



Common Object Request Broker Architecture (CORBA) Version 3.4

Part 3: CORBA Component Model

OMG Document Number: formal/20110201

Standard document URL: <https://www.omg.org/spec/CORBA/3.4/Components/PDF>

Machine-consumable files: <https://www.omg.org/spec/CORBA/3.3/20110201>

https://www.omg.org/spec/CORBA/3.3/20110201/ccm_dlr.idl

https://www.omg.org/spec/CORBA/3.3/20110201/ccm_dds.idl

https://www.omg.org/spec/CORBA/3.3/20110201/DDS_DefaultQos.xml

https://www.omg.org/spec/CORBA/3.3/20110201/DDS_QosProfile.xsd

Copyright © 1997-2001 Electronic Data Systems Corporation
Copyright © 1997-2001 Hewlett-Packard Company
Copyright © 1997-2001 IBM Corporation
Copyright © 1997-2001 ICON Computing
Copyright © 1997-2001 i-Logix
Copyright © 1997-2001 IntelliCorp
Copyright © 1997-2001 Microsoft Corporation
Copyright © 2011-2020 Object Management Group
Copyright © 1997-2001 ObjecTime Limited
Copyright © 1997-2001 Oracle Corporation
Copyright © 1997-2001 Platinum Technology, Inc.
Copyright © 1997-2001 Ptech Inc.
Copyright © 1997-2001 Rational Software Corporation
Copyright © 1997-2001 Reich Technologies
Copyright © 1997-2001 Softeam
Copyright © 1997-2001 Sterling Software
Copyright © 1997-2001 Taskon A/S
Copyright © 1997-2001 Unisys Corporation

Use of Specification - Terms, Conditions & Notices

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

Licenses

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.

Subject to all of the terms and conditions below, the owners of the copyright in this International Standard hereby grant you a fully-paid up, non-exclusive, nontransferable, perpetual, worldwide license (without the right to sublicense), to use this International Standard to create and distribute software and special purpose specifications that are based upon this International Standard, and to use, copy, and distribute this International Standard as provided under the Copyright Act; provided that: (1) both the copyright notice identified above and this permission notice appear on any copies of this International Standard; (2) the use of the specifications is for informational purposes and will not be copied or posted on any network computer or broadcast in any media and will not be otherwise resold or transferred for commercial purposes; and (3) no modifications are made to this International Standard. This limited permission automatically terminates without notice if you breach any of these terms or conditions. Upon termination, you will destroy immediately any copies of the specifications in your possession or control.

Patents

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.

General Use Restrictions

Any unauthorized use of this International Standard may violate copyright laws, trademark laws, and communications regulations and statutes. 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.

Disclaimer Of Warranty

WHILE THIS PUBLICATION IS BELIEVED TO BE ACCURATE, IT IS PROVIDED "AS IS" AND MAY CONTAIN ERRORS OR MISPRINTS. THE OBJECT MANAGEMENT GROUP AND THE COMPANIES LISTED ABOVE MAKE NO WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, WITH REGARD TO THIS PUBLICATION, INCLUDING BUT NOT LIMITED TO ANY WARRANTY OF TITLE OR OWNERSHIP, IMPLIED WARRANTY OF MERCHANTABILITY OR WARRANTY OF FITNESS FOR A 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 DIRECT, INDIRECT, INCIDENTAL, SPECIAL, CONSEQUENTIAL, RELIANCE OR COVER DAMAGES, INCLUDING LOSS OF PROFITS, REVENUE, DATA OR USE, INCURRED BY ANY USER OR ANY THIRD PARTY IN CONNECTION WITH THE FURNISHING, PERFORMANCE, OR USE OF THIS MATERIAL, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

The entire risk as to the quality and performance of software developed using this International Standard is borne by you. This disclaimer of warranty constitutes an essential part of the license granted to you to use this International Standard.

Restricted Rights Legend

Use, duplication or disclosure by the U.S. Government is subject to the restrictions set forth in subparagraph (c) (1) (ii) of The Rights in Technical Data and Computer Software Clause at DFARS 252.227-7013 or in subparagraph (c)(1) and (2) of the Commercial Computer Software - Restricted Rights clauses at 48 C.F.R. 52.227-19 or as specified in 48 C.F.R. 227-7202-2 of the DoD F.A.R. Supplement and its successors, or as specified in 48 C.F.R. 12.212 of the Federal Acquisition Regulations and its successors, as applicable. The specification copyright owners are as indicated above and may be contacted through the Object Management Group, 9C Medway Road, PMB 274, Medway, MA 01757, U.S.A.

Trademarks

CORBA[®], CORBA logos[®], FIBO[®], Financial Industry Business Ontology[®], FINANCIAL INSTRUMENT GLOBAL IDENTIFIER[®], IIOF[®], IMM[®], Model Driven Architecture[®], MDA[®], Object Management Group[®], OMG[®], OMG Logo[®], SoaML[®], SOAML[®], SysML[®], UAF[®], Unified Modeling Language[®], UML[®], UML Cube Logo[®], VSIPL[®], and XMI[®] are registered trademarks of the Object Management Group, Inc.

For a complete list of trademarks, see: https://www.omg.org/legal/tm_list.htm. All other products or company names mentioned are used for identification purposes only and may be trademarks of their respective owners. Compliance

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.

Software developed under the terms of this license may claim compliance or conformance with this International Standard if and only if the software compliance is of a nature fully matching the applicable compliance points as stated in the specification. Software developed only partially matching the applicable compliance points may claim only that the software was based on this International Standard, but may not claim compliance or conformance with this International Standard. In the event that testing suites are implemented or approved by Object Management Group, Inc., software developed using this International Standard may claim compliance or conformance with the specification only if the software satisfactorily completes the testing suites.

OMG's Issue Reporting Procedure

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 <https://www.omg.org>, under Documents, Report a Bug/Issue (https://www.omg.org/report_issue.htm).

Table of Contents

1 Scope	1
2 Conformance and Compliance	1
3 References	3
3.1 Normative References	3
3.2 Non-normative References	4
4 Terms and Definitions	5
4.1 Terms Defined in this International Standard	5
4.2 Keywords for Requirements Statements	7
5 Symbols (and abbreviated terms)	7
6 Component Model	9
6.1 Introduction	9
6.2 Component Model	9
6.2.1 Component Levels	9
6.2.2 Ports	9
6.2.3 Components and Facets	10
6.2.4 Component Identity	11
6.2.5 Component Homes	11
6.3 Component Definition	11
6.4 Component Declaration	11
6.4.1 Basic Components	11
6.4.2 Equivalent IDL	12
6.4.3 Component Body	13
6.5 Facets and Navigation	13
6.5.1 Equivalent IDL	13
6.5.2 Semantics of Facet References	14
6.5.3 Navigation	14

6.5.4	Provided References and Component Identity	17
6.5.5	Supported interfaces	18
6.6	Receptacles	20
6.6.1	Equivalent IDL	20
6.6.2	Behavior	21
6.6.3	Receptacles Interface	22
6.7	Events	25
6.7.1	Event types	25
6.7.2	EventConsumer Interface	26
6.7.3	Event Service Provided by Container	27
6.7.4	Event Sources—Publishers and Emitters	27
6.7.5	Publisher	28
6.7.6	Emitters	29
6.7.7	Event Sinks	30
6.7.8	Events interface	30
6.8	Homes	34
6.8.1	Equivalent Interfaces	34
6.8.2	Primary Key Declarations	36
6.8.3	Explicit Operations in Home Definitions	37
6.8.4	Home inheritance	38
6.8.5	Semantics of Home Operations	39
6.8.6	CCMHome Interface	41
6.8.7	KeylessCCMHome Interface	42
6.9	Home Finders	42
6.10	Component Configuration	44
6.10.1	Exclusive Configuration and Operational Life Cycle Phases	45
6.11	Configuration with Attributes	46
6.11.1	Attribute Configurators	46
6.11.2	Factory-based Configuration	47
6.12	Component Inheritance	49
6.12.1	CCMObject Interface	50
6.13	Conformance Requirements	51
6.13.1	A Note on Tools	53
6.13.2	Changes to Object Services	53
7	Generic Interaction Support	55
7.1	Introduction	55
7.2	Simple Generic Interaction Support	55

7.2.1 Overview	55
7.2.2 Extended Ports	56
7.2.3 Connectors	57
7.3 Generic Interaction Support with Templates	60
7.3.1 Template Support Overview	60
7.3.2 Template Modules	61
7.4 IDL3+ Grammar	63
7.5 Programming Model for Connectors	63
7.5.1 Interface CCMObject	64
7.5.2 configuration_complete	64
7.5.3 Interface KeylessCCMHome	64
7.5.4 Interface HomeConfiguration	65
7.5.5 Equivalent IDL (w.r.t Equivalent IDL section in CCM)	65
7.5.6 Connector Implementation Interfaces	65
7.6 Connector Deployment and Configuration	65
7.6.1 Integration of Connectors in D&C	66
7.6.2 Component Data Model	66
7.6.3 Execution Data Model	70
7.6.4 Connector Configuration	71
7.6.5 CCM Meta-model Extension to support Generic Interactions	71
8 OMG CIDL Syntax and Semantics	81
8.1 Overview	81
8.2 Lexical Conventions	81
8.2.1 Keywords	82
8.3 OMG CIDL Grammar	82
8.4 OMG CIDL Specification	84
8.5 Composition Definition	84
8.5.1 Life Cycle Category and Constraints	85
8.6 Home Executor Definition	85
8.7 Home Implementation Declaration	86
8.8 Storage Home Binding	87
8.9 Home Persistence Declaration	87
8.10 Executor Definition	87
8.11 Segment Definition	88

8.12 Segment Persistence Declaration	88
8.13 Facet Declaration	89
8.14 Feature Delegation Specification	89
8.15 Abstract Storage Home Delegation Specification	90
8.16 Executor Delegation Specification	91
8.17 Abstract Spec Declaration	92
8.18 Proxy Home Declaration	92
9 CCM Implementation Framework	93
9.1 Introduction	93
9.2 Component Implementation Framework (CIF) Architecture	93
9.2.1 Component Implementation Definition Language (CIDL)	93
9.2.2 Component persistence and behavior	93
9.2.3 Implementing a CORBA Component	93
9.2.4 Behavioral elements: Executors	94
9.2.5 Unit of implementation: Composition	94
9.2.6 Composition structure	95
9.2.7 Compositions with Managed Storage	101
9.2.8 Relationship between Home Executor and Abstract Storage Home	102
9.2.9 Executor Definition	115
9.2.10 Proxy Homes	122
9.2.11 Component Object References	123
9.3 Language Mapping	125
9.3.1 Overview	125
9.3.2 Common Interfaces	126
9.3.3 Mapping Rules	127
10 The Container Programming Model	135
10.1 General	135
10.2 Introduction	135
10.2.1 External API Types	136
10.2.2 Container API Type	137
10.2.3 CORBA Usage Model	137
10.2.4 Component Categories	137
10.3 The Server Programming Environment	138

10.3.1 Component Containers	138
10.3.2 CORBA Usage Model	139
10.3.3 Component Factories	140
10.3.4 Component Activation	140
10.3.5 Servant Lifetime Management	140
10.3.6 Transactions	141
10.3.7 Security	143
10.3.8 Events	143
10.3.9 Persistence	144
10.3.10 Application Operation Invocation	146
10.3.11 Component Implementations	146
10.3.12 Component Levels	146
10.3.13 Component Categories	147
10.4 Server Programming Interfaces - Basic Components	150
10.4.1 Component Interfaces	150
10.4.2 Interfaces Common to both Container API Types	151
10.4.3 Interfaces Supported by the Session Container API Type	156
10.4.4 Interfaces Supported by the Entity Container API Type	158
10.5 Server Programming Interfaces - Extended Components	160
10.5.1 Interfaces Common to both Container API Types	161
10.5.2 Interfaces Supported by the Session Container API Type	163
10.5.3 Interfaces Supported by the Entity Container API Type	164
10.6 The Client Programming Model	170
10.6.1 Component-aware Clients	171
10.6.2 Component-unaware Clients	174
11 Integrating with Enterprise JavaBeans	177
11.1 Introduction	177
11.2 Enterprise JavaBeans Compatibility Objectives and Requirements	178
11.3 CORBA Component Views for EJBs	179
11.3.1 Mapping of EJB to Component IDL Definitions	179
11.3.2 Translation of CORBA Component requests into EJB requests	183
11.3.3 Interoperability of the View	184
11.3.4 CORBA Component view Example	186
11.4 EJB Views for CORBA Components	188
11.4.1 Mapping of Component IDL to Enterprise JavaBeans specifications	188
11.4.2 Translation of EJB Requests into CORBA Component Requests	190
11.4.3 Interoperability of the View	192
11.4.4 Example	194

11.5 Compliance with the Interoperability of Integration Views	195
11.6 Comparing CCM and EJB	195
11.6.1 The Home Interfaces	196
11.6.2 The Component Interfaces	197
11.6.3 The Callback Interfaces	199
11.6.4 The Context Interfaces	200
11.6.5 The Transaction Interfaces	201
11.6.6 The Metadata Interfaces	202
12 Interface Repository Metamodel	203
12.1 Introduction	203
12.1.1 BaseIDL Package	203
12.1.2 ComponentIDL Package	214
12.2 Conformance Criteria	222
12.2.1 Conformance Points	223
12.3 MOF DTDs and IDL for the Interface Repository Metamodel	223
12.3.1 XMI DTD	223
12.3.2 IDL for the BaseIDL Package	248
12.3.3 IDL for the ComponentIDL Package	270
13 CIF Metamodel	289
13.1 CIF Package	289
13.2 Classes and Associations	289
13.2.1 ComponentImplDef	290
13.2.2 SegmentDef	291
13.2.3 ArtifactDef	291
13.2.4 Policy	291
13.2.5 HomeImplDef	292
13.3 Conformance Criteria	293
13.3.1 Conformance Points	293
13.4 MOF DTDs and IDL for the CIF Metamodel	293
13.4.1 XMI DTD	294
13.4.2 IDL for the CIF Package	294
14 Lightweight CCM Profile	301
14.1 Summary	301

14.2	Changes associated with excluding support for persistence	302
14.3	Changes associated with excluding support for introspection, navigation and type-specific operations redundant with generic operations ...	304
14.4	Changes associated with excluding support for segmentation	305
14.5	Changes associated with excluding support for transactions	306
14.6	Changes associated with excluding support for security	306
14.7	Changes associated with excluding support for configurators	307
14.8	Changes associated with excluding support for proxy homes	307
14.9	Changes associated with excluding support for home finders	307
14.10	Changes adding additional restrictions to the extended model not represented by exclusions above	308
15	Deployment PSM for CCM	309
15.1	Introduction	309
15.2	Overview	309
15.3	Definition of Meta-Concepts	310
15.3.1	Component	310
15.3.2	ImplementationArtifact	311
15.3.3	Package	311
15.4	PIM to PSM for CCM Transformation	311
15.4.1	ComponentInterfaceDescription	311
15.4.2	PlanSubcomponentPortEndpoint	312
15.4.3	Application	312
15.4.4	RepositoryManager	313
15.4.5	SatisfierProperty	313
15.5	PSM for CCM to PSM for CCM for IDL Transformation	313
15.5.1	Generic Transformation Rules	313
15.5.2	Special Transformation Rules	315
15.5.3	Mapping to IDL	316
15.6	PSM for CCM to PSM for CCM for XML Transformation	316
15.6.1	Generic Transformation Rules	316
15.6.2	Special Transformation Rules	317
15.6.3	Transformation Exceptions and Extensions	321
15.6.4	Interpretation of Relative References	322
15.6.5	Mapping to XML	323

15.7 Miscellaneous	323
15.7.1 Entry Points	323
15.7.2 Homes	324
15.7.3 Valuetype Factories	324
15.7.4 Discovery and Initialization	324
15.7.5 Location	325
15.7.6 Segmentation	325
15.8 Migration Issues	326
15.8.1 Component Implementations	326
15.8.2 Component and Assembly Packages and Metadata	326
15.8.3 Component Deployment Systems	326
15.9 Metadata Vocabulary	327
15.9.1 Implementation Selection Requirements	327
15.9.2 Monolithic Implementation Resource Requirements	327
16 Deployment IDL for CCM	329
16.1 Overview	329
17 XML Schema for CCM	343
17.1 Introduction	343
17.2 Schema	343
Annex A - Acknowledgements	365

Preface

About the Object Management Group

Founded in 1989, the Object Management Group, Inc. (OMG) is an open membership, not-for-profit computer industry standards consortium that produces and maintains computer industry specifications for interoperable, portable and reusable enterprise applications in distributed, heterogeneous environments. Membership includes Information Technology vendors, end users, government agencies and academia.

