

packet's destination IP NET\_ID in a routing table and forwarding based on the information in the table. But it is *routing protocols*, and *not* IP, that populate the routing tables with routing information. There are three routing protocols commonly associated with IP and the Internet, namely, RIP, OSPF, and BGP.

OSPF and RIP are primarily used to provide routing within a particular domain, such as within a corporate network or within an ISP's network. Since the routing is *inside* of the domain, these protocols are generically referred to as *interior gateways protocols*.

The Routing Information Protocol version 2 (RIP-2), described in [RFC 2453](#), describes how routers will exchange routing table information using a distance-vector algorithm. With RIP, neighboring routers periodically exchange their entire routing tables. RIP uses hop count as the metric of a path's cost, and a path is limited to 16 hops. Unfortunately, RIP has become increasingly inefficient on the Internet as the network continues its fast rate of growth. Current routing protocols for many of today's LANs are based upon RIP, including those associated with NetWare, AppleTalk, VINES, and DECnet. The IANA maintains a list of [RIP message types](#).

The Open Shortest Path First (OSPF) protocol is a link state routing algorithm that is more robust than RIP, converges faster, requires less network bandwidth, and is better able to scale to larger networks. With OSPF, a router broadcasts only changes in its links' status rather than entire routing tables. OSPF Version 2, described in [RFC 1583](#), is rapidly replacing RIP in the Internet.

The Border Gateway Protocol version 4 (BGP-4) is an *exterior gateway protocol* because it is used to provide routing information between Internet routing domains. BGP is a distance vector protocol, like RIP, but unlike almost all other distance vector protocols, BGP tables store the actual route to the destination network. BGP-4 also supports policy-based routing, which allows a network's administrator to create routing policies based on political, security, legal, or economic issues rather than technical ones. BGP-4 also supports CIDR. BGP-4 is described in [RFC 1771](#), while [RFC 1268](#) describes use of BGP in the Internet. In addition, the IANA maintains a list of [BGP parameters](#).

Figure 1 shows the protocol relationship of RIP, OSPF, and BGP to IP. A RIP message is carried in a UDP datagram which, in turn, is carried in an IP packet. An OSPF message, on the other hand, is carried directly in an IP datagram. BGP messages, in a total departure, are carried in TCP segments over IP. Although all of the TCP/IP books mentioned above discuss IP routing to some level of detail, *Routing in the Internet* by Christian Huitema is one of the best available references on this specific subject.

### 3.2.5. ICMP

The Internet Control Message Protocol, described in [RFC 792](#), is an adjunct to IP that notifies the sender of IP datagram about abnormal events. This collateral protocol is particularly important in the connectionless environment of IP.

The commonly employed ICMP message types include:

- *Destination Unreachable*: Indicates that a packet cannot be delivered because the destination host cannot be reached. The reason for the non-delivery may be that the host or network is unreachable or unknown, the protocol or port is unknown or unusable, fragmentation is required but not allowed (DF-flag is set), or the network or host is unreachable for this type of service.
- *Echo and Echo Reply*: These two messages are used to check whether hosts are reachable on the network. One host sends an Echo message to the other, optionally containing some data, and the receiving host responds with an Echo Reply containing the same data. These messages are the basis for the Ping command.
- *Parameter Problem*: Indicates that a router or host encountered a problem with some aspect of the packet's

Header.

- *Redirect*: Used by a host or router to let the sending host know that packets should be forwarded to another address. *For security reasons, Redirect messages should usually be blocked at the firewall.*
- *Source Quench*: Sent by a router to indicate that it is experiencing congestion (usually due to limited buffer space and is discarding datagrams).
- *TTL Exceeded*: Indicates that a datagram has been discarded because the TTL field reached 0 or because the entire packet was not received before the fragmentation timer expired.
- *Timestamp and Timestamp Reply*: These messages are similar to the Echo messages, but place a timestamp (with millisecond granularity) in the message, yielding a measure of how long remote systems spend buffering and processing datagrams, and providing a mechanism so that hosts can synchronize their clocks.

ICMP messages are carried in IP packets. The IANA maintains a complete list of [ICMP parameters](#).

### 3.2.6. IP version 6

The official version of IP that has been in use since the early 1980s is *version 4*. Due to the tremendous growth of the Internet and new emerging applications, it was recognized that a new version of IP was becoming necessary. In late 1995, IP version 6 (IPv6) was entered into the Internet Standards Track. The primary description of IPv6 is contained in [RFC 1883](#) and a number of related specifications, including [ICMPv6](#).

IPv6 is designed as an evolution from IPv4, rather than a radical change. Primary areas of change relate to:

- Increasing the IP address size to 128 bits
- Better support for traffic types with different quality-of-service objectives
- Extensions to support authentication, data integrity, and data confidentiality

For more information about IPv6, check out:


- [IPng: Internet Protocol Next Generation](#) by Scott Bradner and Allison Mankin (Addison-Wesley, 1996)
- [IPv6: The New Internet Protocol](#) by Christian Huitema (Prentice-Hall, 1996).
- ["IPv6: The Next Generation Internet Protocol"](#) by Gary Kessler.
- [IPng and the TCP/IP Protocols](#) by Stephen Thomas (John Wiley & Sons, 1996)
- [IPng Working Group page \(IETF\)](#)
- [IP Next Generation Web Page \(Sun\)](#)
- [6bone Web Page \(LBL\)](#)

### 3.3. The Transport Layer Protocols

The TCP/IP protocol suite comprises two protocols that correspond roughly to the OSI Transport and Session Layers; these protocols are called the Transmission Control Protocol and the User Datagram Protocol (UDP). One can argue that it is a misnomer to refer to "TCP/IP applications," as most such applications actually run over TCP or UDP, as shown in [Figure 1](#).

Higher-layer applications are referred to by a port identifier in TCP/UDP messages. The port identifier and IP address together form a *socket*, and the end-to-end communication between two hosts is uniquely identified on the Internet by the four-tuple (source port, source address, destination port, destination address). *Well-known port numbers* denote the server side of a connection and include:





**FIGURE 4. TCP segment format.**

The TCP data unit is called a *segment*; the name is due to the fact that TCP does not recognize messages, per se, but merely sends a block of bytes from the byte stream between sender and receiver. The fields of the segment (Figure 4) are:

- **Source Port and Destination Port:** Identify the source and destination ports to identify the end-to-end connection and higher-layer application.
- **Sequence Number:** Contains the sequence number of this segment's first data byte in the overall connection byte stream; since the sequence number refers to a byte count rather than a segment count, sequence numbers in contiguous TCP segments are not numbered sequentially.
- **Acknowledgment Number:** Used by the sender to acknowledge receipt of data; this field indicates the sequence number of the next byte expected from the receiver.
- **Data Offset:** Points to the first data byte in this segment; this field, then, indicates the segment header length.
- **Control Flags:** A set of flags that control certain aspects of the TCP virtual connection. The flags include:
  - **Urgent Pointer Field Significant (URG):** When set, indicates that the current segment contains urgent (or high-priority) data and that the Urgent Pointer field value is valid.
  - **Acknowledgment Field Significant (ACK):** When set, indicates that the value contained in the Acknowledgment Number field is valid. This bit is usually set, except during the first message during connection establishment.
  - **Push Function (PSH):** Used when the transmitting application wants to force TCP to immediately transmit the data that is currently buffered without waiting for the buffer to fill; useful for transmitting small units of data.
  - **Reset Connection (RST):** When set, immediately terminates the end-to-end TCP connection.
  - **Synchronize Sequence Numbers (SYN):** Set in the initial segments used to establish a connection, indicating that the segments carry the initial sequence number.
  - **Finish (FIN):** Set to request normal termination of the TCP connection in the direction this segment is traveling; completely closing the connection requires one FIN segment in each direction.
- **Window:** Used for flow control, contains the value of the *receive window size* which is the number of transmitted bytes that the sender of this segment is willing to accept from the receiver.
- **Checksum:** Provides rudimentary bit error detection for the segment (including the header and data).
- **Urgent Pointer:** Urgent data is information that has been marked as high-priority by a higher layer application; this data, in turn, usually bypasses normal TCP buffering and is placed in a segment between the header and "normal" data. The Urgent Pointer, valid when the URG flag is set, indicates the position of the first octet of nonexpedited data in the segment.
- **Options:** Used at connection establishment to negotiate a variety of options; maximum segment size (MSS) is the most commonly used option and, if absent, defaults to an MSS of 536. The IANA maintains a list of all TCP Option Numbers.

### 3.3.2. UDP

UDP, described in [RFC 768](#), provides an end-to-end datagram (connectionless) service. Some applications, such as those that involve a simple query and response, are better suited to the datagram service of UDP because there is no time lost to virtual circuit establishment and termination. UDP's primary function is to add a port number to the IP address to provide a socket for the application.




FIGURE 5. UDP datagram format.

The fields of a UDP datagram (Figure 5) are:

- **Source Port:** Identifies the UDP port at the source side of the connection; use of this field is optional in UDP and may be set to 0.
- **Destination Port:** Identifies the destination port of the end-to-end connection.
- **Length:** Indicates the total length of the UDP datagram.
- **Checksum:** Provides rudimentary bit error detection for the datagram (including the header and data).

### 3.4. Applications

The TCP/IP Application Layer protocols support the applications and utilities that are the Internet. Commonly used protocols include:


- **Telnet:** Short for *Telecommunication Network*, a virtual terminal protocol allowing a user logged on to one TCP/IP host to access other hosts on the network ([RFC.854](#)).

- **FTP:** The File Transfer Protocol allows a user to transfer files between local and remote host computers ([RFC 959](#)).
- **Archie:** A utility that allows a user to search all registered anonymous FTP sites for files on a specified topic.
- **Gopher:** A tool that allows users to search through data repositories using a menu-driven, hierarchical interface, with links to other sites ([RFC 1436](#)).
- **SMTP:** The Simple Mail Transfer Protocol is the standard protocol for the exchange of electronic mail over the Internet ([RFC 821](#)). SMTP is used between e-mail servers on the Internet or to allow an e-mail client to send mail to a server. [RFC 822](#) specifically describes the mail message body format, and [RFCs 1521](#) and [1522](#) describe MIME (Multipurpose Internet Mail Extensions). Reference books on electronic mail systems include [!%@:: Addressing and Networks](#) by D. Frey and R. Adams (O'Reilly & Associates, 1993) and [THE INTERNET MESSAGE: Closing the Book With Electronic Mail](#) by M. Rose (PTR Prentice Hall, 1993).
- **HTTP:** The Hypertext Transfer Protocol is the basis for exchange of information over the World Wide Web (WWW). Various versions of HTTP are in use over the Internet, with HTTP version 1.0 ([RFC 1945](#)) being the most current. WWW pages are written in the Hypertext Markup Language (HTML), an ASCII-based, platform-independent formatting language ([RFC 1866](#)).
- **Finger:** Used to determine the status of other hosts and/or users ([RFC 1288](#)).
- **POP:** The Post Office Protocol defines a simple interface between a user's mail client software and an e-mail server; POP is used to download mail from the server to the client and allows the user to manage their mailboxes. The current version is POP3 ([RFC 1460](#)).
- **DNS:** The Domain Name System (described in slightly more detail in [Section 3.2.2](#) above) defines the structure of Internet names and their association with IP addresses, as well as the association of mail and name servers with domains.
- **SNMP:** The Simple Network Management Protocol defines procedures and management information databases for managing TCP/IP-based network devices. SNMP ([RFC 1157](#)) is widely deployed in local and wide area network. SNMP Version 2 (SNMPv2, [RFC 1441](#)) adds security mechanisms that are missing in SNMP, but is also very complex; widespread use of SNMPv2 has yet to be seen. Additional information on SNMP and TCP/IP-based network management can be found in [SNMP](#) by S. Feit (McGraw-Hill, 1994) and [THE SIMPLE BOOK: An Introduction to Internet Management](#), 2/e, by M. Rose (PTR Prentice Hall, 1994).
- **Ping:** The Packet Internet Groper, a utility that allows a user at one system to determine the status of other hosts and the latency in getting a message to that host. Uses ICMP Echo messages.
- **Whois/NICNAME:** Utilities that search databases for information about Internet domains and domain contact information ([RFC 954](#)).
- **Traceroute:** A tool that displays the route that packets will take when traveling to a remote host.

A guide to using most of these applications can be found in "A Primer on Internet and TCP/IP Tools and Utilities" (FYI 30/[RFC 2151](#)) by Gary Kessler & Steve Shepard (also available in [HTML](#), [Postscript](#), and [Word](#)).

### 3.5. Summary

As this discussion has shown, *TCP/IP* is not merely a pair of communication protocols but is a suite of protocols, applications, and utilities. Increasingly, these protocols are referred to as the *Internet Protocol Suite*, but the older name will not disappear anytime soon.






FIGURE 6. TCP/IP protocol suite architecture.

Figure 6 shows the relationship between the various protocol layers of TCP/IP. Applications and utilities reside in host, o end-communicating, systems. TCP provides a reliable, virtual circuit connection between the two hosts. (UDP, not shown, provides an end-to-end datagram connection at this layer.) IP provides a datagram (DG) transport service over any intervening subnetworks, including local and wide area networks. The underlying subnetwork may employ nearly any common local or wide area network technology.

Note that the term *gateway* is used for the device interconnecting the two subnets, a device usually called a *router* in LAN environments or *intermediate system* in OSI environments. In OSI terminology, a *gateway* is used to provide protocol conversion between two networks and/or applications.

4. Other Information Sources

This memo has only provided background information about the TCP/IP protocols and the Internet. There is a wide rang of additional information that the reader can access to further use and understand the tools and scope of the Internet. The real fun begins now!

Internet specifications, standards, reports, humor, and tutorials are distributed as Request for Comments (RFC) documents. RFCs are all freely available on-line, and most are available in ASCII text format.

Internet standards are documented in a subset of the RFCs, identified with an "STD" designation. RFC 2026 describes the Internet standards process and STD 1 always contains the official list of Internet standards.

For Your Information (FYI) documents are another RFC subset, specifically providing background information for the Internet community. The FYI notes are described in RFC 1150.

Frequently Asked Question (FAQ) lists may be found for a number of topics, ranging from ISDN and cryptography to the Internet and Gopher. Two such FAQs are of particular interest to Internet users: "FYI on Questions and Answers - Answers to Commonly Asked 'New Internet User' Questions" (RFC 1594) and "FYI on Questions and Answers: Answers to Commonly Asked 'Experienced Internet User' Questions" (RFC 1207). All three of these documents point to even more information sources.

### 5. Acronyms and Abbreviations

ARP	Address Resolution Protocol
ARPANET	Advanced Research Projects Agency Network
ASCII	American Standard Code for Information Interchange
ATM	Asynchronous Transfer Mode
BGP	Border Gateway Protocol
BSD	Berkeley Software Development
CCITT	International Telegraph and Telephone Consultative Committee
CIX	Commercial Internet Exchange
DARPA	Defense Advanced Research Projects Agency
DNS	Domain Name System
DoD	U.S. Department of Defense
FAQ	Frequently Asked Questions lists
FDDI	Fiber Distributed Data Interface
FTP	File Transfer Protocol
FYI	For Your Information series of RFCs
GOSIP	U.S. Government Open Systems Interconnection Profile
HTML	Hypertext Markup Language
HTTP	Hypertext Transfer Protocol
IAB	Internet Activities Board
IANA	Internet Assigned Numbers Authority
ICMP	Internet Control Message Protocol
IESG	Internet Engineering Steering Group
IETF	Internet Engineering Task Force
IP	Internet Protocol
ISO	International Organization for Standardization
ISOC	Internet Society
ITU-T	International Telecommunication Union Telecommunication Standardization Sector
MAC	Medium (or media) access control
Mbps	Megabits (millions of bits) per second
NICNAME	Network Information Center name service
NSF	National Science Foundation
NSFNET	National Science Foundation Network



OSI	Open Systems Interconnection
OSPF	Open Shortest Path First
PPP	Point-to-Point Protocol
RARP	Reverse Address Resolution Protocol
RIP	Routing Information Protocol
RFC	Request For Comments
SLIP	Serial Line IP
SMDS	Switched Multimegabit Data Service
SMTP	Simple Mail Transfer Protocol
SNMP	Simple Network Management Protocol
STD	Internet Standards series of RFCs
TCP	Transmission Control Protocol
TLD	Top-level domain
UDP	User Datagram Protocol

#### 6. Author's Address

Gary C. Kessler  
Hill Associates  
106 Highpoint Center  
Colchester, VT 05446

+1 802-879-3375 (home office)  
+1 630-604-5529 (fax)

E-mail: [kumquat@hill.com](mailto:kumquat@hill.com) or [kumquat@sover.net](mailto:kumquat@sover.net)  
<http://www.hill.com> or <http://www.sover.net/~kessfam>

Copyright © 1995, 1996, 1997, 1998, 1999, 2000, 2001 Hill Associates, Inc. All Rights Reserved.

Questions or Comments? [Send us your feedback!](#)  
[Home | About Us | Education & Training | Professional Services](#)  
[E-store | Our People | Staff Publications | News | Careers](#)

---

# The Common Object Request Broker: Architecture and Specification

---

---

Revision 2.6  
December 2001

---

---

Copyright 1998, 1999, Alcatel  
Copyright 1997, 1998, 1999 BEA Systems, Inc.  
Copyright 1995, 1996 BNR Europe Ltd.  
Copyright 1998, Borland International  
Copyright 1998, Cooperative Research Centre for Distributed Systems Technology (DSTC Pty Ltd)  
Copyright 2001, Concept Five Technologies  
Copyright 1991, 1992, 1995, 1996, Digital Equipment Corporation  
Copyright 2001, Eternal Systems, Inc.  
Copyright 1995, 1996, 1998, Expersoft Corporation  
Copyright 1996, 1997 FUJITSU LIMITED  
Copyright 1996, Genesis Development Corporation  
Copyright 1989- 2001, Hewlett-Packard Company  
Copyright 2001, HighComm  
Copyright 1998, 1999, Highlander Communications, L.C.  
Copyright 1991, 1992, 1995, 1996 HyperDesk Corporation  
Copyright 1998, 1999, Inprise Corporation  
Copyright 1996 - 2001, International Business Machines Corporation  
Copyright 1995, 1996 ICL, plc  
Copyright 1998 - 2001, Inprise Corporation  
Copyright 1998, International Computers, Ltd.  
Copyright 1995 - 2001, IONA Technologies, Ltd.  
Copyright 1998 - 2001, Lockheed Martin Federal Systems, Inc.  
Copyright 1998, 1999, 2001, Lucent Technologies, Inc.  
Copyright 1996, 1997 Micro Focus Limited  
Copyright 1991, 1992, 1995, 1996 NCR Corporation  
Copyright 1998, NEC Corporation  
Copyright 1998, Netscape Communications Corporation  
Copyright 1998, 1999, Nortel Networks  
Copyright 1998, 1999, Northern Telecom Corporation  
Copyright 1995, 1996, 1998, Novell USG  
Copyright 1991, 1992, 1995, 1996 by Object Design, Inc.  
Copyright 1991- 2001 Object Management Group, Inc.  
Copyright 1998, 1999, 2001, Objective Interface Systems, Inc.  
Copyright 1998, 1999, Object-Oriented Concepts, Inc.  
Copyright 1998, 2001, Oracle Corporation  
Copyright 1998, PeerLogic, Inc.  
Copyright 1996, Siemens Nixdorf Informationssysteme AG  
Copyright 1991 - 2001, Sun Microsystems, Inc.  
Copyright 1995, 1996, SunSoft, Inc.  
Copyright 1996, Sybase, Inc.  
Copyright 1998, Telefónica Investigación y Desarrollo S.A. Unipersonal  
Copyright 1998, TIBCO, Inc.  
Copyright 1998, 1999, Tri-Pacific Software, Inc.  
Copyright 1996, Visual Edge Software, Ltd.

The companies listed above have granted to the Object Management Group, Inc. (OMG) a nonexclusive, royalty-free, paid up, worldwide license to copy and distribute this document and to modify this document and distribute copies of the modified version. Each of the copyright holders listed above has agreed that no person shall be deemed to have infringed the copyright in the included material of any such copyright holder by reason of having used the specification set forth herein or having conformed any computer software to the specification.

---

## PATENT

The attention of adopters is directed to the possibility that compliance with or adoption of OMG specifications may require use of an invention covered by patent rights. OMG shall not be responsible for identifying patents for which a license may be required by any OMG specification, or for conducting legal inquiries into the legal validity or scope of those patents that are brought to its attention. OMG specifications are prospective and advisory only. Prospective users are responsible for protecting themselves against liability for infringement of patents.

## NOTICE

The information contained in this document is subject to change without notice. The material in this document details an Object Management Group specification in accordance with the license and notices set forth on this page. This document does not represent a commitment to implement any portion of this specification in any company's products.

WHILE THE INFORMATION IN THIS PUBLICATION IS BELIEVED TO BE ACCURATE, THE OBJECT MANAGEMENT GROUP AND THE COMPANIES LISTED ABOVE MAKE NO WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, WITH REGARD TO THIS MATERIAL INCLUDING, BUT NOT LIMITED TO ANY WARRANTY OF TITLE OR OWNERSHIP, IMPLIED WARRANTY OF MERCHANTABILITY OR WARRANTY OF FITNESS FOR PARTICULAR PURPOSE OR USE. In no event shall The Object Management Group or any of the companies listed above be liable for errors contained herein or for indirect, incidental, special, consequential, reliance or cover damages, including loss of profits, revenue, data or use, incurred by any user or any third party. The copyright holders listed above acknowledge that the Object Management Group (acting itself or through its designees) is and shall at all times be the sole entity that may authorize developers, suppliers and sellers of computer software to use certification marks, trademarks or other special designations to indicate compliance with these materials. This document contains information which is protected by copyright. All Rights Reserved. No part of this work covered by copyright herein may be reproduced or used in any form or by any means--graphic, electronic, or mechanical, including photocopying, recording, taping, or information storage and retrieval systems--without permission of the copyright owner.

RESTRICTED RIGHTS LEGEND. Use, duplication, or disclosure by government is subject to restrictions as set forth in subdivision (c) (1) (ii) of the Right in Technical Data and Computer Software Clause at DFARS 252.227.7013 OMG<sup>®</sup> and Object Management are registered trademarks of the Object Management Group, Inc. Object Request Broker, OMG IDL, ORB, CORBA, CORBAfacilities, CORBAservices, COSS, and IIOP are trademarks of the Object Management Group, Inc. X/Open is a trademark of X/Open Company Ltd.

## ISSUE REPORTING

All OMG specifications are subject to continuous review and improvement. As part of this process we encourage readers to report any ambiguities, inconsistencies, or inaccuracies they may find by completing the Issue Reporting Form listed on the main web page <http://www.omg.org>, under Documents & Specifications, Report a Bug/Issue.



