the attributes from its superclass, and so on until the root, **Top**, is reached (from which all classes derive their attributes). Figure 4-9 shows the object class **Organization** as a subclass of **GDS-Top**. As such, it inherits its attributes from **GDS-Top**, which in turn inherits from its superclass, **Top**.

Figure 4–9. A Partial Representation of the Object Class Table



The OCT also contains the unique object identifier of each class in the DIT. These numbers are defined by various standards authorities and in the X.500 standards documents mentioned previously. The AT also contains the predefined object identifiers for each attribute in the directory. These object identifiers are defined in the header files that are included as part of the GDS API. Table 4-4 shows some examples of object identifiers for directory classes as defined in the X.500 standard.

Table 4–4. Object Identifiers for Selected Directory Classes

Object Class Type	Object Identifier
Alias	85.6.1
Application-Entity	85.6.12
Application-Process	85.6.11
Country	85.6.2
Device	85.6.14
DSA	85.6.13
Group-of-Names	85.6.9
Locality	85.6.3
Organization	85.6.4
Organizational-Person	85.6.7
Organizational-Role	85.6.8
Organizational-Unit	85.6.5
Person	85.6.6
Residential-Person	85.6.10
Top	85.6.0

Note: The object identifiers in Table 4-4 stem from the root **{joint-iso-ccitt(2) ds(5) objectClass(6)}**.

Another important feature of the OCT is the distinction made between mandatory and optional attributes for each object class. This distinction is based on recommendations from X.500 standards documents. These documents (Recommendations X.520 and X.521) define selected object classes and associated attribute types by using ASN.1 notation. Most object classes have one or more mandatory attributes associated with them for use by implementors who want to comply with the X.500 standards recommendations. In addition, optional attributes are defined.

The following example provides a flavor of ASN.1 notation; it shows how the object class **country** is described in Recommendation X.520 (*The Directory: Selected Object Classes*).

```
country OBJECT-CLASS
  SUBCLASS of top
  MUST CONTAIN {
    countryName}
  MAY CONTAIN {
    description,
    searchGuide}
  ::= {objectClass 2}
```

This ASN.1 definition defines **Country** as a subclass of superclass **Top**. The class, **Country**, must contain the mandatory attribute **countryName** (or **Country-Name** as defined in the GDS standard schema) and can contain the optional attributes **Description** and **Search-Guide**. In addition, the DCE implementation adds two more attributes, **CDS-Cell** and **CDS-Replica**, to incorporate other aspects of the DCE environment that are implementation specific.

Country is assigned the object identifier **2.5.6.2**. This number distinguishes it from the other object classes defined by the standard. The **Top** superclass is designated as **2.5.6.0**. The first three numbers, **2.5.6**, identify the object class as a member of a discrete set of object classes defined by X.500. The last number in the object identifier distinguishes objects within that discrete set. Alias, a subclass of **Top**, is assigned the number **2.5.6.1**. **Country** is assigned the number **2.5.6.2**, and so on. **GDS-Top** has no object identifier because it is implementation specific and thus not identified by the standard.

4.4.4 The Attribute Table

The attributes that make up the entries in the GDS standard schema are defined in the AT. (Refer to the *OSF DCE GDS Administration Guide and Reference* for a complete listing of the AT.) The object identifiers are in the range from **85.4.0** through **85.4.35** as defined by the X.500 standard, **86.5.2.0** through **86.5.2.10** as defined by the X.400 standard, and there are additional object identifiers for GDS-specific attributes.

Table 4-5 shows a partial listing of the AT for the GDS standard schema.

Note: The access class for every attribute listed in Table 4-5 is 0 (zero).

Table 4–5. Attribute Table Entries

Acr. of Attr.	Obj. ID	Name of Attribute	Lower Bound	Upper Bound	Max. No. of Val.	Syntax	Phon. Flag	Index Level
OCL	85.4.0	Object-Class	1	28	0	2	0	0
AON	85.4.1	Aliased- Object-Name	1	1024	1	1	0	0
KNI	85.4.2	Knowledge- Information	1	1024	0	4	0	0
CN	85.4.3	Common-Name	1	64	2	4	1	1
SN	85.4.4	Surname	1	64	2	4	1	0
SER	85.4.5	Serial-Number	1	64	2	5	0	0
C	85.4.6	Country-Name	2	2	1	1010	1	1
L	85.4.7	Locality-Name	1	128	2	4	1	1
SPN	85.4.8	State-or- Province-Name	1	128	2	4	1	0

The columns with the headings Lower Bound and Upper Bound specify the range of the number of bytes (or octets) that the value of an attribute can contain. The schema puts constraints on the number of values that an attribute can contain in the Maximum Number of Values column.

The Syntax column describes how the data is represented and relates to ASN.1 syntax definitions for attributes. For example, a sample of ASN.1 notation for the **Common-Name** attribute follows:

```
commonName ATTRIBUTE
  WITH ATTRIBUTE-SYNTAX
    caseignoreStringSyntax
      (SIZE(1..ub-common-name))
  ::= (attributeType 3)
```

The **Common-Name** attribute is defined as case insensitive. The size of the string is from 1 to the upper bound defined by the schema for the **Common-Name** attribute in the Upper Bound column (in this case, 64 bytes or octets).

Note also that the **Common-Name** attribute is assigned the number 3 by the standard. This corresponds to the **3** in the object identifier **85.4.3**.

The other columns in the AT refer to the phonetic matching flag, security access classes, and index level.

As mentioned previously for object classes, object identifier values specified in the AT are defined as constants in the GDS header files.

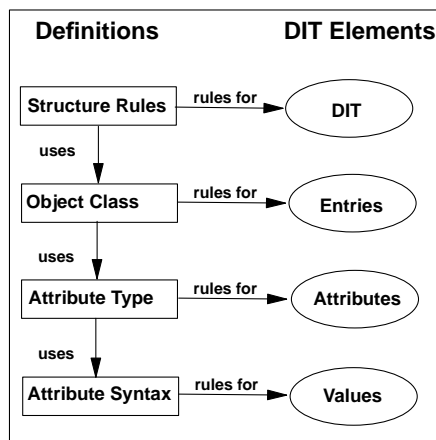
4.4.5 Defining Subclasses

The ability to define subclasses is a powerful feature of the directory. Structure rules govern which object classes can be children of which others in the DIT and therefore determine possible name forms.

The directory standard defines a number of standard attribute types and object classes. For example, the attribute types **Common-Name** and **Description**, and the object classes **Country** and **Organizational-Person** are defined. Implementations of the directory standard, such as DCE, define their own schemas following rules stated in the standard with additional attribute types and object classes.

Figure 4-10 shows the relationship between schemas and the directory information model.

Figure 4–10. The Relationship Between Schemas and the DIT



4.5 Abstract Syntax Notation 1

The need for Abstract Syntax Notation 1 (ASN.1) arises because different computer systems represent information in different ways. For example, one computer can use EBCDIC character representation while another can use ASCII. To transfer a file of characters from one system to another, common representation must be used during the transfer. This transfer can be one representation or the other, or some mutually agreed upon representation negotiated by the two systems. Similarly, floating-point values, integers, and other types of data can be stored internally in different ways. To exchange information, a common format must be agreed to before information can be exchanged.

The translation of EBCDIC to ASCII characters can seem like a trivial problem, but that leaves the larger issue of mapping between the many diverse representations that can exist within a network environment. To address this need, the ISO standards committee defined ASN.1 and Basic Encoding Rules.

ASN.1 is based on the idea that the aspects of transferred information that are preserved are type, length, and value. Data types are collections of values distinguished for some reason, such as characters, integers, and floating-point values. Records and

structure types become more complex when they combine several types into a single structure.

ASN.1 provides a way to group types into abstract syntaxes. An abstract syntax is a named group of types. The standard defines abstract syntax as the notation rules that are independent of the encoding technique used to represent them. Abstract syntax does not specify how to represent values of types, but merely defines the types that make up the group of types.

Abstract syntaxes are not enough to define how values of the data types in a specific abstract syntax are to be represented during communications. For this reason, ISO further defines a *transfer syntax* for each abstract syntax. A transfer syntax is a set of rules for encoding values of some specified group of types.

4.5.1 ASN.1 Types

ASN.1 is similar to a high-level programming language. Unlike other high-level languages, ASN.1 has no executable statements. It includes only language constructs required to define types and values.

ASN.1 defines a number of built-in types. Users of ASN.1 can then define their own types based on the built-in types provided by the language. The ASN.1 standard defines four categories of types that are commonly used in defining application interfaces such as XOM and XDS:

- ASN.1 simple types
- ASN.1 useful types
- ASN.1 character string types
- ASN.1 type constructors

ASN.1 simple types are Bit String, Boolean, Integer, Null, Object Identifier, Octet String, and Real. Table 4-6 shows the relationship of OM syntaxes (syntaxes defined in XOM API) to ASN.1 simple types. (Refer to Chapter 17 for the complete set of tables for the four categories of ASN.1 types.) As shown in the table, for every ASN.1 type except Real, there is an OM syntax that is functionally equivalent to it. The simple types are listed in the first column of the table; the corresponding syntaxes are listed in the second column.

Table 4–6. Syntax for the Simple ASN.1 Types

ASN.1 Type	OM Syntax
Bit String	String(OM_S_BIT_STRING)
Boolean	OM_S_BOOLEAN
Integer	OM_S_INTEGER
Null	OM_S_NULL
Object Identifier	String(OM_S_OBJECT_IDENTIFIER_STRING)
Octet String	String(OM_S_OCTET_STRING)
Real	None ¹

1 A future edition of XOM can define a syntax corresponding to this type.

An example will illustrate how OM syntaxes are used to define the syntax of values for various attributes. One of the simplest of the ASN.1 types is Boolean. There are only two possible values for a Boolean type: TRUE and FALSE. The **DS_FROM_ENTRY** OM attribute of the **DS_C_ENTRY_INFO** object class has a value syntax of **OM_S_BOOLEAN**. **OM_S_BOOLEAN** is the C language representation for the OM syntax that corresponds to the ASN.1 Boolean type. The value of the **DS_FROM_ENTRY** OM attribute indicates whether information from the directory was extracted from the specified object's entry (TRUE), or from a copy of the entry (FALSE). The actual C language definition for **OM_S_BOOLEAN** is made in the XOM API header file **xom.h**.

4.5.2 Basic Encoding Rules

It is possible to define a single transfer syntax that is powerful enough to encode values drawn from a number of abstract syntaxes. ISO defines a set of rules for encoding values of many different types for ASN.1. This set of encoding rules is called *basic encoding rules* (BER). It is so powerful that values from any abstract syntax described by using ASN.1 can be encoded by using the transfer syntax defined by BER.

Although other transfer syntaxes could be used for representing values from ASN.1, BER is used most often.

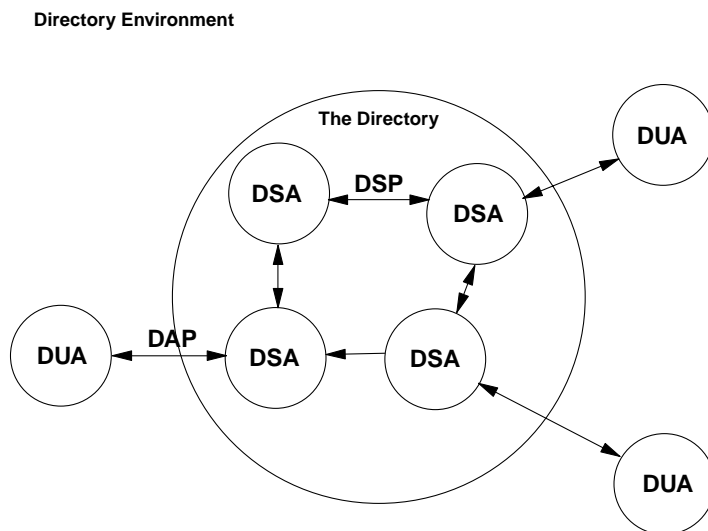
4.6 GDS as a Distributed Service

When present in a DCE cell, GDS can serve two basic functions. First, it can provide a high-level, worldwide directory service by tying together independent DCE cells. Second, it can be used as an additional directory service to CDS for storing object names and attributes in a central place.

The GDS database contains information that can be distributed over several GDS servers. In addition, copies of information can be stored in multiple GDS servers, and the information can also be cached locally. The unit of replication in GDS is the directory entry; whole subtrees can be also replicated.

The information belonging to the DIB is shared between several Directory Service Agents (DSAs). A DSA is a process that runs on a GDS server machine and manages the GDS database. DSAs cooperate to perform directory service operations, with each DSA knowing a fraction of the total directory information, as shown in Figure 4-11. DSAs are a combination of local database functions and a remote interface to the clients of users and other DSAs. DSAs can cooperate to execute operations. This cooperation often involves the navigation of operations through the network.

Figure 4–11. The Relationship Between the DSA and the DUA



Users access the directory via Directory User Agents (DUAs). DUAs make requests of DSAs on behalf of users requesting directory service operations. The manner in which DUAs communicate with DSAs is defined by the X.500 standard. For communications between DUAs and DSAs, the directory access protocol (DAP) is defined. For communications between DSAs in a distributed directory, the standard defines the directory system protocol (DSP).

4.6.1 The Directory Access Protocol

The directory standard defines directory functions in the DAP. The directory functions can be divided into three general categories: read, search, and modify.

Read operations involve the retrieval of information from specific named entries. This allows a general name-to-attributes mapping analogous to the White Pages phone directory.

Search operations involve the general browsing and relational searching of information. Search operations support human interaction with the directory service and is analogous to that of the “Yellow Pages” telephone directory.

Modify operations involve the modification of information in the directory.

4.6.2 The Directory System Protocol

The DSA can interact with other DSAs to provide services by using the DSP. DSP is a protocol defined by the directory standard to allow DSAs to communicate with one another. DSP provides two methods of distributed request resolution: referral and chaining.

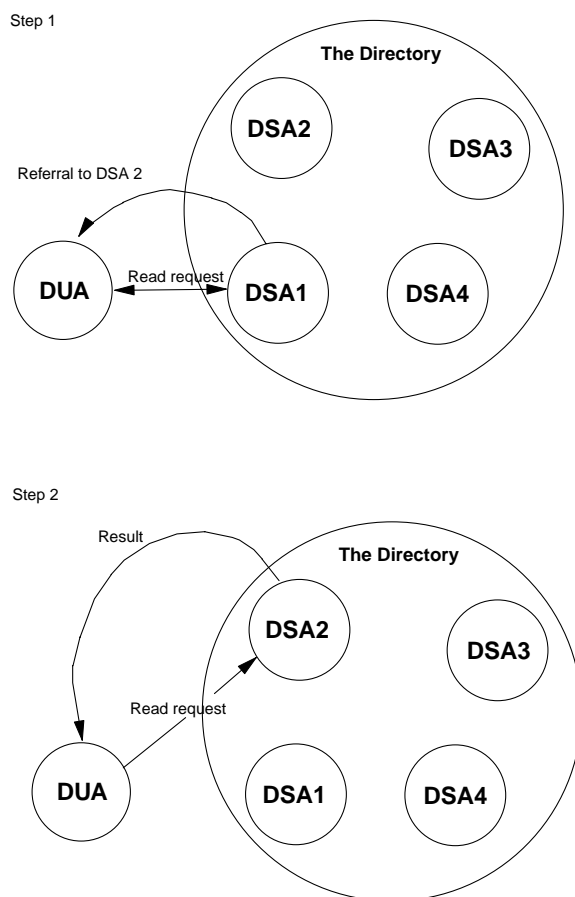
4.6.3 Referral

In some cases, a DSA may not be able to provide service to a DUA because the required information is held elsewhere in the network. A DSA can simply choose to inform the DUA or the calling DSA where the information can be found. This is called *referral* and can occur because of the user’s preference or the DSA’s circumstances.

Referrals are possible because the distinguished name provided by the DUA identifies where in the DIT the requested entry is located. DSAs use their knowledge of the DIT to inform the DUA of the DSA that holds the requested information.

Figure 4-12 shows an example of a referral. **DSA1** passes a referral to **DSA2** back to the DUA. The DUA then makes a request to **DSA2**.

Figure 4–12. An Example of a Referral



4.6.4 Chaining

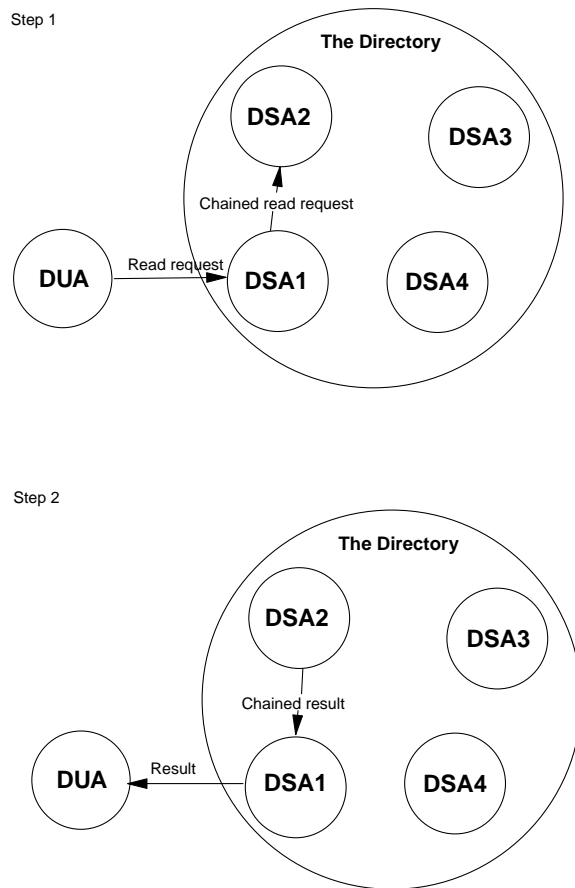
If a request received from a DUA cannot be fulfilled by the receiving DSA, that DSA can send a referral back to the initiating DUA over DAP. Alternatively, the DSA can chain the request over DSP, asking another DSA to perform the requested function. That DSA can perform the function or can send back a referral of its own. In either

case, the first DSA eventually responds to the originating DUA with either the results of the completed operation or a referral.

Chaining can go deeper than one level. To prevent lengthy searches, a user can request no chaining or specify a limit on the total elapsed time for an operation.

Figure 4-13 shows an example of chaining. The DUA makes a request of **DSA1**. **DSA1** is unable to service the request and passes it to **DSA2**. **DSA2** services the request, passes the result back to **DSA1**, and **DSA1** passes the result back to the DUA.

Figure 4-13. An Example of Chaining



4.6.5 The Directory User Agent Cache

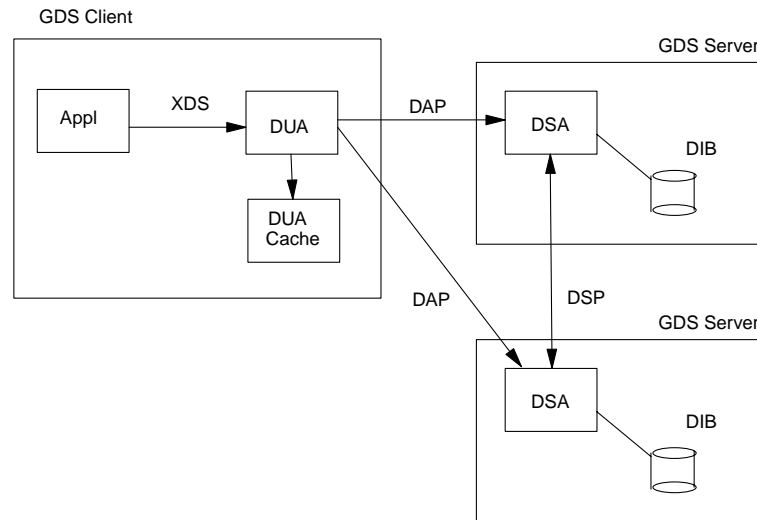
The DUA cache is a process that keeps a cache of information obtained from DSAs. One DUA cache runs on each client machine and is used by all users on that client. The DUA cache contains copies of recently accessed object entries and information about DSAs. The user specifies which information should be cached. It is also possible to bypass the DUA cache to obtain information directly from a DSA. This is desirable, for example, when the user wants to make sure the information obtained is up-to-date.

The shadow update and cache update are processes that update replicated information in DSAs and DUA caches. These processes run as needed and then terminate. The shadow update process runs on the GDS server machine; the cache update process runs on GDS client machines.

When an application program makes a directory service call by using XDS API, the call is handed to the DUA library. The DUA first looks in the DUA cache (if requested by the user) to see if the requested information is already available on the local machine. If it is not, the DUA queries a DSA. If the DSA has the requested information, it returns the results to the DUA. If it does not, the query can proceed either by using chaining or a referral. In either case, different DSAs are queried until the information is found. It is cached (if requested by the user) in the DUA cache and the results are returned to the application program.

Figure 4-14 shows the interaction between an application program, via the XDS interface, and the GDS client and server. The GDS client and server use DAP to communicate. GDS servers use DSP to communicate with one another. DAP and DSP perform functions similar to the functions that DCE RPC protocols perform in other DCE services.

Figure 4–14. GDS Components



4.6.5.1 Placing Entries in the Local DUA Cache

A special object OM class, **DSX_C_GDS_CONTEXT**, is provided in the GDS package to allow an application program to manage the placement of entries in the local DUA cache as a result of a directory request.

DSX_C_GDS_CONTEXT inherits the OM attributes of its superclasses **OM_C_OBJECT** and **DS_C_CONTEXT**. To enable caching entries, the **DS_DONT_USE_COPY** OM attribute of **DS_C_CONTEXT** must be set to a value of **OM_FALSE**, indicating that a directory request can access copies of directory entries maintained in other DSAs or copies cached locally.

DSX_C_GDS_CONTEXT has the following private extension OM attributes in addition to the OM attributes inherited from **DS_C_CONTEXT**:

- **DSX_DUAFIRST**
- **DSX_DONT_STORE**
- **DSX_NORMAL_CLASS**

- **DSX_PRIV_CLASS**
- **DSX_RESIDENT_CLASS**
- **DSX_USED_SA**
- **DSX_DUA_CACHE**

DSX_DUA_FIRST determines where a query operation, such as a search or list, looks first for an entry. The default value is **OM_FALSE**, indicating that the DSA is searched first. If the entry is not found, then the DUA cache is searched.

DSX_DONT_STORE determines if information read from the DSAs by a query function also needs to be stored in the DUA cache. If this OM attribute is set to **OM_TRUE**, nothing is stored in the cache. If this OM attribute is set to **OM_FALSE**, the information read is stored in the DUA cache. The objects returned by **ds_list()** and **ds_compare()** are stored in the cache without their associated attribute information. The objects returned by **ds_read()** and **ds_search()** are stored in the cache with all their cachable attributes; these are all public attributes that do not exceed 4 Kilobytes in length.

The three different memory classes that the user can specify for a cached entry are **DSX_NORMAL_CLASS**, **DSX_PRIV_CLASS**, and **DSX_RESIDENT_CLASS**.

DSX_NORMAL_CLASS assigns the entry to the class of normal objects. If the number of entries in this class exceeds a maximum value, the entry that is not accessed for the longest period of time is removed from the DUA cache.

DSX_PRIV_CLASS assigns the entry to the class of privileged objects. Entries can be removed from the class in the same way as normal objects. However, by setting this area of memory aside to be used sparingly, the user can protect entries from deletion.

DSX_RESIDENT_CLASS assigns the entry to the class of resident objects. An entry in this class is never removed automatically. It must be explicitly removed by using an XDS **ds_remove_entry()** function applied directly to the cache; that is, **DSX_DUA_CACHE** and **DSX_USED_SA** are set to **OM_TRUE** and **OM_FALSE**, respectively.

Tables 4-7 through 4-9 show the possible conditions that result when **DSX_DUA_CACHE** and **DSX_USED_SA** are set to **OM_TRUE**.

Table 4–7. Cache Attributes: Read Cache First

OM Attribute Type	OM_TRUE	OM_FALSE
DSX_DUA_CACHE	X	
DSX_USEDSA	X	
DS_DONT_USE_COPY		X
DSX_DUAFIRST	X	

In the situation presented in Table 4-7, the cache is read first, then the other DSAs. The requested operation is permitted to use copies of entries.

Table 4–8. Cache Attributes: Read DSA First

OM Attribute Type	OM_TRUE	OM_FALSE
DSX_DUA_CACHE	X	
DSX_USEDSA	X	
DS_DONT_USE_COPY		X
DSX_DUAFIRST		X

In the situation presented in Table 4-8, the DSA is read first, then the cache. The requested operation is permitted to use copies of entries.

Table 4–9. Cache Attributes: Read DSA Only

OM Attribute Type	OM_TRUE	OM_FALSE
DSX_DUA_CACHE	X	
DSX_USEDSA	X	
DS_DONT_USE_COPY	X	
DSX_DUAFIRST	N/A	N/A

In the situation presented in Table 4-9, only the DSA is read. The requested operation is not permitted to use copies of entries.

Tables 4-10 through 4-12 show the possible situations when **DSX_DUA_CACHE** and **DSX_USEDSA** are not both set to **OM_TRUE**.

Table 4-10. Cache Attributes: DSX_USEDSA is OM_FALSE

OM Attribute Type	OM_TRUE	OM_FALSE
DSX_DUA_CACHE	X	
DSX_USEDSA		X
DS_DONT_USE_COPY		X

In the situation presented in Table 4-10, the DUA Cache is used exclusively.

Table 4-11. Cache Attributes: DSX_DUA_CACHE is OM_FALSE

OM Attribute Type	OM_TRUE	OM_FALSE
DSX_DUA_CACHE		X
DSX_USEDSA	X	

In the situation presented in Table 4-11, the DSA is used exclusively.

Table 4-12. Cache Attributes: Error

OM Attribute Type	OM_TRUE	OM_FALSE
DSX_DUA_CACHE		X
DSX_USEDSA		X

In the situation presented in Table 4-12, neither the DSA nor the DUA cache is used, and an error is returned.

4.6.5.2 Accessing the DUA Cache Without a GDS Server Present

An application program may need to access the local DUA cache without binding to a GDS server. This section describes the steps that should be included in the

application program. Refer to Chapters 5 and 6 for information on how to use the XDS and XOM API calls **ds_initialize()**, **ds_version()**, **ds_bind()**, **ds_shutdown()**, **om_create()**, **om_remove()**, how to do static initialization of public objects, and how to create private objects.

The steps are as follows:

1. Call **ds_initialize()** as normal.
2. Negotiate the **DSX_GDS_PKG** by using **ds_version()**. This is necessary because **DSX_C_GDS_CONTEXT** is required in order to set the DUA cache service controls.
3. Supply a **DSX_C_GDS_SESSION** object to the **ds_bind()** call, which has no **DS_DSA_NAME** attribute and no **DS_DSA_ADDRESS** attribute present.

There are two ways of achieving this step:

- Supply a public **DSX_C_GDS_SESSION** object (static initialization):

```
OM_descriptor cache_session[] = {  
    OM_OID_DESC(OM_CLASS, DSX_C_GDS_SESSION),  
    OM_NULL_DESCRIPTOR  
};
```

The other attributes of **DSX_C_GDS_SESSION** (**DS_REQUESTOR**, **DSX_PASSWORD**, **DSX_DIR_ID**, **DSX_AUTH_MECHANISM** and **DSX_AUTH_INFO**) can be included if required.

- Supply a private **DSX_C_GDS_SESSION** object (using XOM API function calls):

```
om_create(DSX_C_GDS_SESSION, OM_TRUE, workspace, &cache_session);  
om_remove(cache_session, DS_DSA_NAME, 0, OM_ALL_VALUES);  
om_remove(cache_session, DS_DSA_ADDRESS, 0, OM_ALL_VALUES);
```

Note that an uninitialized session object may not be passed to **ds_bind()**. That means that **OM_FALSE** should not be used with the previous **om_create()** function call.

4. Supply a **DSX_C_GDS_CONTEXT** object to the XDS calls that access the DUA cache. The following service controls must be set to ensure that access is restricted to the DUA cache alone:

```
DSX_DUA_CACHE = OM_TRUE  
DSX_USED_SA   = OM_FALSE
```

(Service controls that access the DSA will result in a **DS_E_BAD_CONTEXT** error.)

There are two ways of achieving this step:

- Supply a public **DSX_C_GDS_CONTEXT** object (static initialization)
 - Supply a private **DSX_C_GDS_CONTEXT** object (using XOM API function calls)
5. Supply the bound **DSX_C_GDS_SESSION** object to the **ds_unbind()** call.
 6. Call **ds_shutdown()** as normal.

4.6.6 GDS Configurations

A GDS machine can be configured in two ways:

- Client Only

A node can contain only the client side of GDS. This node can access remote DSAs and cache information in the DUA cache.

- Client/Server

A machine can be configured with both the GDS client and server. This is the typical configuration for a machine acting as a GDS server. This configuration can be useful even if a node acts mainly as a client because the DSA can be used as a larger, more permanent cache of information contained in remote DSAs.

Note: When a client and server reside on the same machine, access to the directory is optimized. Communications between the DUA and the DSA are by means of interprocess communications (IPC) via shared memory.

4.6.7 GDS Security

A number of authentication mechanisms are supported by GDS. XDS applications must indicate which method is to be used. Since authentication takes place at bind time, it is appropriate to pass the selected authentication mechanism as an argument to **ds_bind()**.

A bind operation can be performed by the application program with or without user credentials. A bind with credentials is referred to as an *authenticated bind* and allows an application program to require a user to specify a distinguished name password as user credentials. A bind without user credentials only permits access to public information in the directory.

A special OM object class, **DSX_C_GDS_SESSION**, is provided in the GDS package to accommodate user credentials and authentication mechanisms. In addition to the OM attributes inherited from its superclass **DS_C_SESSION**, this OM class consists of the following OM attributes:

- **DSX_PASSWORD**

This attribute contains the password for the user credentials.

- **DSX_DIR_ID**

This attribute contains the identifier for distinguishing between several configurations of the directory service within a GDS installation. **DSX_DIR_ID** plays no role in user credentials.

- **DSX_AUTH_MECHANISM**

If this attribute is present, it identifies the selected authentication mechanism. If this attribute is absent, then a bind without credentials (that is, anonymous bind) is attempted.

- **DSX_AUTH_INFO**

This attribute is for future use.

The GDS package also provides the following special OM classes to support access rights to specific OM attributes by directory service users:

- **DSX_C_GDS_ACL**

This attribute describes up to five categories of rights for one or more directory users.

- **DSX_C_GDS_ACL_ITEM**

This attribute specifies the user, or subtree of users, to whom an access right applies.

The five categories of rights correspond to the access rights defined for the directory service as described in the *OSF DCE GDS Administration Guide and Reference*. The categories are as follows:

- Modify Public
- Read Standard
- Modify Standard
- Read Sensitive
- Modify Sensitive

Refer to Chapter 6 for more information on binding with credentials and setting access rights for users. The sample programs in Chapter 7 provide examples of how security features are used in application programs.

4.6.8 GDS API Logging

The GDS API logging facility displays informational and error messages for XDS functions. In addition, the input and output arguments to XDS function calls can also be displayed. For each XDS object, its OM types, syntaxes, and values are displayed recursively. A number of different display formats can be selected for the XDS objects. These are selected by setting the value of the environment variable **XDS_LOG** as shown in Table 4-13.

Logging can be activated dynamically at runtime by setting the environment variable **XDS_LOG**.

Table 4–13. XDS_LOG Values

XDS_LOG Value	Result	Example
Bit 1 = on	Display arguments, messages, results and errors	N/A
Bit 1 = off	Display messages only (all other bits ignored)	N/A
Bit 2 = on	Display result and error objects as private objects	N/A
Bit 2 = off	Display result and error objects as public objects	N/A
Bit 3 = on	Object identifiers displayed as specified in 4th bit	N/A
Bit 3 = off	Object identifiers displayed as symbolic constants	DS_C_SESSION
Bit 4 = on	Object identifiers displayed as dotted-decimal	2.5.4.35
Bit 4 = off	Object identifiers displayed as hexadecimal bytes	\x55\x04\x23
Bit 5 = on	Syntaxes displayed as integers	127
Bit 5 = off	Syntaxes displayed as symbolic constants	OM_S_OBJECT
Bit 6 = on	Types displayed as integers	715
Bit 6 = off	Types displayed as symbolic constants	DS_AVAS

The bits shown in Table 4-13 can be combined. For example, the following command sequence sets **XDS_LOG** to **5** (00101 in binary):

XDS_LOG=5; export XDS_LOG

In this example, the logging facility is directed to display arguments, messages, results, and errors, to convert results and errors into public objects (for display purposes only), and to display object identifiers as hexadecimal bytes; and to display OM syntaxes and OM types as symbolic constants. Normally, **XDS_LOG** should be set to 0. If full tracing is required, then set **XDS_LOG** to 1.

4.6.8.1 Logging Format

The following general display format is used by the logging facility:

```
identifier-name = {
  { type, syntax, value },
  { type, syntax, value },
  .
  .
  .
}; /* identifier-name */
```

where:

- | | |
|---------------|--|
| <i>type</i> | Is the integer or symbolic constant for the specified type. |
| <i>syntax</i> | Is the integer or symbolic constant for the specified syntax. A +L is appended to the syntax label if the OM_S_LOCAL_STRING bit is set in the OM_syntax field. |
| <i>value</i> | Is one of the following: <ul style="list-style-type: none"> • An integer (if <i>syntax</i> is OM_S_INTEGER or OM_S_ENUMERATION). • OM_FALSE or OM_TRUE (if <i>syntax</i> is OM_S_BOOLEAN). • Symbolic constant, dotted-decimal notation, or hexadecimal bytes (if <i>syntax</i> is OM_S_OBJECT_ID_STRING). • Quoted-string (if <i>syntax</i> is any other type of string). • Another object (if <i>syntax</i> is OM_S_OBJECT). |

Note: The terminating NULL descriptor is expected but not displayed.

4.6.8.2 Examples

The following examples show how a selection of XDS objects are displayed by the logging facility.

The following filter selects entries that do not have the value **secret** for the **DS_A_USER_PASSWORD** attribute. The **DS_FILTER_TYPE** has the value **DS_NOT**. It contains a single **DS_C_FILTER_ITEM** attribute. **DS_C_FILTER_ITEM** tests for equality against the **DS_A_USER_PASSWORD** attribute.

```
my_filter = {
  { OM_CLASS, OM_S_OBJECT_ID_STRING, DS_C_FILTER },
  { DS_FILTER_ITEMS, OM_S_OBJECT,
    {
      { OM_CLASS, OM_S_OBJECT_ID_STRING, DS_C_FILTER_ITEM },
      { DS_FILTER_ITEM_TYPE, OM_S_ENUMERATION, 0 },
      { DS_ATTRIBUTE_TYPE, OM_S_OBJECT_ID_STRING, DS_A_USER_PASSWORD },
      { DS_ATTRIBUTE_VALUES, OM_S_OCTET_STRING, "secret" },
    },
  },
  { DS_FILTER_TYPE, OM_S_ENUMERATION, 3 },
}; /* my_filter */
```

The following example shows logging output if the interface logger encounters a NULL pointer. The NULL pointer is flagged as follows:

```
my_session = {
  { OM_CLASS, OM_S_OBJECT_ID_STRING, DS_C_SESSION },
  { DS_DSA_NAME, OM_S_OBJECT, ---WARNING: NULL pointer encountered--- },
}; /* my_session */
```

The following example shows logging output if the interface logger encounters a private object. The private object is displayed as follows:

```
bound_session = {
  { OM_PRIVATE_OBJECT, OM_S_OBJECT_ID_STRING, DS_C_SESSION } ...
}; /* bound_session */
```


The following example shows how a 5-part DSA distinguished name is displayed (/C=de/O=sni/OU=ap/CN=dsa/CN=dsa-m1):

```
dsa_name = {
  { DS_DSA_NAME, OM_S_OBJECT,
    {
      { OM_CLASS, OM_S_OBJECT_ID_STRING, DS_C_DS_DN },
      { DS_RDNS, OM_S_OBJECT,
        {
          { OM_CLASS, OM_S_OBJECT_ID_STRING, DS_C_DS_RDN },
          { DS_AVAS, OM_S_OBJECT,
            {
              { OM_CLASS, OM_S_OBJECT_ID_STRING, DS_C_AVA },
              { DS_ATTRIBUTE_TYPE, OM_S_OBJECT_ID_STRING, DS_A_COUNTRY_NAME },
              { DS_ATTRIBUTE_VALUES, OM_S_PRINTABLE_STRING, "de" },
            },
          },
        },
      },
    },
  { DS_RDNS, OM_S_OBJECT,
    {
      { OM_CLASS, OM_S_OBJECT_ID_STRING, DS_C_DS_RDN },
      { DS_AVAS, OM_S_OBJECT,
        {
          { OM_CLASS, OM_S_OBJECT_ID_STRING, DS_C_AVA },
          { DS_ATTRIBUTE_TYPE, OM_S_OBJECT_ID_STRING, DS_A_ORG_NAME },
          { DS_ATTRIBUTE_VALUES, OM_S_TELETEX_STRING, "sni" },
        },
      },
    },
  },
  { DS_RDNS, OM_S_OBJECT,
    {
      { OM_CLASS, OM_S_OBJECT_ID_STRING, DS_C_DS_RDN },
      { DS_AVAS, OM_S_OBJECT,
        {
          { OM_CLASS, OM_S_OBJECT_ID_STRING, DS_C_AVA },
          { DS_ATTRIBUTE_TYPE, OM_S_OBJECT_ID_STRING, DS_A_ORG_UNIT_NAME },
          { DS_ATTRIBUTE_VALUES, OM_S_TELETEX_STRING, "ap" },
        },
      },
    },
  },
}
```

```
    },
  },
},
{ DS_RDNS, OM_S_OBJECT,
  {
    { OM_CLASS, OM_S_OBJECT_ID_STRING, DS_C_DS_RDN },
    { DS_AVAS, OM_S_OBJECT,
      {
        { OM_CLASS, OM_S_OBJECT_ID_STRING, DS_C_AVA },
        { DS_ATTRIBUTE_TYPE, OM_S_OBJECT_ID_STRING, DS_A_COMMON_NAME },
        { DS_ATTRIBUTE_VALUES, OM_S_TELETEX_STRING, "dsa" },
      },
    },
  },
},
{ DS_RDNS, OM_S_OBJECT,
  {
    { OM_CLASS, OM_S_OBJECT_ID_STRING, DS_C_DS_RDN },
    { DS_AVAS, OM_S_OBJECT,
      {
        { OM_CLASS, OM_S_OBJECT_ID_STRING, DS_C_AVA },
        { DS_ATTRIBUTE_TYPE, OM_S_OBJECT_ID_STRING, DS_A_COMMON_NAME },
        { DS_ATTRIBUTE_VALUES, OM_S_TELETEX_STRING, "dsa-m1" },
      },
    },
  },
},
},
},
}; /* dsa_name */
```

Chapter 5

XOM Programming

XOM API defines a general-purpose interface for use in conjunction with other application-specific APIs for OSI services, such as XDS API to directory services or X.400 Application API to electronic mail service. It presents the application programmer with a uniform information architecture based on the concept of groups, classes, and similar information objects.

This chapter describes some of the basic concepts required to understand and use the XOM API effectively.

The following names refer to the complete XDS example programs, which can be found in Chapter 7:

- **acl.c (acl.h)**
- **example.c (example.h)**
- **teldir.c**

For multithreaded XDS/XOM applications, please refer to Chapter 8. For use of the XDS/XOM convenience functions please refer to Chapter 9.

5.1 OM Objects

The purpose of XOM API is to provide an interface to manage complex information objects. These information objects belong to classes and have attributes associated with them. There are two distinct kinds of classes and attributes that are used throughout the directory service documentation: *directory* classes and attributes and *OM* classes and attributes.

The directory classes and attributes defined for XDS API correspond to entries that make up the objects in the directory. These classes and attributes are defined in the X.500 directory standard and by additional GDS extensions created for DCE. Other APIs, such as the X.400 API, which is the application interface for the industry standard X.400 electronic mail service, define their own set of objects in terms of classes and attributes. OM classes and OM attributes are used to model the objects in the directory.

XOM API provides a common information architecture so that the information objects defined for any API that conforms to this architectural model can be shared. Different application service interfaces can communicate by using this common way of defining objects by means of workspaces. A workspace is simply a common work area where objects defined by a service can be accessed and manipulated. In turn, XOM API provides a set of standard functions that perform common operations on these objects in a workspace. Two different APIs can share information by copying data from one workspace to another.

5.1.1 OM Object Attributes

OM objects are composed of OM attributes. OM objects may contain zero or more OM attributes. Every OM attribute has zero or more values. An attribute comprises an integer that indicates the attribute's value. Each value is accompanied by an integer that indicates that value's syntax.

An OM attribute type is a category into which all the values of an OM attribute are placed on the basis of its purpose. Some OM attributes may either have zero, one, or multiple values. The OM attribute type is used as the name of the OM attribute.

A syntax is a category into which a value is placed on the basis of its form. **OM_S_PRINTABLE_STRING** is an example of a syntax.


An OM attribute value is an information item that can be viewed as a characteristic or property of the OM object of which it is a part.

OM attribute types and syntaxes have integer values and symbolic equivalents assigned to them for ease of use by naming authorities in the various API specifications. The integers that are assigned to the OM attribute type and syntax are fixed, but the attribute values may change. These OM attribute types and syntaxes are defined in the DCE implementation of XDS and XOM APIs in header files that are included with the software along with additional OM attributes specific to the GDS implementation.

Figure 5-1 shows the internal structure of an OM object.

Figure 5–1. The Internal Structure of an OM Object

object



For example, the tables in Figure 5-2 show the OM attributes, syntax, and values for the OM class **DS_C_ENTRY_INFO_SELECTION**, and how the integer values

are mapped to corresponding names in the **xom.h** and **xds.h** header files. The chapters in Part 4 of this guide contain tables for every OM class supported by the directory service. Refer to Chapter 11 for a complete description of **DS_C_ENTRY_INFO_SELECTION** and the accompanying table.

DS_C_ENTRY_INFO_SELECTION is a subclass of **OM_C_OBJECT**. This information is supplied in the description of this OM class in Chapter 19. As such, **DS_C_ENTRY_INFO_SELECTION** inherits the OM attributes of **OM_C_OBJECT**. The only OM attribute of **OM_C_OBJECT** is **OM_CLASS**. **OM_CLASS** identifies the object's OM class, which in this case is **DS_C_ENTRY_INFO_SELECTION**. **DS_C_ENTRY_INFO_SELECTION** identifies information to be extracted from a directory entry and has the following OM attributes, in addition to those inherited from **OM_C_OBJECT**:

- **DS_ALL_ATTRIBUTES**
- **DS_ATTRIBUTES_SELECTED**
- **DS_INFO_TYPE**

As part of an XDS function call, **DS_ALL_ATTRIBUTES** specifies to the directory service whether all the attributes of a directory entry are relevant to the application program. It can take the values **OM_TRUE** or **OM_FALSE**. These values are defined to be of syntax **OM_S_BOOLEAN**. The value **OM_TRUE** indicates that information is requested on all attributes in the directory entry. The value **OM_FALSE** indicates that information is only requested on those attributes that are listed in the OM attribute **DS_ATTRIBUTES_SELECTED**.

DS_ATTRIBUTES_SELECTED lists the types of attributes in the entry from which information is to be extracted. The syntax of the value is specified as **OM_S_OBJECT_IDENTIFIER_STRING**.

OM_S_OBJECT_IDENTIFIER_STRING contains an octet string of integers that are BER encoded object identifiers of the types of OM attributes in the OM attribute list. The value of **DS_ATTRIBUTES_SELECTED** is significant only if the value of **DS_ALL_ATTRIBUTES** is **OM_FALSE**, as described previously.

DS_INFO_TYPE identifies the information that is to be extracted from each OM attribute identified. The syntax of the value is specified as Enum(**DS_Information_Type**). **DS_INFO_TYPE** is an enumerated type that has two possible values: **DS_TYPES_ONLY** and **DS_TYPES_AND_VALUES**.

DS_TYPES_ONLY indicates that only the attribute types in the entry are returned by the directory service operation. **DS_TYPES_AND_VALUES** indicates that both the types and the values of the attributes in the directory entry are returned.

A typical directory service operation, such as a read operation (**ds_read()**), requires the *entry_information_selection* parameter to specify to the directory service the information to be extracted from the directory entry. This *entry_information_selection* parameter is built by the application program as a public object (Section 5.1.4 describes how to create a public object), and is included as a parameter to the **ds_read()** function call, as shown in the following code fragment from **example.c**:

```
/*
 * Public Object ("Descriptor List") for
 * Entry-Information-Selection
 * parameter to ds_read().
 */
OM_descriptor selection[] = {
    OM_OID_DESC(OM_CLASS, DS_C_ENTRY_INFO_SELECTION),
    { DS_ALL_ATTRIBUTES, OM_S_BOOLEAN, { OM_FALSE, NULL } },
    OM_OID_DESC(DS_ATTRIBUTES_SELECTED, DS_A_PHONE_NBR),
    { DS_INFO_TYPE, OM_S_ENUMERATION,
      { DS_TYPES_AND_VALUES, NULL } },
    OM_NULL_DESCRIPTOR
};

CHECK_DS_CALL(ds_read(session, DS_DEFAULT_CONTEXT,
                      name, selection, &result, &invoke_id));
```

Figure 5–2. Mapping the Class Definition of DS_C_ENTRY_INFO_SELECTION

OM Attributes of a OM_C_OBJECT				
Attribute	Value Syntax	Value Length	Value No.	Value Initially
OM_CLASS	String (OM_S_OBJECT_IDENTIFIER_STRING)	-	1	-

OM Attributes of a DS_C_ENTRY_INFO_SELECTION				
Attribute	Value Syntax	Value Length	Value No.	Value Initially
DS_ALL_ATTRIBUTES	OM_S_BOOLEAN	-	1	OM_TRUE
DS_ATTRIBUTES_SELECTED	String (OM_S_OBJECT_IDENTIFIER_STRING)	-	0 or more	-
DS_INFO_TYPE	Enum(DS_Information_Type)	-	1	DS_TYPES_AND_VALUES


```
#define OM_CLASS ((OM_type) 3)
#define OM_S_BOOLEAN ((OM_syntax) 1)
#define OM_S_OBJECT_IDENTIFIER_STRING ((OM_syntax) 6)
#define OM_S_ENUMERATION ((OM_syntax) 10)

enum DS_Information_Type {
    DS_TYPES_ONLY = 0,
    DS_TYPES_AND_VALUES = 1
};
    o
    o
    o

#define DS_ALL_ATTRIBUTES ((OM_TYPE) 707)
#define DS_ATTRIBUTES_SELECTED ((OM_TYPE) 710)
#define DS_INFO_TYPE ((OM_type) 734)
```

sample code from
from the xom.h header file

sample code from
from the xds.h header file

5.1.2 Object Identifiers

OM classes are uniquely identifiable by means of ASN.1 object identifiers. OM classes have mandatory and optional OM attributes. Each OM attribute has a type, value, and syntax. OM objects are instances of OM classes that are uniquely identifiable by means of ASN.1 object identifiers. The syntax of values defined for these OM object classes and OM attributes are representations at a higher level of abstraction so that implementors can provide the necessary high-level language binding for their own implementations of the various application interfaces, such as XDS API.

The DCE implementation uses C language to define the internal representation of OM classes and OM attributes. These definitions are supplied in the header files that are included as part of XDS and the XOM API.

OM classes are defined as symbolic constants that correspond to ASN.1 object identifiers. An ASN.1 object identifier is a sequence of integers that uniquely identifies a specific class. OM attribute type and syntax are defined as integer constants. These standardized definitions provide application programs with a uniform and stable naming environment in which to perform directory operations. Registration authorities are responsible for allocating the unique object identifiers.

The following code fragment from the **xdsbdcp.h** (the basic directory contents package) header file contains the symbolic constant **OMP_O_DS_A_COUNTRY_NAME**:

```
#ifndef dsP_attributeType /* joint-iso-ccitt(2) ds(5) attributeType(4) */
#define dsP_attributeType(X) ("x55x04" #X)
#endif

#define OMP_O_DS_A_COUNTRY_NAME          dsp_attributeType(\x06)
```

It resolves to **2.5.4.6**, which is the object identifier value for the **Country-Name** attribute type as defined in the directory standard. The symbolic constant for the directory object class **Country** resolves to **2.5.6.2**, the corresponding object identifier in the directory standard. OM classes are defined in the header files in the same manner.

5.1.3 C Naming Conventions

In the DCE implementation of XDS and XOM APIs, all object identifiers start with the letters **ds**, **DS**, **MH**, or **OMP**. Note that the interface reserves *all* identifiers starting with the letters **dsP** and **omP** for internal use by implementations of the interface. It also reserves all identifiers starting with the letters **dsX**, **DSX**, **omX**, and **OMX** for vendor-specific extensions of the interface. Applications programmers should not use any identifier starting with these letters.

The C identifiers for interface elements are formed by using the following conventions:

- XDS API function names are specified entirely in lowercase letters and are prefixed by **ds_** (for example, **ds_read()**).
- XOM API function names are specified entirely in lowercase letters and are prefixed by **om_** (for example, **om_get()**).
- C function parameters are derived from the parameter and result names and are specified entirely in lowercase letters. In addition, the names of results have **_return** added as a suffix (for example, **operation_status_return**).
- OM class names are specified entirely in uppercase letters and are prefixed by **DS_C_** and **MH_C_** (for example, **DS_C_AVA**).
- OM attribute names are specified entirely in uppercase letters and are prefixed by **DS_** and **MH_** (for example, **DS_RDNS**).
- OM syntax names are specified entirely in uppercase letters and are prefixed by **OM_S_** (for example, **OM_S_PRINTABLE_STRING**).
- Directory class names are specified entirely in uppercase letters and are prefixed by **DS_O** (for example, **DS_O_ORG_PERSON**).
- Directory attribute names are specified entirely in uppercase letters and are prefixed by **DS_A** (for example, **DS_A_COUNTRY_NAME**).
- Errors are treated as a special case. Constants that are the possible values of the OM attribute **DS_PROBLEM** of a subclass of the OM class **DS_C_ERROR** are specified entirely in uppercase letters and are prefixed by **DS_E_** (for example, **DS_E_BAD_CLASS**).
- The constants in the Value Length and Value Number columns of the OM class definition tables are also assigned identifiers. Where the upper limit in one of these columns is *not* 1, it is given a name that consists of the OM attribute name, prefixed by **DS_VL_** for value length, or **DS_VN_** for value number.

- The sequence of octets for each object identifier is also assigned an identifier for internal use by certain OM macros. These identifiers are all uppercase letters and are prefixed by **OMP_O_**.

Tables 5-1 and 5-2 summarize the XDS and XOM naming conventions.

Table 5–1. C Naming Conventions for XDS

Item	Prefix
Reserved for implementors	dsP
Reserved for interface extensions	dsX
Reserved for interface extensions	DSX
XDS functions	ds_
Error problem values	DS_E_
OM class names	DS_C_, MH_C_
OM attribute names	DS_, MH_
OM value length limits	DS_VL_
OM value number limits	DS_VN_
Other constants	DS_, MH_
Attribute type	DS_A_
Object class	DS_O_

Table 5–2. C Naming Conventions for XOM

Element Type	Prefix
Data type	OM_
Data value	OM_
Data value (class)	OM_C_
Data value (syntax)	OM_S_
Data value component (structure member)	None

Element Type	Prefix
Function	om_
Function parameter	None
Function result	None
Macro	OM_
Reserved for use by implementors	OMP
Reserved for use by implementors	omP
Reserved for proprietary extension	omX
Reserved for proprietary extension	OMX

5.1.4 Public Objects

The ultimate aim of an application program is access to the directory to perform some operation on the contents of the directory. A user may request the telephone number or electronic mail address of a fellow employee. In order to access this information, the application performs a read operation on the directory so that information is extracted about a target object in the directory and manipulated locally within the application.

XDS functions that perform directory operations, such as **ds_read()**, require *public* and *private* objects as input parameters. Typically, a public object is generated by an application program and contains the information required to access a target directory object. This information includes the AVAs and RDNs that make up a distinguished name of an entry in the directory. However, an application program may also generate a private object. Private objects are described in Section 5.1.5.

A public object is created by using OM classes and OM attributes. These OM classes and OM attributes model the target object entry in the directory and provide other information required by the directory service to access the directory.

5.1.4.1 Descriptor Lists

A public object is represented by a sequence of **OM_descriptor** data structures that is built by the application program. A descriptor contains the type, syntax, and value for an OM attribute in a public object.

The data structure **OM_descriptor** is defined in the **xom.h** header file as follows:

```
typedef struct OM_descriptor_struct {
    OM_type           type;
    OM_syntax         syntax;
    union OM_value_union value;
}OM_descriptor;
```

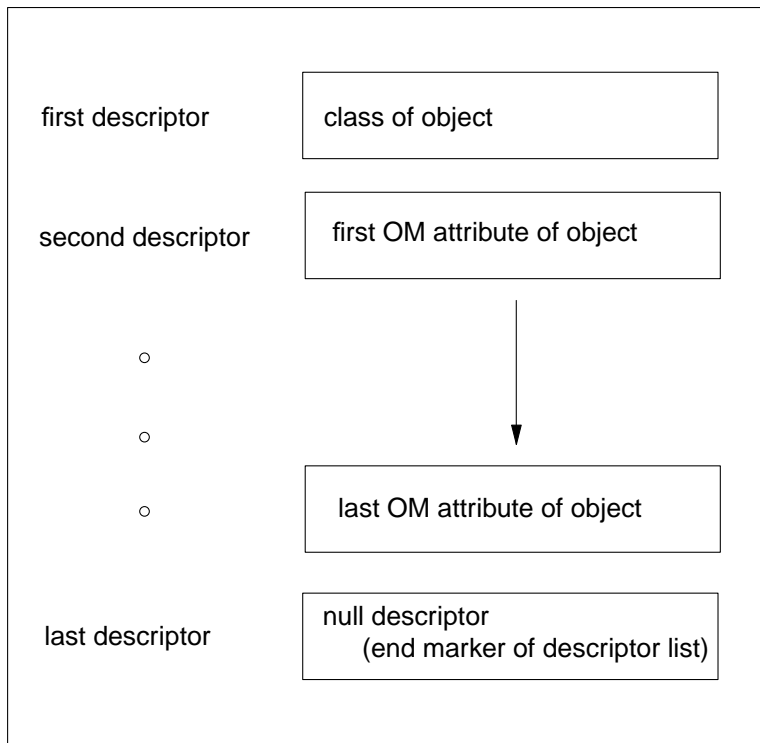
Figure 5-3 shows the representation of a public object in a descriptor list. The first descriptor in the list indicates the object's OM class; the last descriptor is a NULL descriptor that signals the end of the list of OM attributes. In between the first and the last descriptor are the descriptors for the OM attributes of the object.

For example, the following represents the public object **country** in **example.c**:

```
static OM_descriptor    country[] = {
    OM_OID_DESC(OM_CLASS, DS_C_AVA),
    OM_OID_DESC(DS_ATTRIBUTE_TYPE, DS_A_COUNTRY_NAME),
    { DS_ATTRIBUTE_VALUES, OM_S_PRINTABLE_STRING, OM_STRING("US") },
    OM_NULL_DESCRIPTOR
};
```

Figure 5–3. A Representation of a Public Object By Using a Descriptor List

object



The descriptor list is an array of data type **OM_descriptor** that defines the OM class, OM attribute types, syntax, and values that make up a public object.

The first descriptor gives the OM class of the object. The OM class of the object is defined by the OM attribute type, **OM_CLASS**. The **OM_OID_DESC** macro initializes the syntax and value of an object identifier, in this case to OM class **DS_C_AVA**, with the syntax of **OM_S_OBJECT_IDENTIFIER_STRING**. **OM_S_OBJECT_IDENTIFIER_STRING** is an OM syntax type that is assigned by definition in the macro to any OM attribute type and value parameters input to it.

The second descriptor defines the first OM attribute type, **DS_ATTRIBUTE_TYPE**, which has as its value **DS_A_COUNTRY_NAME** and syntax **OM_S_OBJECT_IDENTIFIER_STRING**.

The third descriptor specifies the AVA of an object entry in the directory. The **OM_OID_DESC** macro is not used here because **OM_OID_DESC** is only used to initialize values having **OM_S_OBJECT_IDENTIFIER_STRING** syntax. The OM attribute type is **DS_ATTRIBUTE_VALUES**, the syntax is **OM_S_PRINTABLE_STRING**, and the value is **US**. The **OM_STRING** macro creates a data value for a string data type (data type **OM_string**), in this case **OM_S_PRINTABLE_STRING**. A string is specified in terms of its length or whether or not it terminates with a NULL. (The **OM_STRING** macro is described in Section 5.8.2.4.)

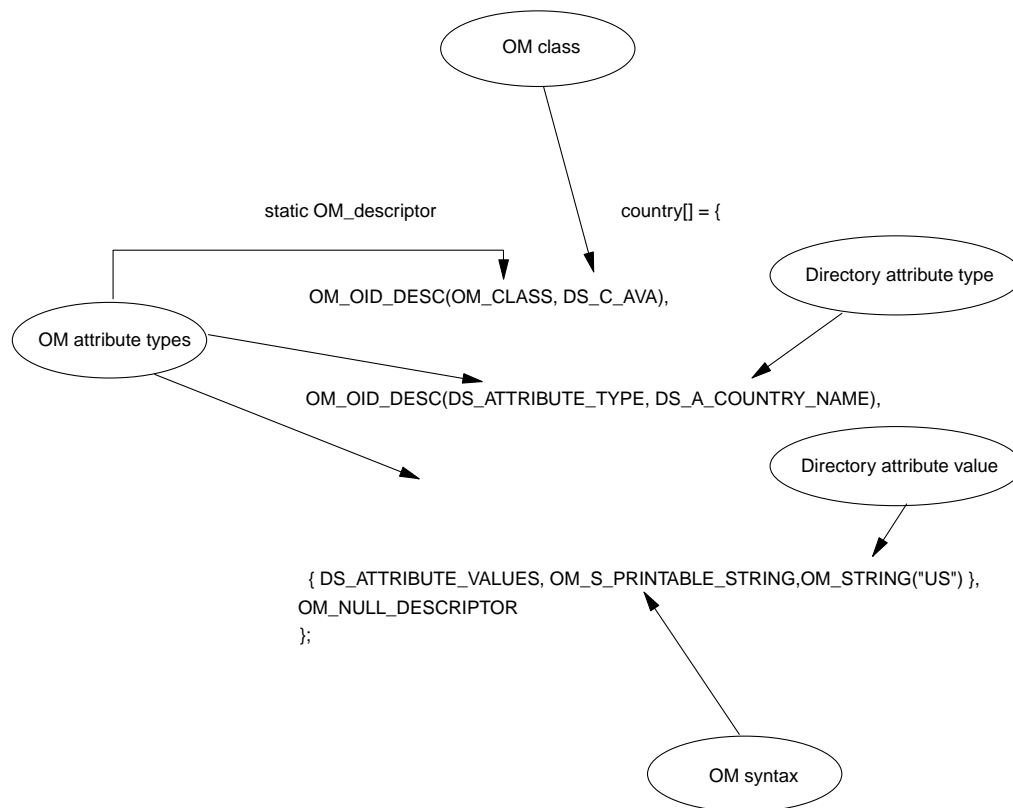
The last descriptor is a NULL descriptor that marks the end of the public object definition. It is defined in the **xom.h** header file as follows:

```
#define OM_NULL_DESCRIPTOR
{ OM_NO_MORE_TYPES, OM_S_NO_MORE_SYNTAXES,
  { { 0, OM_ELEMENTS_UNSPECIFIED } } }
```

OM_NULL_DESCRIPTOR is OM attribute type **OM_NO_MORE_TYPES**, syntax **OM_S_NO_MORE_SYNTAXES**, and value **OM_ELEMENTS_UNSPECIFIED**.

Figure 5-4 shows the composition of a descriptor list representing a public object.

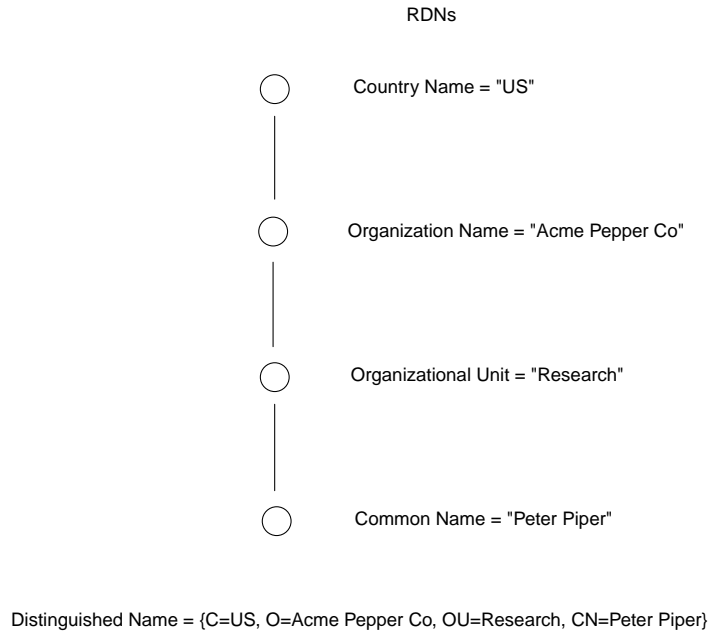
Figure 5-4. A Descriptor List for the Public Object: country



5.1.4.2 Building the Distinguished Name as a Public Object

Recall that RDNs are built from AVAs, and a distinguished name is built from a series of RDNs. In a typical application program, several AVAs are defined in descriptor lists as public objects. These public objects are incorporated into descriptor lists that represent corresponding RDNs. Finally, the RDNs are incorporated into one descriptor list that represents the distinguished name of an object in the directory, as shown in Figure 5-5. This descriptor list is included as one of the input parameters to a directory service function.

Figure 5–5. The Distinguished Name of “Peter Piper” in the DIT



The following code fragment from **example.c** shows how a distinguished name is built as a public object. The public object is the *name* parameter for a subsequent read operation call to the directory. The representation of a distinguished name in the DIT is shown in Figure 5-5.

The first section of code defines the four AVAs. These AVAs make the assertion to the directory service that the attribute values in the distinguished name of **Peter Piper** are valid and can therefore be read from the directory. The country name is **US**, the organization name is **Acme Pepper Co**, the organizational unit name is **Research**, and the common name is **Peter Piper**.

```

/*
 * Public Object ("Descriptor List") for Name parameter to
 * ds_read().
 * Build the Distinguished-Name of Peter Piper
 */

```

```
static OM_descriptor      country[] = {
    OM_OID_DESC(OM_CLASS, DS_C_AVA),
    OM_OID_DESC(DS_ATTRIBUTE_TYPE, DS_A_COUNTRY_NAME),
    { DS_ATTRIBUTE_VALUES, OM_S_PRINTABLE_STRING, OM_STRING("US") },
    OM_NULL_DESCRIPTOR
};

static OM_descriptor      organization[] = {
    OM_OID_DESC(OM_CLASS, DS_C_AVA),
    OM_OID_DESC(DS_ATTRIBUTE_TYPE, DS_A_ORG_NAME),
    { DS_ATTRIBUTE_VALUES, OM_S_TELETEX_STRING,
      OM_STRING("Acme Pepper Co") },
    OM_NULL_DESCRIPTOR
};

static OM_descriptor      organizational_unit[] = {
    OM_OID_DESC(OM_CLASS, DS_C_AVA),
    OM_OID_DESC(DS_ATTRIBUTE_TYPE, DS_A_ORG_UNIT_NAME),
    { DS_ATTRIBUTE_VALUES, OM_S_TELETEX_STRING, OM_STRING("Research") },
    OM_NULL_DESCRIPTOR
};

static OM_descriptor      common_name[] = {
    OM_OID_DESC(OM_CLASS, DS_C_AVA),
    OM_OID_DESC(DS_ATTRIBUTE_TYPE, DS_A_COMMON_NAME),
    { DS_ATTRIBUTE_VALUES, OM_S_TELETEX_STRING, OM_STRING("Peter Piper") },
    OM_NULL_DESCRIPTOR
};
```

The next section of code is nested one level above the previously defined AVAs. Each RDN has a descriptor with OM attribute type **DS_AVAS** (indicating that it is OM attribute type **AVA**), a syntax of **OM_S_OBJECT**, and a value of the name of the descriptor array defined in the previous section of code for an AVA. The **rdn1** descriptor contains the descriptor list for the AVA **country**, the **rdn2** descriptor contains the descriptor list for the AVA **organization**, and so on.

OM_S_OBJECT is a syntax that indicates that its value is a subobject. For example, the value for **DS_AVAS** is the previously defined object **country**. In this manner, a hierarchy of linked objects and subobjects can be constructed.

```
static OM_descriptor      rdn1[] = {
    OM_OID_DESC(OM_CLASS, DS_C_DS_RDN),
```

```

    { DS_AVAS, OM_S_OBJECT, { 0, country } },
    OM_NULL_DESCRIPTOR
};
static OM_descriptor rdn2[] = {
    OM_OID_DESC(OM_CLASS, DS_C_DS_RDN),
    { DS_AVAS, OM_S_OBJECT, { 0, organization } },
    OM_NULL_DESCRIPTOR
};
static OM_descriptor rdn3[] = {
    OM_OID_DESC(OM_CLASS, DS_C_DS_RDN),
    { DS_AVAS, OM_S_OBJECT, { 0, organizational_unit } },
    OM_NULL_DESCRIPTOR
};
static OM_descriptor rdn4[] = {
    OM_OID_DESC(OM_CLASS, DS_C_DS_RDN),
    { DS_AVAS, OM_S_OBJECT, { 0, common_name } },
    OM_NULL_DESCRIPTOR
};

```

The next section of code contains the RDNs that make up the distinguished name, which is stored in the array of descriptors called *name*. It is made up of the OM class **DS_C_DS_DN** (representing a distinguished name) and four RDNs of OM attribute type **DS_RDNS** and syntax **OM_S_OBJECT**.

```

OM_descriptor name[] = {
    OM_OID_DESC(OM_CLASS, DS_C_DS_DN),
    { DS_RDNS, OM_S_OBJECT, { 0, rdn1 } },
    { DS_RDNS, OM_S_OBJECT, { 0, rdn2 } },
    { DS_RDNS, OM_S_OBJECT, { 0, rdn3 } },
    { DS_RDNS, OM_S_OBJECT, { 0, rdn4 } },
    OM_NULL_DESCRIPTOR
};

```

In summary, the distinguished name for **Peter Piper** is stored in the array of descriptors called *name*, which is composed of three nested levels of arrays of descriptors (see Figure 5-6). The definitions for the AVAs are at the innermost level, the definitions for RDNs are at the next level up, and the distinguished name is at the top level.

Figure 5–6. Building a Distinguished Name

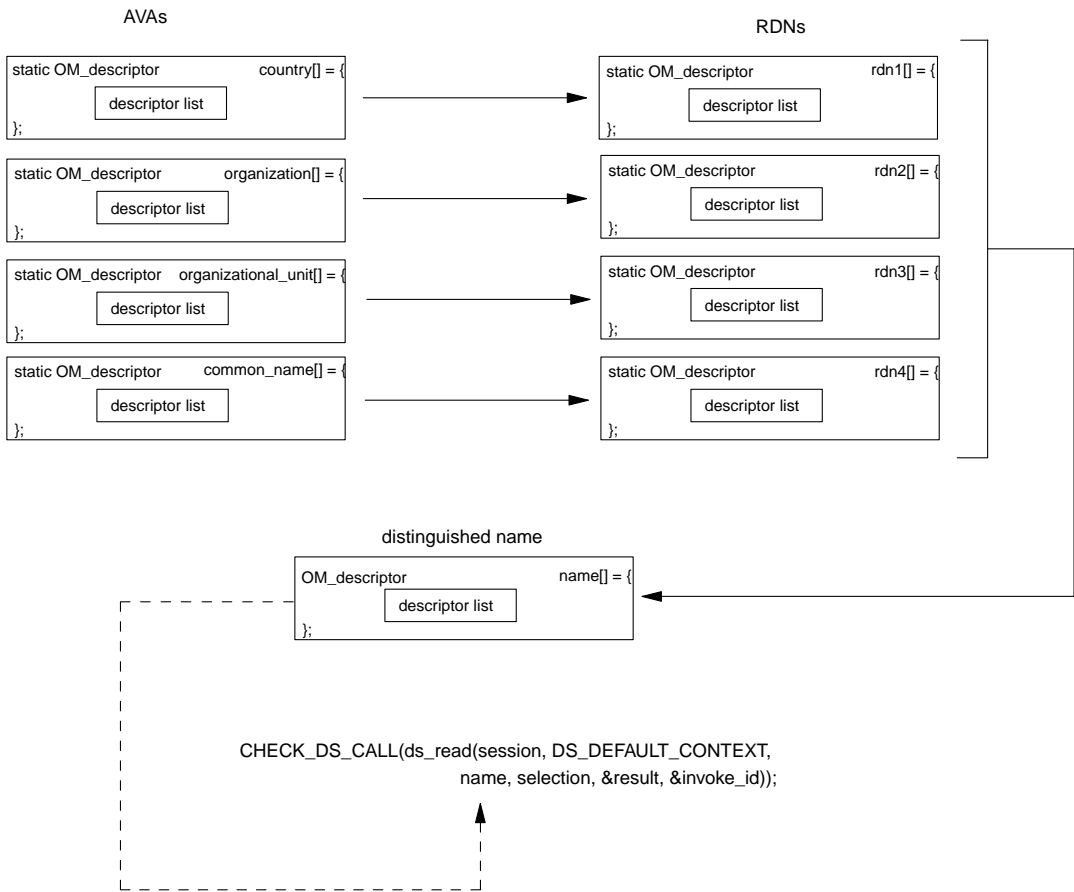
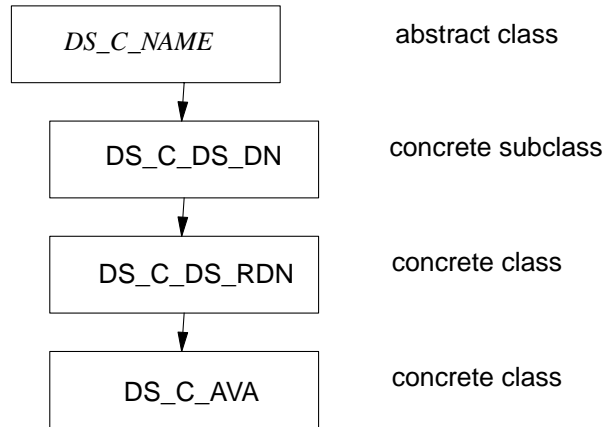


Figure 5-7 shows a more general view of the structure distinguished name.

Figure 5–7. A Simplified View of the Structure of a Distinguished Name



The **name** descriptor defines a public object that is provided as the *name* parameter required by the XDS API read function call, **ds_read()**, as follows. (XDS API function calls are described in detail in Chapter 6.)

```
CHECK_DS_CALL(ds_read(session, DS_DEFAULT_CONTEXT,
                      name, selection, &result, &invoke_id));
```

The result of the **ds_read()** function call is in a private implementation-specific format; it is stored in a workspace and pointed to by **result**. The application program must use XOM function calls (described in Section 5.7) to interpret the data and extract the information. This extraction process involves uncovering the nested data structures in a series of XOM function calls.

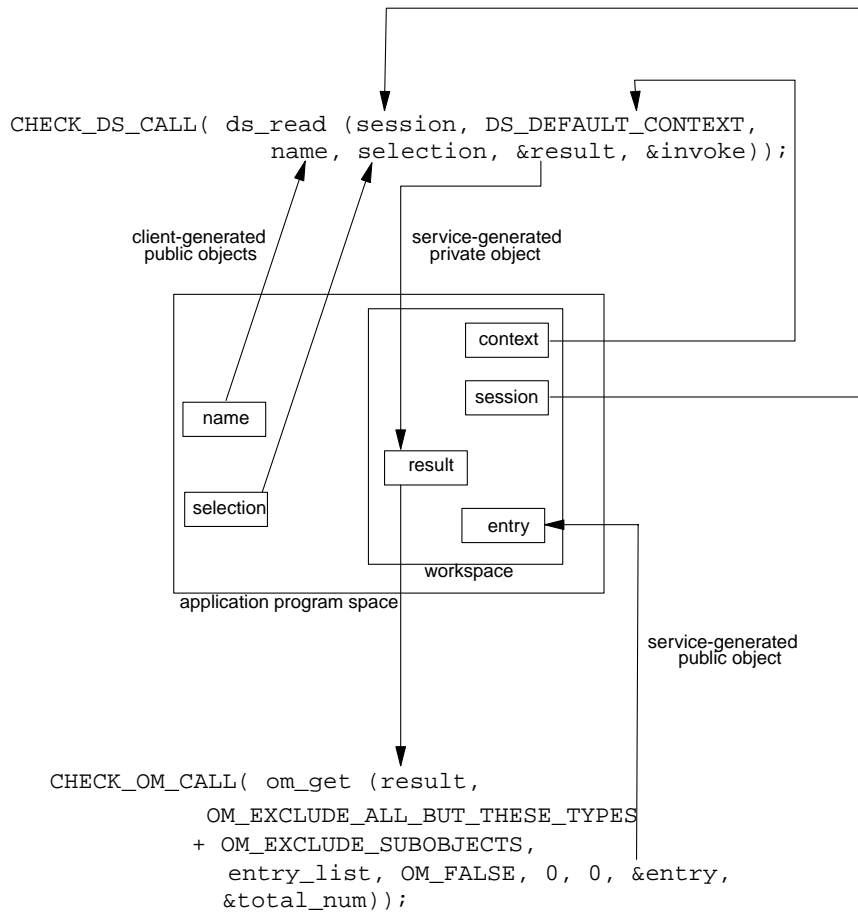
5.1.4.3 Client-Generated and Service-Generated Public Objects

There are two types of public objects: service-generated objects and client-generated objects. The distinguished name object described in the previous section is a client-generated public object because an application program (the client) created the data structure. As the creator of the public object, it is the responsibility of the application program to manage the memory resources allocated for it.

Service-generated public objects are created by the XOM service. Service-generated public objects may be generated as a result of an XOM request. An XOM API function, such as **om_get()**, converts a private object into a service-generated public object. This is necessary because XDS may return a pointer to data in private format that can only be interpreted by XOM functions such as **om_get()**.

For example, Figure 5-8 shows how the read request described in the previous example returns a pointer to an encoded data structure stored in **result**. This encoded data structure, referred to as a *private object* (described in the next section) is one of the input parameters to **om_get()**. The **om_get()** function provides a pointer to a public object (in this case, **entry**) as an output parameter. The public object is a data structure that has been interpreted by **om_get()** and is accessible by the application program (the client). The information requested by the application in the read request is contained in the output parameter **entry**.

Figure 5–8. Client-Generated and Service-Generated Objects



The application program is responsible for managing the storage (memory) for the service-generated public object. This is an important point because it requires that the application issue a series of **om_delete()** calls to delete the service-generated public object from memory. Because the data structures involved with directory service requests can be very large (often involving large subtrees of the DIT), it is imperative that the application programmer build into any application program the efficient management of memory resources.

The following code fragment from **example.h** demonstrates how storage for public and private objects is released by using a series of **om_delete()** function calls after they are no longer needed by the application program. The data (a list of phone numbers associated with the name **Peter Piper** required by the application program) has already been extracted by using a series of **om_get()** function calls, as follows:

```
/* We can now safely release all the private objects
 * and the public objects we no longer need
 */
CHECK_OM_CALL(om_delete(session));
CHECK_OM_CALL(om_delete(result));
CHECK_OM_CALL(om_delete(entry));
CHECK_OM_CALL(om_delete(attributes));
CHECK_DS_CALL(ds_shutdown(workspace));
```

5.1.5 Private Objects

Private objects are created dynamically by the service interface. In Figure 5-8, the **ds_read()** function returns a pointer to the data structure **result** in the workspace. This service-generated data structure is a private object in a private implementation-specific format, which requires a call to **om_get()** to interpret the data. A private object is one of the required input parameters to XOM API functions (such as **om_get()**), as shown in Figure 5-8. Private objects are always service generated.

Table 5-3 compares private and public objects.

Table 5–3. Comparison of Private and Public Objects

Private	Public
Representation is implementation specific	Representation is defined in the API specification
Not directly accessible by the client	Directly accessible by the client
Manipulated by the client by using OM functions	Manipulated by the client by using programming constructs

Private	Public
Created in storage provided by the service	Is a service-generated object if created by the service Is a client-generated object if created by the client in storage provided by the client
Cannot be modified by the client directly, except through the service interface	If a client-generated object, can be modified directly by the client If a service-generated object, cannot be modified directly by the client, except through the service interface
Storage is allocated and released by the service	If a service-generated object, storage is allocated and released by the service If a client-generated object, storage is allocated and released by the client

Private objects can also be used as input to XOM and XDS API functions to improve program efficiency. For example, the output of a **ds_search()** request can be used as input to **ds_read()**. The search request returns the name of each entry in the search. If the application program requires the address and telephone number of each name, a **ds_read()** operation can be performed on each name as a private object.

5.1.6 Object Classes

Objects are categorized into OM classes based on their purpose and internal structure. An object is an instance of its OM class. An OM class is characterized by OM attribute types that may appear in its instances. An OM class is uniquely identified by an ASN.1 object identifier.

Later in this section, it will be shown how OM classes are organized into groups of OM classes, called *packages*, that support some aspect of the directory service.

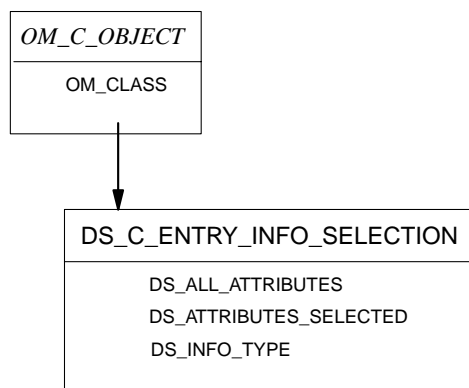
5.1.6.1 OM Class Hierarchy and Inheritance Properties

OM classes are related to each other in a tree hierarchy whose root is a special OM class called **OM_C_OBJECT**. Each of the other OM classes is the immediate subclass of precisely one other OM class. This tree structure is known as the *OM class hierarchy*. It is important because of the property of inheritance. The OM class hierarchy is defined by the XDS/XOM standards. DCE implements this hierarchy for GDS and adds its own set of OM classes defined in the GDS package.

The OM attribute types that may exist in an instance of an OM class, but not in an instance of the OM class above it in the tree hierarchy, are said to be *specific* to that OM class. OM attributes that may appear in an object are those specific to its OM class as well as those inherited from OM classes above it in the tree. OM classes above an instance of an OM class in the tree are *superclasses* of that OM class. OM classes below an instance of an OM class are *subclasses* of that OM class.

For example, as shown in Figure 5-9, **DS_C_ENTRY_INFO_SELECTION** inherits its OM attributes from its superclass **OM_C_OBJECT**. The OM attributes **DS_ALL_ATTRIBUTES**, **DS_ATTRIBUTES_SELECTED**, and **DS_INFO_TYPE** are attributes specific to the OM class **DS_C_ENTRY_INFO_SELECTION**. The **DS_C_ENTRY_INFO_SELECTION** class has no subclasses.

Figure 5–9. The OM Class DS_C_ENTRY_INFO_SELECTION



Another important point about OM class inheritance is that an instance of an OM class is also considered to be an instance of each of its superclasses and may appear wherever the interface requires an instance of any of those superclasses. For example, **DS_C_DS_DN** is a subclass of **DS_C_NAME**. Everywhere in an application program where **DS_C_NAME** is expected at the interface (as a parameter to **ds_read()**, for example), it is permitted to supply **DS_C_DS_DN**.

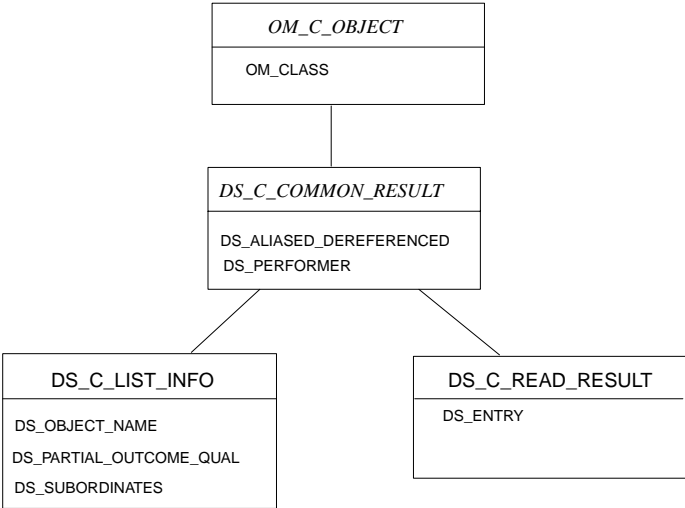
5.1.6.2 Abstract and Concrete Classes

OM classes are defined as being either *abstract* or *concrete*. An abstract OM class is an OM class in which instances are not permitted. An abstract OM class may be defined so that subclasses can share a common set of OM attributes between them.

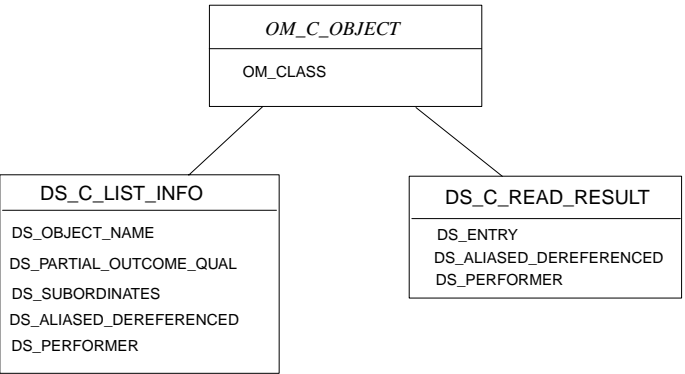
In contrast to abstract OM classes, instances of OM concrete classes are permitted. However, the definition of each OM concrete class may include the restriction that a client not be allowed to create instances of that OM class. For example, consider two alternative means of defining the OM classes used in XDS: **DS_C_LIST_INFO** and **DS_C_READ_RESULT**. **DS_C_LIST_INFO** and **DS_C_READ_RESULT** are subclasses of the abstract OM class **DS_C_COMMON_RESULT**.