OMG member companies write, adopt, and maintain its specifications following a mature, open process. OMG's specifications implement the Model Driven Architecture® (MDA®), maximizing ROI through a full-lifecycle approach to enterprise integration that covers multiple operating systems, programming languages, middleware and networking infrastructures, and software development environments. OMG's specifications include: UML® (Unified Modeling Language™); CORBA® (Common Object Request Broker Architecture); CWM™ (Common Warehouse Meta-model); and industry-specific standards for dozens of vertical markets.

More information on the OMG is available at <https://www.omg.org/>.

OMG Specifications

As noted, OMG specifications address middleware, modeling and vertical domain frameworks. All OMG Formal Specifications are available from this URL: <https://www.omg.org/spec>

All of OMG's formal specifications may be downloaded without charge from our website. (Products implementing OMG specifications are available from individual suppliers.) Copies of specifications, available in PostScript and PDF format, may be obtained from the Specifications Catalog cited above or by contacting the Object Management Group, Inc. at:

OMG Headquarters
9C Medway Road, PMB 274
Milford, MA 01757
USA
Tel: +1-781-444-0404
Fax: +1-781-444-0320
Email: pubs@omg.org

Certain OMG specifications are also available as ISO/IEC standards. Please consult: <http://www.iso.org>

Issues

The reader is encouraged to report and technical or editing issues/problems with this specification to: https://www.omg.org/report_issue.htm

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.

Times/Times New Roman - 10 pt.: Standard body text

Helvetica/Arial - 10 pt. Bold: OMG Interface Definition Language (OMG IDL) and syntax elements.

Courier/New - 10 pt. Bold: Programming language elements.

Helvetica/Arial - 10 pt: Exceptions

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

Issues

The reader is encouraged to report any technical or editing issues/problems with this specification to https://www.omg.org/report_issue.htm.

1 Scope

This specification defines:

- The syntax and semantics of a component model (see “Component Model” on page 9), based on CORBA IDL, and a corresponding meta-model (see “Interface Repository Metamodel” on page 203).
- Generic interaction support allowing to define new interactions in CCM (see “Generic Interaction Support” on page 55).
- A language to describe the structure and state of component implementations (see “OMG CIDL Syntax and Semantics” on page 81), and a corresponding meta-model (see “CIF Metamodel” on page 289).
- A programming model for constructing component implementations (see “CCM Implementation Framework” on page 93).
- A runtime environment for component implementations (see “The Container Programming Model” on page 135).
- Interaction between components and Enterprise Java Beans (see “Integrating with Enterprise JavaBeans” on page 177).
- Meta-data for describing component-based applications, and interfaces for their deployment (see “Deployment PSM for CCM” on page 309).
- A lightweight subset of the component model, programming model and runtime environment (see “Lightweight CCM Profile” on page 301).

2 Conformance and Compliance

The following conformance points are defined:

1. A CORBA COS vendor shall provide the relevant changes to the Lifecycle, Transaction, and Security Services identified in 6.13.2, Changes to Object Services.
2. A CORBA Component vendor shall provide a conforming implementation of the Basic Level of CORBA Components. A Lightweight CORBA Component vendor shall provide a conforming implementation of the Lightweight CCM Profile as specified in item 8 below.
3. A CORBA Component vendor may provide a conforming implementation of the Extended Level of CORBA Components.
4. To be conformant at the Basic level a non-Java product shall implement (at a minimum) the following:
 - the IDL extensions and generation rules to support the client and server side component model for basic level components.
 - CIDL. The multiple segment feature of CIDL (8.11, Segment Definition) need not be supported for basic components.
 - a container for hosting basic level CORBA components.
 - the XML deployment descriptors and associated zip files for basic components.

Such implementations shall work on a CORBA ORB as defined in #1 above.

5. To be conformant at the Basic level a Java product shall implement (at a minimum):

- EJB1.1, including support for the EJB 1.1 XML DTD.
- the java to IDL mapping, also known as RMI/IIOP.
- EJB to IDL mapping as defined in 11.3.2, Translation of CORBA Component requests into EJB requests.

Such implementations shall work in a CORBA interoperable environment, including interoperable support for IIOP CORBA transactions, and CORBA security.

6. To be conformant at the extended level, a product shall implement (at a minimum) the requirements needed to achieve Basic PLUS:

- IDL extensions to support the client and server side component model for extended level components.
- A container for hosting extended level CORBA components.
- The XML deployment descriptors and associated zip files for basic and enhanced level components in the format defined in “15, Deployment PSM for CCM.”

Such implementations shall work on a CORBA ORB as defined in #1 above.

7. The Lightweight CCM profile is a conformance point based on the extended model as defined above. 14, Lightweight CCM Profile defines the specific parts of this CCM specification that are impacted and the normative specific subsetting of CCM. In summary, the following general capabilities (and associated machinery) are excluded from the extended model to define this conformance point:

- Persistence (only session and service components are supported)
- Introspection
- Navigation
- Redundancies, preferring generic over specific
- Segmentation (not allowed for session or service components)
- Transactions
- Security
- Configurators
- Proxy homes
- Home finders
- CIDL
- POA related mandates

8. A CORBA Component vendor may optionally support EJB clients interacting with CORBA Components, by implementing the IDL to EJB mapping as defined in 11.4.2, Translation of EJB Requests into CORBA Component Requests.
9. A CCM framework claiming conformance with the “Generic Interaction Support” part of this specification shall support extended ports and connectors:
 - Extensions of IDL3 to support porttype, mirrorport, and port declarations
 - Extension of IDL3 to support parameterized interfaces (template)

- Extension of D&C PSM for CCM to describe extended ports
- Extension of IDL3 to support connector declaration
- Extension of D&C PSM for CCM to deploy and configure connector fragments

3 References

3.1 Normative References

- [CORBA] Object Management Group, “Common Object Request Broker Architecture,” version 3.0.3, OMG document number formal/04-03-01. Available from <https://www.omg.org/cgi-bin/doc?formal/04-03-01>
- [D+C] Object Management Group, “*Deployment and Configuration of Component-based Distributed Applications Specification*,” version 4.0. OMG document number formal/06-04-02. Available from <https://www.omg.org/cgi-bin/doc?formal/06-04-02>
- [EJB] Java Community Process, “Enterprise JavaBeans Specification - Version 1.1” (with errata), June 8, 2000. Available from <http://jcp.org/aboutJava/communityprocess/maintenance/jsr905>
- [HTTP] R. Fielding, J. Gettys, J. Mogul, H. Frystyk, L. Masinter, P. Leach, T. Berners-Lee, RFC 2616: “*Hypertext Transfer Protocol -- HTTP/1.1*.” June 1999. Available from <http://www.ietf.org/rfc/rfc2616.txt>
- [INS] Object Management Group, “*Naming Service Specification*,” version 1.1. February 2001. OMG document number formal/01-02-65. Available from <https://www.omg.org/cgi-bin/doc?formal/01-02-65>
- [JIDL] Object Management Group, “Java™ to IDL Language Mapping Specification,” version 1.3. September 2003. OMG document number formal/03-09-04. Available from <https://www.omg.org/cgi-bin/doc?formal/03-09-04>
- [MDA] Object Management Group, “MDA Guide - Version 1.0.1” Available from <https://www.omg.org/cgi-bin/doc?omg/03-06-01>
- [MOF1] ISO/IEC 19503:2005 Information technology -- XML Metadata Interchange (XMI)
- [PSS] Object Management Group, “*Persistent State Service*,” version 2.0. September 2002. OMG document number formal/02-09-06. Available from <https://www.omg.org/cgi-bin/doc?formal/02-09-06>
- [QOS4CCM] Quality of Service for CORBA Components (formal/2008-10-02) <https://www.omg.org/spec/QOSCCM/1.1>

- [RFC2119] IETF RFC2119, “Key words for use in RFCs to Indicate Requirement Levels,” S. Bradner, March 1997. Available from <http://ietf.org/rfc/rfc2119>
- [SSS] Object Management Group, “*Security Service Specification*,” version 1.8. March 2002. OMG document number formal/02-03-11. Available from <https://www.omg.org/cgi-bin/doc?formal/02-03-11>
- [TSS] Object Management Group, “*Transaction Service Specification*,” version 1.4. September 2003. OMG document number formal/03-09-02. Available from <https://www.omg.org/cgi-bin/doc?formal/03-09-02>
- UML1] ISO/IEC 19501:2005 Information Technology - Unified Modeling Language (UML)
- [UPC] Object Management Group, “*UML™ Profile for CORBA™ Specification*,” version 1.0. Adopted specification. April 2002. OMG document number formal/02-04-01. Available from <https://www.omg.org/cgi-bin/doc?formal/02-04-01>
- [URI] T. Berners-Lee, R. Fielding, L. Masinter, RFC 2396: “*Uniform Resource Identifiers (URI): Generic Syntax*.” August 1998. Available from <http://www.ietf.org/rfc/rfc2396.txt>
- [XMI] ISO/IEC 19503:2005 Information technology -- XML Metadata Interchange (XMI)
- [XML] World Wide Web Consortium (W3C), “*Extensible Markup Language (XML)*,” version 1.0 (second edition). W3C Recommendation, October 6, 2000. Available from <http://www.w3.org/TR/REC-xml>
- [XSD] World Wide Web Consortium (W3C), “*XML Schema Part 1: Structures*.” W3C Recommendation, May 2, 2001. Available from <http://www.w3.org/TR/xmlschema-1/>
- World Wide Web Consortium (W3C), “*XML Schema Part 2: Datatypes*.” W3C Recommendation, May 2, 2001. Available from <http://www.w3.org/2001/xmlschema-2/>
- [IDL4] Object Management Group, “*Interface Definition Language*”, Version 4.2 formal/2018-01-05 <https://www.omg.org/spec/IDL/4.2>

3.2 Non-normative References

- [LCS] Object Management Group, “*Life Cycle Service Specification*,” version 1.2. September 2002. OMG document number formal/02-09-01. Available from <https://www.omg.org/cgi-bin/doc?formal/02-09-01>
- [NSS] Object Management Group, “*Notification Service Specification*,” version 1.1. October 2004. OMG document number formal/04-10-11. Available from <https://www.omg.org/cgi-bin/doc?formal/04-10-11>

4 Terms and Definitions

4.1 Terms Defined in this International Standard

Basic Component

A basic component is not allowed to inherit from other components, offer facets, receptacles, event sources or sinks. A basic component may only offer attributes.

Component

A specific, named collection of features that can be described by an IDL component definition or a corresponding structure in an Interface Repository.

Connector

Interaction entity between components. A connector is seen at design level as a connection between components and is composed of several fragments (artifacts) at execution level, to realize the interaction.

Component Home

A meta-type that acts as a manager for instances of a specified component type. Component home interfaces provide operations to manage component life cycles, and optionally, to manage associations between component instances and primary key values.

Component-aware Client

A client that is defined using the IDL extensions in the component model.

Composition

Denotes both the set of artifacts that constitute the unit of component implementation, and the definition of this aggregate entity.

Consumer

An event sink.

Container

Containers provide the run-time execution environment for CORBA component implementations. A container is a framework for integrating transactions, security, events and persistence into a component's behavior at runtime.

Emitter

An event source that can be connected to at most one consumer.

Entity Component

A CORBA component with persistent state, identity which is architecturally visible to clients through a primary key, and behavior, which may be transactional.

Equivalent IDL

The client mappings; that is, mappings of the externally-visible component features for component declarations, or home features for home declarations. Implicitly defined by a component definition in IDL.

Equivalent Interface

The interface that manifests the component's or home's surface features to clients, allowing clients to navigate among the component's facets, and to connect to the component's ports, as defined by the component's or home's equivalent IDL.

Event Sink

A named connection point into which events of a specified type may be pushed.

Event Source

A named connection point that emits events of a specified type to one or more interested event consumers, or to an event channel.

Executor

The programming artifact(s) that supply the behavior of a component or a component home.

Extended Component

Extended components may offer any type of port.

Extended Port

Consists of zero or more provided as well as zero or more required interfaces, i.e., closely resembling the UML2 specification of a port.

Facet

A distinct named interface provided by the component for client interaction. The primary vehicle through which a component exposes its functional application behavior to clients during normal execution.

Fragment

Artifact, part of the connector implementation. A fragment corresponds to one executor that can be deployed onto an execution node, co-localized with one component for which it supports the interaction provided by the connector.

Home

A home definition describes an interface for managing instances of a specified component type. A home definition implicitly defines an equivalent interface, which can be described in terms of IDL. A home may be associated with a primary key specification.

Monolithic Executor

An executor consisting of a single artifact.

Multiplex Receptacle

A specialization of a receptacle that allows multiple simultaneous connections.

Primary Key

Primary key values uniquely identify component instances within the scope of the home that manages them.

Port

A surface feature through which clients and other elements of an application environment may interact with a component. The component model supports four basic kinds of ports: facets, receptacles, event sources, event sinks and attributes.

Process Component

A CORBA component with persistent state which is not visible to the client, persistent identity, and behavior, which may be transactional.

Proxy Home

Implements the component home interface specified by a composition definition, but the implementation is not required to be collocated with the container where the components managed by the home are activated.

Publisher

An event source that can be connected to an arbitrary number of consumers, who are said to *subscribe* to the publisher event source.

Receptacle

A named connection point that describes the component's ability to use a reference supplied by some external agent.

Segmented Executor

A set of physically distinct artifacts, a physical partition of the executor. Each segment encapsulates independent state and is capable of being independently activated. Each segment provides at least one facet.

Service Component

A CORBA component with behavior, no state, and no identity.

Session Component

A CORBA component with behavior, transient state, and identity (which is not persistent).

Simplex Receptacle

A specialization of a receptacle that only allows a single connection at a given time.

4.2 Keywords for Requirements Statements

The keywords “must,” “must not,” “shall,” “shall not,” “should,” “should not,” and “may” in this specification are to be interpreted as described in [RFC 2119].

5 Symbols (and abbreviated terms)

API — Application Programming Interface

CCM — CORBA Component Model

CIDL — Component Implementation Definition Language

CIF — Component Implementation Framework

CMT — Container-managed Transaction

CORBA — Common Object Request Broker Architecture

COS — Common Object Services
CRUD — Create, Read, Update, Delete
DII — Dynamic Invocation Interface
DTD — Document Type Definition
EJB — Enterprise Java Beans
GIOP — General Inter-ORB Protocol
IDL — Interface Definition Language
IIOP — Internet Inter-ORB Protocol
IR — Interface Repository
JDK — Java Development Kit
JNDI — Java Naming and Directory Interface
JTA — Java Transaction API
MOF — Meta Object Facility
OMG — Object Management Group
ORB — Object Request Broker
PIM — Platform Independent Model (see [MDA])
POA — Portable Object Adapter
PSDL — Persistent State Definition Language
PSM — Platform Specific Model (see [MDA])
RMI — Remote Method Invocation
SECIOP — Secure Inter-ORB Protocol
SMT — Self-managed Transaction
SSL — Secure Sockets Layer
UML — Unified Modeling Language
URI — Uniform Resource Identifier
URL — Uniform Resource Locator
XMI — XML Metadata Interchange
XML — Extensible Markup Language

6 Component Model

6.1 Introduction

This clause describes the semantics of the CORBA Component Model (CCM) and the conformance requirements for vendors.