# Contents

---

<b>Preface</b> .....	<b>xxxvii</b>
<b>1. The Object Model</b> .....	<b>1-1</b>
1.1 Overview .....	1-1
1.2 Object Semantics .....	1-2
1.2.1 Objects .....	1-2
1.2.2 Requests .....	1-3
1.2.3 Object Creation and Destruction .....	1-4
1.2.4 Types .....	1-4
1.2.4.1 Basic types .....	1-4
1.2.4.2 Constructed types .....	1-5
1.2.5 Interfaces .....	1-6
1.2.6 Value Types .....	1-6
1.2.7 Abstract Interfaces .....	1-7
1.2.8 Operations .....	1-7
1.2.8.1 Parameters .....	1-8
1.2.8.2 Return Result .....	1-8
1.2.8.3 Exceptions .....	1-8
1.2.8.4 Contexts .....	1-8
1.2.8.5 Execution Semantics .....	1-8
1.2.9 Attributes .....	1-9
1.3 Object Implementation .....	1-9
1.3.1 The Execution Model: Performing Services .....	1-9
1.3.2 The Construction Model .....	1-10
<b>2. CORBA Overview</b> .....	<b>2-1</b>
2.1 Structure of an Object Request Broker .....	2-1
2.1.1 Object Request Broker .....	2-6
2.1.2 Clients .....	2-7
2.1.3 Object Implementations .....	2-7
2.1.4 Object References .....	2-8

# Contents

---

2.1.5	OMG Interface Definition Language .....	2-8
2.1.6	Mapping of OMG IDL to Programming Languages .....	2-8
2.1.7	Client Stubs .....	2-9
2.1.8	Dynamic Invocation Interface .....	2-9
2.1.9	Implementation Skeleton .....	2-9
2.1.10	Dynamic Skeleton Interface .....	2-10
2.1.11	Object Adapters .....	2-10
2.1.12	ORB Interface .....	2-10
2.1.13	Interface Repository .....	2-11
2.1.14	Implementation Repository .....	2-11
2.2	Example ORBs .....	2-11
2.2.1	Client- and Implementation-resident ORB .....	2-11
2.2.2	Server-based ORB .....	2-12
2.2.3	System-based ORB .....	2-12
2.2.4	Library-based ORB .....	2-12
2.3	Structure of a Client .....	2-12
2.4	Structure of an Object Implementation .....	2-13
2.5	Structure of an Object Adapter .....	2-15
2.6	CORBA Required Object Adapter .....	2-17
2.6.1	Portable Object Adapter .....	2-17
2.7	The Integration of Foreign Object Systems .....	2-17
3.	<b>OMG IDL Syntax and Semantics .....</b>	<b>3-1</b>
3.1	Overview .....	3-2
3.2	Lexical Conventions .....	3-3
3.2.1	Tokens .....	3-5
3.2.2	Comments .....	3-6
3.2.3	Identifiers .....	3-6
3.2.3.1	Escaped Identifiers .....	3-6
3.2.4	Keywords .....	3-7
3.2.5	Literals .....	3-8
3.2.5.1	Integer Literals .....	3-8
3.2.5.2	Character Literals .....	3-9
3.2.5.3	Floating-point Literals .....	3-10
3.2.5.4	String Literals .....	3-10
3.2.5.5	Fixed-Point Literals .....	3-11
3.3	Preprocessing .....	3-11
3.4	OMG IDL Grammar .....	3-12
3.5	OMG IDL Specification .....	3-16
3.6	Module Declaration .....	3-17
3.7	Interface Declaration .....	3-17
3.7.1	Interface Header .....	3-17
3.7.2	Interface Inheritance Specification .....	3-18
3.7.3	Interface Body .....	3-18

# Contents

	<b>3.7.4 Forward Declaration</b> .....	<b>3-19</b>
	<b>3.7.5 Interface Inheritance</b> .....	<b>3-19</b>
<b>3.8</b>	<b>Value Declaration</b> .....	<b>3-24</b>
	<b>3.8.1 Regular Value Type</b> .....	<b>3-24</b>
	3.8.1.1 Value Header .....	3-24
	3.8.1.2 Value Element .....	3-25
	3.8.1.3 Value Inheritance Specification .....	3-25
	3.8.1.4 State Members .....	3-25
	3.8.1.5 Initializers .....	3-26
	3.8.1.6 Value Type Example .....	3-26
	<b>3.8.2 Boxed Value Type</b> .....	<b>3-26</b>
	<b>3.8.3 Abstract Value Type</b> .....	<b>3-27</b>
	<b>3.8.4 Value Forward Declaration</b> .....	<b>3-28</b>
	<b>3.8.5 Valuetype Inheritance</b> .....	<b>3-28</b>
<b>3.9</b>	<b>Constant Declaration</b> .....	<b>3-29</b>
	<b>3.9.1 Syntax</b> .....	<b>3-29</b>
	<b>3.9.2 Semantics</b> .....	<b>3-30</b>
<b>3.10</b>	<b>Type Declaration</b> .....	<b>3-33</b>
	<b>3.10.1 Basic Types</b> .....	<b>3-34</b>
	3.10.1.1 Integer Types .....	3-35
	3.10.1.2 Floating-Point Types .....	3-36
	3.10.1.3 Char Type .....	3-36
	3.10.1.4 Wide Char Type .....	3-36
	3.10.1.5 Boolean Type .....	3-36
	3.10.1.6 Octet Type .....	3-36
	3.10.1.7 Any Type .....	3-37
	<b>3.10.2 Constructed Types</b> .....	<b>3-37</b>
	3.10.2.1 Structures .....	3-37
	3.10.2.2 Discriminated Unions .....	3-37
	3.10.2.3 Constructed Recursive Types and IForward Declarations .....	3-39
	3.10.2.4 Enumerations .....	3-41
	<b>3.10.3 Template Types</b> .....	<b>3-41</b>
	3.10.3.1 Sequences .....	3-41
	3.10.3.2 Strings .....	3-42
	3.10.3.3 Wstrings .....	3-42
	3.10.3.4 Fixed Type .....	3-43
	<b>3.10.4 Complex Declarator</b> .....	<b>3-43</b>
	3.10.4.1 Arrays .....	3-43
	<b>3.10.5 Native Types</b> .....	<b>3-43</b>
<b>3.11</b>	<b>Exception Declaration</b> .....	<b>3-47</b>
<b>3.12</b>	<b>Operation Declaration</b> .....	<b>3-47</b>
	<b>3.12.1 Operation Attribute</b> .....	<b>3-48</b>
	<b>3.12.2 Parameter Declarations</b> .....	<b>3-48</b>
	<b>3.12.3 Raises Expressions</b> .....	<b>3-49</b>
	<b>3.12.4 Context Expressions</b> .....	<b>3-49</b>
<b>3.13</b>	<b>Attribute Declaration</b> .....	<b>3-50</b>
<b>3.14</b>	<b>CORBA Module</b> .....	<b>3-51</b>



# Contents

---

3.15	Names and Scoping	3-52
3.15.1	Qualified Names	3-52
3.15.2	Scoping Rules and Name Resolution	3-54
3.15.3	Special Scoping Rules for Type Names	3-57
4.	<b>ORB Interface</b>	<b>4-1</b>
4.1	Overview	4-1
4.2	The ORB Operations	4-2
4.2.1	<b>ORB Identity</b>	<b>4-7</b>
4.2.1.1	id	4-7
4.2.2	<b>Converting Object References to Strings</b>	<b>4-8</b>
4.2.2.1	object_to_string	4-8
4.2.2.2	string_to_object	4-8
4.2.3	<b>Getting Service Information</b>	<b>4-8</b>
4.2.3.1	get_service_information	4-8
4.2.4	<b>Thread-Related Operations</b>	<b>4-9</b>
4.2.4.1	work_pending	4-9
4.2.4.2	perform_work	4-9
4.2.4.3	run	4-10
4.2.4.4	shutdown	4-10
4.2.4.5	destroy	4-11
4.3	Object Reference Operations	4-12
4.3.1	<b>Determining the Object Interface</b>	<b>4-13</b>
4.3.1.1	get_interface	4-13
4.3.2	<b>Duplicating and Releasing Copies of Object References</b>	<b>4-14</b>
4.3.2.1	duplicate	4-14
4.3.2.2	release	4-14
4.3.3	<b>Nil Object References</b>	<b>4-14</b>
4.3.3.1	is_nil	4-14
4.3.4	<b>Equivalence Checking Operation</b>	<b>4-15</b>
4.3.4.1	is_a	4-15
4.3.5	<b>Probing for Object Non-Existence</b>	<b>4-15</b>
4.3.5.1	non_existent	4-15
4.3.6	<b>Object Reference Identity</b>	<b>4-16</b>
4.3.6.1	Hashing Object Identifiers	4-16
4.3.6.2	Equivalence Testing	4-16
4.3.7	<b>Type Coercion Considerations</b>	<b>4-17</b>
4.3.8	<b>Getting Policy Associated with the Object</b>	<b>4-17</b>
4.3.8.1	get_policy	4-17
4.3.8.2	get_client_policy	4-18
4.3.8.3	get_policy_overrides	4-19
4.3.9	<b>Overriding Associated Policies on an Object Reference</b>	<b>4-19</b>
4.3.9.1	set_policy_overrides	4-19
4.3.10	<b>Validating Connection</b>	<b>4-20</b>
4.3.10.1	validate_connection	4-20
4.3.11	<b>Getting the Domain Managers Associated with the Object</b>	<b>4-20</b>
4.3.11.1	get_domain_managers	4-20
4.4	ValueBase Operations	4-21

4.5	ORB and OA Initialization and Initial References . . . . .	4-21
4.5.1	ORB Initialization . . . . .	4-22
4.5.2	Obtaining Initial Object References . . . . .	4-23
4.5.3	Configuring Initial Service References . . . . .	4-26
4.5.3.1	ORB-specific Configuration . . . . .	4-26
4.5.3.2	ORBInitRef . . . . .	4-26
4.5.3.3	ORBDefaultInitRef . . . . .	4-27
4.5.3.4	Configuration Effect on resolve_initial_references . . . . .	4-27
4.5.3.5	Configuration Effect on list_initial_services	4-28
4.6	Context Object . . . . .	4-28
4.6.1	Introduction . . . . .	4-28
4.6.2	Context Object Operations . . . . .	4-29
4.6.2.1	get_default_context . . . . .	4-30
4.6.2.2	set_one_value . . . . .	4-30
4.6.2.3	set_values . . . . .	4-30
4.6.2.4	get_values . . . . .	4-31
4.6.2.5	delete_values . . . . .	4-31
4.6.2.6	create_child . . . . .	4-32
4.6.2.7	delete . . . . .	4-32
4.7	Current Object . . . . .	4-32
4.8	Policy Object . . . . .	4-33
4.8.1	Definition of Policy Object . . . . .	4-33
4.8.1.1	Copy . . . . .	4-34
4.8.1.2	Destroy . . . . .	4-34
4.8.1.3	Policy_type . . . . .	4-34
4.8.2	Creation of Policy Objects . . . . .	4-34
4.8.2.1	PolicyErrorCode . . . . .	4-35
4.8.2.2	PolicyError . . . . .	4-35
4.8.2.3	Create_policy . . . . .	4-35
4.8.3	Usages of Policy Objects . . . . .	4-36
4.8.4	Policy Associated with the Execution Environment	4-37
4.8.5	Specification of New Policy Objects . . . . .	4-37
4.8.6	Standard Policies . . . . .	4-39
4.9	Management of Policies . . . . .	4-43
4.9.1	Client Side Policy Management . . . . .	4-43
4.9.2	Server Side Policy Management . . . . .	4-43
4.9.3	Policy Management Interfaces . . . . .	4-44
4.9.3.1	interface PolicyManager . . . . .	4-44
4.9.3.2	interface PolicyCurrent . . . . .	4-46
4.10	Management of Policy Domains . . . . .	4-46
4.10.1	Basic Concepts . . . . .	4-46
4.10.1.1	Policy Domain . . . . .	4-46
4.10.1.2	Policy Domain Manager . . . . .	4-47
4.10.1.3	Policy Objects . . . . .	4-47
4.10.1.4	Object Membership of Policy Domains	4-47
4.10.1.5	Domains Association at Object Reference Creation . . . . .	4-48
4.10.1.6	Implementor's View of Object Creation	4-48
4.10.2	Domain Management Operations . . . . .	4-49

# Contents

---

	4.10.2.7 Domain Manager .....	4-50
	4.10.2.8 Construction Policy .....	4-51
4.11	TypeCodes .....	4-51
4.11.1	The TypeCode Interface .....	4-52
4.11.2	TypeCode Constants .....	4-56
4.11.3	Creating TypeCodes .....	4-57
4.12	Exceptions .....	4-61
4.12.1	Definition of Terms .....	4-61
4.12.2	System Exceptions .....	4-62
4.12.3	Standard System Exception Definitions .....	4-63
4.12.3.1	UNKNOWN .....	4-65
4.12.3.2	BAD_PARAM .....	4-65
4.12.3.3	NO_MEMORY .....	4-65
4.12.3.4	IMP_LIMIT .....	4-66
4.12.3.5	COMM_FAILURE .....	4-66
4.12.3.6	INV_OBJREF .....	4-66
4.12.3.7	NO_PERMISSION .....	4-66
4.12.3.8	INTERNAL .....	4-66
4.12.3.9	MARSHAL .....	4-66
4.12.3.10	INITIALIZE .....	4-67
4.12.3.11	NO_IMPLEMENT .....	4-67
4.12.3.12	BAD_TYPECODE .....	4-67
4.12.3.13	BAD_OPERATION .....	4-67
4.12.3.14	NO_RESOURCES .....	4-67
4.12.3.15	NO_RESPONSE .....	4-67
4.12.3.16	PERSIST_STORE .....	4-67
4.12.3.17	BAD_INV_ORDER .....	4-67
4.12.3.18	TRANSIENT .....	4-68
4.12.3.19	FREE_MEM .....	4-68
4.12.3.20	INV_IDENT .....	4-68
4.12.3.21	INV_FLAG .....	4-68
4.12.3.22	INTF_REPOS .....	4-68
4.12.3.23	BAD_CONTEXT .....	4-68
4.12.3.24	OBJ_ADAPTER .....	4-68
4.12.3.25	DATA_CONVERSION .....	4-68
4.12.3.26	OBJECT_NOT_EXIST .....	4-69
4.12.3.27	TRANSACTION_REQUIRED .....	4-69
4.12.3.28	TRANSACTION_ROLLEDBACK ..	4-69
4.12.3.29	INVALID_TRANSACTION .....	4-69
4.12.3.30	INV_POLICY .....	4-69
4.12.3.31	CODESET_INCOMPATIBLE .....	4-69
4.12.3.32	REBIND .....	4-69
4.12.3.33	TIMEOUT .....	4-70
4.12.3.34	TRANSACTION_UNAVAILABLE ..	4-70
4.12.3.35	TRANSACTION_MODE .....	4-70
4.12.3.36	BAD_QOS .....	4-70
4.12.4	Standard Minor Exception Codes .....	4-70
5.	Value Type Semantics .....	5-1
5.1	Overview .....	5-1
5.2	Architecture .....	5-2
5.2.1	Abstract Values .....	5-3

5.2.2	<b>Operations</b> .....	5-3
5.2.3	<b>Value Type vs. Interfaces</b> .....	5-4
5.2.4	<b>Parameter Passing</b> .....	5-4
5.2.4.1	Value vs. Reference Semantics .....	5-4
5.2.4.2	Sharing Semantics .....	5-4
5.2.4.3	Identity Semantics .....	5-4
5.2.4.4	Any parameter type .....	5-5
5.2.5	<b>Substitutability Issues</b> .....	5-5
5.2.5.1	Value instance -> Interface type .....	5-5
5.2.5.2	Value Instance -> Abstract interface type .....	5-5
5.2.5.3	Value instance -> Value type .....	5-5
5.2.6	<b>Widening/Narrowing</b> .....	5-6
5.2.7	<b>Value Base Type</b> .....	5-6
5.2.8	<b>Life Cycle issues</b> .....	5-7
5.2.8.1	Creation and Factories .....	5-7
5.2.9	<b>Security Considerations</b> .....	5-7
5.2.9.1	Value as Value .....	5-8
5.2.9.2	Value as Object Reference .....	5-8
5.3	<b>Standard Value Box Definitions</b> .....	5-9
5.4	<b>Language Mappings</b> .....	5-9
5.4.1	<b>General Requirements</b> .....	5-9
5.4.2	<b>Language Specific Marshaling</b> .....	5-9
5.4.3	<b>Language Specific Value Factory Requirements</b> .....	5-9
5.4.4	<b>Value Method Implementation</b> .....	5-10
5.5	<b>Custom Marshaling</b> .....	5-10
5.5.1	<b>Implementation of Custom Marshaling</b> .....	5-11
5.5.2	<b>Marshaling Streams</b> .....	5-11
5.6	<b>Access to the Sending Context Run Time</b> .....	5-18
<b>6.</b>	<b>Abstract Interface Semantics</b> .....	<b>6-1</b>
6.1	Overview .....	6-1
6.2	Semantics of Abstract Interfaces .....	6-1
6.3	Usage Guidelines .....	6-3
6.4	Example .....	6-3
6.5	Security Considerations .....	6-4
6.5.1	<b>Passing Values to Trusted Domains</b> .....	6-4
<b>7.</b>	<b>Dynamic Invocation Interface</b> .....	<b>7-1</b>
7.1	Overview .....	7-1
7.1.1	<b>Common Data Structures</b> .....	7-2
7.1.2	<b>Memory Usage</b> .....	7-4
7.1.3	<b>Return Status and Exceptions</b> .....	7-4
7.2	Request Operations .....	7-4
7.2.1	<b>create_request</b> .....	7-5
7.2.2	<b>add_arg</b> .....	7-7
7.2.3	<b>invoke</b> .....	7-8

# Contents

---

	7.2.4	delete	7-8
	7.2.5	send	7-8
	7.2.6	poll_response	7-9
	7.2.7	get_response	7-9
	7.2.8	sendp	7-10
	7.2.9	prepare	7-10
	7.2.10	sendc	7-10
7.3		ORB Operations	7-11
	7.3.1	send_multiple_requests	7-11
	7.3.2	get_next_response and poll_next_response	7-11
7.4		Polling	7-12
	7.4.1	Abstract Valuetype Pollable	7-14
		7.4.1.1 is_ready	7-14
		7.4.1.2 create_pollable_set	7-14
	7.4.2	Abstract Valuetype DIIPollable	7-14
	7.4.3	interface PollableSet	7-14
		7.4.3.1 create_dii_pollable	7-15
		7.4.3.2 add_pollable	7-15
		7.4.3.3 get_ready_pollable	7-15
		7.4.3.4 remove	7-16
		7.4.3.5 number_left	7-16
7.5		List Operations	7-16
	7.5.1	create_list	7-17
	7.5.2	add_item	7-17
	7.5.3	free	7-17
	7.5.4	free_memory	7-18
	7.5.5	get_count	7-18
	7.5.6	create_operation_list	7-18
8.		Dynamic Skeleton Interface	8-1
	8.1	Introduction	8-1
	8.2	Overview	8-2
	8.3	ServerRequestPseudo-Object	8-3
		8.3.1 ExplicitRequest State: ServerRequestPseudo-Object	8-3
	8.4	DSI: Language Mapping	8-4
		8.4.1 ServerRequest's Handling of Operation Parameters	8-4
		8.4.2 Registering Dynamic Implementation Routines	8-5
9.		Dynamic Management of Any Values	9-1
	9.1	Overview	9-1
	9.2	DynAny API	9-3
		9.2.1 Locality and Usage Constraints	9-9
		9.2.2 Creating a DynAny Object	9-9
		9.2.3 The DynAny Interface	9-11
		9.2.3.1 Obtaining the TypeCode associated	

	with a DynAny object .....	9-11
	9.2.3.2 Initializing a DynAny object from another DynAny object .....	9-12
	9.2.3.3 Initializing a DynAny object from an any value .....	9-12
	9.2.3.4 Generating an any value from a DynAny object .....	9-12
	9.2.3.5 Comparing DynAny values .....	9-12
	9.2.3.6 Destroying a DynAny object .....	9-13
	9.2.3.7 Creating a copy of a DynAny object ....	9-13
	9.2.3.8 Accessing a value of some basic type in a DynAny object .....	9-13
	9.2.3.9 Iterating through components of a DynAny	9-15
	<b>9.2.4 The DynFixed Interface .....</b>	<b>9-16</b>
	<b>9.2.5 The DynEnum Interface .....</b>	<b>9-16</b>
	<b>9.2.6 The DynStruct Interface .....</b>	<b>9-17</b>
	<b>9.2.7 The DynUnion interface .....</b>	<b>9-19</b>
	<b>9.2.8 The DynSequence Interface .....</b>	<b>9-21</b>
	<b>9.2.9 The DynArray Interface .....</b>	<b>9-22</b>
	<b>9.2.10 The DynValueCommon Interface .....</b>	<b>9-23</b>
	<b>9.2.11 The DynValue Interface .....</b>	<b>9-24</b>
	<b>9.2.12 The DynValueBox Interface .....</b>	<b>9-24</b>
<b>9.3</b>	<b>Usage in C++ Language .....</b>	<b>9-25</b>
	<b>9.3.1 Dynamic creation of CORBA::Any values .....</b>	<b>9-25</b>
	9.3.1.1 Creating an any that contains a struct ...	9-25
	<b>9.3.2 Dynamic interpretation of CORBA::Any values .....</b>	<b>9-26</b>
	9.3.2.1 Filtering of events .....	9-26
<b>10.</b>	<b>The Interface Repository .....</b>	<b>10-1</b>
	10.1 Overview .....	10-1
	10.2 Scope of an Interface Repository .....	10-2
	10.3 Implementation Dependencies .....	10-4
	10.3.1 Managing Interface Repositories .....	10-4
	10.4 Basics .....	10-5
	10.4.1 Names and Identifiers .....	10-6
	10.4.2 Types and TypeCodes .....	10-6
	10.4.3 Interface Repository Objects .....	10-6
	10.4.4 Structure and Navigation of the Interface Repository .....	10-7
	10.5 Interface Repository Interfaces .....	10-9
	10.5.1 Supporting Type Definitions .....	10-10
	10.5.2 IRObjct .....	10-11
	10.5.2.1 Read Interface .....	10-11
	10.5.2.2 Write Interface .....	10-11
	10.5.3 Contained .....	10-11
	10.5.3.1 Read Interface .....	10-12
	10.5.3.2 Write Interface .....	10-13
	10.5.4 Container .....	10-14
	10.5.4.1 Read Interface .....	10-17