Figure 5-10 shows the relationship of **DS_C_LIST_INFO** and **DS_C_READ_RESULTS** when the abstract OM class **DS_C_COMMON_RESULT** is defined and when it is not defined. It demonstrates that the presence of an abstract OM class enables the programmer to develop applications that process information more efficiently.

Figure 5–10. Comparison of Two Classes With/Without an Abstract OM Class



DS_C_LIST_INFO and DS_C_READ_RESULT with the DS_C_COMMON_RESULT abstract class defined



DS_C_LIST_INFO and DS_C_READ_RESULT without the DS_C_COMMON_RESULT abstract class defined

The following list contains the hierarchy of concrete and abstract OM classes in the directory service package. Abstract OM classes are shown in italics. The indentation shows the class hierarchy; for example, the abstract class *OM_C_OBJECT*

is a superclass of the abstract class *DS_C_COMMON_RESULTS*, which in turn is a superclass of the concrete class **DS_C_COMPARE_RESULT**.

OM_C_OBJECT

- **DS_C_ACCESS_POINT**
- *DS_C_ADDRESS*
 - **DS_C_PRESENTATION_ADDRESS**
- **DS_C_ATTRIBUTE**
 - **DS_C_AVA**
 - **DS_C_ENTRY_MOD**
 - **DS_C_FILTER_ITEM**
- **DS_C_ATTRIBUTE_ERROR**
- **DS_C_ATTRIBUTE_LIST**
 - **DS_C_ENTRY_INFO**
- *DS_C_COMMON_RESULTS*
 - **DS_C_COMPARE_RESULT**
 - **DS_C_LIST_INFO**
 - **DS_C_READ_RESULT**
 - **DS_C_SEARCH_INFO**
- **DS_C_CONTEXT**
- **DS_C_CONTINUATION_REF**
 - **DS_C_REFERRAL**
- **DS_C_ENTRY_INFO_SELECTION**
- **DS_C_ENTRY_MOD_LIST**
- *DS_C_ERROR*
 - **DS_C_ABANDON_FAILED**
 - **DS_C_ATTRIBUTE_PROBLEM**
 - **DS_C_COMMUNICATIONS_ERROR**

- **DS_C_LIBRARY_ERROR**
- **DS_C_NAME_ERROR**
- **DS_C_SECURITY_ERROR**
- **DS_C_SERVICE_ERROR**
- **DS_C_SYSTEM_ERROR**
- **DS_C_UPDATE_ERROR**
- **DS_C_EXT**
- **DS_C_FILTER**
- **DS_C_LIST_INFO_ITEM**
- **DS_C_LIST_RESULT**
- *DS_C_NAME*
 - **DS_C_DS_DN**
- **DS_C_OPERATION_PROGRESS**
- **DS_C_PARTIAL_OUTCOME_QUAL**
- *DS_C_RELATIVE_NAME*
 - **DS_C_DS_RDN**
- **DS_C_SEARCH_RESULT**
- **DS_C_SESSION**

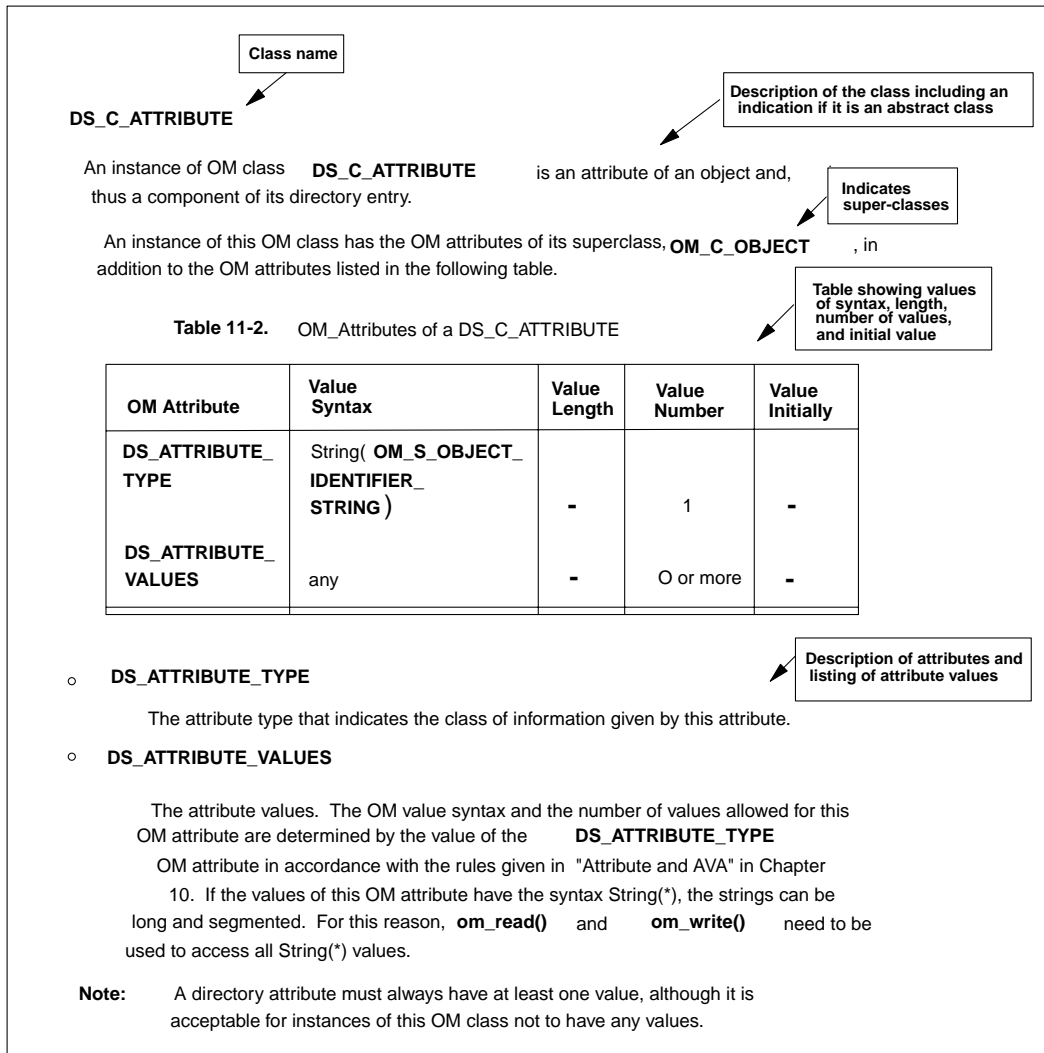
In summary, an OM class is defined with the following elements:

- OM class name (indicated by an object identifier)
- Identity of its immediate superclass
- Definitions of the OM attribute types specific to the OM class
- Indication whether the OM class is abstract or concrete
- Constraints on the OM attributes

A complete description of OM classes, OM attributes, syntaxes, and values that are defined for XDS and XOM APIs are described in Part 4. Tables and textual

descriptions, such as the one shown in Figure 5-11 for the concrete OM class **DS_C_ATTRIBUTE**, are provided for each OM class.

Figure 5–11. Complete Description of Concrete OM Class **DS_C_ATTRIBUTE**



The table shown in Figure 5-11 provides information under the following headings:

- OM Attribute

This is the name of each of the OM attributes.

- Value Syntax

This provides the syntaxes of each of the OM attribute's values.

- Value Length

This describes any constraints on the number of bits, octets, or characters in each value that is a string.

- Value Number

This describes any constraints on the number of values.

- Value Initially

This is any value with which the OM attribute can be initialized.

An OM class can be constrained to contain only one member of a set of OM attributes. In turn, OM attributes can be restricted to having no more than a fixed number of values, either 0 (zero) or 1 as an optional value, or exactly one mandatory value.

An OM attribute's value may be also constrained to a single syntax. That syntax can be further restricted to a subset of defined values.

An object passed as a parameter to an XOM and XDS function call needs to meet a minimum set of conditions, as follows:

- The type of each OM attribute must be specific to the object's OM class or one of its superclasses.
- The number of values of each OM attribute must be within OM class limits.
- The syntax of each value must be among those the OM class permits.
- The number of bits, octets, or characters in each string value must be within OM class limits.

5.2 Packages

A *package* is a collection of OM classes that are grouped together, usually by function. The packages themselves are features that are negotiated with the directory service by using the XDS function **ds_version()**. Consider which OM classes will be required for your application programs and determine the packages that contain these OM classes.

A package is uniquely identified by an ASN.1 object identifier. DCE XDS API supports the following five packages, of which one is mandatory and four are optional:

- The directory service package (mandatory)
- The basic directory contents package (optional)
- The strong authentication package (optional)
- The GDS package (optional)
- The message handling system (MHS) directory user package (optional)

5.2.1 The Directory Service Package

The directory service package is the default package and as such does not require negotiation. The optional packages have to be negotiated with the directory service by using the **ds_version()** function.

The object identifiers for specific packages are defined in header files that are part of the XDS API and XOM API. An object identifier consists of a string of integers. The header files include **#define** preprocessor statements that assign names to the constants in order to make them more readable. For the application programmer, these assignments alleviate the burden of maintaining strings of integers. For example, the object identifiers for the directory service package are defined in **xds.h**. The **xds.h** header file contains OM class and OM attribute names, OM object constants, and defines prototypes for XDS API functions, as shown in the following code fragment from **xds.h**:

```
/* DS package object identifier */
/* {iso(1) identifier-organization(3) icd-ecma(12)
 * member-company(2)
```

```
* dec(1011) xopen(28) dsp(0) } */

#define OMP_O_DS_SERVICE_PKG    "\x2B\x0C\x02\x87\x1C\x00"
```

A **ds_version()** function call must be included within an application program to negotiate the optional features (packages) with the directory service. The first step is to build an array of object identifiers for the optional packages to be negotiated (the basic directory contents package and the GDS package), as shown in the following code fragment from the **acl.h** header file:

```
DS_feature features[] = {
    { OM_STRING(OMP_O_DS_BASIC_DIR_CONTENTS_PKG), OM_TRUE },
    { OM_STRING(OMP_O_DSX_GDS_PKG), OM_TRUE },
    { 0 }
};
```

The **OM_STRING** macro is provided for creating a data value of data type **OM_string** for octet strings and characters. XOM API macros are described in Section 5.8.2.

The array of object identifiers is stored in **features** and passed as an input parameter to **ds_version()**, as shown in the following code fragment from **acl.c**:

```
/* Negotiate the use of the BDC and GDS packages. */

if (ds_version(features) != DS_SUCCESS)
    printf("ds_version() error\n");
```

5.2.2 The Basic Directory Contents Package

The basic directory contents package contains the object identifier definition of directory classes and attribute types as defined by the X.500 standard. These definitions allow the creation and maintenance of directory entries for a number of common objects so that the representation of all such objects is the same throughout the directory. Also included are the definitions of the OM classes and OM attributes required to support the directory attribute types. Chapter 12 describes the basic directory contents package in detail.

The object identifier associated with the basic directory contents package is shown in the following code fragment from the **xdsbdcp.h** header file:

```
/* BDC package object identifier */
/* { iso(1) identifier-organization(3) icd-ecma(12)
 *   member-company (2)
 *   dec(1011) xopen(28) bdc(1) } */

#define OMP_DS_BASIC_DIR_CONTENTS_PKG "\x2B\x0C\x02\x87\x73\x1C\x01"
```

5.2.3 The Strong Authentication Package

The strong authentication package contains the object identifier definition of directory classes and attribute types as defined by the X.500 standard for security purposes. Also included are the definitions of the OM classes and OM attributes required to support these security attribute types. Chapter 13 describes the strong authentication package in detail.

The object identifier associated with the strong authentication package is shown in the following code fragment from the **xdsap.h** header file:

```
/* SA package object identifier */
/* { iso(1) identifier-organization(3) icd-ecma(12)
 *   member-company (2)
 *   dec(1011) xopen(28) sap(2) } */

#define OMP_DS_STRONG_AUTHENT_PKG "\x2B\x0C\x02\x87\x73\x1C\x02"
```

5.2.4 The GDS Package

The GDS package contains the definition of a DCE extension to the XDS API. It contains the definitions of OM classes, OM attributes, and syntaxes to support extended functionality specific to DCE. Chapter 15 describes the GDS package in detail.

The following code fragment from the **xds_gds.h** header file shows the object identifier for the GDS package:

```
/* GDS package object identifier */
/* { iso(1) identifier-organization(3) icd-ecma(12) member-company (2)
/*   siemens-units(1107) sni(1) directory(3) xds-api(100) gdsp(1) } */

#define OMP_O_DSX_GDS_PKG    "\x2B\x0C\x02\x88\x53\x01\x03\x64\x01"
```

5.2.5 The MHS Directory User Package

The MHS directory user package contains definitions to support the use of the directory in accordance with the standard X.400 (1988) User Agents and Message Transfer Agents (MTAs) for name resolution, distribution list expansion, and capability assessment. The definitions are based on the attribute types and syntaxes specified in X.402, Annex A. The definitions of the OM classes and OM attributes required to support these MHS attribute types are also included with this package. Chapter 14 describes the MHS directory user package in detail.

The object identifier associated with the MHS directory user package is shown in the following code fragment from the **xdsmdup.h** header file:

```
/* MDU package object identifier */
/* { iso(1) identifier-organization(3) icd-ecma(12)
/*   member-company (2)
/*   dec(1011) xopen(28) mdup(3) } */

#define OMP_DS_MHS_DIR_USER_PKG "\x2B\x0C\x02\x87\x73\x1C\x03"
```

Part 4 of this guide describes in detail the attributes and data types that make up the OM and directory classes defined in the XDS API packages. Chapter 7 examines in detail how these packages are used in developing the sample application programs.

5.2.6 Package Closure

An OM class can be defined to have an attribute whose OM class is defined in some other package in order to avoid duplication of OM classes. This gives rise to the concept of a package closure. A package closure is the set of all OM classes that need to be supported so that all possible instances of all OM classes can be defined in the package.

5.3 Workspaces

Two application-specific APIs or two different implementations of the same service require work areas, called *workspaces*, to maintain private and public (service-generated) objects. The workspace is required because two implementations of the same service (or different services) can represent private objects differently. Each one has its own workspace. Using the functions provided by XOM API, such as **om_get()** and **om_copy()**, objects can be copied and moved from one workspace to another.

Recall that private objects are returned by a service to a workspace in private implementation-specific format. Using the OM function calls described in Section 5.7, the data can be extracted from the private object for further program processing.

Before a request to the directory can be made by an application program, a workspace must be created by using the appropriate XDS function. An application creates a workspace by performing the XDS API call **ds_initialize()**. Once the workspace is obtained, subsequent XDS API calls, such as **ds_read()**, return a pointer to a private object in the workspace. When program processing is completed, the workspace is destroyed by using the **ds_shutdown()** XDS API function. Implicit in **ds_shutdown()** is a call to the XOM API function **om_delete()** to delete each private object the workspace contains.

The programs in Chapter 7 demonstrate how to initialize and shut down a workspace. The XDS functions **ds_initialize()** and **ds_shutdown()** are described in detail in Chapter 6.

The closures of one or more packages are associated with a workspace. A package can be associated with any number of workspaces. An application program must

obtain a workspace that supports an OM class before it is able to create any instances of that OM class. For example, some of these operations in an application may require involvement with GDS security, ACLs, or the DUA cache. Therefore, in addition to the basic packages provided by the directory service APIs, the workspace would have to support the GDS package. The following code fragment demonstrates how an application initializes a workspace and negotiates the packages to be associated with that workspace:

```
/* Build up an array of object identifiers for the optional */
/* packages to be negotiated.                                     */

DS_feature features[] = {
    { OM_STRING(OMP_O_DS_BASIC_DIR_CONTENTS_PKG), OM_TRUE },
    { OM_STRING(OMP_O_DSX_GDS_PKG), OM_TRUE },
    { 0 }
};

CHECK_DS_CALL((OM_object) !(workspace = ds_initialize()));

CHECK_DS_CALL(ds_version(bdcp_package, workspace));
```

5.4 Storage Management

An object occupies storage. The storage occupied by a public object is allocated by the client, and is, therefore, directly accessible by the client and can be released by the client. The storage occupied by a private object is not accessible by the client and must be managed indirectly by using XOM function calls.

Objects are accessed by an application program via object handles. Object handles are used as input parameters to interface functions by the client and returned as output parameters by the service. The object handle for a public object is simply a pointer to the data structure (an array of descriptors) containing the object OM attributes. The object handle for a private object is a pointer to a data structure that is in private implementation-specific format and, therefore, inaccessible directly by the client.

The client creates a client-generated public object by using normal programming language constructs; for example, static initialization. The client is responsible for managing any storage involved. The service creates service-generated public objects

and allocates the necessary storage. As previously mentioned, the client must destroy service-generated public objects and release the storage by applying the XOM function **om_delete()** to it, as shown in the following code fragment:

```
/* We can now safely release all the private objects
 * and the public objects we no longer need
 */
CHECK_OM_CALL(om_delete(session));
CHECK_OM_CALL(om_delete(result));
CHECK_OM_CALL(om_delete(entry));
CHECK_OM_CALL(om_delete(attributes));
CHECK_DS_CALL(ds_shutdown(workspace));
```

The service also creates private objects for which it allocates storage that must be managed by the application.

One of the input parameters to the **ds_read()** function call is *name*. The *name* parameter is a public object created by the application from a series of nested data structures (RDNs and AVAs) to represent the distinguished name containing **Peter Piper**. When the application no longer needs the public object, it issues the XDS function call **ds_shutdown()** to release the memory resources associated with the public object. The **ds_read()** call returns the pointer to a private object, **result**, deposited in the workspace by the service.

The program goes on to use the XOM function **om_get()** with the input parameter *result* as a pointer to extract attribute values from the returned private object. The **om_get()** call returns the pointer *entry* as a service-generated public object to the program so that the attribute values specified in the call can be accessed by it. Once the value is extracted, the application can continue processing; for example, printing a message to a user with some extracted value like a phone number or postal address. The service-generated public object becomes the responsibility of the application program. The program goes on to release the resources allocated by the service by issuing a series of calls to **om_delete()**, as shown in the following code fragment from **example.h**:

```
/*
 * extract the telephone number(s) of "name" from the result
 *
 * There are 4 stages:
 * (1) get the Entry-Information from the Read-Result.
```

```
* (2) get the Attributes from the Entry-Information.
* (3) get the list of phone numbers.
* (4) scan the list and print each number.
*/

CHECK_OM_CALL(    om_get(result,
                        OM_EXCLUDE_ALL_BUT_THESE_TYPES
+ OM_EXCLUDE_SUBOBJECTS,
                        entry_list, OM_FALSE, 0, 0, &entry,
                        &total_num));

CHECK_OM_CALL(    om_get(entry->value.object.object,
                        OM_EXCLUDE_ALL_BUT_THESE_TYPES
+ OM_EXCLUDE_SUBOBJECTS,
                        attributes_list, OM_FALSE, 0, 0, &attributes,
                        &total_num));

CHECK_OM_CALL(    om_get(attributes->value.object.object,
                        OM_EXCLUDE_ALL_BUT_THESE_TYPES
+ OM_EXCLUDE_SUBOBJECTS,
                        telephone_list, OM_FALSE, 0, 0, &telephones,
                        &total_num));

/* We can now safely release all the private objects
 * and the public objects we no longer need
 */
CHECK_OM_CALL(om_delete(session));
CHECK_OM_CALL(om_delete(result));
CHECK_OM_CALL(om_delete(entry));
CHECK_OM_CALL(om_delete(attributes));
```

If the client possesses a valid handle (or pointer) for an object, it has access to a private object. If the client does not possess an object handle or the handle is not a valid one, a private object is inaccessible to the client and an error is returned to the calling function. In the preceding code fragment, the handles for the objects stored in *entry*, *attributes*, and *telephones* are the pointers *&entry*, *&attributes*, and *&telephones*, respectively.

5.5 OM Syntaxes for Attribute Values

An OM attribute is made up of an integer uniquely defined within a package that indicates the OM attribute's type, an integer giving that value's syntax, and an information item called a *value*. The syntaxes defined by the XOM API standard are closely aligned with ASN.1 types and type constructors.

Some syntaxes are described in the standard in terms of syntax templates.

A syntax template defines a group of related syntaxes. The syntax templates that are defined are as follows:

- Enum(*)
- Object(*)
- String(*)

5.5.1 Enumerated Types

An OM attribute with syntax template Enum(*) is an enumerated type (**OM_S_ENUMERATION**) and has a set of values associated with that OM attribute. For example, one of the OM attributes of the OM class **DS_C_ENTRY_INFO_SELECTION** is **DS_INFO_TYPE**. **DS_INFO_TYPE** is listed in the OM attribute table for **DS_C_ENTRY_INFO_SELECTION** in Chapter 11 as having a value syntax of Enum(**DS_Information_Type**), as shown in Table 5-4. **DS_INFO_TYPE** takes one of the following values:

- **DS_TYPES_ONLY**
- **DS_TYPES_AND_VALUES**

Table 5–4. Description of an OM Attribute By Using Syntax Enum(*)

OM Attributes of DS_C_ENTRY_INFO_SELECTION				
OM Attribute	Value Syntax	Value Length	Value Number	Value Initially
DS_ALL_ATTRIBUTES	OM_S_BOOLEAN	—	1	OM_TRUE
DS_ATTRIBUTES_SELECTED	String(OM_S_OBJECT_IDENTIFIER_STRING)	—	0 or more	—
DS_INFO_TYPE	Enum(DS_Information_Type)	—	1	DS_TYPES_AND_VALUES

The C language representation of the syntax of the OM attribute type **DS_INFO_TYPE** is **OM_S_ENUMERATION** as defined in the **xom.h** header file. The value of the OM attribute is either **DS_TYPES_ONLY** or **DS_TYPES_AND_VALUES**, as shown in the following code fragment from **example.h**:

```
/*
 * Public Object ("Descriptor List") for
 * Entry-Information-Selection
 * parameter to ds_read().
 */
OM_descriptor selection[] = {
    OM_OID_DESC(OM_CLASS, DS_C_ENTRY_INFO_SELECTION),
    { DS_ALL_ATTRIBUTES, OM_S_BOOLEAN, { OM_FALSE, NULL } },
    OM_OID_DESC(DS_ATTRIBUTES_SELECTED, DS_A_PHONE_NBR),
    { DS_INFO_TYPE, OM_S_ENUMERATION,
      { DS_TYPES_AND_VALUES, NULL } },
    OM_NULL_DESCRIPTOR
};
```

5.5.2 Object Types

An OM attribute with syntax template Object(*) has **OM_S_OBJECT** as syntax and a subobject as a value. For example, one of the OM attributes of the OM class **DS_C_DS_DN** is **DS_RDNS**. **DS_RDNS** is listed in the OM attribute table for **DS_C_DS_DN** as having a value syntax of Object(**DS_C_DS_RDN**), as shown in Table 5-5.

Table 5-5. Description of an OM Attribute By Using Syntax Object(*)

OM Attributes of DS_C_DS_DN				
OM Attribute	Value Syntax	Value Length	Value Number	Value Initially
DS_RDNS	Object(DS_C_DS_RDN)	—	0 or more	—

The C language representation of the syntax of the OM attribute type **DS_RDNS** is **OM_S_OBJECT**, as shown in following code fragment from **example.h**:

```
OM_descriptor      name[] = {
    OM_OID_DESC(OM_CLASS, DS_C_DS_DN),
    { DS_RDNS, OM_S_OBJECT, { 0, rdn1 } },
    { DS_RDNS, OM_S_OBJECT, { 0, rdn2 } },
    { DS_RDNS, OM_S_OBJECT, { 0, rdn3 } },
    { DS_RDNS, OM_S_OBJECT, { 0, rdn4 } },
    OM_NULL_DESCRIPTOR
};
```

5.5.3 Strings

An OM attribute with syntax template String(*) specifies the string syntax of its value. A string is categorized as either a *bit string*, an *octet string*, or a *character string*. The bits of a bit string, the octets of an octet string, or the octets of a character string constitute the *elements* of the string. (Refer to Chapter 17 for a list of the syntaxes that form the string group.)

The value length of a string is the number of elements in the string. Any constraints on the value length of a string are specified in the appropriate OM class definitions.

The elements of the string are numbered. The position of the first element is 0 (zero). The positions of successive elements are successive positive integers.

For example, one of the attributes of the OM class **DS_C_ENTRY_INFO_SELECTION** is **DS_ATTRIBUTES_SELECTED**. **DS_ATTRIBUTES_SELECTED** is listed in the OM attribute table for **DS_C_ENTRY_INFO_SELECTION** as having a value syntax of String(**OM_S_OBJECT_IDENTIFIER_STRING**), as shown in Table 5-4.

5.5.4 Other Syntaxes

The other syntaxes are defined as follows:

- **OM_S_BOOLEAN**

A value of this syntax is a Boolean; that is, the value can be **OM_TRUE** or **OM_FALSE**.

- **OM_S_INTEGER**

A value of this syntax is a positive or negative integer.

- **OM_S_NULL**

The one value of this syntax is a valueless placeholder.

5.6 Service Interface Data Types

The local variables within an application program that contain the parameters and results of XDS and XOM API function calls are declared by using a standard set of data types. These data types are defined by **typedef** statements in the **xom.h** header files. Some of the more commonly used data types are described in the following subsections. A complete description of service interface data types is provided in Chapter 18 and in the *OSF DCE Application Development Reference*.

5.6.1 The OM_descriptor Data Type

The **OM_descriptor** data type is used to describe an OM attribute type and value. A data value of this type is a descriptor, which embodies an OM attribute value. An array of descriptors can represent all the values of an object.

OM_descriptor is defined in the **xom.h** header file as follows:

```
/* Descriptor */

typedef struct OM_descriptor_struct {
    OM_type          type;
    OM_syntax        syntax;
    union OM_value_union value;
} OM_descriptor;
```

OM_descriptor is made up of a series of nested data structures, as shown in Figure 5-12.

Figure 5–12. Data Type OM_descriptor_struct

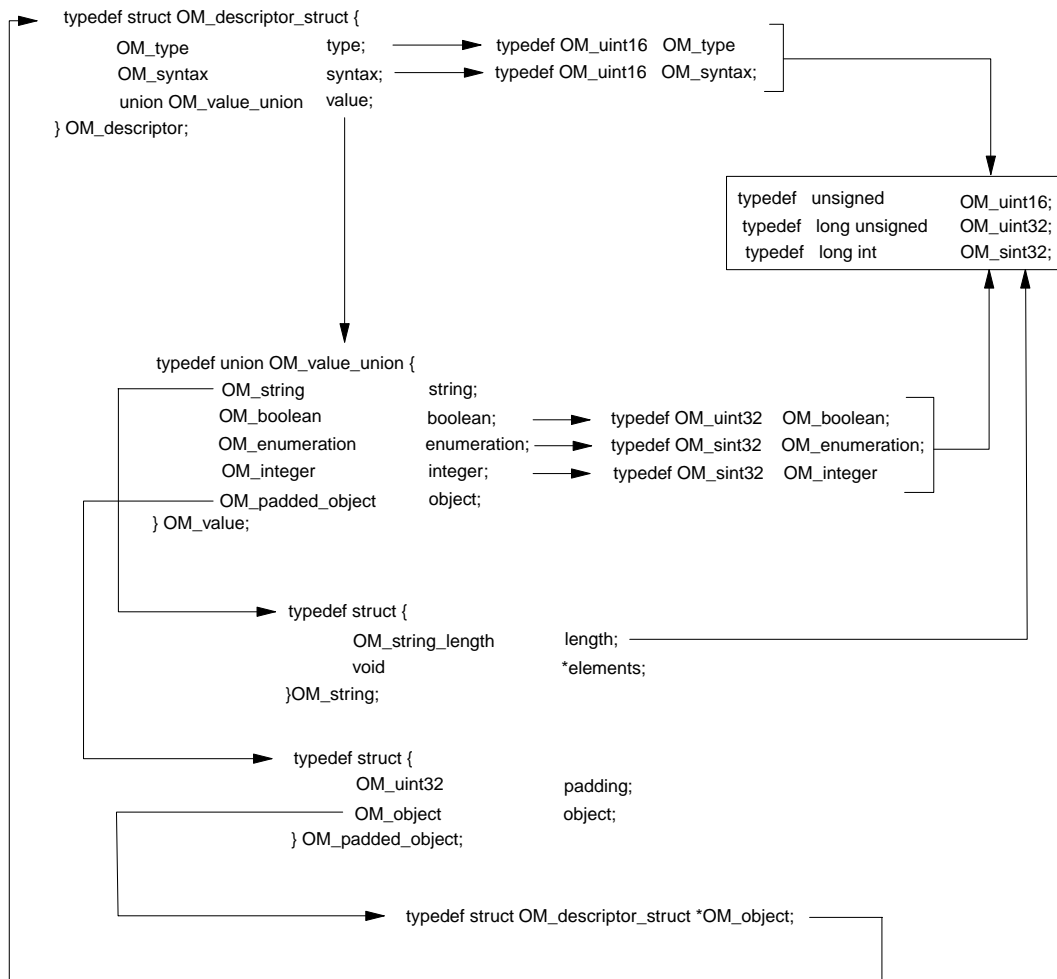


Figure 5-12 shows that **type** and **syntax** are integer constants for an OM attribute type and syntax, as shown in the following code fragment from **example.c**:

```
static OM_descriptor    country[] = {
    OM_OID_DESC(OM_CLASS, DS_C_AVA),
    OM_OID_DESC(DS_ATTRIBUTE_TYPE, DS_A_COUNTRY_NAME),
```

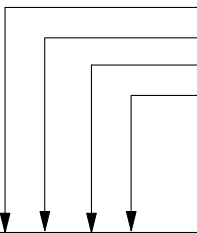
```

{ DS_ATTRIBUTE_VALUES, OM_S_PRINTABLE_STRING, OM_STRING( "US" ) },
OM_NULL_DESCRIPTOR
};

```

The code fragment initializes four descriptors, as shown in Figure 5-13. The type and syntax evaluate to integers for all four descriptors.

Figure 5–13. Initializing Descriptors



```

static OM_descriptor
OM_OID_DESC(OM_CLASS, DS_C_AVA),
OM_OID_DESC(DS_ATTRIBUTE_TYPE, DS_A_COUNTRY_NAME),
{ DS_ATTRIBUTE_VALUES, OM_S_PRINTABLE_STRING, OM_STRING("US") },
OM_NULL_DESCRIPTOR
};
country[] = {

```

Type	Syntax	Value
OM_CLASS = 3	OM_S_OBJECT_IDENTIFIER_STRING = 6	9, DS_C_AVA = \x2B\x0C\x02\x87\x73\x1C\x00\x85\x44
DS_ATTRIBUTE_TYPE = 711	OM_S_OBJECT_IDENTIFIER_STRING = 6	3, DS_A_COUNTRY_NAME = \x55\x04\x06
DS_ATTRIBUTE_VALUES = 713	OM_S_PRINTABLE_STRING = 19	2, "US"
OM_NO_MORE_TYPES = 0	OM_S_NO_MORE_SYNTAXES = 0	0, OM_ELEMENTS_UNSPECIFIED = 0

The **value** component is a little more complex. Figure 5-12 shows that **value** is a union of **OM_value_union**. **OM_value_union** has five members: **string**, **boolean**, **enumeration**, **integer**, and **object**. The members **boolean**, **enumeration**, and **integer** have integer values. The **string** member contains a string of type **OM_string**, which is a structure composed of a length and a pointer to a string of characters. The **object** member is a structure of type **OM_padded_object** that points to another object nested below it. Many OM attributes have other objects as values. These subobjects, in turn, may have other subobjects and so on.

For example, as shown in Figure 5-14, the OM class **DS_C_READ_RESULT** has one OM attribute: **DS_ENTRY**. The syntax of **DS_ENTRY** is **OM_S_OBJECT** with a value of **DS_C_ENTRY_INFO**, indicating that it points to the subobject **DS_C_ENTRY_INFO**. **DS_C_ENTRY_INFO** has the OM attribute

DS_OBJECT_NAME with the syntax **OM_S_OBJECT**, indicating that it points to the subobject **DS_C_NAME**.

Figure 5–14. An Object and a Subordinate Object

OM Class	Attribute	Syntax and Value
DS_C_READ_RESULT	DS_ENTRY	Object(DS_C_ENTRY_INFO)
DS_C_ENTRY_INFO	DS_FROM_ENTRY DS_OBJECT_NAME	OM_S_BOOLEAN Object(DS_C_NAME)

5.6.2 Data Types for XDS API Function Calls

The following code fragment from **example.h** shows how the data types are used to declare the variables that contain the output parameters from the XDS API function calls.

```
int main(void)
{
    DS_status      error;      /* return value from DS functions */
    OM_return_code  return_code; /* return value from OM functions */
    OM_workspace    workspace; /* workspace for objects */
    OM_private_object session; /* session for directory operations */
    OM_private_object result; /* result of read operation */
    OM_sint         invoke_id; /* Invoke-ID of the read operation */
}
```

The code fragment shows the following:

- The **ds_initialize()** call returns a variable of type **OM_workspace** that contains a handle or pointer to a workspace.
- The **ds_bind()** call returns a pointer to a variable of type **OM_private_object**. The private object contains the session information required by all subsequent XDS API calls, except **ds_shutdown()**.

- The **ds_read()** call returns a pointer to the result of a directory read request in a variable of type **OM_private_object**.
- The error handling macros **CHECK_DS_CALL** and **CHECK_OM_CALL**, defined in the **example.h** header file, use the data types **DS_status** and **OM_return_code**, respectively, as return values from XDS and XOM API function calls.

5.6.3 Data Types for XOM API Calls

The following code fragment from **example.h** shows how the data types are used to declare the variables that contain the input and output parameters for the XOM API function calls.

```
/*
 * variables to extract the telephone number(s)
 */
OM_type          entry_list[]      = { DS_ENTRY, 0 };
OM_type          attributes_list[] = { DS_ATTRIBUTES, 0 };
OM_type          telephone_list[]  = { DS_ATTRIBUTE_VALUES, 0 };
OM_public_object entry;
OM_public_object attributes;
OM_public_object telephones;
OM_descriptor    *telephone; /* current phone number */
OM_value_position total_num; /* number of Attribute Descriptors */
```

The code fragment shows the following:

- The series of **om_get()** calls requires a list of OM attribute types that identifies the types of OM attributes to be included in the operation. The variables **entry_list**, **attribute_list**, and **telephone_list** are declared as type **OM_type**.
- The series of **om_get()** calls return pointers to variables of type **OM_public_object**. The **om_get()** call generates public objects that are accessible to the application program.
- Where the variable **total_num** is type **OM_value_position** and is used to hold the number of OM descriptors returned by **om_get()**.

Chapter 17 contains detailed descriptions of all the data types defined by XOM API.

5.7 OM Function Calls

XOM API supports general-purpose OM functions defined by the X/Open standards body that allow an application program to manipulate objects in a workspace. Section 5.7.1 lists the OM function calls and gives a brief description of each. Section 5.7.2 illustrates the use of OM function calls by using the **om_get()** call as an example.

5.7.1 Summary of OM Function Calls

The following list of XOM API function calls contains a brief description of each function. Refer to the appropriate reference page in the *OSF DCE Application Development Reference* for a detailed description of the input and output parameters, return codes, and usage of each function.

- **om_copy()**

Creates an independent copy of an existing private object and all of its subobjects in a specified workspace.

- **om_copy_value()**

Replaces an existing OM attribute value or inserts a new value into a target private object with a copy of an existing OM attribute value found in a source private object.

- **om_create()**

Creates a private object that is an instance of the specified OM class.

- **om_delete()**

Deletes a private or service-generated public object.

- **om_get()**

Creates a new public object that is an exact, but independent, copy of an existing private object; certain exclusions and/or syntax conversion may be requested for the copy.

- **om_instance()**

Tests to determine if an object is an instance of a specified OM class (includes the case when the object is a subclass of that OM class).

- **om_put()**

Places or replaces copies of the attribute values of the source private or public object into the target private object.

- **om_read()**

Reads a segment of a string attribute from a private object.

- **om_remove()**

Removes and discards values of an attribute of a private object.

- **om_write()**

Writes a segment of a string attribute to a private object.

- **om_encode()**

Not supported by DCE XOM API.

- **om_decode()**

Not supported by DCE XOM API.

5.7.2 Using the OM Function Calls

Most application programs require the use of a series of **om_get()** function calls to create service-generated public objects from which the program can extract requested information. For this reason, this section uses the operation of **om_get()** as an example to describe how XOM API functions operate in general.

The following code fragment from **example.h** shows how a series of **om_get()** function calls extract a list of telephone numbers from a workspace. The **ds_read()** function call deposits the private object stored in **result** in the workspace and provides access to it by the pointer **&result**.

```
/*
 * extract the telephone number(s) of "name" from the result
 *
 * There are 4 stages:
 * (1) get the Entry-Information from the Read-Result.
 * (2) get the Attributes from the Entry-Information.
 * (3) get the list of phone numbers.
 * (4) scan the list and print each number.
 */

CHECK_OM_CALL( om_get(result,
    OM_EXCLUDE_ALL_BUT_THESE_TYPES
    + OM_EXCLUDE_SUBOBJECTS,
    entry_list, OM_FALSE, 0, 0, &entry,
    &total_num));

CHECK_OM_CALL( om_get(entry->value.object.object,
    OM_EXCLUDE_ALL_BUT_THESE_TYPES
    + OM_EXCLUDE_SUBOBJECTS,
    attributes_list, OM_FALSE, 0, 0, &attributes,
    &total_num));

CHECK_OM_CALL( om_get(attributes->value.object.object,
    OM_EXCLUDE_ALL_BUT_THESE_TYPES
    + OM_EXCLUDE_SUBOBJECTS,
    telephone_list, OM_FALSE, 0, 0, &telephones,
    &total_num));

/* We can now safely release all the private objects
 * and the public objects we no longer need
 */
CHECK_OM_CALL(om_delete(session));
CHECK_OM_CALL(om_delete(result));
CHECK_OM_CALL(om_delete(entry));
CHECK_OM_CALL(om_delete(attributes));
```

```

CHECK_DS_CALL(ds_shutdown(workspace));

for (telephone = telephones;
     telephone->type != DS_ATTRIBUTE_VALUES;
     telephone++)
{
    if (telephone->type != DS_ATTRIBUTE_VALUES
        || (telephone->syntax & OM_S_SYNTAX) !=
            OM_S_PRINTABLE_STRING)
    {
        (void) fprintf(stderr, "malformed telephone number\n");
        exit(EXIT_FAILURE);
    }

    (void) printf("Telephone number: %s\n",
                  telephone->value.string.elements);
}

CHECK_OM_CALL(om_delete(telephones));

```

The **om_get()** call makes a copy of all or a selected set of parts of a private object. The copy is a service-generated public object that is accessible to the application program. The application program extracts the list of telephone numbers from this copy.

5.7.2.1 Required Input Parameters

The **om_get()** function requires the following input parameters:

- A private object
- A set of exclusions
- A set of OM attributes to be included in the copy
- A flag to indicate whether local string processing is required
- The position of the first value to be copied (the base value)
- The position within each OM attribute that is one beyond the last attribute to be included in the copy (indicating the scope of the copy)

The **om_get()** call returns the following output parameters:

- The public object that is a copy of the private object
- The number of OM attribute descriptors returned in the public object

In the code fragment from **example.h**, the private object **result** is input to **om_get()**.

The next parameter, the *exclusions* parameter, reduces the copy to a prescribed portion of the original. The exclusions apply to the OM attributes of the object, but not to those of subobjects. The possibilities for determining the combinations of types, values, subobjects, and descriptors to be excluded depend on the creativity of the programmer. For a detailed description of all the exclusion possibilities, refer to the *OSF DCE Application Development Reference*. The values chosen for the **om_get()** calls in **example.h** are simplified for clarity. These exclusion values are as follows:

- **OM_EXCLUDE_ALL_BUT_THESE_TYPES**
- **OM_EXCLUDE_SUBOBJECTS**

Each value indicates an exclusion, as defined by **om_get()**, and is chosen from the set of exclusions; alternatively, the single value **OM_NO_EXCLUSIONS** may be chosen, which selects the entire object. Each value, except **OM_NO_EXCLUSIONS**, is represented by a distinct bit, the presence of the value being represented as 1, and its absence as 0 (zero). Multiple exclusions are requested by adding or ORing the values that indicate the individual exclusions.

OM_EXCLUDE_ALL_THESE_TYPES indicates that the OM attributes included are only the ones defined in the list of included types supplied in the next parameter, *entry_list*. **OM_EXCLUDE_SUBOBJECTS** indicates that, for each value whose syntax is **OM_S_OBJECT**, a descriptor containing an object handle for the original private subobject is returned, rather than a public copy of it. This handle makes that subobject accessible for use in subsequent function calls. Exclusion provides a means to examine an object one level at a time. The object the handle points to is used in the next **om_get()** call to get the next level.

The *entry_list* parameter is declared in **example.h** as data type **OM_type** and initialized as a two-cell array with values **DS_ENTRY** and a NULL terminator. **DS_ENTRY** specifies the single OM attribute type included for that **om_get()** call. This call only limits processing to the one directory entry; only one entry was defined previously in the program — the distinguished name of **Peter Piper**. The 0 (zero) marks the end of the OM attribute list.

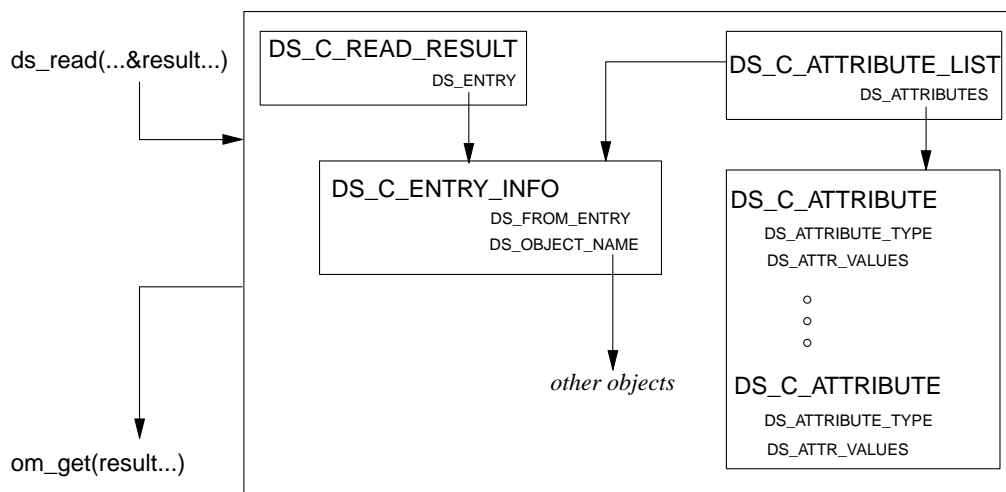
The next parameter, **OM_FALSE**, indicates that mapping to a local string format is not required. The next two parameters set the initial and limiting value to 0 (zero), meaning that no specific values are to be excluded.

The final two parameters are output parameters: *entry*, a pointer to a service-generated public object deposited by **om_get()** in the workspace, and *total_num*, an integer. Both *entry* and *total_num* are available for examination by the application program.

5.7.2.2 Extracting the Data from the Read Result

The *entry* parameter contains the result of processing by **om_get()** of the **read** parameter generated by the **ds_read()** operation. A successful call to **ds_read()** returns an instance of OM class **DS_C_READ_RESULT** in the private object *result*. **DS_C_READ_RESULT** contains the information extracted from the directory entry of the target object. Figure 5-15 shows the relationship of some of the superclasses, subclasses, and the OM attribute of **DS_C_READ_RESULT**. Consider Figure 5-15 as a partial map of the contents of *result*.

Figure 5-15. The Read Result

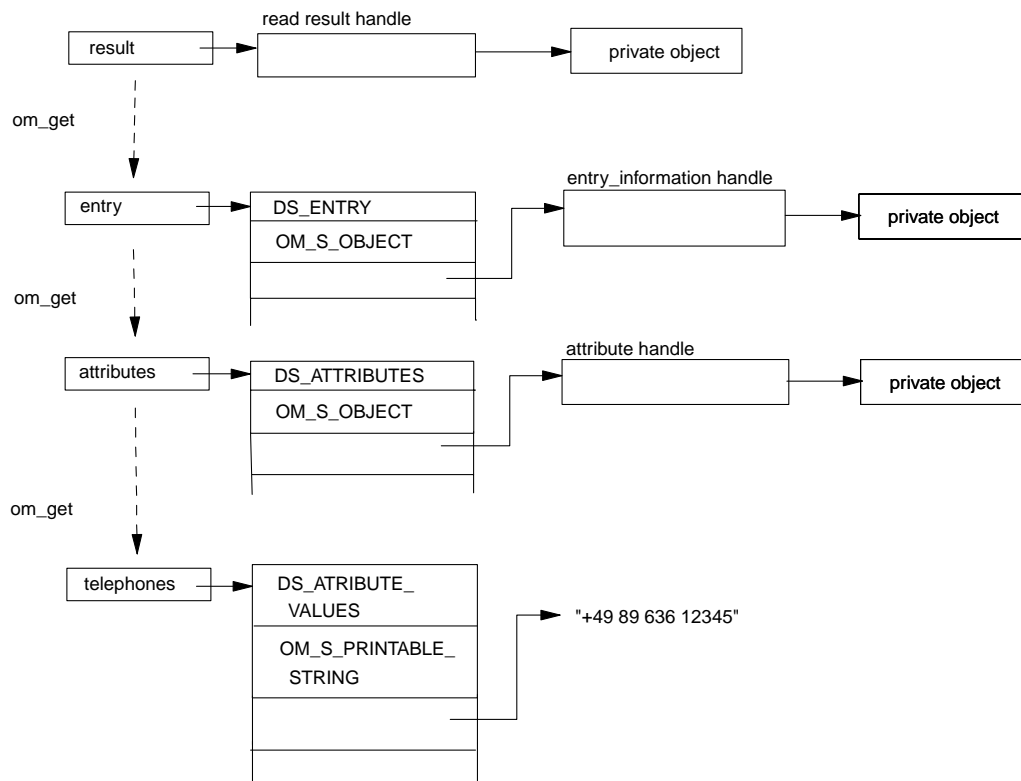


The **om_get()** function call creates a public object to make the information contained in *result* available to the application program. The *entry* parameter is defined as

data type **OM_public_object**. As such, it is composed of several nested layers of subobjects that contain entry information, OM attributes, and OM attribute values, as shown in Figure 5-16. The series of **om_get()** calls removes these layers of objects to extract a list of telephone numbers.

Figure 5-16 also shows that the process of exposing the subobjects continues while the syntax of the subobjects is **OM_S_OBJECT**. In effect, **example.h** is reversing the process of building up a series of public objects as input to **ds_read()**; namely, the distinguished name of **Peter Piper** and the descriptor list for *entry_information_selection*.

Figure 5–16. Extracting Information Using **om_get()**



The following code fragment from **example.c** shows how the syntax of the variable **telephones** is tested for valid syntax; in this case, **OM_S_PRINTABLE_STRING**:

```
for (telephone = telephones;
    telephone->type != DS_ATTRIBUTE_VALUES;
    telephone++)
{
    if (telephone->type != DS_ATTRIBUTE_VALUES ||
        (telephone->syntax & OM_S_SYNTAX) !=
        OM_S_PRINTABLE_STRING)
    {
        (void) fprintf(stderr, "malformed telephone number\n");
        exit(EXIT_FAILURE);
    }
    (void) printf("Telephone number: %s\n",
        telephone->value.string.elements);
}
```

The preceding example determines whether **telephones** is in a format that can be used by the application program as string data that can be printed out, and that the syntax is correct for a list of telephone numbers. Note that the program uses the constant **OM_S_SYNTAX** to mask off the top 6 bits. These bits are special bits that are used by XOM API. (Refer to Chapter 18 for more information on these special bits.)

5.7.2.3 Return Codes

XOM API function calls return a value of type **OM_return_code**, which indicates whether the function succeeded. If the function is successful, the value of **OM_return_code** is set to **OM_SUCCESS**. If the function fails, it returns one of the values listed in Chapter 18. The constants for **OM_return_code** are defined in the **xom.h** header file.

5.8 XOM API Header Files

The XOM API includes the header file **xom.h**. This header file is composed of declarations defining the C workspace interface. It supplies type definitions, symbolic constant definitions, and macro definitions.

5.8.1 XOM Type Definitions and Symbolic Constant Definitions

The **xom.h** header file includes **typedef** statements that define the data types of all OM objects used in the interface. It also provides definitions of symbolic constants used by the interface.

Refer to the **xom.h(4xom)** reference page for more information.

5.8.2 XOM API Macros

XOM API provides several macros that are useful in defining public objects in your application programs. These macros are defined in the **xom.h** header file.

- **OM_IMPORT**

Makes object identifier symbolic constants available within a C source module.

- **OM_EXPORT**

Allocates memory and initializes object identifier symbolic constants within a C source module.

- **OM_OID_DESC**

Initializes the type, syntax, and value of an OM attribute that holds an object identifier.

- **OM_NULL_DESCRIPTOR**

Marks the end of a client-generated public object.

- **OMP_LENGTH**

Calculates the length of an object identifier.

- **OM_STRING**

Creates a data value of a string data type.

5.8.2.1 The **OM_EXPORT** and **OM_IMPORT** Macros

Most application programs find it convenient to export all the names they use from the same C source module. **OM_EXPORT** allocates memory for the constants that represent an object OM class or an object identifier, as shown in the following code fragment from **example.c**:

```
/* Define necessary Object Identifier constants
 */
OM_EXPORT(DS_A_COMMON_NAME)
OM_EXPORT(DS_A_COUNTRY_NAME)
OM_EXPORT(DS_A_ORG_NAME)
OM_EXPORT(DS_A_ORG_UNIT_NAME)
OM_EXPORT(DS_A_PHONE_NBR)
OM_EXPORT(DS_C_AVA)
OM_EXPORT(DS_C_DS_DN)
OM_EXPORT(DS_C_DS_RDN)
OM_EXPORT(DS_C_ENTRY_INFO_SELECTION)
```

In this code fragment, object identifier constants that represent OM classes defined in the **xds.h** and **xdsbdc.h** header files are exported to the main program module. The object identifier constants are defined in **xds.h**, with the **OMP_O** prefix followed by the variable name for the object identifier. The constant itself provides the hexadecimal value of the object identifier string.

The **OM_EXPORT** macro takes the OM class name as input and creates two new data structures: a character string and a structure of type **OM_string**. The structure of type **OM_string** contains a length and a pointer to a string that may be used later in an application program by the **OM_OID_DESC** macro to initialize the value of an object identifier.

OM_IMPORT marks the identifiers as external for the compiler. It is used if **OM_EXPORT** is called in a different file from where its values are referenced.

OM_IMPORT is not used in **example.c** because **OM_EXPORT** is called in the file where the object identifiers are referenced.

5.8.2.2 The **OM_OID_DESC** and **OMP_LENGTH** Macros

The **OM_OID_DESC** macro initializes the type, syntax, and value of an OM attribute that holds an object identifier; in other words, it initializes **OM_descriptor**. It takes as input an OM attribute type and the name of an object identifier. The object identifier should have already been exported to the program module, as shown in the previous section.

The output of the macro is an **OM_descriptor** composed of a type, syntax, and value. The type is the name of the OM class. The syntax is **OM_S_OBJECT_IDENTIFIER**. The value is a two-member structure with the length of the object identifier and a pointer to the actual object identifier string. It is defined as a pointer to **void** so that it can be used as a generic pointer; it can point to any data type.

OM_OID_DESC calls **OMP_LENGTH** to calculate the length of the object identifier string.

The following code fragment from **xom.h** shows the **OM_OID_DESC** and **OMP_LENGTH** macros:

```
/* Private macro to calculate length
 * of an object identifier
 */
#define OMP_LENGTH(oid_string) (sizeof(OMP_O_##oid_string)-1)

/* Macro to initialize the syntax and value
 * of an object identifier
 */
#define OM_OID_DESC(type, oid_name)
    { (type), OM_S_OBJECT_IDENTIFIER_STRING,
      { OMP_LENGTH(oid_name) , OMP_D_##oid_name } }
```

5.8.2.3 The **OM_NULL_DESCRIPTOR** Macro

The **OM_NULL_DESCRIPTOR** macro marks the end of a client-generated public object by setting the type, syntax, and value to **OM_NO_MORE_TYPES**, **OM_S_NO_MORE_SYNTAXES**, and a value of zero length and a NULL string, respectively.

5.8.2.4 The **OM_STRING** Macro

The **OM_STRING** macro creates a string data value. Data strings are of type **OM_string**, as shown in this code fragment from the **xom.h** header file:

```
/* String */

typedef struct {
    OM_string_length    length;
    void                *elements;
} OM_string;

#define OM_STRING(string) \
    { (OM_string_length)(sizeof(string)-1), string }
```

A string is specified in terms of its length or whether or not it terminates with a NULL. **OM_string_length** is the number of octets by which the string is represented, or it is the **OM_LENGTH_UNSPECIFIED** value if the string terminates with a NULL.

The bits of a bit string are represented as a sequence of octets. The first octet stores the number of unused bits in the last octet. The bits in the bit string, beginning with the first bit and proceeding to the trailing bit, are placed in bits 7 to 0 of the second octet. These are followed by bits 7 to 0 of the third octet, then by bits 7 to 0 of each octet in turn, followed by as many bits as are required of the final octet commencing with bit 7.

Chapter 6

XDS Programming

The XDS API defines an application programming interface to directory services in the X/Open Common Applications Environment as defined in *X/Open Portability Guide*. This interface is based on the 1988 CCITT X.500 Series of Recommendations and the ISO 9594 Standard. This joint standard is referred to from this point on simply as X.500.

This chapter describes the purpose and function of XDS API functions in a general way. Refer to the reference pages in the *OSF DCE Application Development Reference* for complete and detailed information on specific function calls.

The sections that follow describe the following types of XDS functions:

- XDS interface management functions

These functions interact with the XDS interface

- Directory connection management functions

These functions initiate, manage, and terminate connections with the directory

- Directory operation functions

These functions perform operations on a directory

Note: The DCE XDS API does not support asynchronous operations from within the same thread. If an application requires asynchronous XDS operations, then it should use multiple threads to achieve this functionality. Please refer to Chapter 8 for information on using the XDS/XOM API in a multithreaded application.

The **ds_abandon()** function is not supported in this release. A **ds_abandon()** call returns a **DS_C_ABANDON_FAILED (DS_E_TOO_LATE)** error. Refer to Chapter 10 for information on abandoning directory operations.

The following names refer to the complete XDS example programs, which can be found in Chapter 7:

- **acl.c (acl.h)**
- **example.c (example.h)**
- **teldir.c**

6.1 XDS Interface Management Functions

XDS API defines a set of functions that only interact with the XDS interface and have no counterpart in the directory standard definition:

- **ds_initialize()**
- **ds_version()**
- **ds_shutdown()**

These interface functions perform operations that involve the initialization, management, and termination of sessions with the XDS interface service.

6.1.1 The `ds_initialize()` Function Call

Every application program must first call `ds_initialize()` to establish a workspace where objects returned by the directory service are deposited. The `ds_initialize()` function must be called before any other directory interface functions are called.

The `ds_initialize()` call returns a handle (or pointer) to a workspace. The application program performs operations on OM objects in this workspace. OM objects created in this workspace can be used as input parameters to the other directory interface functions. In addition, objects returned by the directory service are deposited in the workspace.

Within the following code fragment from **example.c**, a workspace is initialized. (The declaration of the variable *workspace* and the call to `ds_initialize()` are found in different sections of the program.)

```
int main(void)
{
    DS_status      error;      /* return value from DS functions */
    OM_return_code  return_code; /* return value from OM functions */
    OM_workspace    workspace; /* workspace for objects */
    OM_private_object session; /* session for directory operations */
    OM_private_object result; /* result of read operation */
    OM_sint         invoke_id; /* Invoke-ID of the read operation */
    OM_value_position total_num; /* Number of Attribute Descriptors */

    /*
     * Perform the Directory operations:
     * (1) Initialize the directory service and get an OM workspace.
     * (2) bind a default directory session.
     * (3) read the telephone number of "name".
     * (4) terminate the directory session.
     */

    CHECK_DS_CALL((OM_object) !(workspace=ds_initialize()));
```

OM_workspace is a type definition in the **xom.h** header file defined as a pointer to **void**. A void pointer is a generic pointer that may point to any data type. The variable *workspace* is declared as data type **OM_workspace**. The return value is

assigned to the variable *workspace*, and the **CHECK_DS_CALL** macro determines if the call is successful. **CHECK_DS_CALL** is an error-handling macro that is defined in **example.h**.

The **ds_initialize()** call returns a handle to a workspace in which OM objects can be created and manipulated. Only objects created in this workspace can be used as parameters to other directory interface functions. The **ds_initialize()** call returns NULL if it fails.

6.1.2 The **ds_version()** Function Call

The **ds_version()** call negotiates features of the directory interface. These features are collected into packages that define the scope of the service. Packages define such things as object identifiers for directory and OM classes and OM attributes, enumerated types, structures, and OM object constants.

XDS API defines the following packages in separate header files as part of the XDS API software product:

- Directory service package

The directory service package contains the OM classes and OM attributes used to interact with the directory service. This package is contained in the **xds.h** header file.

- Basic directory contents package

The basic directory contents package contains OM classes and OM attributes that represent values of selected attributes and selected objects defined in the X.500 standard. This package is contained in the **xdsbdcp.h** header file.

- Strong authentication package

The strong authentication package contains OM classes and OM attributes that represent values of security attributes and objects defined in the X.500 standard. This package is contained in the **xdssap.h** header file.

- GDS package

The GDS package contains the OM classes and OM attributes that are required for GDS. This package is contained in the **xdsgds.h** header file.

- MHS directory user package

The MHS (message handling system) directory user package contains the OM classes and OM attributes that are required for electronic mail support. This package is contained in the **xdsmdup.h** header file.

The application program, which is the client, uses **ds_version()** to negotiate the scope of the services the directory service will provide to the client. A **ds_version()** function call includes a list of features (or packages) that the client wants to include as part of the interface. The features are object identifiers that represent packages supported by the DCE XDS API. The service returns a list of Boolean values to indicate whether or not the package was successfully negotiated.

These features are assigned to the workspace that an application program initialized (as described in Section 6.1.1). In addition, an application program must include the header files for the appropriate packages as part of the source code.

It is not necessary to negotiate the directory service package. It is a mandatory requirement for XDS API, and as such it is included by default. The other packages listed previously are optional and require negotiation by using **ds_version()**.

The following code fragment from **acl.h** shows how an application builds up an array of object identifiers for the optional packages to be negotiated: the basic directory contents package and the GDS package.

```
static DS_feature features[] = {
    { OM_STRING(OMP_O_DS_BASIC_DIR_CONTENTS_PKG), OM_TRUE },
    { OM_STRING(OMP_O_DSX_GDS_PKG), OM_TRUE },
    { 0 }
};
```

The **OM_STRING** macro is provided for creating a data value of data type **OM_string** for octets strings and characters. The array of object identifiers is stored in *features*, the input parameter to **ds_version()**, as shown in the following code fragment from **acl.c**:

```
/* Negotiate the use of the BDCP and GDS packages.      */  
  
if (ds_version(features,workspace) != DS_SUCCESS)  
    printf("ds_version() error\n");
```

6.1.3 The **ds_shutdown()** Function Call

The **ds_shutdown()** call deletes the workspace established by **ds_initialize()** and enables the directory service to release resources. No other directory functions that reference that workspace may be called after this function.

The following code fragment from **acl.c** demonstrates how the application closes the directory workspace by performing a **ds_shutdown()** call.

```
/* Close the directory workspace.                        */  
  
if (ds_shutdown (workspace) != DS_SUCCESS)  
    printf ("ds_shutdown() error \n");
```

6.2 Directory Connection Management Functions

The following subsections describe the XDS functions that initiate, manage, and terminate connections with the directory service.

6.2.1 A Directory Session

A directory session identifies the DSA to which a directory operation is sent. It also defines the characteristics of a session, such as the distinguished name of the requestor.

An application program can request a session with specific OM attributes tailored for the program's requirements. The application passes an instance of OM class **DC_C_SESSION** with the appropriate OM attributes, or it uses the default parameters

by passing the constant **DS_DEFAULT_SESSION** as a parameter to the **ds_bind()** function call.

6.2.2 The **ds_bind()** Function Call

The **ds_bind()** call establishes a session with the directory. The **ds_bind()** call corresponds to the **DirectoryBind** function in the Abstract Service defined in the X.500 standard.

When a **ds_bind()** call completes successfully, the directory returns a pointer to an OM private object of OM class **DC_C_SESSION**. This parameter is then passed as the first parameter to most interface function calls until a **ds_unbind()** is called to terminate the directory session.

XDS API supports multiple concurrent sessions so that an application can interact with the directory service by using several identities, and interact directly and concurrently with different parts of the directory service.

The following code fragment from **example.c** shows how an application binds to the GDS server (without credentials) by using the default session:

```
CHECK_DS_CALL(ds_bind(DS_DEFAULT_SESSION, workspace, &session));
```

If a user wants to do an authenticated bind and/or wants to specify the directory identifier, an instance of OM class **DSX_C_GDS_SESSION** from the GDS package is required. **DSX_C_GDS_SESSION** identifies a particular link from an application to a DSA. Since **DSX_C_GDS_SESSION** is a subclass of the standard OM class for a session, **DS_C_SESSION**, it may be passed as a parameter to an XDS API function, such as **ds_bind()**, wherever a standard session is expected.

The following code fragment from **acl.c** shows how an application performs an authenticated bind to the GDS:

```
/*
 * Create a default session object.
 */
if ((rc = om_create(DSX_C_GDS_SESSION, OM_TRUE, workspace, &session))
    != OM_SUCCESS)
```

```
    printf("om_create() error %d\n", rc);

/*
 * Alter the default session object to include the following
 * credentials:
 * requestor: /C=de/O=sni/OU=ap/CN=norbert
 * password: "secret"
 * authentication mechanism: simple
 */
if ((rc = om_put(session, OM_REPLACE_ALL, credentials, 0, 0, 0))
    != OM_SUCCESS)
    printf("om_put() error %d\n", rc);

/*
 * Bind with credentials to the default GDS server.
 * The returned session object is stored in the private object
 * variable bound_session and is used for all further XDS
 * function calls.
 */
if (ds_bind(session, workspace, &bound_session) != DS_SUCCESS)
    printf("ds_bind() error\n");
```

The program creates a default session object by using the XOM API function **om_create()** and alters the default session object by using **om_put()**. The bind credentials are initialized in the following code fragment from the **example.h** header file included in the main program module:

```
/* The following descriptor list specifies
 * the bind credentials
 */

static OM_descriptor credentials[] = {
    {DS_REQUESTOR, OM_S_OBJECT, {0, dn_norbert} },
    {DSX_PASSWORD, OM_S_OCTET_STRING, OM_STRING("secret")},
    {DSX_AUTH_MECHANISM, OM_S_ENUMERATION, {DSX_SIMPLE, 0}},
    OM_NULL_DESCRIPTOR
};
```

The *credentials* parameter is provided as an input parameter to the **om_put()** function call to modify the existing session object in the directory service. A private object that is used for all subsequent directory calls is returned to the workspace by **om_put()**.

6.2.3 The **ds_unbind()** Function Call

The **ds_unbind()** call terminates a directory session and makes the *session* parameter unavailable for use with other interface functions. However, the unbound session can be modified by OM functions and used again as a parameter to **ds_bind()**. When the *session* parameter is no longer needed, it should be deleted by using OM functions such as **om_delete()**.

The following code fragment from **example.c** shows how the application closes the connection to the GDS server by using **ds_unbind()**:

```
/* Close the connection to the GDS server. */

if (ds_unbind(bound_session) != DS_SUCCESS)
    printf("ds_unbind() error\n");
```

The **ds_unbind()** call corresponds to the **DirectoryUnbind** function in the Abstract Service defined in the X.500 standard.

6.2.4 Automatic Connection Management

The XDS implementation does not support automatic connection management. A DSA connection is established when **ds_bind()** is called and released when **ds_unbind()** is called.

6.3 XDS Interface Class Definitions

The XDS interface class definitions are described in detail in Chapter 11. The OM attribute types, syntax, and values and inheritance properties are described for each OM class.

A good way to begin to understand how the OM class hierarchy is structured and the relationship between OM classes and OM attributes to the service provided by the directory service package is to look up one of the OM classes listed in Chapter 11.

6.3.1 Example: The **DS_C_FILTER** Class

For example, **DS_C_FILTER** inherits the OM attributes from its superclass **OM_C_OBJECT**, as do all OM classes. **OM_C_OBJECT**, as defined in Chapter 19, has one OM attribute, **OM_CLASS**, which has the value of an object identifier string that identifies the numeric representation of the object's OM class. **DS_C_FILTER**, on the other hand, has several OM attributes.

The purpose of **DS_C_FILTER** is to select or reject an object on the basis of information in its directory entry. It has the following OM attributes:

- **DS_FILTER_ITEMS**
- **DS_FILTERS**
- **DS_FILTER_TYPE**

Two of these OM attributes, **DS_FILTER_ITEMS** and **DS_FILTERS**, have values that are OM object classes themselves. The value of the OM attribute **DS_FILTER_ITEMS** is **DS_C_FILTER_ITEM**, which is an OM class. **DS_C_FILTER_ITEM** is a component of a filter and defines the nature of the filter. The value of the OM attribute **DS_FILTERS** is **DS_C_FILTER**, an OM class. Thus, **DS_FILTERS** defines a collection of filters. The OM attribute **DS_FILTER_TYPE** has a value that is an enumerated type, which takes one of the values **DS_AND**, **DS_OR**, or **DS_NOT**.

Refer to Figure 6-3 for a description of the relationship of **DS_C_FILTER** to its superclass **OM_C_OBJECT** and its attributes.

6.3.2 The **DS_C_CONTEXT** Parameter

The OM class **DS_C_CONTEXT** is the second parameter to every directory service request. **DS_C_CONTEXT** defines the characteristics of the directory service interaction that are specific to a particular directory service operation.

These characteristics are divided into three categories of OM attributes: common parameters, service controls, and local controls.

Common parameters affect the processing of each directory service operation.

Service controls indicate how the directory service should handle requests. Included in this category are decisions about whether or not chaining is permitted, the priority of requests, the scope of referral (to DSAs within a country or within a DMD), and the maximum number of objects about which a function should return information.

Local controls include asynchronous support and automatic continuation; XDS does not currently support asynchronous operations from within the same thread. Applications requiring asynchronous use of the XDS/XOM API should use threads as defined in Chapter 8.

6.4 Directory Class Definitions

The X.500 standards define a number of attribute types and classes. These definitions allow the creation and maintenance of directory entries for a number of common objects so that the representation of all such objects is the same throughout the directory. The basic directory contents package contains OM classes and OM attributes that model the X.500 attribute types and classes.

The X.500 object classes and attributes are defined in the following documents published by CCITT. These are the objects and the associated attributes that will be the targets of directory service operations in your application programs:

- *The Directory: Selected Attributes Types (Recommendation X.520)*
- *The Directory: Selected Object Classes (Recommendation X.521)*

Table 6-1 describes the OM classes, OM attributes, and their object identifiers that model the X.500 objects and attributes. (See Chapter 12 for more tables with the same type of information.)

Table 6–1. Representation of Values for Selected Attribute Types

Attribute Type	OM Value Syntax	Value Length	Multi-valued	Matching Rules
DS_A_ALIASED_OBJECT_NAME	Object(DS_C_NAME)	—	no	E
DS_A_BUSINESS_CATEGORY	String(OM_S_TELETEX_STRING)	1–128	yes	E, S
DS_A_COMMON_NAME	String(OM_S_TELETEX_STRING)	1–64	yes	E, S
DS_A_COUNTRY_NAME	String(OM_S_PRINTABLE_STRING) ¹	2	no	E
DS_A_DESCRIPTION	String(OM_S_TELETEX_STRING)	1–1024	yes	E, S

1 As permitted by ISO 3166.

The tables in Chapter 12 contain similar categories of information as do similar tables for the attributes defined in the directory service package. These information categories include the following:

- OM Value Syntax
- Value Length
- Multivalued
- Matching Rules

The OM Value Syntax column describes the structure of the values of an OM attribute. The Value Length column gives the range of lengths permitted for the string types. The Multivalued column indicates whether the attribute can have multiple values.

The CCITT standards define matching rules that are used for determining whether two values are equal, for ordering two values, or for identifying one value as a substring of another in directory service operations. These are indicated in the Matching Rules column.

The GDS administrator maintains the directory service and determines the structure of the DIT as defined by the GDS schema. The GDS standard (or default) schema is based on the recommendations in the CCITT documents mentioned previously.

Recall that the structure rule table (SRT) of the GDS schema defines the structure of the DIT, the object class table (OCT) defines class inheritance properties, and the attribute table (AT) defines the mandatory and optional attributes for each class. You will find it useful to familiarize yourself with the existing schema when developing an application program that will access the directory. This is because the public objects that your programs will create (by using OM classes and OM attributes) are modeled after objects and attributes in the directory.

6.5 The GDS Package

The GDS software provides functional extensions to the standard in the following areas:

- Authentication
- Access control
- DUA cache

6.5.1 Authentication

An instance of OM class **DSX_C_GDS_SESSION** identifies a particular link from an application program to a DSA. This additional OM class is necessary if the user either wants to specify use of an authentication mechanism (for example, a password), or wants to specify a directory identifier.

6.5.2 Access Control

In addition to authentication (for example, by means of name and password), access protection is required for each object at the attribute level. A telephone number, for example, is an attribute that generally everybody is allowed to read. However, an attribute value such as a *userpassword* usually has restricted access. In addition,

even for attributes that everyone is allowed to read, it may only be acceptable for a small number of people to have authorization to change the values.

Because there can be a multitude of different attributes in the DIT, it is too expensive to define a protection mechanism for each individual attribute type. The directory attribute **DSX_A_ACL** is present for each entry in the DIT. Its syntax is `Object(DSX_C_GDS_ACL)`, referencing the GDS class **DSX_C_GDS_ACL**. These OM classes and attributes have been added to the directory service to specify the category of access to the individual attributes that are granted to users. There are three categories of access: public, standard, and sensitive.

DSX_C_GDS_ACL has five OM attributes that define the read and modify access rights for each of these categories (read access is granted to all users; modify access implicitly grants read access):

- **DSX_MODIFY_PUBLIC**

Specifies the user, or group of users, that can modify attributes classified as public attributes

- **DSX_READ_STANDARD**

Specifies the user, or group of users, that can read attributes classified as standard attributes

- **DSX_MODIFY_STANDARD**

Specifies the user, or group of users, that can modify attributes classified as standard attributes

- **DSX_READ_SENSITIVE**

Specifies the user, or group of users, that can read attributes classified as sensitive attributes

- **DSX_MODIFY_SENSITIVE**

Specifies the user, or group of users, that can modify attributes classified as sensitive attributes

The ACL of the default schema has no access rights when GDS is configured. Every user, including the anonymous user, has read and modify access to all attributes in the schema.

A master entry can be created only by the user who has write access to the naming attribute of the parent node. Thus, the user can create all attributes of the entry. Using the ACL class, the user can establish which objects can be accessed. If the user does not enter an ACL attribute when creating an entry, GDS automatically uses the ACL attribute of the parent node for the new entry.

A master entry can only be deleted by users who have write access to the naming attribute of the entry to be deleted.

A shadow entry created by means of shadow handling (refer to the *OSF DCE GDS Administration Guide and Reference*) has the same ACL attribute as the corresponding master entry. This entry can therefore only be modified and deleted by the user who can also modify and delete the master entry.

6.5.3 DUA Cache

To further optimize access times, frequently requested information is automatically loaded to a section of memory in the client computer, the DUA cache, and can be overwritten again if it is not used within a certain interval of time. The cache may be periodically updated. The GDS administration program specifies the period. It can also specify that certain data is never written to the cache, or that certain data that is transferred must under no circumstances be deleted, unless it is deleted by the user. Because the DUA cache is not subject to any access control, GDS ensures that only the information that everybody is allowed to read is stored.

The GDS package includes the OM class **DSX_C_GDS_CONTEXT** to support additional service controls for caching. **DSX_C_GDS_CONTEXT** is a subclass of **DS_C_CONTEXT**. As such, it inherits all the standard X.500 attributes associated with **DS_C_CONTEXT**, in addition to its own OM attributes related to caching. Refer to Chapter 4 for more information on how to manage the DUA cache by using XDS.

6.5.4 Advanced Administration Operations

GDS makes use of three operational attributes:

- **DSX_A_MASTER_KNOWLEDGE**

Contains the distinguished name of the DSA that holds the master copy of a specific entry

- **DSX_A_ACL**

Used for GDS access control

- **DS_A_USER_PASSWORD** attribute of a **DS_O_DSA** object class

Used by the GDS shadowing mechanism.

The **DSX_A_MASTER_KNOWLEDGE** and **DSX_A_ACL** attributes are present in every GDS entry. When an application requests all attributes, it may prevent any of these three attributes from being returned by setting the **DSX_PREFER_ADM_FUNCS** service control (OM class **DSX_C_GDS_CONTEXT**) to **OM_FALSE**. Certain GDS applications, such as GDS administration, may need these attributes. They can achieve this by setting this service control to **OM_TRUE**.

6.6 Directory Operation Functions

The X.500 standard defines the operations provided by the directory in a document called the *Abstract Service Definition*. DCE implements this standard with XDS API functions calls. The XDS API functions allow an application program to interact with the directory service. The standard divides these interactions into three general categories: read, search, and modify.

The XDS API functions correspond to the Abstract Service functions defined in the X.500 standard, as shown in Table 6-2.

Table 6–2. Mapping of XDS API Functions to the Abstract Services

XDS Function Call	Equivalent Abstract Service
ds_read()	Read
ds_compare()	Compare
ds_list()	List
ds_search()	Search
ds_add_entry()	AddEntry
ds_remove_entry()	RemoveEntry
ds_modify_entry()	ModifyEntry
ds_modify_rdn()	ModifyRDN

6.7 Directory Read Operations

Read functions retrieve information from specific named entries in the directory where names are mapped to attributes. This is analogous to looking up some information about a name in the “White Pages” phone directory.

XDS API implements the following read functions:

- **ds_read()**

The requestor supplies a distinguished name and one or more attribute types. The value(s) of requested attributes or just the attribute type(s) is returned by the DSA.

- **ds_compare()**

The requestor gives a distinguished name and an attribute value assertion (AVA). If the AVA is TRUE for the named entry, a value of TRUE is returned by the DSA.

For example, a typical read operation could request the telephone number of a particular employee. A read request would submit the distinguished name of the

employee with an indication to return its telephone number: `/C=us/O=sni/OU=sales/CN=John Smith`.

6.7.1 Reading an Entry from the Directory

The following sections describe a typical read operation by using the `ds_read()` function call. They include a description of tasks directly related to the read operation. They do not include service-related tasks such as initializing the interface, allocating an OM workspace, and binding to the directory. These tasks are described in Section 6.1. The following sections also do not describe the process of extracting information from the workspace by using XOM functions. Refer to Chapter 5 for a description of how to use XOM functions to access the workspace.

A typical read operation involves the following steps:

1. Define the necessary object identifier constants for the OM classes and OM attributes that will define public objects for input to `ds_read()`, by using the `OM_EXPORT` macro.
2. Declare the variables that will contain the output from the XDS functions to be used in the application.
3. Build public objects (descriptor lists) for the *name* parameter to `ds_read()`.
4. Create a descriptor list for the *selection* parameter to `ds_read()` that selects the type and scope of information in your request.
5. Perform the read operation.

These steps are demonstrated in the following code fragments from `example.c` (refer to Chapter 7 for a complete program listing). The program reads the telephone numbers of a given target entry.

6.7.2 Step 1: Export Object Identifiers for Required Directory Classes and Attributes

Most application programs find it convenient to export all the names they use from the same C source module. In the following code fragment from `example.c`, the

OM_EXPORT macro allocates memory for the constants that represent the OM object classes and directory attributes required for the read operation:

```
/* Define necessary Object Identifier constants
 */
OM_EXPORT(DS_A_COMMON_NAME)
OM_EXPORT(DS_A_COUNTRY_NAME)
OM_EXPORT(DS_A_ORG_NAME)
OM_EXPORT(DS_A_ORG_UNIT_NAME)
OM_EXPORT(DS_A_PHONE_NBR)
OM_EXPORT(DS_C_AVA)
OM_EXPORT(DS_C_DS_DN)
OM_EXPORT(DS_C_DS_RDN)
OM_EXPORT(DS_C_ENTRY_INFO_SELECTION)
```

The **OM_EXPORT** macro performs the following steps:

1. It defines a character array called **OMP_D_** concatenated with the *class_name* input parameter.
2. It initializes this array to the value of a character string called **OMP_O_** concatenated with the *class_name* input parameter. This value has already been defined in a header file.
3. It defines an **OM_string** data structure as the *class_name* input parameter.
4. It initializes the **OM_string** data structure's first component to the length of the array initialized in Step 2, and initializes the second component to a pointer to the value of the array initialized in Step 2.

6.7.3 Step 2: Declare Local Variables

The local variables *session*, *result*, and *invoke_id* are defined in the following code fragment from **example.c**:

```
int main(void)
{
    DS_status      error; /* return value from DS functions */
    OM_return_code return_code; /* return value from OM functions */
    OM_workspace   workspace; /* workspace for objects */
}
```

```
OM_private_object session;    /* session for directory operations*/
OM_private_object result;    /* result of read operation      */
OM_sint            invoke_id; /* Invoke-ID of the read operation */
OM_value_position total_num; /* Number of Attribute Descriptors */
```

These data types are defined in a **typedef** statement in the **xom.h** header file. The *session* and *result* variables are defined as data type **OM_private_object** because they are returned by **ds_bind()** and **ds_read()**, respectively, to the workspace as private objects. Since asynchronous operations (within the same thread) are not supported, the *invoke_id* functionality is redundant. The *invoke_id* parameter must be supplied to the XDS functions as described in the *OSF DCE Application Development Reference*, but its return value should be ignored.

Values in *error* and *return_code* are returned by XOM and XDS functions to indicate whether a call was successful. The *workspace* variable is defined as **OM_workspace** and is used when establishing an OM workspace. The *total_num* variable is defined as **OM_value_position** to indicate the number of attribute descriptors returned in the public object by **om_get()**, based on the inclusion and exclusion parameters specified.

6.7.4 Step 3: Build Public Objects

A **ds_read()** function call can take a public object as an input parameter. A public object is generated by an application program and contains the information required to access a target directory object. This information includes the AVAs and RDNs that make up a distinguished name of an entry in the directory.

A public object is created by using OM classes and OM attributes. These OM classes and OM attributes model the target object entry in the directory and provide other information required by the directory service to access the directory. In this case, the target object entry in the directory is the entry for **Peter Piper**.

Chapter 5 describes how to create the required public objects for the **ds_read()** function call by using macros and data structures defined in the XDS and XOM API header files.

The purpose of building the public objects for AVAs and RDNs is to provide the public objects that represent a distinguished name. The distinguished name public

object is stored in the array of descriptors called *name* and provided as an input parameter to the **ds_read()** function call.

6.7.5 Step 4: Create an Entry-Information-Selection Parameter

The distinguished name for **Peter Piper** is an entry in the directory that the application is designed to access. The *selection* parameter of the **ds_read()** function call tailors its results to obtain just part of the required entry. Information on all attributes, no attributes, or a specific group of attributes can be chosen. Attribute types are always returned, but the attribute values need not be.

The value of the parameter is a public object (descriptor list) that is an instance of OM class **DS_C_ENTRY_INFO_SELECTION**, as shown in the following code fragment from **example.c**:

```
/*
 * Public Object ("Descriptor List") for
 * Entry-Information-Selection
 * parameter to ds_read().
 */
OM_descriptor selection[] = {
    OM_OID_DESC(OM_CLASS, DS_C_ENTRY_INFO_SELECTION),
    { DS_ALL_ATTRIBUTES, OM_S_BOOLEAN, { OM_FALSE, NULL } },
    OM_OID_DESC(DS_ATTRIBUTES_SELECTED, DS_A_PHONE_NBR),
    { DS_INFO_TYPE, OM_S_ENUMERATION,
      { DS_TYPES_AND_VALUES, NULL } },
    OM_NULL_DESCRIPTOR
};
```

DS_C_ENTRY_INFO_SELECTION is a subclass of **OM_C_OBJECT**. (This information is supplied in the description of this class in Chapter 11.) As such, **DS_C_ENTRY_INFO_SELECTION** inherits the OM attributes of **OM_C_OBJECT**. The only OM attribute of **OM_C_OBJECT** is **OM_CLASS**. **OM_CLASS** identifies an object's class, which in this case is **DS_C_ENTRY_INFO_SELECTION**. **DS_C_ENTRY_INFO_SELECTION** identifies information to be extracted from a directory entry and has the following OM attributes:

- **OM_C_CLASS** (inherited from **OM_C_OBJECT**)

- **DS_ALL_ATTRIBUTES**
- **DS_ATTRIBUTES_SELECTED**
- **DS_INFO_TYPE**

As part of a **ds_read()** or **ds_search()** function call, **DS_ALL_ATTRIBUTES** specifies to the directory service those attributes of a directory entry that are relevant to the application program. It can take the values **OM_TRUE** or **OM_FALSE**. These values are defined to be of syntax **OM_S_BOOLEAN**. The value **OM_TRUE** indicates that information is requested on all attributes in the directory entry. The value **OM_FALSE**, used in the preceding sample code fragment, indicates that information is only requested on those attributes that are listed in the OM attribute **DS_ATTRIBUTES_SELECTED**.

DS_ATTRIBUTES_SELECTED lists the types of attributes in the entry from which information is to be extracted. The syntax of the value is specified as **OM_S_OBJECT_IDENTIFIER_STRING**.

OM_S_OBJECT_IDENTIFIER_STRING contains an octet string of BER-encoded integers, which are decimal representations of object identifiers of the types of attributes in the attribute list. In the preceding code fragment, the string value is the attribute name **DS_A_PHONE_NBR** because the purpose of the read call is to read a list of telephone numbers from the directory.