6.2 Component Model

Component is a basic meta-type in CORBA. The component meta-type is an extension and specialization of the object meta-type. Component types are specified in IDL and represented in the Interface Repository. A component is denoted by a component reference, which is represented by an object reference. Correspondingly, a component definition is a specialization and extension of an interface definition.

A component type is a specific, named collection of features that can be described by an IDL component definition or a corresponding structure in an Interface Repository. Although the current specification does not attempt to provide mechanisms to support formal semantic descriptions associated with component definitions, they are designed to be associated with a single well-defined set of behaviors. Although there may be several realizations of the component type for different run-time environments (e.g., OS/hardware platforms, languages, etc.), they should all behave consistently. As an abstraction in a type system, a component type is instantiated to create concrete entities (instances) with state and identity.

A component type encapsulates its internal representation and implementation. Although the component specification includes standard frameworks for component implementation, these frameworks, and any assumptions that they might entail, are completely hidden from clients of the component.

6.2.1 Component Levels

There are two levels of components: *basic* and *extended*. Both are managed by component homes, but they differ in the capabilities they can offer. Basic components essentially provide a simple mechanism to “componentize” a regular CORBA object. Extended components, on the other hand, provide a richer set of functionality.

A basic component is very similar in functionality to an EJB as defined in the Enterprise JavaBeans 1.1 specification. This allows much easier mapping and integration at this level.

6.2.2 Ports

Components support a variety of surface features through which clients and other elements of an application environment may interact with a component. These surface features are called *ports*. The component model supports four basic kinds of ports:

- **Facets**, which are distinct named interfaces provided by the component for client interaction.
- **Receptacles**, which are named connection points that describe the component’s ability to use a reference supplied by some external agent.
- **Event sources**, which are named connection points that emit events of a specified type to one or more interested event consumers, or to an event channel.

- **Event sinks**, which are named connection points into which events of a specified type may be pushed.
- **Attributes**, which are named values exposed through accessor and mutator operations. Attributes are primarily intended to be used for component configuration, although they may be used in a variety of other ways.

Basic components are not allowed to offer facets, receptacles, event sources, and sinks. They may only offer attributes. Extended components may offer any type of port.

6.2.3 Components and Facets

A component can provide multiple object references, called *facets*, which are capable of supporting distinct (i.e., unrelated by inheritance) IDL interfaces. The component has a single distinguished reference whose interface conforms to the component definition. This reference supports an interface, called the component's *equivalent interface*, that manifests the component's surface features to clients. The equivalent interface allows clients to navigate among the component's facets, and to connect to the component's ports.

Basic components cannot support facets, therefore attempts to navigate to other facets will always fail. The equivalent interface of a basic component is the only object available with which a client may interact.

The other interfaces provided by the component are referred to as *facets*. Figure 6.1 illustrates the relationship between the component and its facets.

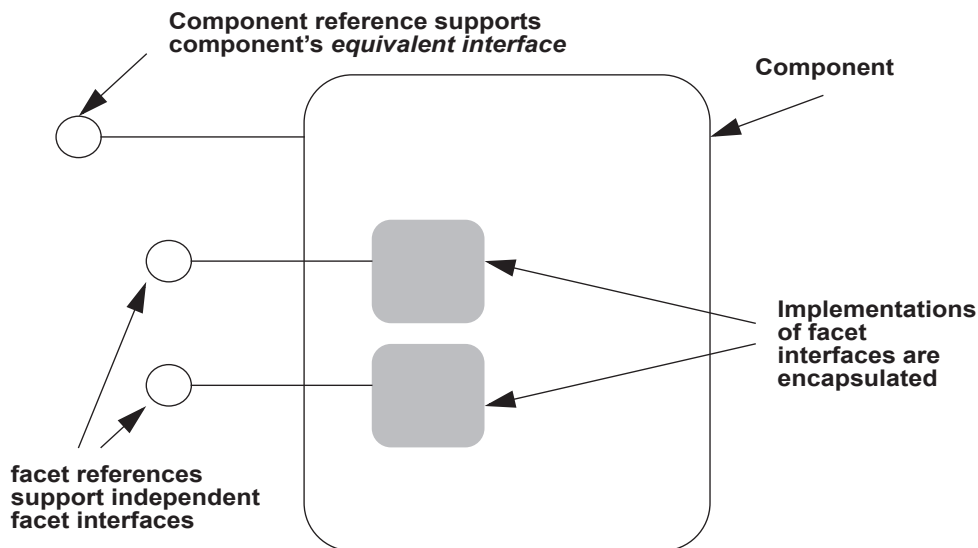


Figure 6.1- Component Interfaces and Facets

The relationship between the component and its facets is characterized by the following observations:

- The implementations of the facet interfaces are encapsulated by the component, and considered to be “parts” of the component. The internal structure of a component is opaque to clients.
- Clients can navigate from any facet to the component equivalent interface, and can obtain any facet from the component equivalent interface.

- Clients can reliably determine whether any two references belong to the same component instance.
- The life cycle of a facet is bounded by the life cycle of its owning component.

6.2.4 Component Identity

A component instance is identified primarily by its component reference, and secondarily by its set of facet references (if any). The component model provides operations to determine whether two references belong to the same component instance, and (as mentioned above) operations to navigate among a component's references. The definition of "same" component instance is ultimately up to the component implementor, in that they may provide a customized implementation of this operation. However, a component framework shall provide standard implementations that constitute *de facto* definitions of "sameness" when they are employed.

Components may also be associated with *primary key values* by a component home. Primary keys are data values exposed to the component's clients that may be used in the context of a component home to identify component instances and obtain references for them. Primary keys are not features of components themselves; the association between a component instance and a particular primary key value is maintained by the home that manages the component.

6.2.5 Component Homes

A *component home* is meta-type that acts as a manager for instances of a specified component type. Component home interfaces provide operations to manage component life cycles, and optionally, to manage associations between component instances and primary key values. A component home may be thought of as a manager for the extent of a type (within the scope of a container). A home must be declared for every component declaration.

Component types are defined in isolation, independent of home types. A home definition, however, must specify exactly one component type that it manages. Multiple different home types can manage the same component type, though they cannot manage the same set of component instances.

At execution time, a component instance is managed by a single home object of a particular type. The operations on the home are roughly equivalent to static or class methods in object-oriented programming languages.

6.3 Component Definition

A component definition in IDL implicitly defines an interface that supports the features defined in the component definition body. It extends the concept of an interface definition to support features that are not supported in interfaces. Component definitions also differ from interface definitions in that they support only single inheritance from other component types.

The IDL grammar for components may be found in Part 1- Interfaces, *OMG IDL Syntax and Semantics*.

6.4 Component Declaration

6.4.1 Basic Components

Basic components cannot avail themselves of certain features in the model. In particular, they cannot inherit from other components, nor can they provide or use interfaces, or make any event declarations. A basic component is declared using a restricted version of a **<component_dcl>**. See CORBA (Part 1), *OMG IDL Syntax and Semantics* clause, "Component" sub clause for the syntax.

To avoid ambiguity between basic and extended definitions, any component declaration that matches the following pattern is a basic component:

```
“component” <identifier> [<supported_interface_spec>]  
“{“ {<attr_dcl> “,”}* “}”
```

Ideally the syntax should explicitly represent these rules. However this can only be achieved by introducing a new keyword to distinguish between basic and extended components. It was felt that an extra keyword would cause problems in the future, as the distinction between basic and extended components gets blurred. This blurring may occur due to future development of both the CORBA Component Model and the Enterprise JavaBeans specifications.

6.4.2 Equivalent IDL

The client mappings; that is, mappings of the externally-visible component features for component declarations are described in terms of *equivalent IDL*.

As described above, the component meta-type is a specialization of the interface meta-type. Each component definition has a corresponding *equivalent interface*. In programming language mappings, components are denoted by object references that support the equivalent interface implied by the component definition.

Since basic components are essentially a profile, no specific rules are defined for them.

6.4.2.1 Simple declaration

For a component declaration with the following form:

```
component component_name { ... };
```

the equivalent interface shall have the following form:

```
interface component_name  
: Components::CCMObject { ... };
```

6.4.2.2 Supported interfaces

For a component declaration with the following form:

```
component <component_name>  
supports <interface_name_1>, <interface_name_2> { ... };
```

the equivalent interface shall have the following form:

```
interface <component_name>  
: Components::CCMObject,  
<interface_name_1>, <interface_name_2> { ... };
```

Supported interfaces are described in detail in “Supported interfaces” on page 18.

6.4.2.3 Inheritance

For a component declaration with the following form:

```
component <component_name> : <base_name> { ... };
```

the equivalent interface shall have the following form:

```
interface <component_name> : <base_name> { ... }
```

6.4.2.4 Inheritance and supported interfaces

For a component declaration with the following form:

```
component <component_name> : <base_name>  
supports <interface_name_1>, <interface_name_2> { ... };
```

the equivalent interface shall have the following form:

```
interface <component_name>  
: <base_name>, <interface_name_1>, <interface_name_2> { ... };
```

6.4.3 Component Body

A component forms a naming scope, nested within the scope in which the component is declared.

Declarations for facets, receptacles, event sources, event sinks and attributes all map onto operations on the component's equivalent interface. These declarations and their meanings are described in detail below.

6.5 Facets and Navigation

A component type may provide several independent interfaces to its clients in the form of facets. Facets are intended to be the primary vehicle through which a component exposes its functional application behavior to clients during normal execution. A component may exhibit zero or more facets.

6.5.1 Equivalent IDL

Facet declarations imply operations on the component interface that provide access to the provided interfaces by their names. A facet declaration of the following form:

```
provides <interface_type> <name>;
```

results in the following operation defined on the equivalent interface:

```
<interface_type> provide_<name> ();
```

The mechanisms for navigating among a component's facets are described in "Navigation" on page 14. The relationships between the component identity and the facet references, and assumptions regarding facet references, are described in "Provided References and Component Identity" on page 17. The implementation of navigation operations are provided by the component implementation framework in generated code; the user-provided implementation of a component type is not responsible for navigation operations. The responsibilities of the component servant framework for supporting navigation operations are described in detail in Clause 8 *OMG CIDL Syntax and Semantics*.

6.5.2 Semantics of Facet References

Clients of a component instance can obtain a reference to a facet by invoking the **provide_<name>** operation on the equivalent interface corresponding to the **provides** declaration in the component definition. The component implementation is responsible for guaranteeing the following behaviors:

- In general, a component instance shall be prepared to return object references for facets throughout the instance's life cycle. A component implementation may, as part of its advertised behavior, return a nil object reference as the result of a **provide_<name>** operation.
- An object reference returned by a **provide_<name>** operation shall support the interface associated with the corresponding **provides** declaration in the component definition. Specifically, when the **_is_a** operation is invoked on the object reference with the **RepositoryId** of the provided interface type, the result shall be **TRUE**, and legal operations of the facet interface shall be able to be invoked on the object reference. If the type specified in the **provides** declaration is **Object**, then there are no constraints on the interface types supported by the reference.

A facet reference provided by a component may support additional interfaces, such as interfaces derived from the declared type, as long as the stated contract is satisfied.

- Facet references must behave properly with respect to component identity and navigation, as defined in “Provided References and Component Identity” on page 17 and “Navigation” below.

6.5.3 Navigation

Navigation among a component's facets may be accomplished in the following ways:

- A client may navigate from any facet reference to the component that provides the reference via **CORBA::Object::get_component**.
- A client may navigate from the component interface to any facet using the generated **provide_<name>** operations on the equivalent interface.
- A client may navigate from the component interface to any facet using the generic **provide_facet** operation on the **Navigation** interface (inherited by all component interfaces through **Components::CCMObject**). Other operations on the **Navigation** interface (i.e., **get_all_facets** and **get_named_facets**) return multiple references, and can also be used for navigation. When using generic navigation operations on **Navigation**, facets are identified by string values that contain their declared names.
- A client may navigate from a facet interface that derives from the **Navigation** interface directly to any other facet on the same component, using **provide_facet**, **get_all_facets**, and **get_named_facets**.
- For components, such as basic components, that do not provide interfaces, only the generic navigation operations are available on the equivalent interface. The behavior of these operations, where there are no facets to navigate to, is defined below.

The detailed descriptions of these mechanisms follow.

6.5.3.1 get_component()

```
module CORBA {  
    interface Object { // PIDL  
        ...  
    }
```

```

        Object get_component ( );
    };
};

```

If the target object reference is itself a component reference (i.e., it denotes the component itself), the **get_component** operation returns the same reference (or another equivalent reference). If the target object reference is a facet reference, the **get_component** operation returns an object reference for the component. If the target reference is neither a component reference nor a provided reference, **get_component** returns a nil reference.

Implementation of get_component

As with other operations on **CORBA::Object**, **get_component** is implemented as a request to the target object. Following the pattern of other **CORBA::Object** operations (i.e., **_interface**, **_is_a**, and **_non_existent**) the operation name in GIOP request corresponding to **get_component** shall be “**_component**.” An implementation of **get_component** is a required element of the CORBA core, even if the ORB does not provide an implementation of CORBA components. Thus component vendors that are not also ORB vendors can rely on the availability of this capability in a compliant ORB.

6.5.3.2 Component-specific provide operations

The **provide_<name>** operation implicitly defined by a **provides** declaration can be invoked to obtain a reference to the facet.

6.5.3.3 Navigation interface on the component

As described in “Component Declaration” on page 11 all component interfaces implicitly inherit directly or indirectly from **CCMObject**, which inherits from **Components::Navigation**. The definition of the **Components::Navigation** interface is as follows:

```

module Components {

    typedef string FeatureName;

    typedef sequence<FeatureName> NameList;

    valuetype PortDescription
    {
        public FeatureName name;
        public CORBA::RepositoryId type_id;
    };

    valuetype FacetDescription : PortDescription
    {
        public Object facet_ref;
    };
    typedef sequence<FacetDescription> FacetDescriptions;

    exception InvalidName { };

    interface Navigation {

```

```

    Object provide_facet (in FeatureName name)
        raises (InvalidName);

    FacetDescriptions get_all_facets();

    FacetDescriptions get_named_facets (in NameList names)
        raises (InvalidName);

    boolean same_component (in Object object_ref);

};
};

```

This interface provides generic navigation capabilities. It is inherited by all component interfaces, and may be optionally inherited by any interface that is explicitly designed to be a facet interface for a component. The descriptions of **Navigation** operations follow.

provide_facet

The **provide_facet** operation returns a reference to the facet denoted by the **name** parameter. The value of the **name** parameter must be identical to the name specified in the provides declaration. The valid names are defined by inherited closure of the actual type of the component; that is, the names of facets of the component type and all of its inherited component types. If the value of the **name** parameter does not correspond to one of the component's facets, the **InvalidName** exception shall be raised. A component that does not provide any facets (e.g., a basic component) will have no valid **name** parameter to this operation and thus shall always raise the **InvalidName** exception.

get_all_facets

The **get_all_facets** operation returns a sequence of value objects, each of which contains the **RepositoryId** of the facet interface and **name** of the facet, along with a reference to the facet. The sequence shall contain descriptions and references for all of the facets in the component's inheritance hierarchy. The order in which these values occur in the sequence is not specified. A component that does not provide any facets (e.g., a basic component) shall return a sequence of length zero.

get_named_facets

The **get_named_facets** operation returns a sequence of described references (identical to the sequence returned by **get_all_facets**), containing descriptions and references for the facets denoted by the **names** parameter. If any name in the **names** parameter is not a valid name for a provided interface on the component, the operation raises the **InvalidName** exception. The order of values in the returned sequence is not specified. A component that does not provide any facets (e.g., a basic component) will have no valid **name** parameter to this operation and thus shall always raise the **InvalidName** exception.