# Contents

---

	10.5.4.2 Write Interface .....	10-18
<b>10.5.5</b>	<b>IDLType .....</b>	<b>10-19</b>
<b>10.5.6</b>	<b>Repository .....</b>	<b>10-20</b>
	10.5.6.1 Read Interface .....	10-21
	10.5.6.2 Write Interface .....	10-21
<b>10.5.7</b>	<b>ModuleDef .....</b>	<b>10-22</b>
<b>10.5.8</b>	<b>ConstantDef .....</b>	<b>10-22</b>
	10.5.8.1 Read Interface .....	10-22
	10.5.8.2 Write Interface .....	10-23
<b>10.5.9</b>	<b>TypedefDef .....</b>	<b>10-23</b>
<b>10.5.10</b>	<b>StructDef .....</b>	<b>10-23</b>
	10.5.10.1 Read Interface .....	10-24
	10.5.10.2 Write Interface .....	10-24
<b>10.5.11</b>	<b>UnionDef .....</b>	<b>10-24</b>
	10.5.11.1 Read Interface .....	10-24
	10.5.11.2 Write Interface .....	10-25
<b>10.5.12</b>	<b>EnumDef .....</b>	<b>10-25</b>
	10.5.12.1 Read Interface .....	10-25
	10.5.12.2 Write Interface .....	10-25
<b>10.5.13</b>	<b>AliasDef .....</b>	<b>10-25</b>
	10.5.13.1 Read Interface .....	10-26
	10.5.13.2 Write Interface .....	10-26
<b>10.5.14</b>	<b>PrimitiveDef .....</b>	<b>10-26</b>
<b>10.5.15</b>	<b>StringDef .....</b>	<b>10-26</b>
<b>10.5.16</b>	<b>WstringDef .....</b>	<b>10-27</b>
<b>10.5.17</b>	<b>FixedDef .....</b>	<b>10-27</b>
<b>10.5.18</b>	<b>SequenceDef .....</b>	<b>10-27</b>
	10.5.18.1 Read Interface .....	10-28
	10.5.18.2 Write Interface .....	10-28
<b>10.5.19</b>	<b>ArrayDef .....</b>	<b>10-28</b>
	10.5.19.1 Read Interface .....	10-28
	10.5.19.2 Write Interface .....	10-28
<b>10.5.20</b>	<b>ExceptionDef .....</b>	<b>10-29</b>
	10.5.20.1 Read Interface .....	10-29
	10.5.20.2 Write Interface .....	10-29
<b>10.5.21</b>	<b>AttributeDef .....</b>	<b>10-29</b>
	10.5.21.1 Read Interface .....	10-30
	10.5.21.2 Write Interface .....	10-30
<b>10.5.22</b>	<b>OperationDef .....</b>	<b>10-30</b>
	10.5.22.1 Read Interface .....	10-31
	10.5.22.2 Write Interface .....	10-32
<b>10.5.23</b>	<b>InterfaceDef .....</b>	<b>10-32</b>
	10.5.23.1 Read Interface .....	10-33
	10.5.23.2 Write Interface .....	10-34
<b>10.5.24</b>	<b>AbstractInterfaceDef .....</b>	<b>10-34</b>
	10.5.24.1 Read Interface .....	10-34
	10.5.24.2 Write Interface .....	10-35
<b>10.5.25</b>	<b>LocalInterfaceDef .....</b>	<b>10-35</b>
	10.5.25.1 Read Interface .....	10-36
	10.5.25.2 Write Interface .....	10-36
<b>10.5.26</b>	<b>ValueMemberDef .....</b>	<b>10-37</b>
	10.5.26.1 Read Interface .....	10-37
	10.5.26.2 Write Interface .....	10-38

10.5.27	<b>ValueDef</b> .....	10-38
10.5.27.1	Read Interface .....	10-40
10.5.27.2	Write Interface .....	10-40
10.5.28	<b>ValueBoxDef</b> .....	10-41
10.5.28.1	Read Interface .....	10-41
10.5.28.2	Write Interface .....	10-41
10.5.29	<b>NativeDef</b> .....	10-41
10.6	<b>RepositoryIds</b> .....	10-42
10.6.1	<b>OMG IDL Format</b> .....	10-42
10.6.2	<b>RMI Hashed Format</b> .....	10-43
10.6.3	<b>DCE UUID Format</b> .....	10-44
10.6.4	<b>LOCAL Format</b> .....	10-45
10.6.5	<b>Pragma Directives for RepositoryId</b> .....	10-45
10.6.5.1	The ID Pragma .....	10-45
10.6.5.2	The Prefix Pragma .....	10-45
10.6.5.3	The Version Pragma .....	10-48
10.6.5.4	Generation of OMG IDL - Format IDs ..	10-49
10.6.6	<b>For More Information</b> .....	10-50
10.6.7	<b>RepositoryIDs for OMG-Specified Types</b> .....	10-50
10.7	<b>OMG IDL for Interface Repository</b> .....	10-51
11.	<b>The Portable Object Adapter</b> .....	11-1
11.1	<b>Overview</b> .....	11-1
11.2	<b>Abstract Model Description</b> .....	11-2
11.2.1	<b>Model Components</b> .....	11-2
11.2.2	<b>Model Architecture</b> .....	11-4
11.2.3	<b>POA Creation</b> .....	11-6
11.2.4	<b>Reference Creation</b> .....	11-7
11.2.5	<b>Object Activation States</b> .....	11-8
11.2.6	<b>Request Processing</b> .....	11-9
11.2.7	<b>Implicit Activation</b> .....	11-10
11.2.8	<b>Multi-threading</b> .....	11-11
11.2.8.1	POA Threading Models .....	11-11
11.2.8.2	Using the Single Thread Model .....	11-11
11.2.8.3	Using the ORB Controlled Model .....	11-12
11.2.8.4	Using the Main Thread Model .....	11-12
11.2.8.5	Limitations When Using Multiple Threads .....	11-12
11.2.9	<b>Dynamic Skeleton Interface</b> .....	11-12
11.2.10	<b>Location Transparency</b> .....	11-14
11.3	<b>Interfaces</b> .....	11-14
11.3.1	<b>The Servant IDL Type</b> .....	11-15
11.3.2	<b>POAManager Interface</b> .....	11-15
11.3.2.1	Processing States .....	11-16
11.3.2.2	activate .....	11-18
11.3.2.3	hold_requests .....	11-18
11.3.2.4	discard_requests .....	11-19
11.3.2.5	deactivate .....	11-19
11.3.2.6	get_state .....	11-20



# Contents

---

11.3.3	<b>AdapterActivator Interface</b> .....	11-20
	11.3.3.1 unknown_adapter .....	11-20
11.3.4	<b>ServantManager Interface</b> .....	11-22
	11.3.4.1 Common Information for Servant Manager Types .....	11-22
11.3.5	<b>ServantActivator Interface</b> .....	11-23
	11.3.5.1 incarnate .....	11-23
	11.3.5.2 etherealize .....	11-24
11.3.6	<b>ServantLocator Interface</b> .....	11-25
	11.3.6.1 preinvoke .....	11-26
	11.3.6.2 postinvoke .....	11-27
	11.3.6.3 ServantLocator and Location Determination .....	11-27
11.3.7	<b>POA Policy Objects</b> .....	11-28
	11.3.7.1 Thread Policy .....	11-28
	11.3.7.2 Lifespan Policy .....	11-29
	11.3.7.3 Object Id Uniqueness Policy .....	11-29
	11.3.7.4 Id Assignment Policy .....	11-30
	11.3.7.5 Servant Retention Policy .....	11-30
	11.3.7.6 Request Processing Policy .....	11-31
	11.3.7.7 Implicit Activation Policy .....	11-32
11.3.8	<b>POA Interface</b> .....	11-33
	11.3.8.1 create_POA .....	11-33
	11.3.8.2 find_POA .....	11-34
	11.3.8.3 destroy .....	11-34
	11.3.8.4 Policy Creation Operations .....	11-35
	11.3.8.5 the_name .....	11-36
	11.3.8.6 the_parent .....	11-36
	11.3.8.7 the_children .....	11-36
	11.3.8.8 the_POAManager .....	11-36
	11.3.8.9 the_activator .....	11-36
	11.3.8.10 get_servant_manager .....	11-37
	11.3.8.11 set_servant_manager .....	11-37
	11.3.8.12 get_servant .....	11-37
	11.3.8.13 set_servant .....	11-37
	11.3.8.14 activate_object .....	11-38
	11.3.8.15 activate_object_with_id .....	11-38
	11.3.8.16 deactivate_object .....	11-38
	11.3.8.17 create_reference .....	11-39
	11.3.8.18 create_reference_with_id .....	11-39
	11.3.8.19 servant_to_id .....	11-40
	11.3.8.20 servant_to_reference .....	11-41
	11.3.8.21 reference_to_servant .....	11-41
	11.3.8.22 reference_to_id .....	11-42
	11.3.8.23 id_to_servant .....	11-42
	11.3.8.24 id_to_reference .....	11-42
	11.3.8.25 id .....	11-42
11.3.9	<b>Current Operations</b> .....	11-43
	11.3.9.1 get_POA .....	11-43
	11.3.9.2 get_object_id .....	11-43
	11.3.9.3 get_reference .....	11-43
	11.3.9.4 get_servant .....	11-44
11.4	<b>IDL for PortableServer Module</b> .....	11-44
11.5	<b>UML Description of PortableServer</b> .....	11-50

11.6	Usage Scenarios . . . . .	11-52
11.6.1	Getting the Root POA . . . . .	11-52
11.6.2	Creating a POA . . . . .	11-53
11.6.3	Explicit Activation with POA-assigned Object Ids	11-53
11.6.4	Explicit Activation with User-assigned Object Ids	11-54
11.6.5	Creating References before Activation. . . . .	11-55
11.6.6	Servant Manager Definition and Creation. . . . .	11-55
11.6.7	Object Activation on Demand. . . . .	11-57
11.6.8	Persistent Objects with POA-assigned Ids. . . . .	11-59
11.6.9	Multiple Object Ids Mapping to a Single Servant	11-59
11.6.10	One Servant for All Objects . . . . .	11-59
11.6.11	Single Servant, Many Objects and Types, Using DSI . . . . .	11-62
12.	<b>Interoperability Overview . . . . .</b>	<b>12-1</b>
12.1	Elements of Interoperability. . . . .	12-1
12.1.1	ORB Interoperability Architecture . . . . .	12-2
12.1.2	Inter-ORB Bridge Support . . . . .	12-2
12.1.3	General Inter-ORB Protocol (GIOP). . . . .	12-3
12.1.4	Internet Inter-ORB Protocol (IIOP). . . . .	12-3
12.1.5	Environment-Specific Inter-ORB Protocols (ESIOPs). . . . .	12-4
12.2	Relationship to Previous Versions of CORBA . . . . .	12-4
12.3	Examples of Interoperability Solutions . . . . .	12-5
12.3.1	Example 1. . . . .	12-5
12.3.2	Example 2. . . . .	12-5
12.3.3	Example 3. . . . .	12-5
12.3.4	Interoperability Compliance. . . . .	12-5
12.4	Motivating Factors . . . . .	12-8
12.4.1	ORB Implementation Diversity . . . . .	12-8
12.4.2	ORB Boundaries . . . . .	12-8
12.4.3	ORBs Vary in Scope, Distance, and Lifetime. . . . .	12-9
12.5	Interoperability Design Goals. . . . .	12-9
12.5.1	Non-Goals. . . . .	12-10
13.	<b>ORB Interoperability Architecture . . . . .</b>	<b>13-1</b>
13.1	Overview . . . . .	13-1
13.1.1	Domains . . . . .	13-2
13.1.2	Bridging Domains . . . . .	13-2
13.2	ORBs and ORB Services . . . . .	13-3
13.2.1	The Nature of ORB Services. . . . .	13-3
13.2.2	ORB Services and Object Requests . . . . .	13-3
13.2.3	Selection of ORB Services. . . . .	13-4
13.3	Domains . . . . .	13-5
13.3.1	Definition of a Domain. . . . .	13-5

# Contents

---

13.3.2	Mapping Between Domains: Bridging .....	13-6
13.4	Interoperability Between ORBs .....	13-7
13.4.1	ORB Services and Domains .....	13-7
13.4.2	ORBs and Domains .....	13-7
13.4.3	Interoperability Approaches .....	13-8
13.4.3.1	Mediated Bridging .....	13-8
13.4.3.2	Immediate Bridging .....	13-9
13.4.3.3	Location of Inter-Domain Functionality .....	13-9
13.4.3.4	Bridging Level .....	13-10
13.4.4	Policy-Mediated Bridging .....	13-10
13.4.5	Configurations of Bridges in Networks .....	13-11
13.5	Object Addressing .....	13-11
13.5.1	Domain-relative Object Referencing .....	13-12
13.5.2	Handling of Referencing Between Domains .....	13-12
13.6	An Information Model for Object References .....	13-14
13.6.1	What Information Do Bridges Need? .....	13-14
13.6.2	Interoperable Object References: IORs .....	13-14
13.6.3	IOR Profiles .....	13-15
13.6.4	Standard IOR Profiles .....	13-17
13.6.4.1	The TAG_INTERNET_IOP Profile .....	13-17
13.6.4.2	The TAG_MULTIPLE_COMPONENTS Profile .....	13-18
13.6.4.3	The TAG_SCCP_IOP Profile .....	13-18
13.6.5	IOR Components .....	13-18
13.6.6	Standard IOR Components .....	13-19
13.6.6.1	TAG_ORB_TYPE Component .....	13-20
13.6.6.2	TAG_ALTERNATE_IOP_ADDRESS Component .....	13-20
13.6.6.3	Other Components .....	13-20
13.6.7	Profile and Component Composition in IORs .....	13-21
13.6.8	IOR Creation and Scope .....	13-22
13.6.9	Stringified Object References .....	13-22
13.6.10	Object URLs .....	13-23
13.6.10.1	corbaloc URL .....	13-24
13.6.10.2	corbaloc:rir URL .....	13-25
13.6.10.3	corbaloc:iiop URL .....	13-26
13.6.10.4	corbaloc Server Implementation .....	13-27
13.6.10.5	corbaname URL .....	13-27
13.6.10.6	Future corbaloc URL Protocols .....	13-27
13.6.10.7	Future URL Schemes .....	13-27
13.7	Service Context .....	13-28
13.7.1	Standard Service Contexts .....	13-29
13.7.2	Service Context Processing Rules .....	13-31
13.8	Coder/Decoder Interfaces .....	13-31
13.8.1	Codec Interface .....	13-31
13.8.1.1	Exceptions .....	13-32
13.8.1.2	Operations .....	13-32
13.8.2	Codec Factory .....	13-33
13.8.2.1	Encoding Structure .....	13-34

13.8.2.2 CodecFactory Interface .....	13-34
13.9 Feature Support and GIOP Versions .....	13-35
13.10 Code Set Conversion .....	13-36
<b>13.10.1 Character Processing Terminology .....</b>	<b>13-36</b>
13.10.1.1 Character Set .....	13-36
13.10.1.2 Coded Character Set, or Code Set .....	13-36
13.10.1.3 Code Set Classifications .....	13-37
13.10.1.4 Narrow and Wide Characters .....	13-37
13.10.1.5 Char Data and Wchar Data .....	13-38
13.10.1.6 Byte-Oriented Code Set .....	13-38
13.10.1.7 Multi-Byte Character Strings .....	13-38
13.10.1.8 Non-Byte-Oriented Code Set .....	13-38
13.10.1.9 Char and Wchar Transmission Code Set (TCS-C and TCS-W) .....	13-38
13.10.1.10 Process Code Set and File Code Set ..	13-38
13.10.1.11 Native Code Set .....	13-39
13.10.1.12 Transmission Code Set .....	13-39
13.10.1.13 Conversion Code Set (CCS) .....	13-39
<b>13.10.2 Code Set Conversion Framework .....</b>	<b>13-39</b>
13.10.2.1 Requirements .....	13-39
13.10.2.2 Overview of the Conversion Framework .....	13-40
13.10.2.3 ORB Databases and Code Set Converters .....	13-41
13.10.2.4 CodeSet Component of IOR Multi-Component Profile .....	13-42
13.10.2.5 GIOP Code Set Service Context .....	13-43
13.10.2.6 Code Set Negotiation .....	13-44
<b>13.10.3 Mapping to Generic Character Environments ..</b>	<b>13-47</b>
13.10.3.1 Describing Generic Interfaces .....	13-48
13.10.3.2 Interoperation .....	13-48
<b>13.10.4 Example of Generic Environment Mapping ....</b>	<b>13-48</b>
13.10.4.1 Generic Mappings .....	13-49
13.10.4.2 Interoperation and Generic Mappings ..	13-49
<b>13.10.5 Relevant OSFM Registry Interfaces .....</b>	<b>13-49</b>
13.10.5.1 Character and Code Set Registry .....	13-49
13.10.5.2 Access Routines .....	13-50
<b>14. Building Inter-ORB Bridges .....</b>	<b>14-1</b>
14.1 Introduction .....	14-1
14.2 In-Line and Request-Level Bridging .....	14-2
14.2.1 In-line Bridging .....	14-3
14.2.2 Request-level Bridging .....	14-3
14.2.3 Collocated ORBs .....	14-4
14.3 Proxy Creation and Management .....	14-5
14.4 Interface-specific Bridges and Generic Bridges .....	14-6
14.5 Building Generic Request-Level Bridges .....	14-6
14.6 Bridging Non-Referencing Domains .....	14-7
14.7 Bootstrapping Bridges .....	14-7

# Contents

---

<b>15. General Inter-ORB Protocol</b> .....	<b>15-1</b>
15.1 Goals of the General Inter-ORB Protocol .....	15-2
15.2 GIOP Overview .....	15-2
15.2.1 Common Data Representation (CDR) .....	15-3
15.2.2 GIOP Message Overview .....	15-3
15.2.3 GIOP Message Transfer .....	15-4
15.3 CDR Transfer Syntax .....	15-4
15.3.1 Primitive Types .....	15-5
15.3.1.1 Alignment .....	15-5
15.3.1.2 Integer Data Types .....	15-6
15.3.1.3 Floating Point Data Types .....	15-7
15.3.1.4 Octet .....	15-10
15.3.1.5 Boolean .....	15-10
15.3.1.6 Character Types .....	15-10
15.3.2 OMG IDL Constructed Types .....	15-11
15.3.2.1 Alignment .....	15-11
15.3.2.2 Struct .....	15-12
15.3.2.3 Union .....	15-12
15.3.2.4 Array .....	15-12
15.3.2.5 Sequence .....	15-12
15.3.2.6 Enum .....	15-12
15.3.2.7 Strings and Wide Strings .....	15-12
15.3.2.8 Fixed-Point Decimal Type .....	15-13
15.3.3 Encapsulation .....	15-14
15.3.4 Value Types .....	15-15
15.3.4.1 Partial Type Information and Versioning .....	15-16
15.3.4.2 Example .....	15-17
15.3.4.3 Scope of the Indirections .....	15-19
15.3.4.4 Null Values .....	15-19
15.3.4.5 Other Encoding Information .....	15-19
15.3.4.6 Fragmentation .....	15-19
15.3.4.7 Notation .....	15-22
15.3.4.8 The Format .....	15-22
15.3.5 Pseudo-Object Types .....	15-23
15.3.5.1 TypeCode .....	15-23
15.3.5.2 Any .....	15-29
15.3.5.3 Principal .....	15-29
15.3.5.4 Context .....	15-29
15.3.5.5 Exception .....	15-29
15.3.6 Object References .....	15-30
15.3.7 Abstract Interfaces .....	15-30
15.4 GIOP Message Formats .....	15-30
15.4.1 GIOP Message Header .....	15-31
15.4.2 Request Message .....	15-33
15.4.2.1 Request Header .....	15-33
15.4.2.2 Request Body .....	15-36
15.4.3 Reply Message .....	15-37
15.4.3.1 Reply Header .....	15-37
15.4.3.2 Reply Body .....	15-38
15.4.4 CancelRequest Message .....	15-40
15.4.4.1 Cancel Request Header .....	15-40

	15.4.5 <b>LocateRequest Message</b> .....	15-41
	15.4.5.1 LocateRequest Header .....	15-41
	15.4.6 <b>LocateReply Message</b> .....	15-42
	15.4.6.1 Locate Reply Header .....	15-42
	15.4.6.2 LocateReply Body .....	15-43
	15.4.6.3 Handling ForwardRequest Exception from ServantLocator .....	15-44
	15.4.7 <b>CloseConnection Message</b> .....	15-44
	15.4.8 <b>MessageError Message</b> .....	15-44
	15.4.9 <b>Fragment Message</b> .....	15-44
15.5	<b>GIOP Message Transport</b> .....	15-46
	15.5.1 <b>Connection Management</b> .....	15-46
	15.5.1.1 Connection Closure .....	15-47
	15.5.1.2 Multiplexing Connections .....	15-48
	15.5.2 <b>Message Ordering</b> .....	15-48
15.6	<b>Object Location</b> .....	15-48
15.7	<b>Internet Inter-ORB Protocol (IIOP)</b> .....	15-50
	15.7.1 <b>TCP/IP Connection Usage</b> .....	15-51
	15.7.2 <b>IIOP IOR Profiles</b> .....	15-51
	15.7.3 <b>IIOP IOR Profile Components</b> .....	15-54
15.8	<b>Bi-Directional GIOP</b> .....	15-55
	15.8.1 <b>Bi-Directional IIOP</b> .....	15-57
	15.8.1.1 IIOP/SSL considerations .....	15-58
15.9	<b>Bi-directional GIOP policy</b> .....	15-58
15.10	<b>OMG IDL</b> .....	15-59
	15.10.1 <b>GIOP Module</b> .....	15-59
	15.10.2 <b>IIOP Module</b> .....	15-63
	15.10.3 <b>BiDirPolicy Module</b> .....	15-64
<b>16.</b>	<b>The DCE ESIOP</b> .....	<b>16-1</b>
	16.1 <b>Goals of the DCE Common Inter-ORB Protocol</b> .....	16-1
	16.2 <b>DCE Common Inter-ORB Protocol Overview</b> .....	16-2
	16.2.1 <b>DCE-CIOP RPC</b> .....	16-2
	16.2.2 <b>DCE-CIOP Data Representation</b> .....	16-3
	16.2.3 <b>DCE-CIOP Messages</b> .....	16-4
	16.2.4 <b>Interoperable Object Reference (IOR)</b> .....	16-5
	16.3 <b>DCE-CIOP Message Transport</b> .....	16-5
	16.3.1 <b>Pipe-based Interface</b> .....	16-6
	16.3.1.1 Invoke .....	16-8
	16.3.1.2 Locate .....	16-8
	16.3.2 <b>Array-based Interface</b> .....	16-8
	16.3.2.1 Invoke .....	16-10
	16.3.2.2 Locate .....	16-11
	16.4 <b>DCE-CIOP Message Formats</b> .....	16-11
	16.4.1 <b>DCE_CIOP Invoke Request Message</b> .....	16-11
	16.4.1.1 Invoke request header .....	16-11
	16.4.1.2 Invoke request body .....	16-12

# Contents

---

16.4.2	<b>DCE-CIOP Invoke Response Message</b> .....	16-12
	16.4.2.1 Invoke response header .....	16-13
	16.4.2.2 Invoke Response Body .....	16-13
16.4.3	<b>DCE-CIOP Locate Request Message</b> .....	16-14
	16.4.3.1 Locate Request Header .....	16-14
16.4.4	<b>DCE-CIOP Locate Response Message</b> .....	16-15
	16.4.4.1 Locate Response Header .....	16-15
	16.4.4.2 Locate Response Body .....	16-16
16.5	<b>DCE-CIOP Object References</b> .....	16-16
	16.5.1 <b>DCE-CIOP String Binding Component</b> .....	16-17
	16.5.2 <b>DCE-CIOP Binding Name Component</b> .....	16-18
	16.5.2.1 BindingNameComponent .....	16-18
	16.5.3 <b>DCE-CIOP No Pipes Component</b> .....	16-19
	16.5.4 <b>Complete Object Key Component</b> .....	16-19
	16.5.5 <b>Endpoint ID Position Component</b> .....	16-20
	16.5.6 <b>Location Policy Component</b> .....	16-20
16.6	<b>DCE-CIOP Object Location</b> .....	16-21
	16.6.1 <b>Location Mechanism Overview</b> .....	16-22
	16.6.2 <b>Activation</b> .....	16-23
	16.6.3 <b>Basic Location Algorithm</b> .....	16-23
	16.6.4 <b>Use of the Location Policy and the Endpoint ID</b> .....	16-24
	16.6.4.1 Current location policy .....	16-24
	16.6.4.2 Original location policy .....	16-24
	16.6.4.3 Original Endpoint ID .....	16-24
16.7	<b>OMG IDL for the DCE CIOP Module</b> .....	16-25
16.8	<b>References for this Chapter</b> .....	16-26
<b>17.</b>	<b>Interworking Architecture</b> .....	<b>17-1</b>
	17.1 <b>Purpose of the Interworking Architecture</b> .....	17-2
	17.1.1 <b>Comparing COM Objects to CORBA Objects</b> ..	17-2
	17.2 <b>Interworking Object Model</b> .....	17-3
	17.2.1 <b>Relationship to CORBA Object Model</b> .....	17-3
	17.2.2 <b>Relationship to the OLE/COM Model</b> .....	17-4
	17.2.3 <b>Basic Description of the Interworking Model</b> ...	17-4
	17.3 <b>Interworking Mapping Issues</b> .....	17-8
	17.4 <b>Interface Mapping</b> .....	17-8
	17.4.1 <b>CORBA/COM</b> .....	17-9
	17.4.2 <b>CORBA/Automation</b> .....	17-9
	17.4.3 <b>COM/CORBA</b> .....	17-10
	17.4.4 <b>Automation/CORBA</b> .....	17-10
	17.5 <b>Interface Composition Mappings</b> .....	17-11
	17.5.1 <b>CORBA/COM</b> .....	17-11
	17.5.1.1 <b>COM/CORBA</b> .....	17-12
	17.5.1.2 <b>CORBA/Automation</b> .....	17-12
	17.5.1.3 <b>Automation/CORBA</b> .....	17-13
	17.5.2 <b>Detailed Mapping Rules</b> .....	17-13
	17.5.2.1 <b>Ordering Rules for the CORBA-&gt;MIDL</b>	