DS_INFO_TYPE identifies what information is to be extracted from each attribute identified. The syntax of the value is specified as Enum(**DS_Information_Type**). **DS_INFO_TYPE** is an enumerated type that has two possible values: **DS_TYPES_ONLY** and **DS_TYPES_AND_VALUES**. **DS_TYPES_ONLY** indicates that only the attribute types of the selected attributes in the entry are returned by the directory service operation. **DS_TYPES_AND_VALUES** indicates that both the attribute types and the attribute values of the selected attributes in the entry are returned. The code fragment from **example.c** shown previously defines the value of **DS_INFO_TYPE** as **DS_TYPES_AND_VALUES** because the program wants to get the actual telephone numbers.

6.7.6 Step 5: Perform the Read Operation

The following code fragment from **example.c** shows the **ds_read()** function call and the XDS calls that precede it:

```
/*
 * Perform the Directory operations:
 * (1) Initialize the directory service
 *     and get an OM workspace.
 * (2) bind a default directory session.
 * (3) read the telephone number of "name".
 * (4) terminate the directory session.
 */

CHECK_DS_CALL((OM_object) !(workspace = ds_initialize()));

CHECK_DS_CALL(ds_version(bdcp_package, workspace));

CHECK_DS_CALL(ds_bind(DS_DEFAULT_SESSION, workspace,
                      &session));

CHECK_DS_CALL(ds_read (session, DS_DEFAULT_CONTEXT,
                      name, selection, &result, &invoke_id));
```

CHECK_DS_CALL is an error-checking macro defined in the **example.h** header file that is included by **example.c**. The **ds_read()** call returns a return code of type **DS_status** to indicate whether or not the read operation completed successfully. If the call was successful, **ds_read()** returns the value **DS_SUCCESS**. If the call fails, it returns an error code. (Refer to Chapter 11 for a comprehensive list of error codes.) **CHECK_DS_CALL** interprets this return value and returns successfully to the program or branches to an error-handling routine.

The *session* input parameter is a private object generated by **ds_bind()** prior to the **ds_read()** call, as shown in the preceding code fragment. **DS_DEFAULT_CONTEXT** describes the characteristics of a directory service interaction. Most XDS API function calls require these two input parameters because they define the operating parameters of a session with a GDS server. (Sessions are described in Section 6.2.1; contexts are described in Section 6.3.2.)

The result of a directory service request is returned in a private object (in this case, *result*) that is appropriate to the type of operation. The result of the operation is returned in a single OM object. The components of the result are represented by OM attributes in the operations result object:

- **DS_C_COMPARE_RESULT**

Returned by **ds_compare()**

- **DS_C_LIST_RESULT**

Returned by **ds_list()**

- **DS_C_READ_RESULT**

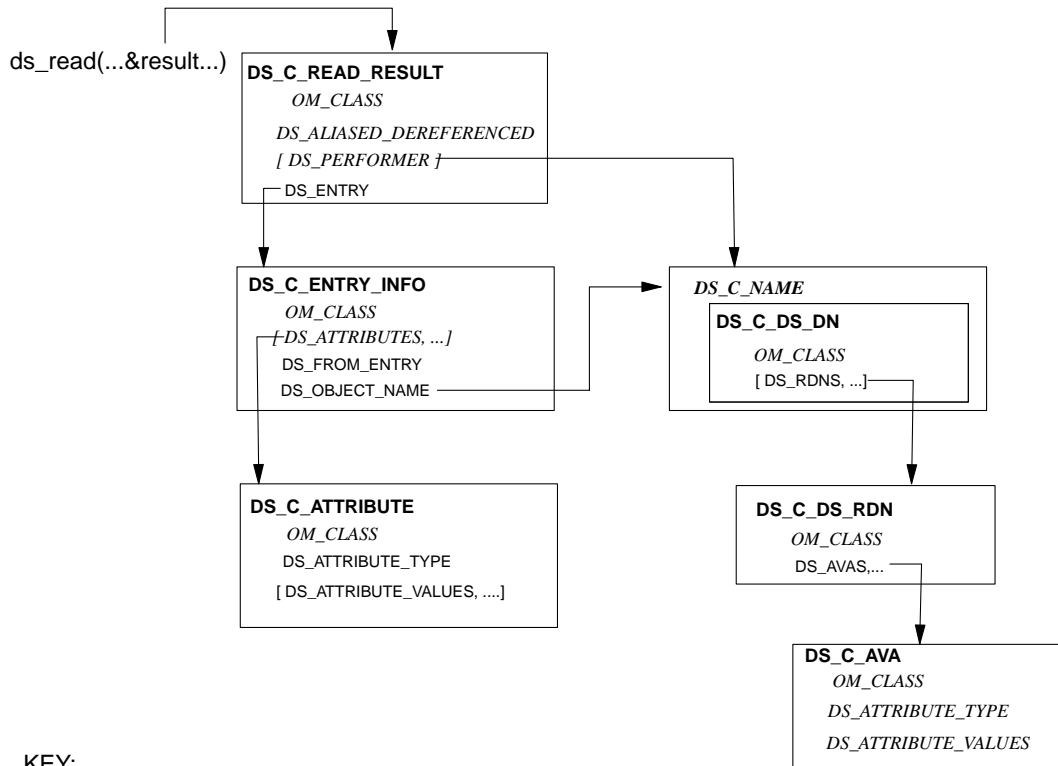
Returned by **ds_read()**

- **DS_C_SEARCH_RESULT**

Returned by **ds_search()**

The OM class returned by **ds_read()** is **DS_C_READ_RESULT**. The OM class returned by the **ds_compare()** call is **DS_C_COMPARE_RESULT**, and so on. (Refer to the reference pages in the *OSF DCE Application Development Reference* for a description of the OM classes associated with a particular function call; refer to Chapter 11 for full descriptions of the OM attributes, syntaxes, and values associated with these OM classes.)

The superclasses, subclasses, and OM attributes for **DS_C_READ_RESULT** are shown in Figure 6-1.

Figure 6–1. Output from `ds_read()`: `DS_C_READ_RESULT`

KEY:

—▶ points to subobjects

BOLD OM class

BOLD and ITALICS Abstract OM class

ITALICS inherited OM attribute

[] optional OM attribute

, ... multi-valued OM attribute

The *result* value is returned to the workspace in private implementation-specific format. As such, it cannot be read directly by an application program, but it requires a series of `om_get()` function calls to extract the requested information from it. The following code fragment from **example.c** shows how a series of `om_get()` calls extracts the list of telephone numbers associated with the distinguished name for **Peter Piper**. Chapter 5 describes this extraction process in detail.

```
/*
 * extract the telephone number(s) of "name" from the result
 *
 * There are 4 stages:
 * (1) get the Entry-Information from the Read-Result.
 * (2) get the Attributes from the Entry-Information.
 * (3) get the list of phone numbers.
 * (4) scan the list and print each number.
 */

CHECK_OM_CALL( om_get()(result,
                        OM_EXCLUDE_ALL_BUT_THESE_TYPES
                        + OM_EXCLUDE_SUBOBJECTS,
                        entry_list, OM_FALSE, 0, 0, &entry,
                        &total_num));

CHECK_OM_CALL( om_get()(entry->value.object.object,
                        OM_EXCLUDE_ALL_BUT_THESE_TYPES
                        + OM_EXCLUDE_SUBOBJECTS,
                        attributes_list, OM_FALSE, 0, 0, &attributes,
                        &total_num));

CHECK_OM_CALL( om_get()(attributes->value.object.object,
                        OM_EXCLUDE_ALL_BUT_THESE_TYPES
                        + OM_EXCLUDE_SUBOBJECTS,
                        telephone_list, OM_FALSE, 0, 0, &telephones,
                        &total_num));
```

6.8 Directory Search Operations

Search functions can be used to browse through the Directory Information Tree (DIT). For example, a search request could supply the distinguished name of an entry and request a list of the distinguished names of the children of that entry.

XDS API implements the following search operations:

- **ds_list()**

The requestor supplies a distinguished name. The directory service returns a list of the immediate subordinates of the named entry.

- **ds_search()**

The requestor supplies a search criterion known as a *filter*. The user names a subtree of the DIT, specifies some target attribute types, and formulates an expression by combining a number of attributes by using logical AND, OR, or NOT operators. The directory service returns information from all of the entries within the specified portion of the DIT that matches the filter. Section 6.8.6 includes a description of how filters are used in **acl.c**.

6.8.1 Searching the Directory

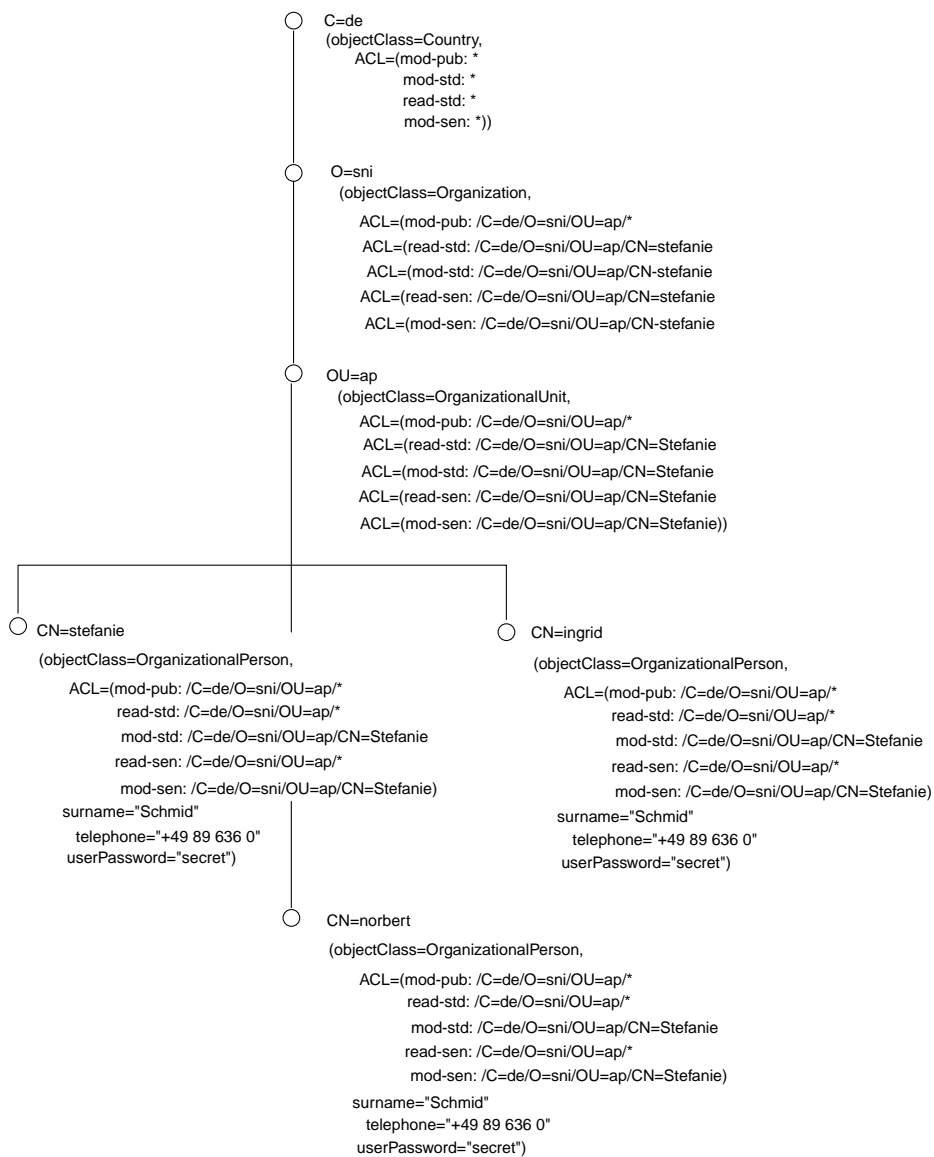
This section describes a typical search operation by using the **ds_search()** function call. It only includes the tasks directly related to the search operation and does not include tasks related to the XDS interface or other directory operations.

A typical search operation involves the following steps:

1. Using the **OM_EXPORT** macro, define the necessary object identifier constants for the OM classes and OM attributes that will define public objects for input to **ds_search()**.
2. Declare the variables that will contain the output from the XDS functions that will be used in the application.
3. Build public objects (descriptor lists) for the *name* parameter to **ds_search()**.
4. Specify the portion of the DIT to be searched.
5. Create a descriptor list for the *filter* parameter to **ds_search()** that designates which entries are to be eliminated from the search.
6. Create a descriptor list for the *selection* parameter to **ds_search()** that selects the type and scope of information in your request.
7. Perform the search operation.

These steps are demonstrated in the following code fragments from **acl.h**. The program includes a search operation. In order to perform the operation, the program assumes the directory contains the subtree shown in Figure 6-2.

Figure 6–2. Subtree for the acl.h Sample Program



6.8.2 Step 1: Export Object Identifiers

Most application programs find it convenient to export all the names they use from the same C source module. In the following **acl.c** fragment, the **OM_EXPORT** macro allocates memory for the constants that represent the object OM classes and OM attributes required for the search operation:

```
/* The application must export the object identifiers it */
/* requires.                                           */

OM_EXPORT (DS_C_AVA)
OM_EXPORT (DS_C_DS_RDN)
OM_EXPORT (DS_C_DS_DN)
OM_EXPORT (DS_C_ENTRY_INFO_SELECTION)
OM_EXPORT (DS_C_ATTRIBUTE)
OM_EXPORT (DS_C_ATTRIBUTE_LIST)
OM_EXPORT (DS_C_FILTER)
OM_EXPORT (DS_C_FILTER_ITEM)
OM_EXPORT (DSX_C_GDS_SESSION)
OM_EXPORT (DSX_C_GDS_CONTEXT)
OM_EXPORT (DSX_C_GDS_ACL)
OM_EXPORT (DSX_C_GDS_ACL_ITEM)

OM_EXPORT (DS_A_COUNTRY_NAME)
OM_EXPORT (DS_A_ORG_NAME)
OM_EXPORT (DS_A_ORG_UNIT_NAME)
OM_EXPORT (DS_A_COMMON_NAME)
OM_EXPORT (DS_A_LOCALITY_NAME)
OM_EXPORT (DS_A_OBJECT_CLASS)
OM_EXPORT (DS_A_USER_PASSWORD)
OM_EXPORT (DS_A_PHONE_NBR)
OM_EXPORT (DS_A_SURNAME)
OM_EXPORT (DS_A_ACL)
OM_EXPORT (DS_TYPELESS_RDN)

OM_EXPORT (DS_O_TOP)
OM_EXPORT (DS_O_COUNTRY)
OM_EXPORT (DS_O_ORG)
OM_EXPORT (DS_O_ORG_UNIT)
OM_EXPORT (DS_O_PERSON)
```

```
OM_EXPORT (DS_O_ORG__PERSON)
```

The **OM_EXPORT** macro takes the OM class name as input and creates two new data structures: a character string and structure of type **OM_string**. The structure of type **OM_string** contains a length and a pointer that are used in Step 3 to initialize the value of the object identifier.

6.8.3 Step 2: Declare Local Variables

The local variables are defined in the following code fragment from **acl.c**:

```
OM_workspace      workspace;      /* workspace for objects          */
OM_private_object session;        /* Session object.                */
OM_private_object bound_session; /* Holds the Session object which */
                                /* is returned by ds_bind()        */
OM_public_object  context;        /* Context object.                */
OM_private_object result;         /* Holds the search result object. */
OM_sint           invoke_id;      /* Integer for the invoke id       */
                                /* returned by ds_search().        */
                                /* (this parameter must be present */
                                /* even though it is ignored).     */
OM_type           sinfo_list[] = { DS_SEARCH_INFO, 0 };
OM_type           entry_list[] = { DS_ENTRIES, 0 };
                                /* Lists of types to be extracted */
OM_public_object  sinfo;          /* Search-Info object from result. */
OM_public_object  entry;          /* Entry object from search info.  */
OM_value_position total_num;      /* Number of descriptors returned. */
OM_return_code    rc;             /* XOM function return code.       */
register int       i;
char              user_name[MAX_DN_LEN];
                                /* Holds requestor's name.        */
char              entry_string[MAX_DN_LEN + 7] = "[?r??] ";
                                /* Holds entry details.           */
```

The data types shown in this code fragment are defined in a **typedef** statement in the **xom.h** header file. Since asynchronous operations (within the same thread) are not supported, the *invoke_id* functionality is redundant. The *invoke_id* parameter must be

supplied to the XDS functions as described in the *OSF DCE Application Development Reference*, but its return value should be ignored.

6.8.4 Step 3: Build Public Objects for the name Parameter to `ds_search()`

The public objects required by the search operation are defined in the **acl.h** header file. The *name* input parameter in the **ds_search()** function call in **acl.c** is the representation of the distinguished name for the root of the DIT. The following code fragment from **acl.c** shows how the descriptor list for the distinguished name is initialized:

```
static OM_descriptor dn_root[] = {
    OM_OID_DESC(OM_CLASS, DS_C_DS_DN),
    OM_NULL_DESCRIPTOR
};
```

6.8.5 Step 4: Specify the Portion of the DIT To Be Searched

The **ds_search()** call requires the *subset* input parameter. The *subset* parameter specifies the portion of the DIT to be searched. It takes the value of one of the following symbolic constants, which are defined in the **xds.h** header file:

- **DS_BASE_OBJECT**, meaning to search just the given object entry
- **DS_ONE_LEVEL**, meaning to search just the immediate subordinates of the given object entry
- **DS_WHOLE_SUBTREE**, meaning to search the given object and all its subordinates

The *subset* parameter in **acl.c** takes the value **DS_WHOLE_SUBTREE**.

6.8.6 Step 5: Create a Filter

The *filter* input parameter is used to eliminate entries from the search that are not wanted. Information is only returned on entries that satisfy the filter.

DS_C_FILTER inherits the attributes from its superclass **OM_C_OBJECT**, as do all OM classes. **OM_C_OBJECT** (as defined in Chapter 11) has one OM attribute, **OM_CLASS**, which has the value of an object identifier string that identifies the numeric representation of the object's OM class. **DS_C_FILTER**, on the other hand, has several OM attributes.

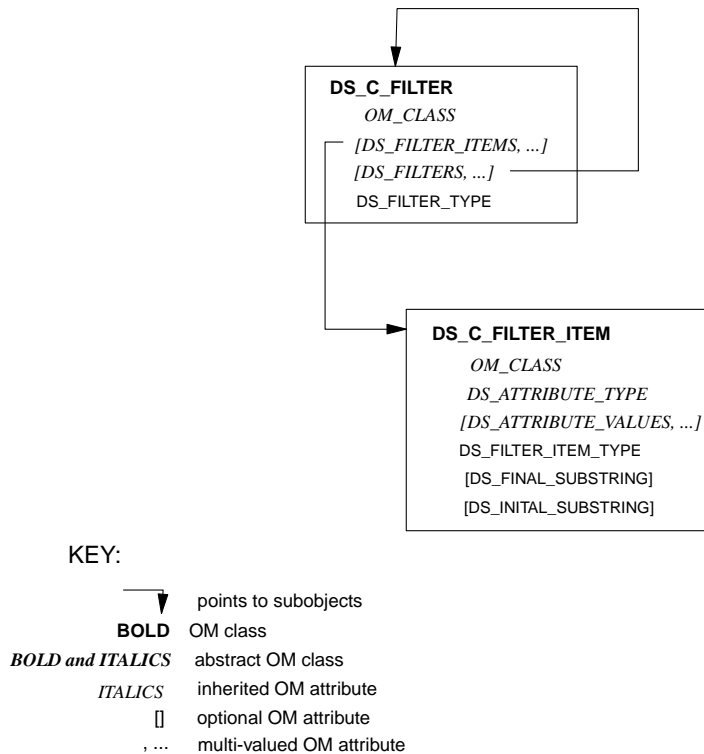
The purpose of **DS_C_FILTER** is to select or reject an object on the basis of information in its directory entry. It has the following OM attributes:

- **DS_FILTER_ITEMS**
- **DS_FILTERS**
- **DS_FILTER_TYPE**

Two of these OM attributes, **DS_FILTER_ITEMS** and **DS_FILTERS**, have values that are OM object classes themselves. The OM attribute **DS_FILTER_ITEMS** has the value OM class **DS_C_FILTER_ITEM**. **DS_C_FILTER_ITEM** is a component of a filter and defines the nature of the filter. The OM attribute **DS_FILTERS** has the value of OM class **DS_C_FILTER** and thus defines a collection of filters. The OM attribute **DS_FILTER_TYPE** has a value that is an enumerated type, which takes one of the values **DS_AND**, **DS_OR**, or **DS_NOT**.

Figure 6-3 shows the relationship of **DS_C_FILTER** to its superclass **OM_C_OBJECT**, and its attributes.

Figure 6–3. OM Class DS_C_FILTER



The **DS_NO_FILTER** constant can be used as the value of this parameter if all entries are searched and no entries are eliminated. This corresponds to a filter with a **DS_FILTER_TYPE** value of **DS_AND**, and no values of the **DS_FILTERS** or **DS_FILTER_ITEMS** OM attributes.

The following code fragment from **acl.c** shows the descriptor list for a filter:

```

/* The following descriptor list specifies a filter */
/* for search : */
/* (Present: objectClass) */

static OM_descriptor filter_item[] = {
    OM_OID_DESC(OM_CLASS, DS_C_FILTER_ITEM),
    {DS_FILTER_ITEM_TYPE, OM_S_ENUMERATION, {DS_PRESENT, 0} },

```

```
    OM_OID_DESC(DS_ATTRIBUTE_TYPE, DS_A_OBJECT_CLASS),
    OM_NULL_DESCRIPTOR
};

static OM_descriptor filter[] = {
    OM_OID_DESC(OM_CLASS, DS_C_FILTER),
    {DS_FILTER_ITEMS, OM_S_OBJECT, {0, filter_item} },
    {DS_FILTER_TYPE, OM_S_ENUMERATION, {DS_AND, 0} },
    OM_NULL_DESCRIPTOR
};
```

6.8.7 Step 6: Create an Entry-Information-Selection Parameter

The **ds_search()** call requires a *selection* input parameter to specify what information from the entry is requested. The *selection* parameter of the **ds_search()** call in **acl.h** requests information on all attributes, as shown in the following code fragment:

```
static OM_descriptor selection_acl[] = {
    OM_OID_DESC(OM_CLASS, DS_C_ENTRY_INFO_SELECTION),
    {DS_ALL_ATTRIBUTES, OM_S_BOOLEAN, OM_FALSE},
    OM_OID_DESC(DS_ATTRIBUTES_SELECTED, DSX_A_ACL),
    {DS_INFO_TYPE, OM_S_ENUMERATION, DS_TYPES_AND_VALUES},
    OM_NULL_DESCRIPTOR
};
```

As shown in the code fragment, **DS_ALL_ATTRIBUTES** has a syntax of **OM_S_BOOLEAN** that is set to **OM_FALSE**, indicating that only the requested attributes of the entry are to be returned. The ACL attribute's types and values are selected. **DS_INFO_TYPE** has a value of **DS_TYPES_AND_VALUES**, indicating that both the attribute types and the attribute values in the entry are returned.

6.8.8 Step 7: Perform the Search Operation

The following code fragment from **acl.c** shows the **ds_search()** function call:


```

/* Search the whole subtree below root.
 * The filter selects entries with an object-class attribute.
 * The selection extracts the ACL attribute from each
 * selected entry.
 * The results are returned in the private object "result".
 *
 * NOTE: Since every entry contains an object-class attribute the
 *       filter performs no function other than to demonstrate how
 *       filters may be used.
 */
if(ds_search(bound_session, context, dn_root, DS_WHOLE_SUBTREE,
  filter, OM_FALSE, selection_acl, &result, &invoke_id) != DS_SUCCESS)
  printf("ds_search() error\n");

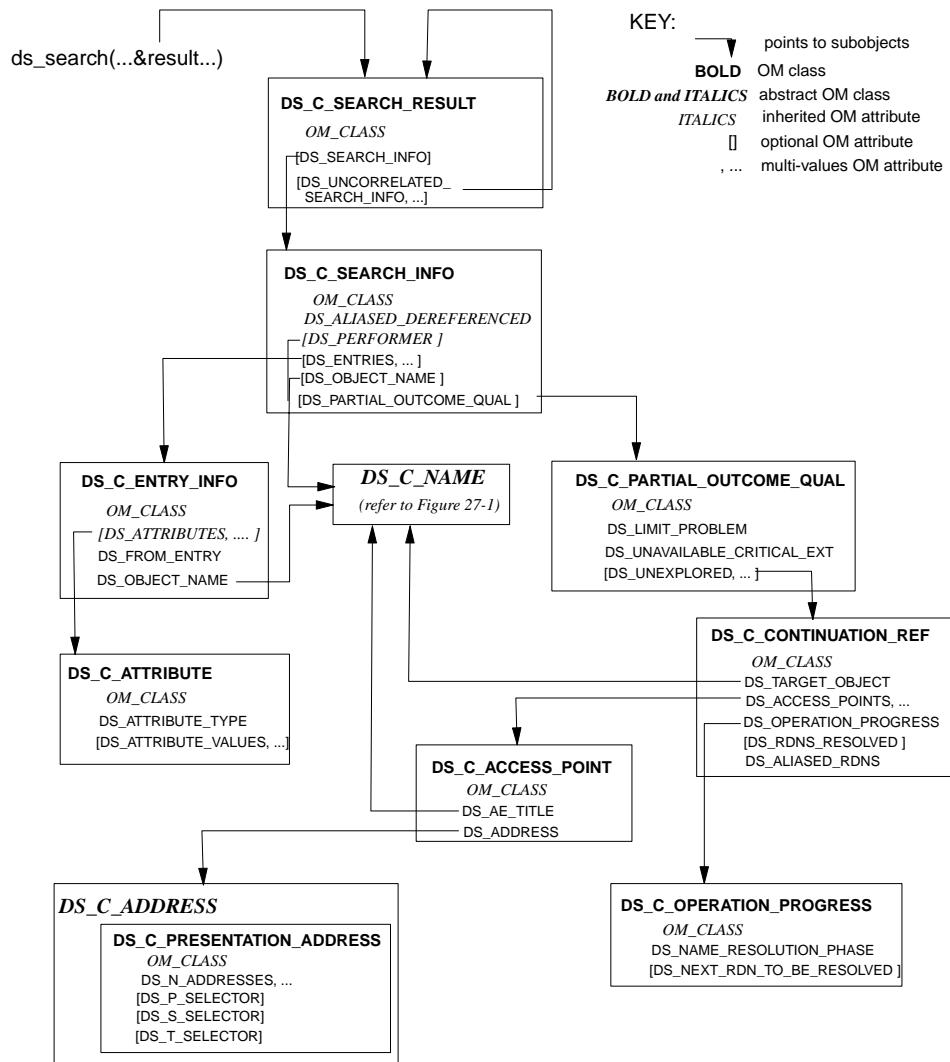
```

The **ds_search()** call returns the value **DS_SUCCESS** if the call successfully completes. Otherwise, it returns an error code. (Refer to Chapter 11 for a comprehensive list of error codes.)

The result of the search operation is returned to the workspace in a private object **result**. This result is returned as a single OM object. The components of the result are represented by OM attributes in the operation's **result** object.

The OM class returned by **ds_search()** is **DS_C_SEARCH_RESULT**. The superclasses, subclasses, and attributes for **DS_C_SEARCH_RESULT** are shown in Figure 6-4.

Figure 6–4. OM Class DS_C_SEARCH_RESULT



The **result** object is returned to the workspace in a private implementation-specific format. As such, it cannot be read directly by an application program, but requires a series of **om_get()** function calls to extract the requested information.

6.9 Directory Modify Operations

Modify functions alter information in the directory. For example, if an employee of an organizational unit transfers to a new organizational unit, a typical modify request would modify the **OU** name attribute in the person's directory entry to reflect the change.

XDS API implements the following modify functions:

- **ds_modify_entry()**

The requestor gives a distinguished name and a list of modifications to the named entry. The directory service carries out the specified changes if the user requesting the change has proper access rights.

- **ds_add_entry()**

The requestor gives a distinguished name and values for a new entry. The entry is added as a leaf node in the DIT if the user requesting the change has proper access rights.

- **ds_remove_entry()**

The requestor gives a distinguished name. The entry with that name is removed if the user requesting the change has proper access rights.

- **ds_modify_rdn()**

The requestor gives a distinguished name and a new RDN for the entry. The directory changes the entry's RDN if the user requesting the change has proper access rights.

Note that **ds_add_entry()**, **ds_remove_entry()**, and **ds_modify_rdn()** only apply to leaf entries. They are not intended to provide a general facility for building and manipulating the DIT.

6.9.1 Modifying Directory Entries

This section describes a modification and subsequent listing of the DIT by using the **ds_add_entry()**, **ds_list()**, and **ds_remove_entry()** function calls. It includes a

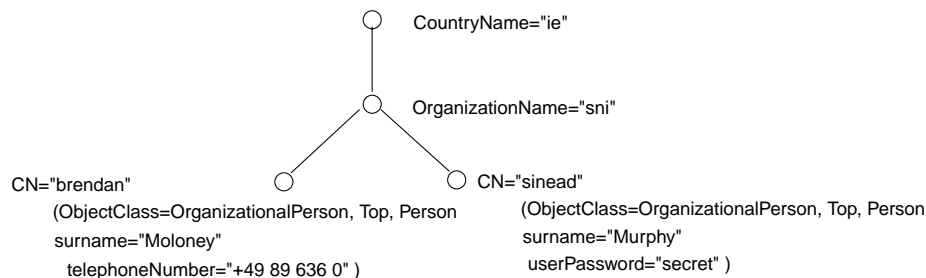
description of tasks directly related to these operations and does not include service-related tasks. It does not include a **ds_modify_entry()** function call. The modify operation is used in the context of the X.500 *Abstract Service Definition*.

A typical operation to add, remove, or list an entry involves following the same basic steps that were defined previously for the read and search operations:

1. Using the **OM_EXPORT** macro, define the necessary object identifier constants for the OM classes and OM attributes that will define public objects for input to the function calls.
2. Declare the variables that will contain the output from the XDS functions you will use in your application.
3. Build public objects (descriptor lists) for the *name* parameters to the function calls.
4. Create descriptor lists for the attributes to be added, removed, or listed.
5. Perform the operations.

These steps are demonstrated in the following code fragments. The program adds two entries to the directory, then a list operation is performed on their superior entry, and finally the two entries are removed from the directory. The directory tree shown in Figure 6-5 is used in the program.

Figure 6–5. A Sample Directory Tree



6.9.2 Step 1: Export Object Identifiers for Required Directory Classes and Attributes

In the following code fragment, the **OM_EXPORT** macro allocates memory for the constants that represent the object classes and attributes required for the add, list, and remove operations:

```
/* The application has to export the object identifiers */
/* it requires. */

OM_EXPORT (DS_C_AVA)
OM_EXPORT (DS_C_DS_RDN)
OM_EXPORT (DS_C_DS_DN)
OM_EXPORT (DS_C_ENTRY_INFO_SELECTION)
OM_EXPORT (DS_C_ATTRIBUTE)
OM_EXPORT (DS_C_ATTRIBUTE_LIST)

OM_EXPORT (DS_A_COUNTRY_NAME)
OM_EXPORT (DS_A_ORG_NAME)
OM_EXPORT (DS_A_ORG_UNIT_NAME)
OM_EXPORT (DS_A_COMMON_NAME)
OM_EXPORT (DS_A_OBJECT_CLASS)
OM_EXPORT (DS_A_PHONE_NBR)
OM_EXPORT (DS_A_USER_PASSWORD)
OM_EXPORT (DS_A_SURNAME)

OM_EXPORT (DS_O_TOP)
OM_EXPORT (DS_O_PERSON)
OM_EXPORT (DS_O_ORG_PERSON)
```

6.9.3 Step 2: Declare Local Variables

The local variables *bound_session*, *result*, and *invoke_id* are defined in the following sample code fragment:

```
OM_private_object bound_session; /* Holds the Session object */
                                /* which is returned by */
                                /* ds_bind(). */
```

```
OM_private_object result;      /* Holds the list result      */
                                /* object.                      */
OM_sint          invoke_id;    /* Integer for the invoke id */
                                /* returned by ds_search(). */
                                /* This parameter must be   */
                                /* present even though it is */
                                /* ignored.                  */
```

These data types are defined in **typedef** statements in the **xom.h** header file. The *bound_session* and *result* variables are defined as data type **OM_private_object** because they are returned by **ds_bind()** and **ds_list()** operations to the workspace as private objects. Since asynchronous operations (within the same thread) are not supported, the *invoke_id* parameter functionality is redundant. The *invoke_id* parameter must be supplied to the XDS functions as described in the *OSF DCE Application Development Reference*, but its return value should be ignored.

6.9.4 Step 3: Build Public Objects

The public objects required by the **ds_add_entry()**, **ds_list()**, and **ds_remove_entry()** operations are defined in the following code fragment:

```
/* Build up descriptor lists for the following distinguished names: */
/*      C=ie/O=sni                                                    */
/*      C=ie/O=sni/OU=ap/CN=brendan                                   */
/*      C=ie/O=sni/OU=ap/CN=sinead                                    */

static OM_descriptor  ava_ie[] = {
    OM_OID_DESC(OM_CLASS, DS_C_AVA),
    OM_OID_DESC(DS_ATTRIBUTE_TYPE, DS_A_COUNTRY_NAME),
    {DS_ATTRIBUTE_VALUES, OM_S_PRINTABLE_STRING, OM_STRING("ie")},
    OM_NULL_DESCRIPTOR
};

static OM_descriptor  ava_sni[] = {
    OM_OID_DESC(OM_CLASS, DS_C_AVA),
    OM_OID_DESC(DS_ATTRIBUTE_TYPE, DS_A_ORG_NAME),
    {DS_ATTRIBUTE_VALUES, OM_S_TELETEX_STRING, OM_STRING("sni")},
    OM_NULL_DESCRIPTOR
};
```

```

static OM_descriptor    ava_ap[] = {
    OM_OID_DESC(OM_CLASS, DS_C_AVA),
    OM_OID_DESC(DS_ATTRIBUTE_TYPE, DS_A_ORG_UNIT_NAME),
    {DS_ATTRIBUTE_VALUES, OM_S_TELETEX_STRING, OM_STRING("ap")},
    OM_NULL_DESCRIPTOR
};

static OM_descriptor    ava_brendan[] = {
    OM_OID_DESC(OM_CLASS, DS_C_AVA),
    OM_OID_DESC(DS_ATTRIBUTE_TYPE, DS_A_COMMON_NAME),
    {DS_ATTRIBUTE_VALUES, OM_S_TELETEX_STRING, OM_STRING("brendan")},
    OM_NULL_DESCRIPTOR
};

static OM_descriptor    ava_sinead[] = {
    OM_OID_DESC(OM_CLASS, DS_C_AVA),
    OM_OID_DESC(DS_ATTRIBUTE_TYPE, DS_A_COMMON_NAME),
    {DS_ATTRIBUTE_VALUES, OM_S_TELETEX_STRING, OM_STRING("sinead")},
    OM_NULL_DESCRIPTOR
};

static OM_descriptor    rdn_ie[] = {
    OM_OID_DESC(OM_CLASS, DS_C_DS_RDN),
    {DS_AVAS, OM_S_OBJECT, {0, ava_ie}},
    OM_NULL_DESCRIPTOR
};

static OM_descriptor    rdn_sni[] = {
    OM_OID_DESC(OM_CLASS, DS_C_DS_RDN),
    {DS_AVAS, OM_S_OBJECT, {0, ava_sni}},
    OM_NULL_DESCRIPTOR
};

static OM_descriptor    rdn_ap[] = {
    OM_OID_DESC(OM_CLASS, DS_C_DS_RDN),
    {DS_AVAS, OM_S_OBJECT, {0, ava_ap}},
    OM_NULL_DESCRIPTOR
};

static OM_descriptor    rdn_brendan[] = {
    OM_OID_DESC(OM_CLASS, DS_C_DS_RDN),
    {DS_AVAS, OM_S_OBJECT, {0, ava_brendan}},
    OM_NULL_DESCRIPTOR
};

static OM_descriptor    rdn_sinead[] = {
    OM_OID_DESC(OM_CLASS, DS_C_DS_RDN),

```

```
        {DS_AVAS, OM_S_OBJECT, {0, ava_sinead}},
        OM_NULL_DESCRIPTOR
    };

    static OM_descriptor dn_ap[] = {
        OM_OID_DESC(OM_CLASS, DS_C_DS_DN),
        {DS_RDNS, OM_S_OBJECT, {0, rdn_ie}},
        {DS_RDNS, OM_S_OBJECT, {0, rdn_sni}},
        {DS_RDNS, OM_S_OBJECT, {0, rdn_ap}},
        OM_NULL_DESCRIPTOR
    };

    static OM_descriptor dn_brendan[] = {
        OM_OID_DESC(OM_CLASS, DS_C_DS_DN),
        {DS_RDNS, OM_S_OBJECT, {0, rdn_ie}},
        {DS_RDNS, OM_S_OBJECT, {0, rdn_sni}},
        {DS_RDNS, OM_S_OBJECT, {0, rdn_ap}},
        {DS_RDNS, OM_S_OBJECT, {0, rdn_brendan}},
        OM_NULL_DESCRIPTOR
    };

    static OM_descriptor dn_sinead[] = {
        OM_OID_DESC(OM_CLASS, DS_C_DS_DN),
        {DS_RDNS, OM_S_OBJECT, {0, rdn_ie}},
        {DS_RDNS, OM_S_OBJECT, {0, rdn_sni}},
        {DS_RDNS, OM_S_OBJECT, {0, rdn_ap}},
        {DS_RDNS, OM_S_OBJECT, {0, rdn_sinead}},
        OM_NULL_DESCRIPTOR
    };
};
```

6.9.5 Step 4: Create Descriptor Lists for Attributes

The following code fragments show how the attribute lists are created for the attributes to be added to the directory.

First, initialize the public object **object_class** to contain the representation of the classes in the DIT that are common to both **Organizational-Person** entries, **Top, Person**, and **Organizational-Person**:


```

/* Build up an array of object identifiers for the      */
/* attributes to be added to the directory.             */

static OM_descriptor    object_class[] = {
    OM_OID_DESC(OM_CLASS, DS_C_ATTRIBUTE),
    OM_OID_DESC(DS_ATTRIBUTE_TYPE, DS_A_OBJECT_CLASS),
    OM_OID_DESC(DS_ATTRIBUTE_VALUES, DS_O_TOP),
    OM_OID_DESC(DS_ATTRIBUTE_VALUES, DS_O_PERSON),
    OM_OID_DESC(DS_ATTRIBUTE_VALUES, DS_O_ORG_PERSON),
    OM_NULL_DESCRIPTOR
};

```

Next, initialize the public objects that represent the attributes to be added. These are **surname** and **telephone** for the distinguished name of Brendan, and **surname2** and **password** for the distinguished name of Sinead:

```

static OM_descriptor    telephone[] = {
    OM_OID_DESC(OM_CLASS, DS_C_ATTRIBUTE),
    OM_OID_DESC(DS_ATTRIBUTE_TYPE, DS_A_PHONE_NBR),
    {DS_ATTRIBUTE_VALUES, OM_S_PRINTABLE_STRING,
    OM_STRING("+49 89 636 0")},
    OM_NULL_DESCRIPTOR
};

static OM_descriptor    surname[] = {
    OM_OID_DESC(OM_CLASS, DS_C_AVA),
    OM_OID_DESC(DS_ATTRIBUTE_TYPE, DS_A_SURNAME),
    {DS_ATTRIBUTE_VALUES, OM_S_TELETEX_STRING,
    OM_STRING("Moloney")},
    OM_NULL_DESCRIPTOR
};

static OM_descriptor    surname2[] = {
    OM_OID_DESC(OM_CLASS, DS_C_AVA),
    OM_OID_DESC(DS_ATTRIBUTE_TYPE, DS_A_SURNAME),
    {DS_ATTRIBUTE_VALUES, OM_S_TELETEX_STRING,
    OM_STRING("Murphy")},
    OM_NULL_DESCRIPTOR
};

```

```
static OM_descriptor    password[] = {
    OM_OID_DESC(OM_CLASS, DS_C_AVA),
    OM_OID_DESC(DS_ATTRIBUTE_TYPE, DS_A_USER_PASSWORD),
    {DS_ATTRIBUTE_VALUES, OM_S_OCTET_STRING,
    OM_STRING("secret")},
    OM_NULL_DESCRIPTOR
};
```

Finally, initialize the public objects that represent the list of attributes to be added to the directory. These are *attr_list1* for the distinguished name Brendan, and *attr_list2* for the distinguished name Sinead:

```
static OM_descriptor    attr_list1[] = {
    OM_OID_DESC(OM_CLASS, DS_C_ATTRIBUTE_LIST),
    {DS_ATTRIBUTES, OM_S_OBJECT, {0, object_class} },
    {DS_ATTRIBUTES, OM_S_OBJECT, {0, surname} },
    {DS_ATTRIBUTES, OM_S_OBJECT, {0, telephone} },
    OM_NULL_DESCRIPTOR
};

static OM_descriptor    attr_list2[] = {
    OM_OID_DESC(OM_CLASS, DS_C_ATTRIBUTE_LIST),
    {DS_ATTRIBUTES, OM_S_OBJECT, {0, object_class} },
    {DS_ATTRIBUTES, OM_S_OBJECT, {0, surname2} },
    {DS_ATTRIBUTES, OM_S_OBJECT, {0, password} },
    OM_NULL_DESCRIPTOR
};
```

The *attr_list1* variable contains the public objects **surname** and **telephone**, which are the C representations of the attributes of the distinguished name **/C=ie/O=sni/OU=ap/CN=Brendan** that are added to the directory. The *attr_list2* variable contains the public objects first **surname2** and **password**, which are the C representations of the attributes of the distinguished name **/C=ie/O=sni/OU=ap/CN=Sinead**.

6.9.6 Step 5: Perform the Operations

The following code fragments show the **ds_add_entry()**, **ds_list()**, and the **ds_remove_entry()** calls.

First, the two **ds_add_entry()** function calls add the attribute lists contained in *attr_list1* and *attr_list2* to the distinguished names represented by **dn_brendan** and **dn_sinead**, respectively:

```
/* Add two entries to the GDS server.                                */

if (ds_add_entry(bound_session, DS_DEFAULT_CONTEXT,
    dn_brendan, attr_list1,
    &invoke_id) != DS_SUCCESS)
    printf("ds_add_entry() error\n");

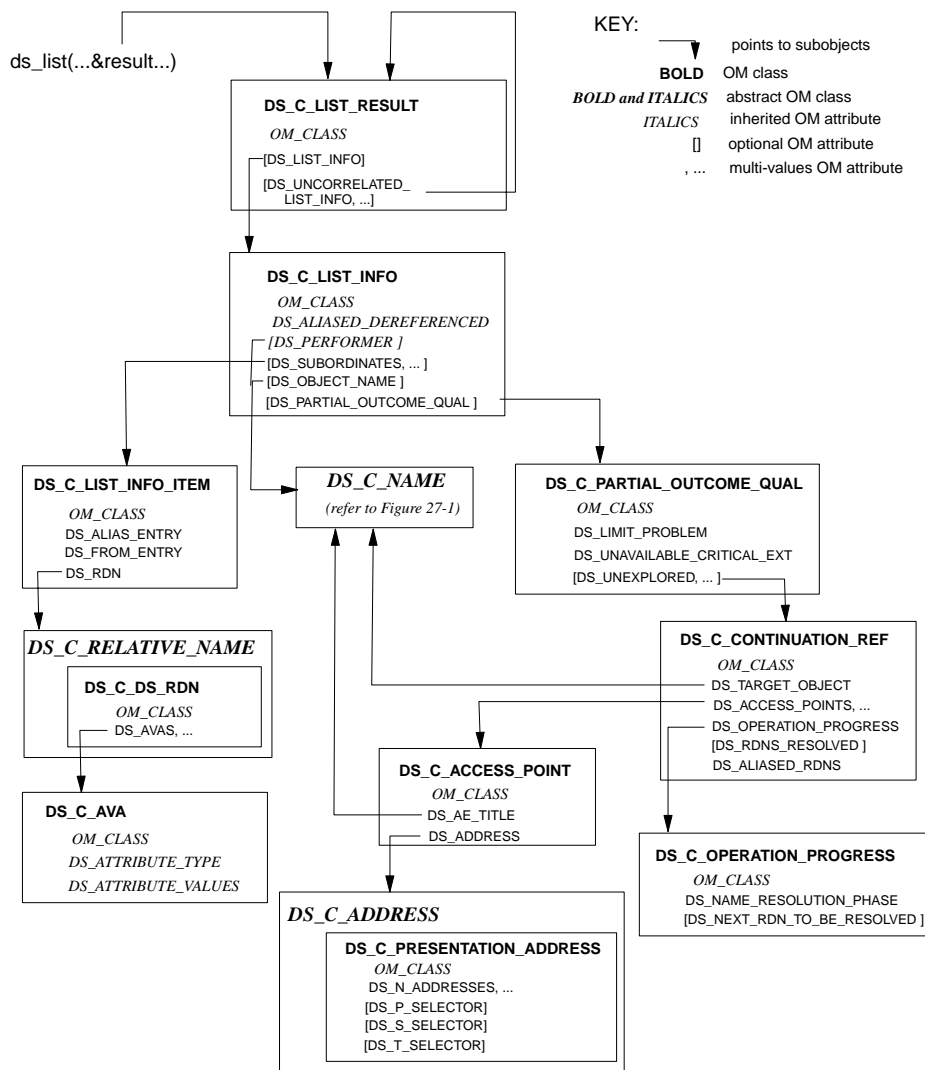
if (ds_add_entry(bound_session, DS_DEFAULT_CONTEXT,
    dn_sinead, attr_list2,
    &invoke_id) != DS_SUCCESS)
    printf("ds_add_entry() error\n");
```

Next, list all the subordinates of the object referenced by the distinguished name /**C=ie/O=sni/OU=ap**:

```
if (ds_list(bound_session, DS_DEFAULT_CONTEXT, dn_ap,
    &result, &invoke_id)
    != DS_SUCCESS)
    printf("ds_list() error\n");
```

The **ds_list()** call returns the result in the private object **result** to the workspace. The components of **result** are represented by OM attributes in the OM class **DS_C_LIST_RESULT** (as shown in Figure 6-6) and can only be read by a series of **om_get()** calls.

Figure 6–6. OM Class DS_C_LIST_RESULT



Finally, remove the two entries from the directory:

```

if (ds_remove_entry(bound_session, DS_DEFAULT_CONTEXT,
    dn_brendan, &invoke_id)

```

```
    != DS_SUCCESS)
    printf("ds_remove_entry() error\n");

if (ds_remove_entry(bound_session, DS_DEFAULT_CONTEXT,
    dn_sinead, &invoke_id)
    != DS_SUCCESS)
    printf("ds_remove_entry() error\n");
```

6.10 Return Codes

XDS API function calls return a value of type **DS_status**, with the exception of **ds_initialize()** which returns a value of type **OM_workspace**. If the function is successful, then **DS_status** returns with a value of **DS_SUCCESS**. If the function does not complete successfully, then **DS_status** takes either the error constant **DS_NO_WORKSPACE** or one of the private error objects described in Chapter 11.

Chapter 7

Sample Application Programs

This chapter contains three sample programs and the header files that are included in them (in parentheses), as follows:

- **example.c** (**example.h**)
- **acl.c** (**acl.h**)
- **teldir.c**

Most of the concepts that you will need to know to understand and use these programs are discussed in previous chapters in this guide. The programs are arranged so that the simplest program (**example.c**) is presented first and the most complex program (**teldir.c**) is presented last. The three programs demonstrate basic XDS and XOM API principles and concepts in operation. The **teldir.c** program is considerably more complex and uses a more sophisticated approach. It allows the user to enter values dynamically; for example, a surname and phone number.

For a sample XDS application that uses threads, please refer to Chapter 8. The **acl.c** sample program is presented again in Chapter 9, this time using the XDS/XOM convenience routines.

7.1 General Programming Guidelines

Writing an application program by using XDS and XOM APIs involves the following general steps before you begin coding:

1. Select the interface functions that you will need for your application and determine the parameters for the function calls.
2. Check for abstract OM classes and superclasses of objects that you will manipulate for inherited OM attributes in Part 4.
3. Find the correct symbolic constants of the appropriate packages; these can be found in the header files included with the GDS API, such as **xdsbdc.h**.
4. Determine the error handling required.

7.2 The **example.c** Program

The **example.c** program uses XDS API in synchronous mode to read a telephone number or numbers of a distinguished name. The program consists of the following general steps:

1. Define the required object identifier constants.
2. Declare the variables involved with directory service operations (Steps 3, 4, 7, 8, and 9).
3. Build the distinguished name of **Peter Piper** as a public object for the input parameter to **ds_read()**.
4. Build a public object for the *selection* parameter to **ds_read()**.
5. Declare the variables to extract the telephone numbers by using **om_get()**.
6. Initialize the directory service and get an OM workspace.
7. Pull in the required packages.
8. Bind to a default directory session.
9. Perform the read operation to extract the telephone number of a distinguished name from the directory.
10. Terminate the directory service session.

11. Extract the telephone number(s) by using a series of **om_get()** calls.
12. Release the storage occupied by private and public objects that are no longer needed.
13. Print the telephone number string.
14. Release the storage occupied by public objects containing telephone numbers.
15. Continue processing and exit.

Step 1 uses the **OM_EXPORT** macro to allocate memory for the object identifier constants that represent an OM class or OM attribute. These constants are the OM attribute values that are used to build the public objects that are required as input to **ds_read()**.

Step 2 declares the variables for directory service operations and error handling. The *session* and *workspace* variables are required for binding a session to a server and creating a workspace into which **ds_read()** can deposit the results of the read operation on the directory.

The *result* variable is a pointer that is returned by **ds_read()** to the workspace. The information stored in *result* is in implementation-specific private format that is not accessible directly by the application program. Subsequent **om_get()** calls extract the telephone number(s) requested by the program from *result* and store the information in the variable *telephones* (declared in Step 5).

The *error* and *return_code* variables are used by the program for error handling. The *error* variable is used for processing the return code from XDS API function calls. The *return_code* variable is used by the error handling header file **example.h** for processing return codes from **om_get()** function calls.

Step 3 builds the public object representing the distinguished name of **Peter Piper**. The program uses statically defined public objects to demonstrate the basic principles of building public objects. However, a more sophisticated approach is presented in the last sample program in this chapter, **teldir.c**. The **teldir.c** program dynamically defines a public object from a user-supplied name in DCE string format.

In the program **example.c**, the process starts with the definition of an array of descriptor lists as AVAs. These AVAs are public objects that are included in the definition of RDNs. The RDNs, in turn, are included in the distinguished name of **Peter Piper** stored in **name**. Using the same method of static definition, Step 4

defines the **DS_C_ENTRY_INFO_SELECTION** public object and stores it in the variable *selection*. The *name* and *selection* variables are required as input parameters to **ds_read()**. This process is described in detail in Chapter 6.

Step 5 declares the variables required by **om_get()** to extract the telephone number(s) from *result*. The *entry_list*, *attributes_list*, and *telephone_list* variables are of type **OM_type** and are initialized to the values of the OM attribute types **DS_ENTRY**, **DS_ATTRIBUTES**, and **DS_ATTRIBUTE_VALUES**, respectively. **DS_ENTRY** contains the selected list of entries, **DS_ATTRIBUTES** contains the selected list of attribute types, and **DS_ATTRIBUTE_VALUES** contains the actual values of the telephone numbers.

The *entry*, *attributes*, and *telephones* variables are of type **OM_public_object** because they store the output parameters of **om_get()**. The **om_get()** call makes these objects available to the application program as public object data types. The program must remove layers of objects and subobjects to get at the actual string data values of the telephone numbers.

The *telephones* variable contains the actual string values of the telephone number(s). It is a descriptor in the array of descriptors that make up the public object that contains the actual string data that the program wants to extract from the directory.

Step 6 initializes the directory service and gets an OM workspace in which **ds_read()** deposits the result of the read operation.

Step 7 pulls the basic directory contents package into the program because it contains features that are required by the program but not included in the default package (the directory service package).

Step 8 binds the session to the default session. An application program can bind with a specifically tailored session object by using an instance of OM class **DS_C_SESSION**. In most cases, however, it is sufficient to use the constant **DS_DEFAULT_SESSION**. **DS_DEFAULT_SESSION** uses the default values of **DS_C_SESSION** and the values of specific OM attributes that are set locally in the cache. These OM attributes are **DS_DSA_ADDRESS** (the address of the default DSA) and **DS_DSA_NAME** (the distinguished name of the default DSA). It is the responsibility of local administrators to make sure that these default values are set correctly in the cache.

Step 9 performs the read operation and deposits the result in the workspace in *result*. The *Sresult* variable is one of the input parameters for the **om_get()** function call.

The **session** variable and the **DS_DEFAULT_CONTEXT** constant are the *session* and *context* parameters required to be present in the **ds_read()** function call.

The *name* variable holds the public object representing the distinguished name of **Peter Piper**; the *selection* variable contains the public object indicating which attributes and values are selected by the read operation from the entry. The *invoke_id* parameter is not used by the DCE implementation of XDS and is ignored.

Step 10 terminates the directory session.

Step 11 uses a series of **om_get()** calls to extract the telephone number(s). The first **om_get()** extracts the information about the entry from *result* and puts it in *entry*. The second **om_get()** extracts the attribute types from *entry* and puts them in *attributes*. The third **om_get()** extracts the actual values of the telephone numbers from *attributes* and puts them in *telephones*. The *telephones* variable contains the string data values of the telephone number(s).

Step 12 releases the storage occupied by the private and public objects that are no longer needed. The program has the data values in *telephone* that it needs to continue processing. A **ds_shutdown()** call is issued to shut down the interface established by **ds_initialize()**.

Step 13 prints out each telephone number associated with the distinguished name **Peter Piper** in the directory, or returns an error message if the number is not in the correct format. It checks for an attribute with type **DS_ATTRIBUTE_VALUES** and a syntax of **OM_S_PRINTABLE_STRING**, the proper syntax for a telephone number. The constant **OM_S_SYNTAX** is used to mask the six high-order bits in the syntax because they are used internally by the XOM service.

Step 14 releases the storage occupied by *telephones* because it is no longer needed.

Step 15 continues processing and exits.

7.2.1 The example.c Code

The following code is a listing of the **example.c** program:

```
/*
 * sample application that uses XDS in synchronous mode
 *
 * This program reads the telephone number(s) of a given target name.
 */

#ifdef THREADSAFE
#include <pthread.h>
#endif

#include <stdio.h>

#include <xom.h>
#include <xds.h>
#include <xdsbdcp.h>

#include "example.h"          /* possible Error Handling header */

/* Step 1 */
*
* Define necessary Object Identifier constants
*/
OM_EXPORT(DS_A_COMMON_NAME)
OM_EXPORT(DS_A_COUNTRY_NAME)
OM_EXPORT(DS_A_ORG_NAME)
OM_EXPORT(DS_A_ORG_UNIT_NAME)
OM_EXPORT(DS_A_PHONE_NBR)
OM_EXPORT(DS_C_AVA)
OM_EXPORT(DS_C_DS_DN)
OM_EXPORT(DS_C_DS_RDN)
OM_EXPORT(DS_C_ENTRY_INFO_SELECTION)

/* Step 2 */

int main(void)
{
    DS_status      error;          /* return value from DS functions */
    OM_return_code  return_code;   /* return value from OM functions */
    OM_workspace    workspace;     /* workspace for objects */
    OM_private_object session;     /* session for directory operations */
}
```

```

OM_private_object  result;      /* result of read operation      */
OM_sint            invoke_id;   /* Invoke-ID of the read operation */
OM_value_position  total_num;   /* Number of Attribute Descriptors */

static DS_feature bdcg_package[] = {
    { OM_STRING(OMP_O_DS_BASIC_DIR_CONTENTS_PKG), OM_TRUE },
    { { (OM_uint32)0, (void *)0 }, OM_FALSE },
};

/* Step 3 */
*
* Public Object ("Descriptor List") for Name parameter to ds_read().
* Build the Distinguished-Name of Peter Piper.
*/

static OM_descriptor country[] = {
    OM_OID_DESC(OM_CLASS, DS_C_AVA),
    OM_OID_DESC(DS_ATTRIBUTE_TYPE, DS_A_COUNTRY_NAME),
    { DS_ATTRIBUTE_VALUES, OM_S_PRINTABLE_STRING, OM_STRING("US") },
    OM_NULL_DESCRIPTOR
};

static OM_descriptor organization[] = {
    OM_OID_DESC(OM_CLASS, DS_C_AVA),
    OM_OID_DESC(DS_ATTRIBUTE_TYPE, DS_A_ORG_NAME),
    { DS_ATTRIBUTE_VALUES, OM_S_TELETEX_STRING,
      OM_STRING("Acme Pepper Co") },
    OM_NULL_DESCRIPTOR
};

static OM_descriptor organizational_unit[] = {
    OM_OID_DESC(OM_CLASS, DS_C_AVA),
    OM_OID_DESC(DS_ATTRIBUTE_TYPE, DS_A_ORG_UNIT_NAME),
    { DS_ATTRIBUTE_VALUES, OM_S_TELETEX_STRING, OM_STRING("Research") },
    OM_NULL_DESCRIPTOR
};

static OM_descriptor common_name[] = {
    OM_OID_DESC(OM_CLASS, DS_C_AVA),
    OM_OID_DESC(DS_ATTRIBUTE_TYPE, DS_A_COMMON_NAME),
    { DS_ATTRIBUTE_VALUES, OM_S_TELETEX_STRING, OM_STRING("Peter Piper") },
    OM_NULL_DESCRIPTOR
};

```

```
static OM_descriptor rdn1[] = {
    OM_OID_DESC(OM_CLASS, DS_C_DS_RDN),
    { DS_AVAS, OM_S_OBJECT, { 0, country } },
    OM_NULL_DESCRIPTOR
};

static OM_descriptor rdn2[] = {
    OM_OID_DESC(OM_CLASS, DS_C_DS_RDN),
    { DS_AVAS, OM_S_OBJECT, { 0, organization } },
    OM_NULL_DESCRIPTOR
};

static OM_descriptor rdn3[] = {
    OM_OID_DESC(OM_CLASS, DS_C_DS_RDN),
    { DS_AVAS, OM_S_OBJECT, { 0, organizational_unit } },
    OM_NULL_DESCRIPTOR
};

static OM_descriptor rdn4[] = {
    OM_OID_DESC(OM_CLASS, DS_C_DS_RDN),
    { DS_AVAS, OM_S_OBJECT, { 0, common_name } },
    OM_NULL_DESCRIPTOR
};

OM_descriptor    name[] = {
    OM_OID_DESC(OM_CLASS, DS_C_DS_DN),
    { DS_RDNS, OM_S_OBJECT, { 0, rdn1 } },
    { DS_RDNS, OM_S_OBJECT, { 0, rdn2 } },
    { DS_RDNS, OM_S_OBJECT, { 0, rdn3 } },
    { DS_RDNS, OM_S_OBJECT, { 0, rdn4 } },
    OM_NULL_DESCRIPTOR
};

/* Step 4 */

/*
 *
 * Public Object ("Descriptor List") for
 * Entry-Information-Selection parameter to ds_read().
 */
OM_descriptor selection[] = {
    OM_OID_DESC(OM_CLASS, DS_C_ENTRY_INFO_SELECTION),
```

```
    { DS_ALL_ATTRIBUTES, OM_S_BOOLEAN, { OM_FALSE, NULL } },
    OM_OID_DESC(DS_ATTRIBUTES_SELECTED, DS_A_PHONE_NBR),
    { DS_INFO_TYPE, OM_S_ENUMERATION, { DS_TYPES_AND_VALUES, NULL } },
    OM_NULL_DESCRIPTOR
};

/* Step 5 */

/*
 * variables to extract the telephone number(s)
 */
OM_type          entry_list[]      = { DS_ENTRY, 0 };
OM_type          attributes_list[] = { DS_ATTRIBUTES, 0 };
OM_type          telephone_list[]  = { DS_ATTRIBUTE_VALUES, 0 };
OM_public_object entry;
OM_public_object attributes;
OM_public_object telephones;
OM_descriptor    *telephone; /* current phone number */

/*
 * Perform the directory service operations:
 * (1) Initialize the directory service and get a workspace
 * (2) bind a default directory session.
 * (3) read the telephone number of "name".
 * (4) terminate the directory session.
 */

/* Step 6 */

CHECK_DS_CALL((OM_object) !(workspace=ds_initialize()));

/* Step 7 */

CHECK_DS_CALL(ds_version(bdcp_package, workspace));

/* Step 8 */

CHECK_DS_CALL(ds_bind(DS_DEFAULT_SESSION, workspace, &session));

/* Step 9 */
```

```
CHECK_DS_CALL(ds_read(session, DS_DEFAULT_CONTEXT, name,
                      selection, &result, &invoke_id));

/*
 * NOTE: should check here for Attribute-Error (no-such-attribute)
 * in case the "name" doesn't have a telephone.
 * Then for all other cases call error_handler
 */

/* Step 10 */

CHECK_DS_CALL(ds_unbind(session));

/* Step 11 */

/*
 * extract the telephone number(s) of "name" from the result
 *
 * There are 4 stages:
 * (1) get the Entry-Information from the Read-Result.
 * (2) get the Attributes from the Entry-Information.
 * (3) get the list of phone numbers.
 * (4) scan the list and print each number.
 */

CHECK_OM_CALL(om_get(result,
                     OM_EXCLUDE_ALL_BUT_THESE_TYPES
                     + OM_EXCLUDE_SUBOBJECTS,
                     entry_list, OM_FALSE, 0, 0, &entry,
                     &total_num));

CHECK_OM_CALL(om_get(entry->value.object.object,
                     OM_EXCLUDE_ALL_BUT_THESE_TYPES
                     + OM_EXCLUDE_SUBOBJECTS,
                     attributes_list, OM_FALSE, 0, 0,
                     &attributes, &total_num));

CHECK_OM_CALL(om_get(attributes->value.object.object,
                     OM_EXCLUDE_ALL_BUT_THESE_TYPES
                     + OM_EXCLUDE_SUBOBJECTS,
```



```
        telephone_list, OM_FALSE, 0, 0,
        &telephones, &total_num));

/* Step 12 */

/* We can now safely release all the private objects
 * and the public objects we no longer need
 */
CHECK_OM_CALL(om_delete(session));
CHECK_OM_CALL(om_delete(result));
CHECK_OM_CALL(om_delete(entry));
CHECK_OM_CALL(om_delete(attributes));
CHECK_DS_CALL(ds_shutdown(workspace));
/* Step 13 */

for (telephone = telephones;
     telephone->type == DS_ATTRIBUTE_VALUES;
     telephone++)
{
    if (telephone->type != DS_ATTRIBUTE_VALUES
        || (telephone->syntax & OM_S_SYNTAX) != OM_S_PRINTABLE_STRING)
    {
        (void) fprintf(stderr, "malformed telephone number\n");
        exit(EXIT_FAILURE);
    }

    (void) printf("Telephone number: %.*s\n",
                  telephone->value.string.length,
                  telephone->value.string.elements);
}

/* Step 14 */

CHECK_OM_CALL(om_delete(telephones));

/* Step 15 */

/* more application-specific processing can occur here...
 */
```

```
/* ... and finally exit. */
exit(EXIT_SUCCESS);
}
```

7.2.2 Error Handling

The **example.c** program includes the header file **example.h** for error handling of XDS and XOM function calls. The **example.h** program contains two error-handling functions: **CHECK_DS_CALL** for handling XDS API errors, and **CHECK_OM_CALL** for handling XOM API errors. Note that **CHECK_DS_CALL** and **CHECK_OM_CALL** are created specifically for **example.c** and are not part of the XDS or XOM APIs. They are included to demonstrate a possible method for error handling.

XDS and XOM API functions return a status code. In **example.c**, *error* contains the status code for XDS API functions. If the call is successful, the function returns **DS_SUCCESS**. Otherwise, one of the error codes described in Chapter 11 is returned.

The *return_code* variable contains the status code for XOM API functions. If the call is successful, the function returns **OM_SUCCESS**. Otherwise, one of the error codes described in Chapter 18 is returned.

The contents of **example.h** are as follows:

```
/*
 * define some convenient exit codes
 */

#define EXIT_FAILURE 1
#define EXIT_SUCCESS 0

/*
 * declare an error handling function and
 * an error checking macro for DS
 */

void handle_ds_error(DS_status error);
```

```
#define CHECK_DS_CALL(function_call)
    error = (function_call);
    if (error != DS_SUCCESS)
        handle_ds_error(error);

/*
 * declare an error handling function and
 * an error checking macro for OM
 */

void handle_om_error(OM_return_code return_code);

#define CHECK_OM_CALL(function_call)
    return_code = (function_call);
    if (return_code != OM_SUCCESS)
        handle_om_error(return_code);

/*
 * the error handling code
 *
 * NOTE: any errors arising in these functions are ignored.
 */

void handle_ds_error(DS_status error)
{
    (void) fprintf(stderr, "DS error has occurred\n");

    (void) om_delete((OM_object) error);

/* At this point, the error has been reported and storage cleaned up,
 * so the handler could return to the main program now for it to take
 * recovery action. But we choose the simple option ...
 */

    exit(EXIT_FAILURE);
}

void handle_om_error(OM_return_code return_code)
{
```

```
(void) fprintf(stderr, "OM error %d has occurred\n", return_code);

/* At this point, the error has been reported and storage cleaned up,
 * so the handler could return to the main program now for it to take
 * recovery action. But we choose the simple option ...
 */

exit(EXIT_FAILURE);
}
```

7.3 The acl.c Program

The **acl.c** file is a program that displays the ACLs on each entry in the directory for a specific user. The permissions are presented in a form similar to UNIX file permissions. In addition, each entry is flagged as either a master or a shadow copy.

The distinguished name of the user requesting the access permissions is **/C=de/O=sni/OU=ap/CN=norbert**. The results of the request are presented in the following format:

[ABCD] *<entry's distinguished name>*

where:

- A** is one of the following:
- **m** (master copy)
 - **s** (shadow copy)
- B** is one of the following:
- **r** (read access to public attributes)
 - **w** (write access to public attributes)
 - **-** (no access to public attributes)
- C** is one of the following:
- **r** (read access to standard attributes)

- **w** (write access to standard attributes)
- **-** (no access to standard attributes))

D is one of the following:

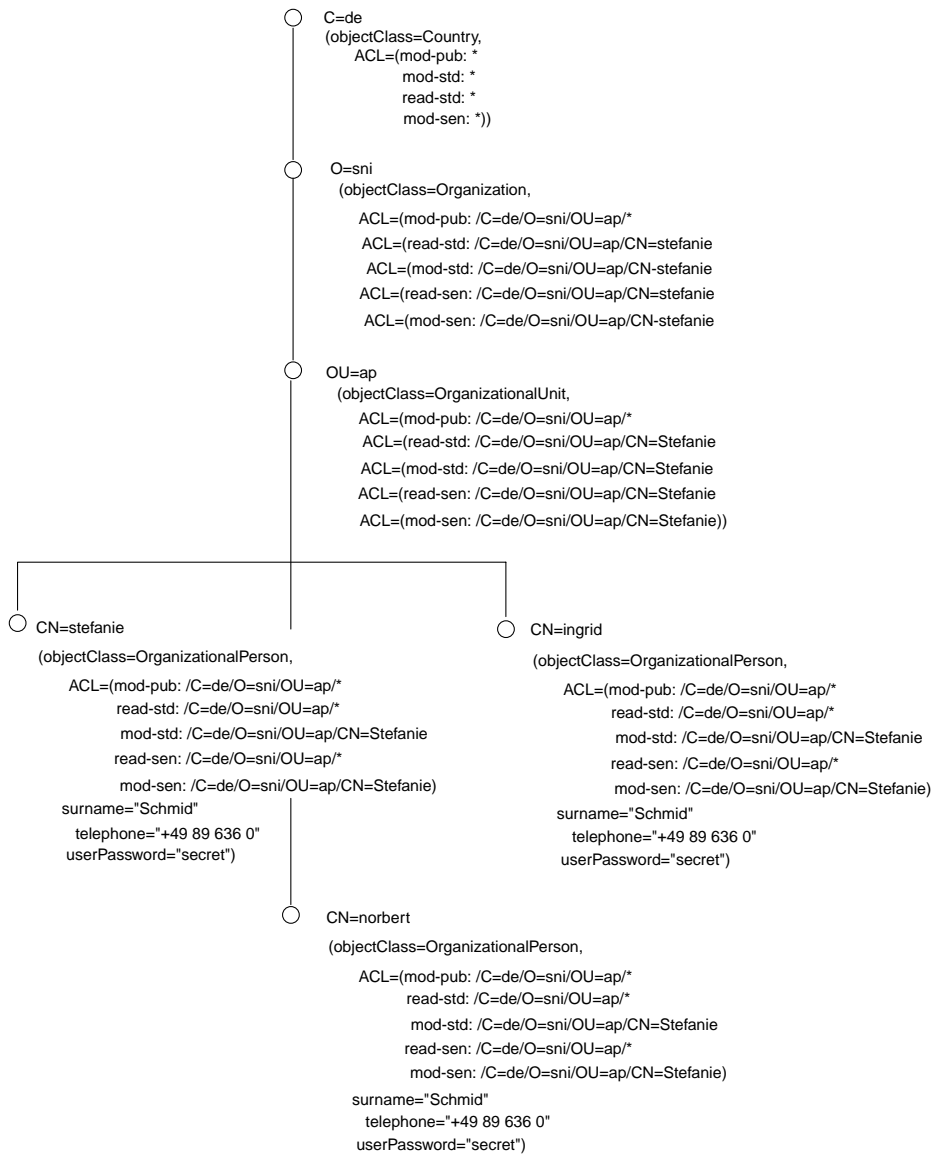
- **r** (read access to sensitive attributes)
- **w** (write access to sensitive attributes)
- **-** (no access to sensitive attributes)

For example, the following result means that the entry **/C=de/O=sni** is a master copy, and that the user who is making the request (**/C=de/O=sni/OU=ap/CN=norbert**) has write access to its public attributes, read access to its standard attributes, and no access to its sensitive attributes:

[mwr-] /C=de/O=sni

The program requires that the user perform an authenticated bind to the directory service. The user's credentials must already exist in the directory. For this reason, the tree of six entries shown in Figure 7-1 is added to the directory each time the program runs, and is removed again afterward.

Figure 7–1. Entries With User Credentials Added to the Directory Tree



The program consists of the following steps:

1. Export the required object identifiers (see **acl.h** in Section 7.3.2).
2. Build the descriptor lists for objects required by the program (see **acl.h** in Section 7.3.2).
3. Initialize a workspace.
4. Negotiate use of the basic directory contents and GDS packages.
5. Add a fixed tree of entries to the directory to permit an authenticated bind.
6. Create a default session object.
7. Alter the default session object to include the credentials of the requestor (**/C=de/O=sni/OU=ap/CN=norbert**).
8. Bind with credentials to the default GDS server.
9. Create a default context object and alter it to include shadow entries.
10. Search the whole subtree below **root** and extract the ACL attribute from each selected entry.
11. Close the connection to the GDS server.
12. Remove the user's credentials from the directory.
13. Extract the components from the search result.
14. Examine each entry and print the entry details.
15. Close the XDS workspace.

Step 1 through Step 4, Step 6 through Step 8, Step 12, and Step 15 are similar to those performed for the previous sample application **example.c**.

Step 5 is included so that the appropriate entries will exist in the directory when the program attempts to access the access permissions.

The default session object created in Step 9 uses **om_create()** to create an instance of a default session object, and it uses **om_put()** to put in the appropriate user credentials. The *credentials* parameter is a descriptor list defined in **acl.h** header file.

Step 10 used the same method as Step 9 to alter the default context to include shadow entries. Using **om_create()** and **om_put()**, the OM attribute **DS_DONT_USE_COPY** is set to a value of **OM_FALSE** to indicate that copies of entries maintained in other DSAs and copies cached locally (that is, shadow copies)

can be used. The *use_copy* parameter is a descriptor list defined in the **acl.h** header file.

Step 11 uses **ds_search()** to search the subtree below **root** to find and extract the ACL attributes from the selected entries defined in the *selection_acl* parameter. The *selection_acl* variable is a descriptor list defined in **acl.h**. The results are returned to the workspace in *result*.

Step 13 and Step 14 extract the components from *result* and examine each entry by using a series of **om_get()** calls, as described in the previous section for **example.c**.

7.3.1 The acl.c Code

The following code is a listing of the **acl.c** program.

```
/*
 *
 *  COPYRIGHT (C) SIEMENS NIXDORF INFORMATIONSSYSTEME AG 1991
 *
 *      ALL RIGHTS RESERVED
 *
 */
/*****

/*
 *
 *  This sample program displays the access permissions (ACL) on each
 *  entry in the directory for a specific user. The permissions are
 *  presented in a form similar to the UNIX file permissions.
 *  In addition, each entry is flagged as either a master
 *  or a shadow copy.
 *
 *  The distinguished name of the user performing the check is:
 *
 *      /C=de/O=sni/OU=ap/CN=norbert
 *
 *  The results are presented in the following format :
 *
 *      [ABCD] <entry's distinguished name>
 *
 */
```



```

*      A:  'm' master copy
*          's' shadow copy
*
*      B:  'r' read access to public attributes
*          'w' write access to public attributes
*          '-' no access to public attributes
*
*      C:  'r' read access to standard attributes
*          'w' write access to standard attributes
*          '-' no access to standard attributes
*
*      D:  'r' read access to sensitive attributes
*          'w' write access to sensitive attributes
*          '-' no access to sensitive attributes
*
* For example, the following result means that the entry '/C=de/O=sni'
* is a master copy and that the requesting user
* (/C=de/O=sni/OU=ap/CN=norbert) has write access to its public
* attributes, read access to its standard
* attributes and no access to its sensitive attributes.
*
*      [mwr-] /C=de/O=sni
*
* The program requires that the specific user perform an authenticated
* bind to the directory. In order to achieve this the user's
* credentials must already exist in the directory.
* Therefore the following tree of 6 entries is added to the directory
* each time the program runs, and removed again afterwards.
*
*      O C=de
*      | (objectClass=Country,
*      |   ACL=(mod-pub: *
*      |         read-std:*
*      |         mod-std: *
*      |         read-sen:*
*      |         mod-sen: *))
*      |
*      |
*      O O=sni
*      | (objectClass=Organization,

```

```

* | ACL=(mod-pub: /C=de/O=sni/OU=ap/*
* | read-std: /C=de/O=sni/OU=ap/CN=stefanie
* | mod-std: /C=de/O=sni/OU=ap/CN=stefanie
* | read-sen: /C=de/O=sni/OU=ap/CN=stefanie
* | mod-sen: /C=de/O=sni/OU=ap/CN=stefanie))
* |
* | O OU=ap
* | (objectClass=OrganizationalUnit,
* | ACL=(mod-pub: /C=de/O=sni/OU=ap/*
* | read-std: /C=de/O=sni/OU=ap/CN=stefanie
* | mod-std: /C=de/O=sni/OU=ap/CN=stefanie
* | read-sen: /C=de/O=sni/OU=ap/CN=stefanie
* | mod-sen: /C=de/O=sni/OU=ap/CN=stefanie))
* |
* |
* | +-----+-----+
* | | |
* | | O CN=ingrid
* | | (objectClass=OrganizationalPerson,
* | | ACL=(mod-pub: /C=de/O=sni/OU=ap/*
* | | read-std: /C=de/O=sni/OU=ap/*
* | | mod-std: /C=de/O=sni/OU=ap/CN=stefanie
* | | read-sen: /C=de/O=sni/OU=ap/*
* | | mod-sen: /C=de/O=sni/OU=ap/CN=stefanie),
* | | surname="Schmid",
* | | telephone="+49 89 636 0",
* | | userPassword="secret")
* | |
* | | O CN=norbert
* | | (objectClass=OrganizationalPerson,
* | | ACL=(mod-pub: /C=de/O=sni/OU=ap/*
* | | read-std: /C=de/O=sni/OU=ap/*
* | | mod-std: /C=de/O=sni/OU=ap/CN=stefanie
* | | read-sen: /C=de/O=sni/OU=ap/*
* | | mod-sen: /C=de/O=sni/OU=ap/CN=stefanie),
* | | surname="Schmid",
* | | telephone="+49 89 636 0",
* | | userPassword="secret")
* | |
* | O CN=stefanie

```

```

*      (objectClass=OrganizationalPerson,
*      ACL=(mod-pub: /C=de/O=sni/OU=ap/*
*      read-std: /C=de/O=sni/OU=ap/*
*      mod-std: /C=de/O=sni/OU=ap/CN=stefanie
*      read-sen: /C=de/O=sni/OU=ap/*
*      mod-sen: /C=de/O=sni/OU=ap/CN=stefanie),
*      surname="Schmid",
*      telephone="+49 89 636 0",
*      userPassword="secret")
*
*
*/

#ifdef THREADSAFE
#include <pthread.h>
#endif

#include <xom.h>
#include <xds.h>
#include <xdsbdc.h>
#include <xdsbdc.h>
#include <xdsbdc.h>
#include <xdsbdc.h>
#include "acl.h"      /* static initialization of data structures. */

void
main(
    int  argc,
    char *argv[]
)
{
    OM_workspace      workspace;      /* workspace for objects      */
    OM_private_object session;        /* Session object.            */
    OM_private_object bound_session; /* Holds the Session object which
                                     /* is returned by ds_bind()    */
    OM_private_object context;        /* Context object.            */
    OM_private_object result;         /* Holds the search result object.
                                     /* returned by ds_search().    */
    OM_sint           invoke_id;      /* Integer for the invoke id
                                     /* (this parameter must be present
                                     /* even though it is ignored).