The **same_component** operation on **Navigation** is described in "Provided References and Component Identity" on page 17.

6.5.3.4 Navigation interface on facet interfaces

Any interface that is designed to be used as a facet interface on a component may optionally inherit from the **Navigation** interface. When the navigation operations (i.e., **provide_facet**, **get_all_facets**, and **get_named_facets**) are invoked on the facet reference, the operations shall return the same results as if they had been invoked on the component interface that provided the target facet. The skeletons generated by the Component Implementation Framework shall provide implementations of these operations that will delegate to the component interface.

This option allows navigation from one facet to another to be performed in a single request, rather than a pair of requests (to get the component reference and navigate from there to the desired facet). To illustrate, consider the following component definition:

```
module example {
  interface foo : Components::Navigation {... };
  interface bar { ... };
  component baz {
    provides foo a;
    provides bar b;
  };
};
```

A client could navigate from a to b as follows:

```
foo myFoo;
// assume myFoo holds a reference to a foo provided by a baz
baz myBaz = bazHelper.narrow(myFoo.get_component());
bar myBar = myBaz.provide_b();
```

Or, it could navigate directly:

```
foo myFoo;
// assume myFoo holds a reference to a foo provided by a baz
bar myBar = barHelper.narrow(myFoo.provide_interface("b"));
```

6.5.4 Provided References and Component Identity

The **same_component** operation on the **Navigation** interface allows clients to determine reliably whether two references belong to the same component instance, that is, whether the references are facets of or directly denote the same component instance. The component implementation is ultimately responsible for determining what the “same component instance” means. The skeletons generated by the Component Implementation Framework shall provide an implementation of **same_component**, where “same instance” is defined in terms of opaque identity values supplied by the component implementation or the container in the container context. User-supplied implementations can provide different semantics.

If a facet interface inherits the **Navigation** interface, then the **same_component** operation on the provided interface shall give the same results as the **same_component** operation on the component interface that owns the provided interface. The skeletons generated by the Component Implementation Framework shall provide an implementation of **same_component** for facets that inherit the **Navigation** interface.

6.5.5 Supported interfaces

A component definition may optionally support one or more interfaces, or in the case of extended components, inherit from a component that supports one or more interfaces. When a component definition header includes a supports clause as follows:

```
component <component_name> supports <interface_name> { ... };
```

the equivalent interface inherits both **CCMObject** and any supported interfaces, as follows:

```
interface <component_name>  
: Components::CCMObject, <interface_name> { ... };
```

The component implementation shall supply implementations of operations defined on supported interfaces. Clients shall be able to widen a reference of the component's equivalent interface type to the type of any of the supported interfaces. Clients shall also be able to narrow a reference of type **CCMObject** to the type of any of the component's supported interfaces.

For example, given the following IDL:

```
module M {  
    interface I {  
        void op();  
    };  
    component A supports I {  
        provides I foo;  
    };  
    home AManager manages A { };  
};
```

The AManager interface shall be derived from KeylessCCMHome, supporting the create_component operation, which returns a reference of type CCMObject. This reference shall be able to be narrowed directly from CCMObject to I:

```
// java  
...  
M.AManager aHome = ...; // get A's home  
org.omg.Components.CCMObject myComp = aHome.create_component();  
M.I myI = M.IHelper.narrow(myComp);  
// must succeed
```

For example, given the following IDL:

```
module M {  
    interface I {  
        void op();  
    };  
    component A supports I {  
        provides I foo;  
    };  
    component B : A { ... };  
    home BHome manages B {};  
};
```

The equivalent IDL is:

```
module M {
    interface I {
        void op();
    };
    interface A :
        org.omg.Components.CCMAbstractObject, I { ... };
    interface B : A { ... };
};
```

which allows the following usage:

```
M.BHome bHome = ... // get B's home
M.B myB = bHome.create();
myB.op(); // I's operations are supported
           // directly on B's interface
```

The supports mechanism provides programming convenience for light-weight components that only need to implement a single operational interface. A client can invoke operations from the supported interface directly on the component reference, without narrowing or navigation:

```
M.A myA = aHome.create();
myA.op();
```

as opposed to

```
M.A myA = aHome.create();
M.I myI = myA.provide_foo();
myI.op();
```

or, assuming that the client has A's home, but doesn't statically know about A's interface or home interface:

```
org.omg.Components.KeylessCCMHome genericHome =
... // get A's home;
org.omg.Components.CCMAbstractObject myComp =
genericHome.create_component();
```

```
M.I myI = M.IHelper.narrow(myComp);
myI.op();
```

as opposed to

```
org.omg.CORBA.Object obj = myComp.provide_interface("foo");
M.I myI = M.IHelper.narrow(obj);
myI.op();
```

This mechanism allows component-unaware clients to receive a reference to a component (passed as type `CORBA::Object`) and use the supported interface.

6.6 Receptacles

A component definition can describe the ability to accept object references upon which the component may invoke operations. When a component accepts an object reference in this manner, the relationship between the component and the referent object is called a *connection*; they are said to be *connected*. The conceptual point of connection is called a *receptacle*. A receptacle is an abstraction that is concretely manifested on a component as a set of operations for establishing and managing connections. A component may exhibit zero or more receptacles.

Receptacles are intended as a mechanical device for expressing a wide variety of relationships that may exist at higher levels of abstraction. As such, receptacles have no inherent higher-order semantics, such as implying ownership, or that certain operations will be transient across connections.

6.6.1 Equivalent IDL

A **uses** declaration of the following form:

```
uses <interface_type> <receptacle_name>;
```

results in the following equivalent operations defined in the component interface:

```
void connect_<receptacle_name> ( in <interface_type> conxn ) raises (  
    Components::AlreadyConnected,  
    Components::InvalidConnection );
```

```
<interface_type> disconnect_<receptacle_name> ( )  
    raises ( Components::NoConnection );
```

```
<interface_type> get_connection_<receptacle_name> ( );
```

A **uses multiple** declaration of the following form:

```
uses multiple <interface_type> <receptacle_name>;
```

results in the following equivalent operations defined in the component interface:

```
struct <receptacle_name>Connection {  
    <interface_type> objref;  
    Components::Cookie ck;  
    };  
sequence <<receptacle_name>Connection> <receptacle_name>Connections;
```

```
Components::Cookie  
    connect_<receptacle_name> ( in <interface_type> connection ) raises (  
        Components::ExceededConnectionLimit,  
        Components::InvalidConnection  
    );
```

```
<interface_type> disconnect_<receptacle_name> (  
    in Components::Cookie ck)  
    raises ( Components::InvalidConnection );
```

```
<receptacle_name>Connections get_connections_<receptacle_name> ( );
```

6.6.2 Behavior

6.6.2.1 Connect operations

Operations of the form **connect_<receptacle_name>** are implemented in part by the component implementor, and in part by generated code in the component servant framework. The responsibilities of the component implementation and servant framework for implementing connect operations are described in detail in Clause 8 *OMG CIDL Syntax and Semantics*. The receptacle holds a copy of the object reference passed as a parameter. The component may invoke operations on this reference according to its design. How and when the component invokes operations on the reference is entirely the prerogative of the component implementation. The receptacle shall hold a copy of the reference until it is explicitly disconnected.

Simplex receptacles

If a receptacle's **uses** declaration does not include the optional **multiple** keyword, then only a single connection to the receptacle may exist at a given time. If a client invokes a connect operation when a connection already exists, the connection operation shall raise the **AlreadyConnected** exception.

The component implementation may refuse to accept the connection for arbitrary reasons. If it does so, the connection operation shall raise the **InvalidConnection** exception.

Multiplex receptacles

If a receptacle's **uses** declaration includes the optional **multiple** keyword, then multiple connections to the receptacle may exist simultaneously. The component implementation may choose to establish a limit on the number of simultaneous connections allowed. If an invocation of a connect operation attempts to exceed this limit, the operation shall raise the **ExceededConnectionLimit** exception.

The component implementation may refuse to accept the connection for arbitrary reasons. If it does so, the connection operation shall raise the **InvalidConnection** exception.

Connect operations for multiplex receptacles return values of type **Components::Cookie**. Cookie values are used to identify the connection for subsequent disconnect operations. Cookie values are generated by the receptacle implementation (the responsibility of the supplier of the component-enabled ORB, not the component implementor). Likewise, cookie equivalence is determined by the implementation of the receptacle implementation.

The client invoking connection operations is responsible for retaining cookie values and properly associating them with connected object references, if the client needs to subsequently disconnect specific references. Cookie values must be unique within the scope of the receptacle that created them. If a cookie value is passed to a disconnect operation on a different receptacle than that which created it, results are undefined.

Cookie values are described in detail in "Cookie type" on page 22."

Cookie values are required because object references cannot be reliably tested for equivalence.

6.6.2.2 Disconnect operations

Operations of the form **disconnect_receptacle_name** terminate the relationship between the component and the connected object reference.

Simplex receptacles

If a connection exists, the disconnect operation will return the connected object reference. If no connection exists, the operation shall raise a `NoConnection` exception.

Multiplex receptacles

The `disconnect_receptacle_name` operation of a multiplex receptacle takes a parameter of type `Components::Cookie`. The `ck` parameter must be a value previously returned by the `connect_receptacle_name` operation on the same receptacle. It is the responsibility of the client to associate cookies with object references they connect and disconnect. If the cookie value is not recognized by the receptacle implementation as being associated with an existing connection, the `disconnect_receptacle_name` operation shall raise an `InvalidConnection` exception.

6.6.2.3 get_connection and get_connections operations

Simplex receptacles

Simplex receptacles have operations named `get_connection_receptacle_name`. If the receptacle is currently connected, this operation returns the connected object reference. If there is no current connection, the operation returns a `nil` object reference.

Multiplex receptacles

Multiplex receptacles have operations named `get_connections_receptacle_name`. This operation returns a sequence of structures, where each structure contains a connected object reference and its associated cookie value. The sequence contains a description of all of the connections that exist at the time of the invocation. If there are no connections, the sequence length will be zero.

6.6.2.4 Cookie type

The `Cookie` valuetype is defined by the following IDL:

```
module Components {
    valuetype Cookie {
        private CORBA::OctetSeq cookieValue;
    };
};
```

Cookie values are created by multiplex receptacles, and are used to correlate a connect operation with a disconnect operation on multiplex receptacles.

Implementations of component-enabled ORBs may employ value type derived from `Cookie`, but any derived cookie types shall be truncatable to `Cookie`, and the information preserved in the `cookieValue` octet sequence shall be sufficient for the receptacle implementation to identify the cookie and its associated connected reference.

6.6.3 Receptacles Interface

The `Receptacles` interface provides generic operations for connecting to a component's receptacles. The `CCMObject` interface is derived from `Receptacles`. For components, such as basic components, that do not use interfaces, only the generic receptacles operations are available on the equivalent interface. The default behavior in such cases is defined below.

The **Receptacles** interfaces is defined by the following IDL.

```
module Components {  
  
    valuetype ConnectionDescription {  
        public Cookie ck;  
        public Object objref;  
    };  
    typedef sequence<ConnectionDescription> ConnectionDescriptions;  
  
    valuetype ReceptacleDescription : PortDescription  
    {  
        public boolean is_multiple;  
        public ConnectionDescriptions connections;  
    };  
    typedef sequence<ReceptacleDescription> ReceptacleDescriptions;  
  
    exception ExceededConnectionLimit { };  
  
    exception CookieRequired { };  
  
    interface Receptacles {  
  
        Cookie connect ( in FeatureName name, in Object connection )  
        raises (  
            InvalidName,  
            InvalidConnection,  
            AlreadyConnected,  
            ExceededConnectionLimit);  
  
        Object disconnect (  
            in FeatureName name,  
            in Cookie ck) raises (  
                InvalidName,  
                InvalidConnection,  
                CookieRequired,  
                NoConnection);  
  
        ConnectionDescriptions get_connections (  
            in FeatureName name) raises (InvalidName);  
  
        ReceptacleDescriptions get_all_receptacles ();  
  
        ReceptacleDescriptions get_named_receptacles (  
            in NameList names) raises(InvalidName);  
    };  
};
```

connect

The **connect** operation connects the object reference specified by the **connection** parameter to the receptacle specified by the **name** parameter on the target component. If the specified receptacle is a multiplex receptacle, the operation returns a cookie value that can be used subsequently to disconnect the object reference. If the receptacle is a simplex receptacle, the return value is a nil. The following exceptions may be raised:

- If the **name** parameter does not specify a valid receptacle name, then the **InvalidName** exception is raised.
- If the receptacle is a simplex receptacle and it is already connected, then the **AlreadyConnected** exception is raised.
- If the object reference in the **connection** parameter does not support the interface declared in the receptacle's **uses** statement, the **InvalidConnection** exception is raised.
- If the receptacle is a multiplex receptacle and the implementation-defined limit to the number of connections is exceeded, the **ExceededConnectionLimit** exception is raised.
- A component that does not have any receptacles (e.g., a basic component) will have no valid **name** parameter to this operation and thus shall always raise the **InvalidName** exception.

disconnect

If the receptacle identified by the **name** parameter is a simplex receptacle, the operation will disassociate any object reference currently connected to the receptacle. The cookie value in the **ck** parameter is ignored. If the receptacle identified by the **name** parameter is a multiplex receptacle, the **disconnect** operation disassociates the object reference associated with the cookie value (i.e., the object reference that was connected by the operation that created the cookie value) from the receptacle. In both cases, the **disconnect** operation returns the previously connected object reference. The following exceptions may be raised:

- If the **name** parameter does not specify a valid receptacle name, then the **InvalidName** exception is raised.
- If the receptacle is a simplex receptacle and there is no current connection, then the **NoConnection** exception is raised.
- If the receptacle is a multiplex receptacle and the cookie value in the **ck** parameter does not denote an existing connection on the receptacle, the **InvalidConnection** exception is raised.
- If the receptacle is a multiplex receptacle and a null value is specified in the **ck** parameter, the **CookieRequired** exception is raised.
- A component that does not have any receptacles (e.g., a basic component) will have no valid **name** parameter to this operation and thus shall always raise the **InvalidName** exception.

get_connections

The **get_connections** operation returns a sequence of **ConnectionDescription** structs. Each struct contains an object reference connected to the receptacle named in the **name** parameter, and a cookie value that denotes the connection. If the **name** parameter does not specify a valid receptacle name, then the **InvalidName** exception is raised. A component that does not have any receptacles (e.g., a basic component) will have no valid **name** parameter to this operation and thus shall always raise the **InvalidName** exception.

get_all_receptacles

The **get_all_receptacles** operation returns information about all receptacle ports in the component's inheritance hierarchy as a sequence of **ReceptacleDescription** values. The order in which these values occur in the sequence is not specified. For components that do not have any receptacles (e.g., a basic component), this operation returns a sequence of length zero.

get_named_receptacles

The **get_named_receptacles** operation returns information about all receptacle ports denoted by the **names** parameter as a sequence of **ReceptacleDescription** values. The order in which these values occur in the sequence is not specified. If any name in the **names** parameter is not a valid name for a receptacle in the component's inheritance hierarchy, the operation raises the **InvalidName** exception. A component that does not provide any receptacles (e.g., a basic component) will have no valid name parameter to this operation and thus shall always raise the **InvalidName** exception.

6.7 Events

The CORBA component model supports a publish/subscribe event model. The event model for CORBA components is designed to be compatible with CORBA notification, as defined in <http://www.omg.org/spec/NOT/>. The interfaces exposed by the component event model provide a simple programming interface whose semantics can be mapped onto a subset of CORBA notification semantics.

6.7.1 Event types

IDL contains event type declarations, which are a restricted form of value type declarations. They are for the use in the CORBA Component event model.