	Transformation .....	17-13
	17.5.2.2 Ordering Rules for the CORBA->Automation Transformation ..	17-13
<b>17.5.3</b>	<b>Example of Applying Ordering Rules .....</b>	<b>17-14</b>
<b>17.5.4</b>	<b>Mapping Interface Identity.....</b>	<b>17-16</b>
	17.5.4.1 Mapping Interface Repository IDs to COM IIDs .....	17-17
	17.5.4.2 Mapping COM IIDs to CORBA Interface IDs .....	17-18
<b>17.6</b>	<b>Object Identity, Binding, and Life Cycle .....</b>	<b>17-18</b>
	<b>17.6.1 Object Identity Issues .....</b>	<b>17-19</b>
	17.6.1.1 CORBA Object Identity and Reference Properties .....	17-19
	17.6.1.2 COM Object Identity and Reference Properties .....	17-19
	<b>17.6.2 Binding and Life Cycle .....</b>	<b>17-20</b>
	17.6.2.1 Lifetime Comparison .....	17-20
	17.6.2.2 Binding Existing CORBA Objects to COM Views .....	17-21
	17.6.2.3 Binding COM Objects to CORBA Views ..	17-22
	17.6.2.4 COM View of CORBA Life Cycle ....	17-22
	17.6.2.5 CORBA View of COM/Automation Life Cycle .....	17-23
<b>17.7</b>	<b>Interworking Interfaces .....</b>	<b>17-23</b>
	<b>17.7.1 SimpleFactory Interface .....</b>	<b>17-23</b>
	<b>17.7.2 IMonikerProvider Interface and Moniker Use ..</b>	<b>17-23</b>
	<b>17.7.3 ICORBAFactory Interface .....</b>	<b>17-24</b>
	<b>17.7.4 IForeignObject Interface.....</b>	<b>17-26</b>
	<b>17.7.5 ICORBAObject Interface .....</b>	<b>17-27</b>
	<b>17.7.6 ICORBAObject2 .....</b>	<b>17-28</b>
	<b>17.7.7 IORBObject Interface.....</b>	<b>17-28</b>
	<b>17.7.8 Naming Conventions for View Components ....</b>	<b>17-30</b>
	17.7.8.1 Naming the COM View Interface .....	17-30
	17.7.8.2 Tag for the Automation Interface Id ...	17-30
	17.7.8.3 Naming the Automation View Dispatch Interface .....	17-30
	17.7.8.4 Naming the Automation View Dual Interface .....	17-31
	17.7.8.5 Naming the Program Id for the COM Class .....	17-31
	17.7.8.6 Naming the Class Id for the COM Class .....	17-32
<b>17.8</b>	<b>Distribution .....</b>	<b>17-32</b>
	<b>17.8.1 Bridge Locality.....</b>	<b>17-32</b>
	<b>17.8.2 Distribution Architecture .....</b>	<b>17-33</b>
<b>17.9</b>	<b>Interworking Targets .....</b>	<b>17-34</b>
<b>17.10</b>	<b>Compliance to COM/CORBA Interworking.....</b>	<b>17-34</b>
	<b>17.10.1 Products Subject to Compliance.....</b>	<b>17-34</b>
	17.10.1.1 Interworking solutions .....	17-34
	17.10.1.2 Mapping solutions .....	17-35



# Contents

---

17.10.1.3	Mapped components	17-35
17.10.2	Compliance Points	17-36
<b>18.</b>	<b>Mapping: COM and CORBA</b>	<b>18-1</b>
18.1	Data Type Mapping	18-1
18.2	CORBA to COM Data Type Mapping	18-2
18.2.1	Mapping for Basic Data Types	18-2
18.2.2	Mapping for Constants	18-2
18.2.3	Mapping for Enumerators	18-3
18.2.4	Mapping for String Types	18-4
18.2.4.1	Mapping for Unbounded String Types	18-4
18.2.4.2	Mapping for Bounded String Types	18-5
18.2.5	Mapping for Struct Types	18-5
18.2.6	Mapping for Union Types	18-6
18.2.7	Mapping for Sequence Types	18-8
18.2.7.1	Mapping for Unbounded Sequence Types	18-8
18.2.7.2	Mapping for Bounded Sequence Types	18-8
18.2.8	Mapping for Array Types	18-9
18.2.9	Mapping for the any Type	18-9
18.2.10	Interface Mapping	18-11
18.2.10.1	Mapping for interface identifiers	18-11
18.2.10.2	Mapping for exception types	18-11
18.2.10.3	Mapping for Nested Types	18-21
18.2.10.4	Mapping for Operations	18-22
18.2.10.5	Mapping for Oneway Operations	18-24
18.2.10.6	Mapping for Attributes	18-24
18.2.10.7	Indirection Levels for Operation Parameters	18-26
18.2.11	Inheritance Mapping	18-26
18.2.12	Mapping for Pseudo-Objects	18-29
18.2.12.1	Mapping for TypeCode pseudo-object	18-29
18.2.12.2	Mapping for context pseudo-object	18-31
18.2.12.3	Mapping for principal pseudo-object	18-32
18.2.13	Interface Repository Mapping	18-32
18.3	COM to CORBA Data Type Mapping	18-33
18.3.1	Mapping for Basic Data Types	18-33
18.3.2	Mapping for Constants	18-34
18.3.3	Mapping for Enumerators	18-34
18.3.4	Mapping for String Types	18-35
18.3.4.1	Mapping for unbounded string types	18-35
18.3.4.2	Mapping for bounded string types	18-36
18.3.4.3	Mapping for Unicode Unbounded String Types	18-36
18.3.4.4	Mapping for unicode bound string types	18-37
18.3.5	Mapping for Structure Types	18-37
18.3.6	Mapping for Union Types	18-38
18.3.6.1	Mapping for Encapsulated Unions	18-38
18.3.6.2	Mapping for nonencapsulated unions	18-39
18.3.7	Mapping for Array Types	18-40
18.3.7.1	Mapping for nonfixed arrays	18-40

18.3.7.2 Mapping for SAFEARRAY .....	18-40
<b>18.3.8 Mapping for VARIANT.....</b>	<b>18-41</b>
<b>18.3.9 Mapping for Pointers.....</b>	<b>18-43</b>
<b>18.3.10 Interface Mapping .....</b>	<b>18-44</b>
18.3.10.1 Mapping for Interface Identifiers ....	18-44
18.3.10.2 Mapping for COM Errors .....	18-44
18.3.10.3 Mapping of Nested Data Types .....	18-47
18.3.10.4 Mapping of Names .....	18-47
18.3.10.5 Mapping for Operations .....	18-47
18.3.10.6 Mapping for Properties .....	18-48
<b>18.3.11 Mapping for Read-Only Attributes .....</b>	<b>18-49</b>
<b>18.3.12 Mapping for Read-Write Attributes .....</b>	<b>18-49</b>
18.3.12.1 Inheritance Mapping .....	18-50
18.3.12.2 Type Library Mapping .....	18-52
<b>19. Mapping: Automation and CORBA .....</b>	<b>19-1</b>
19.1 Mapping CORBA Objects to Automation .....	19-2
19.1.1 Architectural Overview.....	19-2
19.1.2 Main Features of the Mapping .....	19-3
19.2 Mapping for Interfaces.....	19-3
19.2.1 Mapping for Attributes and Operations .....	19-4
19.2.2 Mapping for OMG IDL Single Inheritance.....	19-5
19.2.3 Mapping of OMG IDL Multiple Inheritance....	19-6
19.3 Mapping for Basic Data Types .....	19-9
19.3.1 Basic Automation Types .....	19-9
19.3.2 Special Cases of Basic Data Type Mapping .....	19-10
19.3.2.1 Converting Automation long to CORBA unsigned long .....	19-10
19.3.2.2 Demoting CORBA unsigned long to Automation long .....	19-11
19.3.2.3 Demoting Automation long to CORBA unsigned short .....	19-11
19.3.2.4 Converting Automation boolean to CORBA boolean and CORBA boolean to Automation boolean .....	19-11
19.3.3 Mapping for Strings .....	19-11
19.4 IDL to ODL Mapping.....	19-12
19.4.1 A Complete IDL to ODL Mapping for the Basic Data Types .....	19-12
19.5 Mapping for Object References .....	19-15
19.5.1 Type Mapping .....	19-15
19.5.2 Object Reference Parameters and IForeignObject.....	19-16
19.6 Mapping for Enumerated Types .....	19-17
19.7 Mapping for Arrays and Sequences .....	19-18
19.8 Mapping for CORBA Complex Types .....	19-19
19.8.1 Mapping for Structure Types .....	19-20
19.8.2 Mapping for Union Types .....	19-21

# Contents

---

19.8.3	Mapping for TypeCodes .....	19-22
19.8.4	Mapping for anys .....	19-24
19.8.5	Mapping for Typedefs .....	19-25
19.8.6	Mapping for Constants .....	19-25
19.8.7	Getting Initial CORBA Object References .....	19-26
19.8.8	Creating Initial in Parameters for Complex Types	19-27
	19.8.8.1 ITypeFactory Interface .....	19-29
	19.8.8.2 DIObjectInfo Interface .....	19-29
19.8.9	Mapping CORBA Exceptions to Automation Exceptions .....	19-30
	19.8.9.1 Overview of Automation Exception Handling .....	19-30
	19.8.9.2 CORBA Exceptions .....	19-30
	19.8.9.3 CORBA User Exceptions .....	19-31
	19.8.9.4 Operations that Raise User Exceptions ..	19-32
	19.8.9.5 CORBA System Exceptions .....	19-33
	19.8.9.6 Operations that raise system exceptions	19-34
19.8.10	Conventions for Naming Components of the Automation View .....	19-36
19.8.11	Naming Conventions for Pseudo-Structs, Pseudo- Unions, and Pseudo-Exceptions .....	19-36
19.8.12	Automation View Interface as a Dispatch Interface (Nondual) .....	19-36
19.8.13	Aggregation of Automation Views .....	19-38
19.8.14	DII and DSI .....	19-38
19.9	Mapping Automation Objects as CORBA Objects .....	19-38
19.9.1	Architectural Overview .....	19-38
19.9.2	Main Features of the Mapping .....	19-39
19.9.3	Getting Initial Object References .....	19-40
19.9.4	Mapping for Interfaces .....	19-40
19.9.5	Mapping for Inheritance .....	19-40
19.9.6	Mapping for ODL Properties and Methods .....	19-41
19.9.7	Mapping for Automation Basic Data Types .....	19-42
	19.9.7.1 Basic automation types .....	19-42
19.9.8	Conversion Errors .....	19-43
19.9.9	Special Cases of Data Type Conversion .....	19-43
	19.9.9.1 Translating COM::Currency to Automation CURRENCY .....	19-43
	19.9.9.2 Translating CORBA double to Automation DATE .....	19-43
	19.9.9.3 Translating CORBA boolean to Automation boolean and Automation boolean to CORBA boolean .....	19-43
19.9.10	A Complete OMG IDL to ODL Mapping for the Basic Data Types .....	19-44
19.9.11	Mapping for Object References .....	19-46
19.9.12	Mapping for Enumerated Types .....	19-47
19.9.13	Mapping for SafeArrays .....	19-48
	19.9.13.1 Multidimensional SafeArrays .....	19-48
19.9.14	Mapping for Typedefs .....	19-48

19.9.15	Mapping for VARIANTS .....	19-48
19.9.16	Mapping Automation Exceptions to CORBA ...	19-49
19.10	Older Automation Controllers .....	19-49
19.10.1	Mapping for OMG IDL Arrays and Sequences to Collections .....	19-49
19.11	Example Mappings .....	19-51
19.11.1	Mapping the OMG Naming Service to Automation .....	19-51
19.11.2	Mapping a COM Service to OMG IDL .....	19-51
19.11.3	Mapping an OMG Object Service to Automation	19-55
20.	Interoperability with non-CORBA Systems .....	20-1
20.1	Introduction .....	20-1
20.1.1	COM/CORBA Part A .....	20-2
20.2	Conformance Issues .....	20-2
20.2.1	Performance Issues .....	20-3
20.2.2	Scalability Issues .....	20-3
20.2.3	CORBA Clients for DCOM Servers .....	20-3
20.3	Locality of the Bridge .....	20-4
20.4	Extent Definition .....	20-5
20.4.1	Marshaling Constraints .....	20-6
20.4.2	Marshaling Key .....	20-6
20.4.3	Extent Format .....	20-7
20.4.3.1	DVO_EXTENT .....	20-8
20.4.3.2	DVO_IFACE .....	20-8
20.4.3.3	DVO_IMPLDATA .....	20-8
20.4.3.4	DVO_BLOB .....	20-8
20.5	Request/Reply Extent Semantics .....	20-8
20.6	Consistency .....	20-9
20.6.1	IValueObject .....	20-10
20.6.2	ISynchronize and DISynchronize .....	20-11
20.6.2.1	Mode Property .....	20-11
20.6.2.2	SyncNow Method .....	20-11
20.6.2.3	ReCopy Method .....	20-11
20.7	DCOM Value Objects .....	20-11
20.7.1	Passing Automation Compound Types as DCOM Value Objects .....	20-11
20.7.2	Passing CORBA-Defined Pseudo-Objects as DCOM Value Objects .....	20-12
20.7.3	IForeignObject .....	20-12
20.7.4	DIForeignComplexType .....	20-12
20.7.5	DIForeignException .....	20-12
20.7.6	DISystemException .....	20-12
20.7.7	DICORBAUserException .....	20-13
20.7.8	DICORBAStruct .....	20-13
20.7.9	DICORBAUnion .....	20-13

# Contents

---

20.7.10	DICORBATypeCode and ICORBATypeCode . . .	20-13
20.7.11	DICORBAAny . . . . .	20-14
20.7.12	ICORBAAny . . . . .	20-15
20.7.13	User Exceptions In COM . . . . .	20-15
20.8	Chain Avoidance . . . . .	20-16
20.8.1	CORBA Chain Avoidance . . . . .	20-16
20.8.2	COM Chain Avoidance . . . . .	20-17
20.9	Chain Bypass . . . . .	20-19
20.9.1	CORBA Chain Bypass . . . . .	20-19
20.9.2	COM Chain Bypass . . . . .	20-20
20.10	Thread Identification . . . . .	20-21
<b>21.</b>	<b>Portable Interceptors . . . . .</b>	<b>21-1</b>
21.1	Introduction . . . . .	21-1
21.1.1	Object Creation . . . . .	21-2
21.1.2	Client Sends Request . . . . .	21-3
21.1.3	Server Receives Request . . . . .	21-4
21.1.4	Server Sends Reply . . . . .	21-4
21.1.5	Client Receives Reply . . . . .	21-5
21.2	Interceptor Interface . . . . .	21-5
21.3	Request Interceptors . . . . .	21-6
21.3.1	Design Principles . . . . .	21-6
21.3.2	General Flow Rules . . . . .	21-7
21.3.3	The Flow Stack Visual Model . . . . .	21-8
21.3.4	The Request Interceptor Points . . . . .	21-8
21.3.5	Client-Side Interceptor . . . . .	21-9
21.3.6	Client-Side Interception Points . . . . .	21-9
	21.3.6.1 send_request . . . . .	21-9
	21.3.6.2 send_poll . . . . .	21-9
	21.3.6.3 receive_reply . . . . .	21-10
	21.3.6.4 receive_exception . . . . .	21-10
	21.3.6.5 receive_other . . . . .	21-11
21.3.7	Client-Side Interception Point Flow . . . . .	21-11
	21.3.7.1 Client-side Flow Rules . . . . .	21-11
	21.3.7.2 Additional Client-side Details . . . . .	21-12
	21.3.7.3 Client-side Flow Examples . . . . .	21-12
21.3.8	Server-Side Interceptor . . . . .	21-14
21.3.9	Server-Side Interception Points . . . . .	21-14
	21.3.9.1 receive_request_service_contexts . . . . .	21-14
	21.3.9.2 receive_request . . . . .	21-15
	21.3.9.3 send_reply . . . . .	21-15
	21.3.9.4 send_exception . . . . .	21-16
	21.3.9.5 send_other . . . . .	21-16
21.3.10	Server-Side Interception Point Flow . . . . .	21-17
	21.3.10.1 Server-side Flow Rules . . . . .	21-17
	21.3.10.2 Additional Server-side Details . . . . .	21-17
	21.3.10.3 Server-side Flow Examples . . . . .	21-18
21.3.11	Request Information . . . . .	21-20

	<b>21.3.12 RequestInfo Interface</b> .....	<b>21-21</b>
	21.3.12.1 request_id .....	21-21
	21.3.12.2 operation .....	21-21
	21.3.12.3 arguments .....	21-21
	21.3.12.4 exceptions .....	21-22
	21.3.12.5 contexts .....	21-22
	21.3.12.6 operation_context .....	21-22
	21.3.12.7 result .....	21-22
	21.3.12.8 response_expected .....	21-23
	21.3.12.9 sync_scope .....	21-23
	21.3.12.10 reply_status .....	21-23
	21.3.12.11 forward_reference .....	21-24
	21.3.12.12 get_slot .....	21-24
	21.3.12.13 get_request_service_context .....	21-25
	21.3.12.14 get_reply_service_context .....	21-25
	<b>21.3.13 ClientRequestInfo Interface</b> .....	<b>21-25</b>
	21.3.13.1 target .....	1-27
	21.3.13.2 effective_target .....	21-27
	21.3.13.3 effective_profile .....	21-27
	21.3.13.4 received_exception .....	21-27
	21.3.13.5 received_exception_id .....	21-27
	21.3.13.6 get_effective_component .....	21-27
	21.3.13.7 get_effective_components .....	21-28
	21.3.13.8 get_request_policy .....	21-28
	21.3.13.9 add_request_service_context .....	21-28
	<b>21.3.14 ServerRequestInfo Interface</b> .....	<b>21-29</b>
	21.3.14.1 sending_exception .....	21-30
	21.3.14.2 object_id .....	21-30
	21.3.14.3 adapter_id .....	21-31
	21.3.14.4 target_most_derived_interface .....	21-31
	21.3.14.5 get_server_policy .....	21-31
	21.3.14.6 set_slot .....	21-31
	21.3.14.7 target_is_a .....	21-31
	21.3.14.8 add_reply_service_context .....	21-32
	<b>21.3.15 ForwardRequest Exception</b> .....	<b>21-32</b>
21.4	<b>Portable Interceptor Current</b> .....	<b>21-33</b>
	<b>21.4.1 Overview</b> .....	<b>21-33</b>
	<b>21.4.2 Obtaining the Portable Interceptor Current</b> .....	<b>21-33</b>
	<b>21.4.3 Portable Interceptor Current Interface</b> .....	<b>21-33</b>
	21.4.3.1 get_slot .....	21-34
	21.4.3.2 set_slot .....	21-34
	<b>21.4.4 Use of Portable Interceptor Current</b> .....	<b>21-34</b>
	21.4.4.1 Client-side use of PICurrent .....	21-34
	21.4.4.2 Example of PICurrent to Handle Client-side Recursion .....	21-35
	21.4.4.3 Server-side use of PICurrent .....	21-36
	21.4.4.4 Request Scope vs Thread Scope .....	21-37
	21.4.4.5 Flow of PICurrent between Scopes .....	21-37
	21.4.4.6 Notes on PICurrent and Scopes .....	21-39
21.5	<b>IOR Interceptor</b> .....	<b>21-39</b>
	<b>21.5.1 Overview</b> .....	<b>21-39</b>
	<b>21.5.2 IORInterceptor Interface</b> .....	<b>21-39</b>
	21.5.2.1 establish_components .....	21-40

# Contents

---

21.5.3	<b>IORInfo Interface</b> .....	<b>21-40</b>
21.5.3.1	get_effective_policy .....	21-40
21.5.3.2	add_ior_component .....	21-41
21.5.3.3	add_ior_component_to_profile .....	21-41
21.6	<b>PolicyFactory</b> .....	<b>21-42</b>
21.6.1	<b>PolicyFactory Interface</b> .....	<b>21-42</b>
21.6.1.1	create_policy .....	21-42
21.7	<b>Registering Interceptors</b> .....	<b>21-42</b>
21.7.1	<b>ORBInitializer Interface</b> .....	<b>21-43</b>
21.7.1.1	pre_init .....	21-43
21.7.1.2	post_init .....	21-43
21.7.2	<b>ORBInitInfo Interface</b> .....	<b>21-43</b>
21.7.2.1	DuplicateName Exception .....	21-44
21.7.2.2	InvalidName Exception .....	21-44
21.7.2.3	arguments .....	21-45
21.7.2.4	orb_id .....	21-45
21.7.2.5	codec_factory .....	21-45
21.7.2.6	register_initial_reference .....	21-45
21.7.2.7	resolve_initial_references .....	21-45
21.7.2.8	add_client_request_interceptor .....	21-45
21.7.2.9	add_server_request_interceptor .....	21-46
21.7.2.10	add_ior_interceptor .....	21-46
21.7.2.11	allocate_slot_id .....	21-46
21.7.2.12	register_policy_factory .....	21-46
21.7.3	<b>register_orb_initializer Operation</b> .....	<b>21-47</b>
21.7.3.1	Mappings of register_orb_initializer ...	21-47
21.7.4	<b>Notes about Registering Interceptors</b> .....	<b>21-49</b>
21.8	<b>Dynamic Initial References</b> .....	<b>21-49</b>
21.8.1	<b>register_initial_reference</b> .....	<b>21-49</b>
21.9	<b>Module Dynamic</b> .....	<b>21-50</b>
21.9.1	<b>NVList PIDL Represented by ParameterList IDL</b> .....	<b>21-50</b>
21.9.2	<b>ContextList PIDL Represented by ContextList IDL</b> .....	<b>21-50</b>
21.9.3	<b>ExceptionList PIDL Represented by ExceptionList IDL</b> .....	<b>21-51</b>
21.9.4	<b>Context PIDL Represented by RequestContext IDL</b> .....	<b>21-51</b>
21.10	<b>Portable Interceptor IDL</b> .....	<b>21-51</b>
22.	<b>CORBA Messaging</b> .....	<b>22-1</b>
22.1	<b>Section I - Introduction</b> .....	<b>22-2</b>
22.2	<b>Messaging Quality of Service</b> .....	<b>22-2</b>
22.2.1	<b>Rebind Support</b> .....	<b>22-5</b>
22.2.1.1	typedef short RebindMode .....	22-5
22.2.1.2	interface RebindPolicy .....	22-5
22.2.2	<b>Synchronization Scope</b> .....	<b>22-6</b>
22.2.2.1	typedef short SyncScope .....	22-6
22.2.2.2	interface SyncScopePolicy .....	22-7