```

```
OM_type      sinfo_list[] = { DS_SEARCH_INFO, 0 };
OM_type      entry_list[] = { DS_ENTRIES, 0 };
                                /* Lists of types to be extracted */
OM_public_object sinfo;        /* Search-Info object from result. */
OM_public_object entry;        /* Entry object from search info. */
OM_value_position total_num;    /* Number of descriptors returned. */
OM_return_code rc;             /* XOM function return code. */
register int i;
char          user_name[MAX_DN_LEN];
                                /* Holds requestor's name. */
char          entry_string[MAX_DN_LEN + 7] = "[?r??] ";
                                /* Holds entry details. */

/* Step 3 (see acl.h program code for Steps 1 and 2)
 *
 * Initialise a directory workspace for use by XOM.
 */
if ((workspace = ds_initialize()) == (OM_workspace)0)
    printf("ds_initialize() error\n");

/* Step 4
 *
 * Negotiate the use of the BDCP and GDS packages.
 */
if (ds_version(features, workspace) != DS_SUCCESS)
    printf("ds_version() error\n");

/* Step 5
 *
 * Add a fixed tree of entries to the directory in order to permit
 * an authenticated bind by: /C=de/O=sni/OU=ap/CN=norbert
 */
if (! add_tree(workspace))
    printf("add_tree() error\n");

/* Step 6
 *
 * Create a default session object.
 */
```

```
if ((rc = om_create(DSX_C_GDS_SESSION,OM_TRUE,workspace,&session))
    != OM_SUCCESS)
    printf("om_create() error %d\n", rc);

/* Step 7
 *
 * Alter the default session object to include the following
 * credentials: requestor: /C=de/O=sni/OU=ap/CN=norbert
 * password: "secret"
 * authentication mechanism: simple
 */
if ((rc = om_put(session, OM_REPLACE_ALL, credentials, 0 ,0, 0))
    != OM_SUCCESS)
    printf("om_put() error %d\n", rc);

/* Step 8
 *
 * Bind with credentials to the default GDS server. The returned
 * session object is stored in the private object variable
 * bound_session and is used for all further XDS function calls.
 */
if (ds_bind(session, workspace, &bound_session) != DS_SUCCESS)
    printf("ds_bind() error\n");

/* Step 9
 *
 * Create a default context object.
 */
if ((rc = om_create(DSX_C_GDS_CONTEXT,OM_TRUE,workspace,&context))
    != OM_SUCCESS)
    printf("om_create() error %d\n", rc);

/*
 * Alter the default context object to include 'shadow' entries.
 */
if ((rc = om_put(context, OM_REPLACE_ALL, use_copy, 0 ,0, 0))
    != OM_SUCCESS)
    printf("om_put() error %d\n", rc);

/* Step 10
```

```
*
* Search the whole subtree below root. The filter selects
* entries with an object-class attribute. The selection
* extracts the ACL attribute from each selected entry.
* The results are returned in the private object 'result'.
*
* NOTE: Since every entry contains an object-class attribute the
*       filter performs no function other than to demonstrate how
*       filters may be used.
*/
if (ds_search(bound_session, context, dn_root, DS_WHOLE_SUBTREE,
             filter, OM_FALSE, selection_acl, &result, &invoke_id)
    != DS_SUCCESS)
    printf("ds_search() error\n");

/* Step 11
*
* Close the connection to the GDS server.
*/
if (ds_unbind(bound_session) != DS_SUCCESS)
    printf("ds_unbind() error\n");

/* Step 12
*
* Remove the user's credentials from the directory.
*/
if (!remove_tree(workspace, session))
    printf("remove_tree() error\n");

/* Step 13
*
* Extract components from the search result by means of om_get().
*/
if ((rc = om_get(result,
                 OM_EXCLUDE_ALL_BUT_THESE_TYPES + OM_EXCLUDE_SUBOBJECTS,
                 sinfo_list, OM_FALSE, 0, 0, &sinfo, &total_num))
    != OM_SUCCESS)
    printf("om_get(Search-Result) error %d\n", rc);
```

```

if ((rc = om_get(sinfo->value.object.object,
                OM_EXCLUDE_ALL_BUT_THESE_TYPES + OM_EXCLUDE_SUBOBJECTS,
                entry_list, OM_FALSE, 0, 0, &entry, &total_num))
    != OM_SUCCESS)
    printf("om_get(Search-Info) error %d\n", rc);

/*
 * Convert the requestor's distinguished name to string format.
 */
if (! xds_name_to_string(dn_norbert, user_name))
    printf("xds_name_to_string() error\n");

printf("User:  %s\nTotal: %d\n", user_name, total_num);

/* Step 14
 *
 * Examine each entry and print the entry details.
 */
for (i = 0; i < total_num; i++) {
    if (process_entry_info((entry+i)->value.object.object,
                          entry_string, user_name))
        printf("%s\n", entry_string);
}

/* Step 15
 *
 * Close the directory workspace.
 */
if (ds_shutdown(workspace) != DS_SUCCESS)
    printf("ds_shutdown() error\n");
}

/*
 * Add the tree of entries described above.
 */
int
add_tree(
    OM_workspace workspace
)
{

```

```
OM_private_object session;    /* Holds the Session object which */
                               /* is returned by ds_bind()      */
OM_sint             invoke_id; /* Integer for the invoke id  */
int                 error = 0;

/* Bind (without credentials) to the default GDS server.          */

if (ds_bind(DS_DEFAULT_SESSION, workspace, &session) != DS_SUCCESS)
    error++;

/* Add entries to the GDS server.                                  */

ds_add_entry(session, DS_DEFAULT_CONTEXT, dn_de, alist_C,
              &invoke_id);

if (ds_add_entry(session, DS_DEFAULT_CONTEXT, dn_sni, alist_O,
                  &invoke_id) != DS_SUCCESS)
    error++;

if (ds_add_entry(session, DS_DEFAULT_CONTEXT, dn_ap, alist_OU,
                  &invoke_id) != DS_SUCCESS)
    error++;

if (ds_add_entry(session, DS_DEFAULT_CONTEXT, dn_stefanie, alist_OP,
                  &invoke_id) != DS_SUCCESS)
    error++;

if (ds_add_entry(session, DS_DEFAULT_CONTEXT, dn_norbert, alist_OP,
                  &invoke_id) != DS_SUCCESS)
    error++;

if (ds_add_entry(session, DS_DEFAULT_CONTEXT, dn_ingrid, alist_OP,
                  &invoke_id) != DS_SUCCESS)
    error++;

/* Close the connection to the GDS server.                        */

if (ds_unbind(session) != DS_SUCCESS)
    error++;
```



```

        return (error?0:1);
    }

    /*
     * Remove the tree of entries described above.
     */
    int
    remove_tree(
        OM_workspace workspace,
        OM_private_object session
    )
    {
        OM_private_object bound_session; /* Holds Session object which */
                                         /* is returned by ds_bind() */
        OM_sint          invoke_id;      /* Integer for the invoke id */
        int               error = 0;

        /* Bind (with credentials) to the default GDS server. */

        if (ds_bind(session, workspace, &bound_session) != DS_SUCCESS)
            error++;

        /* Remove entries from the GDS server. */

        if (ds_remove_entry(bound_session, DS_DEFAULT_CONTEXT, dn_ingrid,
                             &invoke_id) != DS_SUCCESS)
            error++;

        if (ds_remove_entry(bound_session, DS_DEFAULT_CONTEXT, dn_stefanie,
                             &invoke_id) != DS_SUCCESS)
            error++;

        if (ds_remove_entry(bound_session, DS_DEFAULT_CONTEXT, dn_norbert,
                             &invoke_id) != DS_SUCCESS)
            error++;

        if (ds_remove_entry(bound_session, DS_DEFAULT_CONTEXT, dn_ap,
                             &invoke_id) != DS_SUCCESS)
            error++;
    }

```

```
        if (ds_remove_entry(bound_session, DS_DEFAULT_CONTEXT, dn_sni,
                            &invoke_id) != DS_SUCCESS)

            error++;

        ds_remove_entry(bound_session, DS_DEFAULT_CONTEXT, dn_de,
                        &invoke_id);

/* Close the connection to the GDS server.                                */

        if (ds_unbind(bound_session) != DS_SUCCESS)
            error++;

        return (error?0:1);
    }

/*
 * Convert a distinguished name in XDS format (OM_descriptor
 * lists) to string format.
 */
int
xds_name_to_string(
    OM_public_object  name,      /* Xds distinguished name.      */
    char              *string    /* String distinguished name.    */
)
{
    register OM_object  dn = name;
    register OM_object  rdn;
    register OM_object  ava;
    register char       *sp = string;
    int                 error = 0;

    while ((dn->type != OM_NO_MORE_TYPES) && (! error)) {
        if ((dn->type == DS_RDNS) &&
            ((dn->syntax & OM_S_SYNTAX) == OM_S_OBJECT)) {
            rdn = dn->value.object.object;

            while ((rdn->type != OM_NO_MORE_TYPES) && (! error)) {
                if ((rdn->type == DS_AVAS) &&
                    ((rdn->syntax & OM_S_SYNTAX) == OM_S_OBJECT)) {
                    ava = rdn->value.object.object;
```

```
while ((ava->type != OM_NO_MORE_TYPES) &&
      (! error)) {
    if ((ava->type == DS_ATTRIBUTE_TYPE) &&
        ((ava->syntax & OM_S_SYNTAX) ==
         OM_S_OBJECT_IDENTIFIER_STRING)) {

        *sp++ = '/';
        if (strncmp(ava->value.string.elements,
                    DS_A_COUNTRY_NAME.elements,
                    ava->value.string.length) == 0)
            *sp++ = 'C';

        else if (strncmp(ava->value.string.elements,
                        DS_A_ORG_NAME.elements,
                        ava->value.string.length) == 0)
            *sp++ = 'O';

        else if (strncmp(ava->value.string.elements,
                        DS_A_ORG_UNIT_NAME.elements,
                        ava->value.string.length) == 0)
            *sp++ = 'O', *sp++ = 'U';

        else if (strncmp(ava->value.string.elements,
                        DS_A_COMMON_NAME.elements,
                        ava->value.string.length) == 0)
            *sp++ = 'C', *sp++ = 'N';

        else if (strncmp(ava->value.string.elements,
                        DS_A_LOCALITY_NAME.elements,
                        ava->value.string.length) == 0)
            *sp++ = 'L';

        else if (strncmp(ava->value.string.elements,
                        DSX_TYPELESS_RDN.elements,
                        ava->value.string.length) != 0) {
            error++;
            continue;
        }
    }
}
```

```
        if (*(sp-1) != '/'); /* no '=' if typeless*/
        *sp++ = '=';
    }
    if (ava->type == DS_ATTRIBUTE_VALUES) {
        switch(ava->syntax & OM_S_SYNTAX) {
            case OM_S_PRINTABLE_STRING :
            case OM_S_TELETEX_STRING :
                strncpy(sp, ava->value.string.elements,
                    ava->value.string.length);
                sp += ava->value.string.length;
                break;

            default:
                error++;
                continue;
        }
    }
    ava++;
}
}
rdn++;
}
}
dn++;
}
*sp = '\0';

return (error?0:1);
}

/*
 * Extract information about an entry from the Entry-Info object:
 * whether the entry is a master-copy, its ACL permissions and
 * its distinguished name.
 * Build up a string based on this information.
 */
int
process_entry_info(
    OM_private_object entry,
```

```

char          *entry_string,
char          *user_name
)
{
    OM_return_code    rc;          /* Return code from XOM function. */
    OM_public_object  ei_attrs;    /* Components from Entry-Info.   */
    OM_public_object  attr;        /* Directory attribute.           */
    OM_public_object  acl;         /* ACL attribute value.           */
    OM_public_object  acl_item;    /* ACL item component.            */
    OM_value_position total_attrs; /* Number of attributes returned. */
    register int      i;
    register int      interp;
    register int      error = 0;
    register int      found_acl = 0;
    static OM_type     ei_attr_list[] = { DS_FROM_ENTRY,
                                           DS_OBJECT_NAME,
                                           DS_ATTRIBUTES,
                                           0 };
                                           /* Attributes to be extracted. */

    /*
     * Extract three attributes from each Entry-Info object.
     */
    if ((rc = om_get(entry, OM_EXCLUDE_ALL_BUT_THESE_TYPES,
                    ei_attr_list, OM_FALSE, 0, 0, &ei_attrs, &total_attrs))
        != OM_SUCCESS) {
        error++;
        printf("om_get(Entry-Info) error %d, rc);\n");
    }

    for (i = 0; ((i < total_attrs) && (! error)); i++, ei_attrs++) {

        /*
         * Determine if current entry is a master-copy or a shadow-copy.
         */
        if ((ei_attrs->type == DS_FROM_ENTRY) &&
            ((ei_attrs->syntax & OM_S_SYNTAX) == OM_S_BOOLEAN))
            if (ei_attrs->value.boolean == OM_TRUE)
                entry_string[1] = 'm';
            else if (ei_attrs->value.boolean == OM_FALSE)

```

```
        entry_string[1] = 's';
    else
        entry_string[1] = '?';

    if ((ei_attrs->type == DS_ATTRIBUTES) &&
        ((ei_attrs->syntax & OM_S_SYNTAX) == OM_S_OBJECT)) {
        attr = ei_attrs->value.object.object;

        while ((attr->type != OM_NO_MORE_TYPES) && (! error)) {

            /*
             * Check that the attribute is an ACL attribute.
             */
            if ((attr->type == DS_ATTRIBUTE_TYPE) &&
                ((attr->syntax & OM_S_SYNTAX) ==
                 OM_S_OBJECT_IDENTIFIER_STRING)) {
                if (strncmp(attr->value.string.elements,
                            DSX_A_ACL.elements,
                            attr->value.string.length) == 0)
                    found_acl++;
            }

            /*
             * Examine the ACL. Check each permission for
             * the current user.
             */
            if ((found_acl) &&
                (attr->type == DS_ATTRIBUTE_VALUES) &&
                ((attr->syntax & OM_S_SYNTAX) == OM_S_OBJECT)) {

                acl = attr->value.object.object;

                entry_string[2] = 'r';
                entry_string[3] = '-';
                entry_string[4] = '-';

                while (acl->type != OM_NO_MORE_TYPES) {

                    if ((acl->syntax & OM_S_SYNTAX) == OM_S_OBJECT)
                        acl_item = acl->value.object.object;
```

```
switch (acl->type) {

case OM_CLASS:
    break;

case DSX_MODIFY_PUBLIC:
    if (permitted_access(user_name, acl_item))
        entry_string[2] = 'w';
    break;

case DSX_READ_STANDARD:
    if (permitted_access(user_name, acl_item))
        entry_string[3] = 'r';
    break;

case DSX_MODIFY_STANDARD:
    if (permitted_access(user_name, acl_item))
        entry_string[3] = 'w';
    break;

case DSX_READ_SENSITIVE:
    if (permitted_access(user_name, acl_item))
        entry_string[4] = 'r';
    break;

case DSX_MODIFY_SENSITIVE:
    if (permitted_access(user_name, acl_item))
        entry_string[4] = 'w';
    break;
}
acl++;
    }
}
attr++;
}

/*
 * Convert the entry's distinguished name to a string format.
```

```
        */
        if ((ei_attrs->type == DS_OBJECT_NAME) &&
            ((ei_attrs->syntax & OM_S_SYNTAX) == OM_S_OBJECT))
            if (! xds_name_to_string(ei_attrs->value.object.object,
                                    &entry_string[7])) {
                error++;
                printf("xds_name_to_string() error\n");
            }
    }

    return (error?0:1);
}

/*
 * Check if a user is permitted access based on the ACL supplied.
 *
 */
int
permitted_access(
    char          *user_name,
    OM_public_object acl_item
)
{
    char acl_name[MAX_DN_LEN];
    int interpretation;
    int acl_present = 0;
    int access = 0;
    int acl_name_length;

    while (acl_item->type != OM_NO_MORE_TYPES) {

        switch (acl_item->type) {
            case OM_CLASS:
                break;

            case DSX_INTERPRETATION:
                interpretation = acl_item->value.boolean;
                break;
        }
    }
}
```



```

    case DSX_USER:
        xds_name_to_string(acl_item->value.object.object, acl_name);

        if (interpretation == DSX_SINGLE_OBJECT) {
            if (strcmp(acl_name, user_name) == 0)
                access = 1;
        }
        else if (interpretation == DSX_ROOT_OF_SUBTREE) {
            if ((acl_name_length = strlen(acl_name)) == 0)
                access = 1;
            else if (strncmp(acl_name, user_name,
                            acl_name_length) == 0)
                access = 1;
        }
        break;
    }
    acl_item++;
}

return (access);
}

```

7.3.2 The **acl.h** Header File

The **acl.h** header file performs the following:

1. It exports the object identifiers that **acl.c** requires.
2. It builds the descriptor lists for the following distinguished names:

```

root
C=de
C=de/O=sni
C=de/O=sni/OU=ap
C=de/O=sni/OU=ap/CN=stefanie
C=de/O=sni/OU=ap/CN=norbert
C=de/O=sni/OU=ap/CN=ingrid

```

3. It builds the object identifiers for attributes to be added to the directory.
4. It builds a descriptor list for the attribute types and values that are to be selected.
5. It builds the descriptor list for bind credentials.
6. It builds the descriptor list for context.
7. It builds the descriptor list for optional packages that are to be negotiated.
8. It builds the descriptor list for search filters.

7.3.3 The **acl.h** Code

The following code is a listing of the **acl.h** file:

```
/*
 *
 *  COPYRIGHT (C) SIEMENS NIXDORF INFORMATIONSSYSTEME AG 1991
 *
 *      ALL RIGHTS RESERVED
 *
 */
*****/

#ifndef ACL_HEADER
#define ACL_HEADER

#define MAX_DN_LEN 100
/* max length of a distinguished name in string format*/

/* The application must export the object identifiers it requires. */

OM_EXPORT (DS_C_AVA)
OM_EXPORT (DS_C_DS_RDN)
OM_EXPORT (DS_C_DS_DN)
OM_EXPORT (DS_C_ENTRY_INFO_SELECTION)
OM_EXPORT (DS_C_ATTRIBUTE)
OM_EXPORT (DS_C_ATTRIBUTE_LIST)
OM_EXPORT (DS_C_FILTER)
OM_EXPORT (DS_C_FILTER_ITEM)
OM_EXPORT (DSX_C_GDS_SESSION)
OM_EXPORT (DSX_C_GDS_CONTEXT)
```

```

OM_EXPORT (DSX_C_GDS_ACL)
OM_EXPORT (DSX_C_GDS_ACL_ITEM)

OM_EXPORT (DS_A_COUNTRY_NAME)
OM_EXPORT (DS_A_ORG_NAME)
OM_EXPORT (DS_A_ORG_UNIT_NAME)
OM_EXPORT (DS_A_COMMON_NAME)
OM_EXPORT (DS_A_LOCALITY_NAME)
OM_EXPORT (DS_A_OBJECT_CLASS)
OM_EXPORT (DS_A_USER_PASSWORD)
OM_EXPORT (DS_A_PHONE_NBR)
OM_EXPORT (DS_A_SURNAME)
OM_EXPORT (DSX_A_ACL)
OM_EXPORT (DSX_TYPELESS_RDN)

OM_EXPORT (DS_O_TOP)
OM_EXPORT (DS_O_COUNTRY)
OM_EXPORT (DS_O_ORG)
OM_EXPORT (DS_O_ORG_UNIT)
OM_EXPORT (DS_O_PERSON)
OM_EXPORT (DS_O_ORG_PERSON)

/* Build up descriptor lists for the following distinguished names: */
/*      root                                                         */
/*      /C=de                                                         */
/*      /C=de/O=sni                                                   */
/*      /C=de/O=sni/OU=ap                                             */
/*      /C=de/O=sni/OU=ap/CN=stefanie                                */
/*      /C=de/O=sni/OU=ap/CN=norbert                                 */
/*      /C=de/O=sni/OU=ap/CN=ingrid                                  */

static OM_descriptor   ava_de[] = {
    OM_OID_DESC(OM_CLASS, DS_C_AVA),
    OM_OID_DESC(DS_ATTRIBUTE_TYPE, DS_A_COUNTRY_NAME),
    {DS_ATTRIBUTE_VALUES, OM_S_PRINTABLE_STRING, OM_STRING("de")},
    OM_NULL_DESCRIPTOR
};

static OM_descriptor   ava_sni[] = {
    OM_OID_DESC(OM_CLASS, DS_C_AVA),

```

```
    OM_OID_DESC(DS_ATTRIBUTE_TYPE, DS_A_ORG_NAME),
    {DS_ATTRIBUTE_VALUES, OM_S_TELETEX_STRING, OM_STRING("sni")},
    OM_NULL_DESCRIPTOR
};

static OM_descriptor    ava_ap[] = {
    OM_OID_DESC(OM_CLASS, DS_C_AVA),
    OM_OID_DESC(DS_ATTRIBUTE_TYPE, DS_A_ORG_UNIT_NAME),
    {DS_ATTRIBUTE_VALUES, OM_S_TELETEX_STRING, OM_STRING("ap")},
    OM_NULL_DESCRIPTOR
};

static OM_descriptor    ava_stefanie[] = {
    OM_OID_DESC(OM_CLASS, DS_C_AVA),
    OM_OID_DESC(DS_ATTRIBUTE_TYPE, DS_A_COMMON_NAME),
    {DS_ATTRIBUTE_VALUES, OM_S_TELETEX_STRING, OM_STRING("stefanie")},
    OM_NULL_DESCRIPTOR
};

static OM_descriptor    ava_norbert[] = {
    OM_OID_DESC(OM_CLASS, DS_C_AVA),
    OM_OID_DESC(DS_ATTRIBUTE_TYPE, DS_A_COMMON_NAME),
    {DS_ATTRIBUTE_VALUES, OM_S_TELETEX_STRING, OM_STRING("norbert")},
    OM_NULL_DESCRIPTOR
};

static OM_descriptor    ava_ingrid[] = {
    OM_OID_DESC(OM_CLASS, DS_C_AVA),
    OM_OID_DESC(DS_ATTRIBUTE_TYPE, DS_A_COMMON_NAME),
    {DS_ATTRIBUTE_VALUES, OM_S_TELETEX_STRING, OM_STRING("ingrid")},
    OM_NULL_DESCRIPTOR
};

static OM_descriptor    rdn_de[] = {
    OM_OID_DESC(OM_CLASS, DS_C_DS_RDN),
    {DS_AVAS, OM_S_OBJECT, {0, ava_de}},
    OM_NULL_DESCRIPTOR
};

static OM_descriptor    rdn_sni[] = {
    OM_OID_DESC(OM_CLASS, DS_C_DS_RDN),
    {DS_AVAS, OM_S_OBJECT, {0, ava_sni}},
    OM_NULL_DESCRIPTOR
};

static OM_descriptor    rdn_ap[] = {
    OM_OID_DESC(OM_CLASS, DS_C_DS_RDN),
```

```

        {DS_AVAS, OM_S_OBJECT, {0, ava_ap}},
        OM_NULL_DESCRIPTOR
};

static OM_descriptor    rdn_stefanie[] = {
        OM_OID_DESC(OM_CLASS, DS_C_DS_RDN),
        {DS_AVAS, OM_S_OBJECT, {0, ava_stefanie}},
        OM_NULL_DESCRIPTOR
};

static OM_descriptor    rdn_norbert[] = {
        OM_OID_DESC(OM_CLASS, DS_C_DS_RDN),
        {DS_AVAS, OM_S_OBJECT, {0, ava_norbert}},
        OM_NULL_DESCRIPTOR
};

static OM_descriptor    rdn_ingrid[] = {
        OM_OID_DESC(OM_CLASS, DS_C_DS_RDN),
        {DS_AVAS, OM_S_OBJECT, {0, ava_ingrid}},
        OM_NULL_DESCRIPTOR
};

static OM_descriptor    dn_root[] = {
        OM_OID_DESC(OM_CLASS, DS_C_DS_DN),
        OM_NULL_DESCRIPTOR
};

static OM_descriptor    dn_de[] = {
        OM_OID_DESC(OM_CLASS, DS_C_DS_DN),
        {DS_RDNS, OM_S_OBJECT, {0, rdn_de}},
        OM_NULL_DESCRIPTOR
};

static OM_descriptor    dn_sni[] = {
        OM_OID_DESC(OM_CLASS, DS_C_DS_DN),
        {DS_RDNS, OM_S_OBJECT, {0, rdn_de}},
        {DS_RDNS, OM_S_OBJECT, {0, rdn_sni}},
        OM_NULL_DESCRIPTOR
};

static OM_descriptor    dn_ap[] = {
        OM_OID_DESC(OM_CLASS, DS_C_DS_DN),
        {DS_RDNS, OM_S_OBJECT, {0, rdn_de}},
        {DS_RDNS, OM_S_OBJECT, {0, rdn_sni}},
        {DS_RDNS, OM_S_OBJECT, {0, rdn_ap}},
        OM_NULL_DESCRIPTOR
};

```

```
static OM_descriptor dn_stefanie[] = {
    OM_OID_DESC(OM_CLASS, DS_C_DS_DN),
    {DS_RDNS, OM_S_OBJECT, {0, rdn_de}},
    {DS_RDNS, OM_S_OBJECT, {0, rdn_sni}},
    {DS_RDNS, OM_S_OBJECT, {0, rdn_ap}},
    {DS_RDNS, OM_S_OBJECT, {0, rdn_stefanie}},
    OM_NULL_DESCRIPTOR
};

static OM_descriptor dn_norbert[] = {
    OM_OID_DESC(OM_CLASS, DS_C_DS_DN),
    {DS_RDNS, OM_S_OBJECT, {0, rdn_de}},
    {DS_RDNS, OM_S_OBJECT, {0, rdn_sni}},
    {DS_RDNS, OM_S_OBJECT, {0, rdn_ap}},
    {DS_RDNS, OM_S_OBJECT, {0, rdn_norbert}},
    OM_NULL_DESCRIPTOR
};

static OM_descriptor dn_ingrid[] = {
    OM_OID_DESC(OM_CLASS, DS_C_DS_DN),
    {DS_RDNS, OM_S_OBJECT, {0, rdn_de}},
    {DS_RDNS, OM_S_OBJECT, {0, rdn_sni}},
    {DS_RDNS, OM_S_OBJECT, {0, rdn_ap}},
    {DS_RDNS, OM_S_OBJECT, {0, rdn_ingrid}},
    OM_NULL_DESCRIPTOR
};

/* Build up an array of object identifiers for the attributes to be */
/* added to the directory.                                          */

static OM_descriptor  obj_class_C[] = {
    OM_OID_DESC(OM_CLASS, DS_C_ATTRIBUTE),
    OM_OID_DESC(DS_ATTRIBUTE_TYPE, DS_A_OBJECT_CLASS),
    OM_OID_DESC(DS_ATTRIBUTE_VALUES, DS_O_TOP),
    OM_OID_DESC(DS_ATTRIBUTE_VALUES, DS_O_COUNTRY),
    OM_NULL_DESCRIPTOR
};

static OM_descriptor  obj_class_O[] = {
    OM_OID_DESC(OM_CLASS, DS_C_ATTRIBUTE),
    OM_OID_DESC(DS_ATTRIBUTE_TYPE, DS_A_OBJECT_CLASS),
    OM_OID_DESC(DS_ATTRIBUTE_VALUES, DS_O_TOP),
```

```

        OM_OID_DESC(DS_ATTRIBUTE_VALUES, DS_O_ORG),
        OM_NULL_DESCRIPTOR
};

static OM_descriptor    obj_class_OU[] = {
    OM_OID_DESC(OM_CLASS, DS_C_ATTRIBUTE),
    OM_OID_DESC(DS_ATTRIBUTE_TYPE, DS_A_OBJECT_CLASS),
    OM_OID_DESC(DS_ATTRIBUTE_VALUES, DS_O_TOP),
    OM_OID_DESC(DS_ATTRIBUTE_VALUES, DS_O_ORG_UNIT),
    OM_NULL_DESCRIPTOR
};

static OM_descriptor    obj_class_OP[] = {
    OM_OID_DESC(OM_CLASS, DS_C_ATTRIBUTE),
    OM_OID_DESC(DS_ATTRIBUTE_TYPE, DS_A_OBJECT_CLASS),
    OM_OID_DESC(DS_ATTRIBUTE_VALUES, DS_O_TOP),
    OM_OID_DESC(DS_ATTRIBUTE_VALUES, DS_O_PERSON),
    OM_OID_DESC(DS_ATTRIBUTE_VALUES, DS_O_ORG_PERSON),
    OM_NULL_DESCRIPTOR
};

static OM_descriptor    att_phone_num[] = {
    OM_OID_DESC(OM_CLASS, DS_C_ATTRIBUTE),
    OM_OID_DESC(DS_ATTRIBUTE_TYPE, DS_A_PHONE_NBR),
    {DS_ATTRIBUTE_VALUES, OM_S_PRINTABLE_STRING,
     OM_STRING("+49 89 636 0")},
    OM_NULL_DESCRIPTOR
};

static OM_descriptor    att_password[] = {
    OM_OID_DESC(OM_CLASS, DS_C_ATTRIBUTE),
    OM_OID_DESC(DS_ATTRIBUTE_TYPE, DS_A_USER_PASSWORD),
    {DS_ATTRIBUTE_VALUES, OM_S_OCTET_STRING, OM_STRING("secret")},
    OM_NULL_DESCRIPTOR
};

static OM_descriptor    att_surname[] = {
    OM_OID_DESC(OM_CLASS, DS_C_ATTRIBUTE),
    OM_OID_DESC(DS_ATTRIBUTE_TYPE, DS_A_SURNAME),
    {DS_ATTRIBUTE_VALUES, OM_S_TELETEX_STRING, OM_STRING("Schmid")},
    OM_NULL_DESCRIPTOR
};

```

```
static OM_descriptor    acl_item_root[] = {
    OM_OID_DESC(OM_CLASS, DSX_C_GDS_ACL_ITEM),
    {DSX_INTERPRETATION, OM_S_ENUMERATION, {DSX_ROOT_OF_SUBTREE, 0}},
    {DSX_USER, OM_S_OBJECT, {0, dn_root}},
    OM_NULL_DESCRIPTOR
};

static OM_descriptor    acl_item_ap[] = {
    OM_OID_DESC(OM_CLASS, DSX_C_GDS_ACL_ITEM),
    {DSX_INTERPRETATION, OM_S_ENUMERATION, {DSX_ROOT_OF_SUBTREE, 0}},
    {DSX_USER, OM_S_OBJECT, {0, dn_ap}},
    OM_NULL_DESCRIPTOR
};

static OM_descriptor    acl_item_stefanie[] = {
    OM_OID_DESC(OM_CLASS, DSX_C_GDS_ACL_ITEM),
    {DSX_INTERPRETATION, OM_S_ENUMERATION, {DSX_SINGLE_OBJECT, 0}},
    {DSX_USER, OM_S_OBJECT, {0, dn_stefanie}},
    OM_NULL_DESCRIPTOR
};

static OM_descriptor    acl1[] = {
    OM_OID_DESC(OM_CLASS, DSX_C_GDS_ACL),
    {DSX_MODIFY_PUBLIC, OM_S_OBJECT, {0, acl_item_root}},
    {DSX_READ_STANDARD, OM_S_OBJECT, {0, acl_item_stefanie}},
    {DSX_MODIFY_STANDARD, OM_S_OBJECT, {0, acl_item_stefanie}},
    {DSX_READ_SENSITIVE, OM_S_OBJECT, {0, acl_item_stefanie}},
    {DSX_MODIFY_SENSITIVE, OM_S_OBJECT, {0, acl_item_stefanie}},
    OM_NULL_DESCRIPTOR
};

static OM_descriptor    acl2[] = {
    OM_OID_DESC(OM_CLASS, DSX_C_GDS_ACL),
    {DSX_MODIFY_PUBLIC, OM_S_OBJECT, {0, acl_item_ap}},
    {DSX_READ_STANDARD, OM_S_OBJECT, {0, acl_item_ap}},
    {DSX_MODIFY_STANDARD, OM_S_OBJECT, {0, acl_item_stefanie}},
    {DSX_READ_SENSITIVE, OM_S_OBJECT, {0, acl_item_ap}},
    {DSX_MODIFY_SENSITIVE, OM_S_OBJECT, {0, acl_item_stefanie}},
    OM_NULL_DESCRIPTOR
};
```



```

};

static OM_descriptor    att_acl1[] = {
    OM_OID_DESC(OM_CLASS, DS_C_ATTRIBUTE),
    OM_OID_DESC(DS_ATTRIBUTE_TYPE, DSX_A_ACL),
    {DS_ATTRIBUTE_VALUES, OM_S_OBJECT, {0, acl1} },
    OM_NULL_DESCRIPTOR
};

static OM_descriptor    att_acl2[] = {
    OM_OID_DESC(OM_CLASS, DS_C_ATTRIBUTE),
    OM_OID_DESC(DS_ATTRIBUTE_TYPE, DSX_A_ACL),
    {DS_ATTRIBUTE_VALUES, OM_S_OBJECT, {0, acl2} },
    OM_NULL_DESCRIPTOR
};

static OM_descriptor    alist_C[] = {
    OM_OID_DESC(OM_CLASS, DS_C_ATTRIBUTE_LIST),
    {DS_ATTRIBUTES, OM_S_OBJECT, {0, obj_class_C} },
    OM_NULL_DESCRIPTOR
};

static OM_descriptor    alist_O[] = {
    OM_OID_DESC(OM_CLASS, DS_C_ATTRIBUTE_LIST),
    {DS_ATTRIBUTES, OM_S_OBJECT, {0, obj_class_O} },
    {DS_ATTRIBUTES, OM_S_OBJECT, {0, att_acl1} },
    OM_NULL_DESCRIPTOR
};

static OM_descriptor    alist_OU[] = {
    OM_OID_DESC(OM_CLASS, DS_C_ATTRIBUTE_LIST),
    {DS_ATTRIBUTES, OM_S_OBJECT, {0, obj_class_OU} },
    OM_NULL_DESCRIPTOR
};

static OM_descriptor    alist_OP[] = {
    OM_OID_DESC(OM_CLASS, DS_C_ATTRIBUTE_LIST),
    {DS_ATTRIBUTES, OM_S_OBJECT, {0, obj_class_OP} },
    {DS_ATTRIBUTES, OM_S_OBJECT, {0, att_acl2} },
    {DS_ATTRIBUTES, OM_S_OBJECT, {0, att_surname} },

```

```
    {DS_ATTRIBUTES, OM_S_OBJECT, {0, att_phone_num} },
    {DS_ATTRIBUTES, OM_S_OBJECT, {0, att_password} },
    OM_NULL_DESCRIPTOR
};

/* The following descriptor list specifies what to return from
/* the entry. The ACL attribute's types and values are selected. */

static OM_descriptor selection_acl[] = {
    OM_OID_DESC(OM_CLASS, DS_C_ENTRY_INFO_SELECTION),
    {DS_ALL_ATTRIBUTES, OM_S_BOOLEAN, OM_FALSE},
    OM_OID_DESC(DS_ATTRIBUTES_SELECTED, DSX_A_ACL),
    {DS_INFO_TYPE, OM_S_ENUMERATION, DS_TYPES_AND_VALUES},
    OM_NULL_DESCRIPTOR
};

/* The following descriptor list specifies the bind credentials */

static OM_descriptor credentials[] = {
    {DS_REQUESTOR, OM_S_OBJECT, {0, dn_norbert} },
    {DSX_PASSWORD, OM_S_OCTET_STRING, OM_STRING("secret")},
    {DSX_AUTH_MECHANISM, OM_S_ENUMERATION, {DSX_SIMPLE,0}},
    OM_NULL_DESCRIPTOR
};

/* The following descriptor list specifies part of the context */

static OM_descriptor use_copy[] = {
    {DS_DONT_USE_COPY, OM_S_BOOLEAN, {OM_FALSE, 0} },
    OM_NULL_DESCRIPTOR
};

/* Build up an array of object identifiers for the optional
/* packages to be negotiated. */

DS_feature features[] = {
    { OM_STRING(OMP_O_DS_BASIC_DIR_CONTENTS_PKG), OM_TRUE },
    { OM_STRING(OMP_O_DSX_GDS_PKG), OM_TRUE },
    { 0 }
}
```

```

};

/* The following descriptor list specifies a filter for search : */
/*      (Present: objectClass)                                     */

static OM_descriptor filter_item[] = {
    OM_OID_DESC(OM_CLASS, DS_C_FILTER_ITEM),
    {DS_FILTER_ITEM_TYPE, OM_S_ENUMERATION, {DS_PRESENT, 0} },
    OM_OID_DESC(DS_ATTRIBUTE_TYPE, DS_A_OBJECT_CLASS),
    OM_NULL_DESCRIPTOR
};

static OM_descriptor filter[] = {
    OM_OID_DESC(OM_CLASS, DS_C_FILTER),
    {DS_FILTER_ITEMS, OM_S_OBJECT, {0, filter_item} },
    {DS_FILTER_TYPE, OM_S_ENUMERATION, {DS_AND, 0} },
    OM_NULL_DESCRIPTOR
};

#endif /* ACL_HEADER */

```

7.4 The teldir.c Program

The sample program **teldir.c** permits a user to add, read, or delete entries in a CDS or GDS namespace in any local or remote DCE cell, assuming that permissions are granted by the ACLs. The entry consists of a person's surname and phone number. Each entry is of class **Organizational-Person**.

The program uses predefined static XDS public objects that are never altered and partially defined static XDS public objects so that values for the surname and phone number can be entered dynamically by a user. It also uses dynamic XDS public objects that are created and filled only as needed by using the **stringToXdsName** function. These techniques are a departure from those used in the first two sample programs where all objects are predefined.

7.4.1 Predefined Static Public Objects

The predefined static object classes and attributes are shown in the following code fragment:

```
/*
 * To hold the attributes we want to attach to the name being added.
 * One attribute is the class of the object (DS_O_ORG_PERSON), the
 * rest of the attributes are the surname (required for all objects
 * of class DS_O_ORG_PERSON) and phone number. In addition, we need
 * an object to hold all this information to pass it into
 * ds_add_entry().
 */
static OM_descriptor xdsObjectClass[] = {

    /* This object is an attribute--an object class. */
    OM_OID_DESC( OM_CLASS,          DS_C_ATTRIBUTE ),
    OM_OID_DESC( DS_ATTRIBUTE_TYPE,  DS_A_OBJECT_CLASS ),

    /* Not only must the class be listed, but also all */
    /* its superclasses.                               */
    OM_OID_DESC( DS_ATTRIBUTE_VALUES, DS_O_TOP ),
    OM_OID_DESC( DS_ATTRIBUTE_VALUES, DS_O_PERSON ),
    OM_OID_DESC( DS_ATTRIBUTE_VALUES, DS_O_ORG_PERSON ),

    /* Null terminator */
    OM_NULL_DESCRIPTOR
};

static OM_descriptor xdsAttributesToAdd[] = {

    /* This object is an attribute list. */
    OM_OID_DESC( OM_CLASS, DS_C_ATTRIBUTE_LIST ),

    /* These are "pointers" to the attributes in the list. */
    { DS_ATTRIBUTES, OM_S_OBJECT, { 0, xdsObjectClass } },
    { DS_ATTRIBUTES, OM_S_OBJECT, { 0, xdsSurname } },
    { DS_ATTRIBUTES, OM_S_OBJECT, { 0, xdsPhoneNum } },

    /* Null terminator */
    OM_NULL_DESCRIPTOR
};
```

```
};

/*
 * To hold the list of attributes we want to read.
 */
static OM_descriptor xdsAttributeSelection[] = {

    /* This is an entry information selection. */
    OM_OID_DESC( OM_CLASS, DS_C_ENTRY_INFO_SELECTION ),

    /* No, we don't want all attributes. */
    { DS_ALL_ATTRIBUTES, OM_S_BOOLEAN, OM_FALSE },

    /* These are the ones we want to read. */
    OM_OID_DESC( DS_ATTRIBUTES_SELECTED, DS_A_SURNAME ),
    OM_OID_DESC( DS_ATTRIBUTES_SELECTED, DS_A_PHONE_NBR ),

    /* Give us both the types and their values. */
    { DS_INFO_TYPE, OM_S_ENUMERATION, { DS_TYPES_AND_VALUES, NULL } },

    /* Null terminator */
    OM_NULL_DESCRIPTOR
};
```

7.4.2 Partially Defined Static Public Objects

The program partially defines static XDS objects with placeholders so that values for the surname and telephone number entered by the user can be added later, as shown in the following code fragment:

```
static OM_descriptor xdsSurname[] = {

    /* This object is an attribute--a surname. */
    OM_OID_DESC( OM_CLASS,          DS_C_ATTRIBUTE ),
    OM_OID_DESC( DS_ATTRIBUTE_TYPE, DS_A_SURNAME ),

    /* No default--so we need a placeholder for the actual surname.  */
};
```

```
    OM_NULL_DESCRIPTOR,

    /* Null terminator */
    OM_NULL_DESCRIPTOR
};

static OM_descriptor xdsPhoneNum[] = {

    /* This object is an attribute--a telephone number. */
    OM_OID_DESC( OM_CLASS,          DS_C_ATTRIBUTE ),
    OM_OID_DESC( DS_ATTRIBUTE_TYPE, DS_A_PHONE_NBR ),

    /* By default, phone numbers are unlisted.  If the user specifies */
    /* an actual phone number, it will go into this position.          */
    { DS_ATTRIBUTE_VALUES, OM_S_PRINTABLE_STRING,
      OM_STRING( "unlisted" ) },

    /* Null terminator */
    OM_NULL_DESCRIPTOR
};
```

The program prompts the user for the surname of the person whose number will be changed and uses the **FILL_OMD_STRING** macro to fill in values, as shown in the following code fragment:

```
if ( operation == 'a' ) {
    /* add operation requires additional input */
    /*
     * Get the person's real name from the user and place it in the
     * XDS object already defined at the
     * top of the program (xdsSurname).
     * We are requiring a name, so we will loop until we get one.
     */
    do {
        printf( "What is this person's surname? " );
        gets( newSurname );
    } while ( *newSurname == '\0' );
    FILL_OMD_STRING( xdsSurname, 2, DS_ATTRIBUTE_VALUES,
                     OM_S_TELETEX_STRING, newSurname )
}
```

7.4.3 Dynamically Defined Public Objects

The program uses the function **stringToXdsName** to convert the DCE name entered by a user into an XDS name object of OM class **DS_C_DS_DN**, which is the representation of a distinguished name. In the other two sample programs, arrays of descriptor lists are statically declared to represent the AVAs and RDNs that make up the public object that represents a distinguished name. The function **stringToXdsName** parses the DCE name and dynamically converts it to a public object.

For example, the following code fragment shows how space for a **DS_C_AVA** object is allocated and its entries are filled by using the **FILL_OMD_XOM_STRING** and **FILL_OMD_NULL** macros:

```
/*
 * Allocate space for a DS_C_AVA object and fill in its entries:
 *     DS_C_AVA class identifier
 *     AVA's type
 *     AVA's value
 *     null terminator
 */
ava = (OM_descriptor *)malloc( sizeof(OM_descriptor) * 4 );
if( ava == NULL )                /* malloc() failed */
    return OM_MEMORY_INSUFFICIENT;
FILL_OMD_XOM_STRING( ava, 0, OM_CLASS, OM_S_OBJECT_IDENTIFIER_STRING,
                    DS_C_AVA )
splitNamePiece( start, &type, &value );
FILL_OMD_XOM_STRING( ava, 1, DS_ATTRIBUTE_TYPE,
                    OM_S_OBJECT_IDENTIFIER_STRING, type )
FILL_OMD_STRING( ava, 2, DS_ATTRIBUTE_VALUES, OM_S_PRINTABLE_STRING,
                value )
FILL_OMD_NULL( ava, 3 )
```

The program uses the same method to build the RDNs that make up the distinguished name. The distinguished name is NULL terminated by using the **FILL_OMD_NULL** macro, and the location of the new public object is provided

for the calling routine (main) in the pointer **xdsNameObj**, as shown in the following code fragment:

```
/* Add the DS_C_RDN object to the DS_C_DS_DN object. */
FILL_OMD_STRUCT( dsdn, index, DS_RDNS, OM_S_OBJECT, rdn )
    }

/*
 * Null terminate the DS_C_DS_DN, tell the calling routine
 * where to find it, and return.
 */
FILL_OMD_NULL( dsdn, index )
*xdsNameObj = dsdn;
return( OM_SUCCESS );

} /* end stringToXdsName() */
```

7.4.4 Main Program Procedural Steps

The program consists of the following general steps:

1. Examine the command-line argument to determine the type of operation (read, add, or delete entry) that the user wants to perform.
2. Initialize a workspace.
3. Pull in the packages with the required XDS features.
4. Prompt the user for the name entry on which the operation will be performed.
5. Convert the DCE-formatted user input string to an XDS object name.
6. Bind (without credentials) to the default server.
7. Perform the requested operation (read, add, or delete entry).
8. Perform error handling.
9. Unbind from the server.
10. Shut down the workspace, releasing resources back to the system.

Step 1 simply involves determining which of the three options—**r** (read), **a** (add), or **d** (delete)—the user has entered. Step 2 initializes a workspace, an operation required by XDS API for every application program. Step 3 is required because additional features not present in the directory service package need to be used by the application program. An additional package, the basic directory contents package, is defined in *featureList* as a static XDS object earlier in the program.

In Step 4, the user is prompted for the DCE-formatted name, which is the distinguished name of the person on whose telephone number the operation is to be performed. The name must be a fully or partially qualified name that begins with either the */...* or */.:* prefix. An example of a fully qualified, or global, name is */.../C=de/O=sni/OU=ap/CN=klaus*. An example of a partially qualified, or cell, name is */.:/brad/sni/com*. Additional information is requested in Step 5 if the user requests an add operation.

Step 5 converts the DCE-formatted name to an XDS object name (public object) by using the **stringToXdsName()** function call. This function builds an XDS public object that represents the distinguished name entered by the user.

Step 6 binds the session to the default server without credentials; username and password are not required.

In Step 7, the requested operation is performed by using XDS API functions calls. For an add operation, **ds_add_entry()** is performed; for a read operation, **ds_read()** is performed; and for a delete operation, **ds_remove_entry()** is performed. The read operation requires a series of XOM API **om_get()** function calls to extract the surname and phone number from the workspace. (For a detailed description of the XDS and XOM API function calls, refer to Chapters 5 and 6.)

Step 8 and Step 9 are required for every XDS API application program in order to clean up before the program exits. The session is unbound from the server, and the public and private objects are released to the system that provided the memory allocated for them.

7.4.5 The **teldir.c** Code

The following is a listing of the file **teldir.c**:

```
/*
 * This sample program behaves like a simple telephone directory.
 * It permits a user to add, read or delete entries in a GDS
 * namespace or to a CDS namespace in any local or remote DCE cell
 * (assuming that permissions are granted by the ACLs).
 *
 * Each entry is of class Organizational-Person and simply contains
 * a person's surname and their phone number.
 *
 * The addition of an entry is followed by a read to verify that the
 * information was entered properly.
 *
 * All valid names should begin with one of the following symbols:
 *     /...      Fully qualified name (from global root).
 *               such as  /.../C=de/O=sni/OU=ap/CN=klaus
 *
 *     /.:       Partially qualified name (from local cell root).
 *               such as  /.:/brad/sni/com
 *
 * This program demonstrates the following techniques:
 * - Using completely static XDS public objects (predefined at the top
 *   of the program and never altered).  See xdsObjectClass,
 *   xdsAttributesToAdd, and xdsAttributeSelection below.
 * - Using partially static XDS public objects (predefined at the top
 *   of the program but altered later).  See xdsSurname and xdsPhoneNum
 *   below.  See also the macros whose names begin with "FILL_OMD_".
 * - Using dynamic XDS public objects (created and filled in only as
 *   needed).  See the function stringToXdsName() below.
 * - Parsing DCE-style names and converting them into XDS objects.  See
 *   the function stringToXdsName() below.
 * - Getting the value of an attribute from an object read from the
 *   namespace using ds_read().  See the function extractValue() below.
 * - Getting the numeric value of an error (type DS_status) returned by
 *   one of the XDS calls.  See the function handleDSError() below.
 */

#ifdef THREADSAFE
#include <pthread.h>
#endif

#include <stdio.h>
```

```

#include <stdlib.h>
#include <string.h>

#include <xom.h>
#include <xds.h>
#include <xdsbdc.h>
#include <xdsbdc.h>
#include <xdsbdc.h>
#include <xdsbdc.h>

OM_EXPORT( DS_A_COMMON_NAME )
OM_EXPORT( DS_A_COUNTRY_NAME )
OM_EXPORT( DS_A_LOCALITY_NAME )
OM_EXPORT( DS_A_OBJECT_CLASS )
OM_EXPORT( DS_A_ORG_UNIT_NAME )
OM_EXPORT( DS_A_ORG_NAME )
OM_EXPORT( DS_A_SURNAME )
OM_EXPORT( DS_A_PHONE_NBR )
OM_EXPORT( DS_A_TITLE )
OM_EXPORT( DS_C_ATTRIBUTE )
OM_EXPORT( DS_C_ATTRIBUTE_LIST )
OM_EXPORT( DS_C_AVA )
OM_EXPORT( DS_C_DS_DN )
OM_EXPORT( DS_C_DS_RDN )
OM_EXPORT( DS_C_ENTRY_INFO_SELECTION )
OM_EXPORT( DS_O_ORG_PERSON )
OM_EXPORT( DS_O_PERSON )
OM_EXPORT( DS_O_TOP )
OM_EXPORT( DSX_TYPELESS_RDN ) /* For "typeless" pieces of a name, as */
                               /* found in cells with bind-style names*/
                               /* and in the CDS namespace.          */

#define MAX_NAME_LEN 1024

/* These values can be found in */
/* the "Directory Class Definitions" chapter. */
/* (One byte must be added for the null terminator.) */
#define MAX_PHONE_LEN 33
#define MAX_SURNAME_LEN 66

/*****

```

```
* Macros for help filling in static XDS objects.
*****/
/* Put NULL value (equivalent to OM_NULL_DESCRIPTOR) in object */
#define FILL_OMD_NULL( desc, index )
    desc[index].type = OM_NO_MORE_TYPES;
    desc[index].syntax = OM_S_NO_MORE_SYNTAXES;
    desc[index].value.object.padding = 0;
    desc[index].value.object.object = OM_ELEMENTS_UNSPECIFIED;

/* Put C-style (null-terminated) string in object */
#define FILL_OMD_STRING( desc, index, typ, syntax, val )
    desc[index].type = typ;
    desc[index].syntax = syntax;
    desc[index].value.string.length = (OM_string_length)
        strlen( val );
    desc[index].value.string.elements = val;

/* Put XOM string in object */
#define FILL_OMD_XOM_STRING( desc, index, typ, syntax, val )
    desc[index].type = typ;
    desc[index].syntax = syntax;
    desc[index].value.string.length = val.length;
    desc[index].value.string.elements = val.elements;

/* Put other value in object */
#define FILL_OMD_STRUCT( desc, index, typ, syntax, val )
    desc[index].type = typ;
    desc[index].syntax = syntax;
    desc[index].value.object.padding = 0;
    desc[index].value.object.object = val;

/*****
 * Static XDS objects.
 *****/
/*
 * To identify which packages we need for this program. We only need
 * the basic package because we are not doing anything fancy with
 * session parameters, etc.
 */
DS_feature featureList[] = {
```

```

        { OM_STRING(OMP_O_DS_BASIC_DIR_CONTENTS_PKG), OM_TRUE },
        { 0 }
};

/*
 * To hold the attributes we want to attach to the name being added.
 * One attribute is the class of the object (DS_O_ORG_PERSON), the
 * rest of the attributes are the surname (required for all objects
 * of class DS_O_ORG_PERSON) and phone number. In addition, we need
 * an object to hold all this information to pass it
 * into ds_add_entry().
 */
static OM_descriptor xdsObjectClass[] = {

    /* This object is an attribute--an object class. */
    OM_OID_DESC( OM_CLASS,          DS_C_ATTRIBUTE ),
    OM_OID_DESC( DS_ATTRIBUTE_TYPE, DS_A_OBJECT_CLASS ),

    /* Not only must the class be listed, but also all */
    /* its superclasses. */
    OM_OID_DESC( DS_ATTRIBUTE_VALUES, DS_O_TOP ),
    OM_OID_DESC( DS_ATTRIBUTE_VALUES, DS_O_PERSON ),
    OM_OID_DESC( DS_ATTRIBUTE_VALUES, DS_O_ORG_PERSON ),

    /* Null terminator */
    OM_NULL_DESCRIPTOR
};

static OM_descriptor xdsSurname[] = {

    /* This object is an attribute--a surname. */
    OM_OID_DESC( OM_CLASS,          DS_C_ATTRIBUTE ),
    OM_OID_DESC( DS_ATTRIBUTE_TYPE, DS_A_SURNAME ),

    /* No default--so we need a placeholder for the actual surname. */
    OM_NULL_DESCRIPTOR,

    /* Null terminator */
    OM_NULL_DESCRIPTOR
};

```

```
static OM_descriptor xdsPhoneNum[] = {

    /* This object is an attribute--a telephone number. */
    OM_OID_DESC( OM_CLASS,          DS_C_ATTRIBUTE ),
    OM_OID_DESC( DS_ATTRIBUTE_TYPE, DS_A_PHONE_NBR ),

    /* By default, phone numbers are unlisted.  If the user specifies */
    /* an actual phone number, it will go into this position.          */
    { DS_ATTRIBUTE_VALUES, OM_S_PRINTABLE_STRING,
      OM_STRING( "unlisted" ) },

    /* Null terminator */
    OM_NULL_DESCRIPTOR
};

static OM_descriptor xdsAttributesToAdd[] = {

    /* This object is an attribute list. */
    OM_OID_DESC( OM_CLASS, DS_C_ATTRIBUTE_LIST ),

    /* These are "pointers" to the attributes in the list. */
    { DS_ATTRIBUTES, OM_S_OBJECT, { 0, xdsObjectClass } },
    { DS_ATTRIBUTES, OM_S_OBJECT, { 0, xdsSurname } },
    { DS_ATTRIBUTES, OM_S_OBJECT, { 0, xdsPhoneNum } },

    /* Null terminator */
    OM_NULL_DESCRIPTOR
};

/*
 * To hold the list of attributes we want to read.
 */
static OM_descriptor xdsAttributeSelection[] = {

    /* This is an entry information selection. */
    OM_OID_DESC( OM_CLASS, DS_C_ENTRY_INFO_SELECTION ),

    /* No, we don't want all attributes. */
    { DS_ALL_ATTRIBUTES, OM_S_BOOLEAN, OM_FALSE },
}
```

```

/* These are the ones we want to read. */
OM_OID_DESC( DS_ATTRIBUTES_SELECTED, DS_A_SURNAME ),
OM_OID_DESC( DS_ATTRIBUTES_SELECTED, DS_A_PHONE_NBR ),

/* Give us both the types and their values. */
{ DS_INFO_TYPE, OM_S_ENUMERATION, { DS_TYPES_AND_VALUES, NULL } },

/* Null terminator */
OM_NULL_DESCRIPTOR
};

/*****
 * dce_cf_get_cell_name()
 * Use this dummy function if CDS is not available.
 *****/
void
dce_cf_get_cell_name(
    char **          cellname,
    unsigned long *   status
)
{
    fprintf( stderr, "CDS unavailable: dce_cf_get_cell_name() error\n" );
    *status = 1;
} /* end dce_cf_get_cell_name() */

/*****
 * showUsage()
 * Display "usage" information.
 *****/
void
showUsage(
    char * cmd          /* In--Name of command being called */
)
{
    fprintf( stderr, "\nusage:  %s [option]\n\n", cmd );
    fprintf( stderr, "option:  -a : add an entry\n" );
}

```

```
        fprintf( stderr, "          -r : read an entry\n" );
        fprintf( stderr, "          -d : delete an entry\n" );

    } /* end showUsage() */

    /*****
    * numNamePieces()
    * Returns the number of pieces in a string name.
    *****/
    int
    numNamePieces(
        char *      string /* In--String whose pieces are to be counted */
    )
    {
        int         count; /* Number of pieces found */
        char *      currSep; /* Pointer to separator between pieces */

        if( string == NULL ) /* If nothing there, no pieces */
            return( 0 );
        count = 1;           /* Otherwise, there's at least one */

        /*
        * If the first character is a /, it's not really separating
        * two pieces so we want to ignore it here.
        */
        if( *string == '/' )
            currSep = string + 1;
        else
            currSep = string;

        /* How many pieces are there? */
        while( (currSep = strchr( currSep, '/' )) != NULL ) {
            count++;
            currSep++; /* Begin at next character */
        }

        return( count );
    } /* end numNamePieces() */
```



```

/*****
 * splitNamePiece()
 * Divides a piece of a name (string) into its XDS attribute type
 * and value.
 *****/
void
splitNamePiece(
    char *      string, /* In--String to be broken down */
    OM_string * type,   /* Out--XDS type of this piece of the name */
    char **     value   /* Out--Pointer to beginning of the value */
)
/* part of string */
{
    char *      equalSign; /* Location of the = within string */

    /*
     * If the string contains an equal sign, this is probably a
     * typed name. Check for all the attribute types allowed by
     * the default schema.
     */
    if( (equalSign = strchr( string, '=' )) != NULL ) {

        *value = equalSign + 1;

        if(( strcmp( string, "C=", 2 ) == 0 ) ||
           ( strcmp( string, "c=", 2 ) == 0 ))
            *type = DS_A_COUNTRY_NAME;

        else if(( strcmp( string, "O=", 2 ) == 0 ) ||
                 ( strcmp( string, "o=", 2 ) == 0 ))
            *type = DS_A_ORG_NAME;

        else if(( strcmp( string, "OU=", 3 ) == 0 ) ||
                 ( strcmp( string, "ou=", 3 ) == 0 ))
            *type = DS_A_ORG_UNIT_NAME;

        else if(( strcmp( string, "LN=", 3 ) == 0 ) ||
                 ( strcmp( string, "ln=", 3 ) == 0 ))
            *type = DS_A_LOCALITY_NAME;
    }
}

```

```
        else if(( strcmp( string, "CN=", 3 ) == 0 ) ||
                ( strcmp( string, "cn=", 3 ) == 0 ))
            *type = DS_A_COMMON_NAME;

/*
 * If this did not appear to be a type allowed by the
 * default schema, consider the whole string as the
 * value (whose type is "typeless").
 */
else {
    *type = DSX_TYPELESS_RDN;
    *value = string;
}

/*
 * If the string does not contain an equal sign, this is a
 * typeless name.
 */
else {
    *type = DSX_TYPELESS_RDN;
    *value = string;
}

} /* end splitNamePiece() */

/*****
 * extractValue()
 * Pulls the value of a particular attribute from a private object
 * that was received using ds_read().
 * Returns:
 * OM_SUCCESS                If successful.
 * OM_NO_SUCH_OBJECT         If no values for the attribute
 *                           were found.
 * other                     Any value returned by one of the
 *                           om_get() calls.
 *****/
OM_return_code
```

```

extractValue(
    OM_private_object    object,      /* In--Object to extract from */
    OM_string *          attribute, /* In--Attribute to extract   */
    char *               value       /* Out--Value found           */
)
{
    OM_public_object    attrList;
    OM_public_object    attrType;
    OM_public_object    attrValue;
    OM_public_object    entry;
    int                 i;
    OM_return_code       omStatus;
    OM_value_position    total;
    OM_value_position    totalAttributes;
    OM_type              xdsIncludedTypes[] = { 0, /* Place holder */
                                                0 }; /* Null terminator*/

    /*
     * Get the entry from the object returned by ds_read().
     */
    xdsIncludedTypes[0] = DS_ENTRY;
    omStatus = om_get( object, /* Object to extract from */
                     OM_EXCLUDE_ALL_BUT_THESE_TYPES+OM_EXCLUDE_SUBOBJECTS,
                     /* Only want what is in */
                     /* xdsIncludedTypes, don't*/
                     /* include subobjects */
                     xdsIncludedTypes, /* What to get */
                     OM_FALSE, /* Currently ignored */
                     OM_ALL_VALUES, /* Start with first value */
                     OM_ALL_VALUES, /* End with last value */
                     &entry, /* Put the entry here */
                     &total ); /* Put number of attribut */
                                /* descriptors here */

    if( omStatus != OM_SUCCESS ) {
        fprintf( stderr, "om_get( entry ) returned error %d\n",
                 omStatus );
        return( omStatus );
    }

    if( total <= 0 ) { /* Make sure something was returned */
        fprintf( stderr,

```

```
        "Number of descriptors returned by om_get( entry )
        was %d\n", total );
    return( OM_NO_SUCH_OBJECT );
}

/*
 * Get the attribute list from the entry.
 */
xdsIncludedTypes[0] = DS_ATTRIBUTES;
omStatus = om_get( entry->value.object.object,
                  OM_EXCLUDE_ALL_BUT_THESE_TYPES+OM_EXCLUDE_SUBOBJECTS,
                  xdsIncludedTypes, OM_FALSE, OM_ALL_VALUES,
                  OM_ALL_VALUES, &attrList, &totalAttributes );
if( omStatus != OM_SUCCESS ) {
    fprintf( stderr, "om_get( attrList ) returned error %d\n",
            omStatus );
    return( omStatus );
}
if( total <= 0 ) { /* Make sure something was returned */
    fprintf( stderr,
        "Number of descriptors returned by om_get( attrList )
        was %d\n", total );
    return( OM_NO_SUCH_OBJECT );
}

/*
 * Search the list for the attribute with the proper type.
 */
for( i = 0; i < totalAttributes; i++ ) {
    xdsIncludedTypes[0] = DS_ATTRIBUTE_TYPE;
    omStatus = om_get( (attrList+i)->value.object.object,
                      OM_EXCLUDE_ALL_BUT_THESE_TYPES+OM_EXCLUDE_SUBOBJECTS,
                      xdsIncludedTypes, OM_FALSE, OM_ALL_VALUES,
                      OM_ALL_VALUES, &attrType, &total );
    if( omStatus != OM_SUCCESS ) {
        fprintf( stderr, "om_get( attrType ) [i = %d] returned
            error %d\n", i, omStatus );
        return( omStatus );
    }
}
if( total <= 0 ) { /* Make sure something was returned */
```

```

        fprintf( stderr,
        "Number of descriptors returned by om_get( attrType )
        [i = %d] was %d\n", i, total );
        return( OM_NO_SUCH_OBJECT );
    }
    if( strcmp( attrType->value.string.elements,
        attribute->elements,
        attribute->length ) == 0 )
        break;    /* If we found a match, quit looking. */
    }
    if( i == totalAttributes ) {    /* Verify that we found a match. */
        fprintf( stderr,
        "%s: extractValue() could not find requested attribute\n" );
        return( OM_NOT_PRESENT );
    }

    /*
    * Get the attribute value from the corresponding item in the
    * attribute list.
    */
    xdsIncludedTypes[0] = DS_ATTRIBUTE_VALUES;
    omStatus = om_get( (attrList+i)->value.object.object,
        OM_EXCLUDE_ALL_BUT_THESE_TYPES+OM_EXCLUDE_SUBOBJECTS,
        xdsIncludedTypes, OM_FALSE, OM_ALL_VALUES,
        OM_ALL_VALUES, &attrValue, &total );
    if( omStatus != OM_SUCCESS ) {
        fprintf( stderr, "om_get( attrValue ) returned error %d\n",
        omStatus );
        return( omStatus );
    }
    if( total <= 0 ) { /* Make sure something was returned */
        fprintf( stderr,
        "Number of descriptors returned by om_get( attrValue )
        was %d\n", total );
        return( OM_NO_SUCH_OBJECT );
    }

    /*
    * Copy the value into the buffer for return to the caller.
    */

```

```
strncpy( value, attrValue->value.string.elements,
        attrValue->value.string.length );
value[attrValue->value.string.length] = '\0';

/*
 * Free up the resources we don't need any more and return.
 */
om_delete( attrValue );
om_delete( attrType );
om_delete( attrList );
om_delete( entry );
return( OM_SUCCESS );

} /* end extractValue() */

/*****
 * stringToXdsName()
 *
 * Converts a string that is a DCE name to an XDS name object (class
 * DS_C_DS_DN). Returns one of the following:
 *
 *     OM_SUCCESS             If successful.
 *     OM_MEMORY_INSUFFICIENT If a malloc fails.
 *     OM_PERMANENT_ERROR     If the name is not in a valid format.
 *     OM_SYSTEM_ERROR       If the local cell's name cannot
 *                           be determined.
 *
 * Technically, the space obtained here through malloc() needs
 * to be returned to the system when it is no longer needed.
 * If this was a more complex application, this function would
 * probably malloc all the space it needs at once and require
 * calling routines to free the space when finished with it.
 *****/
OM_return_code
stringToXdsName(
    char * origString, /* In--String name to be converted */
    OM_object * xdsNameObj /* Out--Pointer to XDS name object */
)
{
    OM_descriptor * ava;          /* DS_C_AVA object */
    char * cellName;             /* Name of this cell */
}
```

```

OM_object      dsdn;      /* DS_C_DS_DN object */
char *         end;       /* End of name piece */
int            index;     /* Index into DS_C_DS_DN object */
int            numberOfPieces; /* Number of pieces in the name */
unsigned long  rc;        /* Return code for some functions */
OM_descriptor * rdn;      /* DS_C_RDN object */
char *         start;     /* Beginning of piece of name */
char *         string;    /* Copy of origString that we can use */
OM_string      type;     /* Type of one piece of the name */
char *         value;    /* Piece of the name */

/*
 * A DS_C_AVA object only contains pointers to the strings that
 * represent the pieces of the name, not the contents of the
 * strings themselves. So we'll make a copy of the string passed
 * in to guarantee that these pieces survive in case the programmer
 * alters or reuses the original string.
 *
 * In addition, all valid names should begin with one of the
 * following symbols:
 *
 *      /...      Fully qualified name (from global root). For
 *                these, we need to ignore the /...
 *
 *      /.:      Partially qualified name (from local cell root).
 *                For these, we must replace the /.: with the name
 *                of the local cell name
 *
 * If we see anything else, we'll return with an error. (Notice
 * that /: is a valid DCE name, but refers to the file system's
 * namespace. Filenames cannot be accessed through
 * CDS, GDS, or XDS.)
 */
if( strcmp( origString, "/.../", 5 ) == 0 ) {
    string = (char *)malloc( strlen(origString+5) + 1 );
    if( string == NULL ) /* malloc() failed */
        return OM_MEMORY_INSUFFICIENT;
    strcpy( string, origString+5 );
}
else if( strcmp( origString, "/./", 4 ) == 0 ) {
    dce_cf_get_cell_name( &cellName, &rc );
    if( rc != 0 ) /* Could not get cell name */
        return OM_SYSTEM_ERROR;
}

```

```
/*
 * The cell name will have /.../ on the front, so we will
 * skip over it as we add it to the string (by always
 * starting at its fifth character).
 */
string = (char *)malloc( strlen
                        (origString+4) + strlen(cellName+5) + 2 );
if( string == NULL )           /* malloc() failed */
    return OM_MEMORY_INSUFFICIENT;
strcpy( string, cellName+5 );
strcat( string, "/" );
strcat( string, origString+4 );
}
else                          /* Invalid name format */
    return OM_PERMANENT_ERROR;

/*
 * Count the number of pieces in the name that will have to
 * be dealt with.
 */
numberOfPieces = numNamePieces( string );

/*
 * Allocate memory for the DS_C_DS_DN object. We will need an
 * OM_descriptor for each name piece, one for the class
 * identifier, and one for the null terminator.
 */
dsdn = (OM_object)malloc(
    (numberOfPieces + 2) * sizeof(OM_descriptor) );
if( dsdn == NULL )           /* malloc() failed */
    return OM_MEMORY_INSUFFICIENT;

/*
 * Initialize it as a DS_C_DS_DN object by placing that class
 * identifier in the first position.
 */
FILL_OMD_XOM_STRING( dsdn, 0, OM_CLASS,
                    OM_S_OBJECT_IDENTIFIER_STRING, DS_C_DS_DN )
```



```

/*
 * For each piece of string, do the following:
 *     Break off the next piece of the string
 *     Build a DS_C_AVA object to show the type and value
 *         of this piece of the name
 *     Wrap the DS_C_AVA up in a DS_C_RDN object
 *     Add the DS_C_RDN to the DS_C_DS_DN object
 */
for( start=string, index=1 ; index <= numberOfPieces ;
    index++, start=end+1 ) {

    /*
     * Find the next delimiter and replace it with a null byte
     * so the piece of the name is effectively separated from
     * the rest of the string.
     */
    end = strchr( start, '/' );
    if( end != NULL )
        *end = '\0';
    else
        /* If this is the last piece, there won't be */
        /* a '/' at the end, just a null byte.          */
        end = strchr( start, '\0' );

    /*
     * Allocate space for a DS_C_AVA object and fill in its entries:
     * DS_C_AVA class identifier
     * AVA's type
     * AVA's value
     * null terminator
     */
    ava = (OM_descriptor *)malloc( sizeof(OM_descriptor) * 4 );
    if( ava == NULL )
        /* malloc() failed */
        return OM_MEMORY_INSUFFICIENT;
    FILL_OMD_XOM_STRING( ava, 0, OM_CLASS,
                        OM_S_OBJECT_IDENTIFIER_STRING, DS_C_AVA )
    splitNamePiece( start, &type, &value );
    FILL_OMD_XOM_STRING( ava, 1, DS_ATTRIBUTE_TYPE,
                        OM_S_OBJECT_IDENTIFIER_STRING, type )
    FILL_OMD_STRING( ava, 2, DS_ATTRIBUTE_VALUES,
                    OM_S_PRINTABLE_STRING, value )
}

```

```
FILL_OMD_NULL( ava, 3 )

/*
 * Allocate space for a DS_C_RDN object and fill in its entries:
 * DS_C_RDN class identifier
 * AVA it contains
 * null terminator
 */
rdn = (OM_descriptor *)malloc( sizeof(OM_descriptor) * 3 );
if( rdn == NULL ) /* malloc() failed */
    return OM_MEMORY_INSUFFICIENT;
FILL_OMD_XOM_STRING( rdn, 0, OM_CLASS,
                    OM_S_OBJECT_IDENTIFIER_STRING, DS_C_DS_RDN )
FILL_OMD_STRUCT( rdn, 1, DS_AVAS, OM_S_OBJECT, ava )
FILL_OMD_NULL( rdn, 2 )

/* Add the DS_C_RDN object to the DS_C_DS_DN object. */
FILL_OMD_STRUCT( dsdn, index, DS_RDNS, OM_S_OBJECT, rdn )
}

/*
 * Null terminate the DS_C_DS_DN, tell the calling routine
 * where to find it, and return.
 */
FILL_OMD_NULL( dsdn, index )
*xdsNameObj = dsdn;
return( OM_SUCCESS );

} /* end stringToXdsName() */

/*****
 * handleDSError()
 *
 * Extracts the error number from a DS_status return code, prints it
 * in an error message, then terminates the program.
 *****/
void
handleDSError(
    char *      header, /* In--Name of function whose return code */
                  /* is being checked */
    )
```

```

    DS_status    returnCode /* In--Return code to be checked */
)
{
    OM_type          includeDSProblem[] = { DS_PROBLEM,
                                           0 };

    OM_return_code    omStatus;
    OM_public_object   problem;
    OM_value_position  total;

    /*
     * A DS_status return code is an object. It will be one of the
     * subclasses of the class DS_C_ERROR. What we want from it is
     * the value of the attribute DS_PROBLEM.
     */
    omStatus = om_get( returnCode,
                      OM_EXCLUDE_ALL_BUT_THESE_TYPES+OM_EXCLUDE_SUBOBJECTS,
                      includeDSProblem,
                      OM_FALSE,
                      OM_ALL_VALUES,
                      OM_ALL_VALUES,
                      &problem,
                      &total );

    /*
     * Make sure we successfully extracted the problem number and print
     * the error message before quitting.
     */
    if( (omStatus == OM_SUCCESS) && (total > 0) )
        printf( "%s returned error %d\n", header,
                problem->value.enumeration );
    else
        printf( "%s failed for unknown reason\n", header );

    exit( 1 );
}

/*****
 * Main program
 */

```

```
void
main(
    int  argc,
    char * argv[]
)
{
    DS_status      dsStatus;
    OM_sint        invokeID;
    char           newName[MAX_NAME_LEN];
    char           newPhoneNum[MAX_PHONE_LEN];
    char           newSurname[MAX_SURNAME_LEN];
    OM_return_code  omStatus;
    char           phoneNumRead[MAX_PHONE_LEN];
    int            rc = 0;
    OM_private_object readResult;
    OM_private_object session;
    char           surnameRead[MAX_SURNAME_LEN];
    OM_object       xdsName;
    OM_workspace    xdsWorkspace;
    int            operation;

    /* Step 1
     *
     * Examine command-line argument.
     */
    operation = getopt( argc, argv, "rad" );
    if ( (operation == '?') || (operation == EOF) ) {
        showUsage( argv[0] );
        exit( 1 );
    }

    /* Step 2
     *
     * Initialize the XDS workspace.
     */
    xdsWorkspace = ds_initialize( );
    if( xdsWorkspace == NULL ) {
        fprintf( stderr, "ds_initialize() failed\n" );
        exit( 1 );
    }
}
```

```
/* Step 3
 *
 * Pull in the packages that contain the XDS features we need.
 */
dsStatus = ds_version( featureList, xdsWorkspace );
if( dsStatus != DS_SUCCESS )
    handleDSError( "ds_version()", dsStatus );

/* Step 4
 *
 * Find out what name the user wants to use in the namespace and
 * convert it to an XDS object. We do this conversion dynamically
 * (not using static structures defined at the top of the program)
 * because we don't know how long the name will be.
 */
switch( operation ) {
case 'r' :
    printf( "What name do you want to read? " );
    break;
case 'a' :
    printf( "What name do you want to add? " );
    break;
case 'd' :
    printf( "What name do you want to delete? " );
    break;
}

/* Step 5 */

gets( newName );
omStatus = stringToXdsName( newName, &xdsName );
if( omStatus != OM_SUCCESS ) {
    fprintf( stderr, "stringToXdsName() failed with OM error %d\n",
            omStatus );
    exit( 1 );
}

if ( operation == 'a' ) {
    /* add operation requires additional input */
}
```

```
/*
 * Get the person's real name from the user and place it in
 * the XDS object already defined at the top of the program
 * (xdsSurname). We are requiring a name, so we will loop
 * until we get one.
 */
do {
    printf( "What is this person's surname? " );
    gets( newSurname );
} while ( *newSurname == '\0' );
    FILL_OMD_STRING( xdsSurname, 2, DS_ATTRIBUTE_VALUES,
        OM_S_TELETEX_STRING, newSurname )

/*
 * Get the person's phone number from the user and place it
 * in the XDS object already defined at the top of the
 * program (xdsPhoneNum). A phone number is not required,
 * so if none is given we will use the default already
 * stored in the structure.
 */
printf( "What is this person's phone number? " );
gets( newPhoneNum );
if( *newPhoneNum != '\0' ) {
    FILL_OMD_STRING( xdsPhoneNum, 2, DS_ATTRIBUTE_VALUES,
        OM_S_PRINTABLE_STRING, newPhoneNum )
}
}

/* Step 6
 *
 * Open the session with the namespace:
 * bind (without credentials) to the default server.
 */
dsStatus = ds_bind( DS_DEFAULT_SESSION, xdsWorkspace, &session );
if( dsStatus != DS_SUCCESS )
    handleDSError( "ds_bind()", dsStatus );

/* Step 7 */
```

```

switch( operation ) {    /* perform the requested operation */

/*
 * Add entry to the namespace.  The xdsSurname and xdsPhoneNum
 * objects are already contained within an attribute list object
 * (xdsAttributesToAdd).
 */
case 'a' :
    dsStatus = ds_add_entry( session, DS_DEFAULT_CONTEXT, xdsName,
                             xdsAttributesToAdd, &invokeID );
    if( dsStatus != DS_SUCCESS )
        handleDSError( "ds_add_entry()", dsStatus );

    /* FALL THROUGH */

/*
 * Read the entry of the name supplied.
 */
case 'r' :
    dsStatus = ds_read( session, DS_DEFAULT_CONTEXT, xdsName,
                        xdsAttributeSelection, &readResult, &invokeID );
    if( dsStatus != DS_SUCCESS )
        handleDSError( "ds_read()", dsStatus );

    /*
     * Get each attribute from the object read and print them.
     */
    omStatus = extractValue( readResult, &DS_A_SURNAME,
                             surnameRead );
    if( omStatus != OM_SUCCESS ) {
        printf( "*** Surname could not be read\n" );
        strcpy( surnameRead, "(unknown)" );
        rc = 1;
    }
    omStatus = extractValue( readResult, &DS_A_PHONE_NBR,
                             phoneNumRead );
    if( omStatus != OM_SUCCESS ) {
        printf( "*** Phone number could not be read\n" );
        strcpy( phoneNumRead, "(unknown)" );
        rc = 1;
    }
}

```

```
    }
    printf( "The phone number for %s is %s.\n", surnameRead,
           phoneNumRead );

    break;

/*
 * delete the entry from the namespace.
 */
case 'd' :
    dsStatus = ds_remove_entry( session, DS_DEFAULT_CONTEXT,
                               xdsName, &invokeID );
    if( dsStatus != DS_SUCCESS )
        handleDSError( "ds_remove_entry()", dsStatus );
    else
        printf( "The entry has been deleted.\n" );
    break;
}

/*
 * Clean up and exit.
 */
/* Step 8 */
dsStatus = ds_unbind( session );
if( dsStatus != DS_SUCCESS )
    handleDSError( "ds_unbind()", dsStatus );

/* Step 9 */
dsStatus = ds_shutdown( xdsWorkspace );
if( dsStatus != DS_SUCCESS )
    handleDSError( "ds_shutdown()", dsStatus );

exit( rc );

} /* end main() */
```


Chapter 8

Using Threads With The XDS/XOM API

Some programs work well when they are structured as multiple flows of control. Other programs may show better performance when they are multithreaded, allowing the multiple threads to be mapped to multiple processors when they are available.

GDS application programs can contain multiple threads of control. For example, a GDS application can need to query several GDS servers. This can be achieved more efficiently by using separate threads simultaneously to query the different servers.

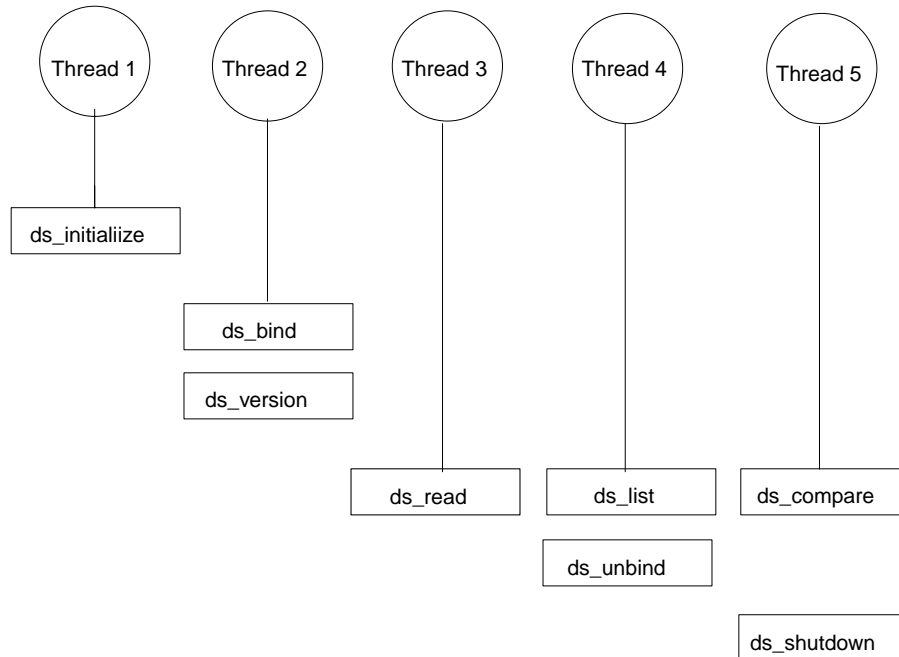
GDS supports multithreaded applications. Writing multithreaded applications over GDS imposes new requirements on programmers: they must manage the threads, synchronize threads' access to global resources, and make choices about thread scheduling and priorities.

This chapter describes a simple GDS application that uses threads. (Refer to the **(3thr)** reference pages for more information on DCE threads.)

The XDS/XOM API calls do not change when they are making use of DCE threads in an application program. The service underneath XDS/XOM API is designed to be both *thread-safe*, to allow multiple threads to safely access shared data, and *cancel-safe*, to handle unexpected cancellation of a thread in an application program.

Figure 8-1 shows an example of how an application can issue XDS/XOM calls from within different threads.

Figure 8–1. Issuing XDS/XOM Calls from Within Different Threads



The order of thread completion is not defined; however, XDS/XOM has an inherent ordering. Multithreaded XDS applications must adhere to the following order of execution:

1. **ds_initialize()**
2. **ds_version()** (optional)
3. **ds_bind()**

4. Other XDS calls in sequence or parallel from multiple threads
5. **ds_unbind()**
6. **ds_shutdown()**

Multithreaded XOM applications must adhere to the following order of execution:

1. **ds_initialize()**
2. XOM calls in sequence or parallel from multiple threads
3. **ds_shutdown()**

The XDS/XOM API returns an appropriate error code if these sequences are not adhered to. For example the following errors are returned:

DS_E_BUSY

If **ds_unbind()** is called while there are still outstanding operations, or if **ds_shutdown()** is called before all directory connections have been released by **ds_unbind()**.

OM_NO_SUCH_WORKSPACE

If any XOM API calls are made before calling **ds_initialize()**, or if a call to **ds_shutdown()** completes while there are outstanding XOM operations on the same workspace. In the latter case, these XOM operations will not be performed.

8.1 Overview of Sample Threads Program

The sample program is called **thradd**. The **thradd** program is a multithreaded XDS application that adds entries to a GDS directory. Each thread performs a **ds_add_entry()** call. The information for each entry to be added is read from an input file.

The **thradd** program can also be used to reset the directory to its original state. This is achieved by invoking **thradd** with a **-d** command-line argument. In this case, **thradd** uses the same input file and calls **ds_remove_entry()** for each entry. The **ds_remove_entry()** calls are also done in separate threads.

To keep the program short and clear, it works with a fixed tree for the upper nodes (**/C=it/O=sni/OU=ap**), to which the entries described in the input file are added. This fixed upper tree is added to the directory by **thradd**. The input file contains the common name, the surname, and the phone number of each **Organizational-Person** entry to be added.

For simplicity, only **pthread_join()** is used for synchronization purposes; mutexes are not used.

The **thradd** program can be enhanced to satisfy the following scenarios:

- As a server program for interactive directory actions from different users. The **thradd** program simulates a server program that gets requests from different users to add entries to a directory. In the case of **thradd**, the users' interactive input is simulated through the entries in the input file. Each line of input represents a different directory entry, and **thradd** uses a separate thread for each line.
- Initialization of the directory with data from file. The **thradd** program could be enhanced to read generic attribute information for a variety of directory object classes from a file, and to add the corresponding entries to the directory.

8.1.1 User Interface

The **thradd** program is called from the command line as follows:

```
thradd [-d] [-f file_name]
```

where:

- d** Causes the entries in the file and the tree **/C=it/O=sni/OU=ap** to be deleted; otherwise, they are added.
- f *file_name*** Specifies the name of the input file. If no input file is specified, then a default filename of **thradd.dat** is used.

8.1.2 Input File Format

The input file can contain any number of lines. Each line represents a directory entry of an organizational person. Each line must contain the following three attributes for each entry:

<common name> <surname> <phone number>

The attributes must be strings without space characters. Lines containing less than three strings are rejected by the program; any input on a line after the first three strings is ignored and can be used for comments. The attributes are separated by one or more space characters.

The input strings are not verified for their relevant attribute syntax. A wrong attribute syntax will result in either a **ds_add_entry()** error or a **ds_remove_entry()** error.

The following would be a valid input file for **thradd**:

Anna	Meister	010101
Erwin	Reiter	020202
Gerhard	Schulz	030303
Gottfried	Schmid	040404
Heidrun	Blum	050505
Hermann	Meier	060606
Josefa	Fischer	070707
Jutta	Arndt	080808
Leopold	Huber	090909
Magdalena	Schuster	101010
Margot	Junge	111111

8.1.3 Program Output

The **thradd** program writes messages to **stdout** for every action done by a thread. The order of the output can differ from the order in the input file; it depends on the execution of the different threads.

Errors are reported to **stderr**.

8.1.4 Prerequisites

The directory must be active before running **thradd**. If you are running **thradd** in *adding* mode then the directory should not contain a node **/C=it**. The **thradd** program should always be invoked twice with the same input file: first without and then with option **-d**. This guarantees that the directory is reset to its original state. The GDS administration program **gdsditadm** can be used to verify the directory contents after adding entries.

8.2 Description of Sample Program

The **thradd** program has a similar structure to the sample XDS programs in the previous chapter. Therefore, only a short general outline of the program is given here. The thread specifics are described in detail in the next section.

The static descriptors for the fixed tree (that is, **/C=it/O=sni/OU=ap**) are declared in the **thradd.h** header file. Listings of both the **thradd.c** application and the **thradd.h** header file are included in later sections of this chapter.

The main routine scans the command-line options, initializes the XDS workspace and, if working in adding mode, binds to the default GDS server without credentials, adds the fixed tree of upper nodes, and then unbinds from the directory.

The program then binds to the default GDS server without credentials. Each line of the input file is processed in turn by a **while** loop (until the end of the file is reached). The **while** loop contains two **for** loops. The first **for** loop creates a separate thread for each line of the input file, up to a maximum of **MAX_THREAD_NO** of threads.

The **add_or_remove()** procedure, which adds or removes an entry to/from the directory, is the starting point of each thread's processing.

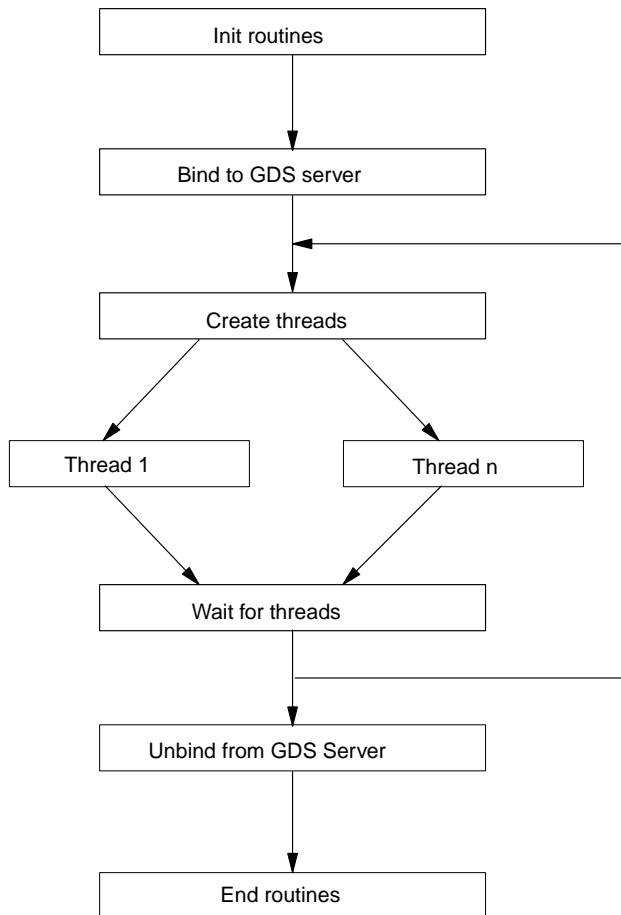
The second **for** loop waits for termination of the threads and then releases the resources used by the threads.