Since the underlying implementation of the component event mechanism provided by the container is CORBA notification, event values shall be inserted into instances of the **any** type. The resulting **any** values shall be inserted into a CORBA notification structured event. The mapping between a component event and a notification event is implemented by the container.

6.7.1.1 Equivalent IDL

For the declaration of event types of the following form:

```
module <module_name> {  
    valuetype A { <A_state_members> };  
    eventtype B : A { <B_state_members> };  
    eventtype C : B { <C_state_members> };  
};
```

The following equivalent IDL is implied:

```
module <module_name> {  
  
    valuetype A { <A_state_members> };  
  
    valuetype B : A, ::Components::EventBase {  
        <B_state_members>  
    };  
};
```

```

interface BConsumer : ::Components::EventConsumerBase {
    void push_B (in B the_b);
};

valuetype C : B {
    <C_state_members>
};

interface CConsumer : BConsumer {
    void push_C (in C the_c);
};
};

```

As shown above the first event type in the inheritance chain introduces the inheritance from **Components::EventBase** into the inheritance chain for the equivalent value types. The same rule applies for the equivalent consumer interfaces and **Components::EventConsumerBase**. Consumer interfaces are in the same inheritance relation as the event types, where they origin.

6.7.1.2 EventBase

The module **Components** contains the following abstract value type definition:

```

module Components {
    abstract valuetype EventBase { };
};

```

It serves as base type for value types derived via the Equivalent IDL mapping for event types.

To ensure proper transmission of value type events, this specification makes the following clarifications to the semantics of value types when inserted into **any**s:

When an **any** containing a value type is received as a parameter in an ORB-mediated operation, the value contained in the **any** shall be preserved, regardless of whether the receiving execution context is capable of constructing the value (in its original form or a truncated form), or not. If the receiving context attempts to extract the value, the extraction may fail, or the extracted value may be truncated. The value contained in the **any** shall remain unchanged, and shall retain its integrity if the **any** is passed as a parameter to another execution context.

6.7.2 EventConsumer Interface

The component event model is a push model. The basic mechanics of this push model are defined by consumer interfaces. Event sources hold references to consumer interfaces and invoke various forms of push operations to send events.

Component event sources hold references to consumer interfaces and push to them. Component event sinks provide consumer references, into which other entities (e.g., channels, clients, other component event sources) push events.

Event consumer interfaces are derived from the **Components::EventConsumerBase** interface, which is defined as follows:

```

module Components {
    exception BadEventType {
        CORBA::RepositoryId expected_event_type;
    };
};

```

```

};
interface EventConsumerBase {
    void push_event(in EventBase evt) raises (BadEventType);
};
};

```

Type-specific event consumer interfaces are derived from the **EventConsumerBase** interface. Event source and sink declarations in component definitions cause type-specific consumer interfaces to be generated for the event types used in the declarations.

The **push_event** operation pushes the event denoted by the **evt** parameter to the consumer. The consumer may choose to constrain the type of event it accepts. If the actual type of the **evt** parameter is not acceptable to the consumer, the **BadEventType** exception shall be raised. The **expected_event_type** member of the exception contains the **RepositoryId** of the type expected by the consumer.

Note that this exception can only be raised by the consumer upon whose reference the **push_event** operation was invoked. The consumer may be a proxy for an event or notification channel with an arbitrary number of subscribers. If any of those subscribers raise any exceptions, they will not be propagated back to the original event source (i.e., the component).

6.7.3 Event Service Provided by Container

Container implementations provide event services to components and their clients. Component implementations obtain event services from the container during initialization, and mediate client access to those event services. The container implementation is free to provide any mechanism that supports the required semantics. The container is responsible for configuring the mechanism and determining the specific quality of service and routing policies to be employed when delivering events.

6.7.4 Event Sources—Publishers and Emitters

An event source embodies the potential for the component to generate events of a specified type, and provides mechanisms for associating consumers with sources.

There are two categories of event sources, *emitters* and *publishers*. Both are implemented using event channels supplied by the container. An emitter can be connected to at most one proxy provider by the container. A publisher can be connected through the channel to an arbitrary number of consumers, who are said to *subscribe* to the publisher event source. A component may exhibit zero or more emitters and publishers.

A *publisher* event source has the following characteristics:

- The equivalent operations for publishers allow multiple subscribers (i.e., consumers) to connect to the same source simultaneously.
- Subscriptions to a publisher are delegated to an event channel supplied by the container at run time. The component is guaranteed to be the only source publishing to that event channel.

An *emitter* event source has the following characteristics:

- The equivalent operations for emitters allow only one consumer to be connected to the emitter at a time.

- The events pushed from an emitter are delegated to an event channel supplied by the container at run time. Other event sources, however, may use the same channel. Events pushed from an emitter are then pushed by the container into the consumer interface supplied as a parameter to the `connect_<source>` operation.

In general, emitters are not intended to be exposed to clients. Rather, they are intended to be used for configuration purposes. It is expected that emitters will be connected at the time of component initialization and configuration to consumer interfaces that are proxies for event channels that may be shared between arbitrary clients, components, and other system elements.

In contrast, publishers are intended to provide clients with direct access to a particular event stream being generated by the component (embodied by the publisher event source). It is our intent that clients subscribe directly to the publisher source.

6.7.5 Publisher

6.7.5.1 Equivalent IDL

For an event source declaration of the following form:

```
module <module_name> {
    component <component_name> {
        publishes <event_type> <source_name>;
    };
};
```

The following equivalent IDL is implied:

```
module <module_name> {
    interface <component_name> : Components::CCMObject {
        Components::Cookie subscribe_<source_name> (
            in <event_type>Consumer consumer)
            raises (Components::ExceededConnectionLimit);

        <event_type>Consumer unsubscribe_<source_name> (
            in Components::Cookie ck)
            raises (Components::InvalidConnection);
    };
};
```

6.7.5.2 Event publisher operations

subscribe_<source_name>

The **subscribe_<source_name>** operation connects the consumer parameter to an event channel provided to the component implementation by the container. The component shall be the only publisher to that channel. If the implementation of the component or the channel place an arbitrary limit on the number of subscriptions that can be supported simultaneously, and the invocation of the subscribe operation would cause that limit to be exceeded, the operation raises the `ExceededConnectionLimit` exception. The **Cookie** value returned by the operation identifies the subscription formed by the association of the subscriber with the publisher event source. This value can be used subsequently in an invocation of **unsubscribe_<source_name>** to disassociate the subscriber from the publisher.

unsubscribe_<source_name>

The **unsubscribe_<source_name>** operation destroys the subscription identified by the **ck** parameter value, returning the reference to the subscriber. If the **ck** parameter value does not identify an existing subscription to the publisher event source, the operation shall raise an `InvalidConnection` exception.

6.7.6 Emitters

6.7.6.1 Equivalent IDL

For an event source declaration of the following form:

```
module <module_name> {  
    component <component_name> {emits <event_type> <source_name>;}  
};
```

The following equivalent IDL is implied:

```
module <module_name> {  
    interface <component_name> : Components::CCMObject {  
        void connect_<source_name> (  
            in <event_type>Consumer consumer)  
            raises (Components::AlreadyConnected);  
  
        <event_type>Consumer disconnect_<source_name>()  
            raises (Components::NoConnection);  
    };  
};
```

6.7.6.2 Event emitter operations

connect_<source_name>

The **connect_<source_name>** operation connects the event consumer denoted by the consumer parameter to the event emitter. If the emitter is already connected to a consumer, the operation shall raise the `AlreadyConnected` exception.

disconnect_<source_name>

The **disconnect_<source_name>** operation destroys any existing connection by disassociating the consumer from the emitter. The reference to the previously connected consumer is returned. If there was no existing connection, the operation raises the `NoConnection` exception.

The following observations and constraints apply to the equivalent IDL for event source declarations:

- The need for a typed event consumer interface requires the definition of a module scope to guarantee that the interface name for the event subscriber is unique. The module (whose name is formed by appending the string “EventConsumers” to the component type name) is defined in the same scope as the component’s equivalent interface. The module is opened before the equivalent interface definition to provide forward declarations for consumer interfaces. It is re-opened after the equivalent interface definition to define the consumer interfaces.

- The name of a consumer interface is formed by appending the string “Consumer” to the name of the event type. One consumer interface type is implied for each unique event type used in event source and event sink declarations in the component definition.

6.7.7 Event Sinks

An event sink embodies the potential for the component to receive events of a specified type. An event sink is, in essence, a special-purpose facet whose type is an event consumer. External entities, such as clients or configuration services, can obtain the reference for the consumer interface associated with the sink.

Unlike event sources, event sinks do not distinguish between *connection* and *subscription*. The consumer interface may be associated with an arbitrary number of event sources, unbeknownst to the component that supplies the event sink. The component event model provides no inherent mechanism for the component to control which events sources may be pushing to its sinks. By exporting an event sink, the component is, in effect, declaring its willingness to accept events pushed from arbitrary sources. A component may exhibit zero or more consumers.

If a component implementation needs control over which sources can push to a particular sink it owns, the sink should not be exposed as a port on the component. Rather, the component implementation can create a consumer internally and explicitly connect or subscribe it to sources.

6.7.7.1 Equivalent IDL

For an event sink declaration of the following form:

```
module <module_name> {
    component <component_name> {
        consumes <event_type> <sink_name>;
    };
};
```

The following equivalent IDL is implied:

```
module <module_name> {
    interface <component_name> : Components::CCMObject {
        <event_type>Consumer get_consumer_<sink_name>();
    };
};
```

6.7.7.2 Event sink operations

The `get_consumer_<sink_name>` operation returns a reference that supports the consumer interface specific to the declared event type.

6.7.8 Events interface

The **Events** interface provides generic access to event sources and sinks on a component. **CCMObject** is derived from **Events**. For components, such as basic components, that do not declare participation in events, only the generic **Events** operations are available on the equivalent interface. The default behavior in such cases is described below.

The **Events** interface is described as follows:

```

module Components {

    exception InvalidName { };
    exception InvalidConnection { };
    exception AlreadyConnected { };
    exception NoConnection { };

    valuetype ConsumerDescription : PortDescription
    {
        public EventConsumerBase consumer;
    };
    typedef sequence<ConsumerDescription> ConsumerDescriptions;

    valuetype EmitterDescription : PortDescription
    {
        public EventConsumerBase consumer;
    };
    typedef sequence<EmitterDescription> EmitterDescriptions;

    valuetype SubscriberDescription
    {
        public Cookie ck;
        public EventConsumerBase consumer;
    };
    typedef sequence<SubscriberDescription> SubscriberDescriptions;

    valuetype PublisherDescription : PortDescription
    {
        public SubscriberDescriptions consumers;
    };
    typedef sequence<PublisherDescription> PublisherDescriptions;

    interface Events {
        EventConsumerBase get_consumer (in FeatureName sink_name)
            raises (InvalidName);
        Cookie subscribe (in FeatureName publisher_name,
            in EventConsumerBase subscriber)
            raises (InvalidName, InvalidConnection,
                ExceededConnectionLimit);
        EventConsumerBase unsubscribe (in FeatureName publisher_name,
            in Cookie ck)
            raises (InvalidName, InvalidConnection);
        void connect_consumer (in FeatureName emitter_name,
            in EventConsumerBase consumer)
            raises (InvalidName, AlreadyConnected,
                InvalidConnection);
        EventConsumerBase disconnect_consumer (
            in FeatureName source_name)
            raises (InvalidName, NoConnection);
        ConsumerDescriptions get_all_consumers ();
        ConsumerDescriptions get_named_consumers (

```

```

        in NameList names)
        raises (InvalidName);
    EmitterDescriptions get_all_emitters ();
    EmitterDescriptions get_named_emitters (in NameList names)
        raises (InvalidName);
    PublisherDescriptions get_all_publishers ();
    PublisherDescriptions get_named_publishers (in NameList names)
        raises (InvalidName);
};
};

```

get_consumer

The **get_consumer** operation returns the **EventConsumerBase** interface for the sink specified by the **sink_name** parameter. If the **sink_name** parameter does not specify a valid event sink on the component, the operation raises the **InvalidName** exception. A component that does not have any sinks (e.g., a basic component) will have no valid **sink_name** parameter to this operation and thus shall always raise the **InvalidName** exception.

subscribe

The **subscribe** operation associates the subscriber denoted by the **subscriber** parameter with the event source specified by the **publisher_name** parameter. If the **publisher_name** parameter does not specify a valid event publisher on the component, the operation raises the **InvalidName** exception. The cookie return value can be used to unsubscribe from the source. A component that does not have any event sources (e.g., a basic component) will have no valid **publisher_name** parameter to this operation and thus shall always raise the **InvalidName** exception. If the object reference in the **subscriber** parameter does not support the consumer interface of the eventtype declared in the **publishes** statement, the **InvalidConnection** exception is raised. If the implementation-defined limit to the number of subscribers is exceeded, the **ExceededConnectionLimit** exception is raised.

unsubscribe

The **unsubscribe** operation disassociates the subscriber associated with **ck** parameter with the event source specified by the **publisher_name** parameter, and returns the reference to the subscriber. If the **publisher_name** parameter does not specify a valid event source on the component, the operation raises the **InvalidName** exception. If the **ck** parameter does not identify a current subscription on the source, the operation raises the **InvalidConnection** exception. A component that does not have any event sources (e.g., a basic component) will have no valid **publisher_name** parameter to this operation and thus shall always raise the **InvalidName** exception.

connect_consumer

The **connect_consumer** operation associates the consumer denoted by the **consumer** parameter with the event source specified by the **emitter_name** parameter. If the **emitter_name** parameter does not specify a valid event emitter on the component, the operation raises the **InvalidName** exception. If a consumer is already connected to the emitter, the operation raises the **AlreadyConnected** exception. If the object reference in the **consumer** parameter does not support the consumer interface of the eventtype declared in the **emits** statement, the **InvalidConnection** exception is raised. The cookie return value can be used to disconnect from the source. A component that does not have any event sources (e.g., a basic component) will have no valid **emitter_name** parameter to this operation and thus shall always raise the **InvalidName** exception.

disconnect_consumer

The **disconnect_consumer** operation disassociates the currently connected consumer from the event source specified by the **emitter_name** parameter, returning a reference to the disconnected consumer. If the **emitter_name** parameter does not specify a valid event source on the component, the operation raises the **InvalidName** exception. If there is no consumer connected to the emitter, the operation raises the **NoConnection** exception. A component that does not have any event sources (e.g., a basic component) will have no valid **emitter_name** parameter to this operation and thus shall always raise the **InvalidName** exception.

get_all_consumers

The **get_all_consumers** operation returns information about all consumer ports in the component's inheritance hierarchy as a sequence of **ConsumerDescription** values. The order in which these values occur in the sequence is not specified. For components that do not consume any events (e.g., a basic component), this operation returns a sequence of length zero.

get_named_consumers

The **get_named_consumers** operation returns information about all consumer ports denoted by the **names** parameter as a sequence of **ConsumerDescription** values. The order in which these values occur in the sequence is not specified. If any name in the **names** parameter is not a valid name for an event sink in the component's inheritance hierarchy, the operation raises the **InvalidName** exception. A component that does not provide any consumers (e.g., a basic component) will have no valid name parameter to this operation and thus shall always raise the **InvalidName** exception.

get_all_emitters

The **get_all_emitters** operation returns information about all emitter ports in the component's inheritance hierarchy as a sequence of **EmitterDescription** values. The order in which these values occur in the sequence is not specified. For components that do not emit any events (e.g., a basic component), this operation returns a sequence of length zero.

get_named_emitters

The **get_named_emitters** operation returns information about all emitter ports denoted by the **names** parameter as a sequence of **EmitterDescription** values. The order in which these values occur in the sequence is not specified. If any name in the **names** parameter is not a valid name for an emitter port in the component's inheritance hierarchy, the operation raises the **InvalidName** exception. A component that does not provide any emitters (e.g., a basic component) will have no valid name parameter to this operation and thus shall always raise the **InvalidName** exception.

get_all_publishers

The **get_all_publishers** operation returns information about all publisher ports in the component's inheritance hierarchy as a sequence of **PublisherDescription** values. The order in which these values occur in the sequence is not specified. For components that do not publish any events (e.g., a basic component), this operation returns a sequence of length zero.

get_named_publishers

The **get_named_publishers** operation returns information about all publisher ports denoted by the **names** parameter as a sequence of **PublisherDescription** values. The order in which these values occur in the sequence is not specified. If any name in the **names** parameter is not a valid name for a publisher port in the component's inheritance hierarchy, the operation raises the **InvalidName** exception. A component that does not provide any publishers (e.g., a basic component) will have no valid name parameter to this operation and thus shall always raise the **InvalidName** exception.