# Contents

---

21.5.3	<b>IORInfo Interface</b> .....	21-40
	21.5.3.1 get_effective_policy .....	21-40
	21.5.3.2 add_ior_component .....	21-41
	21.5.3.3 add_ior_component_to_profile .....	21-41
21.6	<b>PolicyFactory</b> .....	21-42
21.6.1	<b>PolicyFactory Interface</b> .....	21-42
	21.6.1.1 create_policy .....	21-42
21.7	<b>Registering Interceptors</b> .....	21-42
21.7.1	<b>ORBInitializer Interface</b> .....	21-43
	21.7.1.1 pre_init .....	21-43
	21.7.1.2 post_init .....	21-43
21.7.2	<b>ORBInitInfo Interface</b> .....	21-43
	21.7.2.1 DuplicateName Exception .....	21-44
	21.7.2.2 InvalidName Exception .....	21-44
	21.7.2.3 arguments .....	21-45
	21.7.2.4 orb_id .....	21-45
	21.7.2.5 codec_factory .....	21-45
	21.7.2.6 register_initial_reference .....	21-45
	21.7.2.7 resolve_initial_references .....	21-45
	21.7.2.8 add_client_request_interceptor .....	21-45
	21.7.2.9 add_server_request_interceptor .....	21-46
	21.7.2.10 add_ior_interceptor .....	21-46
	21.7.2.11 allocate_slot_id .....	21-46
	21.7.2.12 register_policy_factory .....	21-46
21.7.3	<b>register_orb_initializer Operation</b> .....	21-47
	21.7.3.1 Mappings of register_orb_initializer ...	21-47
21.7.4	<b>Notes about Registering Interceptors</b> .....	21-49
21.8	<b>Dynamic Initial References</b> .....	21-49
21.8.1	<b>register_initial_reference</b> .....	21-49
21.9	<b>Module Dynamic</b> .....	21-50
21.9.1	<b>NVList PIDL Represented by ParameterList IDL</b> .....	21-50
21.9.2	<b>ContextList PIDL Represented by ContextList IDL</b> .....	21-50
21.9.3	<b>ExceptionList PIDL Represented by ExceptionList IDL</b> .....	21-51
21.9.4	<b>Context PIDL Represented by RequestContext IDL</b> .....	21-51
21.10	<b>Portable Interceptor IDL</b> .....	21-51
22.	<b>CORBA Messaging</b> .....	22-1
22.1	<b>Section I - Introduction</b> .....	22-2
22.2	<b>Messaging Quality of Service</b> .....	22-2
22.2.1	<b>Rebind Support</b> .....	22-5
	22.2.1.1 typedef short RebindMode .....	22-5
	22.2.1.2 interface RebindPolicy .....	22-5
22.2.2	<b>Synchronization Scope</b> .....	22-6
	22.2.2.1 typedef short SyncScope .....	22-6
	22.2.2.2 interface SyncScopePolicy .....	22-7



	<b>22.2.3 Request and Reply Priority</b> . . . . .	<b>22-7</b>
	22.2.3.1 struct PriorityRange . . . . .	22-7
	22.2.3.2 interface RequestPriorityPolicy . . . . .	22-7
	22.2.3.3 interface ReplyPriorityPolicy . . . . .	22-8
	<b>22.2.4 Request and Reply Timeout</b> . . . . .	<b>22-8</b>
	22.2.4.1 interface RequestStartTimePolicy . . . . .	22-8
	22.2.4.2 interface RequestEndTimePolicy . . . . .	22-9
	22.2.4.3 interface ReplyStartTimePolicy . . . . .	22-9
	22.2.4.4 interface ReplyEndTimePolicy . . . . .	22-9
	22.2.4.5 interface RelativeRequestTimeoutPolicy . . . . .	22-9
	22.2.4.6 interface RelativeRoundtripTimeout Policy . . . . .	22-10
	<b>22.2.5 Routing</b> . . . . .	<b>22-10</b>
	22.2.5.1 typedef short RoutingType . . . . .	22-10
	22.2.5.2 struct RoutingTypeRange . . . . .	22-10
	22.2.5.3 interface RoutingPolicy . . . . .	22-11
	22.2.5.4 interface MaxHopsPolicy . . . . .	22-11
	<b>22.2.6 Queue Ordering</b> . . . . .	<b>22-11</b>
	22.2.6.1 typedef short Ordering . . . . .	22-11
	22.2.6.2 interface QueueOrderPolicy . . . . .	22-12
22.3	Propagation of Messaging QoS . . . . .	22-12
	<b>22.3.1 Structures</b> . . . . .	<b>22-12</b>
	<b>22.3.2 Messaging QoS Profile Component</b> . . . . .	<b>22-13</b>
	<b>22.3.3 Messaging QoS Service Context</b> . . . . .	<b>22-13</b>
22.4	Section II - Introduction . . . . .	22-13
22.5	Running Example . . . . .	22-15
22.6	Async Operation Mapping . . . . .	22-16
	<b>22.6.1 Callback Model Signatures (sendc)</b> . . . . .	<b>22-16</b>
	22.6.1.1 Implied-IDL for Operations . . . . .	22-16
	22.6.1.2 Implied-IDL for Attributes . . . . .	22-17
	22.6.1.3 Example . . . . .	22-17
	<b>22.6.2 Polling Model Signatures (sendp)</b> . . . . .	<b>22-18</b>
	22.6.2.1 Implied-IDL for Operations . . . . .	22-18
	22.6.2.2 Implied-IDL for Attributes . . . . .	22-19
	22.6.2.3 Example . . . . .	22-19
22.7	Exception Delivery in the Callback Model . . . . .	22-20
	<b>22.7.1 Generic ExceptionHolder Value</b> . . . . .	<b>22-20</b>
	<b>22.7.2 Type-Specific ExceptionHolder Mapping</b> . . . . .	<b>22-21</b>
	<b>22.7.3 Example</b> . . . . .	<b>22-21</b>
22.8	Type-Specific ReplyHandler Mapping . . . . .	22-22
	<b>22.8.1 ReplyHandler Operations for         NO_EXCEPTION Replies</b> . . . . .	<b>22-23</b>
	<b>22.8.2 ReplyHandler Operations for Exceptional         Replies</b> . . . . .	<b>22-24</b>
	<b>22.8.3 Example</b> . . . . .	<b>22-24</b>
22.9	Generic Poller Value . . . . .	22-25
	<b>22.9.1 operation_target</b> . . . . .	<b>22-26</b>
	<b>22.9.2 operation_name</b> . . . . .	<b>22-26</b>
	<b>22.9.3 associated_handler</b> . . . . .	<b>22-26</b>

# Contents

---

22.9.4	is_from_poller .....	22-26
22.10	Type-Specific Poller Mapping .....	22-26
22.10.1	Basic Type-Specific Poller .....	22-27
22.10.1.1	Poller operations for Interface operations .....	22-27
22.10.1.2	Poller operations for Interface attributes .....	22-28
22.10.2	Persistent Type-Specific Poller .....	22-29
22.10.3	Example .....	22-29
22.11	Example Programmer Usage .....	22-30
22.11.1	Example Programmer Usage (Examples Mapped to C++) .....	22-30
22.11.2	Client-Side C++ Example for the Asynchronous Method Signatures .....	22-31
22.11.3	Client-Side C++ Example of the Callback Model .....	22-32
22.11.3.1	C++ Example of Generated ExceptionHandler .....	22-32
22.11.3.2	C++ Example of Generated ReplyHandler .....	22-32
22.11.3.3	C++ Example of User-Implemented ReplyHandler .....	22-34
22.11.3.4	C++ Example of Callback Client Program .....	22-38
22.11.4	Client-Side C++ Example of the Polling Model .....	22-39
22.11.4.1	C++ Example of Generated Poller .....	22-39
22.11.4.2	C++ Example of Polling Client Program .....	22-40
22.11.4.3	C++ Example of Using PollableSet in a Client Program .....	22-42
22.11.5	Server Side .....	22-44
22.12	Section III - Introduction .....	22-45
22.13	Routing Object References .....	22-46
22.14	Message Routing .....	22-47
22.14.1	Structures .....	22-49
22.14.1.1	MessageBody .....	22-49
22.14.1.2	RequestMessage .....	22-49
22.14.1.3	ReplyDestination .....	22-50
22.14.1.4	RequestInfo .....	22-50
22.14.2	Interfaces .....	22-51
22.14.2.1	ReplyHandler .....	22-51
22.14.2.2	Router .....	22-51
22.14.2.3	send_request .....	22-51
22.14.2.4	send_multiple_requests .....	22-51
22.14.2.5	UntypedReplyHandler .....	22-51
22.14.2.6	reply .....	22-51
22.14.2.7	PersistentRequest .....	22-52
22.14.2.8	readonly attribute reply_available .....	22-52
22.14.2.9	get_reply .....	22-52
22.14.2.10	attribute associated_handler .....	22-52
22.14.2.11	PersistentRequestRouter .....	22-53
22.14.2.12	create_persistent_request .....	22-53

22.14.3	<b>Routing Protocol</b> .....	22-53
22.14.3.1	Invoking Client .....	22-54
22.14.3.2	Initial Request Router .....	22-55
22.14.3.3	Request Routing Algorithm .....	22-55
22.14.3.4	Intermediate Request Router .....	22-56
22.14.3.5	Target Router .....	22-56
22.14.3.6	Replying to a Type-specific ReplyHandler .....	22-58
22.14.3.7	Replying to an UntypedReplyHandler .....	22-58
22.14.3.8	Handling of Service Contexts .....	22-58
22.14.3.9	Handling LOCATION_FORWARD Replies .....	22-59
22.14.3.10	Routing of Replies .....	22-59
22.14.3.11	UntypedReplyHandler .....	22-59
22.15	<b>Router Administration</b> .....	22-60
22.15.1	<b>Constants</b> .....	22-63
22.15.1.1	typedef short RegistrationState .....	22-63
22.15.2	<b>Exceptions</b> .....	22-64
22.15.2.1	exception InvalidState .....	22-64
22.15.3	<b>Valuetypes</b> .....	22-64
22.15.3.1	RetryPolicy .....	22-64
22.15.3.2	ImmediateSuspend .....	22-64
22.15.3.3	UnlimitedPing .....	22-64
22.15.3.4	LimitedPing .....	22-64
22.15.3.5	DecayPolicy .....	22-65
22.15.3.6	ResumePolicy .....	22-65
22.15.4	<b>Interfaces</b> .....	22-65
22.15.4.1	RouterAdmin .....	22-65
22.15.4.2	register_destination .....	22-65
22.15.4.3	suspend_destination .....	22-65
22.15.4.4	resume_destination .....	22-65
22.15.4.5	unregister_destination .....	22-66
23.	<b>Minimum CORBA</b> .....	23-1
23.1	Introduction .....	23-2
23.2	IDL .....	23-2
23.3	CORBA Omitted Features .....	23-2
23.4	ORB Interface Omissions .....	23-3
23.4.1	<b>ORB</b> .....	23-3
23.4.2	<b>Object</b> .....	23-4
23.4.3	<b>ConstructionPolicy</b> .....	23-4
23.5	Dynamic Invocation Interface .....	23-5
23.6	Dynamic Skeleton Interface .....	23-5
23.7	Dynamic Any .....	23-5
23.8	Interface Repository .....	23-5
23.8.1	<b>TypeCode</b> .....	23-5
23.9	Portable Object Adapter .....	23-6
23.9.1	<b>Interfaces</b> .....	23-6
23.9.1.1	POA .....	23-6

# Contents

---

23.9.1.2	Current	23-6
23.9.1.3	Policy interfaces	23-7
23.9.1.4	POAManager	23-7
23.9.1.5	AdapterActivator	23-7
23.9.1.6	ServantManagers	23-7
<b>23.9.2</b>	<b>Policies</b>	<b>23-7</b>
23.9.2.1	ThreadPolicy	23-7
23.9.2.2	LifespanPolicy	23-8
23.9.2.3	ObjectIdUniquenessPolicy	23-8
23.9.2.4	IdAssignmentPolicy	23-8
23.9.2.5	ServantRetentionPolicy	23-8
23.9.2.6	RequestProcessingPolicy	23-8
23.9.2.7	ImplicitActivationPolicy	23-9
23.10	Interoperability	23-9
23.10.1	DCE Interoperability	23-9
23.11	COM/CORBA Interworking	23-10
23.12	Interceptors	23-10
23.13	Language Mappings	23-10
23.13.1	C++ Mapping Specific Issues	23-10
23.13.2	Java Mapping Specific Issues	23-10
23.14	minimumCORBA OMG IDL	23-11
23.14.1	ORB Interface	23-11
23.14.2	Dynamic Invocation Interface	23-14
23.14.3	Dynamic Skeleton Interface	23-14
23.14.4	Dynamic Management of Any Values	23-14
23.14.5	Interface Repository	23-14
23.14.6	Portable Object Adapter	23-22
23.14.7	Interceptors	23-29
<b>24.</b>	<b>Real-Time CORBA</b>	<b>24-1</b>
24.1	Goals of the Specification	24-2
24.2	Extending CORBA	24-3
24.3	Approach to Real-Time CORBA	24-3
24.3.1	The Nature of Real-Time	24-3
24.3.2	Meeting Real-Time Requirements	24-4
24.3.3	activities	24-4
24.3.4	End-to-End Predictability	24-5
24.3.5	Management of Resources	24-6
24.4	Compatibility	24-6
24.4.1	Interoperability	24-6
24.4.2	Portability	24-7
24.4.3	CORBA - Real-Time CORBA Interworking	24-7
24.5	Real-Time CORBA Architectural Overview	24-7
24.5.1	Real-Time CORBA Modules	24-8
24.5.2	Real-Time ORB	24-8
24.5.3	Thread Scheduling	24-9

# Contents

24.5.4	Real-Time CORBA Priority .....	24-9
24.5.5	Native Priority and PriorityMappings .....	24-9
24.5.6	Real-Time CORBA Current .....	24-9
24.5.7	Priority Models .....	24-10
24.5.8	Real-Time CORBA Mutexes and Priority Inheritance 24-10	
24.5.9	Threadpools .....	24-10
24.5.10	Priority Banded Connections .....	24-11
24.5.11	Non-Multiplexed Connections .....	24-11
24.5.12	Invocation Timeouts .....	24-11
24.5.13	Client and Server Protocol Configuration .....	24-11
24.5.14	Real-Time CORBA Configuration .....	24-11
24.5.15	Scheduling Service .....	24-12
24.6	Real-Time ORB .....	24-12
24.6.1	Real-Time ORB Initialization .....	24-13
24.6.2	Real-Time CORBA System Exceptions .....	24-13
24.7	Real-Time POA .....	24-14
24.8	Native Thread Priorities .....	24-15
24.9	CORBA Priority .....	24-16
24.10	CORBA Priority Mappings .....	24-16
24.10.1	C Language binding for PriorityMapping .....	24-17
24.10.2	C++ Language binding for PriorityMapping .....	24-17
24.10.3	Ada Language binding for PriorityMapping .....	24-18
24.10.4	Java Language binding for PriorityMapping .....	24-18
24.10.5	Semantics .....	24-18
24.11	Real-Time Current .....	24-19
24.12	Real-Time CORBA Priority Models .....	24-20
24.12.1	PriorityModelPolicy .....	24-20
24.12.2	Scope of PriorityModelPolicy .....	24-21
24.12.3	Client Propagated Priority Model .....	24-22
24.12.4	Server Declared Priority Model .....	24-23
24.12.5	Setting Server Priority on a per-Object Reference Basis .....	24-23
24.13	Priority Transforms .....	24-25
24.13.1	C Language Binding for PriorityTransform .....	24-26
24.13.2	C++ Language Binding for PriorityTransform .....	24-26
24.13.3	Ada Language binding for PriorityTransform .....	24-27
24.13.4	Java Language binding for PriorityTransform .....	24-27
24.13.5	Semantics .....	24-27
24.14	Mutex Interface .....	24-28
24.15	Threadpools .....	24-29
24.15.1	Creation of Threadpool without Lanes .....	24-31
24.15.2	Creation of Threadpool with Lanes .....	24-32
24.15.3	Request Buffering .....	24-32

# Contents

---

24.15.4	Scope of ThreadpoolPolicy .....	24-33
24.16	Implicit and Explicit Binding .....	24-33
24.17	Priority Banded Connections .....	24-34
24.17.1	Scope of PriorityBandedConnectionPolicy .....	24-35
24.17.2	Binding of Priority Banded Connection .....	24-36
24.18	PrivateConnectionPolicy .....	24-37
24.19	Invocation Timeout .....	24-38
24.20	Protocol Configuration .....	24-38
24.20.1	ServerProtocolPolicy .....	24-39
24.20.2	Scope of ServerProtocolPolicy .....	24-41
24.20.3	ClientProtocolPolicy .....	24-41
24.20.4	Scope of ClientProtocolPolicy .....	24-42
24.20.5	Protocol Configuration Semantics .....	24-42
24.21	Consolidated IDL .....	24-43
24.22	Introduction .....	24-48
24.23	IDL .....	24-49
24.24	Semantics .....	24-50
24.25	Example .....	24-51
24.25.1	Server C++ Example Code .....	24-51
24.25.2	Client C++ Example Code .....	24-52
24.25.3	Explanation of Example .....	24-53
<b>25.</b>	<b>Fault Tolerant CORBA .....</b>	<b>25-1</b>
25.1	Fault Tolerant CORBA .....	25-1
25.1.1	Fault Tolerance for Diverse Applications .....	25-1
25.1.2	Objectives .....	25-2
25.1.3	Basic Concepts .....	25-3
25.1.3.1	Replication and Object Groups .....	25-3
25.1.3.2	Fault Tolerance Domains .....	25-3
25.1.3.3	Fault Tolerance Properties .....	25-3
25.1.3.4	Strong Replica Consistency .....	25-4
25.1.4	Architectural Overview .....	25-4
25.1.4.1	Fault Tolerance Property Management .....	25-6
25.1.4.2	Replication Management .....	25-6
25.1.4.3	Fault Detection and Notification .....	25-7
25.1.4.4	Logging and Recovery .....	25-7
25.1.5	Requirements .....	25-8
25.1.6	Limitations .....	25-11
25.2	Basic Fault Tolerance Mechanisms .....	25-12
25.2.1	Overview .....	25-12
25.2.2	Interoperable Object Group References .....	25-13
25.2.2.1	TAG_FT_GROUP Component .....	25-14
25.2.2.2	TAG_FT_PRIMARY Component .....	25-16
25.2.3	Interoperable Object Group Reference Operations .....	25-16

	25.2.3.1 get_interface .....	25-17
	25.2.3.2 is_a .....	25-17
	25.2.3.3 is_nil .....	25-17
	25.2.3.4 non_existent .....	25-17
	25.2.3.5 is_equivalent .....	25-17
	25.2.3.6 hash .....	25-18
	25.2.3.7 create_request .....	25-18
	25.2.3.8 get_policy .....	25-18
	25.2.3.9 get_domain_managers .....	25-18
	25.2.3.10 set_policy_overrides .....	25-18
<b>25.2.4</b>	<b>Modes of Profile Addressing .....</b>	<b>25-18</b>
	25.2.4.1 Profiles That Address Object Group Members .....	25-18
	25.2.4.2 Profiles That Address Gateways .....	25-19
	25.2.4.3 Choice of Profile Addressing Mode .....	25-19
<b>25.2.5</b>	<b>Accessing Server Object Groups .....</b>	<b>25-19</b>
	25.2.5.1 Access via IOP Directly to the Primary Member .....	25-20
	25.2.5.2 Access via IOP and a Gateway .....	25-20
	25.2.5.3 Access via a Multicast Group Communication Protocol .....	25-20
<b>25.2.6</b>	<b>Extensions to CORBA Failover Semantics .....</b>	<b>25-21</b>
<b>25.2.7</b>	<b>Most Recent Object Group Reference .....</b>	<b>25-22</b>
	25.2.7.1 FT_GROUP_VERSION Service Context .....	25-22
<b>25.2.8</b>	<b>Transparent Reinvocation .....</b>	<b>25-23</b>
	25.2.8.1 FT_REQUEST Service Context .....	25-24
	25.2.8.2 Request Duration Policy .....	25-26
	25.2.8.3 Fault Handling for GIOP Messages .....	25-26
<b>25.2.9</b>	<b>Transport Heartbeats .....</b>	<b>25-27</b>
	25.2.9.1 TAG_FT_HEARTBEAT_ENABLED Component .....	25-28
	25.2.9.2 Heartbeat Policy .....	25-28
	25.2.9.3 Heartbeat Enabled Policy .....	25-30
<b>25.3</b>	<b>Replication Management .....</b>	<b>25-31</b>
	<b>25.3.1 Overview .....</b>	<b>25-31</b>
	<b>25.3.2 Fault Tolerance Properties .....</b>	<b>25-32</b>
	25.3.2.1 ReplicationStyle .....	25-32
	25.3.2.2 MembershipStyle .....	25-33
	25.3.2.3 ConsistencyStyle .....	25-34
	25.3.2.4 FaultMonitoringStyle .....	25-35
	25.3.2.5 FaultMonitoringGranularity .....	25-35
	25.3.2.6 Factories .....	25-36
	25.3.2.7 InitialNumberReplicas .....	25-36
	25.3.2.8 MinimumNumberReplicas .....	25-36
	<b>25.3.3 FaultMonitoringIntervalAndTimeout .....</b>	<b>25-37</b>
	<b>25.3.4 CheckpointInterval .....</b>	<b>25-37</b>
	<b>25.3.5 Common Types .....</b>	<b>25-38</b>
	25.3.5.1 Identifiers .....	25-40
	25.3.5.2 Exceptions .....	25-42
	<b>25.3.6 Replication Manager .....</b>	<b>25-44</b>
	25.3.6.1 Operations .....	25-44
	<b>25.3.7 PropertyManager .....</b>	<b>25-45</b>
	25.3.7.1 Operations .....	25-46