When the entire input file has been processed, **thradd** closes the connection to the GDS server and, if working in *removing* mode, removes the fixed tree of upper nodes (that is, **/C=it/O=sni/OU=ap**).

Finally, the XDS workspace is closed.

Figure 8-2 shows the program flow.

Figure 8–2. Program Flow for the thradd Sample Program



8.3 Detailed Description of Thread Specifics

The program consists of the following general steps:

1. Include the header file **pthread.h**.
2. Define a parameter block structure type for the thread start routine.
3. Declare arrays for thread handles and parameter blocks.
4. Read the input file line by line.
5. Update the parameter block.
6. Create the thread.
7. Wait for the termination of the thread.
8. Release the resources used by the thread.
9. Define the thread start routine.
10. Declare local variables needed for descriptors for the objects read from the input file.

The following paragraphs describe the corresponding step numbers from the program listing in the next section:

Step 1 includes the header file **pthread.h**, which is required for thread programming.

Step 2 defines a parameter block structure type for the thread start routine. A thread start routine must have exactly one parameter. However, **add_or_remove()** requires three parameters (session object, input line, and operating mode). The structure **pb_add_or_remove** is defined as the parameter block for these components. Therefore, the single parameter block contains the three parameters required by **add_or_remove()**.

Step 3 declares arrays for thread handles and parameter blocks. The routine that creates the thread (**main**, in this case) must maintain the following information for each thread:

- A thread handle of type **pthread_t** to identify the thread for join and detach calls.

- A thread-specific parameter block that cannot be accessed by any other thread. This makes sure that a parameter for one thread is not overwritten by another thread.

Step 4 reads the input file line by line. A thread is created for each line. A maximum **MAX_THREAD_NO** of threads is created in parallel. The program then waits for the termination of the created threads so that it can release the resources used by these threads, allowing it to create new threads for remaining input lines (if any).

The absolute maximum number of threads working in parallel depends on system limits; for **thradd**, a value of 10 was chosen (see **thradd.h**), which is well below the maximum on most systems.

Step 5 updates the parameter block. For each thread, a different element of the array of parameter blocks is used.

Step 6 creates the thread. The thread is created by using the function **pthread_create()**. The function has the following parameters:

- The thread handle (output) is stored in an element of the array of type **pthread_t**.
- For the thread characteristics, the default **pthread_attr_default** is used.
- The start routine for this thread is **add_or_remove()**.
- The parameter passed to **add_or_remove()** is a pointer to an element of the array of parameter blocks.

Step 7 waits for the termination of the thread. The **pthread_join()** routine is called with the thread handle as the input parameter. The program waits for the termination of the thread. If the thread has already terminated, then **pthread_join()** returns immediately. The second parameter of **pthread_join()** contains the return value of the start function; here it is a dummy value because **add_or_remove()** returns a **void**. The **add_or_remove()** routine is designed as a **void** function because the calling routine does not have to deal with error cases. The **add_or_remove()** routine prints status messages itself to show the processing order of the threads. Usually, a status should be returned to the application.

Step 8 releases the resources used by the thread. The thread handle is used as input for the function **pthread_detach()**, which releases the resources (for example, memory) used by the thread.

Step 9 defines the thread start routine. As previously mentioned, the thread start routine must have exactly one parameter. In this case, it is a pointer to the parameter block structure defined in Step 2.

Step 10 declares local variables needed for descriptors for the objects read from the input file. These descriptors are variables and are declared as automatic because of the reentrancy requirement. In the previous sample programs, descriptors were generally declared static. For this example, this is only possible for the constant descriptors declared in **thradd.h**.

Of course, this example shows only a small part of the possibilities of multithreaded XDS programming. For example, each thread could make its own bind, which would be useful if more than one GDS server was involved.

8.4 The thradd.c Code

The following code is a listing of the **thradd.c** program:

```
/*
 * The program operates in two modes; it adds or removes entries of
 * object type organizational person to/from a directory. The
 * information about the entries is read from a file.
 *
 * The program requires that a tree exists in the directory.
 * Therefore, each time the program runs, the following tree of 3
 * entries is added to or removed from the directory, according
 * to the operation mode.
 *
 *      O C=it
 *      | (objectClass=Country)
 *      |
 *      O O=sni
 *      | (objectClass=Organization)
 *      |
 *      O OU=ap
 *      (objectClass=OrganizationalUnit)
 *
 * Information about the organizational persons to be added or
```

```
* removed is read from the input file. It may contain any number
* of lines, where each line must have the following syntax:
*
*     <common name>   <surname>   <phone number>
*     Each item must be a string without a space.
*
* Lines containing less than 3 strings are rejected by the
* program. The program does not check to see if the strings conform
* to the appropriate attribute syntax; that is a wrong attribute
* syntax will lead to a ds_add_entry error, or to a
* ds_remove_entry error.
*
* Usage: thradd [-d] [-f<file_name>]
*   -d           If the option -d is set, the entries in the
*                file and the tree described above are removed,
*                otherwise they are added.
*   -f<file_name> The option -f specifies the name of the input
*                file. If left out, the default "thradd.dat"
*                is used.
*/

/* Step 1 */

/*
 * Header file for thread programming:
 */
#include <pthread.h>

#include <stdio.h>
#include <xom.h>
#include <xds.h>
#include <xdsbdc.h>
#include <xdsbdc.h>
#include <xdsbdc.h>
#include <xdsbdc.h>
#include "thradd.h"                                /* static data structures. */

/* Step 2 */

/*
 * typedef for parameter block of function add_or_remove
```

```
* (this is necessary because the start function of a thread
* takes only 1 parameter). The following 3 parameters are
* passed to add_or_remove:
*
*   Input - Session object from the ds_bind call
*   Input - Buffer with the entry information
*   Input - "adding" or "removing" mode ?
*/
typedef struct {
    OM_private_object session;
    char                line[MAX_LINE_LEN+1];
    int                 do_remove;
} pb_add_or_remove;

/*
 * static constants:
 */
/* Default name for input file containing entry information.
 */
static char fn_default[] = "thradd.dat";

/*
 * function declarations:
 */
char *own_fgets(char *s, int n, FILE *f);
void add_or_remove(pb_add_or_remove *pb);

int
main(
    int  argc,
    char *argv[]
)
{
    OM_workspace      workspace;          /* workspace for objects */
    OM_private_object bound_session;      /* Holds the Session      */
                                          /* returned by ds_bind() */

    FILE              *fp;                /* pointer for input file*/
    int                do_remove = FALSE; /* "adding" or "removing"*/
    int                error      = FALSE; /* error in options ?    */
    int                is_eof     = FALSE; /* EOF input file reached*/
}
```

```
int          thread_count;      /* no. of created threads*/
char         *file_name;       /* ptr to input file_name*/

/* Step 3 */

pthread_t     threads[MAX_THREAD_NO]; /* thread table */
pb_add_or_remove param_block[MAX_THREAD_NO]; /* 1 param block*/
/* for start routine per thread */

int          dummy;
int          c;
int          i;
extern char   *optarg; /* external variable used by getopt */
extern int    optind; /* external variable used by getopt */

/*
 * scan options -d and -f
 */
file_name = fn_default;

while ((c=getopt(argc, argv, "df:")) != EOF)
{
    switch (c)
    {
        case 'd':
            do_remove = TRUE;
            break;
        case 'f':
            file_name = optarg;
            break;
        default:
            error = TRUE;
            break;
    }
}

if (error)
{
    printf("usage: %s [-d] [-f<file_name>]\n", argv[0]);
    return(FAILURE);
}
```

```
if (( fp = fopen(file_name, "r")) == (FILE *) NULL)
{
    printf("fopen() error, file name: %s\n", file_name);
    return(FAILURE);
}

/*
 * Initialize a directory workspace for use by XOM.
 */
if ((workspace = ds_initialize()) == (OM_workspace)0)
    printf("ds_initialize() error\n");

/*
 * Negotiate the use of the BDCP and GDS packages.
 */
if (ds_version(features, workspace) != DS_SUCCESS)
    printf("ds_version() error\n");

/*
 * Add the fixed tree of entries, if in adding mode
 */
if (!do_remove)
    if (add_entries(workspace))
        printf("add_entries() error\n");

/*
 * Bind to the default GDS server.
 * The returned session object is stored in the private
 * object variable bound_session and is used for further
 * XDS function calls.
 */
if (ds_bind(DS_DEFAULT_SESSION, workspace, &bound_session)
    != DS_SUCCESS)
    printf("ds_bind() error\n");

/* Step 4 */

/*
 * Add or remove entries described in input file. This is done
```

```
    * in parallel, creating up to MAX_THREAD_NO threads at a time.
    */
while (!is_eof)
{
    for (thread_count=0; thread_count<MAX_THREAD_NO;
        thread_count++)
    {

/* Step 5 */

        /*
         * Prepare parameter block:
         */
        is_eof = (own_fgets(param_block[thread_count].line,
                           MAX_LINE_LEN, fp) == NULL);

        if (is_eof)
            break;

        param_block[thread_count].session = bound_session;
        param_block[thread_count].do_remove = do_remove;

/* Step 6 */

        /*
         * Create thread with start routine add_or_remove:
         */
        if (pthread_create(&threads[thread_count],
                           pthread_attr_default,
                           (pthread_startroutine_t) add_or_remove,
                           (pthread_addr_t) &param_block[thread_count])
            != SUCCESS)
            printf("pthread_create() error\n");
    } /* end for */

    /*
     * Wait for termination of the created threads and release
     * resources:
     */
    for (i=0; i<thread_count; i++)
    {
```

```
/* Step 7 */

/*
 * Wait for termination of thread
 * (If thread has terminated already, the function
 * returns immediately):
 */
if (pthread_join(threads[i], (pthread_addr_t) &dummy)
    != SUCCESS)
    printf("pthread_join() error\n");

/* Step 8 */

/*
 * Release resources used by the thread:
 */
if (pthread_detach(&threads[i]) != SUCCESS)
    printf("pthread_detach() error\n");
} /* end for */
} /* end while */

/*
 * Close the connection to the GDS server.
 */
if (ds_unbind(bound_session) != DS_SUCCESS)
    printf("ds_unbind() error\n");

if (om_delete(bound_session) != OM_SUCCESS)
    printf("om_delete() error\n");

/*
 * Remove the tree from the directory, if removing mode
 */
if (do_remove)
    if (remove_entries(workspace))
        printf("remove_entries() error\n");

/*
 * Close the directory workspace.
 */
```



```
    */
    if (ds_shutdown(workspace) != DS_SUCCESS)
        printf("ds_shutdown() error\n");

    fclose(fp);
    return(SUCCESS);
} /* end main() */

/* Step 9 */

/*
 * Handle (add or remove) a directory entry
 */
void
add_or_remove(
    pb_add_or_remove    *pb    /* parameter information    */
)
{
    /*
     * further local variables:
     */
    char                common_name[MAX_AT_LEN+1];
    char                phone_num[MAX_AT_LEN+1];
    char                surname[MAX_AT_LEN+1];
    OM_sint             invoke_id;

/* Step 10 */

    /*
     * local variables for descriptors for objects read from file
     */
    OM_descriptor       ava_genop[] = {
        OM_OID_DESC(OM_CLASS, DS_C_AVA),
        OM_OID_DESC(DS_ATTRIBUTE_TYPE, DS_A_COMMON_NAME),
        OM_NULL_DESCRIPTOR,    /* place holder */
        OM_NULL_DESCRIPTOR
    };

    OM_descriptor       rdn_genop[] = {
```

```
        OM_OID_DESC(OM_CLASS, DS_C_DS_RDN),
        OM_NULL_DESCRIPTOR, /* place holder */
        OM_NULL_DESCRIPTOR
    };

    OM_descriptor dn_genop[] = {
        OM_OID_DESC(OM_CLASS, DS_C_DS_DN),
        {DS_RDNS, OM_S_OBJECT, {0, rdn_it}},
        {DS_RDNS, OM_S_OBJECT, {0, rdn_sni}},
        {DS_RDNS, OM_S_OBJECT, {0, rdn_ap}},
        OM_NULL_DESCRIPTOR, /* place holder */
        OM_NULL_DESCRIPTOR
    };

    OM_descriptor att_phone_num[] = {
        OM_OID_DESC(OM_CLASS, DS_C_ATTRIBUTE),
        OM_OID_DESC(DS_ATTRIBUTE_TYPE, DS_A_PHONE_NBR),
        OM_NULL_DESCRIPTOR, /* place holder */
        OM_NULL_DESCRIPTOR
    };

    OM_descriptor att_surname[] = {
        OM_OID_DESC(OM_CLASS, DS_C_ATTRIBUTE),
        OM_OID_DESC(DS_ATTRIBUTE_TYPE, DS_A_SURNAME),
        OM_NULL_DESCRIPTOR, /* place holder */
        OM_NULL_DESCRIPTOR
    };

    OM_descriptor alist_OP[] = {
        OM_OID_DESC(OM_CLASS, DS_C_ATTRIBUTE_LIST),
        {DS_ATTRIBUTES, OM_S_OBJECT, {0, obj_class_OP}},
        OM_NULL_DESCRIPTOR, /* place holder */
        OM_NULL_DESCRIPTOR, /* place holder */
        OM_NULL_DESCRIPTOR
    };

    rdn_genop[1].type = DS_AVAS;
    rdn_genop[1].syntax = OM_S_OBJECT;
    rdn_genop[1].value.object.padding = 0;
```

```
rdn_genop[1].value.object.object = ava_genop;

dn_genop[4].type = DS_RDNS;
dn_genop[4].syntax = OM_S_OBJECT;
dn_genop[4].value.object.padding = 0;
dn_genop[4].value.object.object = rdn_genop;

alist_OP[2].type = DS_ATTRIBUTES;
alist_OP[2].syntax = OM_S_OBJECT;
alist_OP[2].value.object.padding = 0;
alist_OP[2].value.object.object = att_surname;

alist_OP[3].type = DS_ATTRIBUTES;
alist_OP[3].syntax = OM_S_OBJECT;
alist_OP[3].value.object.padding = 0;
alist_OP[3].value.object.object = att_phone_num;

if (sscanf(pb->line, "%s %s %s", common_name,
           surname, phone_num) != 3)
{
    printf("invalid input line: >%s<\n", pb->line);
    return;
}
/*
 * Fill descriptor for common name
 */
ava_genop[2].type = DS_ATTRIBUTE_VALUES;
ava_genop[2].syntax = OM_S_PRINTABLE_STRING;
ava_genop[2].value.string.length =
    (OM_string_length)strlen(common_name);
ava_genop[2].value.string.elements = common_name;

if (!pb->do_remove)    /* add */
{
    /*
     * Fill descriptors for surname and phone number
     */
    att_surname[2].type = DS_ATTRIBUTE_VALUES;
    att_surname[2].syntax = OM_S_TELETEX_STRING;
    att_surname[2].value.string.length =
```

```
        (OM_string_length)strlen(surname);
att_surname[2].value.string.elements = surname;

att_phone_num[2].type = DS_ATTRIBUTE_VALUES;
att_phone_num[2].syntax = OM_S_PRINTABLE_STRING;
att_phone_num[2].value.string.length =
        (OM_string_length)strlen(phone_num);
att_phone_num[2].value.string.elements = phone_num;

/*
 * add entry
 */
if (ds_add_entry(pb->session, DS_DEFAULT_CONTEXT, dn_genop,
        alist_OP, &invoke_id) != DS_SUCCESS)
    printf("ds_add_entry() error: %s %s %s\n",
        common_name, surname, phone_num);
else
    printf("entry added: %s %s %s\n",
        common_name, surname, phone_num);
}
else /* remove */
{
    /*
     * remove entry
     */
    if (ds_remove_entry(pb->session, DS_DEFAULT_CONTEXT,
        dn_genop, &invoke_id) != DS_SUCCESS)
        printf("ds_remove_entry() error: %s\n", common_name);
    else
        printf("entry removed: %s\n", common_name);
} /* end if */
} /* end add_or_remove() */

/*
 * Add the tree of entries described above.
 */
int
add_entries(
    OM_workspace workspace /* In--XDS workspace */
)
```

```
{
    OM_private_object bound_session; /* Holds Session object */
                                   /* returned by ds_bind() */
    OM_sint          invoke_id;
    int              error = FALSE;

    /* Bind (without credentials) to the default GDS server */

    if (ds_bind(DS_DEFAULT_SESSION, workspace, &bound_session)
        != DS_SUCCESS)
        error = TRUE;

    /* Add entries to the GDS server */

    if (ds_add_entry(bound_session, DS_DEFAULT_CONTEXT, dn_it,
                    alist_C, &invoke_id) != DS_SUCCESS)
        error = TRUE;

    if (ds_add_entry(bound_session, DS_DEFAULT_CONTEXT, dn_sni,
                    alist_O, &invoke_id) != DS_SUCCESS)
        error = TRUE;

    if (ds_add_entry(bound_session, DS_DEFAULT_CONTEXT, dn_ap,
                    alist_OU, &invoke_id) != DS_SUCCESS)
        error = TRUE;

    /* Close the connection to the GDS server */

    if (ds_unbind(bound_session) != DS_SUCCESS)
        error = TRUE;

    if (om_delete(bound_session) != OM_SUCCESS)
        error = TRUE;

    return (error);
}

/*
```

```
* Remove the tree of entries described above.
*/
int
remove_entries(
    OM_workspace workspace      /* In--XDS workspace */
)
{
    OM_private_object bound_session; /* Holds Session object */
                                      /* returned by ds_bind() */
    OM_sint             invoke_id;
    int                 error = FALSE;

    /* Bind to the default GDS server */

    if (ds_bind(DS_DEFAULT_SESSION, workspace, &bound_session)
        != DS_SUCCESS)
        error = TRUE;

    /* Remove entries from the GDS server */

    if (ds_remove_entry(bound_session, DS_DEFAULT_CONTEXT,
                        dn_ap, &invoke_id) != DS_SUCCESS)
        error = TRUE;

    if (ds_remove_entry(bound_session, DS_DEFAULT_CONTEXT,
                        dn_sni, &invoke_id) != DS_SUCCESS)
        error = TRUE;

    if (ds_remove_entry(bound_session, DS_DEFAULT_CONTEXT,
                        dn_it, &invoke_id) != DS_SUCCESS)
        error = TRUE;

    /* Close the connection to the GDS server */

    if (ds_unbind(bound_session) != DS_SUCCESS)
        error = TRUE;

    if (om_delete(bound_session) != OM_SUCCESS)
        error = TRUE;
}
```

```
        return (error);
    }

    /*
     * read one line with fgets and overwrite new line by
     * a null character
     */

    char *
    own_fgets(
        char *s,      /* OUT--string read */
        int n,        /* IN--maximum number of chars to be read */
        FILE *f       /* IN---input file */
    )
    {
        char *result;
        int i = 0;

        result = fgets(s, n, f);
        if (result != NULL)
        {
            i = strlen(s);
            if (s[i-1] == '\n')
                s[i-1] = '\0';
        }
        return (result);
    }
```

8.5 The thradd.h Header File

The following code is a listing of the **thradd.h** header file:

```
#ifndef THRADD_H
#define THRADD_H

#ifdef TRUE
#define TRUE (1)
```

```
#endif

#ifndef FALSE
#define FALSE (0)
#endif

#define SUCCESS 0
#define FAILURE 1
#define MAX_LINE_LEN 100 /* max length of line in input file */
#define MAX_AT_LEN 100 /* max length of an attribute value */
#define MAX_THREAD_NO 10 /* max number of threads created */

/* The application must export the object
   identifiers it requires.
*/

OM_EXPORT (DS_C_AVA)
OM_EXPORT (DS_C_DS_RDN)
OM_EXPORT (DS_C_DS_DN)
OM_EXPORT (DS_C_ATTRIBUTE)
OM_EXPORT (DS_C_ATTRIBUTE_LIST)

OM_EXPORT (DS_A_COUNTRY_NAME)
OM_EXPORT (DS_A_ORG_NAME)
OM_EXPORT (DS_A_ORG_UNIT_NAME)
OM_EXPORT (DS_A_COMMON_NAME)
OM_EXPORT (DS_A_OBJECT_CLASS)
OM_EXPORT (DS_A_PHONE_NBR)
OM_EXPORT (DS_A_SURNAME)

OM_EXPORT (DS_O_TOP)
OM_EXPORT (DS_O_COUNTRY)
OM_EXPORT (DS_O_ORG)
OM_EXPORT (DS_O_ORG_UNIT)
OM_EXPORT (DS_O_PERSON)
OM_EXPORT (DS_O_ORG_PERSON)

/* Build descriptor lists for the following */
```



```
/* distinguished names: */
/*      root              */
/*      /C=it             */
/*      /C=it/O=sni       */
/*      /C=it/O=sni/OU=ap */

static OM_descriptor  ava_it[] = {
    OM_OID_DESC(OM_CLASS, DS_C_AVA),
    OM_OID_DESC(DS_ATTRIBUTE_TYPE, DS_A_COUNTRY_NAME),
    {DS_ATTRIBUTE_VALUES, OM_S_PRINTABLE_STRING, OM_STRING("it")},
    OM_NULL_DESCRIPTOR
};

static OM_descriptor  ava_sni[] = {
    OM_OID_DESC(OM_CLASS, DS_C_AVA),
    OM_OID_DESC(DS_ATTRIBUTE_TYPE, DS_A_ORG_NAME),
    {DS_ATTRIBUTE_VALUES, OM_S_TELETEX_STRING, OM_STRING("sni")},
    OM_NULL_DESCRIPTOR
};

static OM_descriptor  ava_ap[] = {
    OM_OID_DESC(OM_CLASS, DS_C_AVA),
    OM_OID_DESC(DS_ATTRIBUTE_TYPE, DS_A_ORG_UNIT_NAME),
    {DS_ATTRIBUTE_VALUES, OM_S_TELETEX_STRING, OM_STRING("ap")},
    OM_NULL_DESCRIPTOR
};

static OM_descriptor  rdn_it[] = {
    OM_OID_DESC(OM_CLASS, DS_C_DS_RDN),
    {DS_AVAS, OM_S_OBJECT, {0, ava_it}},
    OM_NULL_DESCRIPTOR
};

static OM_descriptor  rdn_sni[] = {
    OM_OID_DESC(OM_CLASS, DS_C_DS_RDN),
    {DS_AVAS, OM_S_OBJECT, {0, ava_sni}},
    OM_NULL_DESCRIPTOR
};

static OM_descriptor  rdn_ap[] = {
    OM_OID_DESC(OM_CLASS, DS_C_DS_RDN),
    {DS_AVAS, OM_S_OBJECT, {0, ava_ap}},
    OM_NULL_DESCRIPTOR
};

static OM_descriptor dn_root[] = {
```

```
        OM_OID_DESC(OM_CLASS, DS_C_DS_DN),
        OM_NULL_DESCRIPTOR
};
static OM_descriptor dn_it[] = {
    OM_OID_DESC(OM_CLASS, DS_C_DS_DN),
    {DS_RDNS, OM_S_OBJECT, {0, rdn_it}},
    OM_NULL_DESCRIPTOR
};
static OM_descriptor dn_sni[] = {
    OM_OID_DESC(OM_CLASS, DS_C_DS_DN),
    {DS_RDNS, OM_S_OBJECT, {0, rdn_it}},
    {DS_RDNS, OM_S_OBJECT, {0, rdn_sni}},
    OM_NULL_DESCRIPTOR
};
static OM_descriptor dn_ap[] = {
    OM_OID_DESC(OM_CLASS, DS_C_DS_DN),
    {DS_RDNS, OM_S_OBJECT, {0, rdn_it}},
    {DS_RDNS, OM_S_OBJECT, {0, rdn_sni}},
    {DS_RDNS, OM_S_OBJECT, {0, rdn_ap}},
    OM_NULL_DESCRIPTOR
};

/* Build up an array of object identifiers for the */
/* attributes to be added to the directory.      */

static OM_descriptor  obj_class_C[] = {
    OM_OID_DESC(OM_CLASS, DS_C_ATTRIBUTE),
    OM_OID_DESC(DS_ATTRIBUTE_TYPE, DS_A_OBJECT_CLASS),
    OM_OID_DESC(DS_ATTRIBUTE_VALUES, DS_O_TOP),
    OM_OID_DESC(DS_ATTRIBUTE_VALUES, DS_O_COUNTRY),
    OM_NULL_DESCRIPTOR
};
static OM_descriptor  obj_class_O[] = {
    OM_OID_DESC(OM_CLASS, DS_C_ATTRIBUTE),
    OM_OID_DESC(DS_ATTRIBUTE_TYPE, DS_A_OBJECT_CLASS),
    OM_OID_DESC(DS_ATTRIBUTE_VALUES, DS_O_TOP),
    OM_OID_DESC(DS_ATTRIBUTE_VALUES, DS_O_ORG),
    OM_NULL_DESCRIPTOR
};
```

```

static OM_descriptor    obj_class_OU[] = {
    OM_OID_DESC(OM_CLASS, DS_C_ATTRIBUTE),
    OM_OID_DESC(DS_ATTRIBUTE_TYPE, DS_A_OBJECT_CLASS),
    OM_OID_DESC(DS_ATTRIBUTE_VALUES, DS_O_TOP),
    OM_OID_DESC(DS_ATTRIBUTE_VALUES, DS_O_ORG_UNIT),
    OM_NULL_DESCRIPTOR
};

static OM_descriptor    obj_class_OP[] = {
    OM_OID_DESC(OM_CLASS, DS_C_ATTRIBUTE),
    OM_OID_DESC(DS_ATTRIBUTE_TYPE, DS_A_OBJECT_CLASS),
    OM_OID_DESC(DS_ATTRIBUTE_VALUES, DS_O_TOP),
    OM_OID_DESC(DS_ATTRIBUTE_VALUES, DS_O_PERSON),
    OM_OID_DESC(DS_ATTRIBUTE_VALUES, DS_O_ORG_PERSON),
    OM_NULL_DESCRIPTOR
};

static OM_descriptor    alist_C[] = {
    OM_OID_DESC(OM_CLASS, DS_C_ATTRIBUTE_LIST),
    {DS_ATTRIBUTES, OM_S_OBJECT, {0, obj_class_C} },
    OM_NULL_DESCRIPTOR
};

static OM_descriptor    alist_O[] = {
    OM_OID_DESC(OM_CLASS, DS_C_ATTRIBUTE_LIST),
    {DS_ATTRIBUTES, OM_S_OBJECT, {0, obj_class_O} },
    OM_NULL_DESCRIPTOR
};

static OM_descriptor    alist_OU[] = {
    OM_OID_DESC(OM_CLASS, DS_C_ATTRIBUTE_LIST),
    {DS_ATTRIBUTES, OM_S_OBJECT, {0, obj_class_OU} },
    OM_NULL_DESCRIPTOR
};

/* Build up an array of object identifiers for the */
/* optional packages to be negotiated.             */

static DS_feature features[] = {
    { OM_STRING(OMP_O_DS_BASIC_DIR_CONTENTS_PKG), OM_TRUE },
    { OM_STRING(OMP_O_DSX_GDS_PKG), OM_TRUE },

```

```
        { 0 }  
    };  
  
#endif /* THRADD_H */
```

Chapter 9

XDS/XOM Convenience Routines

This chapter describes functions that are available to XDS/XOM programmers to help simplify and speed up the development of XDS applications. The convenience functions target two main areas, as follows:

- Filling, comparing, and extracting objects
- Converting objects to and from strings

The following six convenience functions are provided:

- **dsX_extract_attr_values()**
- **omX_fill()**
- **omX_fill_oid()**
- **omX_extract()**
- **omX_string_to_object()**
- **omX_object_to_string()**

Refer to the ***(3xds)** and ***(3xom)** reference pages for detailed descriptions of these functions.

To demonstrate the power of the convenience functions, the **acl.c** sample program from Chapter 7 is presented again here, after being modified to make use of these functions. The modified sample program is called **acl2.c**.

9.1 String Handling

The convenience functions provide the ability to specify OM objects in string format by means of abbreviations. These abbreviations are defined in the XOM object information file **xoiscema**.

X.500 attribute types can be specified as abbreviations or object identifier strings. The mapping of the attribute abbreviations and object identifier strings to BER encoded object identifiers and the associated attribute syntaxes is determined by the XOM object information module with the help of the **xoiscema** file. For valid attribute abbreviations, please refer to the **xoiscema** file in the following directory:

dce_local_path>/var/adm/directory/gds/adm

It is important that any schema changes to the DSA are reflected in the **xoiscema** file.

The convenience functions are able to handle strings with special syntax. The strings can be broadly classified into the following:

- Strings representing GDS attribute information
- Strings representing structured GDS attribute information
- Strings representing a structured GDS attribute value
- Strings representing a distinguished name (DN)
- Strings representing expressions

9.1.1 Strings Representing GDS Attribute Information

Strings that represent GDS attribute information are used to associate the attributes with their values. They are of the form:

attribute_type = attribute_value

The attribute types can either be specified as abbreviations or object identifier strings. An object identifier string is defined as a series of digits separated by the . (dot) character. If attribute abbreviations are used, they are case insensitive. For example, **cn=schmid** or **85.4.3=schmid**.

In the case of attributes with **OM_S_OBJECT_IDENTIFIER** syntax, the attribute value can also be specified as an abbreviation string. For example, an object class for **Residential Person** can be specified as **OCL=REP** or **OCL='x55x06x0A'**

All leading and trailing whitespace (surrounding the attribute type, the = (equal sign), and the attribute value) is ignored.

The following are the reserved characters for such strings:

- ' Used to enclose the attribute values. If this character is used, all other reserved characters within the quoted string except the \ (backslash) are not interpreted. For example, **cn=henry mueller**
- ;
 Separates multiple values of a recurring attribute. All leading and trailing whitespace (surrounding the semicolon) is ignored. For example, **TN=899898;979779**
- = Associates the attribute with its value.
- *xnn* Specifies hexadecimal data. The two characters *nn* are read as the hexadecimal value.
- \ Used to escape any of the other reserved characters.

9.1.2 Strings Representing Structured GDS Attribute Information

Strings that represent structured GDS attribute information are used to associate the structured attribute and its components with their values. They are of the form:

structured_attribute_type = {*Comp1* = *Value*, *Comp2* = *Value*, ..}

The structured attribute type can be specified as abbreviations or object identifier strings. An object identifier string is defined as a series of digits separated by dots. If attribute abbreviations are used, they are case insensitive. *Comp1*, *Comp2*, and so on, are the components of the structured attribute. They should be specified as abbreviations, as in the following example:

TXN={TN=977999, CC=345, AB=8444}

Recurring values for structured attributes can be specified with the help of the semicolon. An example follows:

TXN={TN=977999, CC=345, AB=8444};{TN=123444,CC=345, AB=8444}

Recurring values for the components should be specified as follows:

TXN={TN=977999; 274424, CC=345, AB=8444}

If any of the components are further structured, they should be enclosed within braces as follows:

FTN={PA={FR=1,TD=1}, PN=67899}

All leading and trailing whitespace, which surrounds the structured attribute type, the component abbreviation, the equal sign, the { (left brace), the , (comma), and the } (right brace), is ignored.

Attributes and components with DN syntax should be specified as follows:

AON={/c=de/o=sni/ou=ap11/cn=mueller}

ACL={MPUB={INT=0, USR={/c=de/o=sni/cn=mueller, sn=schmid}}}

In the case of attributes with **OM_S_OBJECT_IDENTIFIER** syntax, the attribute value can also be specified as an abbreviation string, as shown in the following:

SG={OCL=REP}

SG={OCL='x55\x06\x0A'}

Attributes of type presentation address (OM class **DS_C_PRESENTATION_ADDRESS**) are handled specially, using the PSAP macro utility. The value for such an attribute can be specified as follows:

PSA={TS=Server, NA='TCP/IP!internet=127.0.0.1+port=12345'}

The *local_string* parameter should be set to **OM_TRUE** in the convenience function being used. Here, the network address (NA) is specified with a special syntax. Refer to the *OSF DCE GDS Administration Guide and Reference* for further information.

The following are the reserved characters for strings with structured attribute information:

- ' Used to enclose the attribute values. If this character is used, all other reserved characters within the quoted string except the backslash are not interpreted. For example, **cn='henry mueller'**
- / Specifies an attribute value with DN syntax. For example, **AON = {c=de/o=sni/ou=ap22/cn=mayer}**
- { Indicates the start of a structured attribute value block.
- } Indicates the end of a structured attribute value block.
- , Separates the components of a structured attribute. For example, **TN=977999, CC=345, AB=8444**

It can also be used to specify multiple AVAs in the case of attributes with DN syntax.

- ; Separates multiple values of a recurring attribute or the recurring components of the structured attribute. All leading and trailing whitespace (surrounding the attribute type, the equal sign, the left and right braces, the component abbreviation, the component value and the semicolon) is ignored. The following is an example:

TXN={TN=977999,CC=345,AB=8444};{TN=53533,CC=242,AB=44242}

- = Associates the components with their values, and associates the components to the structured attribute.

<code>\xnn</code>	Used to specify hexadecimal data. The two characters <i>nn</i> are read as the hexadecimal value.
<code>\</code>	Used to escape any of the other reserved characters.

9.1.3 Strings Representing a Structured GDS Attribute Value

Strings are used to represent the structured GDS attribute value. Only one structured attribute value can be specified.

They are of the form:

Comp1 = Value, Comp2 = Value,

Comp1, *Comp2*, and so on, are the components of the structured attribute. They should be specified as abbreviations. For example, to specify a value for **DS_C_TELEX_NBR** class, the string format is the following:

TN=977999, CC=345, AB=8444

Recurring values for the components can be specified as shown in the following:

TN=977999; 274424, CC=345, AB=8444

If any of the components are further structured, they should be enclosed within braces as follows:

FTP={FR=1,TD=1}, PN=67899

Components with DN syntax can be specified as follows:

MPUB={INT=0, USR={/c=de/o=sni/cn=mueller, sn=schmid}}

Components of type presentation address (OM class **DS_C_PRESENTATION_ADDRESS**) are handled specially, using the PSAP macro utility. The value for the components can be specified as follows:

TS=Server, NA='TCP/IP!internet=127.0.0.1+port=12345'

The *local_string* parameter should be set to **OM_TRUE** in the convenience function being used. Here, the NA is specified with a special syntax. Refer to the *OSF DCE GDS Administration Guide and Reference* for further information.

The reserved characters for such strings are the same as those for strings representing structured attribute information (Section 9.1.2).

9.1.4 Strings Representing a Distinguished Name

Strings are used to represent the DN of the object. They are of the form:

/attribute_type = naming_attribute_value

or

/attribute_value/attribute_value

The attribute types can be specified as abbreviations or object identifier strings. An object identifier string is defined as a series of digits separated by dots. If attribute abbreviations are used, they are case insensitive. Multiple AVAs are represented by separating the naming attribute values with commas.

The first RDN can also be specified as the DCE global root string */...*, which is a sequence of the slash followed by three dots. In this case, the */...* string is simply ignored and the rest of the string is processed. Three examples follow:

/c=de/o=sni/ou=ap11, l=munich/85.4.3=schmid

/c=us/o=osf/ou=abc/subsystems/server/xyz

/.../c=us/o=osf/ou=abc/subsystems/server/xyz

The first nonspace character should always be the slash. All leading and trailing whitespace (surrounding the slash, the attribute type, the equal sign and the attribute value) is ignored.

The following are the reserved characters:

- ' Used to enclose the naming attribute values. If this character is used, all other reserved characters within the quoted string except the backslash are not interpreted. For example, **cn='henry mueller'**.

/	Used as a delimiter between RDNs.
,	Specifies multiple AVAs. All leading and trailing whitespace surrounding the comma is ignored. An example follows: /c=de/o=dbp/ou=dap11/cn=schmid, ou=ap11
=	Associates the object with its naming attribute value.
\ <i>xnn</i>	Used to specify hexadecimal data. The two characters <i>nn</i> are read as the hexadecimal value.
\	Used to escape any of the other reserved characters.

9.1.5 Strings Representing Expressions

Strings are used to specify an SQL-like expression in a search operation. For example, consider the following:

(CN~=schmid) && (OCL=ORP || OCL=REP) && !(SN=ronnie)

This is used to search for anybody who is an organizational person or a residential person, whose name approximately matches **schmid** but whose surname is not **ronnie**.

Object identifiers can also be used instead of attribute abbreviations. The object identifier string is a series of numbers separated by dots.

All leading and trailing whitespace (surrounding the attribute types, the operators, and the attribute values) is ignored.

If spaces are part of the attribute value, then the complete attribute value must be enclosed in quotes.

Additionally, the presence of an attribute can also be tested in either of the following ways:

c = de && cn

c = de && cn = *

The following are the reserved characters:

'	Used to indicate the start/end of an attribute value string. Can be used when spaces are part of the data. If this character is used, all other reserved characters within the quoted string except the backslash are not interpreted. An example follows: OU=sni && cn='Henri Mueller' && tn=89989
/	Used to specify an attribute value with DN syntax. An example follows: AON = {/c=de/o=sni/ou=ap22/cn=mayer}
=	Used to associate the attribute with its value.
&&	Used to logically AND two conditions.
	Used to logically OR two conditions.
!	Used to logically NEGATE a condition.
~=	Used to specify phonetic matching during a search operation.
>	Used to match values greater than a specified value.
>=	Used to match values greater than or equal to a specified value.
<	Used to match values less than a specified value.
<=	Used to match values less than or equal to a specified value.
*	Used to specify substrings during search.
(Used for nesting of filters.
)	Used for nesting of filters.
{	Indicates the start of a structured attribute value block.
}	Indicates the end of a structured attribute value block.
,	Separates the components of a structured attribute. For example, TN=977999, CC=345, AB=8444 It can also be used to specify multiple AVAs in the case of attributes with DN syntax.
\xnn	Used to specify hexadecimal data. The two characters <i>nn</i> are read as the hexadecimal value.
\	Used to escape any of the other reserved characters.

During evaluation of complex expressions during search operations, the following precedence of operators prevail:

1. ()
2. !
3. &&
4. ||

The () operators have the highest precedence, and || the lowest.

9.2 The **acl2.c** Program

The **acl2.c** file is a program that performs the same functionality as **acl.c** described in Chapter 7. Please refer to Chapter 7 for a complete description of the program's functionality, including outputs. The purpose of **acl2.c** and **acl2.h** is to show how the XDS/XOM convenience functions can be used to reduce the complexity of a real application.

The program consists of the following steps:

1. Export the required object identifiers. (See the **acl2.h** description in Section 9.2.2.)
2. Define the string expressions for the directory entry names and their attributes. (See the **acl2.h** description in Section 9.2.2.)
3. Initialize a workspace.
4. Negotiate use of the basic directory contents and GDS packages.
5. Build the name objects for the entries to be added to the directory.
6. Build the attribute objects for the entries to be added to the directory.
7. Add the fixed tree of entries to the directory in order to permit an authenticated bind.
8. Create a default session object.
9. Alter the default session object to include the credentials of the requestor (/C=**de/O=sni/OU=ap/CN=norbert**).

10. Bind with credentials to the default GDS server.
11. Create a default context object and alter it to include shadow entries.
12. Build filter, name, and entry information selection objects to be used for the search process.
13. Search the whole subtree below **root** and extract the ACL attribute from each selected entry.
14. Close the connection to the GDS server.
15. Remove the user's credentials from the directory.
16. Release the memory used for application-created objects.
17. Extract the components from the search result.
18. Examine each entry and print the entry details.
19. Close the XDS workspace.

In comparison to the **acl.c** program in Chapter 7, the following points should be noted:

- Step 1 has not changed significantly. The number of object identifiers, which the **acl2.c** needs to be exported, has been reduced.
- Step 2 has been completely revised. In fact, the header file has been reduced substantially. This is as a result of removing all the static descriptor lists for the directory names and attributes and replacing them with string expressions.
- Steps 3 and 4 are the same as before.
- Steps 5 and 6 are new steps that make use of the convenience functions **omX_string_to_object()**, **omX_fill_oid()**, and **omX_fill()**.
- Steps 7 through 10 are the same as Steps 5 through 8.
- Step 11 is the same as Step 9, but with an additional call to build an object to specify the use of shadow entries. A convenience function is used for this purpose. This replaces a static descriptor list definition from the old header file.
- Step 12 is new. It calls several convenience functions to create objects that are used by **ds_search()**. These objects were statically declared in the header file.
- Steps 13 through 15 are the same as Steps 10 through 12 from the old code.
- Step 16 is a new step to release memory that has been allocated by the convenience functions when creating objects.

- Step 17 replaces Step 13 from the old program with a call to the convenience function **omX_extract()** to extract the required components from the search result.
- Step 18 is the same as Step 14 in the old program, but with an additional call to free the memory allocated by **omX_extract()** in the previous step.
- Step 19 is the same as Step 15 in the old code.

9.2.1 The acl2.c Code

The following code is a listing of the **acl2.c** program:

```
/*
 *
 *  COPYRIGHT (C) SIEMENS NIXDORF INFORMATIONSSYSTEME AG 1991
 *
 *      ALL RIGHTS RESERVED
 *
 */
*****/

/*
 * This sample program displays the access permissions (ACL) on each
 * entry in the directory for a specific user. The permissions are
 * presented in a form similar to the UNIX file permissions. In
 * addition, each entry is flagged as either a master or a shadow copy.
 *
 * The distinguished name of the user performing the check is:
 *
 *      /C=de/O=sni/OU=ap/CN=norbert
 *
 * The results are presented in the following format:
 *
 *      [ABCD] <entry's distinguished name>
 *
 *      A:  'm' master copy
 *          's' shadow copy
 *
 *      B:  'r' read access to public attributes
 *          'w' write access to public attributes
 */
```



```

*      '-' no access to public attributes
*
*      C: 'r' read access to standard attributes
*          'w' write access to standard attributes
*          '-' no access to standard attributes
*
*      D: 'r' read access to sensitive attributes
*          'w' write access to sensitive attributes
*          '-' no access to sensitive attributes
*
* For example, the following result means that the entry
* '/C=de/O=sni' is a master copy and that the requesting user
* (/C=de/O=sni/OU=ap/CN=norbert) has write access to its public
* attributes, read access to its standard attributes and no access
* to its sensitive attributes.
*
*      [mwr-] /C=de/O=sni
*
* The program requires that the specific user perform an authenticated
* bind to the directory. In order to achieve this the user's
* credentials must already exist in the directory. Therefore the
* following tree of 6 entries is added to the directory each time the
* program runs, and removed again afterwards.
*
*      O C=de
*      | (objectClass=Country,
*      |   ACL=(mod-pub: *
*      |       read-std:*
*      |       mod-std: *
*      |       read-sen:*
*      |       mod-sen: *))
*      |
*      O O=sni
*      | (objectClass=Organization,
*      |   ACL=(mod-pub: /C=de/O=sni/OU=ap/*
*      |       read-std:/C=de/O=sni/OU=ap/CN=stefanie
*      |       mod-std: /C=de/O=sni/OU=ap/CN=stefanie
*      |       read-sen:/C=de/O=sni/OU=ap/CN=stefanie
*      |       mod-sen: /C=de/O=sni/OU=ap/CN=stefanie))

```

```

*      |
*      O OU=ap
*      | (objectClass=OrganizationalUnit,
*      |   ACL=(mod-pub: /C=de/O=sni/OU=ap/*
*      |         read-std:/C=de/O=sni/OU=ap/CN=stefanie
*      |         mod-std: /C=de/O=sni/OU=ap/CN=stefanie
*      |         read-sen:/C=de/O=sni/OU=ap/CN=stefanie
*      |         mod-sen: /C=de/O=sni/OU=ap/CN=stefanie))
*      |
*      +-----+-----+
*      |      |      |
*      |      O CN=ingrid
*      |      | (objectClass=OrganizationalPerson,
*      |      |   ACL=(mod-pub: /C=de/O=sni/OU=ap/*
*      |      |         read-std:/C=de/O=sni/OU=ap/*
*      |      |         mod-std: /C=de/O=sni/OU=ap/CN=stefanie
*      |      |         read-sen:/C=de/O=sni/OU=ap/*
*      |      |         mod-sen: /C=de/O=sni/OU=ap/CN=stefanie),
*      |      |   surname="Schmid",
*      |      |   telephone="+49 89 636 0",
*      |      |   userPassword="secret")
*      |      |
*      |      O CN=norbert
*      |      | (objectClass=OrganizationalPerson,
*      |      |   ACL=(mod-pub: /C=de/O=sni/OU=ap/*
*      |      |         read-std:/C=de/O=sni/OU=ap/*
*      |      |         mod-std: /C=de/O=sni/OU=ap/CN=stefanie
*      |      |         read-sen:/C=de/O=sni/OU=ap/*
*      |      |         mod-sen: /C=de/O=sni/OU=ap/CN=stefanie),
*      |      |   surname="Schmid",
*      |      |   telephone="+49 89 636 0",
*      |      |   userPassword="secret")
*      |      |
*      O CN=stefanie
*      | (objectClass=OrganizationalPerson,
*      |   ACL=(mod-pub: /C=de/O=sni/OU=ap/*
*      |         read-std:/C=de/O=sni/OU=ap/*
*      |         mod-std: /C=de/O=sni/OU=ap/CN=stefanie
*      |         read-sen:/C=de/O=sni/OU=ap/*
*      |         mod-sen: /C=de/O=sni/OU=ap/CN=stefanie),

```

```

*      surname="Schmid",
*      telephone="+49 89 636 0",
*      userPassword="secret")
*
*
* In this version of the program, instead of providing client-generated
* public objects, the XOM Convenience Functions are used for creating
* objects. They are also used for extracting information from service
* generated objects.
*/

#ifdef THREADSAFE
#include <pthread.h>
#endif

#include <stdio.h>
#include <xom.h>
#include <xds.h>
#include <xdsbdcp.h>
#include <xdsfds.h>
#include <xdsfds.h>
#include <xdsext.h>      /* convenience functions header file */
#include <xomext.h>      /* convenience functions header file */
#include "acl2.h"

void
main(
    int  argc,
    char *argv[]
)
{
    OM_workspace      workspace;      /* Workspace for objects */
    OM_private_object session;        /* Session object. */
    OM_private_object bound_session; /* Holds the Session object which */
                                     /* is returned by ds_bind() */
    OM_private_object context;        /* Context object. */
    OM_private_object result;         /* Holds the search result object.*/
    OM_sint           invoke_id;      /* Integer for the invoke id */
                                     /* returned by ds_search(). */
                                     /* (this parameter must be present*/

```

```
/* even though it is ignored). */
OM_type      navigation_path[] = { DS_SEARCH_INFO, 0 };
/* List of OM types to the target */
/* object - of the search result */
OM_type      entry_list[] = { DS_ENTRIES, 0 };
/* List of types to be extracted */
OM_public_object entry; /* Entry object from search info. */
OM_value_position total_num; /* Number of descriptors returned.*/
OM_return_code rc; /* XOM function return code. */
register int i;
char user_name[MAX_DN_LEN] = DN_NORBERT;
/* Holds the requestor's name - */
/* "/C=de/O=sni/OU=ap/CN=norbert" */
char entry_string[MAX_DN_LEN + 7] = "[?r??] ";
/* Holds entry details. */
struct entry entry_array[6]; /* List of entry names and attrs */
OM_object credentials; /* Credentials part of session obj*/
OM_object use_copy; /* Specifies whether to use shadow*/
/* entries, in context object */
OM_object filter; /* Filter - for search operation */
OM_object dn_root; /* Name object for "/" */
OM_object selection_acl; /* Entry Information */
/* Selection obj */

static char *name_list[] =
{ DN_DE, DN_SNI, DN_AP, DN_STEFANIE,
  DN_NORBERT, DN_INGRID };
/* Array of names to be added */

static char *C_attr_list[] = { OBJ_CLASS_C };
static char *O_attr_list[] = { OBJ_CLASS_O, ATT_ACL1 };
static char *OU_attr_list[] = { OBJ_CLASS_OU };
static char *OP_attr_list[] = { OBJ_CLASS_OP, ATT_ACL2,
  ATT_SURNAME, ATT_PHONE_NUM, ATT_PASSWORD };
/* Attribute lists, in string fmt */

static char *dn_root_str = DN_ROOT;
static char *filter_str = FILTER;

/* Step 3
```

```

*
* Initialize a directory workspace for use by XOM.
*/
if ((workspace = ds_initialize()) == (OM_workspace)0)
    printf("ds_initialize() error\n");

/* Step 4
*
* Negotiate the use of the BDC and GDS packages.
*/
if (ds_version(features, workspace) != DS_SUCCESS)
    printf("ds_version() error\n");

/* Step 5
*
* Build name objects for entries to be added to the directory.
*/
for (i = 0; i < NO_OF_ENTRIES; i++)
    if (! build_name_object(workspace, name_list[i],
                           &(entry_array[i].name)))
        printf("build_name_object() error\n");

/* Step 6
*
* Build attribute objects for entries to be added to the directory
*/
if ((! build_attr_list_object(workspace, NO_C_ATTRS, C_attr_list,
                              &entry_array[0].attr_list)) ||
    (! build_attr_list_object(workspace, NO_O_ATTRS, O_attr_list,
                              &entry_array[1].attr_list)) ||
    (! build_attr_list_object(workspace, NO_OU_ATTRS, OU_attr_list,
                              &entry_array[2].attr_list)) ||
    (! build_attr_list_object(workspace, NO_OP_ATTRS, OP_attr_list,
                              &entry_array[3].attr_list)))
    printf("build_attr_list_object() error\n");

/*
* These entries also have the OP attribute list.
*/
entry_array[4].attr_list = entry_array[3].attr_list;

```

```
entry_array[5].attr_list = entry_array[3].attr_list;

/* Step 7
 *
 * Add a fixed tree of entries to the directory in order to permit
 * an authenticated bind by: /C=de/O=sni/OU=ap/CN=norbert
 */
if (! add_tree(workspace, entry_array, NO_OF_ENTRIES))
    printf("add_tree() error\n");

/* Step 8
 *
 * Create a default session object.
 */
if ((rc = om_create(DSX_C_GDS_SESSION, OM_TRUE, workspace, &session))
    != OM_SUCCESS)
    printf("om_create() error %d\n", rc);

/* Step 9
 *
 * Build an object with the following credentials:
 * requestor: /C=de/O=sni/OU=ap/CN=norbert
 * password: "secret"
 * authentication mechanism: simple
 */
if (! build_credentials_object(entry_array[4].name, &credentials))
    printf("build_credentials_object() error\n");

/*
 * Alter the default session object to include the credentials
 */
if ((rc = om_put(session, OM_REPLACE_ALL, credentials, 0, 0, 0))
    != OM_SUCCESS)
    printf("om_put() error %d\n", rc);

/* Step 10
 *
 * Bind with credentials to the default GDS server. The
 * returned session object is stored in the private object variable
 * bound_session and is used for all further XDS function calls.
```

```
*/
if (ds_bind(session, workspace, &bound_session) != DS_SUCCESS)
    printf("ds_bind() error\n");

/* Step 11
 *
 * Create a default context object.
 */
if ((rc = om_create(DSX_C_GDS_CONTEXT, OM_TRUE, workspace, &context))
    != OM_SUCCESS)
    printf("om_create() error %d\n", rc);

/*
 * Build an object specifying that shadow entries should be used.
 */
if (! build_use_copy_object(&use_copy))
    printf("build_use_copy_object() error\n");

/*
 * Alter the default context object to include 'shadow' entries.
 */
if ((rc = om_put(context, OM_REPLACE_ALL, use_copy, 0, 0, 0))
    != OM_SUCCESS) printf("om_put() error %d\n", rc);

/* Step 12
 *
 * Build a filter object, specifying presence of object class attr.
 */
if (! build_filter_object(workspace, filter_str, &filter))
    printf("build_filter_object() error\n");

/*
 * Build a root name object, name = "/"
 */
if (! build_name_object(workspace, dn_root_str, &dn_root))
    printf("build_name_object() error\n");

/*
 * Build an entry information selection object,
 * selecting acl attributes.
```

```
*/
if (! build_selection_object(&selection_acl))
    printf("build_selection_object() error\n");

/* Step 13
*
* Search the whole subtree below root. The filter selects entries
* with an object-class attribute. The selection extracts the ACL
* attribute from each selected entry. The results are returned in
* the private object 'result'.
*
* NOTE: Since every entry contains an object-class attribute the
*       filter performs no function other than to demonstrate how
*       filters may be used.
*/

if (ds_search(bound_session, context, dn_root, DS_WHOLE_SUBTREE, filter,
             OM_FALSE, selection_acl, &result, &invoke_id) != DS_SUCCESS)
    printf("ds_search() error\n");

/* Step 14
*
* Close the connection to the GDS server.
*/
if (ds_unbind(bound_session) != DS_SUCCESS)
    printf("ds_unbind() error\n");

/* Step 15
*
* Remove the user's credentials from the directory.
*/
if (! remove_tree(workspace, session, entry_array, NO_OF_ENTRIES))
    printf("remove_tree() error\n");

/* Step 16
*
* Free the name and attribute objects
* which make up the directory entries.
*/
if (! free_entry_list(entry_array))
```



```
    printf("free_entry_list() error\n");

/*
 * Free public objects which were created.
 */
free(selection_acl);
free(use_copy);
free(credentials);

if ((om_delete(filter) != OM_SUCCESS) ||
    (om_delete(dn_root) != OM_SUCCESS))
    printf("om_delete() error\n");

/* Step 17
 *
 * Extract components from the search result by means of the XOM
 * Convenience Function, omX_extract()
 */
if ((rc = omX_extract(result, navigation_path,
                     OM_EXCLUDE_ALL_BUT_THESE_TYPES + OM_EXCLUDE_SUBOBJECTS,
                     entry_list, OM_FALSE, 0, 0, &entry, &total_num))
    != OM_SUCCESS)
    printf("omX_extract(Search-Result) error %d\n", rc);

/*
 * Requestor's name = "/C=de/O=sni/OU=ap/CN=norbert"
 */
printf("User:  %s\nTotal: %d\n", user_name, total_num);

/* Step 18
 *
 * Examine each entry and print the entry details.
 */
for (i = 0; i < total_num; i++) {
    if (process_entry_info((entry+i)->value.object.object,
                          entry_string, user_name))
        printf("%s\n", entry_string);
}

/*
```

```
    * Now free the entry object (returned from omX_extract() ).
    */
    if (om_delete(entry) != OM_SUCCESS)
        printf("om_delete() error\n");

/* Step 19
 *
 * Close the directory workspace.
 */
    if (ds_shutdown(workspace) != DS_SUCCESS)
        printf("ds_shutdown() error\n");
}

/*
 * Add the tree of entries described above.
 */
int
add_tree(
    OM_workspace workspace,
    struct entry elist[],
    int          no_entries
)
{
    OM_private_object session;      /* Holds the Session object which */
                                   /* is returned by ds_bind()      */
    OM_sint          invoke_id;     /* Integer for the invoke id    */
    int              error = 0;
    int              i;

/*
 * Bind (without credentials) to the default GDS server.
 */
    if (ds_bind(DS_DEFAULT_SESSION, workspace, &session) != DS_SUCCESS)
        error++;

/*
 * Add entries to the GDS server.
 */
    for (i = 0; i < no_entries; i++)
        if (ds_add_entry(session, DS_DEFAULT_CONTEXT, elist[i].name,
```

```

        elist[i].attr_list, &invoke_id) != DS_SUCCESS) {
/* Ignore error if adding country - possibly already there */
    if (i != 0) error++;
}

/*
 * Close the connection to the GDS server.
 */
if (ds_unbind(session) != DS_SUCCESS)
    error++;

return (error?0:1);
}

/*
 * Remove the tree of entries described above.
 */
int
remove_tree(
    OM_workspace      workspace,
    OM_private_object session,
    struct entry      elist[],
    int               no_entries
)
{
    OM_private_object bound_session; /* Holds the Session object which */
                                     /* is returned by ds_bind() */
    OM_sint           invoke_id;     /* Integer for the invoke id */
    int               i;
    int               error = 0;

/*
 * Bind (without credentials) to the default GDS server.
 */
if (ds_bind(session, workspace, &bound_session) != DS_SUCCESS)
    error++;

/*
 * Remove entries from the GDS server.
 */

```

```
for (i = no_entries-1; i >= 0; i--)
    if (ds_remove_entry(bound_session, DS_DEFAULT_CONTEXT,
        elist[i].name, &invoke_id) != DS_SUCCESS) {
        /* Ignore error if removing country - possibly has entries */
        /* below it */
        if (i != 0) error++;
    }

/*
 * Close the connection to the GDS server.
 */
if (ds_unbind(bound_session) != DS_SUCCESS)
    error++;

return (error?0:1);
}

/*
 * Extract information about an entry from the Entry-Info object: whether
 * the entry is a master-copy, its ACL permissions and its distinguished
 * name. Build up a string based on this information.
 */
int
process_entry_info(
    OM_private_object entry,
    char *entry_string,
    char *user_name
)
{
    OM_return_code rc;          /* Return code from XOM function. */
    OM_public_object ei_attrs;  /* Components from Entry-Info. */
    OM_public_object attr;      /* Directory attribute. */
    OM_public_object acl;       /* ACL attribute parts. */
    OM_public_object acl_vals;  /* ACL attribute value. */
    OM_public_object acl_item;  /* ACL item component. */
    OM_value_position total_attrs; /* Number of attributes returned. */
    OM_value_position total_acls; /* Number of acl values returned. */
    register int i;
    register int interp;
    register int error = 0;
```

```

register int      found_acl = 0;
static OM_type    ei_attr_list[] = { DS_FROM_ENTRY,
                                      DS_OBJECT_NAME,
                                      0 };
                                      /* Attributes to be extracted. */

OM_string         entry_str;

/*
 * Extract occurrences of DS_FROM_ENTRY, and DS_OBJECT_NAME
 * from each Entry-Info object.
 */
if ((rc = om_get(entry, OM_EXCLUDE_ALL_BUT_THESE_TYPES,
                ei_attr_list, OM_FALSE, 0, 0, &ei_attrs, &total_attrs))
    != OM_SUCCESS) {
    error++;
    printf("om_get(Entry-Info) error %d\n", rc);
}

for (i = 0; ((i < total_attrs) && (! error)); i++, ei_attrs++) {

    /*
     * Determine if current entry is a master-copy or a shadow-copy.
     */
    if ((ei_attrs->type == DS_FROM_ENTRY) &&
        ((ei_attrs->syntax & OM_S_SYNTAX) == OM_S_BOOLEAN))
        if (ei_attrs->value.boolean == OM_TRUE)
            entry_string[1] = 'm';
        else if (ei_attrs->value.boolean == OM_FALSE)
            entry_string[1] = 's';
        else
            entry_string[1] = '?';

    /*
     * Convert the entry's distinguished name to a string format.
     */
    entry_str.elements = &entry_string[7];
    entry_str.length = MAX_DN_LEN;
    if ((ei_attrs->type == DS_OBJECT_NAME) &&
        ((ei_attrs->syntax & OM_S_SYNTAX) == OM_S_OBJECT))
        if ((rc = omX_object_to_string(ei_attrs->value.object.object,

```

```
        OM_FALSE, &entry_str)) != OM_SUCCESS) {
            error++;
            printf("omX_object_to_string() error\n");
        }
    }

    /*
     * Now extract occurrences of attributes, where the attribute
     * type is ACL from the Entry-Info object.
     */
    dsX_extract_attr_values(entry, DSX_A_ACL, OM_TRUE,
        &acl_vals, &total_acls);

    for (i = 0; ((i < total_acls) && (! error)); i++) {
        acl = acl_vals[i].value.object.object;

        /*
         * Examine the ACL. Check each permission for the current user.
         */

        entry_string[2] = 'r';
        entry_string[3] = '-';
        entry_string[4] = '-';

        while (acl->type != OM_NO_MORE_TYPES) {

            if ((acl->syntax & OM_S_SYNTAX) == OM_S_OBJECT)
                acl_item = acl->value.object.object;

            switch (acl->type) {

                case OM_CLASS:
                    break;

                case DSX_MODIFY_PUBLIC:
                    if (permitted_access(user_name, acl_item))
                        entry_string[2] = 'w';
                    break;

                case DSX_READ_STANDARD:
```

```
        if (permitted_access(user_name, acl_item))
            entry_string[3] = 'r';
        break;

    case DSX_MODIFY_STANDARD:
        if (permitted_access(user_name, acl_item))
            entry_string[3] = 'w';
        break;

    case DSX_READ_SENSITIVE:
        if (permitted_access(user_name, acl_item))
            entry_string[4] = 'r';
        break;

    case DSX_MODIFY_SENSITIVE:
        if (permitted_access(user_name, acl_item))
            entry_string[4] = 'w';
        break;
    }
    acl++;
}

}
/*
 * Now free acl_vals.
 */
if (total_acls > 0)
    if ((rc = om_delete(acl_vals)) != OM_SUCCESS) {
        error++;
        printf("om_delete() error, rc = %d\n", rc);
    }

return (error?0:1);
}

/*
 * Check if a user is permitted access based on the ACL supplied.
 */
int
permitted_access(
    char          *user_name,
```

```
    OM_public_object  acl_item
)
{
    char                acl_name[MAX_DN_LEN];
    OM_string           acl_name_str;
    int                 interpretation;
    int                 acl_present = 0;
    int                 access = 0;
    int                 acl_name_length;
    OM_return_code      rc;

    while (acl_item->type != OM_NO_MORE_TYPES) {

        switch (acl_item->type) {
        case OM_CLASS:
            break;

        case DSX_INTERPRETATION:
            interpretation = acl_item->value.boolean;
            break;

        case DSX_USER:
            acl_name_str.elements = acl_name;
            if ((rc = omX_object_to_string(acl_item->value.object.object,
                OM_FALSE, &acl_name_str)) == OM_SUCCESS) {
                if (interpretation == DSX_SINGLE_OBJECT) {
                    if (strcmp(acl_name, user_name) == 0)
                        access = 1;
                }
                else if (interpretation == DSX_ROOT_OF_SUBTREE) {
                    if ((acl_name_length = strlen(acl_name)) == 0)
                        access = 1;
                    else if (strncmp(acl_name, user_name, acl_name_length)
                        == 0)
                        access = 1;
                }
            }
            break;
        }
        acl_item++;
    }
}
```



```
    }

    return (access);
}

/*
 * Build a name object from a name string using the XOM
 * Convenience Function omX_string_to_object().
 */
int
build_name_object(
    OM_workspace    workspace,
    char            *name,
    OM_private_object *name_obj
)
{
    OM_integer      err_pos;
    OM_integer      err_type;
    OM_return_code   rc;
    OM_string        name_str;
    int              error = 0;

    name_str.length = strlen(name);
    name_str.elements = name;
    if ((rc = omX_string_to_object(workspace, &name_str, DS_C_DS_DN,
                                   OM_TRUE, name_obj, &err_pos, &err_type)) != OM_SUCCESS)
        error++;

    return (error?0:1);
}

/*
 * Build an attribute list object given a list of attribute strings.
 * Use the XOM Convenience Function omX_string_to_object() to build
 * an attribute object from an attribute string, and omX_fill() to
 * create the other OM descriptor required.
 */
int
build_attr_list_object(
```

```
OM_workspace workspace,
OM_integer no_attrs,
char *attr_str_array[],
OM_object *attr_list_obj
)
{
    OM_integer err_pos;
    OM_integer err_type;
    OM_object attr;
    OM_object alist;
    OM_string attr_str;
    OM_return_code rc;
    OM_descriptor null_desc = OM_NULL_DESCRIPTOR;
    int error = 0;
    int i;

    /*
     * Allocate space for class descriptor, null descriptor and
     * one descriptor for each attribute.
     */
    if ((alist =
        (OM_descriptor *)malloc((2+no_attrs) * sizeof(OM_descriptor)))
        == 0)
        error++;

    if ((rc = omX_fill_oid(OM_CLASS, DS_C_ATTRIBUTE_LIST, &alist[0]))
        != OM_SUCCESS)
        error++;

    for (i = 1; i <= no_attrs; i++) {

        attr_str.length = strlen(attr_str_array[i-1]);
        attr_str.elements = attr_str_array[i-1];
        if ((rc = omX_string_to_object(workspace, &attr_str, DS_C_ATTRIBUTE,
            OM_TRUE, &attr, &err_pos, &err_type)) != OM_SUCCESS)
            error++;

        if ((rc = omX_fill(DS_ATTRIBUTES, OM_S_OBJECT, 0, attr, &alist[i]))
            != OM_SUCCESS)
            error++;
    }
}
```

```
    }

    alist[i] = null_desc;

    *attr_list_obj = alist;
    return (error?0:1);
}

/*
 * Build an entry info selection object using the XOM Convenience
 * Functions omX_fill() and omX_fill_oid() to fill the OM descriptors.
 */
int
build_selection_object(
    OM_object *selection_obj
)
{
    OM_integer      err_pos;
    OM_integer      err_type;
    OM_object       desc;
    OM_object       sel;
    OM_return_code   rc;
    OM_descriptor    null_desc = OM_NULL_DESCRIPTOR;
    int             error = 0;

    /*
     * Allocate space for class descriptor, null descriptor and one
     * descriptor for each attribute.
     */
    if ((sel = (OM_descriptor *)malloc((5) * sizeof(OM_descriptor))) == 0)
        error++;

    if ((rc = omX_fill_oid(OM_CLASS, DS_C_ENTRY_INFO_SELECTION, &sel[0]))
        != OM_SUCCESS)
        error++;

    if ((rc = omX_fill(DS_ALL_ATTRIBUTES, OM_S_BOOLEAN, OM_FALSE, 0,
                      &sel[1])) != OM_SUCCESS)
        error++;
}
```

```
if ((rc = omX_fill_oid(DS_ATTRIBUTES_SELECTED, DSX_A_ACL,
                      &sel[2])) != OM_SUCCESS)
    error++;

if ((rc = omX_fill(DS_INFO_TYPE, OM_S_ENUMERATION, DS_TYPES_AND_VALUES,
                  0, &sel[3])) != OM_SUCCESS)
    error++;

sel[4] = null_desc;

*selection_obj = sel;
return (error?0:1);
}

/*
 * Build a credentials object using the XOM Convenience Function
 * omX_fill().
 */
int
build_credentials_object(
    OM_object  name,
    OM_object  *credentials_obj
)
{
    OM_integer  err_pos;
    OM_integer  err_type;
    OM_object  cred;
    OM_return_code  rc;
    OM_descriptor  null_desc = OM_NULL_DESCRIPTOR;
    int  error = 0;

    /*
     * Just allocate space for a null descriptor and two other descriptors,
     * no class descriptor required.
     */
    if ((cred = (OM_descriptor *)malloc((4) * sizeof(OM_descriptor))) == 0)
        error++;

    if ((rc = omX_fill(DS_REQUESTOR, OM_S_OBJECT, 0, name, &cred[0]))
```

```

                                                                    != OM_SUCCESS)

    error++;

    if ((rc = omX_fill(DSX_PASSWORD, OM_S_OCTET_STRING, (sizeof(PASSWD)-1),
                      PASSWD, &cred[1])) != OM_SUCCESS)

    if ((rc = omX_fill(DSX_AUTH_MECHANISM, OM_S_ENUMERATION, DSX_SIMPLE,
                      0, &cred[2])) != OM_SUCCESS)

        error++;

    cred[3] = null_desc;

    *credentials_obj = cred;
    return (error?0:1);
}

/*
 * Build an object setting DS_DONT_USE_COPY to FALSE, using the
 * XOM Convenience Function omX_fill().
 */
int
build_use_copy_object(
    OM_object      *use_copy_obj
)
{
    OM_integer      err_pos;
    OM_integer      err_type;
    OM_object      desc;
    OM_object      copy;
    OM_return_code   rc;
    OM_descriptor    null_desc = OM_NULL_DESCRIPTOR;
    int             error = 0;

    /*
     * Just allocate space for a null descriptor and one other
     * descriptor, no class descriptor required.
     */
    if ((copy = (OM_descriptor *)malloc((2) * sizeof(OM_descriptor))) == 0)
        error++;

```

```
        if ((rc = omX_fill(DS_DONT_USE_COPY, OM_S_BOOLEAN, OM_FALSE, 0,
                          &copy[0])) != OM_SUCCESS)
            error++;

        copy[1] = null_desc;

        *use_copy_obj = copy;
        return (error?0:1);
    }

/*
 * Build a filter object from a filter string using the XOM Convenience
 * Function omX_string_to_object().
 */
int
build_filter_object(
    OM_workspace workspace,
    char    *filter,
    OM_object *filter_obj
)
{
    OM_integer    err_pos;
    OM_integer    err_type;
    OM_string     filter_str;
    OM_return_code rc;
    int           error = 0;

    filter_str.length = strlen(filter);
    filter_str.elements = filter;
    if ((rc = omX_string_to_object(workspace, &filter_str, DS_C_FILTER,
                                   OM_TRUE, filter_obj, &err_pos, &err_type)) != OM_SUCCESS)
        error++;

    return (error?0:1);
}

/*
 * Free the name and attribute list objects in the entry list.  Objects
```

```

    * which have been created using the XOM Convenience Function
    * omX_string_to_object() must be deleted using om_delete().
    */
int
free_entry_list(
    struct entry    entry_array[]
)
{
    OM_object        attr_list_obj;
    int              i, j;
    int              error = 0;

    for (i = 0; i < NO_OF_ENTRIES; i++) {

        /*
         * Delete the service generated public name object .
         */
        if (om_delete(entry_array[i].name) != OM_SUCCESS)
            error++;

        /*
         * The last two attribute lists were the same as the 4th one.
         */
        if (i < NO_OF_ENTRIES-2) {
            attr_list_obj = entry_array[i].attr_list;
            for (j = 0; attr_list_obj[j].type != OM_NO_MORE_TYPES; j++) {
                if (attr_list_obj[j].type == DS_ATTRIBUTES)

                    /*
                     * Delete the service generated public attribute object.
                     */
                    if (om_delete(attr_list_obj[j].value.object.object)
                        != OM_SUCCESS)
                        error++;
            }

            /*
             * Free the whole attribute list object.
             */
            free(attr_list_obj);
        }
    }
}
```

```
    }  
}  
  
return (error?0:1);  
}
```

9.2.2 The **acl2.h** Header File

The **acl2.h** header file performs the following:

1. It exports the object identifiers that **acl2.c** requires.
2. It declares a structure to contain the name and attributes of directory entries.
3. It defines abbreviated names for the directory entries.
4. It defines abbreviated names for the directory attributes.
5. It builds the descriptor list for optional packages that are to be negotiated.

The following code is a listing of the **acl2.h** file:

```
/*  
*****  
*  
*   COPYRIGHT (C) SIEMENS NIXDORF INFORMATIONSSYSTEME AG 1991   *  
*           ALL RIGHTS RESERVED                                   *  
*  
*****  
*/  
  
#ifndef _ACL2_H  
#define _ACL2_H  
  
#define MAX_DN_LEN 100 /* max length of a distinguished name in */  
                        /* string format.                          */  
  
/* Step 1 */  
  
/* The application must export the object identifiers it requires. */
```



```

OM_EXPORT (DS_C_DS_DN)
OM_EXPORT (DS_C_ENTRY_INFO_SELECTION)
OM_EXPORT (DS_C_ATTRIBUTE)
OM_EXPORT (DS_C_ATTRIBUTE_LIST)
OM_EXPORT (DS_C_FILTER)
OM_EXPORT (DSX_C_GDS_SESSION)
OM_EXPORT (DSX_C_GDS_CONTEXT)
OM_EXPORT (DSX_A_ACL)

/* Structure to contain the name and attribute list          */
/* of a directory entry.                                     */

struct entry {
    OM_private_object  name;
    OM_object          attr_list;
} Entry;

/* Step 2 */
/*
 * Names of directory entries, in string format.
 */
#define DN_ROOT          "/"
#define DN_DE            "/C=de"
#define DN_SNI           "/C=de/O=sni"
#define DN_AP            "/C=de/O=sni/OU=ap"
#define DN_STEFANIE      "/C=de/O=sni/OU=ap/CN=stefanie"
#define DN_NORBERT       "/C=de/O=sni/OU=ap/CN=norbert"
#define DN_INGRID        "/C=de/O=sni/OU=ap/CN=ingrid"

/*
 * Attributes, in string format.
 */
#define OBJ_CLASS_C       "OCL = TOP; C"
#define OBJ_CLASS_O       "OCL = TOP; ORG"
#define OBJ_CLASS_OU      "OCL = TOP; OU"
#define OBJ_CLASS_OP      "OCL = TOP; PER; ORP"
#define ATT_PHONE_NUM     "TN = '+49 89 636 0' "
#define ATT_PASSWORD      "UP = secret"
#define ATT_SURNAME       "SN = Schmid"
#define ATT_ACL1          "ACL={MPUB = {INT = 1,USR = {/}}, \

```

```
RSTD = {INT = 0,USR = {/C=de/O=sni/OU=ap/CN=stefanie}},\
MSTD = {INT = 0,USR = {/C=de/O=sni/OU=ap/CN=stefanie}},\
RSEN = {INT = 0,USR = {/C=de/O=sni/OU=ap/CN=stefanie}},\
MSEN = {INT = 0,USR = {/C=de/O=sni/OU=ap/CN=stefanie}}}"
#define ATT_ACL2  "ACL={MPUB = {INT = 1,USR = {/C=de/O=sni/OU=ap}},\
RSTD = {INT = 1,USR = {/C=de/O=sni/OU=ap}},\
MSTD = {INT = 0,USR = {/C=de/O=sni/OU=ap/CN=stefanie}},\
RSEN = {INT = 1,USR = {/C=de/O=sni/OU=ap}},\
MSEN = {INT = 0,USR = {/C=de/O=sni/OU=ap/CN=stefanie}}}"

/* Other strings.                                     */
#define PASSWD      "secret"
#define FILTER      "OCL"

#define NO_OF_ENTRIES  6  /* 6 entries to be added          */
#define NO_C_ATTRS     1  /* 1 attr in Country attribute list      */
#define NO_O_ATTRS     2  /* 2 attr in Org attribute list          */
#define NO_OU_ATTRS    1  /* 1 attr in Org-Unit attribute list     */
#define NO_OP_ATTRS    5  /* 5 attr in Org-Person attribute list*/

/* Build up an array of object identifiers for the optional */
/* packages to be negotiated.                                */
DS_feature features[] = {
    { OM_STRING(OMP_O_DS_BASIC_DIR_CONTENTS_PKG), OM_TRUE },
    { OM_STRING(OMP_O_DSX_GDS_PKG), OM_TRUE },
    { 0 }
};

#endif /* _ACL2_H */
```

9.2.3 Example Strings

This section contains examples of input strings to **omX_string_to_object()** and some examples of strings that can be returned by **omX_object_to_string()**.

9.2.3.1 Input Strings to **omX_string_to_object()**

The following are examples of strings that can be handled by the **omX_string_to_object()** function.

9.2.3.1.1 Example 1

To create a **DS_C_DS_DN** object (root), use strings like the following:

```
/
/...
```

9.2.3.1.2 Example 2

To create other **DS_C_DS_DN** objects, use strings like the following:

```
/c=de/o=sni/ou=ap11/cn=naik,sn=naik
/c=de/o=sni/ou=ap11/85.4.3=naik,sn=naik
/c=de/o=sni/ou=ap11/cn=naik,sn=na\x69k
/c=de/o=sni/ou=ap11/cn=naik,loc=Muenchen\,8000
/c=de/o=sni/ou=ap11/cn=naik,loc='Muenchen,8000'
/ C = de / O = sni / Ou = ap11/CN=naik, SN=naik
```

9.2.3.1.3 Example 3

To create a **DS_C_DS_DN** object (DCE name), use a string like the following:

```
/.../c=us/o=osf/ou=abc/subsystems/server/xyz
```

9.2.3.1.4 Example 4

To create a **DS_C_DS_RDN** object, use strings like the following:

```
cn=naik,sn=naik
cn=naik,sn=na\x69k
```

CN = naik, SN = naik

9.2.3.1.5 Example 5

To create a **DS_C_DS_RDN** object (DCE name), use a string like the following:

server

9.2.3.1.6 Example 6

To create a **DS_C_ATTRIBUTE** object (containing, for example, **Common-Name**), use strings like the following:

cn=bhavesh naik
CN = bhavesh naik
85.4.3=bhavesh nai\x69k

9.2.3.1.7 Example 7

To create a **DS_C_ATTRIBUTE** object (containing an object class with multiple values of **Residential-Person** and **Organizational-Person**), use strings like the following:

OCL=REP;ORP
OCL = '\x55\x06\x0a' ; '\x55\x06\x07'

9.2.3.1.8 Example 8

To create a **DS_C_ATTRIBUTE** object (containing a GDS structured attribute like **Telex-Number** or **Owner**), use strings like the following:

TXN={TN=12345,CC=678,AB=90}
TXN = { TN = 12345, CC = 678, AB = 90}
own={/c=de/o=sni/ou=ap11};{/c=de/o=sni/ou=ap22}
pa={pa='Wilhelm Riehl Str.85';'Munich'}

9.2.3.1.9 Example 9

To create a **DSX_C_GDS_ACL** object, use a string like the following:

```
MPUB={INT=0, USR={/c=de/o=sni/cn=naik, sn=bhaves}}
```

9.2.3.1.10 Example 10

To create a **DS_C_PRESENTATION_ADDRESS** object, use a string like the following:

```
TS=Server,NA='TCP/IP!internet=127.0.0.1+port=25015'
```

9.2.3.1.11 Example 11

To create a **DS_C_FILTER** object, use strings like the following:

```
c  
!c  
C = de && CN = 'bha\x76esh naik'  
c=de&&cn~=mueller  
c = de && (cn = 'a*' || cn = b* || cn = c* )  
ACL={MPUB={INT=0,USR={/c=de/o=sni/cn=naik, sn=bhaves}}}  
c = de || cn = *aa*bb*cc*  
(cn ~=naik)&&((OCL=ORP)||((OCL=REP))&& !(SN='bhaves naik')&&(L=*))
```

9.2.3.1.12 Example 12

The following is an example of the error return when an erroneous string is supplied:

```
/c=de/o=sni,=de
```

The **OM_return_code** would be **OM_WRONG_VALUE_MAKEUP**.

The **error_type** would be **OMX_MISSING_ABBRV**.

The **error_position** would be 13.

9.2.3.2 Strings Returned by **omX_object_to_string()**

The following are examples of strings returned by the **omX_object_to_string()** function.

9.2.3.2.1 Example 1

If a **DS_C_DS_DN** object is supplied, the following might be returned:

```
/
/C=de/O=sni/OU=ap11/CN=naik,SN=naik
/C=de/O=sni/OU=ap11/CN=naik,LOC=Muenchen\,8000
```

9.2.3.2.2 Example 2

If a **DS_C_DS_RDN** object is supplied, the following might be returned:

```
CN=naik,SN=naik
server
```

9.2.3.2.3 Example 3

If a **DS_C_ATTRIBUTE** object is supplied, the following might be returned:

```
CN=bhavesh naik
OCL=REP;ORP
TXN={AB=90,CC=678,TN=12345}
OWN={/C=de/O=sni/OU=ap11};{/C=de/O=sni/OU=ap22}
```

9.2.3.2.4 Example 4

If a **DSX_C_GDS_ACL** object is supplied, the following might be returned:

```
MPUB={INT=0,USR={/C=de/O=sni/CN=naik,SN=bhavesh}}
```

9.2.3.2.5 Example 5

If a **DS_C_NAME_ERROR** object is supplied with **DS_PROBLEM** of **DS_E_NO_SUCH_OBJECT**, the following might be returned:

The specified name does not match the name of any object
in the directory

9.2.3.2.6 Example 6

If a **DS_C_ATTRIBUTE_ERROR** object is supplied
with **DS_C_ATTRIBUTE_PROBLEM** containing
DS_E_ATTRIBUTE_OR_VALUE_EXISTS, the following might
be returned:

An attempt is made to add an attribute or value that already
exists. Violating Attribute - Telephone-Number

Part 4

XDS/XOM Supplementary Information

Part 4 contains reference material for the X/Open Object Management (XOM) programming interface.

Chapter 10

XDS Interface Description

The XDS interface comprises a number of functions, together with many OM classes of OM objects, which are used as the parameters and results of the functions. Both the functions and the OM objects are based closely on the abstract service that is specified in the standards (see *The Directory: Abstract Service Definition*, ISO 9594-3, CCITT X.511).

The interface models the directory interactions as service requests made through a number of interface functions, which take a number of input parameters. Each valid request causes an operation within the directory service, which eventually returns a status and any result of the operation.

All interactions between the user and the directory service belong to a session, which is represented by an OM object passed as the first parameter to most interface functions.

The other parameters to the functions include a context and various service-specific parameters. The context includes a number of parameters that are common to many functions, and that seldom change from operation to operation.

Each of the components of this model are described in the following sections in this chapter along with other features of the interface, such as security.

10.1 XDS Conformance to Standards

The XDS interface defines an API that application programs can use to access the functionality of the underlying directory service. The DCE XDS API conforms to the *X/Open CAE Specification, API to directory services (XDS)* (November 1991).

The DCE XDS implementation supports the following features:

- A synchronous interface. Asynchronous functionality can be achieved by using threads as described in Chapter 8.
- All synchronous interface functions are supported. The two asynchronous-specific functions are handled as follows:

— **ds_abandon()**

This call does not issue a directory service abandon operation. It returns with a **DS_C_ABANDON_FAILED (DS_E_TOO_LATE)** error. For details on abandoning operations see Section 10.10.3.

— **ds_receive_result()**

If there are any outstanding operations (when multiple threads issue XDS calls in parallel), this function returns **DS_SUCCESS** with the *completion_flag_return* parameter set to **DS_OUTSTANDING_OPERATIONS**. If no XDS calls are outstanding, this function returns **DS_SUCCESS** with the *completion_flag_return* parameter set to **DS_NO_OUTSTANDING_OPERATION**.

- Automatic connection management is not provided. The **ds_bind()** and **ds_unbind()** functions always try, respectively, to set up and release directory service connections immediately.
- The **DS_FILE_DESCRIPTOR** attribute of the **DS_C_SESSION** object is not used.
- The default values for OM attributes in the **DS_C_CONTEXT** and **DS_C_SESSION** objects are described in Chapter 11.