6.8 Homes

An IDL specification may include home definitions. A home definition describes an interface for managing instances of a specified component type. The salient characteristics of a home definition are as follows:

- A home definition implicitly defines an equivalent interface, which can be described in terms of IDL.
- The presence of a primary key specification in a home definition causes home's equivalent interface to contain a set of implicitly defined operations whose signatures are determined by the types of the primary key and the managed component. These operations are specified in "Home definitions with primary keys" on page 35.

6.8.1 Equivalent Interfaces

Every home definition implicitly defines a set of operations whose names are the same for all homes, but whose signatures are specific to the component type managed by the home and, if present, the primary key type specified by the home.

Because the same operation names are used for these operations on different homes, the implicit operations cannot be inherited. The specification for home equivalent interfaces accommodates this constraint. A home definition results in the definition of three interfaces, called the *explicit* interface, the *implicit* interface, and the *equivalent* interface. The name of the explicit interface has the form `<home_name>Explicit`, where `<home_name>` is the declared name of the home definition. Similarly, the name of the implicit interface has the form `<home_name>Implicit`, and the name of the equivalent interface is simply the name of the home definition, with the form `<home_name>`. All of the operations defined explicitly on the home (including explicitly-defined factory and finder operations) are represented on the explicit interface. The operations that are implicitly defined by the home definition are exported by the implicit interface. The equivalent interface inherits both the explicit and implicit interfaces, forming the interface presented to programmer using the home.

The same names are used for implicit operations in order to provide clients with a simple, uniform view of the basic life cycle operations—creation, finding, and destruction. The signatures differ to make the operations specific to the storage type (and, if present, primary key) associated with the home. These two goals—uniformity and type safety—are admittedly conflicting, and the resulting complexity of equivalent home interfaces reflects this conflict. Note that this complexity manifests itself in generated interfaces and their inheritance relationships; the model seen by the client programmer is relatively simple.

6.8.1.1 Home definitions with no primary key

Given a home definition of the following form:

```
home <home_name> manages <component_type> {  
    <explicit_operations>  
};
```

The resulting explicit, implicit, and equivalent local interfaces have the following forms:

```
interface <home_name>Explicit : Components::CCMHome {  
    <equivalent_explicit_operations>  
};  
  
interface <home_name>Implicit : Components::KeylessCCMHome {  
    <component_type> create() raises(CreateFailure);  
};
```

```
interface <home_name> : <home_name>Explicit, <home_name>Implicit { };
```

where *<equivalent_explicit_operations>* are the operations defined in the home declaration (*<explicit_operations>*), with factory and finder operations transformed to their equivalent operations, as described in “Explicit Operations in Home Definitions” on page 37.”

create

This operation creates a new component instance of the type managed by the home. The `CreateFailure` exception is raised if any application errors are encountered in home creation.

6.8.1.2 Home definitions with primary keys

Given a home of the following form:

```
home <home_name> manages <component_type> primarykey <key_type> {  
    <explicit_operations>  
};
```

The resulting explicit, implicit, and equivalent interfaces have the following forms:

```
interface <home_name>Explicit : Components::CCMHome {  
    <equivalent_explicit_operations>  
};
```

```
interface <home_name>Implicit {  
    <component_type> create (in <key_type> key)  
    raises (Components::CreateFailure, Components::DuplicateKeyValue,  
           Components::InvalidKey);  
  
    <component_type> find_by_primary_key (in <key_type> key)  
    raises (Components::FinderFailure, Components::UnknownKeyValue,  
           Components::InvalidKey);  
  
    void remove (in <key_type> key)  
    raises (Components::RemoveFailure, Components::UnknownKeyValue,  
           Components::InvalidKey);  
  
    <key_type> get_primary_key (in <component_type> comp);  
};
```

```
interface <home_name> : <home_name>Explicit, <home_name>Implicit { };
```

where *<equivalent_explicit_operations>* are the operations defined in the home declaration (*<explicit_operations>*), with factory and finder operations transformed to their equivalent operations, as described in “Explicit Operations in Home Definitions” on page 37.

create

This operation creates a new component associated with the specified primary key value, returning a reference to the component. If the specified key value is already associated with an existing component managed by the storage home, the operation raises a `DuplicateKeyValue` exception. If the key value was not a well-formed, legal value, the operation shall raise the `InvalidKey` exception. All other error conditions may raise the `CreateFailure` exception.

find_by_primary_key

This operation returns a reference to the component identified by the primary key value. If the key value does not identify an existing component managed by the home, an `UnknownKeyValue` exception is raised. If the key value was not a well-formed, legal value, the operation shall raise the `InvalidKey` exception. All other error conditions may raise the `FinderFailure` exception.

remove

This operation removes the component identified by the specified key value. Subsequent requests to any of the component's facets shall raise an `OBJECT_NOT_EXIST` system exception. If the specified key value does not identify an existing component managed by the home, the operation shall raise an `UnknownKeyValue` exception. If the key value was not a well-formed, legal value, the operation shall raise the `InvalidKey` exception. All other error conditions may raise the `RemoveFailure` exception.

6.8.1.3 Supported interfaces

A home definition may optionally support one or more interfaces. When a home definition header includes a **supports** clause as follows:

```
home <home_name> supports <interface_name>
                        manages <component_type> {
    <explicit_operations>
};
```

the resulting explicit interface inherits both **CCMHome** and any supported interfaces, as follows:

```
interface <home_name>Explicit : Components::CCMHome,
    <interface_name> {
    <equivalent_explicit_operations>
};
```

The home implementation shall supply implementations of operations defined on supported interfaces. Clients shall be able to widen a reference of the home's resulting explicit or equivalent interface type to the type of any of the supported interfaces. Clients shall also be able to narrow a reference of type **CCMHome** to the type of any of the home's supported interfaces.

6.8.2 Primary Key Declarations

Primary key values shall uniquely identify component instances within the scope of the home that manages them. Two component instances cannot exist on the same home with the same primary key value.

Different home types that manage the same component type may specify different primary key types. Consequently, a primary key type is not inherently related to the component type, and vice versa. A home definition determines the association between a component type and a primary key type. The home implementation is responsible for maintaining the association between specific primary key values and specific component identities.

Note that this discussion pertains to component definitions as abstractions. A particular implementation of a component type may be cognizant of, and dependent upon, the primary keys associated with its instances. Such dependencies, however, are not exposed on the surface of the component type. A particular implementation of a component type may be designed to be manageable by different home interfaces with different primary keys, or it may be inextricably bound to a particular home definition. Generally, an implementation of a component type and the implementation of its associated home are inter-dependent, although this is not absolutely necessary.

6.8.2.1 Primary key type constraints

Primary key and types are subject to the following constraints:

- A primary key type must be a value type derived from **Components::PrimaryKeyBase**.
- A primary key type must be a concrete type with at least one public state member.
- A primary key type may not contain private state members.
- A primary key type may not contain any members whose type is a CORBA interface reference type, including references for interfaces, abstract interfaces, and local interfaces.
- These constraints apply recursively to the types of all of the members; that is, members that are structs, unions, value types, sequences or arrays may not contain interface reference types. If the type of a member is a value type or contains a value type, it must meet all of the above constraints.

6.8.2.2 PrimaryKeyBase

The base type for all primary keys is the abstract value type **Components::PrimaryKeyBase**. The definition of **PrimaryKeyBase** is as follows:

```
module Components {  
    abstract valuetype PrimaryKeyBase { };  
};
```

6.8.3 Explicit Operations in Home Definitions

A home body may include zero or more operation declarations, where the operation may be a *factory* operation, a *finder* operation, or a normal operation or attribute.

6.8.3.1 Factory operations

A factory operation is denoted by the **factory** keyword. A factory operation has a corresponding equivalent operation on the home's explicit interface. Given a factory declaration of the following form:

```
home <home_name> manages <component_type> {  
    factory <factory_operation_name> (<parameters>)  
    raises (<exceptions>);  
};
```

The equivalent operation on the explicit interface is as follows:

```
<component_type> <factory_operation_name> ( <parameters> )  
  raises (Components::CreateFailure, <exceptions> );
```

A factory operation is required to support creation semantics; that is, the reference returned by the operation shall identify a component that did not exist prior to the operation's invocation. Factory operations are required to raise `CreateFailure` and may raise other exceptions.

6.8.3.2 Finder operations

A finder operation is denoted by the **finder** keyword. A finder operation has a corresponding equivalent operation on the home's explicit interface. Given a finder declaration of the following form:

```
home <home_name> manages <component_type> {  
  finder <finder_operation_name> (<parameters>) raises (<exceptions>);  
};
```

The equivalent operation on the explicit interface is as follows:

```
<component_type> <finder_operation_name> ( <parameters> )  
  raises (Components::FinderFailure, <exceptions> );
```

A finder operation shall support the following semantics. The reference returned by the operation shall identify a previously-existing component managed by the home. The operation implementation determines which component's reference to return based on the values of the operation's parameters. Finder operations are required to raise `FinderFailure` and may raise other exceptions.

6.8.3.3 Miscellaneous exports

All of the exports, other than factory and finder operations, that appear in a home definition are duplicated exactly on the home's explicit interface.

6.8.4 Home inheritance

Given a derived home definition of the following form:

```
home <home_name> : <base_home_name> manages <component_type> {  
  <explicit_operations>  
};
```

The resulting explicit interface has the following form:

```
interface <home_name>Explicit : <base_home_name>Explicit {  
  <equivalent_explicit_operations>  
};
```

Given a derived home definition supporting one or more interfaces, as follows:

```
home <home_name> : <base_home_name>  
  supports <interface_name>
```

```

    manages <component_type> {
      <explicit_operations>
    };

```

The resulting explicit interface has the following form:

```

interface <home_name>Explicit : <base_home_name>Explicit, <interface_name> {
  <equivalent_explicit_operations>
};

```

where *<equivalent_explicit_operations>* are the operations defined in the home declaration (*<explicit_operations>*), with factory and finder operations transformed to their equivalent operations, as described in “Explicit Operations in Home Definitions” on page 37. The forms of the implicit and equivalent interfaces are identical to the corresponding forms for non-derived storage homes, determined by the presence or absence of a primary key specification.

A home definition with no primary key specification constitutes a pair (H, T) where H is the home type and T is the managed component type. If the home definition includes a primary key specification, it constitutes a triple (H, T, K) , where H and T are as previous and K is the type of the primary key. Given a home definition (H', T') or (H', T', K) , where K is a primary key type specified on H' , such that H' is derived from H , then T' must be identical to T or derived (directly or indirectly) from T .

Given a base home definition with a primary key (H, T, K) , and a derived home definition with no primary key (H', T') , such that H' is derived from H , then the definition of H' implicitly includes a primary key specification of type K , becoming (H', T', K) . The implicit interface for H' shall have the form specified for an implicit interface of a home with primary key K and component type T' .

Given a base home definition (H, T, K) , noting that K may have been explicitly declared in the definition of H , or inherited from a base home type, and a home definition (H', T', K') such that H' is derived from H , then T' must be identical to or derived from T and K' must be identical to or derived from K .

Note the following observations regarding these constraints and the structure of inherited equivalent interfaces:

- If a home definition does not specify a primary key directly in its header, but it is derived from a home definition that does specify a primary key, the derived home inherits the association with that primary key type, precisely as if it had explicitly specified that type in its header. This inheritance is transitive. For the purposes of the following discussion, home definitions that inherit a primary key type are considered to have specified that primary key type, even though it did not explicitly appear in the definition header.
- Operations on **CCMHome** are inherited by all home equivalent interfaces. These operations apply equally to homes with and without primary keys.
- Operations on **KeylessCCMHome** are inherited by all homes that do not specify primary keys.
- Implicitly-defined operations (i.e., that appear on the implicit interface) are only visible to the equivalent interface for the specific home type that implies their definitions. Implicitly-defined operations on a base home type are not inherited by a derived home type. Note that the implicit operations for a derived home may be identical in form to the corresponding operations on the base type, but they are defined in a different name scope.
- Explicitly-defined operations (i.e., that appear on the explicit interface) are inherited by derived home types.

6.8.5 Semantics of Home Operations

Operations in home interfaces fall into two categories:

1. Operations that are defined by the component model. Default implementations of these operations must, in some cases, be supplied by the component-enabled ORB product, without requiring user programming or intervention. Implementations of these operations must have predictable, uniform behaviors. Hence, the required semantics for these operations are specified in detail. For convenience, we will refer to these operations as *orthodox* operations.
2. Operations that are defined by the user. The semantics of these operations are defined by the user-supplied implementation. Few assumptions can be made regarding the behavior of such operations. For convenience, we will refer to these operations as *heterodox* operations.

Orthodox operations include the following:

- Operations defined on **CCMHome** and **KeylessCCMHome**.
- Operations that appear on the implicit interface for any home.

Heterodox operations include the following:

- Operations that appear in the body of the home definition, including factory operations, finder operations, and normal IDL operations and attributes.

6.8.5.1 Orthodox operations

Because of the inheritance structure described in “Home inheritance” on page 38 problems relating to polymorphism in orthodox operations are limited. For the purposes of determining key uniqueness and mapping key values to components in orthodox operations, equality of value types (given the constraints on primary key types specified in “Primary key type constraints” on page 37) are defined as follows:

- Only the state of the primary key type specified in the home definition (which is also the actual parameter type in operations using primary keys) shall be used for the purposes of determining equality. If the type of the actual parameter to the operation is more derived than the formal type, the behavior of the underlying implementation of the operation shall be as if the value were truncated to the formal type before comparison. This applies to all value types that may be contained in the closure of the membership graph of the actual parameter value; that is, if the type of a member of the actual parameter value is a value type, only the state that constitutes the member’s declared type is compared for equality.
- Two values are equal if their types are precisely equivalent and the values of all of their public state members are equal. This applies recursively to members that are value types.
- If the values being compared constitute a graph of values, the two values are equal only if the graphs are isomorphic.
- Union members are equal if both the discriminator values and the values of the union member denoted by the discriminator are precisely equal.
- Members that are sequences or arrays are considered equal if all of their members are precisely equal, where order is significant.

6.8.5.2 Heterodox operations

Polymorphism in heterodox operations is somewhat more problematic, as they are inherited by homes that may specify more-derived component and primary key types. Assume a home definition (H, T, K) , with an explicit factory operation f that takes a parameter of type K , and a home definition (H', T', K') , such that H' is derived from H , T' is derived from T , and K' is derived from K . The operation f (whose parameter type is K) is inherited by equivalent interface for H' . It may be the intended behavior of the designer that the actual type of the parameter to invocations of f on H' should be K' ,

exploiting the polymorphism implied by inheritance of K by K' . Alternatively, it may be the intended behavior of the designer that actual parameter values of either K or K' are legitimate, and the implementation of the operation determines what the appropriate semantics of operation are with respect to key equality.

This specification does not attempt to define semantics for polymorphic equality. Instead, we define the behavior of operations on home that depend on primary key values in terms of abstract tests for equality that are provided by the implementation of the heterodox operations.