# Contents

---

	25.3.7.2 get_properties .....	25-49
<b>25.3.8</b>	<b>ObjectGroupManager</b> .....	<b>25-49</b>
	25.3.8.1 Operations .....	25-50
<b>25.3.9</b>	<b>GenericFactory</b> .....	<b>25-56</b>
	25.3.9.1 Identifiers .....	25-59
	25.3.9.2 Operations .....	25-59
<b>25.3.10</b>	<b>Obtaining the Reference for the Replication Manager</b> .....	<b>25-61</b>
<b>25.3.11</b>	<b>Use Cases</b> .....	<b>25-61</b>
	25.3.11.1 Infrastructure-Controlled Membership Style .....	25-61
	25.3.11.2 Application-Controlled Membership Style .....	25-63
	25.3.11.3 Unreplicated Object Creation and Deletion .....	25-65
<b>25.4</b>	<b>Fault Management</b> .....	<b>25-66</b>
<b>25.4.1</b>	<b>Overview</b> .....	<b>25-66</b>
<b>25.4.2</b>	<b>Architecture</b> .....	<b>25-67</b>
	25.4.2.1 Fault Detection .....	25-68
	25.4.2.2 Fault Notification .....	25-68
	25.4.2.3 Fault Analysis .....	25-68
	25.4.2.4 Scalability .....	25-68
	25.4.2.5 Deployment of Fault Detectors .....	25-69
<b>25.4.3</b>	<b>Connecting Fault Detectors to Applications</b> ....	<b>25-70</b>
<b>25.4.4</b>	<b>Pull-Based Monitoring</b> .....	<b>25-71</b>
	25.4.4.1 PULL Fault Monitoring Style .....	25-71
	25.4.4.2 PullMonitorable Interface .....	25-71
<b>25.4.5</b>	<b>Fault Event Types</b> .....	<b>25-72</b>
	25.4.5.1 ObjectCrashFault .....	25-72
<b>25.4.6</b>	<b>Fault Notifier</b> .....	<b>25-73</b>
	25.4.6.1 Identifiers .....	25-75
	25.4.6.2 Operations .....	25-75
	25.4.6.3 Filtering .....	25-77
	25.4.6.4 Mapping of the Fault Notifier to the CosNotification Service .....	25-78
<b>25.4.7</b>	<b>Use Cases</b> .....	<b>25-79</b>
	25.4.7.1 The Fault Detector as a Fault Notification Supplier .....	25-79
	25.4.7.2 The Replication Manager as a Fault Notification Consumer .....	25-80
<b>25.5</b>	<b>Logging &amp; Recovery Management</b> .....	<b>25-81</b>
<b>25.5.1</b>	<b>Overview</b> .....	<b>25-81</b>
<b>25.5.2</b>	<b>Logging Mechanism</b> .....	<b>25-81</b>
<b>25.5.3</b>	<b>Recovery Mechanism</b> .....	<b>25-82</b>
<b>25.5.4</b>	<b>Checkpointable and Updateable Interfaces</b> ....	<b>25-84</b>
	25.5.4.1 Identifiers .....	25-85
	25.5.4.2 Exceptions .....	25-85
	25.5.4.3 Operations .....	25-86
	25.5.4.4 set_update .....	25-87
<b>25.5.5</b>	<b>Use Case</b> .....	<b>25-87</b>
	25.5.5.1 Infrastructure-Controlled Consistency Style .....	25-87



<b>26. Secure Interoperability</b>	<b>26-1</b>
26.1 Overview	26-2
26.1.1 Assumptions	26-3
26.2 Protocol Message Definitions	26-4
26.2.1 The Security Attribute Service Context Element	26-4
26.2.2 SAS context_data Message Body Types	26-5
26.2.2.1 EstablishContext Message Format	26-5
26.2.2.2 ContextError Message Format	26-7
26.2.2.3 CompleteEstablishContext Message Format	26-7
26.2.2.4 MessageInContext Message Format	26-9
26.2.3 Authorization Token Format	26-10
26.2.3.1 Extensions of the IETF AC Profile for CSiv2	26-11
26.2.4 Client Authentication Token Format	26-11
26.2.4.1 Username Password GSS Mechanism (GSSUP)	26-12
26.2.5 Identity Token Format	26-14
26.2.6 Principal Names and Distinguished Names	26-15
26.3 Security Attribute Service Protocol	26-16
26.3.1 Compound Mechanisms	26-16
26.3.1.1 Context Validation	26-17
26.3.1.2 Legend for Request Principal Interpretations	26-18
26.3.1.3 Anonymous Identity Assertion	26-19
26.3.1.4 Presumed Trust	26-19
26.3.1.5 Failed Trust Evaluations	26-19
26.3.1.6 Request Principal Interpretations	26-20
26.3.2 Session Semantics	26-21
26.3.2.1 Negotiation of Statefulness	26-21
26.3.2.2 Stateful/Reusable Contexts	26-22
26.3.3 TSS State Machine	26-23
26.3.3.1 TSS State Machine Actions	26-25
26.3.4 CSS State Machine	26-27
26.3.4.1 CSS State Machine Actions	26-30
26.3.5 ContextError Values and Exceptions	26-30
26.4 Transport Security Mechanisms	26-31
26.4.1 Transport Layer Interoperability	26-31
26.4.2 Transport Mechanism Configuration	26-31
26.4.2.1 Recommended SSL/TLS Ciphersuites	26-31
26.5 Interoperable Object References	26-32
26.5.1 Target Security Configuration	26-32
26.5.1.1 AssociationOptions Type	26-33
26.5.1.2 Transport Address	26-35
26.5.1.3 TAG_TLS_SEC_TRANS	26-35
26.5.1.4 TAG_SECIOP_SEC_TRANS	26-37
26.5.1.5 TAG_CSI_SEC_MECH_LIST	26-38
26.5.1.6 TAG_NULL_TAG	26-43
26.5.2 Client-side Mechanism Selection	26-43
26.5.3 Client-Side Requirements and Location Binding	26-44

# Contents

---

	26.5.3.1 Comments on Establishing Trust in Client	26-45
26.6	Conformance Levels	26-45
26.6.1	Conformance Level 0	26-45
	26.6.1.1 Transport-Layer Requirements	26-45
	26.6.1.2 Service Context Protocol Requirements	26-46
	26.6.1.3 Interoperable Object References (IORs)	26-47
26.6.2	Conformance Level 1	26-47
	26.6.2.1 Authorization Tokens	26-47
26.6.3	Conformance Level 2	26-47
	26.6.3.1 Authorization-Token-Based Delegation	26-47
26.6.4	Stateful Conformance	26-48
26.7	Sample Message Flows and Scenarios	26-48
26.7.1	Confidentiality, Trust in Server, and Trust in Client Established in the Connection	26-49
	26.7.1.1 Sample IOR Configuration	26-50
26.7.2	Confidentiality and Trust in Server Established in the Connection - Stateless Trust in Client Established in Service Context	26-51
	26.7.2.1 Sample IOR Configuration	26-52
26.7.3	Confidentiality, Trust in Server, and Trust in Client Established in the Connection - Stateless Trust Association Established in Service Context	26-53
	26.7.3.1 Sample IOR Configuration	26-54
	26.7.3.2 Validating the Trusted Server	26-54
	26.7.3.3 Presuming the Security of the Connection	26-55
26.7.4	Confidentiality, Trust in Server, and Trust in Client Established in the Connection - Stateless Forward Trust Association Established in Service Context	26-56
	26.7.4.1 Sample IOR Configuration	26-57
26.8	References for this Chapter	26-57
26.9	IDL	26-58
26.9.1	Module IOP	26-58
	26.9.1.1 New Types Defined for CSIv2	26-58
26.9.2	Module GSSUP - Username/Password GSSAPI Token Formats	26-58
26.9.3	Module CSI - Common Secure Interoperability	26-59
26.9.4	Module CSIIOP - CSIv2 IOR Component Tag Definitions	26-63
	Appendix A - OMG IDL Tags	A-1
	Glossary	1
	Index	1

## *Preface*

---

### *About This Document*

Under the terms of the collaboration between OMG and X/Open Co Ltd., this document is a candidate for endorsement by X/Open, initially as a Preliminary Specification and later as a full CAE Specification. The collaboration between OMG and X/Open Co Ltd. ensures joint review and cohesive support for emerging object-based specifications.

X/Open Preliminary Specifications undergo close scrutiny through a review process at X/Open before publication and are inherently stable specifications. Upgrade to full CAE Specification, after a reasonable interval, takes place following further review by X/Open. This further review considers the implementation experience of members and the full implications of conformance and branding.

### *Object Management Group*

The Object Management Group, Inc. (OMG) is an international organization supported by over 800 members, including information system vendors, software developers and users. Founded in 1989, the OMG promotes the theory and practice of object-oriented technology in software development. The organization's charter includes the establishment of industry guidelines and object management specifications to provide a common framework for application development. Primary goals are the reusability, portability, and interoperability of object-based software in distributed, heterogeneous environments. Conformance to these specifications will make it possible to develop a heterogeneous applications environment across all major hardware platforms and operating systems.

OMG's objectives are to foster the growth of object technology and influence its direction by establishing the Object Management Architecture (OMA). The OMA provides the conceptual infrastructure upon which all OMG specifications are based.

---

## *X/Open*

X/Open is an independent, worldwide, open systems organization supported by most of the world's largest information system suppliers, user organizations and software companies. Its mission is to bring to users greater value from computing, through the practical implementation of open systems. X/Open's strategy for achieving its mission is to combine existing and emerging standards into a comprehensive, integrated systems environment called the Common Applications Environment (CAE).

The components of the CAE are defined in X/Open CAE specifications. These contain, among other things, an evolving portfolio of practical application programming interfaces (APIs), which significantly enhance portability of application programs at the source code level. The APIs also enhance the interoperability of applications by providing definitions of, and references to, protocols and protocol profiles.

The X/Open specifications are also supported by an extensive set of conformance tests and by the X/Open trademark (XPG brand), which is licensed by X/Open and is carried only on products that comply with the CAE specifications.

## *Intended Audience*

The architecture and specifications described in this manual are aimed at software designers and developers who want to produce applications that comply with OMG standards for the Object Request Broker (ORB). The benefit of compliance is, in general, to be able to produce interoperable applications that are based on distributed, interoperating objects. As defined by the Object Management Group (OMG) in the *Object Management Architecture Guide*, the ORB provides the mechanisms by which objects transparently make requests and receive responses. Hence, the ORB provides interoperability between applications on different machines in heterogeneous distributed environments and seamlessly interconnects multiple object systems.

## *Context of CORBA*

The key to understanding the structure of the CORBA architecture is the Reference Model, which consists of the following components:

- **Object Request Broker**, which enables objects to transparently make and receive requests and responses in a distributed environment. It is the foundation for building applications from distributed objects and for interoperability between applications in hetero- and homogeneous environments. The architecture and specifications of the Object Request Broker are described in this manual.
- **Object Services**, a collection of services (interfaces and objects) that support basic functions for using and implementing objects. Services are necessary to construct any distributed application and are always independent of application domains. For example, the Life Cycle Service defines conventions for creating, deleting, copying, and moving objects; it does not dictate how the objects are implemented in an application. Specifications for Object Services are contained in *CORBA services: Common Object Services Specification*.

- 
- **Common Facilities**, a collection of services that many applications may share, but which are not as fundamental as the Object Services. For instance, a system management or electronic mail facility could be classified as a common facility. Information about Common Facilities will be contained in *CORBAfacilities: Common Facilities Architecture*.
  - **Application Objects**, which are products of a single vendor or in-house development group that controls their interfaces. Application Objects correspond to the traditional notion of applications, so they are not standardized by OMG. Instead, Application Objects constitute the uppermost layer of the Reference Model.

The Object Request Broker, then, is the core of the Reference Model. It is like a telephone exchange, providing the basic mechanism for making and receiving calls. Combined with the Object Services, it ensures meaningful communication between CORBA-compliant applications.

## *Associated Documents*

The CORBA documentation set includes the following books:

- *Object Management Architecture Guide* defines the OMG's technical objectives and terminology and describes the conceptual models upon which OMG standards are based. It also provides information about the policies and procedures of OMG, such as how standards are proposed, evaluated, and accepted.
- *CORBA: Common Object Request Broker Architecture and Specification* contains the architecture and specifications for the Object Request Broker.
- *CORBAservices: Common Object Services Specification* contains specifications for the Object Services.
- *CORBAfacilities: Common Facilities Architecture* contains the architecture for Common Facilities.

OMG collects information for each book in the documentation set by issuing Requests for Information, Requests for Proposals, and Requests for Comment and, with its membership, evaluating the responses. Specifications are adopted as standards only when representatives of the OMG membership accept them as such by vote.

To obtain books in the documentation set, or other OMG publications, refer to the enclosed subscription card or contact the Object Management Group, Inc. at:

OMG Headquarters  
250 First Avenue, Suite 201  
Needham, MA 02494  
USA  
Tel: +1-781-444-0404  
Fax: +1-781-444-0320  
pubs@omg.org  
<http://www.omg.org>

---

## *Definition of CORBA Compliance*

The minimum required for a CORBA-compliant system is adherence to the specifications in CORBA Core and one mapping. Each additional language mapping is a separate, optional compliance point. Optional means users aren't required to implement these points if they are unnecessary at their site, but if implemented, they must adhere to the *CORBA* specifications to be called CORBA-compliant. For instance, if a vendor supports C++, their ORB must comply with the OMG IDL to C++ binding specified in the *C++ Language Mapping Specification*.

Interoperability and Interworking are separate compliance points. For detailed information about Interworking compliance, refer to "Compliance to COM/CORBA Interworking" on page 17-34.

As described in the *OMA Guide*, the OMG's Core Object Model consists of a core and components. Likewise, the body of *CORBA* specifications is divided into core and component-like specifications. The structure of this manual reflects that division.

The *CORBA* core specifications are categorized as follows:

**CORBA Core**, as specified in Chapters 1-11

**CORBA Interoperability**, as specified in Chapters 12-16

**CORBA Interworking**, as specified in Chapters 17-21

**CORBA Quality of Service**, as specified in Chapters 22-26

---

**Note** – The *CORBA* Language Mappings have been separated from the *CORBA* Core and each language mapping is its own separate book. Refer to *CORBA* Language Mappings at the OMG Formal Document web area for this information.

---

## *Structure of This Manual*

This manual is divided into the categories of Core, Interoperability, and Interworking. These divisions reflect the compliance points of *CORBA*. In addition to this preface, *CORBA: Common Object Request Broker Architecture and Specification* contains the following chapters:

### **Core**

**Chapter 1 - The Object Model** describes the computation model that underlies the *CORBA* architecture.

**Chapter 2 - *CORBA* Overview** contains the overall structure of the ORB architecture and includes information about *CORBA* interfaces and implementations.

---

**Chapter 3 - OMG IDL Syntax and Semantics** details the OMG interface definition language (OMG IDL), which is the language used to describe the interfaces that client objects call and object implementations provide.

**Chapter 4 - ORB Interface** defines the interface to the ORB functions that do not depend on object adapters: these operations are the same for all ORBs and object implementations.

**Chapter 5 - Value Type Semantics** describes the semantics of passing an object by value, which is similar to that of standard programming languages.

**Chapter 6 - Abstract Interface Semantics** explains an IDL abstract interface, which provides the capability to defer the determination of whether an object is passed by reference or by value until runtime.

**Chapter 7 - The Dynamic Invocation Interface** details the DII, the client's side of the interface that allows dynamic creation and invocation of request to objects.

**Chapter 8 -- The Dynamic Skeleton Interface** describes the DSI, the server's-side interface that can deliver requests from an ORB to an object implementation that does not have compile-time knowledge of the type of the object it is implementing. DSI is the server's analogue of the client's Dynamic Invocation Interface (DII).

**Chapter 9 - Dynamic Management of Any Values** details the interface for the Dynamic Any type. This interface allows statically-typed programming languages such as C and Java to create or receive values of type Any without compile-time knowledge that the typer contained in the Any.

**Chapter 10 - Interface Repository** explains the component of the ORB that manages and provides access to a collection of object definitions.

**Chapter 11 - Portable Object Adapter** defines a group of IDL interfaces than an implementation uses to access ORB functions.

## **Interoperability**

**Chapter 12 - Interoperability Overview** describes the interoperability architecture and introduces the subjects pertaining to interoperability: inter-ORB bridges; general and Internet inter-ORB protocols (GIOP and IIOP); and environment-specific, inter-ORB protocols (ESIOPs).

**Chapter 13 - ORB Interoperability Architecture** introduces the framework of ORB interoperability, including information about domains; approaches to inter-ORB bridges; what it means to be compliant with ORB interoperability; and ORB Services and Requests.

**Chapter 14 - Building Inter-ORB Bridges** explains how to build bridges for an implementation of interoperating ORBs.

**Chapter 15 - General Inter-ORB Protocol** describes the general inter-ORB protocol (GIOP) and includes information about the GIOP's goals, syntax, format, transport, and object location. This chapter also includes information about the Internet inter-ORB protocol (IIOP).

---

**Chapter 16 - DCE ESIOP - Environment-Specific Inter-ORB Protocol (ESIOP)** details a protocol for the OSF DCE environment. The protocol is called the DCE Environment Inter-ORB Protocol (DCE ESIOP).

## **Interworking**

**Chapter 17 - Interworking Architecture** describes the architecture for communication between two object management systems: Microsoft's COM (including OLE) and the OMG's CORBA.

**Chapter 18 - Mapping: COM and CORBA** explains the data type and interface mapping between COM and CORBA. The mappings are described in the context of both Win16 and Win32 COM.

**Chapter 19 - Mapping: OLE Automation and CORBA** details the two-way mapping between OLE Automation (in ODL) and CORBA (in OMG IDL).

Note: Chapter 19 also includes an appendix describing solutions that vendors might implement to support existing and older OLE Automation controllers and an appendix that provides an example of how the Naming Service could be mapped to an OLE Automation interface according to the Interworking specification.

**Chapter 20 - Interoperability with non-CORBA Systems** describes the effective access to CORBA servers through DCOM and the reverse.

**Chapter 21 - Portable Interceptors** defines ORB operations that allow services such as security to be inserted in the invocation path.

## **Quality of Service (QoS)**

**Chapter 22 - CORBA Messaging** includes three general topics: Quality of Service, Asynchronous Method Invocations (to include Time-Independent or "Persistent" Requests), and the specification of interoperable Routing interfaces to support the transport of requests asynchronously from the handling of their replies.

**Chapter 23 - Minimum CORBA** describes minimumCORBA, a subset of CORBA designed for systems with limited resources.

**Chapter 24 - Real-Time CORBA** defines an optional set of extensions to CORBA tailored to equip ORBs to be used as a component of a Real-Time system.

**Chapter 25 - Fault Tolerant CORBA** describes Fault Tolerant systems, basic fault tolerance mechanisms, replication management, and logging and recovery management.

**Chapter 26 - Common Secure Interoperability** defines the CORBA Security Attribute Service (SAS) protocol and its use within the CSIv2 architecture to address the requirements of CORBA security for interoperable authentication, delegation, and privileges.



---

## *Typographical Conventions*

The type styles shown below are used in this document to distinguish programming statements from ordinary English. However, these conventions are not used in tables or section headings where no distinction is necessary.

**Helvetica bold** - OMG Interface Definition Language (OMG IDL) and syntax elements.

**Courier bold** - Programming language elements.

Helvetica - Exceptions

Terms that appear in *italics* are defined in the glossary. Italic text also represents the name of a document, specification, or other publication.

## *Acknowledgements*

The following companies submitted and/or supported parts of the specifications that were approved by the Object Management Group to become *CORBA*:

- Adiron, LLC
- Alcatel
- BEA Systems, Inc.
- BNR Europe Ltd.
- Borland International, Inc.
- Compaq Computer Corporation
- Concept Five Technologies
- Cooperative Research Centre for Distributed Systems Technology (DSTC)
- Defense Information Systems Agency
- Digital Equipment Corporation
- Ericsson
- Eternal Systems, Inc.
- Expersoft Corporation
- France Telecom
- FUJITSU LIMITED
- Genesis Development Corporation
- Gensym Corporation
- Hewlett-Packard Company
- HighComm
- Highlander Communications, L.C.
- Humboldt-University
- HyperDesk Corporation
- ICL, Plc.
- Inprise Corporation
- International Business Machines Corporation
- International Computers, Inc.

- 
- IONA Technologies, Plc.
  - Lockheed Martin Federal Systems, Inc.
  - Lucent Technologies, Inc.
  - Micro Focus Limited
  - MITRE Corporation
  - Motorola, Inc.
  - NCR Corporation
  - NEC Corporation
  - Netscape Communications Corporation
  - Nortel Networks
  - Northern Telecom Corporation
  - Novell, Inc.
  - Object Design, Inc.
  - Objective Interface Systems, Inc.
  - Object-Oriented Concepts, Inc.
  - OC Systems, Inc.
  - Open Group - Open Software Foundation
  - Oracle Corporation
  - PeerLogic, Inc.
  - Persistence Software, Inc.
  - Promia, Inc.
  - Siemens Nixdorf Informationssysteme AG
  - SPAWAR Systems Center
  - Sun Microsystems, Inc.
  - SunSoft, Inc.
  - Sybase, Inc.
  - Telefónica Investigación y Desarrollo S.A. Unipersonal
  - TIBCO, Inc.
  - Tivoli Systems, Inc.
  - Tri-Pacific Software, Inc.
  - University of California, Santa Barbara
  - University of Rhode Island
  - Visual Edge Software, Ltd.
  - Washington University

In addition to the preceding contributors, the OMG would like to acknowledge Mark Linton at Silicon Graphics and Doug Lea at the State University of New York at Oswego for their work on the C++ mapping.

## *References*

IDL Type Extensions RFP, March 1995. OMG TC Document 95-1-35.

---

The Common Object Request Broker: Architecture and Specification, Revision 2.2, February 1998.

CORBA services: Common Object Services Specification, Revised Edition, OMG TC Document 95-3-31.

COBOL Language Mapping RFP, December 1995. OMG TC document 95-12-10.

COBOL 85 ANSI X3.23-1985 / ISO 1989-1985.

IEEE Standard for Binary Floating-Point Arithmetic, ANIS/IEEE Std 754-1985.

XDR: External Data Representation Standard, RFC1832, R. Srinivasan, Sun Microsystems, August 1995.

OSF Character and Code Set Registry, OSF DCE SIG RFC 40.1 (Public Version), S. (Martin) O'Donnell, June 1994.

RPC Runtime Support For 118N Characters — Functional Specification, OSF DCE SIG RFC 41.2, M. Romagna, R. Mackey, November 1994.

X/Open System Interface Definitions, Issue 4 Version 2, 1995.



## Contents

This chapter contains the following sections.

Section Title	Page
“Elements of Interoperability”	12-1
“Relationship to Previous Versions of CORBA”	12-4
“Examples of Interoperability Solutions”	12-5
“Motivating Factors”	12-8
“Interoperability Design Goals”	12-9

ORB interoperability specifies a comprehensive, flexible approach to supporting networks of objects that are distributed across and managed by multiple, heterogeneous CORBA-compliant ORBs. The approach to “interORBability” is universal, because its elements can be combined in many ways to satisfy a very broad range of needs.

## 12.1 Elements of Interoperability

The elements of interoperability are as follows:

- ORB interoperability architecture
- Inter-ORB bridge support
- General and Internet inter-ORB Protocols (GIOPs and IOPs)

In addition, the architecture accommodates environment-specific inter-ORB protocols (ESIOPs) that are optimized for particular environments such as DCE.

### 12.1.1 ORB Interoperability Architecture

The ORB Interoperability Architecture provides a conceptual framework for defining the elements of interoperability and for identifying its compliance points. It also characterizes new mechanisms and specifies conventions necessary to achieve interoperability between independently produced ORBs.

Specifically, the architecture introduces the concepts of *immediate* and *mediated bridging* of ORB domains. The Internet Inter-ORB Protocol (IIOP) forms the common basis for broad-scope mediated bridging. The inter-ORB bridge support can be used to implement both immediate bridges and to build “half-bridges” to mediated bridge domains.

By use of bridging techniques, ORBs can interoperate without knowing any details of that ORB’s implementation, such as what particular IPC or protocols (such as ESIOPs) are used to implement the *CORBA* specification.

The IIOP may be used in bridging two or more ORBs by implementing “half bridges” that communicate using the IIOP. This approach works for both stand-alone ORBs, and networked ones that use an ESIOp.

The IIOP may also be used to implement an ORB’s internal messaging, if desired. Since ORBs are not required to use the IIOP internally, the goal of not requiring prior knowledge of each others’ implementation is fully satisfied.

### 12.1.2 Inter-ORB Bridge Support

The interoperability architecture clearly identifies the role of different kinds of domains for ORB-specific information. Such domains can include object reference domains, type domains, security domains (e.g., the scope of a *Principal* identifier), a transaction domain, and more.

Where two ORBs are in the same domain, they can communicate directly. In many cases, this is the preferable approach. This is not always true, however, since organizations often need to establish local control domains.