- Support for local strings. XDS supports the mapping from/to local string formats. The programmer can request this feature when using the following XDS/XOM functions:

- **dsX_extract_attr_values()**
- **omX_extract()**
- **omX_object_to_string()**
- **omX_string_to_object()**
- **om_get()**
- **om_read()**

The programmer controls this mapping through the *local_strings* Boolean parameter. To request conversion, set this parameter to **OM_TRUE**. The mappings currently supported are as follows:

- T.61 String to/from ISO 8859-1 (that is, LATIN-1)

For details on these mappings, refer to the *OSF DCE GDS Administration Guide and Reference*.

On input, when requesting conversion of LATIN-1 characters to T.61 format, you should only use the T.61 subset; otherwise, an error is returned.

DCE XDS supports five packages, of which one is mandatory and four are optional. Use of the optional packages is negotiated by using **ds_version()**. The packages are as follows:

- The directory service package (as defined in Chapter 11), which also includes the errors. This package is mandatory.
- The basic directory contents package (as defined in Chapter 12). This package is optional.
- The strong authentication package (as defined in Chapter 13). This package is optional.
- The message handling system directory user package (as defined in Chapter 14). This package is optional.
- The GDS package (as defined in Chapter 15). This package is optional.

None of the OM classes defined in these five packages are encodable. Thus, DCE XDS application programmers do not require the use of the XOM functions **om_encode()** and **om_decode()**, which are not supported by the DCE XOM API.

10.2 The XDS Functions

As mentioned already, the standards define abstract services that requestors use to interact with the directory. Each of these abstract services maps to a single function call, and the detailed specifications are given in the XDS reference pages. The services and the function calls to which they map are as follows:

DirectoryBind

Maps to **ds_bind()**

DirectoryUnbind

Maps to **ds_unbind()**

Read

Maps to **ds_read()**

Compare

Maps to **ds_compare()**

Abandon

Maps to **ds_abandon()**

List

Maps to **ds_list()**

Search

Maps to **ds_search()**

AddEntry

Maps to **ds_add_entry()**

RemoveEntry

Maps to **ds_remove_entry()**

ModifyEntry Maps to **ds_modify_entry()**

ModifyRDN Maps to **ds_modify_rdn()**

There is a function called **ds_receive_result()**, which has no counterpart in the abstract service. It is used with asynchronous operations. (See the **xds_intro(3xds)** reference page for information on how the asynchronous functions **ds_abandon()** and **ds_receive_result()** are handled by the DCE XDS API.)

The **ds_initialize()**, **ds_shutdown()**, and **ds_version()** functions are used to control the XDS API and do not initiate any directory operations.

The interface functions are summarized in Table 10-1.

Table 10–1. The XDS Interface Functions

Name	Description
ds_abandon()	Abandons the result of a pending asynchronous operation. This function is not supported. See xds_intro(3xds) .
ds_add_entry()	Adds a leaf entry to the DIT.
ds_bind()	Opens a session with a DUA (Directory User Agent), which in turn connects to a DSA.
ds_compare()	Compares a purported attribute value with the attribute value stored in the DIB for a particular entry.
ds_initialize()	Initializes the XDS interface.
ds_list()	Enumerates the names of the immediate subordinates of a particular directory entry.
ds_modify_entry()	Atomically performs modification to a directory entry.
ds_modify_rdn()	Changes the RDN of a leaf entry.
ds_read()	Queries information on a particular directory entry by name.
ds_receive_result()	Retrieves the result of an asynchronously executed function. See xds_intro(3xds) .
ds_remove_entry()	Removes a leaf entry from the DIT.
ds_search()	Finds entries of interest in a portion of the DIT.
ds_shutdown()	Discards a workspace.
ds_unbind()	Unbinds from a directory session.
ds_version()	Negotiates features of the interface and service.

10.3 The XDS Negotiation Sequence

The interface has an initialization and shutdown sequence that permits the negotiation of optional features. This involves the **ds_initialize()**, **ds_version()**, and **ds_shutdown()** functions.

Every application program must first call **ds_initialize()**, which returns a workspace. This workspace supports the standard directory service package (see Chapter 11).

The workspace can be extended to support the optional basic directory contents package (see Chapter 12), the strong authentication package (see Chapter 13), the GDS package (see Chapter 15), or the MHS directory user package (see Chapter 14). These packages are identified by means of OSI object identifiers, and these object identifiers are supplied to **ds_version()** to incorporate the extensions into the workspace.

After a workspace with the required features is negotiated in this way, the application can use the workspace as required. It can create and manipulate OM objects by using the OM functions, and it can start one or more directory sessions by using **ds_bind()**.

After completing its tasks, terminating all its directory sessions by using **ds_unbind()**, and releasing all its OM objects by using **om_delete()**, the application needs to ensure that resources associated with the interface are freed by calling **ds_shutdown()**.

It is possible to retain access to service-generated public objects after **ds_shutdown()** is called, or to start another cycle by calling **ds_initialize()** if so required by the application design.

10.4 The session Parameter

A session identifies the DUA and the suite of DSAs to which a particular directory operation is sent. It contains some **DirectoryBindArguments**, such as the distinguished name of the requestor. The *session* parameter is passed as the first parameter to most interface functions.

A session is described by an OM object of OM class **DS_C_SESSION**. It is created, and appropriate parameter values can be set with the OM functions. A

directory session then starts with **ds_bind()** and later terminates with **ds_unbind()**. A session with default parameters can be started by passing the constant **DS_DEFAULT_SESSION** as the **DS_C_SESSION** parameter to **ds_bind()**.

The **ds_bind()** function must be called before **DS_C_SESSION** can be used as a parameter to any other function in this interface. After **ds_unbind()** is called, **ds_bind()** must be called again if another session is to be started.

The interface supports multiple concurrent sessions so that an application implemented as a single process, such as a server in a client/server model, can interact with the directory by using several identities, and a process can interact directly and concurrently with different parts of the directory.

Details of the OM class **DS_C_SESSION** are given in Chapter 11.

10.5 The context Parameter

The context defines the characteristics of the directory interaction that are specific to a particular directory operation; nevertheless, the same characteristics are often used for many operations. Since these parameters are presumed to be relatively static for a given directory user during a particular directory interaction, these parameters are collected into an OM object of OM class **DS_C_CONTEXT**, which is supplied as the second parameter of each directory service request. This reduces the number of parameters passed to each function.

The context includes many administrative details, such as the **CommonArguments** defined in the abstract service, which affect the processing of each directory operation. These details include a number of **ServiceControls**, which allow control over some aspects of the service. The **ServiceControls** include options such as **preferChaining**, **chainingProhibited**, **localScope**, **dontUseCopy**, and **dontDereferenceAliases**, together with **priority**, **timeLimit**, **sizeLimit**, and **scopeOfReferral**. Each of these is mapped onto an OM attribute in the context (see Chapter 11).

The effect of passing the *context* parameter is as if its contents were passed as a group of additional parameters for every function call. The value of each component of the context is determined when the interface function is called, and it remains fixed throughout the operation.

All OM attributes in the class **DS_C_CONTEXT** have default values, some of which are administered locally. The constant **DS_DEFAULT_CONTEXT** can be passed as the value of the **DS_C_CONTEXT** parameter to the interface functions, and it has the same effect as a context OM object created with default values. The context must be a private object, unless it is **DS_DEFAULT_CONTEXT**.

(See Chapter 11 for detailed specifications of the OM class **DS_C_CONTEXT**.)

10.6 The XDS Function Arguments

The abstract service defines specific parameters for each operation. These are mapped onto corresponding parameters to each interface function, which are also called *input parameters*. Although each service has different parameters, some specific parameters recur in several operations and these are briefly introduced here. (For complete details of these parameters, see Chapter 11.)

All parameters that are OM objects can generally be supplied to the interface functions as public objects (that is, descriptor lists) or as private objects. Private objects must be created in the workspace that is returned by **ds_initialize()**. In some cases, constants can be supplied instead of OM objects.

Note: Wherever a function can accept an instance of a particular OM class as the value of a parameter, it also accepts an instance of any subclass of the OM class. For example, most functions have a *name* parameter, which accepts values of OM class **DS_C_NAME**. It is always acceptable to supply an instance of the subclass **DS_C_DS_DN** as the value of the parameter.

10.6.1 Attribute and Attribute Value Assertion

Each directory attribute is represented in the interface by an OM object of OM class **DS_C_ATTRIBUTE**. The type of the directory attribute is represented by an OM attribute, **DS_ATTRIBUTE_TYPE**, within the OM object. The values of the directory attribute are expressed as the values of the OM attribute **DS_ATTRIBUTE_VALUES**.

The representation of the attribute value depends on the attribute type and is determined as indicated in the following list. The list describes the way in which an application program must supply values to the interface; for example, in the *changes* parameter to **ds_modify_entry()**. The interface follows the same rules when returning attribute values to the application; for example, in the **ds_read()** result.

- The first possibility is that the attribute type and the representation of the corresponding values can be defined in a package; for example, the selected attribute types from the standards that are defined in the basic directory contents package in Chapter 12 and the strong authentication package in Chapter 13. In this case, attribute values are represented as specified. Additional directory attribute types and their OM representations are defined by the GDS package.
- If the attribute type is not known and the value is an ASN.1 simple type such as **IntegerType**, the representation is the corresponding type specified in Chapter 17.
- If the attribute type is not known and the value is an ASN.1 structured type, the value is represented in the Basic Encoding Rules (BER) with OM syntax `String(OM_S_ENCODING_STRING)`.

Note: The distinguished encoding specified in the standards (see Clause 8.7 of *The Directory: Authentication Framework*, ISO 9594-8, CCITT X.500) must be used if the request is to be signed.

Where attribute values have OM syntax `String(*)`, they can be long segmented strings, and the functions **om_read()** and **om_write()** need to be used to access them.

An attribute value assertion (AVA) is an assertion about the value of an attribute of an entry, and it can be TRUE, FALSE, or undefined. It consists of an attribute type and a single value. In general, the AVA is TRUE if one of the values of the given attribute in the entry matches the given value. An AVA is represented in the interface by an instance of OM class **DS_C_AVA**, which is a subclass of **DS_C_ATTRIBUTE** and can only have one value.

Information used by **ds_add_entry()** to construct a new directory entry is represented by an OM object of OM class **DS_C_ATTRIBUTE_LIST**, which contains a single multivalued OM attribute whose values are OM objects of OM class **DS_C_ATTRIBUTE**.

10.6.2 The selection Parameter

The *selection* parameter of the **ds_read()** and **ds_search()** operations tailors its results to obtain just part of the required entry. Information on all attributes, no attributes, or a specific group of attributes can be chosen. Attribute types are always returned, but the attribute values are not necessarily returned.

The value of the parameter is an instance of OM class **DS_C_ENTRY_INFO_SELECTION**, but one of the constants in the following list can be used in simple cases:

- To verify the existence of an entry for the purported name, use the constant **DS_SELECT_NO_ATTRIBUTES**.
- To return just the types of all attributes, use the constant **DS_SELECT_ALL_TYPES**.
- To return the types and values of all attributes, use the constant **DS_SELECT_ALL_TYPES_AND_VALUES**.

To choose a particular set of attributes, create a new instance of the OM class **DS_C_ENTRY_INFO_SELECTION** and set the appropriate OM attribute values by using the OM functions.

10.6.3 The name Parameter

Most operations take a *name* parameter to specify the target of the operation. The name is represented by an instance of one of the subclasses of the OM class **DS_C_NAME**. The DCE XDS API defines the subclass **DS_C_DS_DN** to represent distinguished names and other names.

For directory interrogations, any aliases in the name are dereferenced, unless prohibited by the **DS_DONT_DEREFERENCE_ALIASES** service control. However, for modify operations, this service control is ignored if set, and aliases are never dereferenced.

RDNs are represented by an instance of one of the subclasses of the OM class **DS_C_RELATIVE_NAME**. The DCE XDS API defines the subclass **DS_C_DS_RDN** to represent RDNs.

10.7 XDS Function Call Results

All XDS functions return a **DS_status**, which is the C function result; most return data in an *invoke_id* parameter, which identifies the particular invocation, and the interrogation operations each return data in the *result* parameter. The *invoke_id* and *result* values are returned using pointers that are supplied as parameters of the C function. These three types of function results are introduced in the following subsections.

All OM objects returned by interface functions (results and errors) are private objects in the workspace returned by **ds_initialize()**.

10.7.1 The *invoke_id* Parameter

All interface functions that invoke a directory service operation return an *invoke_id* parameter, which is an integer that identifies the particular invocation of an operation. Since asynchronous operations (within the same thread) are not supported, the *invoke_id* return value is no longer relevant for operations. DCE application programmers must still supply this parameter as described in the XDS reference pages, but they should ignore the value returned.

10.7.2 The *result* Parameter

Directory service interrogation operations return a *result* value only if they succeed. All errors from these operations, including directory access protocol (DAP) errors, are reported in **DS_status** (see Section 10.7.3), as are errors from all other operations.

The result of an interrogation is returned in a private object whose OM class is appropriate to the particular operation. The format of directory operation results is driven by the abstract service. To simplify processing, the result of a single operation is returned in a single OM object, which corresponds to the abstract result defined in the standards. The components of the result of an operation are represented by OM attributes in the operation's result object. All information contained in the abstract service result is made available to the application program. The result is inspected using the functions provided in the object management API, **om_get()**.

Only the interrogation operations produce results, and each type of interrogation has a specific OM class of OM object for its result. These OM classes are as follows (see Chapter 11 for their definitions):

- **DS_C_COMPARE_RESULT**
- **DS_C_LIST_RESULT**
- **DS_C_READ_RESULT**
- **DS_C_SEARCH_RESULT**

The results of the different operations share several common components, including the **CommonResults** defined in the standards (see *The Directory: Abstract Service Definition*, ISO 9594-3, CCITT X.511) by inheriting OM attributes from the superclass **DS_C_COMMON_RESULTS**. An additional common component is the full DN of the target object, after all aliases are dereferenced.

The actual OM class of the result can always be a subclass of that named in order to allow flexibility for extensions. Thus, **om_instance()** always needs to be used when testing the OM class.

Any attribute values in the result are represented as discussed in Section 10.6.1.

10.7.3 The DS_status Return Value

Every interface function returns a **DS_status** value, which is either the constant **DS_SUCCESS** or an error. Errors are represented by private objects whose OM class is a subclass of **DS_C_ERROR**. Details of all errors are given in Chapter 11.

Other results of functions are not valid unless the status result has the value **DS_SUCCESS**.

10.8 Synchronous Operations

Since asynchronous use of the interface (within the same thread) is not supported, the value of the **DS_ASYNCHRONOUS** OM attribute in **DS_C_CONTEXT** is always **OM_FALSE**, causing all operations within the same thread to be synchronous.

In synchronous mode, all functions wait until the operation is complete before returning. The thread of control is blocked within the interface after calling a function, and it can use the result immediately after the function returns.

Implementations define a limit on the number of asynchronous operations that can be outstanding at any one time on any one session. The limit is given by the implementation-defined constant **DS_MAX_OUTSTANDING_OPERATIONS**. It always has the value 0 (zero) because asynchronous operations within the same thread are not supported.

All errors occurring during a synchronous request are reported when the function returns. (See Chapter 11 for complete details of error handling.)

The **DS_FILE_DESCRIPTOR** OM attribute of **DS_C_SESSION** is not used by the DCE XDS API and is always set to **DS_NO_VALID_FILE_DESCRIPTOR**.

10.9 Security and XDS

The X/Open XDS specifications do not define a security interface because this can put constraints on security features of existing directory implementations.

DCE GDS provides an extension to the XDS API for security support. This is achieved at the XDS API level through a new **DSX_C_GDS_SESSION** session object that contains information on the security mechanism that should be used. Simple authentication through the use of name and password, and external authentication based on DCE security, are supported. (See Chapter 15 for additional information.)

10.10 Other Features of the XDS Interface

The following subsections describe these features of the interface:

- Automatic Connection Management
- Automatic Continuation and Referral Handling

10.10.1 Automatic Connection Management

An implementation can provide automatic management of the association or connection between the user and the directory service, making and releasing connections at its discretion.

The DCE XDS implementation does not support automatic connection management. A DSA connection is established when **ds_bind()** is called and released when **ds_unbind()** is called.

10.10.2 Automatic Continuation and Referral Handling

The interface provides automatic handling of continuation references and referrals in order to reduce the burden on application programs. These facilities can be inhibited to meet special needs.

A *continuation reference* describes how the performance of all or part of an operation can be continued at a different DSA or DSAs. A single continuation reference returned as the entire response to an operation is called a *referral* and is classified as an error. One or more continuation references can also be returned as part of **DS_PARTIAL_OUTCOME_QUAL** returned from a **ds_list()** or **ds_search()** operation.

A DSA returns a referral if it has administrative, operational, or technical reasons for preferring not to chain. It can return a referral if **DS_CHAINING_PROHIB** is set in the **DS_C_CONTEXT**, or it can report a service error (**DS_E_CHAINING_REQUIRED**) instead.

By default, the implementation uses any continuation references it receives to try to contact the other DSA or DSAs, enabling it to make further progress in the operation, whenever practical. It only returns the result, or an error, to the application after it has made this attempt. Note that continuation references can still be returned to the application; for example, if the relevant DSA cannot be contacted.

The default behavior is the simplest for most applications but, if necessary, the application can cause all continuation references to be returned to it. It does this by setting the value of the OM attribute **DS_AUTOMATIC_CONTINUATION** in the **DS_C_CONTEXT** to **OM_FALSE**.

10.10.3 Abandoning Operations

The XDS user can abandon a directory operation when operating in multithreaded mode. An operation is abandoned by calling **pthread_cancel()** to cancel the thread that issued the directory operation. General cancelability must be enabled; otherwise, the cancelability will be ignored.

XDS will react as follows, depending on when the cancel is delivered:

- Before interaction with the DSA
 - Nothing is sent to the DSA.
 - The exception **pthread_cancel_e** is reraised.
- While waiting for a response from the DSA
 - An **ABANDON** message is sent to the DSA.
 - The exception **pthread_cancel_e** is reraised.
- After the result has arrived, but before a point has been reached when it is committed to be passed back to the user
 - The result is thrown away.
 - The exception **pthread_cancel_e** is reraised.
- After the point where result return is committed place
 - The cancel is ignored.
 - The result is returned normally.

It is the responsibility of the user to handle the cancel exception in the last case and, if necessary, to discard the result.

Chapter 11

XDS Class Definitions

When referring to classes and attributes in the directory service, the chapters in Parts 3 and 4 make a clear distinction between OM classes and directory classes, and between OM attributes and directory attributes. In both cases, the former is a construct of the closely associated Object Management interface, while the latter is a construct of the directory service to which XDS provides access. The terms *object class* and *attribute* indicate the directory constructs, while the phrases *OM class* and *OM attribute* indicate the Object Management constructs.

11.1 Introduction to OM Classes

This chapter defines, in alphabetical order, the OM classes that constitute the directory service package. This package incorporates the OM classes for the errors that may be returned at the XDS interface. The object identifier associated with this package is

```
{iso(1) identified-organization(3) icd-ecma(0012) member-company(2)
dec(1011) xopen(28) dsp(0)}
```

It takes the following encoding:

`\x2B\xC\x2\x87\x73\x1C\x0`

This object identifier is represented by the constant **DS_SERVICE_PKG**.

The object management notation is briefly described in the following text. See Chapters 17 through 19 for more information on object management.

Each OM class is described in a separate section, which identifies the OM attributes specific to that OM class. The OM classes and OM attributes for each OM class are listed in alphabetical order. The OM attributes that can be found in an instance of an OM class are those OM attributes specific to that OM class, as well as those inherited from each of its superclasses (see Chapter 5). The OM class-specific OM attributes are defined in a table. The table indicates the name of each OM attribute, the syntax of each of its values, any restrictions on the length (in bits, octets (bytes), or characters) of each value, any restrictions upon the number of values, and the value, if any, **om_create()** supplies.

The constants that represent the OM classes and OM attributes in the C binding are defined in the **xds.h(4xds)** header file.

11.2 XDS Errors

Errors are reported to the application program by means of **DS_status**, which is a result of every function. (The **DS_status** is *the* function result in the C language binding for most functions.) A function that completes successfully returns the value **DS_SUCCESS**, whereas one that is not successful returns an error. The error is a private object containing details of the problem that occurred. The error constant **DS_NO_WORKSPACE** can be returned by all directory service functions, except **ds_initialize()**. **DS_NO_WORKSPACE** is returned if **ds_initialize()** is not invoked before calling any other directory service function.

Errors are classified into ten OM classes. The standards (see *The Directory: Abstract Service Definition*, ISO 9594-3, CCITT X.511) classify errors into eight different groups, as follows:

- Abandoned

- Abandon Failed
- Attribute Error
- Name Error
- Referral
- Security Error
- Service Error
- Update Error

The directory service interface never returns an Abandoned error. The interface also defines three more kinds of errors, as follows:

- **DS_C_LIBRARY_ERROR**
- **DS_C_COMMUNICATIONS_ERROR**
- **DS_C_SYSTEM_ERROR**

Each of these kinds of errors is represented by an OM class. These OM classes are detailed in subsequent sections of this chapter. All of them inherit the OM attribute **DS_PROBLEM** from their superclass *DS_C_ERROR*, which is described in this chapter. The values that **DS_PROBLEM** can take are listed in the relevant subsections of this chapter. For a description of these errors, refer to the *OSF DCE Problem Determination Guide*. The error OM classes defined in this chapter are part of the directory service package.

The **ds_bind()** operation returns a Security Error or a Service Error. All other operations can also return the same errors as **ds_bind()**. Such errors can arise in the course of following an automatic referral list.

DS_C_REFERRAL is not a real error, and it is not a subclass of *DS_C_ERROR*, although it is reported in the same way as a **DS_status** result. A **DS_C_ATTRIBUTE_ERROR**, also not a subclass of *DS_C_ERROR*, is special because it can report several problems at once. Each one is reported in **DS_C_ATTRIBUTE_PROBLEM**, which is a subclass of *DS_C_ERROR*.

11.3 OM Class Hierarchy

This section shows the hierarchical organization of the OM classes defined in this chapter and, as a result, shows which OM classes inherit additional OM attributes from their superclasses. In the following list, subclassification is indicated by indentation, and the names of abstract classes are in italics. Thus, for example, the concrete class **DS_C_PRESENTATION_ADDRESS** is an immediate subclass of the abstract class *DS_C_ADDRESS*, which in turn is an immediate subclass of the abstract class *OM_C_OBJECT*. (*OM_C_OBJECT* is defined in Chapter 19 of this guide.)

OM_C_OBJECT

- **DS_C_ACCESS_POINT**
- *DS_C_ADDRESS*
 - **DS_C_PRESENTATION_ADDRESS**
- **DS_C_ATTRIBUTE**
 - **DS_C_AVA**
 - **DS_C_ENTRY_MOD**
 - **DS_C_FILTER_ITEM**
- **DS_C_ATTRIBUTE_ERROR**
- **DS_C_ATTRIBUTE_LIST**
 - **DS_C_ENTRY_INFO**
- *DS_C_COMMON_RESULTS*
 - **DS_C_COMPARE_RESULT**
 - **DS_C_LIST_INFO**
 - **DS_C_READ_RESULT**
 - **DS_C_SEARCH_INFO**
- **DS_C_CONTEXT**
- **DS_C_CONTINUATION_REF**
 - **DS_C_REFERRAL**
- **DS_C_ENTRY_INFO_SELECTION**

- **DS_C_ENTRY_MOD_LIST**
- *DS_C_ERROR*
 - **DS_C_ABANDON_FAILED**
 - **DS_C_ATTRIBUTE_PROBLEM**
 - **DS_C_COMMUNICATIONS_ERROR**
 - **DS_C_LIBRARY_ERROR**
 - **DS_C_NAME_ERROR**
 - **DS_C_SECURITY_ERROR**
 - **DS_C_SERVICE_ERROR**
 - **DS_C_SYSTEM_ERROR**
 - **DS_C_UPDATE_ERROR**
- **DS_C_EXT**
- **DS_C_FILTER**
- **DS_C_LIST_INFO_ITEM**
- **DS_C_LIST_RESULT**
- *DS_C_NAME*
 - **DS_C_DS_DN**
- **DS_C_OPERATION_PROGRESS**
- **DS_C_PARTIAL_OUTCOME_QUAL**
- *DS_C_RELATIVE_NAME*
 - **DS_C_DS_RDN**
- **DS_C_SEARCH_RESULT**
- **DS_C_SESSION**

None of the classes in the preceding list are encodable using **om_encode()** and **om_decode()**. The application is not permitted to create or modify instances of some OM classes because these OM classes are only returned by the interface and never supplied to it. These OM classes are as follows:

- **DS_C_ACCESS_POINT**
- **DS_C_ATTRIBUTE_ERROR**
- **DS_C_COMPARE_RESULT**
- **DS_C_CONTINUATION_REF**
- All subclasses of *DS_C_ERROR*
- **DS_C_LIST_INFO**
- **DS_C_LIST_INFO_ITEM**
- **DS_C_LIST_RESULT**
- **DS_C_OPERATION_PROGRESS**
- **DS_C_PARTIAL_OUTCOME_QUAL**
- **DS_C_READ_RESULT**
- **DS_C_REFERRAL**
- **DS_C_SEARCH_INFO**
- **DS_C_SEARCH_RESULT**

11.4 DS_C_ABANDON_FAILED

An instance of OM class **DS_C_ABANDON_FAILED** reports a problem encountered during an attempt to abandon an operation.

An application is not permitted to create or modify instances of this OM class. An instance of this OM class has the OM attributes of its superclasses, *OM_C_OBJECT* and *DS_C_ERROR*, and no additional OM attributes.

The OM attribute **DS_PROBLEM**, which is inherited from the superclass *DS_C_ERROR*, identifies the problem. Its value is one of the following:

- **DS_E_CANNOT_ABANDON**
- **DS_E_NO_SUCH_OPERATION**
- **DS_E_TOO_LATE**

A **ds_abandon()** XDS call always returns a **DS_E_TOO_LATE** error for the **DS_C_ABANDON_FAILED** OM class. Refer to Chapter 10 for information on abandoning directory operations.

11.5 DS_C_ACCESS_POINT

An instance of OM class **DS_C_ACCESS_POINT** identifies a particular point at which a DSA can be accessed.

An application is not permitted to create or modify instances of this OM class. An instance of this OM class has the OM attributes of its superclass, *OM_C_OBJECT*, in addition to the OM attributes listed in Table 11-1.

Table 11-1. OM Attributes of DS_C_ACCESS_POINT

OM Attribute	Value Syntax	Value Length	Value Number	Value Initially
DS_ADDRESS	Object(<i>DS_C_ADDRESS</i>)	—	1	—
DS_AE_TITLE	Object(<i>DS_C_NAME</i>)	—	1	—

- **DS_ADDRESS**

This attribute indicates the address of the DSA to be used when communicating with it.

- **DS_AE_TITLE**

This attribute indicates the name of the DSA.

11.6 DS_C_ADDRESS

The OM class *DS_C_ADDRESS* represents the address of a particular entity or service, such as a DSA.

It is an abstract class that has the OM attributes of its superclass, *OM_C_OBJECT*, and no other OM attributes.

An address is an unambiguous name, label, or number that identifies the location of the entity or service. All addresses are represented as instances of some subclass of this OM class.

The only subclass defined by the DCE XDS API is **DS_C_PRESENTATION_ADDRESS**, which is the presentation address of an OSI application entity used for OSI communications with this subclass.

11.7 DS_C_ATTRIBUTE

An instance of OM class **DS_C_ATTRIBUTE** is an attribute of an object, and is thus a component of its directory entry.

An instance of this OM class has the OM attributes of its superclass, *OM_C_OBJECT*, in addition to the OM attributes listed in Table 11-2.

Table 11-2. OM Attributes of DS_C_ATTRIBUTE

OM Attribute	Value Syntax	Value Length	Value Number	Value Initially
DS_ATTRIBUTE_TYPE	String(OM_S_OBJECT_IDENTIFIER_STRING)	—	1	—
DS_ATTRIBUTE_VALUES	Any	—	0 or more	—

- **DS_ATTRIBUTE_TYPE**

The attribute type that indicates the class of information given by this attribute.

- **DS_ATTRIBUTE_VALUES**

The attribute values. The OM value syntax and the number of values allowed for this OM attribute are determined by the value of the **DS_ATTRIBUTE_TYPE** OM attribute in accordance with the rules given in Chapter 10.

If the values of this OM attribute have the syntax `String(*)`, the strings can be long and segmented. For this reason, **om_read()** and **om_write()** need to be used to access all `String(*)` values.

Note: A directory attribute must always have at least one value, although it is acceptable for instances of this OM class not to have any values.

11.8 DS_C_ATTRIBUTE_ERROR

An instance of OM class **DS_C_ATTRIBUTE_ERROR** reports an attribute-related directory service error.

An application is not permitted to create or modify instances of this OM class. An instance of this OM class has the OM attributes of its superclass, *OM_C_OBJECT*, in addition to the OM attributes listed in Table 11-3.

Table 11-3. OM Attributes of DS_C_ATTRIBUTE_ERROR

OM Attribute	Value Syntax	Value Length	Value Number	Value Initially
DS_OBJECT_NAME	Object(<i>DS_C_NAME</i>)	—	1	—
DS_PROBLEMS	Object(DS_C_ATTRIBUTE_PROBLEM)	—	1 or more	—

- **DS_OBJECT_NAME**

This attribute contains the name of the directory entry to which the operation is applied when the failure occurs.

- **DS_PROBLEMS**

This attribute documents the attribute-related problems encountered. Uniquely, a **DS_C_ATTRIBUTE_ERROR** can report several problems at once. All problems are related to the preceding object.

11.9 DS_C_ATTRIBUTE_LIST

An instance of OM class **DS_C_ATTRIBUTE_LIST** is a list of directory attributes.

An instance of this OM class has the OM attributes of its superclass, *OM_C_OBJECT*, in addition to the OM attribute listed in Table 11-4.

Table 11-4. OM Attribute of DS_C_ATTRIBUTE_LIST

OM Attribute	Value Syntax	Value Length	Value Number	Value Initially
DS_ATTRIBUTES	Object(DS_C_ATTRIBUTE)	—	0 or more	—

- **DS_ATTRIBUTES**

This attribute indicates the attributes that constitute a new object’s directory entry, or those selected from an existing entry.

11.10 DS_C_ATTRIBUTE_PROBLEM

An instance of OM class **DS_C_ATTRIBUTE_PROBLEM** documents one attribute-related problem encountered while performing an operation as requested on a particular occasion.

An application is not permitted to create or modify instances of this OM class. An instance of this OM class has the OM attributes of its superclasses, *OM_C_OBJECT* and *DS_C_ERROR*, in addition to the OM attributes listed in Table 11-5.

Table 11–5. OM Attributes of DS_C_ATTRIBUTE_PROBLEM

OM Attribute	Value Syntax	Value Length	Value Number	Value Initially
DS_ATTRIBUTE_TYPE	String(OM_S_OBJECT_IDENTIFIER_STRING)	—	1	—
DS_ATTRIBUTE_VALUE	Any	—	0 or 1	—

- **DS_ATTRIBUTE_TYPE**

This attribute identifies the type of attribute with which the problem is associated.

- **DS_ATTRIBUTE_VALUE**

This attribute specifies the attribute value with which the problem is associated. Its syntax is determined by the value of **DS_ATTRIBUTE_TYPE**. This OM attribute is present if it is necessary to avoid ambiguity.

The OM attribute **DS_PROBLEM**, which is inherited from the superclass *DS_C_ERROR*, identifies the problem. Its value is one of the following:

- **DS_E_ATTRIBUTE_OR_VALUE_EXISTS**
- **DS_E_CONSTRAINT_VIOLATION**
- **DS_E_INAPPROP_MATCHING**
- **DS_E_INVALID_ATTRIBUTE_SYNTAX**
- **DS_E_NO_SUCH_ATTRIBUTE_OR_VALUE**
- **DS_E_UNDEFINED_ATTRIBUTE_TYPE**

11.11 DS_C_AVA

An instance of OM class **DS_C_AVA** (attribute value assertion) is a proposition concerning the values of a directory entry.

An instance of this OM class has the OM attributes of its superclasses, *OM_C_OBJECT* and **DS_C ATTRIBUTE**, and no other OM attributes. An additional restriction on this OM class is that there must be exactly one value of the OM attribute **DS_ATTRIBUTE_VALUES**. The **DS_ATTRIBUTE_TYPE** remains single valued. The OM value syntax of **DS_ATTRIBUTE_VALUES** must conform to the rules outlined in Chapter 10.

11.12 DS_C_COMMON_RESULTS

The OM class *DS_C_COMMON_RESULTS* comprises results that are returned by, and are common to, the directory interrogation operations.

It is an abstract OM class, which has the OM attributes of its superclass, *OM_C_OBJECT*, in addition to the OM attributes listed in Table 11-6.

Table 11-6. OM Attributes of DS_C_COMMON_RESULTS

OM Attribute	Value Syntax	Value Length	Value Number	Value Initially
DS_ALIAS_DEREFERENCED	OM_S_BOOLEAN	—	1	—
DS_PERFORMER	Object(<i>DS_C_NAME</i>)	—	0 or 1	—

- **DS_ALIAS_DEREFERENCED**

This attribute indicates whether the name of the target object that is passed as a function argument includes an alias that is dereferenced to determine the DN.

- **DS_PERFORMER**

When present, this attribute gives the DN of the performer of a particular operation. It can be present when the result is signed, and it holds the name of the DSA that signed the result. The DCE directory service does not support the optional feature of signed results; therefore, this OM attribute is never present.

11.13 DS_C_COMMUNICATIONS_ERROR

An instance of OM class **DS_C_COMMUNICATIONS_ERROR** reports an error occurring in the other OSI services supporting the directory service.

An application is not permitted to create or modify instances of this OM class. An instance of this OM class has the OM attributes of its superclasses, *OM_C_OBJECT* and *DS_C_ERROR*, and no additional OM attributes.

Communications errors include those arising in remote operation, association control, presentation, session, and transport.

The OM attribute **DS_PROBLEM**, which is inherited from the superclass *DS_C_ERROR*, identifies the problem. Its value is **DS_E_COMMUNICATIONS_PROBLEM**.

11.14 DS_C_COMPARE_RESULT

An instance of OM class **DS_C_COMPARE_RESULT** comprises the results of a successful call to **ds_compare()**.

An application is not permitted to create or modify instances of this OM class. An instance of this OM class has the OM attributes of its superclasses, *OM_C_OBJECT* and *DS_C_COMMON_RESULTS*, in addition to the OM attributes listed in Table 11-7.

Table 11-7. OM Attributes of DS_C_COMPARE_RESULT

OM Attribute	Value Syntax	Value Length	Value Number	Value Initially
DS_FROM_ENTRY	OM_S_BOOLEAN	—	1	—
DS_MATCHED	OM_S_BOOLEAN	—	1	—
DS_OBJECT_NAME	Object(<i>DS_C_NAME</i>)	—	0 or 1	—

- **DS_FROM_ENTRY**

This attribute indicates whether the assertion is tested against the specified object's entry, rather than a copy of the entry.

- **DS_MATCHED**

This attribute indicates whether the assertion specified as an argument returns the value **OM_TRUE**. It takes the value **OM_TRUE** if the values are compared and matched; otherwise, it takes the value **OM_FALSE**.

- **DS_OBJECT_NAME**

This attribute contains the DN of the target object of the operation. It is present if the OM attribute **DS_ALIAS_DEREFERENCED**, inherited from the superclass *DS_C_COMMON_RESULTS*, is **OM_TRUE**.

11.15 DS_C_CONTEXT

An instance of OM class **DS_C_CONTEXT** comprises per-operation arguments that are accepted by most of the interface functions.

An instance of this OM class has the OM attributes of its superclass, *OM_C_OBJECT*, in addition to the OM attributes listed in Table 11-8.

Table 11–8. OM Attributes of DS_C_CONTEXT

OM Attribute	Value Syntax	Value Length	Value Number	Value Initially
Common Arguments				
DS_EXT	Object(DS_C_EXT)	—	0 or more	—
DS_OPERATION_PROGRESS	Object(DS_C_OPERATION_PROGRESS)	—	1	DS_OPERATION_NOT_STARTED
DS_ALIASED_RDNS	OM_S_INTEGER	—	0 or 1	0
Service Controls				

OM Attribute	Value Syntax	Value Length	Value Number	Value Initially
DS_CHAINING_PROHIB	OM_S_ BOOLEAN	—	1	OM_TRUE
DS_DONT_DEREFERENCE_ALIASES	OM_S_ BOOLEAN	—	1	OM_FALSE
DS_DONT_USE_COPY	OM_S_ BOOLEAN	—	1	OM_TRUE
DS_LOCAL_SCOPE	OM_S_ BOOLEAN	—	1	OM_FALSE
DS_PREFER_CHAINING	OM_S_ BOOLEAN	—	1	OM_FALSE
DS_PRIORITY	Enum(DS_Priority)	—	1	DS_MEDIUM
DS_SCOPE_OF_REFERRAL	Enum(DS_Scope_of_Referral)	—	0 or 1	—
DS_SIZE_LIMIT	OM_S_ INTEGER	—	0 or 1	—
DS_TIME_LIMIT	OM_S_ INTEGER	—	0 or 1	—
Local Controls				
DS_ASYNCHRONOUS	OM_S_ BOOLEAN	—	1	OM_FALSE
DS_AUTOMATIC_CONTINUATION	OM_S_ BOOLEAN	—	1	OM_TRUE

The context gathers several arguments passed to interface functions, which are presumed to be relatively static for a given directory user during a particular directory interaction. The context is passed as an argument to each function that interrogates or updates the directory. Although it is generally assumed that the context is changed infrequently, the value of each argument can be changed between every operation if required. The **DS_ASYNCHRONOUS** argument must not be changed. Each argument is represented by one of the OM attributes of the **DS_C_CONTEXT** OM class.

The context contains the common arguments defined in the standards (see *The Directory: Abstract Service Definition*, ISO 9594-3, CCITT X.511), except that all security information is omitted for reasons discussed in Chapter 10. These are made up of a number of service controls explained in the following text, possible extensions in the **DS_EXT** OM attribute, and operation progress and alias dereferencing information in the **DS_OPERATION_PROGRESS** OM attribute. It also contains a number of arguments that provide local control over the interface.

The OM attributes of the **DS_C_CONTEXT** OM class are as follows:

- Common Arguments

- **DS_EXT**

This attribute represents any future standardized extensions that need to be applied to the directory service operation. The DCE XDS implementation does not evaluate this optional OM attribute.

- **DS_OPERATION_PROGRESS**

This attribute represents the state that the directory service assumes at the start of the operation. This OM attribute normally takes its default value, which is the value **DS_OPERATION_NOT_STARTED** described in the **DS_C_OPERATION_PROGRESS** OM class definition.

- **DS_ALIASED_RDNS**

This attribute indicates to the directory service that the object component of the *operation* parameter is created by dereferencing of an alias on an earlier operation attempt. This value is set in the referral response of the previous operation.

- Service Controls

- **DS_CHAINING_PROHIB**

This attribute indicates that chaining and other methods of distributing the request around the directory service are prohibited.

- **DS_DONT_DEREFERENCE_ALIASES**

This attribute indicates that any alias used to identify the target entry of an operation is not dereferenced. This allows interrogation of alias entries. (Aliases are never dereferenced during updates.)

— **DS_DONT_USE_COPY**

This attribute indicates that the request can only be satisfied by accessing directory entries, and not by using copies of entries. This includes both copies maintained in other DSAs by bilateral agreement, and, copies cached locally.

— **DS_LOCAL_SCOPE**

This attribute indicates that the directory request will be satisfied locally. The meaning of this option is configured by an administrator. This option typically restricts the request to a single DSA or DMD.

— **DS_PREFER_CHAINING**

This attribute indicates that chaining is preferred to referrals when necessary. The directory service is not obliged to follow this preference and can return a referral even if it is set.

— **DS_PRIORITY**

This attribute indicates the priority, relative to other directory requests, according to which the directory service attempts to satisfy the request. This is not a guaranteed service since there is no queuing throughout the directory. Its value must be one of the following:

DS_LOW

DS_MEDIUM

DS_HIGH

— **DS_SCOPE_OF_REFERRAL**

This attribute indicates the part of the directory to which referrals are limited. This includes referral errors and partial outcome qualifiers. Its value must be one of the following:

DS_COUNTRY, meaning DSAs within the country in which the request originates.

DS_DMD, meaning DSAs within the DMD in which the request originates.

DS_SCOPE_OF_REFERRAL is an optional attribute. The lack of this attribute in a **DS_C_CONTEXT** object indicates that the scope is not limited.

— **DS_SIZE_LIMIT**

If present, this attribute indicates the maximum number of objects about which **ds_list()** or **ds_search()** needs to return information. If this limit is exceeded, information is returned about exactly this number of objects. The objects that are chosen are not specified because this can depend on the timing of interactions between DSAs, among other reasons.

— **DS_TIME_LIMIT**

If present, this attribute indicates the maximum elapsed time, in seconds, within which the service needs to be provided (not the processing time devoted to the request). If this limit is reached, a service error (**DS_E_TIME_LIMIT_EXCEEDED**) is returned, except for the **ds_list()** or **ds_search()** operations, which return an arbitrary selection of the accumulated results.

- Local Controls

— **DS_ASYNCHRONOUS** (Optional Functionality)

The interface currently operates synchronously (within the same thread) only, as detailed in Chapter 10. There is only one possible value, as follows:

OM_FALSE, meaning that the operation is performed sequentially (synchronously) with the application being blocked until a result or error is returned.

— **DS_AUTOMATIC_CONTINUATION**

This attribute indicates the requestor's requirement for continuation reference handling, including referrals and those in partial outcome qualifiers. The value is one of the following:

OM_FALSE, meaning that the interface returns all continuation references to the application program.

OM_TRUE, meaning that continuation references are automatically processed, and the subsequent results are returned to the application instead of the continuation references, whenever practical. This is a much simpler option than **OM_FALSE** unless the application has special requirements.

Note: Continuation references can still be returned to the application if, for example, the relevant DSA cannot be contacted.

Applications can assume that an object of OM class **DS_C_CONTEXT**, created with default values of all its OM attributes, works with all the interface functions. The **DS_DEFAULT_CONTEXT** constant can be used as an argument to interface functions instead of creating an OM object with default values.

11.16 DS_C_CONTINUATION_REF

An instance of OM class **DS_C_CONTINUATION_REF** comprises the information that enables a partially completed directory request to be continued; for example, following a referral.

An application is not permitted to create or modify instances of this OM class. An instance of this OM class has the OM attributes of its superclass, *OM_C_OBJECT*, in addition to the OM attributes listed in Table 11-9.

Table 11-9. OM Attributes of DS_C_CONTINUATION_REF

OM Attribute	Value Syntax	Value Length	Value Number	Value Initially
DS_ACCESS_POINTS	Object(DS_C_ACCESS_POINT)	—	1 or more	—
DS_ALIASED_RDNS	OM_S_INTEGER	—	1	—
DS_OPERATION_PROGRESS	Object(DS_C_OPERATION_PROGRESS)	—	1	—

OM Attribute	Value Syntax	Value Length	Value Number	Value Initially
DS_RDNS_RESOLVED	OM_S_INTEGER	—	0 or 1	—
DS_TARGET_OBJECT	Object(<i>DS_C_NAME</i>)	—	1	—

- **DS_ACCESS_POINTS**

This attribute indicates the names and presentation addresses of the DSAs from where the directory request is continued.

- **DS_ALIASED_RDNS**

This attribute indicates how many (if any) of the RDNs in the target name are produced by dereferencing an alias. Its value is 0 (zero) if no aliases are dereferenced. This value needs to be used in the **DS_C_CONTEXT** of any continued operation.

- **DS_OPERATION_PROGRESS**

This attribute indicates the state at which the directory request must be continued. This value needs to be used in the **DS_C_CONTEXT** of any continued operation.

- **DS_RDNS_RESOLVED**

This attribute indicates the number of RDNs in the supplied object name that are resolved (using internal references), and not just assumed to be correct (using cross-references).

- **DS_TARGET_OBJECT**

This attribute indicates the name of the object upon which the continuation must focus.

11.17 DS_C_DS_DN

An instance of OM class **DS_C_DS_DN** represents a name of a directory object.

An instance of this OM class has the OM attributes of its superclasses, *OM_C_OBJECT* and *DS_C_NAME*, in addition to the OM attribute listed in Table 11-10.

Table 11-10. OM Attribute of DS_C_DS_DN

OM Attribute	Value Syntax	Value Length	Value Number	Value Initially
DS_RDNS	Object(DS_C_DS_RDN)	—	0 or more	—

- **DS_RDNS**

This attribute indicates the sequence of RDNs that define the path through the DIT from its root to the object that the **DS_C_DS_DN** indicates. The **DS_C_DS_DN** of the root of the directory is the null name (no **DS_RDNS** values). The order of the values is significant; the first value is closest to the root, and the last value is the RDN of the object.

11.18 DS_C_DS_RDN

An instance of OM class **DS_C_DS_RDN** is a relative distinguished name. An RDN uniquely identifies an immediate subordinate of an object whose entry is displayed in the DIT.

An instance of this OM class has the OM attributes of its superclasses, *OM_C_OBJECT* and *DS_C_RELATIVE_NAME*, in addition to the OM attribute listed in Table 11-11.

Table 11-11. OM Attribute of DS_C_DS_RDN

OM Attribute	Value Syntax	Value Length	Value Number	Value Initially
DS_AVAS	Object(DS_C_AVA)	—	1 or more	—

- **DS_AVAS**

This attribute indicates the **DS_AVAS** that are marked by the DIB as components of the object's RDN. The assertion is TRUE of the object but not of any of its siblings, and the attribute type and value are displayed in the object's directory entry. The order of the **DS_AVAS** is not significant.

11.19 DS_C_ENTRY_INFO

An instance of OM class **DS_C_ENTRY_INFO** contains selected information from a single directory entry.

An instance of this OM class has the OM attributes of its superclasses, *OM_C_OBJECT* and **DS_C_ATTRIBUTE_LIST**, in addition to the OM attributes listed in Table 11-12.

Table 11-12. OM Attributes of DS_C_ENTRY_INFO

OM Attribute	Value Syntax	Value Length	Value Number	Value Initially
DS_FROM_ENTRY	OM_S_BOOLEAN	—	1	—
DS_OBJECT_NAME	Object(<i>DS_C_NAME</i>)	—	1	—

The OM attribute **DS_ATTRIBUTES** is inherited from the superclass **DS_C_ATTRIBUTE_LIST**. It contains the information extracted from the directory entry of the target object. The type of each attribute requested and located is indicated in the list as are its values, if types and values are requested.

The OM class-specific OM attributes are as follows:

- **DS_FROM_ENTRY**

This attribute indicates whether the information is extracted from the specified object's entry, rather than from a copy of the entry.

- **DS_OBJECT_NAME**

This attribute contains the object's DN.

11.20 DS_C_ENTRY_INFO_SELECTION

An instance of OM class **DS_C_ENTRY_INFO_SELECTION** identifies the information to be extracted from a directory entry.

An instance of this OM class has the OM attributes of its superclass, *OM_C_OBJECT*, in addition to the OM attributes listed in Table 11-13.

Table 11-13. OM Attributes of DS_C_ENTRY_INFO_SELECTION

OM Attribute	Value Syntax	Value Length	Value Number	Value Initially
DS_ALL_ATTRIBUTES	OM_S_BOOLEAN	—	1	OM_TRUE
DS_ATTRIBUTES_SELECTED	String(OM_S_OBJECT_IDENTIFIER_STRING)	—	0 or more	—
DS_INFO_TYPE	Enum(DS_Information_Type)	—	1	DS_TYPES_AND_VALUES

- **DS_ALL_ATTRIBUTES**

This attribute indicates which attributes are relevant. It can take one of the following values:

- **OM_FALSE**, meaning that information is only requested on those attributes that are listed in the OM attribute **DS_ATTRIBUTES_SELECTED**.
- **OM_TRUE**, meaning that information is requested on all attributes in the directory entry. Any values of the OM attribute **DS_ATTRIBUTES_SELECTED** are ignored in this case.

- **DS_ATTRIBUTES_SELECTED**

This attribute lists the types of attributes in the entry from which information will be extracted. The value of this OM attribute is used only if the value of **DS_ALL_ATTRIBUTES** is **OM_FALSE**. If an empty list is supplied, no attribute data is returned that could be used to verify the existence of an entry for a DN.

- **DS_INFO_TYPE**

This attribute identifies the information that will be extracted from each attribute identified. It must take one of the following values:

- **DS_TYPES_ONLY**, meaning that only the attribute types of the selected attributes in the entry are returned.
- **DS_TYPES_AND_VALUES**, meaning that both the attribute types and the attribute values of the selected attributes in the entry are returned.

11.21 DS_C_ENTRY_MOD

An instance of OM class **DS_C_ENTRY_MOD** describes a single modification to a specified attribute of a directory entry.

An instance of this OM class has the OM attributes of its superclasses, *OM_C_OBJECT* and **DS_C_ATTRIBUTE**, in addition to the OM attribute listed in Table 11-14.

Table 11-14. OM Attribute of DS_C_ENTRY_MOD

OM Attribute	Value Syntax	Value Length	Value Number	Value Initially
DS_MOD_TYPE	Enum(DS_Modification_Type)	—	1	DS_ADD_ATTRIBUTE

The attribute type to be modified, and the associated values, are specified in the OM attributes **DS_ATTRIBUTE_TYPE** and **DS_ATTRIBUTE_VALUES** that are inherited from the **DS_C_ATTRIBUTE** superclass.

- **DS_MOD_TYPE**

This attribute identifies the type of modification. It must have one of the following values:

- **DS_ADD_ATTRIBUTE**, meaning that the specified attribute is absent and will be added with the specified values.

- **DS_ADD_VALUES**, meaning that the specified attribute is present and that one or more specified values will be added to it.
- **DS_REMOVE_ATTRIBUTE**, meaning that the specified attribute is present and will be removed. Any values present in the OM attribute **DS_ATTRIBUTE_VALUES** are ignored.
- **DS_REMOVE_VALUES**, meaning that the specified attribute is present and that one or more specified values will be removed from it.

11.22 DS_C_ENTRY_MOD_LIST

An instance of OM class **DS_C_ENTRY_MOD_LIST** comprises a sequence of changes to be made to a directory entry.

An instance of this OM class has the OM attributes of its superclass, *OM_C_OBJECT*, in addition to the OM attribute listed in Table 11-15.

Table 11-15. OM Attribute of DS_C_ENTRY_MOD_LIST

OM Attribute	Value Syntax	Value Length	Value Number	Value Initially
DS_CHANGES	Object(DS_C_ENTRY_MOD)	—	1 or more	—

• DS_CHANGES

This attribute identifies the modifications to be made (in the order specified) to the directory entry of the specified object.

11.23 DS_C_ERROR

The OM class *DS_C_ERROR* comprises the parameters common to all errors.

It is an abstract OM class with the OM attributes of its superclass, *OM_C_OBJECT*, in addition to the OM attribute listed in Table 11-16.

Table 11–16. OM Attribute of DS_C_ERROR

OM Attribute	Value Syntax	Value Length	Value Number	Value Initially
DS_PROBLEM	Enum(DS_Problem)	—	1	—

Details of errors are returned in an instance of a subclass of this OM class. Each such subclass represents a particular kind of error, and is one of the following:

- **DS_C_ABANDON_FAILED**
- **DS_C_ATTRIBUTE_PROBLEM**
- **DS_C_COMMUNICATIONS_ERROR**
- **DS_C_LIBRARY_ERROR**
- **DS_C_NAME_ERROR**
- **DS_C_SECURITY_ERROR**
- **DS_C_SERVICE_ERROR**
- **DS_C_SYSTEM_ERROR**
- **DS_C_UPDATE_ERROR**

A number of possible values are defined for these subclasses. DCE XDS does not return other values for error conditions described in this chapter. Information on system errors can be found in Section 11.44. The following is a list of the error values. Each error OM class section defines the possible error values associated with that class. For a description of the errors, refer to the *OSF DCE Problem Determination Guide*.

- **DS_E_ADMIN_LIMIT_EXCEEDED**
- **DS_E_AFFECTS_MULTIPLE_DSAS**
- **DS_E_ALIAS_DEREFERENCING_PROBLEM**
- **DS_E_ALIAS_PROBLEM**
- **DS_E_ATTRIBUTE_OR_VALUE_EXISTS**
- **DS_E_BAD_ARGUMENT**
- **DS_E_BAD_CLASS**

- **DS_E_BAD_CONTEXT**
- **DS_E_BAD_NAME**
- **DS_E_BAD_SESSION**
- **DS_E_BAD_WORKSPACE**
- **DS_E_BUSY**
- **DS_E_CANNOT_ABANDON**
- **DS_E_CHAINING_REQUIRED**
- **DS_E_COMMUNICATIONS_PROBLEM**
- **DS_E_CONSTRAINT_VIOLATION**
- **DS_E_DIT_ERROR**
- **DS_E_ENTRY_EXISTS**
- **DS_E_INAPPROP_AUTHENTICATION**
- **DS_E_INAPPROP_MATCHING**
- **DS_E_INSUFFICIENT_ACCESS_RIGHTS**
- **DS_E_INVALID_ATTRIBUTE_SYNTAX**
- **DS_E_INVALID_ATTRIBUTE_VALUE**
- **DS_E_INVALID_CREDENTIALS**
- **DS_E_INVALID_REF**
- **DS_E_INVALID_SIGNATURE**
- **DS_E_LOOP_DETECTED**
- **DS_E_MISCELLANEOUS**
- **DS_E_MISSING_TYPE**
- **DS_E_MIXED_SYNCHRONOUS**
- **DS_E_NAMING_VIOLATION**
- **DS_E_NO_INFO**
- **DS_E_NO_SUCH_ATTRIBUTE_OR_VALUE**
- **DS_E_NO_SUCH_OBJECT**

- **DS_E_NO_SUCH_OPERATION**
- **DS_E_NOT_ALLOWED_ON_NON_LEAF**
- **DS_E_NOT_ALLOWED_ON_RDN**
- **DS_E_NOT_SUPPORTED**
- **DS_E_OBJECT_CLASS_MOD_PROHIB**
- **DS_E_OBJECT_CLASS_VIOLATION**
- **DS_E_OUT_OF_SCOPE**
- **DS_E_PROTECTION_REQUIRED**
- **DS_E_TIME_LIMIT_EXCEEDED**
- **DS_E_TOO_LATE**
- **DS_E_TOO_MANY_OPERATIONS**
- **DS_E_TOO_MANY_SESSIONS**
- **DS_E_UNABLE_TO_PROCEED**
- **DS_E_UNAVAILABLE**
- **DS_E_UNAVAILABLE_CRIT_EXT**
- **DS_E_UNDEFINED_ATTRIBUTE_TYPE**
- **DS_E_UNWILLING_TO_PERFORM**

11.24 DS_C_EXT

An instance of OM class **DS_C_EXT** indicates that a standardized extension to the directory service is outlined in the standards. Such extensions will only be standardized in post-1988 versions of the standards. Therefore, this OM class is not used by the XDS API and is only included for X/Open conformance purposes.

An instance of this OM class has the OM attributes of its superclass, *OM_C_OBJECT*, in addition to the OM attributes listed in Table 11-17.

Table 11–17. OM Attributes of DS_C_EXT

OM Attribute	Value Syntax	Value Length	Value Number	Value Initially
DS_CRIT	OM_S_BOOLEAN	—	1	OM_FALSE
DS_IDENT	OM_S_INTEGER	—	1	—
DS_ITEM_PARAMETERS	Any	—	1	—

- **DS_CRIT**

This attribute must have one of the following values:

- **OM_FALSE**, meaning that the originator permits the operation to be performed even if the extension is not available.
- **OM_TRUE**, meaning that the originator mandates that the extended operation be performed. If the extended operation is not performed, an error is reported.

- **DS_IDENT**

This attribute identifies the service extension.

- **DS_ITEM_PARAMETERS**

This OM attribute supplies the parameters of the extension. Its syntax is determined by the value of **DS_IDENT**.

11.25 DS_C_FILTER

An instance of OM class **DS_C_FILTER** is used to select or reject an object on the basis of information in its directory entry. At any point in time, an attribute filter has a value relative to every object. The value is FALSE, TRUE, or undefined. The object is selected if, and only if, the filter's value is TRUE.

An instance of this OM class has the OM attributes of its superclass, *OM_C_OBJECT*, in addition to the OM attributes listed in Table 11-18.

Table 11–18. OM Attributes of DS_C_FILTER

OM Attribute	Value Syntax	Value Length	Value Number	Value Initially
DS_FILTER_ITEMS	Object(DS_C_FILTER_ITEM)	—	0 or more	—
DS_FILTERS	Object(DS_C_FILTER)	—	0 or more	—
DS_FILTER_TYPE	Enum(DS_Filter_Type)	—	1	DS_AND

A *filter* is a collection of less elaborate filters and elementary **DS_FILTER_ITEMS**, together with a Boolean operation. The filter value is undefined if, and only if, all the component **DS_FILTERS** and **DS_FILTER_ITEMS** are undefined. Otherwise, the filter has a Boolean value with respect to any directory entry, which can be determined by evaluating each of the nested components and combining their values using the Boolean operation. The components whose values are undefined are ignored.

- **DS_FILTER_ITEMS**

This attribute is a collection of assertions, each relating to just one attribute of a directory entry.

- **DS_FILTERS**

This attribute is a collection of simpler filters.

- **DS_FILTER_TYPE**

This attribute is the filter's type. It can have any of the following values:

- **DS_AND**, meaning that the filter is the logical conjunction of its components. The filter is TRUE unless any of the nested filters or filter items is FALSE. If there are no nested components, the filter is TRUE.
- **DS_OR**, meaning that the filter is the logical disjunction of its components. The filter is FALSE unless any of the nested filters or filter items is TRUE. If there are no nested components, the filter is FALSE.
- **DS_NOT**, meaning that the result of this filter is reversed. There must be exactly one nested filter or filter item. The filter is TRUE if the enclosed

filter or filter item is FALSE, and it is FALSE if the enclosed filter or filter item is TRUE.

11.26 DS_C_FILTER_ITEM

An instance of OM class **DS_C_FILTER_ITEM** is a component of **DS_C_FILTER**. It is an assertion about the existence or values of a single attribute type in a directory entry.

An instance of this OM class has the OM attributes of its superclasses, *OM_C_OBJECT* and **DS_C_ATTRIBUTE**, in addition to the OM attributes listed in Table 11-19.

Table 11-19. OM Attributes of DS_C_FILTER_ITEM

OM Attribute	Value Syntax	Value Length	Value Number	Value Initially
DS_FILTER_ITEM_TYPE	Enum(DS_Filter_Item_Type)	—	1	—
DS_FINAL_SUBSTRING	String(*)	1 or more	0 or 1	—
DS_INITIAL_SUBSTRING	String(*)	1 or more	0 or 1	—

Note: OM attributes **DS_ATTRIBUTE_TYPE** and **DS_ATTRIBUTE_VALUES** are inherited from the superclass **DS_C_ATTRIBUTE**.

The value of the filter item is undefined in the following cases:

- The **DS_ATTRIBUTE_TYPE** is not known.
- None of the **DS_ATTRIBUTE_VALUES** conform to the attribute syntax defined for that attribute type.
- The **DS_FILTER_ITEM_TYPE** uses a matching rule that is not defined for the attribute syntax.

Access control restrictions can also cause the value to be undefined.

- **DS_FILTER_ITEM_TYPE**

This attribute identifies the type of filter item and, thus, the nature of the filter. The filter item can adopt any of the following values:

- **DS_APPROXIMATE_MATCH**, meaning that the filter is TRUE if the directory entry contains at least one value of the specified type that is approximately equal to that specified (the meaning of “approximately equal” is implementation dependent); otherwise, the filter is FALSE.

Rules for approximate matching are defined locally. For example, an approximate match may take into account spelling variations or employ phonetic comparison rules. In the absence of any such capabilities, a DSA needs to treat an approximate match as a test for equality. DCE GDS supports phonetic comparisons. There must be exactly one value of the OM attribute **DS_ATTRIBUTE_VALUES**.

- **DS_EQUALITY**, meaning that the filter is TRUE if the entry contains at least one value of the specified type that is equal to the value specified, according to the equality matching rule in force; otherwise, the filter is FALSE. There must be exactly one value of the OM attribute **DS_ATTRIBUTE_VALUES**.
- **DS_GREATER_OR_EQUAL**, meaning that the filter item is TRUE if, and only if, at least one value of the attribute is greater than or equal to the supplied value. There must be exactly one value of the OM attribute **DS_ATTRIBUTE_VALUES**.
- **DS_LESS_OR_EQUAL**, meaning that the filter item is TRUE if, and only if, at least one value of the attribute is less than or equal to the supplied value. There must be exactly one value of the OM attribute **DS_ATTRIBUTE_VALUES**.
- **DS_PRESENT**, meaning that the filter is TRUE if the entry contains an attribute of the specified type; otherwise, it is FALSE.

Any values of the OM attribute **DS_ATTRIBUTE_VALUES** are ignored.

- **DS_SUBSTRINGS**, meaning that the filter is TRUE if the entry contains at least one value of the specified attribute type that contains all of the specified substrings in the given order; otherwise, the filter is FALSE.

Any number of substrings can be given as values of the OM attribute **DS_ATTRIBUTE_VALUES**. Similarly, no substrings can be specified. There can also be a substring in **DS_INITIAL_SUBSTRING** or **DS_FINAL_SUBSTRING**, or both. The substrings do not overlap, but they can be separated from each other or from the ends of the attribute value by zero or more string elements. However, at least one attribute of type **DS_ATTRIBUTE_VALUES**, **DS_INITIAL_SUBSTRING**, or **DS_FINAL_SUBSTRING** must exist.

- **DS_FINAL_SUBSTRING**

If present, this attribute is the substring that will match the final part of an attribute value in the entry. This attribute can only exist if the **DS_FILTER_ITEM_TYPE** is equal to **DS_SUBSTRINGS**.

- **DS_INITIAL_SUBSTRING**

If present, this attribute is the substring that will match the initial part of an attribute value in the entry.

11.27 DS_C_LIBRARY_ERROR

An instance of OM class **DS_C_LIBRARY_ERROR** reports an error detected by the interface function library.

An application is not permitted to create or modify instances of this OM class. An instance of this OM class has the OM attributes of its superclasses, *OM_C_OBJECT* and *DS_C_ERROR*, and no additional OM attributes.

Each function has several possible errors that can be detected by the library itself and that are returned directly by the subroutine. These errors occur when the library itself is incapable of performing an action, submitting a service request, or deciphering a response from the directory service.

The OM attribute **DS_PROBLEM**, which is inherited from the superclass *DS_C_ERROR*, identifies the particular library error that occurred. (In reference pages, the ERRORS section of each function description lists the errors that the respective function can return.) Its value is one of the following:

- **DS_E_BAD_ARGUMENT**
- **DS_E_BAD_CLASS**
- **DS_E_BAD_CONTEXT**
- **DS_E_BAD_NAME**
- **DS_E_BAD_SESSION**
- **DS_E_MISCELLANEOUS**
- **DS_E_MISSING_TYPE**
- **DS_E_MIXED_SYNCHRONOUS**
- **DS_E_NOT_SUPPORTED**
- **DS_E_TOO_MANY_OPERATIONS**
- **DS_E_TOO_MANY_SESSIONS**

11.28 DS_C_LIST_INFO

An instance of OM class **DS_C_LIST_INFO** is part of the results of **ds_list()**.

An application is not permitted to create or modify instances of this OM class. An instance of this OM class has the OM attributes of its superclasses, *OM_C_OBJECT* and *DS_C_COMMON_RESULTS*, in addition to the OM attributes listed in Table 11-20.

Table 11–20. OM Attributes of DS_C_LIST_INFO

OM Attribute	Value Syntax	Value Length	Value Number	Value Initially
DS_OBJECT_NAME	Object(<i>DS_C_NAME</i>)	—	0 or 1	—
DS_PARTIAL_OUTCOME_QUAL	Object(DS_C_PARTIAL_OUTCOME_QUAL)	—	0 or 1	—
DS_SUBORDINATES	Object(DS_C_LIST_INFO_ITEM)	—	0 or more	—

- **DS_OBJECT_NAME**

This attribute is the DN of the target object of the operation. It is present if the OM attribute **DS_ALIAS_DEREFERENCED**, inherited from the superclass *DS_C_COMMON_RESULTS*, is **OM_TRUE**.

- **DS_PARTIAL_OUTCOME_QUAL**

This OM attribute value is present if the list of subordinates is incomplete. The DSA or DSAs that provided this list did not complete the search for some reason. The partial outcome qualifier contains details of why the search is not completed, and which areas of the directory have not been searched.

- **DS_SUBORDINATES**

This attribute contains information about zero or more subordinate objects identified by **ds_list()**.

11.29 DS_C_LIST_INFO_ITEM

An instance of OM class **DS_C_LIST_INFO_ITEM** comprises details returned by **ds_list()** of a single subordinate object.

An application is not permitted to create or modify instances of this OM class. An instance of this OM class has the OM attributes of its superclass, *OM_C_OBJECT*, in addition to the OM attributes listed in Table 11-21.

Table 11-21. OM Attributes of *DS_C_LIST_INFO_ITEM*

OM Attribute	Value Syntax	Value Length	Value Number	Value Initially
DS_ALIAS_ENTRY	OM_S_BOOLEAN	—	1	—
DS_FROM_ENTRY	OM_S_BOOLEAN	—	1	—
DS_RDN	Object(<i>DS_C_RELATIVE_NAME</i>)	—	1	—

- **DS_ALIAS_ENTRY**

This attribute indicates whether the object is an alias.

- **DS_FROM_ENTRY**

This attribute indicates whether information about the object was obtained directly from its directory entry, rather than from a copy of the entry.

- **DS_RDN**

This attribute contains the RDN of the object. If this is the name of an alias entry, as indicated by **DS_ALIAS_ENTRY**, it is not dereferenced.

11.30 DS_C_LIST_RESULT

An instance of OM class **DS_C_LIST_RESULT** comprises the results of a successful call to **ds_list()**.

An application is not permitted to create or modify instances of this OM class. An instance of this OM class has the OM attributes of its superclass, *OM_C_OBJECT*, in addition to the OM attributes listed in Table 11-22.

Table 11–22. OM Attributes of DS_C_LIST_RESULT

OM Attribute	Value Syntax	Value Length	Value Number	Value Initially
DS_LIST_INFO	Object(DS_C_LIST_INFO)	—	0 or 1	—
DS_UNCORRELATED_LIST_INFO	Object(DS_C_LIST_RESULT)	—	0 or more	—

Note: No instance contains values of both OM attributes.

- **DS_LIST_INFO**

This attribute contains the full results of **ds_list()**, or just part of them.

- **DS_UNCORRELATED_LIST_INFO**

When the DUA requests a protection request of *signed*, the information returned can comprise a number of sets of results originating from, and signed by, different components of the directory. Implementations can reflect this structure by nesting **DS_LIST_RESULT** OM objects as values of this OM attribute. Alternatively, they can collapse all results into a single value of the OM attribute **DS_LIST_INFO**. The DCE directory service does not support the optional feature of signed results; therefore, this OM attribute is never present.

11.31 DS_C_NAME

The OM class *DS_C_NAME* represents a name of an object in the directory, or a part of such a name.

It is an abstract class that has the attributes of its superclass, *OM_C_OBJECT*, and no other OM attributes.

A name uniquely distinguishes the object from all other objects whose entries are displayed in the DIT. However, an object can have more than one name; that is, a name need not be unique. A DN is unique; there are no other DNs that identify

the same object. An RDN is part of a name and only distinguishes the object from others that are its siblings.

Most of the interface functions take a *name* parameter, the value of which must be an instance of one of the subclasses of this OM class. Thus, this OM class is useful for amalgamating all possible representations of names.

The DCE XDS implementation defines one subclass of this OM class and, thus, a single representation for names; that is, **DS_C_DS_DN**, which provides a representation for names, including DNs.

11.32 DS_C_NAME_ERROR

An instance of OM class **DS_C_NAME_ERROR** reports a name-related directory service error.

An application is not permitted to create or modify instances of this OM class. An instance of this OM class has the OM attributes of its superclasses, *OM_C_OBJECT* and *DS_C_ERROR*, in addition to the OM attribute listed in Table 11-23.

Table 11-23. OM Attribute of DS_C_NAME_ERROR

OM Attribute	Value Syntax	Value Length	Value Number	Value Initially
DS_MATCHED	Object(<i>DS_C_NAME</i>)	—	1	—

- **DS_MATCHED**

This attribute identifies the initial part (up to, but excluding, the first RDN that is unrecognized) of the name that is supplied, or of the name resulting from dereferencing an alias. It names the lowest entry (object or alias) in the DIT that is matched.

The OM attribute **DS_PROBLEM**, which is inherited from the superclass *DS_C_ERROR*, identifies the cause of the failure. Its value is one of the following:

- **DS_E_ALIAS_DEREFERENCING_PROBLEM**

- **DS_E_ALIAS_PROBLEM**
- **DS_E_INVALID_ATTRIBUTE_VALUE**
- **DS_E_NO_SUCH_OBJECT**

11.33 DS_C_OPERATION_PROGRESS

An instance of OM class **DS_C_OPERATION_PROGRESS** specifies the progress or processing state of a directory request.

An application is not permitted to create or modify instances of this OM class. An instance of this OM class has the OM attributes of its superclass, *OM_C_OBJECT*, in addition to the OM attributes listed in Table 11-24.

Table 11-24. OM Attributes of DS_C_OPERATION_PROGRESS

OM Attribute	Value Syntax	Value Length	Value Number	Value Initially
DS_NAME_RESOLUTION_PHASE	Enum(DS_Name_Resolution_Phase)	—	1	—
DS_NEXT_RDN_TO_BE_RESOLVED	OM_S_INTEGER	—	0 or 1	—

The target name mentioned as follows is the name upon which processing of the directory request is currently focused.

- **DS_NAME_RESOLUTION_PHASE**

This attribute indicates what phase is reached in handling the target name. It must have one of the following values:

- **DS_COMPLETED**, meaning that the DSA holding the target object is reached.
- **DS_NOT_STARTED**, meaning that so far a DSA is not reached with a naming context containing the initial RDNs of the name.

- **DS_PROCEEDING**, meaning that the initial part of the name has been recognized, although the DSA holding the target object has not yet been reached.

- **DS_NEXT_RDN_TO_BE_RESOLVED**

This attribute indicates to the DSA which of the RDNs in the target name is next to be resolved. It takes the form of an integer in the range from 1 to the number of RDNs in the name. This OM attribute only has a value if the value of **DS_NAME_RESOLUTION_PHASE** is **DS_PROCEEDING**.

The constant **DS_OPERATION_NOT_STARTED** can be used in the **DS_C_CONTEXT** of an operation instead of an instance of this OM class.

11.34 DS_C_PARTIAL_OUTCOME_QUAL

An instance of OM class **DS_C_PARTIAL_OUTCOME_QUAL** explains to what extent the results of a call to **ds_list()** or **ds_search()** are incomplete and why.

An application is not permitted to create or modify instances of this OM class. An instance of this OM class has the OM attributes of its superclass, *OM_C_OBJECT*, in addition to the OM attributes listed in Table 11-25.

Table 11–25. OM Attributes of a DS_C_PARTIAL_OUTCOME_QUAL

OM Attribute	Value Syntax	Value Length	Value Number	Value Initially
DS_LIMIT_PROBLEM	Enum(DS_Limit_Problem)	—	1	—
DS_UNAVAILABLE_CRIT_EXT	OM_S_BOOLEAN	—	1	—
DS_UNEXPLORED	Object(DS_C_CONTINUATION_REF)	—	0 or more	—

- **DS_LIMIT_PROBLEM**

This attribute explains fully or partly why the results are incomplete. It can have one of the following values:

- **DS_ADMIN_LIMIT_EXCEEDED**, meaning that an administrative limit is reached.
- **DS_NO_LIMIT_EXCEEDED**, meaning that there is no limit problem.
- **DS_SIZE_LIMIT_EXCEEDED**, meaning that the maximum number of objects specified as a service control is reached.
- **DS_TIME_LIMIT_EXCEEDED**, meaning that the maximum number of seconds specified as a service control is reached.

- **DS_UNAVAILABLE_CRIT_EXT**

If **OM_TRUE**, this attribute indicates that some part of the directory service cannot provide a requested critical service extension. The user requested one or more standard service extensions by including values of the OM attribute **DS_EXT** in the **DS_C_CONTEXT** supplied for the operation. Furthermore, the user indicated that some of these extensions are essential by setting the OM attribute **DS_CRIT** in the extension to **OM_TRUE**. Some of the critical extensions cannot be performed by one particular DSA or by a number of DSAs. In general, it is not possible to determine which DSA could not perform which particular extension.

- **DS_UNEXPLORED**

This attribute identifies any regions of the directory that are left unexplored in such a way that the directory request can be continued. Only continuation references within the scope specified by the **DS_SCOPE_OF_REFERRAL** service control are included.

11.35 DS_C_PRESENTATION_ADDRESS

An instance of OM class **DS_C_PRESENTATION_ADDRESS** is a presentation address of an OSI application entity, which is used for OSI communications with this instance.

An instance of this OM class has the OM attributes of its superclasses, *OM_C_OBJECT* and *DS_C_ADDRESS*, in addition to the OM attributes listed in Table 11-26.

Table 11-26. OM Attributes of *DS_C_PRESENTATION_ADDRESS*

OM Attribute	Value Syntax	Value Length	Value Number	Value Initially
DS_N_ADDRESSES	String(OM_S_OCTET_STRING)	—	1 or more	—
DS_P_SELECTOR	String(OM_S_OCTET_STRING)	—	0 or 1	—
DS_S_SELECTOR	String(OM_S_OCTET_STRING)	—	0 or 1	—
DS_T_SELECTOR	String(OM_S_OCTET_STRING)	—	0 or 1	—

- **DS_N_ADDRESSES**

This attribute is the network addresses of the application entity.

- **DS_P_SELECTOR**

This attribute is the presentation selector.

- **DS_S_SELECTOR**

This attribute is the session selector.

- **DS_T_SELECTOR**

This attribute is the transport selector.

11.36 DS_C_READ_RESULT

An instance of OM class **DS_C_READ_RESULT** comprises the result of a successful call to **ds_read()**. An application is not permitted to create or modify instances of this

OM class. An instance of this OM class has the OM attributes of its superclasses, *OM_C_OBJECT* and *DS_C_COMMON_RESULTS*, in addition to the OM attribute listed in Table 11-27.

Table 11-27. OM Attribute of DS_C_READ_RESULT

OM Attribute	Value Syntax	Value Length	Value Number	Value Initially
DS_ENTRY	Object(DS_C_ENTRY_INFO)	—	1	—

- **DS_ENTRY**

This attribute contains the information extracted from the directory entry of the target object.

11.37 DS_C_REFERRAL

An instance of OM class **DS_C_REFERRAL** reports failure to perform an operation and redirects the requestor to one or more access points better equipped to perform the operation.

An application is not permitted to create or modify instances of this OM class. An instance of this OM class has the OM attributes of its superclasses, *OM_C_OBJECT* and **DS_C_CONTINUATION_REF**, and no additional OM attributes.

The referral is a continuation reference by means of which the operation can progress.

11.38 DS_C_RELATIVE_NAME

The OM class *DS_C_RELATIVE_NAME* represents the RDNs of objects in the directory. It is an abstract class, which has the attributes of its superclass, *OM_C_OBJECT*, and no other OM attributes.

An RDN is part of a name, and only distinguishes the object from others that are its siblings. This OM class is used to accumulate all possible representations of RDNs. An argument of interface functions that is an RDN, or an OM attribute value that is an RDN is an instance of one of the subclasses of this OM class.

The DCE XDS API defines one subclass of this OM class, and, thus, a single representation for RDNs; that is, **DS_C_DS_RDN**, which provides a representation for RDNs.

11.39 DS_C_SEARCH_INFO

An instance of OM class **DS_C_SEARCH_INFO** is part of the result of **ds_search()**.

An application is not permitted to create or modify instances of this OM class. An instance of this OM class has the OM attributes of its superclasses, *OM_C_OBJECT* and *DS_C_COMMON_RESULTS*, in addition to the OM attributes listed in Table 11-28.

Table 11-28. OM Attributes of DS_C_SEARCH_INFO

OM Attribute	Value Syntax	Value Length	Value Number	Value Initially
DS_ENTRIES	Object(DS_C_ENTRY_INFO)	—	0 or more	—
DS_OBJECT_NAME	Object(<i>DS_C_NAME</i>)	—	0 or 1	—
DS_PARTIAL_OUTCOME_QUAL	Object(DS_C_PARTIAL_OUTCOME_QUAL)	—	0 or 1	—

- **DS_ENTRIES**

This attribute contains information about zero or more objects found by **ds_search()** that matched the given selection criteria.

- **DS_OBJECT_NAME**

This attribute contains the DN of the target object of the operation. It is present if the OM attribute **DS_ALIAS_DEREFERENCED**, inherited from the superclass *DS_C_COMMON_RESULTS*, is **OM_TRUE**.

- **DS_PARTIAL_OUTCOME_QUAL**

This OM attribute value is only present if the list of entries is incomplete. The DSA or DSAs that provided this list did not complete the search for some reason. The partial outcome qualifier contains details of why the search was not completed and which areas of the directory were not searched.

11.40 DS_C_SEARCH_RESULT

An instance of OM class **DS_C_SEARCH_RESULT** comprises the result of a successful call to **ds_search()**.

An application is not permitted to create or modify instances of this OM class. An instance of this OM class has the OM attributes of its superclass, *OM_C_OBJECT*, in addition to the OM attributes listed in Table 11-29.

Table 11-29. OM Attributes of DS_C_SEARCH_RESULT

OM Attribute	Value Syntax	Value Length	Value Number	Value Initially
DS_SEARCH_INFO	Object(DS_C_SEARCH_INFO)	—	0 or 1	—
DS_UNCORRELATED_SEARCH_INFO	Object(DS_C_SEARCH_RESULT)	—	0 or more	—

Note: No instance contains values of both OM attributes.

- **DS_SEARCH_INFO**

This attribute contains the full result of **ds_search()**, or part of the result.

- **DS_UNCORRELATED_SEARCH_INFO**

When the DUA requests a protection request of *signed*, the information returned can comprise a number of sets of results originating from and signed by different components of the directory service. Implementations can reflect this structure by nesting **DS_C_SEARCH_RESULT** OM objects as values of this OM attribute. Alternatively, they can collapse all results into a single value of the OM attribute **DS_SEARCH_INFO**. The DCE directory service does not support the optional feature of signed results; therefore, this OM attribute is never present.

11.41 DS_C_SECURITY_ERROR

An instance of OM class **DS_C_SECURITY_ERROR** reports a security-related directory service error.

An application is not permitted to create or modify instances of this OM class. An instance of this OM class has the OM attributes of its superclasses, **OM_C_OBJECT** and **DS_C_ERROR**, and no additional OM attributes.

The OM attribute **DS_PROBLEM**, which is inherited from the superclass **DS_C_ERROR**, identifies the cause of this failure. Its value is one of the following:

- **DS_E_INAPPROP_AUTHENTICATION**
- **DS_E_INSUFFICIENT_ACCESS_RIGHTS**
- **DS_E_INVALID_CREDENTIALS**
- **DS_E_INVALID_SIGNATURE**
- **DS_E_NO_INFO**
- **DS_E_PROTECTION_REQUIRED**

11.42 DS_C_SERVICE_ERROR

An instance of OM class **DS_C_SERVICE_ERROR** reports a directory service error related to the provision of the service.

An application is not permitted to create or modify instances of this OM class. An instance of this OM class has the OM attributes of its superclasses, *OM_C_OBJECT* and *DS_C_ERROR*, and no additional OM attributes.

The OM attribute **DS_PROBLEM**, which is inherited from the superclass *DS_C_ERROR*, identifies the cause of the failure. Its value is one of the following:

- **DS_E_ADMIN_LIMIT_EXCEEDED**
- **DS_E_BUSY**
- **DS_E_CHAINING_REQUIRED**
- **DS_E_DIT_ERROR**
- **DS_E_INVALID_REF**
- **DS_E_LOOP_DETECTED**
- **DS_E_OUT_OF_SCOPE**
- **DS_E_TIME_LIMIT_EXCEEDED**
- **DS_E_UNABLE_TO_PROCEED**
- **DS_E_UNAVAILABLE**
- **DS_E_UNAVAILABLE_CRIT_EXT**
- **DS_E_UNWILLING_TO_PERFORM**

11.43 DS_C_SESSION

An instance of OM class **DS_C_SESSION** identifies a particular link from the application program to a DUA.

An instance of this OM class has the OM attributes of its superclass, *OM_C_OBJECT*, in addition to the OM attributes listed in Table 11-30.

Table 11–30. OM Attributes of DS_C_SESSION

OM Attribute	Value Syntax	Value Length	Value Number	Value Initially
DS_DSA_ADDRESS	Object(<i>DS_C_ADDRESS</i>)	—	0 or 1	<i>local</i> ¹
DS_DSA_NAME	Object(<i>DS_C_NAME</i>)	—	0 or 1	<i>local</i> ¹
DS_FILE_DESCRIPTOR	OM_S_INTEGER	—	1	See text
DS_REQUESTOR	Object(<i>DS_C_NAME</i>)	—	0 or 1	—

- 1 The default values of these OM attributes are set to the address and name of the default DSA entry in the local cache. If this cache entry is not present, then these OM attributes are absent.

The **DS_C_SESSION** gathers all the information that describes a particular directory interaction. The parameters that will control such a session are set up in an instance of this OM class, which is then passed as an argument to **ds_bind()**. This sets the OM attributes that describe the actual characteristics of this session, and then starts the session. A session started in this way must pass as the first argument to each interface function. The result of modifying an initiated session is unspecified. Finally, **ds_unbind()** is used to terminate the session, after which the parameters can be modified and a new session started using the same instance, if required. Multiple concurrent sessions can run using multiple instances of this OM class.

The OM attributes of a session are as follows:

- **DS_DSA_ADDRESS**

This attribute indicates the address of the default DSA named by **DS_DSA_NAME**.

- **DS_DSA_NAME**

This attribute indicates the DN of the DSA that is used by default to service directory requests.

- **DS_FILE_DESCRIPTOR** (Optional Functionality)

This OM attribute is not used by DCE XDS and is always set to **DS_NO_VALID_FILE_DESCRIPTOR**.