Implementations of heterodox operations, including implementations of key value comparison for equality, are user-supplied. This specification imposes the following constraints on the tests for equality of value types used as keys in heterodox operations:

- For any two actual key values A and B, the comparison results must be the same for all invocations of all operations on the home.
- The comparison behavior must meet the general definition of equivalence; that is, it must be symmetric, reflexive, and transitive.

6.8.6 CCMHome Interface

The definition of the **CCMHome** interface is as follows:

```
module Components {  
  
    typedef unsigned long FailureReason;  
  
    exception CreateFailure { FailureReason reason; };  
  
    exception FinderFailure { FailureReason reason; };  
  
    exception RemoveFailure { FailureReason reason; };  
  
    exception DuplicateKeyValue { };  
  
    exception InvalidKey { };  
  
    exception UnknownKeyValue { };  
  
    interface CCMHome {  
        CORBA::IObject get_component_def();  
        CORBA::IObject get_home_def ();  
        void remove_component ( in CCMObject comp)  
            raises (RemoveFailure);  
    };  
};
```

get_component_def

The **get_component_def** operation returns an object reference that supports the **CORBA::ComponentIR::ComponentDef** interface, describing the component type associated with the home object. In strongly typed languages, the **IObject** returned must be narrowed to **CORBA::ComponentIR::ComponentDef** before use.

get_home_def

The **get_home_def** operation returns an object reference that supports the **CORBA::ComponentIR::HomeDef** interface describing the home type. In strongly typed languages, the **IObject** returned must be narrowed to **CORBA::ComponentIR::HomeDef** before use.

remove_component

The **remove_component** operation causes the component denoted by the reference to cease to exist. Subsequent invocations on the reference will cause an **OBJECT_NOT_EXIST** system exception to be raised. If the component denoted by the parameter does not exist in the container associated with target home object, **remove_component** raises a **BAD_PARAM** system exception. All other application errors raise the **RemoveFailure** exception.

Note – This specification does not define explicitly what the **FailureReason** values are for the **CreateFailure**, **FinderFailure**, and **RemoveFailure** exceptions. These values are currently vendor specific and will be standardized once consensus among vendors is established.

6.8.7 KeylessCCMHome Interface

The definition of the **KeylessCCMHome** interface is as follows:

```
module Components {
    interface KeylessCCMHome {
        CCMObject create_component() raises (CreateFailure);
    };
};
```

create_component

The **create_component** operation creates a new instance of the component type associated with the home object. A home implementation may choose to disable the parameter-less **create_component** operation, in which case it shall raise a **NO_IMPLEMENT** system exception. All other failures raise the **CreateFailure** exception.

6.9 Home Finders

The **HomeFinder** interface is, conceptually, a greatly simplified analog of the **CosLifeCycle::FactoryFinder** interface. Clients can use the **HomeFinder** interface to obtain homes for particular component types, of particularly home types, or homes that are bound to specific names in a naming service.

A reference that supports the **HomeFinder** interface may be obtained from the ORB pseudo-object by invoking **CORBA::ORB::resolve_initial_references**, with the parameter value “**ComponentHomeFinder**.” This requires the following enhancement to the **ORB** interface definition:

```

module CORBA {

    interface ORB {
        Object resolve_initial_references (in ObjectID identifier)
            raises (InvalidName);
    };
};

```

The **HomeFinder** interface is defined by the following IDL:

```

module Components {

    exception HomeNotFound { };

    interface HomeFinder {
        CCMHome find_home_by_component_type (
            in CORBA::RepositoryId comp_repid)raises (HomeNotFound);
        CCMHome find_home_by_home_type (
            in CORBA::RepositoryId home_repid) raises (HomeNotFound);
        CCMHome find_home_by_name (
            in string home_name) raises (HomeNotFound);
    };
};

```

find_home_by_component_type

The **find_home_by_component_type** operation returns a reference, which supports the interface of a home object that manages the component type specified by the **comp_repid** parameter. This parameter contains the repository identifier of the component type required. If there are no homes that manage the specified component type currently registered, the operation shall raise the HomeNotFound exception.

*Little is guaranteed about the home interface returned by this operation. If the definition of the returned home specified a primary key, there is no generic factory operation available on any standard interface (i.e. pre-defined, as opposed to generated type-specific interface) supported by the home. The only generic factory operation that is potentially available is **Components::KeylessCCMHome::create_component**. The client must first attempt to narrow the **CCMHome** reference returned by the **find_home_by_component_type** to **KeylessCCMHome**. Otherwise, the client must have specific out-of-band knowledge regarding the home interface that may be returned, or the client must be sophisticated enough to obtain the **HomeDef** for the home and use the DII to discover and invoke a create operation on a type-specific interface supported by the home.*

find_home_by_home_type

The **find_home_by_home_type** operation returns a reference that supports the interface of the type specified by the repository identifier in the **home_repid** parameter. If there are no homes of this type currently registered, the operation shall raise the HomeNotFound exception.

*The current **LifeCycle.find_factories** operation returns a sequence of factories to the client requiring the client to choose the one that will create the instance. Based on the experience of the submitters, CORBA components defines operations which allows the server to choose the “best” home for the client request based on its knowledge of workload, etc.*

Since the operation returns a reference to **CCMHome**, it must be narrowed to the specific home type before it can be used.

find_home_by_name

The **find_home_by_name** operation returns a home reference bound to the name specified in the **home_name** parameter. This parameter is expected to contain a name in the format described in the *Naming Service* specification (formal/01-02-65), section 2.4, “Stringified Names.” The implementation of this operation may be delegated directly to an implementation of CORBA naming, but it is not required. The semantics of the implementation are considerably less constrained, being defined as follows:

- The implementation is free to maintain multiple bindings for a given name, and to return any reference bound to the name.

It is generally expected that implementations that do not choose to use CORBA naming will do so for reasons of scalability and flexibility, in order, for example, to provide a home which is logically more “local” to the home finder (and thus, the client).

- The client’s expectations regarding the returned reference, other than that it supports the **CCMHome** interface, are not guaranteed or otherwise mediated by the home. The fact that certain names may be expected to provide certain home types or qualities of implementation are outside the scope of this specification.

This is no different than any application of naming services in general. Applications that require clients to be more discriminating are free to use the Trader service, or any other similar mechanism that allows query or negotiation to select an appropriate home. This mechanism is intentionally kept simple.

If the specified name does not map onto a home object registered with the finder, the operation shall raise the **HomeNotFound** exception.

6.10 Component Configuration

The CORBA component model provides mechanisms to support the concept of component *configurability*.

Experience has proven that building re-usable components involves making difficult trade-offs between providing well-defined, reasonably-scoped functionality, and providing enough flexibility and generality to be useful (or re-useful) across a variety of possible applications. Packaging assumptions of the component architecture preclude customizing a component’s behavior by directly altering its implementation or (in most cases) by deriving specialized sub-types. Instead, the model focuses on extension and customization through delegation (e.g., via dependencies expressed with uses declarations) and configuration. Our assumption is that generalized components will typically provide a set of optional behaviors or modalities that can be selected and adjusted for a specific application.

The configuration framework is designed to provide the following capabilities:

- *The ability to define attributes on the component type that are used to establish a component instance’s configuration. Component attributes are intended to be used during a component instance’s initialization to establish its fundamental behavioral properties. Although the component model does not constrain the visibility or use of attributes defined on the component, it is generally assumed that they will not be of interest to the same clients that will use the component after it is configured. Rather, it is intended for use by component factories or by deployment tools in the process of instantiating an assembly of components.*
- *The ability to define a configuration in an environment other than the deployment environment (e.g., an assembly tool), and store that configuration in a component package or assembly package to be used subsequently in deployment.*
- *The ability to define such a configuration without having to instantiate the component type itself.*
- *The ability to associate a pre-defined configuration with a component factory, such that component instances created by that factory will be initialized with the associated configuration.*

- *Support for visual, interactive configuration tools to define configurations. Specifically, the framework allows component implementors to provide a configuration manager associated with the component implementation. The configuration manager interface provides descriptive information to interactive users, constrains configuration options, and performs validity checks on proposed configurations.*

The CORBA component model allows a distinction to be made between interface features that are used primarily for configuration, and interface features that are used primarily by application clients during normal application operation. This distinction, however, is not precise, and enforcement of the distinction is largely the responsibility of the component implementor.

It is the intent of this specification (and a strong recommendation to component implementors and users) that operational interfaces should be either provided interfaces or supported interfaces. Features on the component interface itself, other than provided interfaces, (i.e., receptacles, event sources and sinks) are generally intended to be used for configuration, although there is no structural mechanism for limiting the visibility of the features on a component interface. A mechanism is provided for defining configuration and operational phases in a component's life cycle, and for disabling certain interfaces during each phase.

The distinction between configuration and operational interfaces is often hard to make in practice. For example, we expect that operational clients of a component will want to receive events generated by a component. On the other hand, some applications will want to establish a fixed set of event source and sink connections as part of the overall application structure, and will want to prevent clients from changing those connections. Likewise, the responsibility for configuration may be hard to assign—in some applications the client that creates and configures a component may be the same client that will use it operationally. For this reason, the CORBA component model provides general guidelines and optional mechanisms that may be employed to characterize configuration operations, but does not attempt to define a strict separation of configuration and operational behaviors.

6.10.1 Exclusive Configuration and Operational Life Cycle Phases

A component implementation may be designed to implement an explicit configuration phase of its life cycle, enforcing serialization of configuration and functional operation. If this is the case, the component life cycle is divided into two mutually exclusive phases, the *configuration phase* and the *operational phase*.

The **configuration_complete** operation (inherited from **Components::CCMObject**) is invoked by the agent effecting the configuration to signal the completion of the configuration phase. The **InvalidConfiguration** exception is raised if the state of the component configuration state at the time **configuration_complete** is invoked does not constitute an acceptable configuration state. It is possible that configuration may be a multi-step process, and that the validity of the configuration may not be determined until the configuration process is complete. The **configuration_complete** operation should not return to the caller until either 1) the configuration is deemed invalid, in which case the **InvalidConfiguration** exception is raised, or 2) the component instance has performed whatever work is necessary to consolidate the final configuration and is prepared to accept requests from arbitrary application clients.

*In general, component implementations should defer as much consolidation and integration of configuration state as possible until **configuration_complete** is invoked. In practice, configuring a highly-connected distributed object assembly has proven very difficult, primarily because of subtle ordering dependencies that are difficult to discover and enforce. If possible, a component implementation should not be sensitive to the ordering of operations (interface connections, configuration state changes, etc.) during configuration. This is one of the primary reasons for the definition of **configuration_complete**.*

6.10.1.1 Enforcing exclusion of configuration and operation

The implementation of a component may choose to disable changes to the configuration after **configuration_complete** is invoked, or to disable invocations of operations on provided interfaces until **configuration_complete** is invoked. If an implementation chooses to do either (or both), an attempt to invoke a disabled operation should raise a **BAD_INV_ORDER** system exception.

Alternatively, a component implementation may choose not to distinguish between configuration phase and deployment phase. In this case, invocation of **configuration_complete** will have no effect.

The component implementation framework provides standard mechanisms to support disabling operations during configuration or operation. Certain operations are implemented by the component implementation framework (see the *CCM Implementation Framework*) and may not be disabled.

6.11 Configuration with Attributes

A component's configuration is established primarily through its attributes. An *attribute configuration* is defined to be a description of a set of invocations on a component's attribute set methods, with specified values as parameters.

There are a variety of possible approaches to attribute configuration at run time, depending on the design of the component implementation and the needs of the application and deployment environments. The CORBA component model defines a set of basic mechanisms to support attribute configuration. These mechanisms can be deployed in a number of ways in a component implementation or application.

6.11.1 Attribute Configurators

A configurator is an object that encapsulates a specific attribute configuration that can be reproduced on many instances of a component type. A configurator may invoke any operations on a component that are enabled during its configuration phase. In general, a configurator is intended to invoke attribute set operations on the target component.

6.11.1.1 The Configurator interface

The following interface is supported by all configurators:

```
module Components {  
  
    exception WrongComponentType { };  
  
    interface Configurator {  
        void configure (in CCMObject comp)  
            raises (WrongComponentType);  
    };  
};
```

configure

The **configure** operation establishes its encapsulated configuration on the target component. If the target component is not of the type expected by the configurator, the configure operation shall raise the **WrongComponentType** exception.

6.11.1.2 The StandardConfigurator interface

The **StandardConfigurator** has the following definition:

```
module Components {  
  
    valuetype ConfigValue {  
        public FeatureName name;  
        public any value;  
    };  
  
    typedef sequence<ConfigValue> ConfigValues;  
  
    interface StandardConfigurator : Configurator {  
        void set_configuration (in ConfigValues descr);  
    };  
};
```

The **StandardConfigurator** interface supports the ability to provide the configurator with a set of values defining an attribute configuration.

set_configuration

The **set_configuration** operation accepts a parameter containing a sequence of **ConfigValue** instances, where each **ConfigValue** contains the name of an attribute and a value for that attribute, in the form of an **any**. The **name** member of the **ConfigValue** type contains the unqualified name of the attribute as declared in the component definition IDL. After a configuration has been provided with **set_configuration**, subsequent invocations of **configure** will establish the configuration on the target component by invoking the set operations on the attributes named in the value set, using the corresponding values provided in the **anys**. Invocations on attribute set methods will be made in the order in which the values occur in the sequence.

6.11.2 Factory-based Configuration

Factory operations on home objects may participate in the configuration process in a variety of ways. A factory operation may

- be explicitly implemented to establish a particular configuration.
- apply a configurator to newly-created component instances. The configurator may be supplied by an agent responsible for deploying a component implementation or a component assembly.
- apply an attribute configuration (in the form of a **Components::ConfigValues** sequence) to newly-created instances. The attribute configuration may be supplied to the home object by an agent responsible for deploying a component implementation or a component assembly.
- be explicitly implemented to invoke **configuration_complete** on newly-created component instances, or to leave component instances open for further configuration by clients.
- be directed by an agent responsible for deploying a component implementation or assembly to invoke **configuration_complete** on newly-created instances, or to leave them open for further configuration by clients.

If no attribute configuration is applied by a factory or by a client, the state established by the component implementation's instance initialization mechanism (e.g., the component servant constructor) constitutes the default configuration.

6.11.2.1 HomeConfiguration interface

The implementation of a component type's home object may optionally support the **HomeConfiguration** interface. The **HomeConfiguration** interface is derived from **Components::CCMHome**. In general, the **HomeConfiguration** interface is intended for use by an agent deploying a component implementation into a container, or an agent deploying an assembly.

The **HomeConfiguration** interface allows the caller to provide a **Configurator** object and/or a set of configuration values that will be applied to instances created by factory operations on the home object. It also allows the caller to cause the home object's factory operations to invoke **configuration_complete** on newly-created instances, or to leave them open for further configuration.

The **HomeConfiguration** allows the caller to disable further use of the **HomeConfiguration** interface on the home object.

*The **Configurator** interface and the **HomeConfiguration** interface are designed to promote greater re-use, by allowing a component implementor to offer a wide range of behavioral variations in a component implementation. As stated previously, the CORBA component specification is intended to enable assembling applications from pre-built, off-the-shelf component implementations. An expected part of the assembly process is the customization (read: configuration) of a component implementation, to select from among available behaviors the behaviors suited to the application being assembled. We anticipate that assemblies will need to define configurations for specific component instances in the assembly, but also that they will need to define configurations for a deployed component type, i.e., all of the instances of a component type managed by a particular home object.*

The **HomeConfiguration** interface is defined by the following IDL:

```
module Components {  
  
    interface HomeConfiguration : CCMHome {  
        void set_configurator (in Configurator cfg);  
        void set_configuration_values (  
            in ConfigValues config);  
        void complete_component_configuration (in boolean b);  
        void disable_home_configuration();  
    };  
};
```

set_configurator

This operation establishes a configurator object for the target home object. Factory operations on the home object will apply this configurator to newly-created instances.

set_configuration_values

This operation establishes an attribute configuration for the target home object, as an instance of **Components::ConfigValues**. Factory operations on the home object will apply this configurator to newly-created instances.

complete_component_configuration

This operation determines whether factory operations on the target home object will invoke **configuration_complete** on newly-created instances. If the value of the boolean parameter is **TRUE**, factory operations will invoke **configuration_complete** on component instances after applying any required configurator or configuration values to the instance. If the parameter is **FALSE**, **configuration_complete** will not be invoked.

disable_home_configuration

This operation serves the same function with respect to the home object that the **configuration_complete** operation serves for components. This operation disables further use of operations on the **HomeConfiguration** interface of the target home object. If a client attempts to invoke **HomeConfiguration** operations, the request will raise a **BAD_INV_ORDER** system exception. This operation may also be interpreted by the implementation of the home as demarcation between its own configuration and operational phases, in which case the home implementation may disable operations and attributes on the home interface.