When information in an invocation must leave its domain, the invocation must traverse a bridge. The role of a bridge is to ensure that content and semantics are mapped from the form appropriate to one ORB to that of another, so that users of any given ORB only see their appropriate content and semantics.

The inter-ORB bridge support element specifies ORB APIs and conventions to enable the easy construction of interoperability bridges between ORB domains. Such bridge products could be developed by ORB vendors, Sieves, system integrators, or other third-parties.

Because the extensions required to support Inter-ORB Bridges are largely general in nature, do not impact other ORB operation, and can be used for many other purposes besides building bridges, they are appropriate for all ORBs to support. Other applications include debugging, interposing of objects, implementing objects with interpreters and scripting languages, and dynamically generating implementations.

---

The inter-ORB bridge support can also be used to provide interoperability with non-CORBA systems, such as Microsoft's Component Object Model (COM). The ease of doing this will depend on the extent to which those systems conform to the CORBA Object Model.

### *12.1.3 General Inter-ORB Protocol (GIOP)*

The General Inter-ORB Protocol (GIOP) element specifies a standard transfer syntax (low-level data representation) and a set of message formats for communications between ORBs. The GIOP is specifically built for ORB to ORB interactions and is designed to work directly over any connection-oriented transport protocol that meets a minimal set of assumptions. It does not require or rely on the use of higher level RPC mechanisms. The protocol is simple, scalable and relatively easy to implement. It is designed to allow portable implementations with small memory footprints and reasonable performance, with minimal dependencies on supporting software other than the underlying transport layer.

While versions of the GIOP running on different transports would not be directly interoperable, their commonality would allow easy and efficient bridging between such networking domains.

### *12.1.4 Internet Inter-ORB Protocol (IIOP)*

The Internet Inter-ORB Protocol (IIOP) element specifies how GIOP messages are exchanged using TCP/IP connections. The IIOP specifies a standardized interoperability protocol for the Internet, providing "out of the box" interoperation with other compatible ORBs based on the most popular product- and vendor-neutral transport layer. It can also be used as the protocol between half-bridges (see below).

The protocol is designed to be suitable and appropriate for use by any ORB to interoperate in Internet Protocol domains unless an alternative protocol is necessitated by the specific design center or intended operating environment of the ORB. In that sense it represents the basic inter-ORB protocol for TCP/IP environments, a most pervasive transport layer.

The IIOP's relationship to the GIOP is similar to that of a specific language mapping to OMG IDL; the GIOP may be mapped onto a number of different transports, and specifies the protocol elements that are common to all such mappings. The GIOP by itself, however, does not provide complete interoperability, just as IDL cannot be used to build complete programs. The IIOP and other similar mappings to different transports, are concrete realizations of the abstract GIOP definitions, as shown in Figure 12-1 on page 12-4.




Figure 12-1 Inter-ORB Protocol Relationships.

### 12.1.5 Environment-Specific Inter-ORB Protocols (ESIOPs)

This specification also makes provision for an open-ended set of Environment-Specific Inter-ORB Protocols (ESIOPs). Such protocols would be used for “out of the box” interoperation at user sites where a particular networking or distributing computing infrastructure is already in general use.

Because of the opportunity to leverage and build on facilities provided by the specific environment, ESIOPs might support specialized capabilities such as those relating to security and administration.

While ESIOPs may be optimized for particular environments, all ESIOP specifications will be expected to conform to the general ORB interoperability architecture conventions to enable easy bridging. The inter-ORB bridge support enables bridges to be built between ORB domains that use the IOP and ORB domains that use a particular ESIOP.

## 12.2 Relationship to Previous Versions of CORBA

The ORB Interoperability Architecture builds on Common Object Request Broker Architecture by adding the notion of ORB Services and their domains. (ORB Services are described in Section 13.2, “ORBs and ORB Services,” on page 13-3). The architecture defines the problem of ORB interoperability in terms of bridging between those domains, and defines several ways in which those bridges can be constructed. The bridges can be internal (in-line) and external (request-level) to ORBs.

APIs included in the interoperability specifications include compatible extensions to previous versions of *CORBA* to support request-level bridging:

- A Dynamic Skeleton Interface (DSI) is the basic support needed for building request-level bridges. It is the server-side analogue of the Dynamic Invocation Interface and in the same way it has general applicability beyond bridging. For information about the Dynamic Skeleton Interface, refer to the Dynamic Skeleton Interface chapter.



- APIs for managing object references have been defined, building on the support identified for the Relationship Service. The APIs are defined in Object Reference Operations in the ORB Interface chapter of this book. The Relationship Service is described in the Relationship Service specification; refer to the *CosObjectIdentity Module* section of that specification.

## 12.3 Examples of Interoperability Solutions

The elements of interoperability (Inter-ORB Bridges, General and Internet Inter-ORB Protocols, Environment-Specific Inter-ORB Protocols) can be combined in a variety of ways to satisfy particular product and customer needs. This section provides some examples.

### 12.3.1 Example 1

ORB product A is designed to support objects distributed across a network and provide “out of the box” interoperability with compatible ORBs from other vendors. In addition it allows bridges to be built between it and other ORBs that use environment-specific or proprietary protocols. To accomplish this, ORB A uses the IIOP and provides inter-ORB bridge support.

### 12.3.2 Example 2

ORB product B is designed to provide highly optimized, very high-speed support for objects located on a single machine. For example, to support thousands of Fresco GUI objects operated on at near function-call speeds. In addition, some of the objects will need to be accessible from other machines and objects on other machines will need to be infrequently accessed. To accomplish this, ORB A provides a half-bridge to support the Internet IOP for communication with other “distributed” ORBs.

### 12.3.3 Example 3

ORB product C is optimized to work in a particular operating environment. It uses a particular environment-specific protocol based on distributed computing services that are commonly available at the target customer sites. In addition, ORB C is expected to interoperate with other arbitrary ORBs from other vendors. To accomplish this, ORB C provides inter-ORB bridge support and a companion half-bridge product (supplied by the ORB vendor or some third-party) provides the connection to other ORBs. The half-bridge uses the IIOP to enable interoperability with other compatible ORBs.

### 12.3.4 Interoperability Compliance

An ORB is considered to be interoperability-compliant when it meets the following requirements:

- In the CORBA Core part of this specification, standard APIs are provided by an ORB to enable the construction of request-level inter-ORB bridges. APIs are defined by the Dynamic Invocation Interface, the Dynamic Skeleton Interface, and by the object identity operations described in the Interface Repository chapter of this book.
- An Internet Inter-ORB Protocol (IIOP) (explained in the Building Inter-ORB Bridges chapter) defines a transfer syntax and message formats (described independently as the General Inter-ORB Protocol), and defines how to transfer messages via TCP/IP connections. The IIOP can be supported natively or via a half-bridge.

Support for additional ESIOPs and other proprietary protocols is optional in an interoperability-compliant system. However, any implementation that chooses to use the other protocols defined by the CORBA interoperability specifications must adhere to those specifications to be compliant with CORBA interoperability.

Figure 12-2 on page 12-7 shows examples of interoperable ORB domains that are CORBA-compliant.

These compliance points support a range of interoperability solutions. For example, the standard APIs may be used to construct “half bridges” to the IIOP, relying on another “half bridge” to connect to another ORB. The standard APIs also support construction of “full bridges,” without using the Internet IOP to mediate between separated bridge components. ORBs may also use the Internet IOP internally. In addition, ORBs may use GIOP messages to communicate over other network protocol families (such as Novell or OSI), and provide transport-level bridges to the IIOP.

The GIOP is described separately from the IIOP to allow future specifications to treat it as an independent compliance point.




Figure 12-2 Examples of CORBA Interoperability Compliance

## 12.4 *Motivating Factors*

This section explains the factors that motivated the creation of interoperability specifications.

### 12.4.1 *ORB Implementation Diversity*

Today, there are many different ORB products that address a variety of user needs. A large diversity of implementation techniques is evident. For example, the time for a request ranges over at least 5 orders of magnitude, from a few microseconds to several seconds. The scope ranges from a single application to enterprise networks. Some ORBs have high levels of security, others are more open. Some ORBs are layered on a particular widely used protocol, others use highly optimized, proprietary protocols.

The market for object systems and applications that use them will grow as object systems are able to be applied to more kinds of computing. From application integration to process control, from loosely coupled operating systems to the information superhighway, CORBA-based object systems can be the common infrastructure.

### 12.4.2 *ORB Boundaries*

Even when it is not required by implementation differences, there are other reasons to partition an environment into different ORBs.

For security reasons, it may be important to know that it is not generally possible to access objects in one domain from another. For example, an “internet ORB” may make public information widely available, but a “company ORB” will want to restrict what information can get out. Even if they used the same ORB implementation, these two ORBs would be separate, so that the company could allow access to public objects from inside the company without allowing access to private objects from outside. Even though individual objects should protect themselves, prudent system administrators will want to avoid exposing sensitive objects to attacks from outside the company.

Supporting multiple ORBs also helps handle the difficult problem of testing and upgrading the object system. It would be unwise to test new infrastructure without limiting the set of objects that might be damaged by bugs, and it may be impractical to replace “the ORB” everywhere simultaneously. A new ORB might be tested and deployed in the same environment, interoperating with the existing ORB until either a complete switch is made or it incrementally displaces the existing one.

Management issues may also motivate partitioning an ORB. Just as networks are subdivided into domains to allow decentralized control of databases, configurations, resources, management of the state in an ORB (object reference location and translation information, interface repositories, per-object data) might also be done by creating sub-ORBs.

### 12.4.3 ORBs Vary in Scope, Distance, and Lifetime

Even in a single computing environment produced by a single vendor, there are reasons why some of the objects an application might use would be in one ORB, and others in another ORB. Some objects and services are accessed over long distances, with more global visibility, longer delays, and less reliable communication. Other objects are nearby, are not accessed from elsewhere, and provide higher quality service. By deciding which ORB to use, an implementer sets expectations for the clients of the objects.

One ORB might be used to retain links to information that is expected to accumulate over decades, such as library archives. Another ORB might be used to manage a distributed chess program in which the objects should all be destroyed when the game is over. Although while it is running, it makes sense for “chess ORB” objects to access the “archives ORB,” we would not expect the archives to try to keep a reference to the current board position.

## 12.5 Interoperability Design Goals

Because of the diversity in ORB implementations, multiple approaches to interoperability are required. Options identified in previous versions of *CORBA* include:

- *Protocol Translation*, where a gateway residing somewhere in the system maps requests from the format used by one ORB to that used by another.
- *Reference Embedding*, where invocation using a native object reference delegates to a special object whose job is to forward that invocation to another ORB.
- *Alternative ORBs*, where ORB implementations agree to coexist in the same address space so easily that a client or implementation can transparently use any of them, and pass object references created by one ORB to another ORB without losing functionality.

In general, there is no single protocol that can meet everyone's needs, and there is no single means to interoperate between two different protocols. There are many environments in which multiple protocols exist, and there are ways to bridge between environments that share no protocols.

This specification adopts a flexible architecture that allows a wide variety of ORB implementations to interoperate and that includes both bridging and common protocol elements.

The following goals guided the creation of interoperability specifications:

- The architecture and specifications should allow high-performance, small footprint, lightweight interoperability solutions.
- The design should scale, should not be unduly difficult to implement, and should not unnecessarily restrict implementation choices.

- Interoperability solutions should be able to work with any vendors' existing ORB implementations with respect to their CORBA-compliant core feature set; those implementations are diverse.
- All operations implied by the CORBA object model (i.e., the stringify and destringify operations defined on the **CORBA:ORB** pseudo-object and all the operations on **CORBA:Object**) as well as type management (e.g., narrowing, as needed by the C++ mapping) should be supported.

### *12.5.1 Non-Goals*

The following were taken into account, but were not goals:

- Support for security
- Support for future ORB Services

## Contents

This chapter contains the following sections.

Section Title	Page
“Overview”	13-1
“ORBs and ORB Services”	13-3
“Domains”	13-5
“Interoperability Between ORBs”	13-7
“Object Addressing”	13-11
“An Information Model for Object References”	13-14
“Service Context”	13-28
“Coder/Decoder Interfaces”	13-31
“Feature Support and GIOP Versions”	13-35
“Code Set Conversion”	13-36

## 13.1 Overview

The original Interoperability RFP defines interoperability as the ability for a client on ORB A to invoke an OMG IDL-defined operation on an object on ORB B, where ORB A and ORB B are independently developed. It further identifies general requirements including in particular:

- Ability for two vendors' ORBs to interoperate without prior knowledge of each other's implementation.

- Support of all ORB functionality.
- Preservation of content and semantics of ORB-specific information across ORB boundaries (for example, security).

In effect, the requirement is for invocations between client and server objects to be independent of whether they are on the same or different ORBs, and not to mandate fundamental modifications to existing ORB products.

### 13.1.1 Domains

The CORBA Object Model identifies various distribution transparencies that must be supported within a single ORB environment, such as location transparency. Elements of ORB functionality often correspond directly to such transparencies. Interoperability can be viewed as extending transparencies to span multiple ORBs.

In this architecture a *domain* is a distinct scope, within which certain common characteristics are exhibited and common rules are observed over which a distribution transparency is preserved. Thus, interoperability is fundamentally involved with transparently crossing such domain boundaries.

Domains tend to be either administrative or technological in nature, and need not correspond to the boundaries of an ORB installation. Administrative domains include naming domains, trust groups, resource management domains and other “run-time” characteristics of a system. Technology domains identify common protocols, syntaxes and similar “build-time” characteristics. In many cases, the need for technology domains derives from basic requirements of administrative domains.

Within a single ORB, most domains are likely to have similar scope to that of the ORB itself: common object references, network addresses, security mechanisms, and more. However, it is possible for there to be multiple domains of the same type supported by a given ORB: internal representation on different machine types, or security domains. Conversely, a domain may span several ORBs: similar network addresses may be used by different ORBs, type identifiers may be shared.

### 13.1.2 Bridging Domains

The abstract architecture describes ORB interoperability in terms of the translation required when an object request traverses domain boundaries. Conceptually, a mapping or *bridging mechanism* resides at the boundary between the domains, transforming requests expressed in terms of one domain’s model into the model of the destination domain.

The concrete architecture identifies two approaches to inter-ORB bridging:

- At application level, allowing flexibility and portability.
- At ORB level, built into the ORB itself.



## 13.2 ORBs and ORB Services

The ORB Core is that part of the ORB which provides the basic representation of objects and the communication of requests. The ORB Core therefore supports the minimum functionality to enable a client to invoke an operation on a server object, with (some of) the distribution transparencies required by *CORBA*.

An object request may have implicit attributes which affect the way in which it is communicated - though not the way in which a client makes the request. These attributes include security, transactional capabilities, recovery, and replication. These features are provided by "ORB Services," which will in some ORBs be layered as internal services over the core, or in other cases be incorporated directly into an ORB's core. It is an aim of this specification to allow for new ORB Services to be defined in the future, without the need to modify or enhance this architecture.

Within a single ORB, ORB services required to communicate a request will be implemented and (implicitly) invoked in a private manner. For interoperability between ORBs, the ORB services used in the ORBs, and the correspondence between them, must be identified.

### 13.2.1 The Nature of ORB Services

ORB Services are invoked implicitly in the course of application-level interactions. ORB Services range from fundamental mechanisms such as reference resolution and message encoding to advanced features such as support for security, transactions, or replication.

An ORB Service is often related to a particular transparency. For example, message encoding – the marshaling and unmarshaling of the components of a request into and out of message buffers – provides transparency of the representation of the request. Similarly, reference resolution supports location transparency. Some transparencies, such as security, are supported by a combination of ORB Services and Object Services while others, such as replication, may involve interactions between ORB Services themselves.

ORB Services differ from Object Services in that they are positioned below the application and are invoked transparently to the application code. However, many ORB Services include components which correspond to conventional Object Services in that they are invoked explicitly by the application.

Security is an example of service with both ORB Service and normal Object Service components, the ORB components being those associated with transparently authenticating messages and controlling access to objects while the necessary administration and management functions resemble conventional Object Services.

### 13.2.2 ORB Services and Object Requests

Interoperability between ORBs extends the scope of distribution transparencies and other request attributes to span multiple ORBs. This requires the establishment of relationships between supporting ORB Services in the different ORBs.

In order to discuss how the relationships between ORB Services are established, it is necessary to describe an abstract view of how an operation invocation is communicated from client to server object.

1. The client generates an operation request, using a reference to the server object, explicit parameters, and an implicit invocation context. This is processed by certain ORB Services on the client path.
2. On the server side, corresponding ORB Services process the incoming request, transforming it into a form directly suitable for invoking the operation on the server object.
3. The server object performs the requested operation.
4. Any result of the operation is returned to the client in a similar manner.

The correspondence between client-side and server-side ORB Services need not be one-to-one and in some circumstances may be far more complex. For example, if a client application requests an operation on a replicated server, there may be multiple server-side ORB service instances, possibly interacting with each other.

In other cases, such as security, client-side or server-side ORB Services may interact with Object Services such as authentication servers.

### 13.2.3 Selection of ORB Services

The ORB Services used are determined by:

- Static properties of both client and server objects; for example, whether a server is replicated.
- Dynamic attributes determined by a particular invocation context; for example, whether a request is transactional.
- Administrative policies (e.g., security).

Within a single ORB, private mechanisms (and optimizations) can be used to establish which ORB Services are required and how they are provided. Service selection might in general require negotiation to select protocols or protocol options. The same is true between different ORBs: it is necessary to agree which ORB Services are used, and how each transforms the request. Ultimately, these choices become manifest as one or more protocols between the ORBs or as transformations of requests.

In principle, agreement on the use of each ORB Service can be independent of the others and, in appropriately constructed ORBs, services could be layered in any order or in any grouping. This potentially allows applications to specify selective transparencies according to their requirements, although at this time CORBA provides no way to penetrate its transparencies.

A client ORB must be able to determine which ORB Services must be used in order to invoke operations on a server object. Correspondingly, where a client requires dynamic attributes to be associated with specific invocations, or administrative policies dictate, it must be possible to cause the appropriate ORB Services to be used on client and

server sides of the invocation path. Where this is not possible - because, for example, one ORB does not support the full set of services required - either the interaction cannot proceed or it can only do so with reduced facilities or transparencies.

### 13.3 Domains

From a computational viewpoint, the OMG Object Model identifies various distribution transparencies which ensure that client and server objects are presented with a uniform view of a heterogeneous distributed system. From an engineering viewpoint, however, the system is not wholly uniform. There may be distinctions of location and possibly many others such as processor architecture, networking mechanisms and data representations. Even when a single ORB implementation is used throughout the system, local instances may represent distinct, possibly optimized scopes for some aspects of ORB functionality.




Figure 13-1 Different Kinds of Domains can Coexist.

Interoperability, by definition, introduces further distinctions, notably between the scopes associated with each ORB. To describe both the requirements for interoperability and some of the solutions, this architecture introduces the concept of *domains* to describe the scopes and their implications.

Informally, a domain is a set of objects sharing a common characteristic or abiding by common rules. It is a powerful modelling concept which can simplify the analysis and description of complex systems. There may be many types of domains (e.g., management domains, naming domains, language domains, and technology domains).

#### 13.3.1 Definition of a Domain

Domains allow partitioning of systems into collections of components which have some characteristic in common. In this architecture a domain is a scope in which a collection of objects, said to be members of the domain, is associated with some common characteristic; any object for which the association does not exist, or is undefined, is not a member of the domain. A domain can be modeled as an object and may be itself a member of other domains.

It is the scopes themselves and the object associations or bindings defined within them which characterize a domain. This information is disjoint between domains. However, an object may be a member of several domains, of similar kinds as well as of different kinds, and so the sets of members of domains may overlap.

The concept of a domain boundary is defined as the limit of the scope in which a particular characteristic is valid or meaningful. When a characteristic in one domain is translated to an equivalent in another domain, it is convenient to consider it as traversing the boundary between the two domains.

Domains are generally either administrative or technological in nature. Examples of domains related to ORB interoperability issues are:

- Referencing domain – the scope of an object reference
- Representation domain – the scope of a message transfer syntax and protocol
- Network addressing domain – the scope of a network address
- Network connectivity domain – the potential scope of a network message
- Security domain – the extent of a particular security policy
- Type domain – the scope of a particular type identifier
- Transaction domain – the scope of a given transaction service

Domains can be related in two ways: containment, where a domain is contained within another domain, and federation, where two domains are joined in a manner agreed to and set up by their administrators.

### *13.3.2 Mapping Between Domains: Bridging*

Interoperability between domains is only possible if there is a well-defined mapping between the behaviors of the domains being joined. Conceptually, a mapping mechanism or bridge resides at the boundary between the domains, transforming requests expressed in terms of one domain's model into the model of the destination domain. Note that the use of the term "bridge" in this context is conceptual and refers only to the functionality which performs the required mappings between distinct domains. There are several implementation options for such bridges and these are discussed elsewhere.

For full interoperability, it is essential that all the concepts used in one domain are transformable into concepts in other domains with which interoperability is required, or that if the bridge mechanism filters such a concept out, nothing is lost as far as the supported objects are concerned. In other words, one domain may support a superior service to others, but such a superior functionality will not be available to an application system spanning those domains.

A special case of this requirement is that the object models of the two domains need to be compatible. This specification assumes that both domains are strictly compliant with the CORBA Object Model and the *CORBA* specifications. This includes the use of OMG IDL when defining interfaces, the use of the CORBA Core Interface Repository, and other modifications that were made to *CORBA*. Variances from this model could easily compromise some aspects of interoperability.

## 13.4 Interoperability Between ORBs

An ORB “provides the mechanisms by which objects transparently make and receive requests and responses. In so doing, the ORB provides interoperability between applications on different machines in heterogeneous distributed environments...” ORB interoperability extends this definition to cases in which client and server objects on different ORBs “transparently make and receive requests.”

Note that a direct consequence of this transparency requirement is that bridging must be bidirectional: that is, it must work as effectively for object references passed as parameters as for the target of an object invocation. Were bridging unidirectional (e.g., if one ORB could only be a client to another) then transparency would not have been provided, because object references passed as parameters would not work correctly: ones passed as “callback objects,” for example, could not be used.

Without loss of generality, most of this specification focuses on bridging in only one direction. This is purely to simplify discussions, and does not imply that unidirectional connectivity satisfies basic interoperability requirements.

### 13.4.1 ORB Services and Domains

In this architecture, different aspects of ORB functionality - ORB Services - can be considered independently and associated with different domain types. The architecture does not, however, prescribe any particular decomposition of ORB functionality and interoperability into ORB Services and corresponding domain types. There is a range of possibilities for such a decomposition:

1. The simplest model, for interoperability, is to treat all objects supported by one ORB (or, alternatively, all ORBs of a given type) as comprising one domain. Interoperability between any pair of different domains (or domain types) is then achieved by a specific all-encompassing bridge between the domains. (This is all *CORBA* implies.)
2. More detailed decompositions would identify particular domain types - such as referencing, representation, security, and networking. A core set of domain types would be pre-determined and allowance made for additional domain types to be defined as future requirements dictate (e.g., for new ORB Services).

### 13.4.2 ORBs and Domains

In many respects, issues of interoperability between ORBs are similar to those which can arise with a single type of ORB (e.g., a product). For example:

- Two installations of the ORB may be installed in different security domains, with different Principal identifiers. Requests crossing those security domain boundaries will need to establish locally meaningful Principals for the caller identity, and for any Principals passed as parameters.
- Different installations might assign different type identifiers for equivalent types, and so requests crossing type domain boundaries would need to establish locally meaningful type identifiers (and perhaps more).