- **DS_REQUESTOR**

This attribute is the DN of the user of this directory service session.

Applications can assume that an object of OM class **DS_C_SESSION**, created with default values of all its OM attributes, works with all the interface functions. Local administrators need to ensure that this is the case. Such a session can be created by passing the constant **DS_DEFAULT_SESSION** as an argument to **ds_bind()**.

11.44 DS_C_SYSTEM_ERROR

An instance of OM class **DS_C_SYSTEM_ERROR** reports an error that occurred in the underlying operating system.

An application is not permitted to create or modify instances of this OM class. An instance of this OM class has the OM attributes of its superclasses, *OM_C_OBJECT* and *DS_C_ERROR*, and no additional OM attributes, although there can be additional implementation-defined OM attributes.

The OM attribute **DS_PROBLEM**, which is inherited from the superclass *DS_C_ERROR*, identifies the cause of the failure. Its value is the same as that of **errno** defined in the C language.

The standard names of system errors are defined in Volume 2 of the *X/Open Portability Guide*.

If such an error persists, a **DS_C_LIBRARY_ERROR (DS_E_MISCELLANEOUS)** is reported.

11.45 DS_C_UPDATE_ERROR

An instance of OM class **DS_C_UPDATE_ERROR** reports a directory service error peculiar to a modification operation.

An application is not permitted to create or modify instances of this OM class. An instance of this OM class has the OM attributes of its superclasses, *OM_C_OBJECT* and *DS_C_ERROR*, and no additional OM attributes.

The OM attribute **DS_PROBLEM**, which is inherited from the superclass *DS_C_ERROR*, identifies the cause of the failure. Its value is one of the following:

- **DS_E_AFFECTS_MULTIPLE_DSAS**
- **DS_E_ENTRY_EXISTS**
- **DS_E_NAMING_VIOLATION**
- **DS_E_NOT_ALLOWED_ON_NON_LEAF**
- **DS_E_NOT_ALLOWED_ON_RDN**
- **DS_E_OBJECT_CLASS_MOD_PROHIB**
- **DS_E_OBJECT_CLASS_VIOLATION**

Chapter 12

Basic Directory Contents Package

The standards define a number of attribute types (known as the *selected attribute types*), attribute syntaxes, attribute sets, and object classes (known as the *selected object classes*¹). These definitions allow the creation and maintenance of directory entries for a number of common objects so that the representation of all such objects is the same throughout the directory. They include such objects as **Country**, **Person**, and **Organization**.

This chapter outlines names for each of these items, and defines OM classes to represent those that are not represented directly by OM syntaxes. The attribute values in the directory are not restricted to those discussed in this chapter, and new attribute types and syntaxes can be created at any time. (For further information on how the values of other syntaxes are represented in the interface, see Chapter 10.)

1. These definitions are chiefly in *The Directory: Selected Attribute Types* (ISO 9594-6, CCITT X.520) and *The Directory: Selected Object Classes* (ISO 9594-7, CCITT X.521) with additional material in *The Directory: Overview of Concepts, Models, and Services* (ISO 9594-1, CCITT X.500) and *The Directory: Authentication Framework* (ISO 9594-8, CCITT X.509).

The constants and OM classes in this chapter are defined in addition to those in Chapter 11, since they are not essential to the working of the interface, but instead allow directory entries to be utilized. The definitions belong to the basic directory contents package (BDCP), which is supported by the DCE XDS API following negotiation of its use with **ds_version()**.

Note: The definitions for the GDS package are provided in Chapter 15. The definitions for the strong authentication package are provided in Chapter 13. The definitions for the MHS directory user package are provided in Chapter 14.

The object identifier associated with the BDCP is

```
{iso(1) identified-organization(3) icd-ecma(0012) member-company(2)
dec(1011) xopen(28) bdc(1)}
```

It takes the following encoding:

```
\x2B\xC\x2\x87\x73\x1C\x1
```

This identifier is represented by the constant **DS_BASIC_DIR_CONTENTS_PKG**. The C constants associated with this package are in the **xdsbdc.h** header file. (See the **xdsbdc.h(4xds)** reference page.)

The concepts and notation used are introduced in Chapter 11. They are also fully explained in Chapters 17 through 19.

The selected attribute types are presented first, followed by the selected object classes. Next, the OM class hierarchy and OM class definitions required to support the selected attribute types are presented.

12.1 Selected Attribute Types

This section presents the attribute types, defined in the standards, which are to be used in directory entries. Each directory entry is composed of a number of attributes, each of which comprises an attribute type together with one or more attribute values. The form of each value of an attribute is determined by the attribute syntax associated with the attribute's type.

In the interface, attributes are displayed as instances of OM class **DS_C_ATTRIBUTE** with the attribute type represented as the value of the OM attribute **DS_ATTRIBUTE_TYPE**, and the attribute value (or values) represented as the value (or values) of the OM attribute **DS_ATTRIBUTE_VALUES**. Each attribute type has an object identifier, assigned in the standards, which is the value of the OM attribute **DS_ATTRIBUTE_TYPE**. These object identifiers are represented in the interface by constants with the same name as the directory attribute, and they are prefixed with **DS_A_** so that they can be easily identified.

Table 12-1 shows the names of the attribute types defined in the standards, together with the BER encoding of the object identifiers associated with each of them. Table 12-2 shows the names of the attribute types, together with the OM value syntax that is used in the interface to represent values of that attribute type. Table 12-2 also includes the range of lengths permitted for the string types. This indicates whether the attribute can be multivalued and which matching rules are provided for the syntax. Following the table is a brief description of each attribute.

The standards define matching rules that are used for deciding whether two values are equal (E), for ordering (O) two values, and for identifying one value as a substring (S) of another in directory service operations. Specific matching rules are given in this chapter for certain attributes. In addition, the following general rules apply as indicated:

- All attribute values whose syntax is String(**OM_S_NUMERIC_STRING**), String(**OM_S_PRINTABLE_STRING**), or String(**OM_S_TELETEX_STRING**) are considered insignificant for the following reasons:
 - Differences caused by the presence of spaces preceding the first printing character
 - Spaces following the last printing character
 - More than one consecutive space anywhere within the value
- For all attribute values whose syntax is String(**OM_S_TELETEX_STRING**), differences in the case of alphabetical characters are considered insignificant.

Note: The third and fourth columns of Table 12-1 contain the contents octets of the BER encoding of the object identifier. All these object identifiers stem from the root **{joint-iso-ccitt(2) ds(5) attributeType(4)}**.

Table 12–1. Object Identifiers for Selected Attribute Types

Package	Attribute Type	Object Identifier BER	
		Decimal	Hexadecimal
BDCP	DS_A_ALIASED_OBJECT_NAME	85, 4, 1	\x55\x04\x01
BDCP	DS_A_BUSINESS_CATEGORY	85, 4, 15	\x55\x04\x0F
BDCP	DS_A_COMMON_NAME	85, 4, 3	\x55\x04\x03
BDCP	DS_A_COUNTRY_NAME	85, 4, 6	\x55\x04\x06
BDCP	DS_A_DESCRIPTION	85, 4, 13	\x55\x04\x0D
BDCP	DS_A_DEST_INDICATOR	85, 4, 27	\x55\x04\x1B
BDCP	DS_A_FACSIMILE_PHONE_NBR	85, 4, 23	\x55\x04\x17
BDCP	DS_A_INTERNAT_ISDN_NBR	85, 4, 25	\x55\x04\x19
BDCP	DS_A_KNOWLEDGE_INFO	85, 4, 2	\x55\x04\x02
BDCP	DS_A_LOCALITY_NAME	85, 4, 7	\x55\x04\x07
BDCP	DS_A_MEMBER	85, 4, 31	\x55\x04\x1F
BDCP	DS_A_OBJECT_CLASS	85, 4, 0	\x55\x04\x00
BDCP	DS_A_ORG_NAME	85, 4, 10	\x55\x04\x0A
BDCP	DS_A_ORG_UNIT_NAME	85, 4, 11	\x55\x04\x0B
BDCP	DS_A_OWNER	85, 4, 32	\x55\x04\x20
BDCP	DS_A_PHYS_DELIV_OFF_NAME	85, 4, 19	\x55\x04\x13
BDCP	DS_A_POST_OFFICE_BOX	85, 4, 18	\x55\x04\x12
BDCP	DS_A_POSTAL_ADDRESS	85, 4, 16	\x55\x04\x10
BDCP	DS_A_POSTAL_CODE	85, 4, 17	\x55\x04\x11
BDCP	DS_A_PREF_DELIV_METHOD	85, 4, 28	\x55\x04\x1C
BDCP	DS_A_PRESENTATION_ADDRESS	85, 4, 29	\x55\x04\x1D
BDCP	DS_A_REGISTERED_ADDRESS	85, 4, 26	\x55\x04\x1A
BDCP	DS_A_ROLE_OCCUPANT	85, 4, 33	\x55\x04\x21

Package	Attribute Type	Object Identifier BER	
		Decimal	Hexadecimal
BDCP	DS_A_SEARCH_GUIDE	85, 4, 14	\x55\x04\x0E
BDCP	DS_A_SEE_ALSO	85, 4, 34	\x55\x04\x22
BDCP	DS_A_SERIAL_NBR	85, 4, 5	\x55\x04\x05
BDCP	DS_A_STATE_OR_PROV_NAME	85, 4, 8	\x55\x04\x08
BDCP	DS_A_STREET_ADDRESS	85, 4, 9	\x55\x04\x09
BDCP	DS_A_SUPPORT_APPLIC_CONTEXT	85, 4, 3	\x55\x04\x1E
BDCP	DS_A_SURNAME	85, 4, 4	\x55\x04\x04
BDCP	DS_A_PHONE_NBR	85, 4, 20	\x55\x04\x14
BDCP	DS_A_TELETEX_TERM_IDENT	85, 4, 22	\x55\x04\x16
BDCP	DS_A_TELEX_NBR	85, 4, 21	\x55\x04\x15
BDCP	DS_A_TITLE	85, 4, 12	\x55\x04\x0C
BDCP	DS_A_USER_PASSWORD	85, 4, 35	\x55\x04\x23
BDCP	DS_A_X121_ADDRESS	85, 4, 24	\x55\x04\x18

Table 12–2. Representation of Values for Selected Attribute Types

Attribute Type	OM Value Syntax	Value Length	Multi-valued	Matching Rules
DS_A_ALIASED_OBJECT_NAME	Object(<i>DS_C_NAME</i>)	—	no	E
DS_A_BUSINESS_CATEGORY	String(OM_S_TELETEX_STRING)	1–128	yes	E, S
DS_A_COMMON_NAME	String(OM_S_TELETEX_STRING)	1–64	yes	E, S
DS_A_COUNTRY_NAME	String(OM_S_PRINTABLE_STRING) ¹	2	no	E

Attribute Type	OM Value Syntax	Value Length	Multi-valued	Matching Rules
DS_A_DESCRIPTION	String(OM_S_TELETEX_STRING)	1–1024	yes	E, S
DS_A_DEST_INDICATOR	String(OM_S_PRINTABLE_STRING) ²	1–128	yes	E, S
DS_A_FACSIMILE_PHONE_NBR	Object(DS_C_FACSIMILE_PHONE_NBR)	—	yes	—
DS_A_INTERNAT_ISDN_NBR	String(OM_S_NUMERIC_STRING) ³	1–16	yes	—
DS_A_KNOWLEDGE_INFO	String(OM_S_TELETEX_STRING)	—	yes	E, S
DS_A_LOCALITY_NAME	String(OM_S_TELETEX_STRING)	1–128	yes	E, S
DS_A_MEMBER	Object(DS_C_NAME)	—	yes	E
DS_A_OBJECT_CLASS	String(OM_S_OBJECT_IDENTIFIER_STRING)	—	yes	E
DS_A_ORG_NAME	String(OM_S_TELETEX_STRING)	1–64	yes	E, S
DS_A_ORG_UNIT_NAME	String(OM_S_TELETEX_STRING)	1–64	yes	E, S
DS_A_OWNER	Object(DS_C_NAME)	—	yes	E
DS_A_PHYS_DELIV_OFF_NAME	String(OM_S_TELETEX_STRING)	1–128	yes	E, S
DS_A_POST_OFFICE_BOX	String(OM_S_TELETEX_STRING)	1–40	yes	E, S
DS_A_POSTAL_ADDRESS	Object(DS_C_POSTAL_ADDRESS)	—	yes	E
DS_A_POSTAL_CODE	String(OM_S_TELETEX_STRING)	1–40	yes	E, S

Attribute Type	OM Value Syntax	Value Length	Multi-valued	Matching Rules
DS_A_PREF_DELIV_METHOD	Enum(DS_Preferred_Delivery_Method)	—	yes	—
DS_A_PRESENTATION_ADDRESS	Object(DS_C_PRESENTATION_ADDRESS)	—	no	E
DS_A_REGISTERED_ADDRESS	Object(DS_C_POSTAL_ADDRESS)	—	yes	—
DS_A_ROLE_OCCUPANT	Object(DS_C_NAME)	—	yes	E
DS_A_SEARCH_GUIDE	Object(DS_C_SEARCH_GUIDE)	—	yes	—
DS_A_SEE_ALSO	Object(DS_C_NAME)	—	yes	E
DS_A_SERIAL_NBR	String(OM_S_PRINTABLE_STRING)	1–64	yes	E, S
DS_A_STATE_OR_PROV_NAME	String(OM_S_TELETEX_STRING)	1–128	yes	E, S
DS_A_STREET_ADDRESS	String(OM_S_OBJECT_IDENTIFIER_STRING)	1–128	yes	E, S
DS_A_SUPPORT_APPLIC_CONTEXT	String(OM_S_OBJECT_IDENTIFIER_STRING)	—	yes	E
DS_A_SURNAME	String(OM_S_TELETEX_STRING)	1–64	yes	E, S
DS_A_PHONE_NBR	String(OM_S_PRINTABLE_STRING) ⁴	1–32	yes	E, S
DS_A_TELETEX_TERM_IDENT	Object(DS_C_TELETEX_TERM_IDENT)	—	yes	—
DS_A_TELEX_NBR	Object(DS_C_TELEX_NBR)	—	yes	—
DS_A_TITLE	String(OM_S_TELETEX_STRING)	1–64	yes	E, S

Attribute Type	OM Value Syntax	Value Length	Multi-valued	Matching Rules
DS_A_USER_PASSWORD	String(OM_S_OCTET_STRING)	0–128	yes	—
DS_A_X121_ADDRESS	String(OM_S_NUMERIC_STRING) ⁵	1–15	yes	E, S

- 1 As permitted by ISO 3166.
- 2 As permitted by Recommendations F.1 and F.31.
- 3 As permitted by E.164.
- 4 As permitted by E.123 (for example, +44 582 10101).
- 5 As permitted by X.121.

Throughout the descriptions that follow, the term *object* indicates the directory object whose directory entry contains the corresponding directory attributes.

- **DS_A_ALIASED_OBJECT_NAME**

This attribute occurs only in alias entries. It assigns the distinguished name (DN) of the object, provided with an alias, using the entry in which this attribute occurs. An alias is an alternative to an object's DN. Any object can (but need not) have one or more aliases. The directory service is said to dereference an alias whenever it replaces the alias during name processing with the DN associated with it by means of this attribute.

- **DS_A_BUSINESS_CATEGORY**

This attribute provides descriptions of the businesses in which the object is engaged.

- **DS_A_COMMON_NAME**

This attribute provides the names by which the object is commonly known in the context defined by its position in the DIT. The names can conform to the naming convention of the country or culture with which the object is associated. They can be ambiguous.

- **DS_A_COUNTRY_NAME**

This attribute identifies the country in which the object is located or with which it is associated in some other important way. The matching rules require that differences in the case of alphabetical characters be considered insignificant. It has a length of two characters and its values are those listed in ISO 3166.

- **DS_A_DESCRIPTION**

This attribute gives informative descriptions of the object.

- **DS_A_DEST_INDICATOR**

This attribute provides the country-city pairs by means of which the object can be reached via the public telegram service. The matching rules require that differences in the case of alphabetical characters be considered insignificant.

- **DS_A_FACSIMILE_PHONE_NBR**

This attribute provides the telephone numbers for facsimile terminals (and their parameters, if required) by means of which the object can be reached or with which it is associated in some other important way.

- **DS_A_INTERNAT_ISDN_NBR**

This attribute provides the international ISDN numbers by means of which the object can be reached or with which it is associated in some other important way. The matching rules require that differences caused by the presence of spaces be considered insignificant.

- **DS_A_KNOWLEDGE_INFO**

This attribute occurs only in entries that describe a DSA. It provides a human-intelligible accumulated description of the directory knowledge possessed by the DSA.

- **DS_A_LOCALITY_NAME**

This attribute identifies geographical areas or localities. When used as part of a directory name, it specifies the localities in which the object is located or with which it is associated in some other important way.

- **DS_A_MEMBER**

This attribute gives the names of objects that are considered members of the present object; for example, a distribution list for electronic mail.

- **DS_A_OBJECT_CLASS**

This attribute identifies the object classes to which the object belongs, and it also identifies their superclasses. All such object classes that have object identifiers assigned to them are present, except that object class **DS_O_TOP** need not (but can) be present provided that some other value is present. This attribute must be present in every entry and cannot be modified. For a further discussion, see Section 12.2.

- **DS_A_ORG_NAME**

This attribute identifies organizations. When used as part of a directory name, it specifies an organization with which the object is affiliated. Several values can identify the same organization in different ways.

- **DS_A_ORG_UNIT_NAME**

This attribute identifies organizational units. When used as part of a directory name, it specifies an organizational unit with which the object is affiliated. The units are understood to be parts of the organization that the **DS_A_ORG_NAME** attribute indicates. Several values can identify the same unit in different ways.

- **DS_A_OWNER**

This attribute gives the names of objects that have responsibility for the object.

- **DS_A_PHYS_DELIV_OFF_NAME**

This attribute gives the names of cities, towns, villages, and so on, that contain physical delivery offices through which the object can take delivery of physical mail.

- **DS_A_POST_OFFICE_BOX**

This attribute identifies post office boxes at which the object can take delivery of physical mail. This information is also displayed as part of the **DS_A_POSTAL_ADDRESS** attribute, if it is present.

- **DS_A_POSTAL_ADDRESS**

This attribute gives the postal addresses at which the object can take delivery of physical mail. The matching rules require that differences in the case of alphabetical characters be considered insignificant.

- **DS_A_POSTAL_CODE**

This attribute gives the postal codes that are assigned to areas or buildings through which the object can take delivery of physical mail. This information is also displayed as part of the **DS_A_POSTAL_ADDRESS** attribute, if it is present.

- **DS_A_PREF_DELIV_METHOD**

This attribute gives the object's preferred methods of communication, in the order of preference. The values are as follows:

- **DS_ANY_DELIV_METHOD**, meaning that the object has no preference.
- **DS_G3_FACSIMILE_DELIV**, meaning via the Group 3 facsimile.
- **DS_G4_FACSIMILE_DELIV**, meaning via the Group 4 facsimile.
- **DS_IA5_TERMINAL_DELIV**, meaning via the IA5 text.
- **DS_MHS_DELIV**, meaning via X.400.
- **DS_PHYS_DELIV**, meaning via the postal or other physical delivery system.
- **DS_PHONE_DELIV**, meaning via telephone.
- **DS_TELETEX_DELIV**, meaning via teletex.
- **DS_TELEX_DELIV**, meaning via telex.
- **DS_VIDEOTEX_DELIV**, meaning via videotex.

- **DS_A_PRESENTATION_ADDRESS**

This attribute contains the OSI presentation address of the object, which is an OSI application entity. The matching rule for a presented value to match a value stored in the directory is that the P-Selector, S-Selector, and T-Selector of the two presentation addresses must be equal, and the N-Addresses of the presented value must be a subset of those of the stored value.

- **DS_A_REGISTERED_ADDRESS**

This attribute contains mnemonics by means of which the object can be reached via the public telegram service, according to Recommendation F.1. A mnemonic

identifies an object in the context of a particular city, and it is registered in the country containing the city. The matching rules require that differences in the case of alphabetical characters be considered insignificant.

- **DS_A_ROLE_OCCUPANT**

This attribute occurs only in entries that describe an organizational role. It gives the names of objects that fulfill the organizational role.

- **DS_A_SEARCH_GUIDE**

This attribute contains the criteria that can be used to build filters for conducting searches in which the object is the base object.

- **DS_A_SEE_ALSO**

This attribute contains the names of objects that represent other aspects of the real-world object that the present object represents.

- **DS_A_SERIAL_NBR**

This attribute contains the serial numbers of a device.

- **DS_A_STATE_OR_PROV_NAME**

This attribute specifies a state or province. When used as part of a directory name, it identifies states, provinces, or other geographical regions in which the object is located or with which it is associated in some other important way.

- **DS_A_STREET_ADDRESS**

This attribute identifies a site for the local distribution and physical delivery of mail. When used as part of a directory name, it identifies the street address (for example, street name and house number) at which the object is located or with which it is associated in some other important way.

- **DS_A_SUPPORT_APPLIC_CONTEXT**

This attribute occurs only in entries that describe an OSI application entity. It identifies OSI application contexts supported by the object.

- **DS_A_SURNAME**

This attribute occurs only in entries that describe individuals. The surname by which the individual is commonly known, normally inherited from the individual's parent (or parents) or taken at marriage, as determined by the custom of the country or culture with which the individual is associated.

- **DS_A_PHONE_NBR**

This attribute identifies telephones by means of which the object can be reached or with which it is associated in some other important way. The matching rules require that differences caused by the presence of spaces and dashes be considered insignificant.

- **DS_A_TELETEX_TERM_IDENT**

This attribute contains descriptions of teletex terminals by means of which the object can be reached or with which it is associated in some other important way.

- **DS_A_TELEX_NBR**

This attribute contains descriptions of telex terminals by means of which the object can be reached or with which it is associated in some other important way.

- **DS_A_TITLE**

This attribute identifies positions or functions of the object within its organization.

- **DS_A_USER_PASSWORD**

This attribute contains the passwords assigned to the object.

- **DS_A_X121_ADDRESS**

This attribute identifies points on the public data network at which the object can be reached or with which it is associated in some other important way. The matching rules require that differences caused by the presence of spaces be considered insignificant.

12.2 Selected Object Classes

This section presents the object classes that are defined in the standards. Object classes are groups of directory entries that share certain characteristics. The object

classes are arranged into a lattice, based on the object class **DS_O_TOP**. In a lattice, each element, except a leaf, has one or more immediate subordinates but also has one or more immediate superiors. This contrasts with a tree, where each element has exactly one immediate superior. Object classes closer to **DS_O_TOP** are called *superclasses*, and those further away are called *subclasses*. This relationship is not connected to any other such relationship in this guide.

Each directory entry belongs to an object class, and to all the superclasses of that object class. Each entry has an attribute named **DS_A_OBJECT_CLASS**, which was discussed in the previous section, and which identifies the object classes to which the entry belongs. The values of this attribute are object identifiers, which are represented in the interface by constants with the same name as the object class, prefixed by **DS_O_**.

Associated with each object class are zero or more mandatory and zero or more optional attributes. Each directory entry must contain all the mandatory attributes and can (but need not) contain the optional attributes associated with the object class and its superclasses.

The object classes defined in the standards are shown in Table 12-3, together with their object identifiers.

Note: The third and fourth columns of Table 12-3 contain the contents octets of the BER encoding of the object identifier. All these object identifiers stem from the root **{joint-iso-ccitt(2) ds(5) objectClass(6)}**.

Table 12–3. Object Identifiers for Selected Object Classes

Package	Attribute Type	Object Identifier BER	
		Decimal	Hexadecimal
BDCP	DS_O_ALIAS	85, 6, 1	\x55\x06\x01
BDCP	DS_O_APPLIC_ENTITY	85, 6, 12	\x55\x06\x0C
BDCP	DS_O_APPLIC_PROCESS	85, 6, 11	\x55\x06\x0B
BDCP	DS_O_COUNTRY	85, 6, 2	\x55\x06\x02
BDCP	DS_O_DEVICE	85, 6, 14	\x55\x06\x0E
BDCP	DS_O_DSA	85, 6, 13	\x55\x06\x0D

BDCP	DS_O_GROUP_OF_NAMES	85, 6, 9	\x55\x06\x09
BDCP	DS_O_LOCALITY	85, 6, 3	\x55\x06\x03
BDCP	DS_O_ORG	85, 6, 4	\x55\x06\x04
BDCP	DS_O_ORG_PERSON	85, 6, 7	\x55\x06\x07
BDCP	DS_O_ORG_ROLE	85, 6, 8	\x55\x06\x08
BDCP	DS_O_ORG_UNIT	85, 6, 5	\x55\x06\x05
BDCP	DS_O_PERSON	85, 6, 6	\x55\x06\x06
BDCP	DS_O_RESIDENTIAL_PERSON	85, 6, 10	\x55\x06\x0A
BDCP	DS_O_TOP	85, 6, 0	\x55\x06\x00

12.3 OM Class Hierarchy

The remainder of this chapter defines the additional OM classes used to represent values of the selected attributes described in Section 12.1. Some of the selected attributes are represented by OM classes that are used in the interface itself, and hence are defined in Chapter 11; for example, *DS_C_NAME*.

This section shows the hierarchical organization of the OM classes that are defined in the following sections, and it shows which OM classes inherit additional OM attributes from their OM superclasses. In the following list, subclassification is indicated by indentation, and the names of abstract OM classes are in italics. For example, **DS_C_POSTAL_ADDRESS** is an immediate subclass of the abstract OM class *OM_C_OBJECT*.

OM_C_OBJECT

- **DS_C_FACSIMILE_PHONE_NBR**
- **DS_C_POSTAL_ADDRESS**
- **DS_C_SEARCH_CRITERION**
- **DS_C_SEARCH_GUIDE**
- **DS_C_TELETEX_TERM_IDENT**

- **DS_C_TELEX_NBR**

None of the OM classes in the preceding list are encodable by using **om_encode()** and **om_decode()**.

12.4 DS_C_FACSIMILE_PHONE_NBR

An instance of OM class **DS_C_FACSIMILE_PHONE_NBR** identifies and describes a facsimile terminal, if required.

An instance of this OM class has the OM attributes of its superclass, *OM_C_OBJECT*, in addition to the OM attributes listed in Table 12-4.

Table 12-4. OM Attributes of DS_C_FACSIMILE_PHONE_NBR

OM Attribute	Value Syntax	Value Length	Value Number	Value Initially
DS_PARAMETERS	Object (MH_C_G3_FAX_NBPS) ¹	—	0 or 1	—
DS_PHONE_NBR	String(OM_S_PRINTABLE_STRING) ²	1–32	1	—

1 As defined in the X.400 API specifications.

2 As permitted by E.123 (for example, +44 582 10101).

- **DS_PARAMETERS**

If present, this attribute identifies the facsimile terminal's nonbasic capabilities.

- **DS_PHONE_NBR**

This attribute contains a telephone number by means of which the facsimile terminal is accessed.

12.5 DS_C_POSTAL_ADDRESS

An instance of OM class **DS_C_POSTAL_ADDRESS** is a postal address.

An instance of this OM class has the OM attributes of its superclass, *OM_C_OBJECT*, in addition to the OM attribute listed in Table 12-5.

Table 12-5. OM Attribute of DS_C_POSTAL_ADDRESS

OM Attribute	Value Syntax	Value Length	Value Number	Value Initially
DS_POSTAL_ADDRESS	String(OM_S_TELETEX_STRING)	1-30	1-6	—

- **DS_POSTAL_ADDRESS**

Each value of this OM attribute is one line of the postal address. It typically includes a name, street address, city name, state or province name, postal code, and possibly a country name.

12.6 DS_C_SEARCH_CRITERION

An instance of OM class **DS_C_SEARCH_CRITERION** is a component of a **DS_C_SEARCH_GUIDE** OM object.

An instance of this OM class has the OM attributes of its superclass, *OM_C_OBJECT*, in addition to the OM attributes listed in Table 12-6.

Table 12–6. OM Attributes of DS_C_SEARCH_CRITERION

OM Attribute	Value Syntax	Value Length	Value Number	Value Initially
DS_ATTRIBUTE_TYPE	String(OM_S_OBJECT_IDENTIFIER_STRING)	—	0 or 1	—
DS_CRITERIA	Object(DS_C_SEARCH_CRITERION)	—	0 or more	—
DS_FILTER_ITEM_TYPE	Enum(DS_Filter_Item_Type)	—	0 or 1	—
DS_FILTER_TYPE	Enum(DS_Filter_Type)	—	1	DS_ITEM

A **DS_C_SEARCH_CRITERION** suggests how to build part of a filter to be used when searching the directory. Its meaning depends on the value of its OM attribute **DS_FILTER_TYPE**. If the value is **DS_ITEM**, then the criteria suggests building an instance of OM class **DS_C_FILTER_ITEM**. If **DS_FILTER_TYPE** has any other value, it suggests building an instance of OM class **DS_C_FILTER**.

- **DS_ATTRIBUTE_TYPE**

This attribute indicates the attribute type to be used in the suggested **DS_C_FILTER_ITEM**. This OM attribute is only present when the value of **DS_FILTER_TYPE** is **DS_ITEM**.

- **DS_CRITERIA**

This attribute contains nested search criteria. This OM attribute is not present when the value of **DS_FILTER_TYPE** is **DS_ITEM**.

- **DS_FILTER_ITEM_TYPE**

This attribute indicates the type of suggested filter item. Its value can be one of the following:

- **DS_APPROXIMATE_MATCH**
- **DS_EQUALITY**
- **DS_GREATER_OR_EQUAL**
- **DS_LESS_OR_EQUAL**

— **DS_SUBSTRINGS**

However, the filter item cannot have the value **DS_PRESENT**. This OM attribute is only present when the value of **DS_FILTER_TYPE** is **DS_ITEM**.

• **DS_FILTER_TYPE**

This attribute indicates the type of suggested filter. The value **DS_ITEM** means that the suggested component is a filter item, not a filter. The other values suggest the corresponding type of filter. Its value is one of the following:

— **DS_AND**

— **DS_ITEM**

— **DS_NOT**

— **DS_OR**

12.7 DS_C_SEARCH_GUIDE

An instance of OM class **DS_C_SEARCH_GUIDE** suggests a criteria for searching the directory for particular entries. It can be used to build a **DS_C_FILTER** parameter for **ds_search()** operations that are based on the object in whose entry the search guide occurs.

An instance of this OM class has the OM attributes of its superclass, *OM_C_OBJECT*, in addition to the OM attributes listed in Table 12-7.

Table 12-7. OM Attributes of DS_C_SEARCH_GUIDE

OM Attribute	Value Syntax	Value Length	Value Number	Value Initially
DS_OBJECT_CLASS	String(OM_S_OBJECT_IDENTIFIER_STRING)	—	0 or 1	—
DS_CRITERIA	Object(DS_C_SEARCH_CRITERION)	—	1	—

• **DS_OBJECT_CLASS**

This attribute identifies the object class of the entries to which the search guide applies. If this OM attribute is absent, the search guide applies to objects of any class.

- **DS_CRITERIA**

This attribute contains the suggested search criteria.

12.8 DS_C_TELETEX_TERM_IDENT

An instance of OM class **DS_C_TELETEX_TERM_IDENT** identifies and describes a teletex terminal.

An instance of this OM class has the OM attributes of its superclass, *OM_C_OBJECT*, in addition to the OM attributes listed in Table 12-8.

Table 12-8. OM Attributes of DS_C_TELETEX_TERM_IDENT

OM Attribute	Value Syntax	Value Length	Value Number	Value Initially
DS_PARAMETERS	Object(MH_C_TELETEX_NBPS) ¹	—	0 or 1	—
DS_TELETEX_TERM	String(OM_S_PRINTABLE_STRING) ²	1–1024	1	—

1 As defined in the X.400 API specifications.

2 As permitted by F.200.

- **DS_PARAMETERS**

This attribute identifies the teletex terminal's nonbasic capabilities.

- **DS_TELETEX_TERM**

This attribute identifies the teletex terminal.

12.9 DS_C_TELEX_NBR

An instance of OM class **DS_C_TELEX_NBR** identifies and describes a telex terminal.

An instance of this OM class has the OM attributes of its superclass, *OM_C_OBJECT*, in addition to the OM attributes listed in Table 12-9.

Table 12-9. OM Attributes of DS_C_TELEX_NBR

OM Attribute	Value Syntax	Value Length	Value Number	Value Initially
DS_ANSWERBACK	String(OM_S_PRINTABLE_STRING)	1-8	1	—
DS_COUNTRY_CODE	String(OM_S_PRINTABLE_STRING)	1-4	1	—
DS_TELEX_NBR	String(OM_S_PRINTABLE_STRING)	1-14	1	—

- **DS_ANSWERBACK**

This attribute contains the code with which the telex terminal acknowledges calls placed to it.

- **DS_COUNTRY_CODE**

This attribute contains the identifier of the country through which the telex terminal is accessed.

- **DS_TELEX_NBR**

This attribute contains the number by means of which the telex terminal is addressed.

Chapter 13

Strong Authentication Package

²This chapter describes the strong authentication package (SAP). In addition to the attribute types, attribute syntaxes, and object classes defined in the basic directory contents package, the standards also contain definitions to support authentication mechanisms.

They include such objects as **Strong-Authentication-User**.

This chapter outlines names for each of these items, and it defines OM classes to represent those that are not represented directly by OM syntaxes. The values of attributes in the directory are not restricted to those discussed in this chapter, and new attribute types and syntaxes can be created at any time. (For further information on how the values of other syntaxes are represented in the interface, see Chapter 10.)

The constants and OM classes in this chapter are defined in addition to those in Chapter 11, since they are not essential to the working of the interface, but instead

-
2. These definitions are chiefly in *The Directory: Selected Attribute Types* (ISO 9594-6, CCITT X.520) and *The Directory: Selected Object Classes* (ISO 9594-7, CCITT X.521) with additional material in *The Directory: Overview of Concepts, Models, and Services* (ISO 9594-1, CCITT X.500) and *The Directory: Authentication Framework* (ISO 9594-8, CCITT X.509).

allow directory entries to be utilized. The definitions belong to the SAP, which is supported by the DCE XDS API following negotiation of its use with **ds_version()**.

The object identifier associated with the SAP is

```
{iso(1) identified-organization(3) icd-ecma(0012) member-company(2)
dec(1011) xopen(28) sap(2)}
```

It takes the following encoding:

```
\x2B\xC\x2\x87\x73\x1C\x2
```

This identifier is represented by the constant **DS_STRONG_AUTHENT_PKG**. The C constants associated with this package are in the **xdssap.h** header file.

The concepts and notation used are introduced in Chapter 11. They are also fully explained in Chapters 17 through 19.

The selected attribute types are presented first, followed by the selected object classes. Next, the OM class hierarchy and OM class definitions required to support the selected attribute types are presented.

13.1 SAP Attribute Types

This section presents the additional attribute types defined in the standards that are to be used with the SAP. Each attribute type has an object identifier, which is the value of the OM attribute **DS_ATTRIBUTE_TYPE**. These object identifiers are represented in the interface by constants with the same name as the directory attribute, and they are prefixed with **DS_A_** so that they can be easily identified.

This section contains two tables that are used to indicate the object identifiers for SAP attribute types (see Table 13-1), and the values for SAP attribute types (see Table 13-2), respectively. Following these two tables is a brief description of each attribute. (See Chapter 12 for information on general matching rules).

Note: The third and fourth columns of Table 13-1 contain the contents octets of the BER encoding of the object identifier. All these object identifiers stem from the root **{joint-iso-ccitt(2) ds(5) attributeType(4)}**.

Table 13–1. Object Identifiers for SAP Attribute Types

Package	Attribute Type	Object Identifier BER	
		Decimal	Hexadecimal
SAP	DS_A_AUTHORITY_REVOC_LIST	85, 4, 38	\x55\x04\x26
SAP	DS_A_CA_CERT	85, 4, 37	\x55\x04\x25
SAP	DS_A_CERT_REVOC_LIST	85, 4, 39	\x55\x04\x27
SAP	DS_A_CROSS_CERT_PAIR	85, 4, 40	\x55\x04\x28
SAP	DS_A_USER_CERT	85, 4, 36	\x55\x04\x24

Table 13–2. Representation of Values for SAP Attribute Types

Attribute Type	OM Value Syntax	Value Length	Multi-valued	Matching Rules
DS_A_AUTHORITY_REVOC_LIST	Object(DS_C_CERT_LIST)	—	yes	
DS_A_CA_CERT	Object(DS_C_CERT)	—	yes	
DS_A_CERT_REVOC_LIST	Object(DS_C_CERT_LIST)	—	yes	
DS_A_CROSS_CERT_PAIR	Object(DS_C_CERT_PAIR)	—	yes	
DS_A_USER_CERT	Object(DS_C_CERT)	—	yes	

Throughout the descriptions that follow, the term *object* indicates the directory object whose directory entry contains the corresponding directory attributes.

- **DS_A_AUTHORITY_REVOC_LIST**

This attribute occurs only in entries that describe a certification authority (CA). It lists all the certificates issued to any of the CAs known to this CA, and later revoked. Each value of this OM attribute is signed by the CA.

- **DS_A_CA_CERT**

This attribute specifies the certificates assigned to the object, which is a CA.

- **DS_A_CERT_REVOC_LIST**

This attribute occurs only in entries that describe a CA. It lists the certificates issued by this CA and later revoked. Each value of this OM attribute is signed by the CA.

- **DS_A_CROSS_CERT_PAIR**

This attribute specifies one or two certificates held in the entry of a CA. The first certificate is that of one CA, guaranteed by a second CA; whereas, the second certificate is that of the second CA, guaranteed by the first CA.

- **DS_A_USER_CERT**

This attribute specifies the user certificates assigned to the object, which may be any user certificate, including a CA certificate.

13.2 SAP Object Classes

This section presents the SAP object classes that are defined in the standards. (See Table 13-3).

Note: The third and fourth columns of Table 13-3 contain the contents octets of the BER encoding of the object identifier. All these object identifiers stem from the root **{joint-iso-ccitt(2) ds(5) objectClass(6)}**.

Table 13–3. Object Identifiers for SAP Object Classes

Package	Attribute Type	Object Identifier BER	
		Decimal	Hexadecimal
SAP	DS_O_CERT_AUTHORITY	85, 6, 16	\x55\x06\x10
SAP	DS_O_STRONG_AUTHENT_USER	85, 6, 15	\x55\x06\x0F

13.3 OM Class Hierarchy

The remainder of this chapter defines the additional OM classes used by SAP. This section shows the hierarchical organization of the OM classes that are defined in the following sections, and it shows which OM classes inherit additional OM attributes from their OM superclasses. In the following list, subclassification is indicated by indentation, and the names of abstract OM classes are in italics.

OM_C_OBJECT

- **DS_C_ALGORITHM_IDENT**
- **DS_C_CERT_PAIR**
- *DS_C_SIGNATURE*
 - **DS_C_CERT**
 - **DS_C_CERT_LIST**
 - **DS_C_CERT_SUBLIST**

None of the OM classes in the preceding list are encodable by using **om_encode()** and **om_decode()**.

13.4 DS_C_ALGORITHM_IDENT

An instance of OM class **DS_C_ALGORITHM_IDENT** records the encryption algorithm that an object uses to digitally sign messages, together with the parameters of the algorithm.

An instance of this OM class has the OM attributes of its superclass, *OM_C_OBJECT*, in addition to the OM attributes listed in Table 13-4.

Table 13–4. OM Attributes of DS_C_ALGORITHM_IDENT

OM Attribute	Value Syntax	Value Length	Value Number	Value Initially
DS_ALGORITHM	String(OM_S_OBJECT_IDENTIFIER_STRING)	—	1	—
DS_ALGORITHM_PARAMETERS	any	—	0 or 1	—

- **DS_ALGORITHM**

This attribute specifies an object identifier that uniquely identifies the algorithm used by some object.

- **DS_ALGORITHM_PARAMETERS**

This attribute specifies the values of the algorithm's parameters that are used by the object. The syntax of the parameters is determined by each individual algorithm.

13.5 DS_C_CERT

An instance of OM class **DS_C_CERT** comprises a user's DN, public key, and additional information, all of which is digitally signed by the issuing CA in order to make the certificate unforgeable. The OM attributes associated with *DS_C_SIGNATURE* (a superclass of **DS_C_CERT**) are present.

An instance of this OM class has the OM attributes of its superclasses, *OM_C_OBJECT* and *DS_C_SIGNATURE*, in addition to the OM attributes listed in Table 13-5.

Table 13–5. OM Attributes of DS_C_CERT

OM Attribute	Value Syntax	Value Length	Value Number	Value Initially
DS_SERIAL_NUMBER	OM_S_INTEGER	—	1	—
DS_SUBJECT	Object(<i>DS_C_NAME</i>)	—	1	—
DS_SUBJECT_ALGORITHM	Object(DS_C_ALGORITHM_IDENT)	—	1	—
DS_SUBJECT_PUBLIC_KEY	String(OM_S_BIT_STRING)	—	1	—
DS_VALIDITY_NOT_AFTER	String(OM_S_UTC_TIME_STRING)	0–17	1	—
DS_VALIDITY_NOT_BEFORE	String(OM_S_UTC_TIME_STRING)	0–17	1	—
DS_VERSION	Enum(DS_Version)	—	1	DS_V1988

- **DS_SERIAL_NUMBER**

This attribute distinguishes the certificate from all other certificates that were ever or will be issued by the CA that issued this certificate.

- **DS_SUBJECT**

This attribute specifies the subject's name.

- **DS_SUBJECT_ALGORITHM**

This attribute specifies the algorithm that is used by the subject for encryption, and which is associated with the public key.

- **DS_SUBJECT_PUBLIC_KEY**

This attribute specifies the subject's public key associated with the algorithm.

- **DS_VALIDITY_NOT_AFTER**

This attribute specifies the last day on which the certificate is valid.

- **DS_VALIDITY_NOT_BEFORE**

This attribute specifies the first day on which the certificate is valid.

- **DS_VERSION**

This attribute identifies the certificate's design. Its value is **DS_V1988**, meaning the design that was specified in the 1988 version of the standards.

13.6 DS_C_CERT_LIST

An instance of OM class **DS_C_CERT_LIST** documents the revocation of zero or more certificates. The documentation is provided by the object, which is a CA whose signature is affixed to the instance.

An instance of this OM class has the OM attributes of its superclasses, *OM_C_OBJECT* and *DS_C_SIGNATURE*, in addition to the OM attributes listed in Table 13-6.

Table 13-6. OM Attributes of DS_C_CERT_LIST

OM Attribute	Value Syntax	Value Length	Value Number	Value Initially
DS_LAST_UPDATE	String(OM_S.UTC_TIME_STRING)	0-17	1	—
DS_REVOKED_CERTS	Object(DS_C_CERT_SUBLIST)	—	0 or more	—

- **DS_LAST_UPDATE**

This attribute indicates the time at which the revocation list was updated to its current state.

- **DS_REVOKED_CERTS**

This attribute identifies the revoked certificates.

13.7 DS_C_CERT_PAIR

An instance of OM class **DS_C_CERT_PAIR** contains one or both of a forward and reverse certificate that assists users in building a certification path.

An instance of this OM class has the OM attributes of its superclass, *OM_C_OBJECT*, in addition to the OM attributes listed in Table 13-7.

Table 13-7. OM Attributes of DS_C_CERT_PAIR

OM Attribute	Value Syntax	Value Length	Value Number	Value Initially
DS_FORWARD	Object(DS_C_CERT)	—	0 or 1 ¹	—
DS_REVERSE	Object(DS_C_CERT)	—	0 or 1 ¹	—

¹ At least one of these OM attributes must be present.

- **DS_FORWARD**

This attribute specifies the certificate of a first CA issued by a second CA.

- **DS_REVERSE**

This attribute specifies the certificate of the second CA issued by the first CA.

13.8 DS_C_CERT_SUBLIST

An instance of OM class **DS_C_CERT_SUBLIST** documents the revocation of zero or more certificates issued by the CA whose signature is affixed to the instance.

An instance of this OM class has the OM attributes of its superclasses, *OM_C_OBJECT* and *DS_C_SIGNATURE*, in addition to the OM attributes listed in Table 13-8.

Table 13–8. OM Attributes of DS_C_CERT_SUBLIST

OM Attribute	Value Syntax	Value Length	Value Number	Value Initially
DS_REVOCATION_DATE	String(OM_S_UTC_TIME_STRING)	0–17	0 or more ¹	—
DS_SERIAL_NUMBERS	OM_S_INTEGER	—	0 or more ¹	—

1 The values of these two OM attributes parallel one another and are equal in number.

- **DS_REVOCATION_DATE**

This attribute specifies the epoch at which each of the certificates was revoked. The serial numbers of the certificates are the corresponding values of the OM attribute **DS_SUBJECT**.

- **DS_SERIAL_NUMBERS**

This attribute specifies the serial numbers assigned to the revoked certificates.

13.9 DS_C_SIGNATURE

An instance of the abstract OM class *DS_C_SIGNATURE* contains the algorithm identifier used to produce a digital signature and the name of the object that produced it. The scope of the signature is any instance of any subclass of this OM class.

An instance of this OM class has the OM attributes of its superclass, *OM_C_OBJECT*, in addition to the OM attributes listed in Table 13-9.

Table 13–9. OM Attributes of DS_C_SIGNATURE

OM Attribute	Value Syntax	Value Length	Value Number	Value Initially
DS_ISSUER	Object(<i>DS_C_NAME</i>)	—	1	—
DS_SIGNATURE	Object(DS_C_ALGORITHM_IDENT)	—	1	—
DS_SIGNATURE_VALUE	String(OM_S_OCTET_STRING)	—	1	—

- **DS_ISSUER**

This attribute indicates the name of the object that produced the digital signature.

- **DS_SIGNATURE**

This attribute identifies the algorithm that was used to produce the digital signature, and it identifies any parameters of the algorithm.

- **DS_SIGNATURE_VALUE**

This attribute provides an enciphered summary of the information to which the signature is appended. The summary is produced by means of a one-way hash function, while the enciphering is carried out by using the secret key of the signer.

Chapter 14

MHS Directory User Package

The message handling system (MHS) directory user package (MDUP) contains definitions to support the use of the directory in accordance with the standard 1988 X.400 User Agents and MTAs for name resolution, distribution list (DL) expansion, and capability assessment. The definitions are based on the attribute types and syntaxes specified in *X.402, Annex A*.

The MDUP is an optional package that can be used by the XDS interface. Applications must negotiate use of this package with **ds_version()** before using any of the MDUP features. If an application attempts to use features specific to the package without first negotiating its use, an appropriate error (for example, **OM_NO_SUCH_CLASS**) is returned by the object management (OM) function.

The object identifier associated with the MDUP is

```
{iso(1) identified-organization(3) icd-ecma(0012) member-company(2)
dec(1011) xopen(28) mdup(3)}
```

It takes the following encoding:

\x2B\xC\x2\x87\x73\x1C\x3

This identifier is represented by the constant **DS_MHS_DIR_USER_PKG**. The C constants associated with this package are defined in the **xdsmdup.h**, **xmhp.h**, and **xmsga.h** header files.

The concepts and notation used are first mentioned in Chapter 11. They are also fully explained in Chapters 17 through 19. The attribute types are introduced first, followed by the object classes. Next, the OM class hierarchy and OM class definitions required to support the new attribute types are described.

14.1 MDUP Attribute Types

This section presents additional directory attribute types that are used with the MDUP. Each attribute type has an object identifier, which is the value of the OM attribute **DS_ATTRIBUTE_TYPE**. These object identifiers are represented in the interface by constants with the same name as the directory attribute and are prefixed by **DS_A_** so that they can be easily identified.

This section contains two tables that are used to indicate the object identifiers for MDUP attribute types (see Table 14-1), and the values for MDUP attribute types (see Table 14-2), respectively. Following these two tables is a brief description of each attribute. (See Chapter 12 for information on general matching rules).

Note: The third and fourth columns of Table 14-1 contain the contents octets of the BER encoding of the object identifiers. All these object identifiers stem from the root **{joint-iso-ccitt(2) mhs-motis(6) arch(5) at(2)}**.

Table 14–1. Object Identifiers for MDUP Attribute Types

Package	Attribute Type	Object Identifier BER	
		Decimal	Hexadecimal
MDUP	DS_A_DELIV_CONTENT_LENGTH	86, 5, 2, 0	\x56\x05\x02\x00
MDUP	DS_A_DELIV_CONTENT_TYPES	86, 5, 2, 1	\x56\x05\x02\x01

Package	Attribute Type	Object Identifier BER	
		Decimal	Hexadecimal
MDUP	DS_A_DELIV_EITS	86, 5, 2, 2	\x56\x05\x02\x02
MDUP	DS_A_DL_MEMBERS	86, 5, 2, 3	\x56\x05\x02\x03
MDUP	DS_A_DL_SUBMIT_PERMS	86, 5, 2, 4	\x56\x05\x02\x04
MDUP	DS_A_MESSAGE_STORE	86, 5, 2, 5	\x56\x05\x02\x05
MDUP	DS_A_OR_ADDRESSES	86, 5, 2, 6	\x56\x05\x02\x06
MDUP	DS_A_PREF_DELIV_METHODS	86, 5, 2, 7	\x56\x05\x02\x07
MDUP	DS_A_SUPP_AUTO_ACTIONS	86, 5, 2, 8	\x56\x05\x02\x08
MDUP	DS_A_SUPP_CONTENT_TYPES	86, 5, 2, 9	\x56\x05\x02\x09
MDUP	DS_A_SUPP_OPT_ATTRIBUTES	86, 5, 2, 10	\x56\x05\x02\x0A

Table 14–2. Representation of Values for MDUP Attribute Types

Attribute Type	OM Value Syntax	Value Length	Multi-valued	Matching Rules
DS_A_DELIV_CONTENT_LENGTH	OM_S_INTEGER	—	no	—
DS_A_DELIV_CONTENT_TYPES	String(OM_S_OBJECT_IDENTIFIER_STRING)	—	yes	—

Attribute Type	OM Value Syntax	Value Length	Multi-valued	Matching Rules
DS_A_DELIV_EITS	String(OM_S_OBJECT_IDENTIFIER_STRING)	—	yes	—
DS_A_DL_MEMBERS	Object(DS_C_OR_NAME)	—	yes	—
DS_A_DL_SUBMIT_PERMS	Object(DS_C_DL_SUBMIT_PERMS)	—	yes	—
DS_A_MESSAGE_STORE	String(DS_C_DS_DN)	—	no	—
DS_A_OR_ADDRESSES	Object(MH_C_OR_ADDRESS)	—	yes	—
DS_A_PREF_DELIV_METHODS	Enum(MH_Delivery_Mode)	—	no	E
DS_A_SUPP_AUTO_ACTIONS	String(OM_S_OBJECT_IDENTIFIER_STRING)	—	yes	—
DS_A_SUPP_CONTENT_TYPES	String(OM_S_OBJECT_IDENTIFIER_STRING)	—	yes	—
DS_A_SUPP_OPT_ATTRIBUTES	String(OM_S_OBJECT_IDENTIFIER_STRING)	—	yes	—

Throughout the descriptions that follow, the term *object* indicates the directory object whose directory entry contains the corresponding directory attributes.

- **DS_A_DELIV_CONTENT_LENGTH**

This attribute identifies the maximum content length of the messages whose delivery a user will accept.

- **DS_A_DELIV_CONTENT_TYPES**

This attribute identifies the content types of the messages whose delivery a user will accept.

- **DS_A_DELIV_EITS**

This attribute identifies the encoded information types (EITs) of the messages whose delivery a user will accept.

- **DS_A_DL_MEMBERS**

This attribute identifies the members of a DL.

- **DS_A_DL_SUBMIT_PERMS**

This attribute identifies the users and DLs that may submit messages to a DL.

- **DS_A_MESSAGE_STORE**

This attribute identifies a user's message store (MS) by name.

- **DS_A_OR_ADDRESSES**

This attribute specifies a user's or DL's originator/recipient (O/R) addresses.

- **DS_A_PREF_DELIV_METHODS**

This attribute identifies, in the order of decreasing preference, the methods of delivery a user prefers.

- **DS_A_SUPP_AUTO_ACTIONS**

This attribute identifies the automatic actions that an MS fully supports.

- **DS_A_SUPP_CONTENT_TYPES**

This attribute identifies the content types of the messages whose syntax and semantics an MS fully supports.

- **DS_A_SUPP_OPT_ATTRIBUTES**

This attribute identifies the optional attributes that an MS fully supports.

14.2 MDUP Object Classes

There are five MDUP object classes and their associated object identifiers (see Table 14-3).

Note: The third and fourth columns of Table 14-3 contain the contents octets of the BER encoding of the object identifier. MDUP object identifiers stem from the root **{joint-iso-ccitt(2) mhs-motis(6) arch(5) oc(1)}**.

Table 14–3. Object Identifiers for MDUP Object Classes

Package	Object Class	Object Identifier BER	
		Decimal	Hexadecimal
MDUP	DS_O_MHS_DISTRIBUTION_LIST	86, 5, 1, 0	\x56\x05\x01\x00
MDUP	DS_O_MHS_MESSAGE_STORE	86, 5, 1, 1	\x56\x05\x01\x01
MDUP	DS_O_MHS_MESSAGE_TRANS_AG	86, 5, 1, 2	\x56\x05\x01\x02
MDUP	DS_O_MHS_USER	86, 5, 1, 3	\x56\x05\x01\x03
MDUP	DS_O_MHS_USER_AG	86, 5, 1, 4	\x56\x05\x01\x04

14.3 MDUP OM Class Hierarchy

The remainder of this chapter defines the additional OM classes used by MDUP. This section shows the hierarchical organization of the OM classes that are defined in the following sections, and shows which classes inherit additional OM attributes from their OM superclasses. In the following list, subclassification is indicated by indentation and the names of abstract OM classes are represented in *italic font*.

OM_C_OBJECT

- *MH_C_OR_ADDRESS*
 - *MH_C_OR_NAME*
- *DS_C_DL_SUBMIT_PERMS*

None of the OM classes in the preceding list are encodable by using **om_encode()** and **om_decode()**.

14.4 MH_C_OR_ADDRESS

An instance of class **MH_C_OR_ADDRESS** distinguishes one user or DL from another, and identifies its point of access to the message transfer system (MTS). Every user or DL is assigned one or more MTS access points and thus one or more originator/recipient (O/R) addresses.

The attributes specific to this class are listed in Table 14-4. The 1988 column indicates that the attribute applies only to the 1988 standard.

Table 14-4. Attributes Specific to MH_C_OR_ADDRESS

Attribute	Value Syntax	Value Length	Value Number	1988?
MH_T_ADMD_NAME ¹	String(OM_S_PRINTABLE_STRING)	0–16	0 or 1	—
MH_T_COMMON_NAME	String(OM_S_PRINTABLE_STRING) or String(OM_S_TELETEX_STRING) ²	1–64	0–2	1988
MH_T_COUNTRY_NAME ¹	String(OM_S_PRINTABLE_STRING)	2–3	0 or 1	—

Attribute	Value Syntax	Value Length	Value Number	1988?
MH_T_DOMAIN_TYPE_2	String(OM_S_PRINTABLE_STRING) or String(OM_S_TELETEX_STRING) ^{2,3}	1–8	0–2 ⁴	—
MH_T_DOMAIN_TYPE_3	String(OM_S_PRINTABLE_STRING) or String(OM_S_TELETEX_STRING) ^{2,3}	1–8	0–2 ⁴	—
MH_T_DOMAIN_TYPE_4	String(OM_S_PRINTABLE_STRING) or String(OM_S_TELETEX_STRING) ^{2,3}	1–8	0–2 ⁴	—
MH_T_DOMAIN_VALUE_1	String(OM_S_PRINTABLE_STRING) or String(OM_S_TELETEX_STRING) ^{2,3}	1–128	0–2 ⁴	—
MH_T_DOMAIN_VALUE_2	String(OM_S_PRINTABLE_STRING) or String(OM_S_TELETEX_STRING) ^{2,3}	1–128	0–2 ⁴	—

Attribute	Value Syntax	Value Length	Value Number	1988?
MH_T_DOMAIN_VALUE_3	String(OM_S_PRINTABLE_STRING) or String(OM_S_TELETEX_STRING) ^{2,3}	1–128	0–2 ⁴	—
MH_T_DOMAIN_VALUE_4	String(OM_S_PRINTABLE_STRING) or String(OM_S_TELETEX_STRING) ^{2,3}	1–128	0–2 ⁴	—
MH_T_GENERATION	String(OM_S_PRINTABLE_STRING) or String(OM_S_TELETEX_STRING) ^{2,3}	1–3	0–2 ⁴	—
MH_T_GIVEN_NAME	String(OM_S_PRINTABLE_STRING) or String(OM_S_TELETEX_STRING) ^{2,3}	1–16	0–2 ⁴	—
MH_T_INITIALS	String(OM_S_PRINTABLE_STRING) or String(OM_S_TELETEX_STRING) ^{2,3}	1–5	0–2 ⁴	—
MH_T_ISDN_NUMBER	String(OM_S_NUMERIC_STRING)	1–15	0 or 1	1988
MH_T_ISDN_SUBADDRESS	String(OM_S_NUMERIC_STRING)	1–40	0 or 1 ⁵	1988

Attribute	Value Syntax	Value Length	Value Number	1988?
MH_T_NUMERIC_USER_IDENTIFIER	String(OM_S_NUMERIC_STRING)	1–32	0 or 1	—
MH_T_ORGANIZATION_NAME	String(OM_S_PRINTABLE_STRING) or String(OM_S_PRINTABLE_STRING)	1–64	0–2 ⁴ , 6	—
MH_T_ORGANIZATIONAL_UNIT_NAME_1	String(OM_S_PRINTABLE_STRING) or String(OM_S_TELETEX_STRING) ^{2,3}	1–32	0–2 ⁴	—
MH_T_ORGANIZATIONAL_UNIT_NAME_2	String(OM_S_PRINTABLE_STRING) or String(OM_S_TELETEX_STRING) ^{2,3}	1–32	0–2 ⁴	—
MH_T_ORGANIZATIONAL_UNIT_NAME_3	String(OM_S_PRINTABLE_STRING) or String(OM_S_TELETEX_STRING) ^{2,3}	1–32	0–2 ⁴	—
MH_T_ORGANIZATIONAL_UNIT_NAME_4	String(OM_S_PRINTABLE_STRING) or String(OM_S_TELETEX_STRING) ^{2,3}	1–32	0–2 ⁴	—

Attribute	Value Syntax	Value Length	Value Number	1988?
MH_T_POSTAL_ADDRESS_DETAILS	String(OM_S_PRINTABLE_STRING) or String(OM_S_TELETEX_STRING) ²	1–30	0–2	1988
MH_T_POSTAL_ADDRESS_IN_FULL	String(OM_S_PRINTABLE_STRING)	1–185	0 or 1	1988
MH_T_POSTAL_ADDRESS_IN_LINES	String(OM_S_PRINTABLE_STRING)	1–30	0–6	1988
MH_T_POSTAL_CODE	String(OM_S_PRINTABLE_STRING)	1–16	0 or 1	1988
MH_T_POSTAL_COUNTRY_NAME	String(OM_S_PRINTABLE_STRING)	2–3	0 or 1	1988
MH_T_POSTAL_DELIVERY_POINT_NAME	String(OM_S_PRINTABLE_STRING) or String(OM_S_TELETEX_STRING) ^{2,3}	1–30	0–2	1988
MH_T_POSTAL_DELIV_SYSTEM_NAME	String(OM_S_PRINTABLE_STRING)	1–16	0 or 1	1988
MH_T_POSTAL_GENERAL_DELIV_ADDR	String(OM_S_PRINTABLE_STRING) or String(OM_S_TELETEX_STRING) ²	1–30	0–2	1988

Attribute	Value Syntax	Value Length	Value Number	1988?
MH_T_POSTAL_LOCALE	String(OM_S_PRINTABLE_STRING) or String(OM_S_TELETEX_STRING) ²	1–30	0–2	1988
MH_T_POSTAL_OFFICE_BOX_NUMBER	String(OM_S_PRINTABLE_STRING) or String(OM_S_TELETEX_STRING) ²	1–30	0–2	1988
MH_T_POSTAL_OFFICE_NAME	String(OM_S_PRINTABLE_STRING) or String(OM_S_TELETEX_STRING) ²	1–30	0–2	1988
MH_T_POSTAL_OFFICE_NUMBER	String(OM_S_PRINTABLE_STRING) or String(OM_S_TELETEX_STRING) ²	1–30	0–2	1988
MH_T_POSTAL_ORGANIZATION_NAME	String(OM_S_PRINTABLE_STRING) or String(OM_S_TELETEX_STRING) ²	1–30	0–2	1988

Attribute	Value Syntax	Value Length	Value Number	1988?
MH_T_POSTAL_PATRON_DETAILS	String(OM_S_PRINTABLE_STRING) or String(OM_S_TELETEX_STRING) ²	1–30	0–2	1988
MH_T_POSTAL_PATRON_NAME	String(OM_S_PRINTABLE_STRING) or String(OM_S_TELETEX_STRING) ²	1–30	0–2	1988
MH_T_POSTAL_STREET_ADDRESS	String(OM_S_PRINTABLE_STRING) or String(OM_S_TELETEX_STRING) ²	1–30	0–2	1988
MH_T_PRESENTATION_ADDRESS	Object(DS_C_PRESENTATION_ADDRESS)	—	0 or 1	1988
MH_T_PRMD_NAME	String(OM_S_PRINTABLE_STRING)	1–16	0 or 1	—
MH_T_SURNAME	String(OM_S_PRINTABLE_STRING) or String(OM_S_TELETEX_STRING) ^{2, 3}	1–40	0–2 ⁴	—
MH_T_TERMINAL_IDENTIFIER	String(OM_S_PRINTABLE_STRING)	1–24	0 or 1	—

Attribute	Value Syntax	Value Length	Value Number	1988?
MH_T_TERMINAL_TYPE	Enum(MH_Terminal_Type)	—	0 or 1	1988
MH_T_X121_ADDRESS	String(OM_S_NUMERIC_STRING)	1–15	0 or 1	—

- 1 The value initially is the current session's attribute of the same name.
- 2 If only one value is present in international communications, its syntax is String(**OM_S_PRINTABLE_STRING**). If two values are present, in either domestic or international communications, the syntax of the first is String(**OM_S_PRINTABLE_STRING**), the syntax of the second is String(**OM_S_TELETEX_STRING**), and the two convey the same information such that either can be safely ignored.
For example, Teletex strings allow inclusion of the accented characters commonly used in many countries. Not all input/output devices, however, permit the entry and display of such characters. Printable strings are required internationally to ensure that such device limitations do not prevent communications.
- 3 For 1984, the syntax of the value is String(**OM_S_PRINTABLE_STRING**).
- 4 For 1984, at most one value is present.
- 5 This attribute is present only if the ISDN number attribute is present.
- 6 For 1988, this attribute is required if any organization name is present.

- **MH_T_ADMD_NAME**

This attribute contains the name of the user's or DL's administration management domain (ADMD). It identifies the ADMD relative to the country that the **MH_T_COUNTRY_NAME** attribute indicates. Its values are defined by that country.

Note that the attribute value that comprises a single space is reserved. If permitted by the country that the **MH_T_COUNTRY_NAME** attribute indicates, a single space designates “any”—that is, all ADMDs within the country. This affects both the identification of users and DLs within the country and the routing of messages, probes, and reports to and among the ADMDs of that country.

Regarding the former, it requires that the O/R addresses of users and DLs within the country be chosen so as to ensure their unambiguousness, even in the absence of the actual names of the users' and DLs' ADMDs. Regarding the latter, it permits private management domains (PRMDs) within the country and ADMDs outside the country to route messages, probes, and reports to any of the ADMDs within the country indiscriminately. It also requires that the ADMDs within the country interconnect themselves in such a way that the messages, probes, and reports are conveyed to their destinations.

- **MH_T_COMMON_NAME**

This attribute contains the name commonly used to refer to the user or DL. It identifies the user or DL relative to the entity indicated by another attribute; for example, **MH_T_ORGANIZATION_NAME**. Its values are defined by that entity.

- **MH_T_COUNTRY_NAME**

This attribute contains the name of the user's or DL's country. Its defined values are the numbers that X.121 assigns to the country, or the character pairs that ISO 3166 assigns to it.

- **MH_T_DOMAIN_TYPE_1**

This attribute contains the name of a class of information. Its values are defined by the user's or DL's ADMD and PRMD, if any, in combination.

- **MH_T_DOMAIN_TYPE_2**

This attribute contains the name of a class of information. Its values are defined by the user's or DL's ADMD and PRMD, if any, in combination.

- **MH_T_DOMAIN_TYPE_3**

This attribute contains the name of a class of information. Its values are defined by the user's or DL's ADMD and PRMD, if any, in combination.

- **MH_T_DOMAIN_TYPE_4**

This attribute contains the name of a class of information. Its values are defined by the user's or DL's ADMD and PRMD, if any, in combination.

- **MH_T_DOMAIN_VALUE_1**

This attribute is an instance of the class of information that the **MH_T_DOMAIN_TYPE_1** attribute indicates.

- **MH_T_DOMAIN_VALUE_2**

This attribute is an instance of the class of information that the **MH_T_DOMAIN_TYPE_2** attribute indicates.

- **MH_T_DOMAIN_VALUE_3**

This attribute is an instance of the class of information that the **MH_T_DOMAIN_TYPE_3** attribute indicates.

- **MH_T_DOMAIN_VALUE_4**

This attribute is an instance of the class of information that the **MH_T_DOMAIN_TYPE_4** attribute indicates.

- **MH_T_GENERATION**

This attribute contains the user's generation; for example, **Jnr**.

- **MH_T_GIVEN_NAME**

This attribute contains the user's given name; for example, **Robert**.

- **MH_T_INITIALS**

This attribute contains the initials of all of the user's names except the user's surname; for example, **RE**.

- **MH_T_ISDN_NUMBER**

This attribute contains the ISDN number of the user's terminal. Its values are defined by E.163 and E.164.

- **MH_T_ISDN_SUBADDRESS**

This attribute contains the ISDN subaddress, if any, of the user's terminal. Its values are defined by E.163 and E.164.

- **MH_T_NUMERIC_USER_IDENTIFIER**

This attribute numerically identifies the user or DL relative to the ADMD that the **MH_T_ADMD_NAME** attribute indicates. Its values are defined by that ADMD.

- **MH_T_ORGANIZATION_NAME**

This attribute contains the name of the user's or DL's organization. As a national matter, such names may be assigned by the country that the **MH_T_COUNTRY_NAME** attribute indicates, the ADMD that the **MH_T_ADMD_NAME** attribute indicates, the PRMD that the **MH_T_PRMD_NAME** attribute indicates, or the latter two organizations together.

- **MH_T_ORGANIZATIONAL_UNIT_NAME_1**

This attribute contains the name of a unit (for example, a division or department) of the organization that the **MH_T_ORGANIZATION_NAME** attribute indicates. The attribute's values are defined by that organization.

- **MH_T_ORGANIZATIONAL_UNIT_NAME_2**

This attribute contains the name of a subunit (for example, a division or department) of the unit that the **MH_T_ORGANIZATIONAL_UNIT_NAME_1** attribute indicates. The attribute's values are defined by the latter unit.

- **MH_T_ORGANIZATIONAL_UNIT_NAME_3**

This attribute contains the name of a subunit (for example, a division or department) of the unit that the **DS_A_ORGANIZATIONAL_UNIT_NAME_2** attribute indicates. The attribute's values are defined by the latter unit.

- **MH_T_ORGANIZATIONAL_UNIT_NAME_4**

This attribute contains the name of a subunit (for example, a division or department) of the unit that the **MH_T_ORGANIZATIONAL_UNIT_NAME_3** attribute indicates. The attribute's values are defined by the latter unit.

- **MH_T_POSTAL_ADDRESS_DETAILS**

This attribute contains the means (for example, a room and the floor numbers in a large building) for identifying the exact point at which the user takes delivery of physical messages.

- **MH_T_POSTAL_ADDRESS_IN_FULL**

This attribute contains the free-form and possibly multiline postal address of the user as a single Teletex string with the lines being separated as prescribed for Teletex strings.

- **MH_T_POSTAL_ADDRESS_IN_LINES**

This attribute contains the free-form postal address of the user in a sequence of printable strings, each representing a line of text.

- **MH_T_POSTAL_CODE**

This attribute contains the postal code for the geographical area in which the user takes delivery of physical messages. It identifies the area relative to the country that the **MH_T_POSTAL_COUNTRY_NAME** attribute indicates. Its values are defined by the postal administration of that country.

- **MH_T_POSTAL_COUNTRY_NAME**

This attribute contains the name of the country in which the user takes delivery of physical messages. Its defined values are the numbers that X.121 assigns to the country, or the character pairs that ISO 3166 assigns to it.

- **MH_T_POSTAL_DELIVERY_POINT_NAME**

This attribute identifies the locus of distribution other than that indicated by the **MH_T_POSTAL_OFFICE_NAME** attribute (for example, a geographical area) of the user's physical messages.

- **MH_T_POSTAL_DELIV_SYSTEM_NAME**

This attribute contains the name of the postal delivery system (PDS) through which the user is to receive physical messages. It identifies the PDS relative to the ADMD that the **MH_T_ADMD_NAME** attribute indicates. Its values are defined by that ADMD.

- **MH_T_POSTAL_GENERAL_DELIV_ADDRESS**

This attribute contains the code that the user gives to the post office to collect the physical messages awaiting delivery to the user. The post office is indicated in the **MH_T_POSTAL_OFFICE_NAME** attribute. The values for

the **MH_T_POSTAL_GENERAL_DELIV_ADDRESS** attribute are defined by that post office.

- **MH_T_POSTAL_LOCALE**

This attribute identifies the point of delivery other than that indicated by the following attributes:

- **MH_T_POSTAL_GENERAL_DELIV_ADDR**
- **MH_T_POSTAL_OFFICE_BOX_NUMBER**
- **MH_T_POSTAL_STREET_ADDRESS**.

For example, a building or a hamlet of the user's physical messages.

- **MH_T_POSTAL_OFFICE_BOX_NUMBER**

This attribute contains the number of the post office box by means of which the user takes delivery of physical messages. The box is located at the post office that the **MH_T_POSTAL_OFFICE_NAME** attribute indicates. The attribute's values are defined by that post office.

- **MH_T_POSTAL_OFFICE_NAME**

This attribute contains the name of the municipality (for example, city or village) where the post office is situated, through which the user takes delivery of physical messages.

- **MH_T_POSTAL_OFFICE_NUMBER**

This attribute contains the means of distinguishing among several post offices indicated by the **MH_T_POSTAL_OFFICE_NAME** attribute.

- **MH_T_POSTAL_ORGANIZATION_NAME**

This attribute contains the name of the organization through which the user takes delivery of physical messages.

- **MH_T_POSTAL_PATRON_DETAILS**

This attribute contains additional information (for example, the name of the organizational unit through which the user takes delivery of physical messages) necessary to identify the user for purposes of physical delivery.

- **MH_T_POSTAL_PATRON_NAME**

This attribute contains the name under which the user takes delivery of physical messages.

- **MH_T_POSTAL_STREET_ADDRESS**

This attribute contains the street address (for example, **43 Primrose Lane**) at which the user takes delivery of physical messages.

- **MH_T_PRESENTATION_ADDRESS**

This attribute contains the presentation address of the user's terminal.

- **MH_T_PRMD_NAME**

This attribute contains the name of the user's PRMD. As a national matter, such names may be assigned by the country that the **MH_T_COUNTRY_NAME** attribute indicates or the ADMD that the **MH_T_ADMD_NAME** attribute indicates.

- **MH_T_SURNAME**

This attribute contains the user's surname; for example, **Lee**.

- **MH_T_TERMINAL_IDENTIFIER**

This attribute contains the terminal identifier of the user's terminal; for example, a Telex answer back or a Teletex terminal identifier.

- **MH_T_TERMINAL_TYPE**

This attribute contains the type of the user's terminal. Its value is selected from among the following:

- **MH_TT_G3_FAX**
- **MH_TT_G4_FAX**
- **MH_TT_IA5_TERMINAL**
- **MH_TT_TELETEX**
- **MH_TT_TELEX**
- **MH_TT_VIDEOTEX**

The meaning of each value is indicated by its name.

- **MH_T_X121_ADDRESS**

This attribute contains the network address of the user's terminal. Its values are defined by X.121.

Note: The strings admitted by X.121 include a telephone number preceded by the telephone escape digit (9), and a Telex number preceded by the Telex escape digit (8).

Certain attributes are grouped together for reference as follows:

- **Personal Name Attributes**

These comprise the following:

- **MH_T_GIVEN_NAME**
- **MH_T_INITIALS**
- **MH_T_SURNAME**
- **MH_T_GENERATION**

- **Organizational Unit Name Attributes**

These comprise the following:

- **MH_T_ORGANIZATIONAL_UNIT_NAME_1**
- **MH_T_ORGANIZATIONAL_UNIT_NAME_2**
- **MH_T_ORGANIZATIONAL_UNIT_NAME_3**
- **MH_T_ORGANIZATIONAL_UNIT_NAME_4**

- **Network Address Attributes**

These comprise the following:

- **MH_T_ISDN_NUMBER**
- **MH_T_ISDN_SUBADDRESS**
- **MH_T_PRESENTATION_ADDRESS**
- **MH_T_X121_ADDRESS**

For any i in the interval $[1, 4]$, the domain type i and domain value i attributes constitute a domain-defined attribute (DDA).

Note: The widespread avoidance of DDAs produces more uniform and thus more user-friendly O/R addresses. However, it is anticipated that not all management domains (MDs) will be able to avoid such attributes immediately. The purpose of DDAs is to permit an MD to retain its existing native addressing conventions for a time. It is intended, however, that all MDs migrate away from the use of DDAs, and thus that DDAs are used only for an interim period.

An O/R address may take any of the forms summarized in Table 14-5. Table 14-5 indicates the attributes that may be present in an O/R address of each form. It also indicates whether it is mandatory (M) or conditional (C) that they do so. When applied to a group of attributes (the network address attributes, for example), mandatory means that at least one member of the group must be present, while conditional means that no members of the group need necessarily be present.

The presence or absence in a particular O/R address of conditional attributes is determined as follows. If a user or DL is accessed through a PRMD, the ADMD that the **MH_T_COUNTRY_NAME** and **MH_T_ADMD_NAME** attributes indicate governs whether attributes used to route messages to the PRMD are present, but it imposes no other constraints on attributes. If a user or DL is *not* accessed through a PRMD, the same ADMD governs whether all conditional attributes, except those specific to postal O/R addresses, are present. All conditional attributes specific to postal O/R addresses are present or absent so as to satisfy the postal addressing requirements of the users they identify.

Table 14–5. Forms of Originator/Recipient Address

Attribute	Mnem ¹	Num ²	Spost ³	Upost ⁴	Term ⁵
MH_T_ADMD_NAME	M	M	M	M	C
MH_T_COMMON_NAME	C	—	—	—	—
MH_T_COUNTRY_NAME	M	M	M	M	C
Domain-Defined Attributes	C	C	—	—	C
Network Address Attributes	—	—	—	—	M

Attribute	Mnem ¹	Num ²	Spost ³	Upost ⁴	Term ⁵
MH_T_NUMERIC_USER_IDENTIFIER	—	M	—	—	—
MH_T_ORGANIZATION_NAME	C	—	—	—	—
Organizational Unit Name Attributes	C	—	—	—	—
Personal Name Attributes	C	—	—	—	—
MH_T_POSTAL_ADDRESS_DETAILS	—	—	C	—	—
MH_T_POSTAL_ADDRESS_IN_FULL	—	—	—	M	—
MH_T_POSTAL_CODE	—	—	M	M	—
MH_T_POSTAL_COUNTRY_NAME	—	—	M	M	—
MH_T_POSTAL_DELIVERY_POINT_NAME	—	—	C	—	—
MH_T_POSTAL_DELIV_SYSTEM_NAME	—	—	C	C	—
MH_T_POSTAL_GENERAL_DELIV_ADDR	—	—	C	—	—
MH_T_POSTAL_LOCALE	—	—	C	—	—
MH_T_POSTAL_OFFICE_BOX_NUMBER	—	—	C	—	—
MH_T_POSTAL_OFFICE_NAME	—	—	C	—	—
MH_T_POSTAL_OFFICE_NUMBER	—	—	C	—	—
MH_T_POSTAL_ORGANIZATION_NAME	—	—	C	—	—

Attribute	Mnem ¹	Num ²	Spost ³	Upost ⁴	Term ⁵
MH_T_POSTAL_PATRON_DETAILS	—	—	C	—	—
MH_T_POSTAL_PATRON_NAME	—	—	C	—	—
MH_T_POSTAL_STREET_ADDRESS	—	—	C	—	—
MH_T_PRMD_NAME	C	C ⁶	C	C	C ⁶
MH_T_TERMINAL_IDENTIFIER	—	—	—	—	C
MH_T_TERMINAL_TYPE	—	—	—	—	C

- 1 Mnemonic. X.400 (1984) calls this Form 1 Variant 1.
- 2 Numeric. X.400 (1984) calls this Form 1 Variant 2.
- 3 Structured postal. For 1984 this O/R address form is undefined.
- 4 Unstructured postal. For 1984 this O/R address form is undefined.
- 5 X.400 (1984) calls this Form 1 Variant 3 and Form 2.
- 6 For 1984 this attribute is absent (—). For 1988 it is conditional (C).

- **Mnemonic O/R Address**

This address mnemonically identifies a user or DL. Using the **MH_T_ADMD_NAME** and **MH_T_COUNTRY_NAME** attributes, it identifies an ADMD. Using the **MH_T_COMMON_NAME** attribute or the personal name attributes, the **MH_T_ORGANIZATION_NAME** attribute, the **Organizational-Unit-Name** attributes, the **MH_T_PRMD_NAME** attribute, or a combination of these, and optionally DDAs, it identifies a user or DL relative to the ADMD.

The personal name attributes identify a user or DL relative to the entity indicated by another attribute; for example, **MH_T_ORGANIZATION_NAME**. The **MH_T_SURNAME** attribute will be present if any of the other three personal name attributes are present.

- **Numeric O/R Address**

This address numerically identifies a user or DL. Using the **MH_T_ADMD_NAME** and **MH_T_COUNTRY_NAME** attributes, it identifies an ADMD. Using the **MH_T_NUMERIC_USER_IDENTIFIER** attribute and possibly the **MH_T_PRMD_NAME** attribute, it identifies the user or DL relative to the ADMD. Any DDAs provide information that is additional to that required to identify the user or DL.

- **Postal O/R Address**

This address identifies a user through its postal address. Two kinds of postal O/R address are distinguished, as follows:

- Structured

Said of a postal O/R address that specifies a user's postal address by means of several attributes. The structure of the postal address is described in the following text in some detail.

- Unstructured

Said of a postal O/R address that specifies a user's postal address in a single attribute. The structure of the postal address is left largely unspecified in the following text.

Whether structured or unstructured, a postal O/R address does the following. Using the **MH_T_ADMD_NAME** and **MH_T_COUNTRY_NAME** attributes, it identifies an ADMD. Using the **MH_T_POSTAL_CODE** and **MH_T_POSTAL_COUNTRY_NAME** attributes, it identifies the geographical region in which the user takes delivery of physical messages. Using the **MH_T_POSTAL_DELIV_SYSTEM_NAME** or **MH_T_PRMD_NAME** attribute or both, it also can identify the PDS by means of which the user is to be accessed.

An unstructured postal O/R address also includes the **MH_T_POSTAL_ADDRESS_IN_FULL** attribute. A structured postal O/R address also includes every other postal addressing attribute that the PDS requires to identify the postal patron.

Note: The total number of characters in the values of all attributes, except for **MH_T_ADMD_NAME**, **MH_T_COUNTRY_NAME**, and **MH_T_POSTAL_DELIV_SYSTEM_NAME**, in a postal O/R address should be small enough to permit their rendition in 6 lines of 30 characters, the size of a typical physical envelope window. The rendition algorithm, while defined by the physical delivery access unit (PDAU), is likely to include inserting delimiters (for example, spaces) between some attribute values.

- **Terminal O/R Address**

This address identifies a user by identifying the user's terminal using the network address attributes. It also may identify the ADMD through which the terminal is accessed by using the **MH_T_ADMD_NAME** and **MH_T_COUNTRY_NAME** attributes. The **MH_T_PRMD_NAME** attribute and any DDAs, which will be present only if the **MH_T_ADMD_NAME** and **MH_T_COUNTRY_NAME** attributes are present, provide information additional to that required to identify the user.

If the terminal is a Telematic terminal, it gives the terminal's network address and possibly, using the **MH_T_TERMINAL_TYPE** and **MH_T_TERMINAL_IDENTIFIER** attributes, its terminal type and identifier. If the terminal is a Telex terminal, it gives the terminal's Telex number.

Whenever two O/R addresses are compared for equality, the following differences are ignored:

- Whether an attribute has a value whose syntax is String(**OM_S_PRINTABLE_STRING**), a value whose syntax is String(**OM_S_TELETEX_STRING**), or both.
- Whether a letter in a value of an attribute not used in DDAs is an uppercase or lowercase letter.
- All leading, all trailing, and all but one consecutive embedded space in an attribute value.

Note: An MD may impose additional equivalence rules upon the O/R addresses it assigns to its own users and DLs. It may define, for example, rules concerning punctuation characters in attribute values, the case of letters in attribute values, or the relative order of DDAs.

As a national matter, MDs may impose additional rules regarding any attribute that may have a value whose syntax is String(**OM_S_PRINTABLE_STRING**), a value whose syntax is String(**OM_S_TELETEX_STRING**), or both. In particular, the rules for deriving from a Teletex string the equivalent printable string may be nationally prescribed.

14.5 MH_C_OR_NAME

An instance of class **MH_C_OR_NAME** comprises a directory name, an O/R address, or both. The name is considered present if, and only if, the **MH_T_DIRECTORY_NAME** attribute is present. The address comprises the attributes specific to the **MH_C_OR_ADDRESS** class and is considered present if, and only if, at least one of those attributes is present.

An O/R name's composition is context sensitive. At submission, the name, the address, or both may be present. At transfer, or delivery, the address is present and the name can (but need not) be present. Whether at submission, transfer or delivery, the MTS uses the name, if it is present, only if the address is absent or invalid.

The attribute specific to this class is listed in Table 14-6.