Conversely, not all of these problems need to appear when connecting two ORBs of a different type (e.g., two different products). Examples include:

- They could be administered to share user visible naming domains, so that naming domains do not need bridging.
- They might reuse the same networking infrastructure, so that messages could be sent without needing to bridge different connectivity domains.

Additional problems can arise with ORBs of different types. In particular, they may support different concepts or models, between which there are no direct or natural mappings. CORBA only specifies the application level view of object interactions, and requires that distribution transparencies conceal a whole range of lower level issues. It follows that within any particular ORB, the mechanisms for supporting transparencies are not visible at the application-level and are entirely a matter of implementation choice. So there is no guarantee that any two ORBs support similar internal models or that there is necessarily a straightforward mapping between those models.

These observations suggest that the concept of an ORB (instance) is too coarse or superficial to allow detailed analysis of interoperability issues between ORBs. Indeed, it becomes clear that an ORB instance is an elusive notion: it can perhaps best be characterized as the intersection or coincidence of ORB Service domains.

### 13.4.3 Interoperability Approaches

When an interaction takes place across a domain boundary, a mapping mechanism, or bridge, is required to transform relevant elements of the interaction as they traverse the boundary. There are essentially two approaches to achieving this: mediated bridging and immediate bridging. These approaches are described in the following subsections.




Figure 13-2 Two bridging techniques, different uses of an intermediate form agreed on between the two domains.

#### 13.4.3.1 Mediated Bridging

With mediated bridging, elements of the interaction relevant to the domain are transformed, at the boundary of each domain, between the internal form of that domain and an agreed, common form.

Observations on mediated bridging are as follows:

- The scope of agreement of a common form can range from a private agreement between two particular ORB/domain implementations to a universal standard.

- There can be more than one common form, each oriented or optimized for a different purpose.
- If there is more than one possible common form, then which is used can be static (e.g., administrative policy agreed between ORB vendors, or between system administrators) or dynamic (e.g., established separately for each object, or on each invocation).
- Engineering of this approach can range from in-line specifically compiled (compare to stubs) or generic library code (such as encryption routines), to intermediate bridges to the common form.

#### 13.4.3.2 *Immediate Bridging*

With immediate bridging, elements of the interaction relevant to the domain are transformed, at the boundary of each domain, directly between the internal form of one domain and the internal form of the other.

Observations on immediate bridging are as follows:

- This approach has the potential to be optimal (in that the interaction is not mediated via a third party, and can be specifically engineered for each pair of domains) but sacrifices flexibility and generality of interoperability to achieve this.
- This approach is often applicable when crossing domain boundaries which are purely administrative (i.e., there is no change of technology). For example, when crossing security administration domains between similar ORBs, it is not necessary to use a common intermediate standard.

As a general observation, the two approaches can become almost indistinguishable when private mechanisms are used between ORB/domain implementations.

#### 13.4.3.3 *Location of Inter-Domain Functionality*

Logically, an inter-domain bridge has components in both domains, whether the mediated or immediate bridging approach is used. However, domains can span ORB boundaries and ORBs can span machine and system boundaries; conversely, a machine may support, or a process may have access to more than one ORB (or domain of a given type). From an engineering viewpoint, this means that the components of an inter-domain bridge may be dispersed or co-located, with respect to ORBs or systems. It also means that the distinction between an ORB and a bridge can be a matter of perspective: there is a duality between viewing inter-system messaging as belonging to ORBs, or to bridges.

For example, if a single ORB encompasses two security domains, the inter-domain bridge could be implemented wholly within the ORB and thus be invisible as far as ORB interoperability is concerned. A similar situation arises when a bridge between two ORBs or domains is implemented wholly within a process or system which has access to both. In such cases, the engineering issues of inter-domain bridging are

confined, possibly to a single system or process. If it were practical to implement all bridging in this way, then interactions between systems or processes would be solely within a single domain or ORB.

#### 13.4.3.4 *Bridging Level*

As noted at the start of this section, bridges may be implemented both internally to an ORB and as layers above it. These are called respectively “in-line” and “request-level” bridges.

Request-level bridges use the CORBA APIs, including the Dynamic Skeleton Interface, to receive and issue requests. However, there is an emerging class of “implicit context” which may be associated with some invocations, holding ORB Service information such as transaction and security context information, which is not at this time exposed through general purpose public APIs. (Those APIs expose only OMG IDL-defined operation parameters, not implicit ones.) Rather, the precedent set with the Transaction Service is that special purpose APIs are defined to allow bridging of each kind of context. This means that request-level bridges must be built to specifically understand the implications of bridging such ORB Service domains, and to make the appropriate API calls.

#### 13.4.4 *Policy-Mediated Bridging*

An assumption made through most of this specification is that the existence of domain boundaries should be transparent to requests: that the goal of interoperability is to hide such boundaries. However, if this were always the goal, then there would be no real need for those boundaries in the first place.

Realistically, administrative domain boundaries exist because they reflect ongoing differences in organizational policies or goals. Bridging the domains will in such cases require *policy mediation*. That is, inter-domain traffic will need to be constrained, controlled, or monitored; fully transparent bridging may be highly undesirable. Resource management policies may even need to be applied, restricting some kinds of traffic during certain periods.

Security policies are a particularly rich source of examples: a domain may need to audit external access, or to provide domain-based access control. Only a very few objects, types of objects, or classifications of data might be externally accessible through a “firewall.”

Such policy-mediated bridging requires a bridge that knows something about the traffic being bridged. It could in general be an application-specific policy, and many policy-mediated bridges could be parts of applications. Those might be organization-specific, off-the-shelf, or anywhere in between.

Request-level bridges, which use only public ORB APIs, easily support the addition of policy mediation components, without loss of access to any other system infrastructure that may be needed to identify or enforce the appropriate policies.



### 13.4.5 Configurations of Bridges in Networks

In the case of network-aware ORBs, we anticipate that some ORB protocols will be more frequently bridged to than others, and so will begin to serve the role of “backbone ORBs.” (This is a role that the IIOP is specifically expected to serve.) This use of “backbone topology” is true both on a large scale and a small scale. While a large scale public data network provider could define its own backbone ORB, on a smaller scale, any given institution will probably designate one commercially available ORB as its backbone.




Figure 13-3 An ORB chosen as a backbone will connect other ORBs through bridges, both full-bridges and half-bridges.

Adopting a backbone style architecture is a standard administrative technique for managing networks. It has the consequence of minimizing the number of bridges needed, while at the same time making the ORB topology match typical network organizations. (That is, it allows the number of bridges to be proportional to the number of protocols, rather than combinatorial.)

In large configurations, it will be common to notice that adding ORB bridges doesn't even add any new “hops” to network routes, because the bridges naturally fit in locations where connectivity was already indirect, and augment or supplant the existing network firewalls.

## 13.5 Object Addressing

The Object Model (see Chapter 1, Requests) defines an object reference as an object name that reliably denotes a particular object. An object reference identifies the same object each time the reference is used in a request, and an object may be denoted by multiple, distinct references.

The fundamental ORB interoperability requirement is to allow clients to use such object names to invoke operations on objects in other ORBs. Clients do not need to distinguish between references to objects in a local ORB or in a remote one. Providing this transparency can be quite involved, and naming models are fundamental to it.

This section discusses models for naming entities in multiple domains, and transformations of such names as they cross the domain boundaries. That is, it presents transformations of object reference information as it passes through networks of inter-ORB bridges. It uses the word “ORB” as synonymous with referencing domain; this is purely to simplify the discussion. In other contexts, “ORB” can usefully denote other kinds of domain.

### *13.5.1 Domain-relative Object Referencing*

Since CORBA does not require ORBs to understand object references from other ORBs, when discussing object references from multiple ORBs one must always associate the object reference’s domain (ORB) with the object reference. We use the notation *DO.R0* to denote an object reference *R0* from domain *DO*; this is itself an object reference. This is called “domain-relative” referencing (or addressing) and need not reflect the implementation of object references within any ORB.

At an implementation level, associating an object reference with an ORB is only important at an inter-ORB boundary; that is, inside a bridge. This is simple, since the bridge knows from which ORB each request (or response) came, including any object references embedded in it.

### *13.5.2 Handling of Referencing Between Domains*

When a bridge hands an object reference to an ORB, it must do so in a form understood by that ORB: the object reference must be in the recipient ORB’s native format. Also, in cases where that object originated from some other ORB, the bridge must associate each newly created “proxy” object reference with (what it sees as) the original object reference.

Several basic schemes to solve these two problems exist. These all have advantages in some circumstances; all can be used, and in arbitrary combination with each other, since CORBA object references are opaque to applications. The ramifications of each scheme merits attention, with respect to scaling and administration. The schemes include:

1. *Object Reference Translation Reference Embedding*: The bridge can store the original object reference itself, and pass an entirely different proxy reference into the new domain. The bridge must then manage state on behalf of each bridged object reference, map these references from one ORB’s format to the other’s, and vice versa.

2. *Reference Encapsulation*: The bridge can avoid holding any state at all by conceptually concatenating a domain identifier to the object name. Thus if a reference  $D0.R$ , originating in domain  $D0$ , traversed domains  $D1... D4$  it could be identified in  $D4$  as proxy reference  $d3.d2.d1.d0.R$ , where  $dn$  is the address of  $Dn$  relative to  $Dn+1$ .



Figure 13-4 Reference encapsulation adds domain information during bridging.

3. *Domain Reference Translation*: Like object reference translation, this scheme holds some state in the bridge. However, it supports sharing that state between multiple object references by adding a domain-based route identifier to the proxy (which still holds the original reference, as in the reference encapsulation scheme). It achieves this by providing encoded domain route information each time a domain boundary is traversed; thus if a reference  $D0.R$ , originating in domain  $D0$ , traversed domains  $D1...D4$  it would be identified in  $D4$  as  $(d3, x3).R$ , and in  $D2$  as  $(d1, x1).R$ , and so on, where  $dn$  is the address of  $Dn$  relative to  $Dn+1$ , and  $xn$  identifies the pair  $(dn-1, xn-1)$ .




Figure 13-5 Domain Reference Translation substitutes domain references during bridging.

4. *Reference Canonicalization*: This scheme is like domain reference translation, except that the proxy uses a “well-known” (e.g., global) domain identifier rather than an encoded path. Thus a reference  $R$ , originating in domain  $D0$  would be identified in other domains as  $D0.R$ .

Observations about these approaches to inter-domain reference handling are as follows:

- Naive application of reference encapsulation could lead to arbitrarily large references. A “topology service” could optimize cycles within any given encapsulated reference and eliminate the appearance of references to local objects as alien references.
- A topology service could also optimize the chains of routes used in the domain reference translation scheme. Since the links in such chains are re-used by any path traversing the same sequence of domains, such optimization has particularly high leverage.

- With the general purpose APIs defined in *CORBA*, object reference translation can be supported even by ORBs not specifically intended to support efficient bridging, but this approach involves the most state in intermediate bridges. As with reference encapsulation, a topology service could optimize individual object references. (APIs are defined by the Dynamic Skeleton Interface and Dynamic Invocation Interface)
- The chain of addressing links established with both object and domain reference translation schemes must be represented as state within the network of bridges. There are issues associated with managing this state.
- Reference canonicalization can also be performed with managed hierarchical name spaces such as those now in use on the Internet and X.500 naming.

## 13.6 *An Information Model for Object References*

This section provides a simple, powerful information model for the information found in an object reference. That model is intended to be used directly by developers of bridging technology, and is used in that role by the IIOP, described in the *General Inter-ORB Protocol* chapter, *Object References* section.

### 13.6.1 *What Information Do Bridges Need?*

The following potential information about object references has been identified as critical for use in bridging technologies:

- *Is it null?* Nulls only need to be transmitted and never support operation invocation.
- *What type is it?* Many ORBs require knowledge of an object's type in order to efficiently preserve the integrity of their type systems.
- *What protocols are supported?* Some ORBs support objrefs that in effect live in multiple referencing domains, to allow clients the choice of the most efficient communications facilities available.
- *What ORB Services are available?* As noted in Section 13.2.3, "Selection of ORB Services" on page 13-4, several different ORB Services might be involved in an invocation. Providing information about those services in a standardized way could in many cases reduce or eliminate negotiation overhead in selecting them.

### 13.6.2 *Interoperable Object References: IORs*

To provide the information above, an "Interoperable Object Reference," (IOR) data structure has been provided. This data structure need not be used internally to any given ORB, and is not intended to be visible to application-level ORB programmers. It should be used only when crossing object reference domain boundaries, within bridges.

This data structure is designed to be efficient in typical single-protocol configurations, while not penalizing multiprotocol ones.

```

module IOP {                                     // IDL

    // Standard Protocol Profile tag values

    typedef unsigned long                        ProfileId;

    struct TaggedProfile {
        ProfileId                                tag;
        sequence <octet>                        profile_data;
    };

    // an Interoperable Object Reference is a sequence of
    // object-specific protocol profiles, plus a type ID.

    struct IOR {
        string                                    type_id;
        sequence <TaggedProfile>                profiles;
    };

    // Standard way of representing multicomponent profiles.
    // This would be encapsulated in a TaggedProfile.

    typedef unsigned long ComponentId;
    struct TaggedComponent {
        ComponentId                              tag;
        sequence <octet>                        component_data;
    };
    typedef sequence<TaggedComponent> TaggedComponentSeq;
};

```

### 13.6.3 IOR Profiles

Object references have at least one *tagged profile*. Each profile supports one or more protocols and encapsulates all the basic information the protocols it supports need to identify an object. Any single profile holds enough information to drive a complete invocation using any of the protocols it supports; the content and structure of those profile entries are wholly specified by these protocols.

When a specific protocol is used to convey an object reference passed as a parameter in an IDL operation invocation (or reply), an IOR which reflects, in its contained profiles, the full protocol understanding of the operation client (or server in case of reply) may be sent. A receiving ORB which operates (based on topology and policy information available to it) on profiles rather than the received IOR as a whole, to create a derived reference for use in its own domain of reference, is placing itself as a bridge between reference domains. Interoperability inhibiting situations can arise when an orb sends an IOR with multiple profiles (using one of its supported protocols)

to a receiving orb, and that receiving orb later returns a derived reference to that object, which has had profiles or profile component data removed or transformed from the original IOR contents.

To assist in classifying behavior of ORBS in such bridging roles, two classes of IOR conformance may be associated with the conformance requirements for a given ORB interoperability protocol:

- Full IOR conformance requires that an orb which receives an IOR for an object passed to it through that ORB interoperability protocol, shall recover the original IOR, in its entirety, for passing as a reference to that object from that orb through that same protocol
- Limited-Profile IOR conformance requires that an orb which receives an IOR passed to it through a given ORB interoperability protocol, shall recover all of the standard information contained in the IOR profile for that protocol, whenever passing a reference to that object, using that same protocol, to another ORB.

---

**Note** – Conformance to IIOP versions 1.0, 1.1 and 1.2 only requires support of limited-Profile IOR conformance, specifically for the IIOP IOR profile. However, due to interoperability problems induced by Limited-Profile IOR conformance, it is now deprecated by the CORBA 2.4 specification for an orb to not support Full IOR conformance. Some future IIOP versions could require Full IOR conformance.

---

An ORB may be unable to use any of the profiles provided in an IOR for various reasons which may be broadly categorized as transient ones like temporary network outage, and non-transient ones like unavailability of appropriate protocol software in the ORB. The decision about the category of outage that causes an ORB to be unable to use any profile from an IOR is left up to the ORB. At an appropriate point, when an ORB discovers that it is unable to use any profile in an IOR, depending on whether it considers the reason transient or non-transient, it should raise the standard system exception **TRANSIENT** with standard minor code 2, or **IMP\_LIMIT** with the standard minor code 1.

Each profile has a unique numeric tag, assigned by the OMG. The ones defined here are for the IIOP (see Section 15.7.3, “IIOP IOR Profile Components” on page 15-54) and for use in “multiple component profiles.” Profile tags in the range **0x80000000** through **0xffffffff** are reserved for future use, and are not currently available for assignment.

Null object references are indicated by an empty set of profiles, and by a “Null” type ID (a string which contains only a single terminating character). Type IDs may only be “Null” in any message, requiring the client to use existing knowledge or to consult the object, to determine interface types supported. The type ID is a Repository ID identifying the interface type, and is provided to allow ORBs to preserve strong typing. This identifier is agreed on within the bridge and, for reasons outside the scope of this interoperability specification, needs to have a much broader scope to address various problems in system evolution and maintenance. Type IDs support detection of type equivalence, and in conjunction with an Interface Repository, allow processes to reason about the relationship of the type of the object referred to and any other type.

The type ID, if provided by the server, indicates the most derived type that the server wishes to publish, at the time the reference is generated. The object's actual most derived type may later change to a more derived type. Therefore, the type ID in the IOR can only be interpreted by the client as a hint that the object supports at least the indicated interface. The client can succeed in narrowing the reference to the indicated interface, or to one of its base interfaces, based solely on the type ID in the IOR, but must not fail to narrow the reference without consulting the object via the “\_is\_a” or “\_get\_interface” pseudo-operations.

ORBs claiming to support the Full-IOR conformance are required to preserve all the semantic content of any IOR (including the ordering of each profile and its components), and may only apply transformations which preserve semantics (e.g., changing Byte order for encapsulation).

For example, consider an echo operation for object references:

```
interface Echoer {Object echo(in Object o);};
```

Assume that the method body implementing this “echo” operation simply returns its argument. When a client application invokes the echo operation and passes an arbitrary object reference, if both the client and server ORBs claim support to Full IOR conformance, the reference returned by the operation is guaranteed to have not been semantically altered by either client or server ORB. That is, all its profiles will remain intact and in the same order as they were present when the reference was sent. This requirement for ORBs which claim support for Full-IOR conformance, ensures that, for example, a client can safely store an object reference in a naming service and get that reference back again later without losing information inside the reference.

### 13.6.4 Standard IOR Profiles

```
module IOP {
    const ProfileId          TAG_INTERNET_IOP = 0;
    const ProfileId          TAG_MULTIPLE_COMPONENTS = 1;
    const ProfileId          TAG_SCCP_IOP = 2;

    typedef sequence <TaggedComponent> MultipleComponentProfile;
};
```

#### 13.6.4.1 The TAG\_INTERNET\_IOP Profile

The **TAG\_INTERNET\_IOP** tag identifies profiles that support the Internet Inter-ORB Protocol. The **ProfileBody** of this profile, described in detail in Section 15:7.2, “IOP IOR Profiles” on page 15-51, contains a CDR encapsulation of a structure containing addressing and object identification information used by IOP. Version 1.1 of the **TAG\_INTERNET\_IOP** profile also includes a **sequence<TaggedComponent>** that can contain additional information supporting optional IOP features, ORB services such as security, and future protocol extensions.

Protocols other than IOP (such as ESIOPs and other GIOPs) can share profile information (such as object identity or security information) with IOP by encoding their additional profile information as components in the **TAG\_INTERNET\_IOP** profile. All **TAG\_INTERNET\_IOP** profiles support IOP, regardless of whether they also support additional protocols. Interoperable ORBs are not required to create or understand any other profile, nor are they required to create or understand any of the components defined for other protocols that might share the **TAG\_INTERNET\_IOP** profile with IOP.

The **profile\_data** for the **TAG\_INTERNET\_IOP** profile is a CDR encapsulation of the **IOP::ProfileBody\_1\_1** type, described in Section 15.7.2, "IOP IOR Profiles" on page 15-51.

#### 13.6.4.2 The **TAG\_MULTIPLE\_COMPONENTS** Profile

The **TAG\_MULTIPLE\_COMPONENTS** tag indicates that the value encapsulated is of type **MultipleComponentProfile**. In this case, the profile consists of a list of protocol components, the use of which must be specified by the protocol using this profile. This profile may be used to carry IOR components, as specified in Section 13.6.5, "IOR Components" on page 13-18.

The **profile\_data** for the **TAG\_MULTIPLE\_COMPONENTS** profile is a CDR encapsulation of the **MultipleComponentProfile** type shown above.

#### 13.6.4.3 The **TAG\_SCCP\_IOP** Profile

See the CORBA/IN Interworking specification (dtc/2000-02-02).

### 13.6.5 IOR Components

**TaggedComponents** contained in **TAG\_INTERNET\_IOP** and **TAG\_MULTIPLE\_COMPONENTS** profiles are identified by unique numeric tags using a namespace distinct from that used for profile tags. Component tags are assigned by the OMG.

Specifications of components must include the following information:

- *Component ID*: The compound tag that is obtained from OMG.
- *Structure and encoding*: The syntax of the component data and the encoding rules. If the component value is encoded as a CDR encapsulation, the IDL type that is encapsulated and the GIOP version which is used for encoding the value, if different than GIOP 1.0, must be specified as part of the component definition.
- *Semantics*: How the component data is intended to be used.
- *Protocols*: The protocol for which the component is defined, and whether it is intended that the component be usable by other protocols.
- *At most once*: whether more than one instance of this component can be included in a profile.



Specifications of protocols must describe how the components affect the protocol. In addition, a protocol definition must specify, for each TaggedComponent, whether inclusion of the component in profiles supporting the protocol is required (MANDATORY PRESENCE) or not required (OPTIONAL PRESENCE). An ORB claiming to support Full-IOR conformance shall not drop optional components, once they have been added to a profile.

### 13.6.6 Standard IOR Components

The following are standard IOR components that can be included in **TAG\_INTERNET\_IOP** and **TAG\_MULTIPLE\_COMPONENTS** profiles, and may apply to IIOP, other GIOPs, ESIOPs, or other protocols. An ORB must not drop these components from an existing IOR.

```

module IOP {
    const ComponentId TAG_ORB_TYPE = 0;
    const ComponentId TAG_CODE_SETS = 1;
    const ComponentId TAG_POLICIES = 2;
    const ComponentId TAG_ALTERNATE_IIOPI_ADDRESS = 3;

    const ComponentId TAG_ASSOCIATION_OPTIONS = 13;
    const ComponentId TAG_SEC_NAME = 14;
    const ComponentId TAG_SPKM_1_SEC_MECH = 15;
    const ComponentId TAG_SPKM_2_SEC_MECH = 16;
    const ComponentId TAG_KerberosV5_SEC_MECH = 17;
    const ComponentId TAG_CSI_ECMA_Secret_SEC_MECH = 18;
    const ComponentId TAG_CSI_ECMA_Hybrid_SEC_MECH = 19;
    const ComponentId TAG_SSL_SEC_TRANS = 20;
    const ComponentId TAG_CSI_ECMA_Public_SEC_MECH = 21;
    const ComponentId TAG_GENERIC_SEC_MECH = 22;
    const ComponentId TAG_FIREWALL_TRANS = 23;
    const ComponentId TAG_SCCP_CONTACT_INFO = 24;
    const ComponentId TAG_JAVA_CODEBASE = 25;
    const ComponentId TAG_TRANSACTION_POLICY = 26;
    const ComponentId TAG_MESSAGE_ROUTERS = 30;
    const ComponentId TAG_OTS_POLICY = 31;
    const ComponentId TAG_INV_POLICY = 32;
    const ComponentId TAG_INET_SEC_TRANS = 123;
};

```

The following additional components that can be used by other protocols are specified in the DCE ESIOP chapter of this document and *CORBAServices*, Security Service, in the Security Service for DCE ESIOP section:

```

const ComponentId TAG_COMPLETE_OBJECT_KEY = 5;
const ComponentId TAG_ENDPOINT_ID_POSITION = 6;
const ComponentId TAG_LOCATION_POLICY = 12;
const ComponentId TAG_DCE_STRING_BINDING = 100;
const ComponentId TAG_DCE_BINDING_NAME = 101;
const ComponentId TAG_DCE_NO_PIPES = 102;

```