Table 14-6. Attribute Specific to MH_C_OR_NAME

Attribute	Value Syntax	Value Length	Value Number	Value Initially	1988?
MH_T_DIRECTORY_NAME	Object(<i>DS_C_NAME</i>)	—	0 or 1	—	1988

- **MH_T_DIRECTORY_NAME**

This attribute contains the name assigned to the user or DL by the worldwide X.500 directory.

14.6 DS_C_DL_SUBMIT_PERMS

An instance of OM class **DS_C_DL_SUBMIT_PERMS** characterizes an attribute each of whose values are a submit permission. An instance of this OM class has the OM attributes of its superclass, *OM_C_OBJECT*, and additionally the OM attributes listed in Table 14-7.

Table 14-7. OM Attributes of DS_C_DL_SUBMIT_PERMS

OM Attribute	Value Syntax	Value Length	Value Number	Value Initially
DS_PERM_TYPE	Enum(DS_Permission_Type)	—	1	—
DS_INDIVIDUAL	Object(MH_C_OR_NAME)	—	0 or 1	—
DS_MEMBER_OF_DL	Object(MH_C_OR_NAME)	—	0 or 1	—
DS_PATTERN_MATCH	Object(MH_C_OR_NAME)	—	0 or 1	—
DS_MEMBER_OF_GROUP	Object(DS_C_DS_DN)	—	0 or more	—

- **DS_PERM_TYPE**

This attribute contains the type of the permission specified herein. Its value can be one of the following:

- **DS_PERM_INDIVIDUAL**
- **DS_PERM_MEMBER_OF_DL**
- **DS_PERM_PATTERN_MATCH**
- **DS_PERM_MEMBER_OF_GROUP**

- **DS_INDIVIDUAL**

This attribute contains the user or unexpanded DL, any of whose O/R names is equal to the specified O/R name.

- **DS_MEMBER_OF_DL**

This attribute contains each member of the DL, any of whose O/R names is equal to the specified O/R name, or of each nested DL, recursively.

- **DS_PATTERN_MATCH**

This attribute contains each user or unexpanded DL, any of whose O/R names matches the specified O/R name pattern.

- **DS_MEMBER_OF_GROUP**

This attribute contains each member of the group-of-names whose name is specified, or of each nested group-of-names, recursively.

Note that exactly one of the four name attributes will be present at any time, according to the value of the **DS_PERM_TYPE** attribute.

Chapter 15

GDS Package

The Global Directory Service (GDS) package (GDSP) is an OSF extension to the XDS interface. Applications must negotiate use of this package with **ds_version()** before using any of the additional features. If an application attempts to use features specific to this package without first negotiating its use, then an appropriate error (for example, **OM_NO_SUCH_CLASS**) is returned by the Object Management function.

The object identifier associated with the GDSP is

```
{iso(1) identified-organization(3) icd-ecma(0012) member-company(2)
siemens-units(1107) sni(1) directory(3) xdsapi(100) gdsp(0)}
```

It takes the following encoding:

```
\x2B\xC\x2\x88\x53\x1\x3\x64\x0
```

The identifier is represented by the constant **DSX_GDS_PKG**. The C constants associated with this package are contained in the **xdsgds.h** header file.

The concepts and notation used are first mentioned in Chapter 11. They are also fully explained in Chapters 17 through 19. The attribute types are introduced first, followed by the object classes. Next, the OM class hierarchy and OM class definitions required to support the new attribute types are described.

15.1 GDSP Attribute Types

This section presents the additional directory attribute types that are used with GDSP. Each attribute type has an object identifier, which is the value of the OM attribute **DS_ATTRIBUTE_TYPE**. These object identifiers are represented in the interface by constants with the same name as the directory attribute, and they are prefixed by **DSX_A_** so that they can be easily identified.

This section contains two tables that are used to indicate the object identifiers for GDSP attribute types (see Table 15-1), and the values for GDSP attribute types (see Table 15-2), respectively. Following these two tables is a brief description of each attribute. (See Chapter 12 for information on general matching rules.)

Table 15-1 shows the names of the GDSP attribute types, together with the BER encoding of the object identifiers associated with each of them.

Note: The third column of Table 15-1 contains the contents octets of the BER encoding of the object identifier in hexadecimal. All these object identifiers stem from the root **{iso(1) identified-organization(3) idc-ecma(0012) member-company(2) siemens-units(1107) sni(1) directory(3) attribute-type(4)}**.

Table 15–1. Object Identifiers for GDSP Attribute Types

Package	Attribute Type	Object Identifier BER
		Hexadecimal
GDSP	DSX_A_ACL	\x2B\x0C\x02\x88\x53\x01\x03\x04\x01
GDSP	DSX_A_AT	\x2B\x0C\x02\x88\x53\x01\x03\x04\x06
GDSP	DSX_A_CDS_CELL	\x2B\x0C\x02\x88\x53\x01\x03\x04\x0D

Package	Attribute Type	Object Identifier BER
		Hexadecimal
GDSP	DSX_A_CDS_REPLICA	\x2B\x0C\x02\x88\x53\x01\x03\x04\x0E
GDSP	DSX_A_CLIENT	\x2B\x0C\x02\x88\x53\x01\x03\x04\x0A
GDSP	DSX_A_DEFAULT_DSA	\x2B\x0C\x02\x88\x53\x01\x03\x04\x08
GDSP	DSX_A_DNLIST	\x2B\x0C\x02\x88\x53\x01\x03\x04\x0B
GDSP	DSX_A_LOCAL_DSA	\x2B\x0C\x02\x88\x53\x01\x03\x04\x09
GDSP	DSX_A_MASTER_KNOWLEDGE	\x2B\x0C\x02\x88\x53\x01\x03\x04\x00
GDSP	DSX_A_OCT	\x2B\x0C\x02\x88\x53\x01\x03\x04\x05
GDSP	DSX_A_SHADOWED_BY	\x2B\x0C\x02\x88\x53\x01\x03\x04\x03
GDSP	DSX_A_SHADOWING_JOB	\x2B\x0C\x02\x88\x53\x01\x03\x04\x0C
GDSP	DSX_A_SRT	\x2B\x0C\x02\x88\x53\x01\x03\x04\x04
GDSP	DSX_A_TIME_STAMP	\x2B\x0C\x02\x88\x53\x01\x03\x04\x02

Table 15-2 shows the names of the attribute types, together with the OM value syntax used in the interface to represent values of that attribute type. The table also includes the range of lengths permitted for the string types, indicates whether the attribute can be multivalued, and lists which matching rules are provided for the syntax.

Table 15–2. Representation of Values for GDSP Attribute Types

Attribute Type	OM Value Syntax	Value Length	Multivalued	Matching Rules
DSX_A_ACL	Object(DSX_C_GDS_ACL)	—	no	E
DSX_A_AT	String(OM_S_PRINTABLE_STRING)	1–101	yes	E,S
DSX_A_CDS_CELL	String(OM_S_OCTET_STRING)	1–284	no	E
DSX_A_CDS_REPLICA	String(OM_S_OCTET_STRING)	1–905	yes	E
DSX_A_CLIENT	Only a cache attribute	—	—	—
DSX_A_DEFAULT_DSA	Only a cache attribute	—	—	—
DSX_A_DNLIST	Object(DS_C_DS_DN)	—	yes	E,S
DSX_A_LOCAL_DSA	Only a cache attribute	—	—	—
DSX_A_MASTER_KNOWLEDGE	Object(DS_C_DS_DN)	—	no	E,S
DSX_A_OCT	String(OM_S_PRINTABLE_STRING)	1–397	yes	E,S
DSX_A_SHADOWED_BY	Not used yet	—	—	—
DSX_A_SHADOWING_JOB	Not used yet	—	—	—

Attribute Type	OM Value Syntax	Value Length	Multivalued	Matching Rules
DSX_A_SRT	String(OM_S_PRINTABLE_STRING)	1–29	yes	E,S
DSX_A_TIME_STAMP	String(OM_S_UTC_TIME_STRING)	11–17	no	E,O

Note: With the exception of the **DSX_A_ACL** attribute, the GDSP attributes in Table 15-2 are only to be manipulated through the GDS administration interface (see the *OSF DCE GDS Administration Guide and Reference*.)

Descriptions of the GDSP attributes follow:

- **DSX_A_ACL**

This attribute describes the access rights for one or more directory service users.

- **DSX_A_AT**

This attribute describes the attribute types permitted in GDS. For further information, see the *OSF DCE GDS Administration Guide and Reference*.

- **DSX_A_CDS_CELL** and **DSX_A_CDS_REPLICA**

These two attributes always exist together in the same object. They describe the information necessary for contacting a remote DCE cell.

- **DSX_A_CLIENT**

This attribute only applies to the cache. It identifies an entry that holds the DUA's presentation address. Its OM syntax is **OM_S_PRINTABLE_STRING** and its value is **CLIENT**.

- **DSX_A_DEFAULT_DSA**

This attribute only applies to the cache. It identifies an entry that holds the DN of the DUA's default DSA. Its OM syntax is **OM_S_PRINTABLE_STRING** and its value is **DEFAULT-DSA**.

- **DSX_A_DNLIST**

This attribute is used internally by the GDS DSA.

- **DSX_A_LOCAL_DSA**

This attribute only applies to the cache. It identifies an entry that holds the DN of the DUA's local DSA. Its OM syntax is **OM_S_PRINTABLE_STRING** and its value is **LOCAL-DSA**.

- **DSX_A_MASTER_KNOWLEDGE**

This attribute contains the DN of the DSA that holds the master copy of this entry.

- **DSX_A_OCT**

This attribute describes the object classes supported by the GDS DSA. For further information, see the *OSF DCE GDS Administration Guide and Reference*.

- **DSX_A_SHADOWED_BY** and **DSX_A_SHADOWING_JOB**

These two GDSP attributes are intended for future use.

- **DSX_A_SRT**

This attribute describes the structure of the DNs permitted in GDS.

- **DSX_A_TIME_STAMP**

This attribute is part of the **DSX_O_SCHEMA** object. It contains the creation time of the **DSX_O_SCHEMA** object.

15.2 GDSP Object Classes

The only additional GDSP object class is **DSX_O_SCHEMA** (see Table 15-3). It is stored in GDS as an object directly under root. The most important attributes of the **DSX_O_SCHEMA** object are the three recurring attributes **DSX_A_OCT**, **DSX_A_AT**, and **DSX_A_SRT**. These three objects describe the GDS DIT structure. For a more detailed explanation of the GDSP **DSX_O_SCHEMA** object, see the *OSF DCE GDS Administration Guide and Reference*.

Note: The third column of Table 15-3 contains the contents octets of the BER encoding of the object identifier in hexadecimal. This object identifier stems from the root **{iso(1) identified-organization(3) idc-ecma(0012) member-company(2) siemens-units(1107) sni(1) directory(3) object-class(6)}**.

Table 15–3. Object Identifier for GDSP Object Classes

Package	Attribute Type	Object Identifier BER
		Hexadecimal
GDSP	DSX_O_SCHEMA	\x2B\x0C\x02\x88\x53\x01\x03\x06\x00

15.3 GDSP OM Class Hierarchy

The remainder of this chapter defines the additional OM classes used by GDSP. This section shows the hierarchical organization of the OM classes that are defined in the following sections, and it shows which classes inherit additional OM attributes from their OM superclasses. In the following list, subclassification is indicated by indentation, and the names of abstract OM classes are represented in *italics*.

OM_C_OBJECT (defined in the OM package)

- **DS_C_SESSION** (defined in the directory service package)
 - **DSX_C_GDS_SESSION**
- **DS_C_CONTEXT** (defined in the directory service package)
 - **DSX_C_GDS_CONTEXT**
- **DSX_C_GDS_ACL**
- **DSX_C_GDS_ACL_ITEM**

None of the OM classes in the preceding list are encodable by using **om_encode()** and **om_decode()**.

15.4 DSX_C_GDS_ACL

An instance of OM class **DSX_C_GDS_ACL** describes up to five categories of rights for one or more directory users.

An instance of this OM class has the OM attributes of its superclass, *OM_C_OBJECT*, in addition to the OM attributes listed in Table 15-4.

Table 15-4. OM Attributes of DSX_C_GDS_ACL

OM Attribute	Value Syntax	Value Length	Value Number	Value Initially
DSX_MODIFY_PUBLIC	Object(DSX_C_GDS_ACL_ITEM)	—	0-4	—
DSX_READ_STANDARD	Object(DSX_C_GDS_ACL_ITEM)	—	0-4	—
DSX_MODIFY_STANDARD	Object(DSX_C_GDS_ACL_ITEM)	—	0-4	—
DSX_READ_SENSITIVE	Object(DSX_C_GDS_ACL_ITEM)	—	0-4	—
DSX_MODIFY_SENSITIVE	Object(DSX_C_GDS_ACL_ITEM)	—	0-4	—

The OM attributes of **DSX_C_GDS_ACL** are as follows:

- **DSX_MODIFY_PUBLIC**

This attribute specifies the user, or subtree of users, that can modify attributes classified as public attributes.

- **DSX_READ_STANDARD**

This attribute specifies the user, or subtree of users, that can read attributes classified as standard attributes.

- **DSX_MODIFY_STANDARD**

This attribute specifies the user, or subtree of users, that can modify attributes classified as standard attributes.

- **DSX_READ_SENSITIVE**

This attribute specifies the user, or subtree of users, that can read attributes classified as sensitive attributes.

- **DSX_MODIFY_SENSITIVE**

This attribute specifies the user, or subtree of users, that can modify attributes classified as sensitive attributes.

15.5 DSX_C_GDS_ACL_ITEM

An instance of OM class **DSX_C_GDS_ACL_ITEM** is a component of an instance of OM class **DSX_C_GDS_ACL**. It specifies the user, or subtree of users, to whom an access right applies.

An instance of this OM class has the OM attributes of its superclass, *OM_C_OBJECT*, in addition to the OM attributes listed in Table 15-5.

Table 15-5. OM Attributes of DSX_C_GDS_ACL_ITEM

OM Attribute	Value Syntax	Value Length	Value Number	Value Initially
DSX_INTERPRETATION	Enum(DSX_ Interpretation)	—	1	—
DSX_USER	Object(DS_C_DS_DN)	—	1	—

The OM attributes of a **DSX_C_GDS_ACL_ITEM** are as follows:

- **DSX_INTERPRETATION**

This attribute specifies the scope of the access right. It can have one of the following values:

- **DSX_SINGLE_OBJECT**, meaning that the access right is granted to the user specified in the **DSX_USER** OM attribute.
- **DSX_ROOT_OF_SUBTREE**, meaning that the access right is granted to all users in the subtree below the name specified in the **DSX_USER** OM attribute.

- **DSX_USER**

This attribute is the DN of the user, or subtree of users, to whom an access right applies.

15.6 DSX_C_GDS_CONTEXT

An instance of OM class **DSX_C_GDS_CONTEXT** comprises per-operation arguments that are accepted by most of the interface functions. GDS supports additional service controls that are defined by the **DSX_C_GDS_CONTEXT** OM class.

An instance of this OM class has the OM attributes of its superclasses, *OM_C_OBJECT* and **DS_C_CONTEXT**, in addition to the OM attributes listed in Table 15-6.

Table 15-6. OM Attributes of DSX_C_GDS_CONTEXT

OM Attribute	Value Syntax	Value Length	Value Number	Value Initially
Service Controls				
DSX_DUAFIRST	OM_S_BOOLEAN	—	1	OM_FALSE
DSX_DONT_STORE	OM_S_BOOLEAN	—	1	OM_TRUE
DSX_NORMAL_CLASS	OM_S_BOOLEAN	—	1	OM_FALSE
DSX_PRIV_CLASS	OM_S_BOOLEAN	—	1	OM_FALSE
DSX_RESIDENT_CLASS	OM_S_BOOLEAN	—	1	OM_FALSE

OM Attribute	Value Syntax	Value Length	Value Number	Value Initially
DSX_USEDSA	OM_S_BOOLEAN	—	1	OM_TRUE
DSX_DUA_CACHE	OM_S_BOOLEAN	—	1	OM_FALSE
DSX_PREFER_ADM_FUNCS	OM_S_BOOLEAN	—	1	OM_FALSE
DSX_SIGN_MECHANISM	Enum(DSX_Sign_Mechanism)	—	0–1	—
DSX_PROT_REQUEST	Enum(DSX_Prot_Request)	—	0–1	—

The OM attributes of the **DSX_C_GDS_CONTEXT** OM class are as follows:

- **DSX_DUAFIRST**

This attribute defines whether the DUA cache or the DSA needs to be read first for query operations. The default value is **OM_FALSE**; that is, search the DSA first, if not found then search the DUA cache.

- **DSX_DONT_STORE**

This attribute specifies whether the information read from the DSAs by the query functions also needs to be stored in the DUA cache. When this service control is set to **OM_TRUE** (default value), nothing is stored in the DUA cache.

When this service control is set to **OM_FALSE**, the information read is stored in the DUA cache. The objects returned by **ds_list()** and **ds_compare()** are stored in the cache without their associated attribute information. The objects returned by **ds_read()** and **ds_search()** are stored in the cache with all their *cacheable* attributes; these are all public attributes that do not exceed 4 Kilobytes in length.

This information is only cached when a list of requested attributes is supplied. If all attributes are requested, then nothing is stored in the cache.

The DUA cache categorizes the information stored into three different memory classes. The user specifies the category with the following service controls:

- **DSX_NORMAL_CLASS**

If this attribute is set to **OM_TRUE**, the entry in the DUA cache is assigned to the class of normal objects. If the number of entries in this class exceeds a maximum value, the entry that is not addressed for the longest period of time is removed from the DUA cache.

- **DSX_PRIV_CLASS**

If this attribute is set to **OM_TRUE**, the entry in the DUA cache is assigned to the class of privileged objects. Entries can be removed from the class in the same way as normal objects. By using this memory sparingly, the user can protect entries from deletion.

- **DSX_RESIDENT_CLASS**

If this attribute is set to **OM_TRUE**, the entry in the DUA cache is assigned to the class of resident objects. An entry in this memory class is never removed automatically; instead, it can only be removed with **ds_remove_entry()**. The number of entries is limited; if this limit is exceeded, **ds_add_entry()** reports an error.

Only the service control of one memory class can be set. The **ds_add_entry()** function also evaluates these service control bits if the function is used on the DUA cache.

- **DSX_DUA_CACHE** and **DSX_USEDSA**

These attributes define whether the entries in the DUA cache or in the DSA, or both, need to be used when providing the service. Depending on the values of these attributes, the following situations can arise:

— **DSX_DUA_CACHE** and **DSX_USEDSA**, both **OM_TRUE**

The **ds_add_entry()** and **ds_remove_entry()** functions report an error.

The query functions evaluate the service controls **DS_DONT_USE_COPY** and **DSX_DUAFIRST**. When **DS_DONT_USE_COPY** is **OM_FALSE**, then **DSX_DUAFIRST** determines whether the DUA cache or the DSA is read first. When **DS_DONT_USE_COPY** is **OM_TRUE**, information from the DSA only is read.

— **DSX_DUA_CACHE**, **OM_TRUE** and **DSX_USEDSA**, **OM_FALSE**

The **ds_add_entry()** and **ds_remove_entry()** functions and the query functions only go to the DUA cache.

- **DSX_DUA_CACHE**, **OM_FALSE** and **DSX_USEDSA**, **OM_TRUE**

The **ds_add_entry()** and **ds_remove_entry()** functions and the query functions only go to the DSA.

- **DSX_DUA_CACHE** and **DSX_USEDSA**, both **OM_FALSE**

The **ds_add_entry()** and **ds_remove_entry()** functions and the query functions report an error.

All other functions always operate on the DSA currently connected.

- **DSX_PREFER_ADM_FUNCS**

GDS uses the three following optional attributes:

- **DSX_A_MASTER_KNOWLEDGE**, which contains the DN of the DSA that holds the master copy of an entry.
- **DSX_A_ACL**, which is used for GDS access control.
- **DS_A_USER_PASSWORD** as an attribute of the **DS_O_DSA** object class, which is used by the GDS shadowing mechanism.

The **DSX_A_MASTER_KNOWLEDGE** and **DSX_A_ACL** attributes are present in every GDS entry.

When an application requests all attributes, it can prevent any of these three optional attributes from being returned by setting this service control to **OM_FALSE**.

If GDS applications (for example, GDS administration) require these attributes, they are obtained by setting this service control to **OM_TRUE**.

- **DSX_SIGN_MECHANISM**

This attribute is reserved for future use.

- **DSX_PROT_REQUEST**

This attribute is reserved for future use.

Applications can assume that an object of OM class **DSX_C_GDS_CONTEXT**, created with default values of all its OM attributes, works with all the interface functions. The constant **DS_DEFAULT_CONTEXT** can be used as an argument to functions instead of creating an OM object with default values.

The default **DSX_C_GDS_CONTEXT** is defined in Table 15-7.

Table 15-7. Default DSX_C_GDS_CONTEXT

OM Attribute	Default Value
Common Arguments	
DS_OPERATION_PROGRESS	DS_OPERATION_NOT_STARTED
DS_ALIASED_RDNS	0
Service Controls	
DS_CHAINING_PROHIB	OM_TRUE
DS_DONT_DEREFERENCE_ALIASES	OM_FALSE
DS_DONT_USE_COPY	OM_TRUE
DS_LOCAL_SCOPE	OM_FALSE
DS_PREFER_CHAINING	OM_FALSE
DS_PRIORITY	DS_MEDIUM
Local Controls	
DS_ASYNCHRONOUS	OM_FALSE
DS_AUTOMATIC_CONTINUATION	OM_TRUE
Private Extensions	
DSX_DUAFIRST	OM_FALSE
DSX_DONT_STORE	OM_TRUE
DSX_NORMAL_CLASS	OM_FALSE
DSX_PRIV_CLASS	OM_FALSE
DSX_RESIDENT_CLASS	OM_FALSE

OM Attribute	Default Value
DSX_USEDSA	OM_TRUE
DSX_DUA_CACHE	OM_FALSE
DSX_PREFER_ADM_FUNCS	OM_FALSE
DSX_SIGN_MECHANISM	Absent
DSX_PROT_REQUEST	Absent

15.7 DSX_C_GDS_SESSION

An instance of OM class **DSX_C_GDS_SESSION** identifies a particular link from an application program to a GDS DUA. This additional OM class is necessary if the user either wants to specify an authentication mechanism or wants to specify the GDS directory identifier, or alternatively wants to specify both an authentication mechanism and the directory identifier. **DSX_C_GDS_SESSION** can be passed as an argument to **ds_bind()**.

An instance of this OM class has the OM attributes of its superclasses, *OM_C_OBJECT* and **DS_C_SESSION**, in addition to the OM attributes listed in Table 15-8.

Table 15-8. OM Attributes of DSX_C_GDS_SESSION

OM Attribute	Value Syntax	Value Length	Value Number	Value Initially
DSX_PASSWORD	String(OM_S_OCTET_STRING)	—	0 or 1	—
DSX_DIR_ID	OM_S_INTEGER	—	1	1

OM Attribute	Value Syntax	Value Length	Value Number	Value Initially
DSX_AUTH_MECHANISM	Enum(DSX_Auth_Mechanism)	—	0–1	—
DSX_AUTH_INFO	String(OM_S_OCTET_STRING)	—	0–1	—

The OM attributes of **DSX_C_GDS_SESSION** are as follows:

- **DSX_PASSWORD**

This attribute indicates the password for the user credentials.

- **DSX_DIR_ID**

This attribute contains an identifier for distinguishing between several configurations of the directory service within a GDS installation. The valid range is from 1 to 20.

- **DSX_AUTH_MECHANISM**

If this attribute is present, then it identifies the authentication mechanism that the application requests. If it is absent or has the value **DSX_NONE_AT_ALL**, then a **ds_bind()** without credentials (anonymous bind) is requested. The values that this attribute can take are as follows:

- **DSX_NONE_AT_ALL**. No authentication.
- **DSX_DEFAULT**. The default authentication mechanism is to use DCE authentication, therefore, **DSX_DEFAULT** defaults to **DSX_DCE_AUTH**. The value for **DSX_DEFAULT** can be modified through an environment variable **XDS_DEF_AUTH_MECH**. The value of this environment variable is checked by the XDS software following a **ds_initialize()** function call.
- **DSX_SIMPLE**. This requests simple authentication by using the **DS_REQUESTOR** and **DSX_PASSWORD** attributes of the **DSX_C_GDS_SESSION** object.
- **DSX_SIMPLE_PROT1**. This is reserved for future use.

- **DSX_SIMPLE_PROT2**. This is reserved for future use.
- **DSX_DCE_AUTH**. This requests the use of the DCE authentication mechanism.
- **DSX_STRONG**. This is reserved for future use.

If an authentication mechanism is selected that is not currently supported, then **ds_bind()** returns a **DS_E_NOT_SUPPORTED** error. If the selected authentication mechanism requires the user's credentials that cannot be assembled, then a **DS_E_NO_INFO** error is returned.

- **DSX_AUTH_INFO**

This attribute is reserved for future use.

Applications can assume that an object of OM class **DSX_C_GDS_SESSION**, created with default values of all its OM attributes, works with all the interface functions. Such a session can be created by passing the constant **DS_DEFAULT_SESSION** as an argument to **ds_bind()**, having already negotiated the GDS package.

Table 15-9 defines **DSX_C_GDS_SESSION**.

Table 15–9. Default DSX_C_GDS_SESSION

OM Attribute	Default Value
DS_DSA_ADDRESS	Value obtained from the cache or absent
DS_DSA_NAME	Value obtained from the cache or absent
DS_FILE_DESCRIPTOR	DS_NO_VALID_FILE_DESCRIPTOR
DSX_DIR_ID	1
DSX_AUTH_MECHANISM	Absent
DSX_AUTH_INFO	Absent

Note: The values of **DS_DSA_ADDRESS** and **DS_DSA_NAME** are taken from the cache of Directory ID 1.

Chapter 16

Distributed Management Environment Support

The Distributed Management Environment (DME) Network Management Option (NMO) provides access to network management protocols. One of the protocols it supports is the CMIP protocol. CMIP uses names to identify and locate managed objects and management applications. GDS is used to provide this name to address resolution.

DME has a requirement to support opaque address forms to cater to instances where CMIP is not running over pure OSI protocols. For this purpose, GDS contains some enhancements that are described in this chapter.

To support DME an additional directory object class, and an additional directory attribute were required. Additional OM classes or OM attributes were not necessary. Therefore, GDS supports DME without having to negotiate a specific XDS/DME package. An application must include **xdsdme.h (4xds)** when using the new directory object class and attribute.

The concepts and notation used are first mentioned in Chapter 11. They are also fully explained in Chapters 17 through 19. The attribute types are introduced first, followed by the object classes.

16.1 DME Attribute Types

This section presents the additional directory attribute type that DME uses. Each attribute type has an object identifier, which is the value of the OM attribute **DS_ATTRIBUTE_TYPE**. These object identifiers are represented in the interface by constants with the same name as the directory attribute, and they are prefixed by **DSX_A_** so that they can be easily identified.

This section contains two tables that are used to indicate the object identifier for the DME attribute type (see Table 16-1), and the values for the DME attribute type (see Table 16-2), respectively. Following these two tables is a brief description of the attribute. (See Chapter 12 for information on general matching rules.)

Table 16-1 shows the name of the DME attribute type, together with the BER encoding of the associated object identifier.

Note: The second column of Table 16-1 contains the contents octets of the BER encoding of the object identifier in hexadecimal. This object identifier stems from the root **{iso(1) identified-organization(3) osf(0022) dme(2) components(1) nmo(2) dmeNmoAttributeType(1)}**.

Table 16–1. Object Identifier for DME Attribute Type

Attribute Type	Object Identifier BER
	Hexadecimal
DSX_A_ALTERNATE_ADDRESS	\x2B\x16\x02\x01\x02\x01\x01

Table 16-2 shows the name of the attribute type, together with the OM value syntax used in the interface to represent values of that attribute type. The table also includes the range of lengths permitted for the string types, indicates whether the attribute can be multivalued, and lists which matching rules are provided for the syntax.

Table 16-2. Representation of Values for DME Attribute Types

Attribute Type	OM Value Syntax	Variable Length	Multivalued	Matching Rules
DSX_A_ALTERNATE_ADDRESS	String(OM_S_OCTET_STRING)	1-800	yes	E

The following is a description of the DME attribute:

- **DSX_A_ALTERNATE_ADDRESS**

This attribute is used by DME to store opaque address formats. In Table 16-2, it can be seen that the **AlternateAddress** attribute is stored internally by GDS as an octet string. The application expects the following syntax:

```
AlternateAddress ::= SEQUENCE {
    address      OCTET STRING,
    protocol     SET OF OBJECT IDENTIFIER }
```

For conversion between octet string and a C structure corresponding to this definition, two functions are provided: **gds_encode_alt_addr (3xds)** and **gds_decode_alt_addr (3xds)**.

16.2 DME Object Classes

The only additional DME object class is **DSX_O_DME_NMO_AGENT** (see Table 16-3). This object class has the same structure rules in the default schema as the application entity object class. **DSX_O_DME_NMO_AGENT** is a subclass of **DS_O_APPLIC_ENTITY** (inherits the mandatory **DS_A_PRESENTATION_ADDRESS** and **DS_A_COMMON_NAME** attributes) and contains one attribute, **DSX_A_ALTERNATE_ADDRESS**.

Note: The second column of Table 16-3 contains the contents octets of the BER encoding of the object identifier in hexadecimal. This object identifier stems from the root **{iso(1) identified-organization(3) osf(0022) dme(2) components(1) nmo(2) dmeNmoObjectClass(2)}**.

Table 16–3. Object Identifier for DME Object Class

Attribute Type	Object Identifier BER
	Hexadecimal
DSX_O_DME_NMO_AGENT	\x2B\x16\x02\x01\x02\x02\x01

Chapter 17

Information Syntaxes

This chapter defines the syntaxes permitted for attribute values. The syntaxes are closely aligned with the types and type constructors of ASN.1. The **OM_value** data type specifies how a value of each syntax is represented in the C interface (see Chapter 18).

17.1 Syntax Templates

The names of certain syntaxes are constructed from *syntax templates*. A syntax template is a lexical construct comprising a primary identifier followed by an * (asterisk) enclosed in parentheses, as follows:

identifier (*)

A syntax template encompasses a group of related syntaxes. Any member of the group, without distinction, is indicated by the primary identifier (*identifier*) alone. A particular member is indicated by the template with the asterisk replaced by one of a set of secondary identifiers associated with the template, as follows:

identifier₁ (identifier₂)

17.2 Syntaxes

A variety of syntaxes are defined. Most are functionally equivalent to ASN.1 types, as documented in Sections 17.5 through 17.8.

The following syntaxes are defined:

- **OM_S_BOOLEAN**

A value of this syntax is a Boolean; that is, it can be **OM_TRUE** or **OM_FALSE**.

- **Enum(*)**

A value of any syntax encompassed by this syntax template is one of a set of values associated with the syntax. The only significant characteristic of the values is that they are distinct.

The group of syntaxes encompassed by this template is open-ended. Zero or more members are added to the group by each package definition. The secondary identifiers that indicate the members are also assigned there.

- **OM_S_INTEGER**

A value of this syntax is a positive or negative integer.

- **OM_S_NULL**

The one value of this syntax is a valueless placeholder.

- **Object(*)**

A value of any syntax encompassed by this syntax template is an object, which is any instance of a class associated with the syntax.

The group of syntaxes encompassed by this template is open-ended. One member is added to the group by each class definition. The secondary identifier that indicates the member is the name of the class.

- **String(*)**

A value of any syntax encompassed by this syntax template is a string (as defined in Section 17.3) whose form and meaning are associated with the syntax.

The group of syntaxes encompassed by this template is closed. One syntax is defined for each ASN.1 string type. The secondary identifier that indicates the member is, in general, the first word of the type's name.

17.3 Strings

A *string* is an ordered sequence of zero or more bits, octets, or characters accompanied by the string's length.

The value *length* of a string is the number of bits in a *bit string*, octets in an *octet string*, or characters in a *character string*. Any constraints on the value length of a string are specified in the appropriate class definitions. The length is confined to the range 0 to 2^{32} .

Note: The length of a character string does not necessarily equal the number of characters it comprises because, for example, a single character can be represented by using several octets.

The elements of a string are numbered. The position of the first element is 0 (zero). The positions of successive elements are successive positive integers.

The syntaxes that form the string group are identified in Table 17-1, which gives the secondary identifier assigned to each such syntax.

Note: The identifiers in the first, second, and third columns of Table 17-1 indicate the syntaxes of bit, octet, and character strings, respectively. The String group comprises all syntaxes identified in the table.

Table 17–1. String Syntax Identifiers

Bit String Identifier	Octet String Identifier	Character String Identifier
OM_S_BIT_STRING	OM_S_ENCODING_STRING ¹	OM_S_GENERAL_STRING ²
	OM_S_OBJECT_IDENTIFIER_STRING ³	OM_S_GENERALIZED_TIME_STRING ²
	OM_S_OCTET_STRING	OM_S_GRAPHIC_STRING ²
		OM_S_IA5_STRING ²
		OM_S_NUMERIC_STRING ²
		OM_S_OBJECT_DESCRIPTOR_STRING ²
		OM_S_PRINTABLE_STRING ²
		OM_S_TELETEX_STRING ²
		OM_S.UTC_TIME_STRING ²
		OM_S_VIDEOTEX_STRING ²
		OM_S_VISIBLE_STRING ²

- 1 The octets are those that BER permits for the contents octets of the encoding of a value of any ASN.1 type.
- 2 The characters are those permitted by ASN.1's type of the corresponding name. Values of these syntaxes are represented in their BER-encoded form. The octets by which they are represented are those that BER permits for the contents octets of a primitive encoding of a value of that type.
- 3 The octets are those that BER permits for the contents octets of the encoding of a value of ASN.1's object identifier type.

17.4 Representation of String Values

In the service interface, a string value is represented by a string data type. This is defined in Section 17.3. The length of a string is the number of octets by which it is represented at the interface. It is confined to the range 0 to 2^{32} .

The length of a character does not need to be equal to the number of characters it comprises because, for example, a single character can be represented by using several octets.

It may be necessary to segment large string values when passing them across the interface. A segment is any zero or more contiguous octets of a string value. Segment boundaries are without semantic significance.

17.5 Relationship to ASN.1 Simple Types

As shown in Table 17-2, for every ASN.1 simple type except Real, there is an OM syntax that is functionally equivalent to it. The simple types are listed in the first column of the table; the corresponding syntaxes are listed in the second column.

Table 17-2. Syntax for ASN.1 Simple Types

Type	Syntax
Bit String	String(OM_S_BIT_STRING)
Boolean	OM_S_BOOLEAN
Integer	OM_S_INTEGER
Null	OM_S_NULL
Object Identifier	String(OM_S_OBJECT_IDENTIFIER_STRING)
Octet String	String(OM_S_OCTET_STRING)
Real	None ¹

1 A future edition of XOM may define a syntax corresponding to this type.

17.6 Relationship to ASN.1 Useful Types

As shown in Table 17-3, for every ASN.1 useful type, there is an OM syntax that is functionally equivalent to it. The useful types are listed in the first column of the table; the corresponding syntaxes are listed in the second column.

Table 17-3. Syntaxes for ASN.1 Useful Types

Type	Syntax
External	Object(OM_C_EXTERNAL)
Generalized Time	String(OM_S_GENERALISED_TIME_STRING)
Object Descriptor	String(OM_S_OBJECT_DESCRIPTOR_STRING)
Universal Time	String(OM_S_UTC_TIME_STRING)

17.7 Relationship to ASN.1 Character String Types

As shown in Table 17-4, for every ASN.1 character string type, there is an OM syntax that is functionally equivalent to it. The ASN.1 character string types are listed in the first column of the table; the corresponding syntax is listed in the second column.

Table 17-4. Syntaxes for ASN.1 Character String Types

Type	Syntax
General String	String(OM_S_GENERAL_STRING)
Graphic String	String(OM_S_GRAPHIC_STRING)
IA5 String	String(OM_S_IA5_STRING)
—	String(OM_S_LOCAL_STRING)
Numeric String	String(OM_S_NUMERIC_STRING)
Printable String	String(OM_S_PRINTABLE_STRING)
Teletex String	String(OM_S_TELETEX_STRING)

Type	Syntax
Videotex String	String(OM_S_VIDEOTEX_STRING)
Visible String	String(OM_S_VISIBLE_STRING)

17.8 Relationship to ASN.1 Type Constructors

As shown in Table 17-5, there are functionally equivalent OM syntaxes for some (but not all) ASN.1 type constructors. The constructors are listed in the first column; corresponding syntaxes are listed in the second column.

Table 17-5. Syntaxes for ASN.1 Type Constructors

Type Constructor	Syntax
Any	String(OM_S_ENCODING_STRING)
Choice	OM_S_OBJECT
Enumerated	OM_S_ENUMERATION
Selection	None ¹
Sequence	OM_S_OBJECT
Sequence Of	OM_S_OBJECT
Set	OM_S_OBJECT
Set Of	OM_S_OBJECT
Tagged	None ²

- 1 This type constructor, a purely specification-time phenomenon, has no corresponding syntax.
- 2 This type constructor is used to distinguish the alternatives of a choice or the elements of a sequence or set, a function performed by attribute types.

The effects of the principal type constructors can be achieved, in any of a variety of ways, by using objects-to-group attributes or using attributes-to-group values. An

OM application designer can (but need not) model these constructors as classes of the following kinds:

- Choice

An attribute type can be defined for each alternative, with just one being permitted in an instance of the class.

- Sequence or Set

An attribute type can be defined for each sequence or set element. If an element is optional, then the attribute has zero or one value.

- Sequence Of or Set Of

A single multivalued attribute can be defined.

An ASN.1 definition of an enumerated type component of a structured type is generally mapped to an OM attribute with an OM syntax **OM_S_ENUMERATION** in this interface. Where the ASN.1 component is optional, this is generally indicated by an additional member of the enumeration, rather than by the omission of the OM attribute. This leads to simpler programming in the application.

Chapter 18

XOM Service Interface

This chapter describes the following aspects of the XOM service interface:

- The conformance of the DCE X/Open OSI-Abstract-Data Manipulation (XOM) implementation to the X/Open specification.
- The data types whose data values are the parameters and results of the functions that the service makes available to the client.
- An overview of the functions that the service makes available to the client. For a complete description of these functions, see the corresponding reference pages.
- The return codes that indicate the outcomes (in particular, the exceptions) that the functions can report.

See Chapter 7 for examples of using the XOM interface.

18.1 Standards Conformance

The DCE XOM implementation conforms to the following specification:

X/Open CAE Specification, OSI-Abstract-Data Manipulation (XOM) (November 1991)

The following apply to the DCE XOM implementation:

- Multiple workspaces for XDS objects are supported.
- The OM package is supported.
- The **om_encode()** and **om_decode()** functions are not supported. The transfer of objects between workspaces is not envisaged within the DCE environment. The OM classes used by the DCE XDS/XOM API are not encodable.
- Translation to local character sets is supported.

18.2 XOM Data Types

The data types of the XOM service interface are defined in this section and listed in Table 18-1. These data types are repeated in the XOM reference pages (see **xom.h(4xom)**).

Table 18–1. XOM Service Interface Data Types

Data Type	Description
OM_boolean	Type definition for a Boolean data value.
OM_descriptor	Type definition for describing an attribute type and value.
OM_enumeration	Type definition for an Enumerated data value.
OM_exclusions	Type definition for the <i>exclusions</i> parameter for om_get() .
OM_integer	Type definition for an Integer data value.
OM_modification	Type definition for the <i>modification</i> parameter for om_put() .
OM_object	Type definition for a handle to either a private or a public object.

Data Type	Description
OM_object_identifier	Type definition for an object identifier data value.
OM_private_object	Type definition for a handle to an object in an implementation-defined, or private, representation.
OM_public_object	Type definition for a defined representation of an object that can be directly interrogated by a programmer.
OM_return_code	Type definition for a value returned from all OM functions, indicating either that the function succeeded or why it failed.
OM_string	Type definition for a data value of one of the String syntaxes.
OM_syntax	Type definition for identifying a syntax type.
OM_type	Type definition for identifying an OM attribute type.
OM_type_list	Type definition for enumerating a sequence of OM attribute types.
OM_value	Type definition for representing any data value.
OM_value_length	Type definition for indicating the number of bits, octets, or characters in a string.
OM_value_position	Type definition for designating a particular location within a String data value.
OM_workspace	Type definition for identifying an application-specific API that implements OM, such as directory or message handling.

Some data types are defined in terms of the following *intermediate data types*, whose precise definitions in C are defined by the system:

- **OM_sint**

The positive and negative integers that can be represented in 16 bits

- **OM_sint16**

The positive and negative integers that can be represented in 16 bits

- **OM_sint32**

The positive and negative integers that can be represented in 32 bits

- **OM_uint**

The nonnegative integers that can be represented in 16 bits

- **OM_uint16**

The nonnegative integers that can be represented in 16 bits

- **OM_uint32**

The nonnegative integers that can be represented in 32 bits

Note: The **OM_sint** and **OM_uint** data types are defined by the range of integers they must accommodate. As typically declared in the C interface, they are defined by the range of integers permitted by the host machine's word size. The latter range, however, always encompasses the former.

The type definitions for these data types are as follows:

```
typedef int          OM_sint;  
typedef short       OM_sint16;  
typedef long int    OM_sint32;  
typedef unsigned    OM_uint;  
typedef unsigned short OM_uint16;  
typedef long unsigned OM_uint32;
```

18.2.1 OM_boolean

The C declaration for an **OM_boolean** data value is as follows:

```
typedef OM_uint32 OM_boolean;
```

A data value of this data type is a Boolean; that is, either FALSE or TRUE.

FALSE (**OM_FALSE**) is indicated by 0 (zero). TRUE is indicated by any other integer, although the symbolic constant **OM_TRUE** refers to the integer 1 specifically.

18.2.2 OM_descriptor

The **OM_descriptor** data type is used to describe an attribute type and value. Its C declaration is as follows:

```
typedef struct OM_descriptor_struct
{
    OM_type      type;
    OM_syntax     syntax;
    union OM_value_union value;
} OM_descriptor;
```

Note: Other components are encoded in high bits of the syntax member.

See the **OM_value** data type described in Section 18.2.16 or the **xom.h(4xom)** reference page for a description of the **OM_value_union** structure.

A data value of this type is a descriptor, which embodies an attribute value. An array of descriptors can represent all the values of all the attributes of an object, and is the representation called **OM_public_object**. A descriptor has the following components:

- *type*

An **OM_type** data type. It identifies the data type of the attribute value.

- *syntax*

An **OM_syntax** data type. It identifies the syntax of the attribute value. Components 3 to 7 (that is, the components *long-string* through *private* that follow) are encoded in the high-order bits of this structure member. Therefore, the syntax always needs to be masked with the constant **OM_S_SYNTAX**. An example is the following:

```
my_syntax = my_public_object[3].syntax &
            OM_S_SYNTAX;

my_public_object[4].syntax =
my_syntax + (my_public_object[4].syntax &
            ~OM_S_SYNTAX);
```

- *long-string*

An **OM_boolean** data type. It is **OM_TRUE** only if the descriptor is a service-generated descriptor and the length of the value is greater than an implementation-defined limit.

This component occupies bit 15 (0x8000) of the syntax and is represented by the constant **OM_S_LONG_STRING**.

- *no-value*

An **OM_boolean** data type. It is **OM_TRUE** only if the descriptor is a service-generated descriptor and the value is not present because **OM_EXCLUDE_VALUES** or **OM_EXCLUDE_MULTIPLES** is set in **om_get()**.

This component occupies bit 14 (0x4000) of the syntax and is represented by the constant **OM_S_NO_VALUE**.

- *local-string*

An **OM_boolean** data type, significant only if the syntax is one of the string syntaxes. It is **OM_TRUE** only if the string is represented in an implementation-defined local character set. The local character set may be more amenable for use as keyboard input or display output than the nonlocal character set, and it can include specific treatment of line termination sequences. Certain interface functions can convert information in string syntaxes to or from the local representation, which may result in a loss of information.

This component occupies bit 13 (0x2000) of the syntax and is represented by the constant **OM_S_LOCAL_STRING**. The DCE XOM implementation does not support translation of strings to a local character set.

- *service-generated*

An **OM_boolean** data type. It is **OM_TRUE** only if the descriptor is a service-generated descriptor and the first descriptor of a public object, or the defined part of a private object (see the ***(3xom)** reference pages).

This component occupies bit 12 (0x1000) of the syntax and is represented by the constant **OM_S_SERVICE_GENERATED**.

- *private*

An **OM_boolean** data type. It is **OM_TRUE** only if the descriptor in the service-generated public object contains a reference to the handle of a private subobject, or in the defined part of a private object.

Note: This applies only when the descriptor is a service-generated descriptor. The client need not set this bit in a client-generated descriptor that contains a reference to a private object.

In the C interface, this component occupies bit 11 (0x0800) of the syntax and is represented by the constant **OM_S_PRIVATE**.

- *value*

An **OM_value** data type. It identifies the attribute value.

18.2.3 OM_enumeration

The **OM_enumeration** data type is used to indicate an Enumerated data value. Its C declaration is as follows:

```
typedef OM_sint32 OM_enumeration;
```

A data value of this data type is an attribute value whose syntax is **OM_S_ENUMERATION**.

18.2.4 OM_exclusions

The **OM_exclusions** data type is used for the *exclusions* parameter of **om_get()**. Its C declaration is as follows:

```
typedef OM_uint OM_exclusions;
```

A data value of this data type is an unordered set of one or more values, all of which are distinct. Each value indicates an exclusion, as defined by **om_get()**, and is chosen from the following set:

- **OM_EXCLUDE_ALL_BUT_THESE_TYPES**
- **OM_EXCLUDE_MULTIPLES**
- **OM_EXCLUDE_ALL_BUT_THESE_VALUES**
- **OM_EXCLUDE_VALUES**
- **OM_EXCLUDE_SUBOBJECTS**
- **OM_EXCLUDE_DESCRIPTOR**

Alternatively, the single value **OM_NO_EXCLUSIONS** can be chosen; this selects the entire object.

Each value except **OM_NO_EXCLUSIONS** is represented by a distinct bit. The presence of the value is represented as 1; its absence is represented as 0 (zero). Thus, multiple exclusions are requested by ORing the values that indicate the individual exclusions.

18.2.5 OM_integer

The **OM_integer** data type is used to indicate an integer data value. Its C declaration is as follows:

```
typedef OM_sint32 OM_integer;
```

A data value of this data type is an attribute value whose syntax is **OM_S_INTEGER**.

18.2.6 OM_modification

The **OM_modification** data type is used for the *modification* parameter of **om_put()**. Its C declaration is as follows:

```
typedef OM_uint OM_modification;
```

A data value of this data type indicates a kind of modification, as defined by **om_put()**. It is chosen from the following set:

- **OM_INSERT_AT_BEGINNING**
- **OM_INSERT_AT_CERTAIN_POINT**
- **OM_INSERT_AT_END**
- **OM_REPLACE_ALL**
- **OM_REPLACE_CERTAIN_VALUES**

18.2.7 OM_object

The **OM_object** data type is used as a handle to either a private or a public object. Its C declaration is as follows:

```
typedef struct OM_descriptor_struct *OM_object;
```

A data value of this data type represents an object, which can be either public or private. It is an ordered sequence of one or more instances of the **OM_descriptor** data type. See the **OM_private_object** and **OM_public_object** data types for restrictions on that sequence (Sections 18.2.9 and 18.2.10, respectively).

18.2.8 OM_object_identifier

The **OM_object_identifier** data type is used as an ASN.1 object identifier. Its C declaration is as follows:

```
typedef OM_string OM_object_identifier;
```

A data value of this data type contains an octet string that comprises the contents octets of the BER encoding of an ASN.1 object identifier.

18.2.8.1 C Declaration of Object Identifiers

Every application program that uses a class or another object identifier must explicitly import it into every compilation unit (C source module) that uses it. Each such class or object identifier name must be explicitly exported from just one compilation module. Most application programs find it convenient to export all the names they use from the same compilation unit. Exporting and importing is performed by using the following two macros:

- The importing macro makes the class or other object identifier constants available within a compilation unit.
 - **OM_IMPORT**(*class_name*)
 - **OM_IMPORT**(*OID_name*)
- The exporting macro allocates memory for the constants that represent the class or another object identifier.
 - **OM_EXPORT**(*class_name*)
 - **OM_EXPORT**(*OID_name*)

Object identifiers are defined in the appropriate header files, with the definition identifier having the prefix **OMP_O_** followed by the variable name for the object identifier. The constant itself provides the hexadecimal value of the object identifier string.

18.2.8.2 Use of Object Identifiers in C

The following macro initializes a descriptor:

OM_OID_DESC(*type*, *OID_name*)

It sets the *type* component to that given, sets the *syntax* component to **OM_S_OBJECT_IDENTIFIER_STRING**, and sets the *value* component to the specified object identifier.

The following macro initializes a descriptor to mark the end of a client-allocated public object:

OM_NULL_DESCRIPTOR

For each class, there is a global variable of type **OM_STRING** with the same name; for example, the External class has a variable called **OM_C_EXTERNAL**. This is also the case for other object identifiers; for example, the object identifier for BER rules has a variable called **OM_BER**. This global variable can be supplied as a parameter to functions when required.

This variable is valid only when it is exported by an **OM_EXPORT** macro and imported by an **OM_IMPORT** macro in the compilation units that use it. This variable cannot form part of a descriptor, but the value of its length and elements components can be used. The following code fragment provides examples of the use of the macros and constants.

```
/* Examples of the use of the macros and constants */

#include <xom.h>

OM_IMPORT(OM_C_ENCODING)
OM_IMPORT(OM_CANONICAL_BER)

/* The following sequence must appear in just one compilation
 * unit in place of the above:
 *
 * #include <xom.h>
 *
 * OM_EXPORT(OM_C_ENCODING)
 * OM_EXPORT(OM_CANONICAL_BER)
 */

main( )
{
/* Use #1 - Define a public object of class Encoding
```

```
*          (Note: xxxx is a Message Handling class which
*          can be encoded)
*/
OM_descriptor my_public_object[] = {
    OM_OID_DESC(OM_CLASS, OM_C_ENCODING),
    OM_OID_DESC(OM_OBJECT_CLASS, MA_C_xxxx),
    { OM_OBJECT_ENCODING, OM_S_ENCODING_STRING, \
      some_BER_value },
    OM_OID_DESC(OM_RULES, OM_CANONICAL_BER),
    OM_NULL_DESCRIPTOR
};

/* Use #2 - Pass class Encoding as parameter to om_instance()
*/
return_code = om_instance(my_object, OM_C_ENCODING,
&boolean_result);
}
```

18.2.9 OM_private_object

The **OM_private_object** data type is used as a handle to an object in an implementation-defined or private representation. Its C declaration is as follows:

```
typedef OM_object OM_private_object;
```

A data value of this data type is the designator or handle to a private object. It comprises a single descriptor whose *type* component is **OM_PRIVATE_OBJECT** and whose *syntax* and *value* components are unspecified.

Note: The descriptor's *syntax* and *value* components are essential to the service's proper operation with respect to the private object.

18.2.10 OM_public_object

The **OM_public_object** data type is used to define an object that can be directly accessed by a programmer. Its C declaration is as follows:

```
typedef OM_object OM_public_object;
```

A data value of this data type is a public object. It comprises one or more (usually more) descriptors, all but the last of which represent values of attributes of the object.

The descriptors for the values of a particular attribute with two or more values are adjacent to one another in the sequence. Their order is that of the values they represent. The order of the resulting groups of descriptors is unspecified.

Since the Class attribute specific to the Object class is represented among the descriptors, it must be represented before any other attributes. Regardless of whether or not the Class attribute is present, the syntax field of the first descriptor must have the **OM_S_SERVICE_GENERATED** bit set or cleared appropriately.

The last descriptor signals the end of the sequence of descriptors. The last descriptor's *type* component is **OM_NO_MORE_TYPES** and its *syntax* component is **OM_S_NO_MORE_SYNTAXES**. The last descriptor's *value* component is unspecified.

18.2.11 OM_return_code

The **OM_return_code** data type is used for a value that is returned from all OM functions, indicating either that the function succeeded or why it failed. Its C declaration is as follows:

```
typedef OM_uint OM_return_code;
```

A data value of this data type is the integer in the range 0 to 2^{16} that indicates an outcome of an interface function. It is chosen from the set specified in Section 18.4.

Integers in the narrower range 0 to 2^{15} are used to indicate the return codes they define.

18.2.12 OM_string

The **OM_string** data type is used for a data value of String syntax. Its C declaration is as follows:

```
typedef OM_uint32 OM_string_length;
typedef struct {
    OM_string_length length;
    void *elements;
} OM_string;

#define OM_STRING(string)\
    { (OM_string_length)(sizeof(string)-1), (string) }
```

A data value of this data type is a string; that is, an instance of a String syntax. A string is specified either in terms of its length or whether or not it terminates with NULL.

A string has the following components:


- *length* (**OM_string_length**)

The number of octets by means of which the string is represented, or the **OM_LENGTH_UNSPECIFIED** value if the string terminates with NULL.

- *elements*

The string's elements. The bits of a bit string are represented as a sequence of octets (see Figure 18-1). The first octet stores the number of unused bits in the last octet. The bits in the bit string, commencing with the first bit and proceeding to the trailing bit, are placed in bits 7 to 0 of the second octet. These are followed by bits 7 to 0 of the third octet, then by bits 7 to 0 of each octet in turn, followed by as many bits as are required of the final octet, commencing with bit 7.

Figure 18–1. OM_String Elements



The service supplies a string value with a specified length. The client can supply a string value to the service in either form, either with a specified length or terminated with NULL.

The characters of a character string are represented as any sequence of octets permitted as the primitive contents octets of the BER encoding of an ASN.1 type value. The ASN.1 type defines the type of character string. A 0 (zero) value character follows the characters of the character string, but is not encompassed by the *length* component. Thus, depending on the type of character string, the 0 (zero) value character can delimit the characters of the character string.

The **OM_STRING** macro is provided for creating a data value of this data type, given only the value of its *elements* component. The macro, however, applies to octet strings and character strings, but not to bit strings.

18.2.13 OM_syntax

The **OM_syntax** data type is used to identify a syntax type. Its C declaration is as follows:

```
typedef OM_uint16 OM_syntax;
```

A data value of this data type is an integer in the range 0 to 2^9 that indicates an individual syntax or a set of syntaxes taken together.

The data value is chosen from among the following:

- **OM_S_BIT_STRING**
- **OM_S_BOOLEAN**
- **OM_S_ENCODING_STRING**
- **OM_S_ENUMERATION**
- **OM_S_GENERAL_STRING**
- **OM_S_GENERALIZED_TIME_STRING**
- **OM_S_GRAPHIC_STRING**
- **OM_S_IA5_STRING**
- **OM_S_INTEGER**
- **OM_S_NULL**
- **OM_S_NUMERIC_STRING**
- **OM_S_OBJECT**
- **OM_S_OBJECT_DESCRIPTOR_STRING**
- **OM_S_OBJECT_IDENTIFIER_STRING**
- **OM_S_OCTET_STRING**
- **OM_S_PRINTABLE_STRING**
- **OM_S_TELETEX_STRING**
- **OM_S_VIDEOTEX_STRING**
- **OM_S_VISIBLE_STRING**
- **OM_S_UTC_TIME_STRING**

Integers in the narrower range 0 to 2^9 are used to indicate the syntaxes they define. The integers in the range 2^9 to 2^{10} are reserved for vendor extensions. Wherever possible, the integers used are the same as the corresponding ASN.1 universal class number.

18.2.14 OM_type

The **OM_type** data type is used to identify an OM attribute type. Its C declaration is as follows:

```
typedef OM_uint16 OM_type;
```

A data value of this data type is an integer in the range 0 to 2^{16} that indicates a type in the context of a package. However, the following values in Table 18-2 are assigned meanings by the respective data types.

Table 18–2. Assigning Meanings to Values

Value	Data Type
OM_NO_MORE_TYPES	OM_type_list
OM_PRIVATE_OBJECT	OM_private_object

Integers in the narrower range 0 to 2^{15} are used to indicate the types they define.

18.2.15 OM_type_list

The **OM_type_list** data type is used to enumerate a sequence of OM attribute types. Its C declaration is as follows:

```
typedef OM_type *OM_type_list;
```

A data value of this data type is an ordered sequence of zero or more type numbers, each of which is an instance of the **OM_type** data type.

An additional data value, **OM_NO_MORE_TYPES**, follows and thus delimits the sequence. The C representation of the sequence is an array.

18.2.16 OM_value

The **OM_value** data type is used to represent any data value. Its C declaration is as follows:

```
typedef struct {
    OM_uint32 padding;
    OM_object object;
} OM_padded_object;

typedef union OM_value_union {
    OM_string      string;
    OM_boolean     boolean;
    OM_enumeration enumeration;
    OM_integer     integer;
    OM_padded_object object;
} OM_value;
```

Note: The first type definition (in particular, its **padding** component) aligns the **object** component with the *elements* component of the **string** component in the second type definition. This facilitates initialization in C.

The identifier **OM_value_union** is defined for reasons of compilation order. It is used in the definition of the **OM_descriptor** data type.

A data value of this data type is an attribute value. It has no components if the value's syntax is **OM_S_NO_MORE_SYNTAXES** or **OM_S_NO_VALUE**. Otherwise, it has one of the following components:

- **string**

The value if its syntax is a string syntax

- **boolean**

The value if its syntax is **OM_S_BOOLEAN**

- **enumeration**

The value if its syntax is **OM_S_ENUMERATION**

- **integer**

The value if its syntax is **OM_S_INTEGER**

- **object**

The value if its syntax is **OM_S_OBJECT**

Note: A data value of this data type is only displayed as a component of a descriptor. Thus, it is always accompanied by indicators of the value's syntax. The latter indicator reveals which component is present.

18.2.17 **OM_value_length**

The **OM_value_length** data type is used to indicate the number of bits, octets, or characters in a string. Its C declaration is as follows:

```
typedef OM_uint32 OM_value_length;
```

A data value of this data type is an integer in the range 0 to 2^{32} that represents the number of bits in a bit string, octets in an octet string, or characters in a character string.

Note: This data type is not used in the definition of the interface. It is provided for use by client programmers for defining attribute constraints.

18.2.18 **OM_value_position**

The **OM_value_position** data type is used to indicate an attribute value's position within an attribute. Its C declaration is as follows:

```
typedef OM_uint32 OM_value_position;
```

A data value of this data type is an integer in the range 0 to $2^{32}-1$ that indicates the position of a value within an attribute. However, the value **OM_ALL_VALUES** has the meaning assigned to it by **om_get()**.

18.2.19 OM_workspace

The **OM_workspace** data type is used to identify an application-specific API that implements OM; for example, directory or message handling. Its C declaration is as follows:

```
typedef void *OM_workspace;
```

A data value of this data type is the designator or handle for a workspace.

18.3 XOM Functions

This section provides an overview of the XOM service interface functions as listed in Table 18-3. For a full description of these functions, see the ***(3xom)** reference pages.

Table 18-3. XOM Service Interface Functions

Function	Description
om_copy()	Copies a private object.
om_copy_value()	Copies a string between private objects.
om_create()	Creates a private object.
om_decode()	Not supported by the DCE XOM interface; it returns an OM_FUNCTION_DECLINED error.
om_delete()	Deletes a private or service-generated object.
om_encode()	Not supported by the DCE XOM interface; it returns an OM_FUNCTION_DECLINED error.
om_get()	Gets copies of attribute values from a private object.
om_instance()	Tests an object's class.

Function	Description
om_put()	Puts attribute values into a private object.
om_read()	Reads a segment of a string in a private object.
om_remove()	Removes attribute values from a private object.
om_write()	Writes a segment of a string into a private object.

The purpose and range of capabilities of the service interface functions can be summarized as follows:

- **om_copy()**

This function creates an independent copy of an existing private object and all its subobjects. The copy is placed in the workspace of the original object, or in another workspace specified by the DCE client.

- **om_copy_value()**

This function replaces an existing attribute value or inserts a new value in one private object with a copy of an existing attribute value found in another. Both values must be strings.

- **om_create()**

This function creates a new private object that is an instance of a particular class. The object can be initialized with the attribute values specified as initial in the class definition. The service does not permit the client to explicitly create instances of all classes, but rather only those indicated by a package's definition as having this property.

- **om_delete()**

This function deletes a service-generated public object or makes a private object inaccessible.

- **om_get()**

This function creates a new public object that is an exact, but independent, copy of an existing private object. The client can request certain exclusions, each of which reduces the copy to a part of the original. The client can also request that values be converted from one syntax to another before they are returned.

The copy can exclude attributes of types other than those specified, values at positions other than those specified within an attribute, values of multivalued attributes, copies of (not handles for) subobjects, or all attribute values. Excluding all attribute values reveals only an attribute's presence.

- **om_instance()**

This function determines whether an object is an instance of a particular class. The client can determine an object's class simply by inspection. This function is useful since it reveals that an object is an instance of a particular class, even if the object is an instance of a subclass of that class.

- **om_put()**

This function places or replaces in one private object copies of the attribute values of another public or private object.

The source values can be inserted before any existing destination values, before the value at a specified position in the destination attribute, or after any existing destination values. Alternatively, the source values can be substituted for any existing destination values or for the values at specified positions in the destination attribute.

- **om_read()**

This function reads a segment of a value of an attribute of a private object. The value must be a string. The value can first be converted from one syntax to another. This function enables the client to read an arbitrarily long value without requiring that the service place a copy of the entire value in memory.

- **om_remove()**

This function removes and discards particular values of an attribute of a private object. The attribute itself is removed if no values remain.

- **om_write()**

This function writes a segment of an attribute value to a private object. The value must be a string. The segment can first be converted from one syntax to another. The written segment becomes the value's last segment since any elements beyond it are discarded. The function enables the client to write an arbitrarily long value without having to place a copy of the entire value in memory.

18.4 XOM Return Codes

This section defines the return codes of the service interface, and thus the exceptions that can prevent the successful completion of an interface function.

Refer to the **ERRORS** section of the **(3xom)* references pages for a list of the errors that each function can return. For an explanation of these error codes, refer to the *OSF DCE Problem Determination Guide*.

The return code values are as follows:

0	OM_SUCCESS
1	OM_ENCODING_INVALID
2	OM_FUNCTION_DECLINED
3	OM_FUNCTION_INTERRUPTED
4	OM_MEMORY_INSUFFICIENT
5	OM_NETWORK_ERROR
6	OM_NO_SUCH_CLASS
7	OM_NO_SUCH_EXCLUSION
8	OM_NO_SUCH_MODIFICATION
9	OM_NO_SUCH_OBJECT
10	OM_NO_SUCH_RULES
11	OM_NO_SUCH_SYNTAX
12	OM_NO_SUCH_TYPE
13	OM_NO_SUCH_WORKSPACE

14	OM_NOT_AN_ENCODING
15	OM_NOT_CONCRETE
16	OM_NOT_PRESENT
17	OM_NOT_PRIVATE
18	OM_NOT_THE_SERVICES
19	OM_PERMANENT_ERROR
20	OM_POINTER_INVALID
21	OM_SYSTEM_ERROR
22	OM_TEMPORARY_ERROR
23	OM_TOO_MANY_VALUES
24	OM_VALUES_NOT_ADJACENT
25	OM_WRONG_VALUE_LENGTH
26	OM_WRONG_VALUE_MAKEUP
27	OM_WRONG_VALUE_NUMBER
28	OM_WRONG_VALUE_POSITION
29	OM_WRONG_VALUE_SYNTAX
30	OM_WRONG_VALUE_TYPE

Chapter 19

Object Management Package

This chapter defines the object management package (OMP). The object identifier (referred to as **om**) assigned to the package, as defined by this guide, is the object identifier specified in ASN.1 as

```
{joint-iso-ccitt(2) mhs-motis(6) group(6) white(1) api(2) om(4)}
```

19.1 Class Hierarchy

This section shows the hierarchical organization of the OM classes. Subclassification is indicated by indentation, and the names of abstract classes are in italics. Thus, for example, **OM_C_ENCODING** is an immediate subclass of *OM_C_OBJECT*, an abstract class. The names of classes to which **om_encode()** applies are in boldface. (DCE XOM does not support the encoding of any OM classes.) The **om_create()** function applies to all concrete classes.

- *OM_C_OBJECT*
 - **OM_C_ENCODING**

— OM_C_EXTERNAL

19.2 Class Definitions

The following subsections define the OM classes.

19.2.1 OM_C_ENCODING

An instance of class **OM_C_ENCODING** is an object represented in a form suitable for transmission between workspaces, for transport via a network, or for storage in a file. Encoding can also be a suitable way of indicating to an intermediate service provider (for example, a directory, or message transfer system) an object that it does not recognize.

This class has the attributes of its superclass, *OM_C_OBJECT*, in addition to the specific attributes listed in Table 19-1.

Table 19-1. Attributes Specific to OM_C_ENCODING

Attribute	Value Syntax	Value Length	Value Number	Value Initially
OM_OBJECT_CLASS	String(OM_S_OBJECT_IDENTIFIER_STRING)	—	1	—
OM_OBJECT_ENCODING	String ¹	—	1	—
OM_RULES	String(OM_S_OBJECT_IDENTIFIER_STRING)	—	1	ber

1 If the **Rules** attribute is **ber** or **canonical-ber**, the syntax of the present attribute must be String(OM_S_ENCODING_STRING).

- **OM_OBJECT_CLASS**

This attribute identifies the class of the object that the **Object Encoding** attribute encodes. The class must be concrete.

- **OM_OBJECT_ENCODING**

This attribute is the encoding itself.

- **OM_RULES**

This attribute identifies the set of rules that are followed to produce the **Object Encoding** attribute. Among the defined values of this attribute are those represented as follows:

- **OM_BER**

This value is specified in ASN.1 as

{joint-iso-ccitt(2) asn1(1) basic-encoding(1)}

This value indicates the BER. (See Clause 25.2 of Recommendation X.209, “Specification of Basic Encoding Rules for Abstract Syntax Notation 1 (ASN.1),” *CCITT Blue Book*, Fascicle VIII.4, International Telecommunications Union, 1988. Also published by ISO as *ISO 8825*.)

- **OM_CANONICAL_BER**

This value is specified in ASN.1 as

**{joint-iso-ccitt(2) mhs-motis(6) group(6) white(1) api(2) om(4)
canonical-ber(4)}**

This value indicates the canonical BER. (See Clause 8.7 of Recommendation X.509, “The Directory: Authentication Framework,” *CCITT Blue Book*, International Telecommunications Union, 1988. Also published by ISO as *ISO 9594-8*.)

Note: In general, an instance of this class cannot appear as a value whose syntax is Object (C) if *C* is not **OM_C_ENCODING**, even if the class of the object encoded is *C*.

19.2.2 OM_C_EXTERNAL

An instance of class **OM_C_EXTERNAL** is a data value and one or more information items that describe the data value and identify its data type. This class corresponds to ASN.1's External type, and thus the class and the attributes specific to it are described indirectly in the specification of ASN.1. (See Clause 34 of Recommendation X.208, "Specification of Abstract Syntax Notation 1 (ASN.1)," *CCITT Blue Book*, Fascicle VIII.4, International Telecommunications Union, 1988. Also published by ISO as *ISO 8824*.)

This class has the attributes of its superclass, *OM_C_OBJECT*, in addition to the OM attributes specific to this class that are listed in Table 19-2.

Table 19-2. Attributes Specific to OM_C_EXTERNAL

Attribute	Value Syntax	Value Length	Value Number	Value Initially
OM_ ARBITRARY_ ENCODING	String(OM_S_ BIT_STRING)	—	0 or 1 ¹	—
OM_ASN1_ ENCODING	String(OM_S_ ENCODING_ STRING)	—	0 or 1 ¹	—
OM_DATA_ VALUE_ DESCRIPTOR	String(OM_S_ OBJECT_ DESCRIPTOR_ STRING)	—	0 or 1	—
OM_DIRECT_ REFERENCE	String(OM_S_ OBJECT_ IDENTIFIER_ STRING)	—	0 or 1	—
OM_INDIRECT_ REFERENCE	OM_S_ INTEGER	—	0 or 1	—
OM_OCTET_ ALIGNED_ ENCODING	String(OM_S_ OCTET_ STRING)	—	0 or 1 ¹	—

¹ Only one of these three attributes is present.

- **OM_ARBITRARY_ENCODING**

This attribute is a representation of the data value as a bit string.

- **OM_ASN1_ENCODING**

The data value. This attribute can be present only if the data type is an ASN.1 type.

If this attribute value's syntax is an Object syntax, the data value's representation is that produced by **om_encode()** when its *Object* parameter is the attribute value and its *Rules* parameter is **ber**. Thus, the object's class must be one to which **om_encode()** applies.

- **OM_DATA_VALUE_DESCRIPTOR**

This attribute contains a description of the data value.

- **OM_DIRECT_REFERENCE**

This attribute contains a direct reference to the data type.

- **OM_INDIRECT_REFERENCE**

This attribute contains an indirect reference to the data type.

- **OM_OCTET_ALIGNED_ENCODING**

This attribute contains a representation of the data value as an octet string.

19.2.3 OM_C_OBJECT

The class *OM_C_OBJECT* represents information objects of any variety. This abstract class is distinguished by the fact that it has no superclass and that all other classes are its subclasses.

The attribute specific to this class is listed in Table 19-3.

Table 19–3. Attribute Specific to OM_C_OBJECT

Attribute	Value Syntax	Value Length	Value Number	Value Initially
OM_CLASS	String(OM_S_OBJECT_IDENTIFIER_STRING)	—	1	—

- **OM_CLASS**

This attribute identifies the object's class.

Index

A

- abstract OM class 5-25
- abstract service 10-1
- Abstract Service Definition . . . 6-16
- Abstract Syntax Notation 1 . . . 4-25
- abstract syntaxes 4-25
- access control 6-13
- acl.c 7-14
- acl.h
 - header file 7-35
- acl2.c 9-10
- acl2.h
 - header file 9-36
- address
 - O/R 14-24
- ADMD 14-14
- administration management domain 14-14
- administrative limit exceeded . 11-41
- alias entries 4-12
- API 10-4, 11-2
- approximate match 11-32
- ASN.1 19-3
 - abstract syntaxes 4-25
 - relating to Basic Encoding Rules 4-25
 - sample definition 4-22
 - simple types 4-26
 - transfer syntaxes 4-25
 - types 4-26
- attribute 11-8, 11-10
 - adding 11-24
 - domain-defined 14-22
 - error 11-9
 - list 11-10
 - matching rules 6-12
 - multi-valued 6-12
 - OM syntax 5-39
 - enumerated type 5-39
 - object type 5-41
 - string type 5-41
 - value 6-12
 - OM type 18-17
 - syntax template 5-39
- table 4-22
- type 11-8, 17-7, 18-17
 - describing 18-5
 - directory 4-5
 - mandatory 4-21
 - OM 5-2
 - optional 4-21
 - selected directory 12-2
- value 5-39
 - OM 11-9, 18-7
 - value length 6-12
- Attribute Value Assertion 4-11, 11-11
- authenticated bind 6-7
- authentication 6-13
- automatic connection management 6-9
- automatic continuation 11-18

AVA 4-11, 11-11

B

Basic Encoding Rules 4-27

BER 4-27, 10-9

bind

 authenticated 6-7

 credentials 6-8

Boolean 18-5, 18-18

C

C

 naming conventions 5-8

cache update process 4-33

canonical-ber 19-3

CCITT 19-4

chaining 4-31

chaining prohibited 11-16

character set

 local 18-6

character string 17-3, 18-15, 18-19

 length 18-19

 type 17-6

class

 abstract OM 5-25

 concrete OM 5-25

 OM 19-3, 19-1

 OM hierarchy 5-24

 OM inheritance 5-24

 OM object 5-23

closure

 package 5-35

common results 11-12

communications error 11-13

compare result 11-13

concrete OM class 5-25, 5-26

context 6-10, 10-1, 11-14

 common parameters 6-11

 GDS

 package 15-10

 local controls 6-11

 service controls 6-11

continuation reference 10-14, 11-19

controls

 service 15-12

D

DAP 4-29

DDA 14-22

default

 context 10-8, 11-19

 directory session 10-7

 session 10-7

descriptor list 5-11

 initializing 5-45

 OM_descriptor data structure 5-43

 representation of public object 5-11

DIB 4-4, 4-9

 structure of 4-4

directory 6-37

 access control 6-13

 alias entries 4-12

 attribute table 4-22

 attribute types 4-5, 4-21

 authentication 6-13

- automatic connection management . . . 6-9
 - building a distinguished name 5-14
 - class 14-28
 - class definitions 6-11
 - connection management functions . . . 6-1, 6-6
 - context 6-10
 - defining subclasses 4-24
 - distinguished name 4-9
 - example of entry 4-8
 - filter 6-27
 - GDS Standard Schema . . . 4-14
 - information model 4-3
 - modify operations 6-37
 - modifying entries 6-37
 - name verification 4-14
 - naming attributes 4-16
 - object entries 4-7
 - object identifiers 4-5
 - objects 4-3
 - OCT 4-17
 - operation functions 6-16
 - read operations 6-17
 - reading an entry 6-18
 - relationship between schemas and the DIT 4-24
 - relative distinguished name 4-11
 - search 6-31
 - search criteria 6-27
 - search operation 6-27
 - search operations 6-26
 - selected attribute types . . 6-11
 - selected object classes . . 6-11
 - service functions 6-1
 - service package 6-12
 - session 6-6
 - SRT 4-15
 - structure of the DIB 4-4
 - Directory Access Protocol . . . 4-29
 - Directory Information Base . . . 4-4
 - Directory Information Tree . . . 4-14
 - Directory Service Agent 4-28
 - Directory System Protocol . . . 4-30
 - Directory User Agent 4-29
 - distinguished encoding 10-9
 - distinguished name 4-9, 5-14
 - as a public object 5-14
 - example of distinguished name 4-9
 - relative 4-11
 - Distributed Management Environment 16-1
 - distribution list 14-1
 - DIT 4-9, 4-24
 - GDS Standard Schema . . 4-14
 - DL 14-1
 - DMD 11-18
 - DME support 16-1
 - domain management 14-22
 - domain-defined attribute . . . 14-22
 - DSA 4-28
 - address 11-7, 11-48
 - name 11-48
 - DSP 4-30
 - DUA 4-28
 - cache 4-33, 4-34, 6-15
 - cache" 4-36
- ## E
- EIT 14-5
 - elements 17-3
 - elements, string 18-14
 - encoded information type . . . 14-5
 - encoding 18-16

- entries 11-43, 11-44
- entry
 - modification 11-24
 - list 11-25
- Enum(*) 17-2
- enumerated type 5-39
- enumeration 18-7
- errors
 - directory service
 - attribute 11-9
 - communications 11-13
 - library 11-33
 - name 11-38
 - security 11-46
 - service 11-46
 - system 11-49
 - update 11-50
- example.c 7-2
- extensions 11-28
- external type 19-4

F

- facsimile telephone number 12-16
- filter 6-27, 6-31, 11-29
 - item 11-30
 - item type 11-32
 - type 11-30
- final substring 11-33
- from entry 11-22, 11-36

G

GDS

- as distributed service 4-28
- authenticated bind 4-40
- binding with credentials 4-40
- chaining 4-31
- DAP 4-29
- Directory User Agent cache 4-33
- DSA-DUA relationship 4-28
- DSP 4-30
- DUA 15-15
- DUA cache 4-33
- extension package 5-33, 6-15
- package 6-13
 - context 15-10
 - session 15-15
- referral 4-30
- security 4-40
- Standard Schema
 - attribute table 4-22
 - naming attributes 4-16
 - structured object classes4-16
- standard schema 4-14
 - OCT 4-17
 - SRT 4-15
- XDS API 4-2
- XOM API 4-2
- GDSP (GDS package) 15-15
- Global Directory Service 4-2

H

- header files

XDS API 7-35
 XOM API 5-56
 high priority 11-17

I

identifier 11-29
 information type 11-24
 encoded 14-5
 initial substring 11-33
 integers 18-16
 intermediate data type 18-3
 ISO 19-3, 19-4
 item 11-31

L

length, string 17-3
 length-unspecified 18-14
 limit problem 11-41
 list
 info 11-34, 11-37
 local scope 11-17
 low priority 11-17

M

management domain 14-22
 matched 11-14
 max outstanding operations 10-13

MD 14-22
 MDUP 14-1
 medium priority 11-17
 message handling system 14-1
 message store 14-5
 message transfer agent 14-1
 metacharacters 2-17
 in CDS 2-17
 in DNS 2-17
 in GDS 2-17

MHS
 directory user package 5-34, 14-1
 mnemonic O/R address 14-24
 modification type 11-24
 MS 14-5
 MTA 14-1

N

name 11-37
 maximum sizes 2-20
 resolution phase 11-39
 valid characters 2-14
 naming
 attributes 4-16
 rules 2-14
 network addresses 11-42
 no limit exceeded 11-41
 numeric O/R address 14-25

O

O/R 14-5, 14-27

- address 14-24
- mnemonic address . . . 14-24
- object
 - class attribute 4-7
 - class hierarchy 5-24
 - class table 4-17
 - directory 4-3
 - dynamically-defined static public
7-49
 - encoding OM 19-3
 - entries 4-7, 6-18
 - example of internal structure 5-3
 - identifier
 - directory 4-5
 - OCT 4-20
 - OM class 5-7
 - identifier, XDS package . . 4-6
 - management 5-2
 - name 11-9, 11-14, 11-22, 11-35,
11-45
 - OM class inheritance . . . 5-24
 - partially-defined static public7-47
 - predefined static public . . 7-46
 - private 18-9, 18-17
 - public 5-10, 18-5, 18-9, 18-17
 - representation of public object 5-
11
 - selected attribute types . . . 6-11
 - selected classes 6-11
 - subordinate 5-46, 11-35
 - type 5-41
 - value 5-3
- OCT 4-17
 - acronyms of super class . . 4-18
 - attributes 4-21
 - class inheritance 4-19
 - partial representation of . . 4-18
- OM
 - attribute types 5-2, 5-7
 - classes 5-23, 11-2, 19-3

- abstract 5-25
- concrete 5-25
- defining 5-28
- hierarchy 5-24
- inheritance 5-24
- initial value 5-29
- mapping class definition5-
6
- object identifier 5-7
- OM attribute 5-29
- value length 5-29
- value number 5-29
- value syntax 5-29
- objects 5-2
- syntax 5-2
- value syntax 6-12
- operation
 - directory service 10-1
 - not started 11-16, 11-40
 - progress 11-16, 11-20, 11-39
- optional functionality . . . 11-18, 11-49
- originator/recipient 14-5
- OSI
 - application contexts 12-12
 - application entity 11-8, 12-11
 - communications 11-8
 - presentation address 12-11

P

- package 5-31
 - basic directory contents 5-32, 6-4
 - closure 5-35
 - directory service 5-31, 6-4
 - ds_version 5-32
 - GDS 5-31, 6-4, 6-13

- GDS extension 5-33
- MHS directory user 5-31, 5-34, 6-4
- negotiating features 5-32
- service 11-2
- strong authentication 5-31, 5-33
- XDS 10-3
- partial outcome qualifier 11-35, 11-40, 11-45
- position
 - string 17-3
- postal address 12-17
- postal O/R address 14-25
- prefer chaining 11-17
- presentation
 - address 11-41
 - selector 11-42
- priority 11-17
- private management domain . 14-15
- private object 5-22
- PRMD 14-15
- public object 5-10
 - building for ds_search() . 6-31
 - client-generated 5-19
 - comparison with private objects 5-22
 - creating 6-20
 - dynamically-defined static 7-49
 - partially-defined static . . 7-47
 - predefined static 7-46
 - representation by using descriptor list 5-11
 - service-generated 5-19

R

- RDN 4-11, 11-20, 11-21
 - resolved 11-20
- read result 6-25, 11-42
- referral 4-30, 10-14, 11-43
- relative distinguished name . . 11-36
- relative name 11-43
- requestor 11-49
- return codes 18-14
 - service interface 18-23
- rules
 - OM object encoding 19-3

S

- search
 - criterion 12-17
 - guide 12-19
 - info 11-44
- selected attribute types 12-2
- service
 - controls 15-10, 15-12
 - interface data types 5-42
 - package 11-2
- service-generated descriptor . . 18-7
- session 10-1
 - default directory 10-7
 - directory 6-6
 - GDS
 - package 15-15
 - multiple concurrent 6-7
 - selector 11-42
- shadow update process 4-33
- size limit 11-18

SRT 4-15, 4-16
standards 10-1, 11-2
status
 directory 10-12
storage management 5-36
string 17-3, 18-18
 length 17-3, 18-14
 position 17-3
 type 5-41
string(*) 17-3
strings
 in directory service . . . 17-3
strong authentication package 5-31, 5-33
Structure Rule Table 4-15
structured O/R address . . . 14-25
structured object classes . . . 4-16
subclasses 12-14
substrings 11-32
superclasses 12-14
 OM 12-15, 13-5
syntax
 template 5-39, 17-2

T

target object 11-20
teldir.c 7-45
teletex terminal identifier . . . 12-20
telex number 12-21
terminal O/R address
 14-26
time limit 11-18
 exceeded . . . 11-18, 11-41
transfer syntax 4-25
types
 and values 11-24

U

unstructured O/R address . . 14-25

V

value
 OM attribute 18-7
 OM data 18-18

W

workspace 5-2, 5-35, 18-20

X

X.500
 directory information model 4-3
 naming concepts 4-9
X/Open
 Directory Service 4-2
 OSI-Abstract-Data Manipulation
 4-2
XDS 6-1
 API 4-2, 6-1
 acl.c 7-1
 bind credentials . . 6-8

- directory
 - connection management
 - functions . . . 6-6
- example.c 7-1
- GDS package . . . 6-13
- MHS directory user package . . . 6-4
- modifying entries . 6-37
- performing a read operation . . 6-23
- sample programs . 7-45
- teldir.c 7-45
- XDS interface management
 - functions 6-1, 6-2
- definitions 6-9
- directory read operations . 6-17
- dynamically-defined static public objects 7-49
- header files 7-35
- interface class definitions . 6-9
- interface management functions 6-1
- management functions . . . 6-2
- partially-defined static public objects 7-47
- predefined static public objects 7-46
- programming guidelines . 7-2
- sample programs 7-1
- XOM 5-1
- API 4-2
 - initializing descriptors 5-45
 - macros . . . 6-18, 6-29
 - OM class 5-23
 - OM class hierarchy 5-24
 - OM class inheritance 5-24
 - OM functions . . 5-48
 - workspace 5-2
- header files 5-56
- macros 5-56
- xom.h header file 5-56