If a home object is supplied with both a configurator and a set of configuration values, the order in which **set_configurator** and **set_configuration_values** are invoked determines the order in which the configurator and configuration values will be applied to component instances. If **set_configurator** is invoked before **set_configuration_values**, the configurator will be applied before the configuration values, and vice-versa.

The component implementation framework defines default implementations of factory operations that are automatically generated. These generated implementations will behave as specified here. Component implementors are free to replace the default factory implementations with customized implementations. If a customized home implementation chooses to support the **HomeConfiguration** interface, then the factory operation implementations must behave as specified, with respect to component configuration.

6.12 Component Inheritance

The mechanics of component inheritance are defined by the inheritance relationships of the equivalent IDL component interfaces. The following rules apply to component inheritance:

- All interfaces for non-derived component types are derived from **CCMObject**.
- If a component type directly supports one or more IDL interfaces, the component interface is derived from both **CCMObject** and the supported interfaces.
- A derived component type may not directly support an interface.
- The interface for a derived component type is derived from the interface of its base component type.
- A component type may have at most one base component type.
- The features of a component that are expressed directly on the component interface are inherited as defined by IDL interface inheritance. These include:
 - operations implied by **provides** statements
 - operations implied by **uses** statements
 - operations implied by **emits** statements
 - operations implied by **publishes** statements

- operations implied by **consumes** statements
- attributes

6.12.1 CCMObject Interface

The **CCMObject** interface is defined by the following IDL:

```

module Components {

    valuetype ComponentPortDescription
    {
        public FacetDescriptions facets;
        public ReceptacleDescriptions receptacles;
        public ConsumerDescriptions consumers;
        public EmitterDescriptions emitters;
        public PublisherDescriptions publishers;
    };

    exception NoKeyAvailable { };

    interface CCMObject : Navigation, Receptacles, Events {
        CORBA::IObject get_component_def ( );
        CCMHome get_ccm_home ( );
        PrimaryKeyBase get_primary_key( ) raises (NoKeyAvailable);
        void configuration_complete( ) raises (InvalidConfiguration);
        void remove() raises (RemoveFailure);
        ComponentPortDescription get_all_ports ( );
    };
};

```

get_component_def

This operation returns an **IObject** reference to the component definition in the Interface Repository. The interface repository representation of a component is defined in the *Interface Repository Metamodel*. In strongly typed languages, the **IObject** returned must be narrowed to **CORBA::ComponentIR::ComponentDef** before use.

get_ccm_home

This operation returns a **CCMHome** reference to the home that manages this component.

get_primary_key

This operation is equivalent to the same operation on the component's home interface. It returns a primary key value if the component is being managed by a home that defines a primary key. Otherwise, the **NoKeyAvailable** exception shall be raised.

configuration_complete

This operation is called by a configurator to indicate that the initial component configuration has completed. If the component determines that it is not sufficiently configured to allow normal client access, it raises the `InvalidConfiguration` exception. The component configuration process is described in “Component Configuration” on page 44.

remove

This operation is used to delete a component. Application failures during **remove** may raise the `RemoveFailure` exception.

get_all_ports

The **get_all_ports** operation returns a value of type `ComponentPortDescription` containing information about all facets, receptacles, event sinks, emitted events and published events in the component’s inheritance hierarchy. The order in which the information occurs in these sequences is not specified. If a component does not offer a port of any type, the associated sequence will have length zero.

6.13 Conformance Requirements

This sub clause identifies the conformance points required for compliant implementations of the CORBA Component model. The following conformance points are defined:

1. A CORBA COS vendor shall provide the relevant changes to the Lifecycle, Transaction, and Security Services identified in the following “Changes to Object Services” on page 53.
2. A CORBA ORB vendor need not provide implementations of Components aside from the changes made to the Core to support components. Conversely a CORBA Component vendor need not be a CORBA ORB vendor.
3. A CORBA Component vendor shall provide a conforming implementation of the Basic Level of CORBA Components. A Lightweight CORBA Component vendor shall provide a conforming implementation of the Lightweight CCM Profile as specified in item 8 below.
4. A CORBA Component vendor may provide a conforming implementation of the Extended Level of CORBA Components.
5. To be conformant at the Basic level a non-Java product shall implement (at a minimum) the following:
 - The IDL extensions and generation rules to support the client and server side component model for basic level components.
 - CIDL. The multiple segment feature of CIDL (“Segment Definition” on page 88) need not be supported for basic components.
 - A container for hosting basic level CORBA components.
 - The XML deployment descriptors and associated zip files for basic components.Such implementations shall work on a CORBA ORB as defined in #1 above.
6. To be conformant at the Basic level a Java product shall implement (at a minimum):
 - EJB1.1, including support for the EJB 1.1 XML DTD.
 - The java to IDL mapping, also known as RMI/IIOP.

- EJB to IDL mapping as defined in “Translation of CORBA Component requests into EJB requests” on page 183.

Such implementations shall work in a CORBA interoperable environment, including interoperable support for IIOP, CORBA transactions and CORBA security.

7. To be conformant at the extended level, a product shall implement (at a minimum) the requirements needed to achieve Basic PLUS:
 - IDL extensions to support the client and server side component model for extended level components.
 - A container for hosting extended level CORBA components.
 - The XML deployment descriptors and associated zip files for basic and enhanced level components in the format defined in “Deployment PSM for CCM” on page 309.

Such implementations shall work on a CORBA ORB as defined in #1 above.
8. The Lightweight CCM profile is a conformance point based on the extended model as defined above. “Lightweight CCM Profile” on page 301 defines the specific parts of this CCM specification that are impacted and the normative specific subsetting of CCM. In summary, the following general capabilities (and associated machinery) are excluded from the extended model to define this conformance point:
 - Persistence (only session and service components are supported)
 - Introspection
 - Navigation
 - Redundancies, preferring generic over specific
 - Segmentation (not allowed for session or service components)
 - Transactions
 - Security
 - Configurators
 - Proxy homes
 - Home finders
 - CIDL
 - POA related mandates
9. A CORBA Component vendor may optionally support EJB clients interacting with CORBA Components, by implementing the IDL to EJB mapping as defined in “Translation of EJB Requests into CORBA Component Requests” on page 190.
10. This specification includes extensions to IDL, in the form of new keywords and grammar. Although a CORBA ORB vendor need not be a CORBA Component vendor, and vice-versa, it is important to maintain IDL as a single language. To this end, all compliant products of any conformance points above shall be able to parse any valid IDL definitions. However, it is permitted to raise errors, or to ignore, those parts of the grammar that relate to another conformance point.

Conforming implementations as defined above may also implement any additional features of this specification not required by the above conformance points.

6.13.1 A Note on Tools

Component implementations are expected to be supported by tools. It is not possible to define conformance points for tools, since a particular tool may only support part of the component development and deployment life-cycle. Hence a suite of tools may be needed. The Component architecture contains a number of definitions that are relevant to tools, including zip files and XML formats, as well as IDL interfaces for customization and installation. Although it cannot be enforced, tools are expected to conform to the relevant areas with which they are dealing. For example, a tool that generates implementations for a particular platform is expected to generate XML according to the `<implementation>` clauses in the DTD (defined in Part 1 - CORBA Interfaces, the *Interface Repository* clause).

6.13.2 Changes to Object Services

6.13.2.1 Life Cycle Service

To support the factory design pattern for creating a component instance and to allow the server, rather than a client, to select from a group of functionally equivalent factories based on load or other server-side visible criteria, the following operation is added to the **FactoryFinder** interface of the **CosLifeCycle** module:

```
module CosLifeCycle {
    interface FactoryFinder {
        Factory find_factory (in Key factory_key) raises (noFactory);
    };
};
```

The parameters of the above operation are as defined by **CosLifeCycle** with the following clarifications:

- The **factory_key** parameter is a name conforming to the Interoperable Naming Specification (orbos/98-10-11) for stringified names.
- The **factory_key** parameter is used as an input to the **find_home_by_name** operation on **Components::HomeFinder**.
- The default factory operation on the home is used to obtain a reference which can be narrowed to the **CosLifeCycle::GenericFactory** type.

6.13.2.2 Transaction Service

The following CORBA transaction service interface is changed to a local interface:

- **CosTransactions::Current**

6.13.2.3 Security Service

The following CORBA Security interfaces are changed to local interfaces:

- **SecurityLevel1::Current**
- **SecurityLevel2::PrincipalAuthenticator**
- **SecurityLevel2::Credentials**
- **SecurityLevel2::ReceivedCredentials**

- **SecurityLevel2::AuditChannel**
- **SecurityLevel2::AuditDecision**
- **SecurityLevel2::AccessDecision**
- **SecurityLevel2::QOPPolicy**
- **SecurityLevel2::MechanismPolicy**
- **SecurityLevel2::InvocationCredentialsPolicy**
- **SecurityLevel2::EstablishTrustPolicy**
- **SecurityLevel2::DelegationDirectivePolicy**
- **SecurityLevel2::Current**
- **SecurityReplacable::Vault**
- **SecurityReplacable::SecurityContext**
- **SecurityReplacable::ClientSecurityContext**
- **SecurityReplacable::ServerSecurityContext**

7 Generic Interaction Support

7.1 Introduction

The Generic Interaction Support includes the definition of *extended ports* and *connectors*. Extended ports can be used at component level to specify the programming contracts that the components need to fulfill in order to interact with other components. Connectors are the entities that can be connected to components via these extended ports, in order to actually realize the interactions.

These extensions fall within the scope of adapting CCM model to specialized application domains, in particular embedded and real-time systems. The lightweight CCM specification has defined a profile to meet embedded equipments. QoS for CCM, '[QOS4CCM]' allows providing non-functional services to components and by this mean allows the use of real-time services plugged into the container. This Generic Interaction Support complements these adaptations with the ability to provide interactions or communication patterns (control of flow, synchronous, asynchronous, shared memory ...) very specific to real-time software.

As for non-functional services, connectors can be platform dependent because they deal with specific communication buses (1553, UDP, TCP, direct calls...) or specific semantics (management of buffers, threads, mutex... inside the fragment). For this reason, they are rather intended to be provided by CCM framework providers or platform providers.

7.2 Simple Generic Interaction Support

7.2.1 Overview

The GIS relies on two constructs: extended ports and connectors. Extensions to IDL3¹ are provided to allow defining and using those constructs.

IDL3+ declarations can be easily translated in plain IDL3. The following figure presents the steps of component definition. Only the first step is new and will be detailed in the following sub clauses.

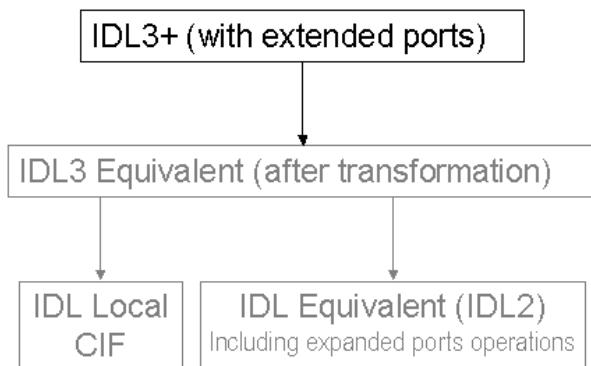


Figure 7.1 - IDL3+ Transformation

1. IDL3 plus its extensions is called IDL3+ in this specification.

As resulting IDL3 is exactly as before, the rest of the transformation is kept unchanged.

The transformation from IDL3+ to equivalent IDL3 shall be done by a tool part of a CCM framework implementing the current specification.

7.2.2 Extended Ports

An extended port is the mean to represent the programming contract that the components need to deal with, in order to interact according to the corresponding interaction pattern. A programming contract can always be expressed by means of interfaces to call and/or interfaces to be called. Extended ports can be thus subsumed in a group of provided/required interfaces, which can be used/provided. In other words, extended ports are just groups of single CCM ports (facet/**provides** and receptacle/**uses**)² and attributes to configure them if needed.

7.2.2.1 IDL3+ Representation

A new keyword **porttype** has been added to IDL3³ to allow defining extended ports. An extended port definition consists in a list of basic ports (**uses** and/or **provides**) and attributes.

A second new keyword **port** allows to set a previously defined extended port to a component.

The following is an example of such definitions:

```
//-----  
// IDL3+  
//-----  
  
interface Data_Pusher {  
    void push(in Data dat);  
};  
  
interface FlowControl {  
    void suspend ();  
    void resume();  
    readonly attribute nb_waiting;  
};  
  
// Extended port definition  
porttype Data_ControlledConsumer {  
    provides Data_Pusher consumer;  
    uses FlowControl control;  
};  
  
// Component declaration with that port  
component C1 {  
    port Data_ControlledConsumer p;  
};
```

-
2. The receptacles correspond to the interfaces that the components will call and the facets, the ones that they will provide to be called.
 3. In this sub clause and the following, the new syntax is just introduced. Formal definition of the new grammar is in sub clause 7.4.

In the original CCM, existing port kinds are seen as groups of matching basic ports (provided/required interfaces, or events sinks/sources). Similarly, it is needed to define inverses of extended ports (i.e., the ones that will “match” them). To avoid duplicated definitions, the keyword **mirrorport** has been introduced for that purpose. A **mirrorport** results in exactly the same number of simple ports as the **port** of the same **porttype**, except that all the **uses** are turned into **provides** and vice-versa.

7.2.2.2 Translation from Extended Ports to Basic Ports

The extensions provided to IDL3 with **porttype**, **port** and **mirrorport** keywords can be directly mapped to usual IDL3 constructs (basic port declarations).

The rules for this transformations are as follows:

- A **provides** in a **port** becomes a **provides** in the equivalent IDL3 declaration of the component;
- A **uses** in a **port** becomes a **uses** in the equivalent IDL3 declaration of the component;
- A **provides** in an **mirrorport** becomes a **uses** in the equivalent IDL3 declaration of the component;
- A **uses** in a **mirrorport** becomes a **provides** in the equivalent IDL3 declaration of the component;
- The name of the basic port is the concatenation of the extended port name and the related basic port name of the **porttype**, separated by '_';
- An **attribute** in a port or a **mirrorport** becomes an **attribute** in the equivalent IDL3 declaration of the component;
- The name of the generated attribute is the concatenation of the extended port name and the related attribute name of the **porttype**, separated by '_'.

Applying these rules, the previous example will result in the following IDL3 declaration:

```
// Resulting IDL3 component definition
component C1 {
    provides Data_Pusher p_consumer;
    uses FlowControl p_control;
};
```

7.2.3 Connectors

Connectors are used to specify an interaction mechanism between components. Connectors can have ports in the same way as components. They can be composed of simple ports (CCM **provides** and **uses**) or extended ports⁴.

The following figure shows a connector as it can be represented at design level:



4. As generally components will be given extended ports by means of keyword **port**, it is very likely that connectors will use **mirrorport** instead.



Figure 7.2 - Logical View of a Connector

The connector will concretely be composed of several parts (called *fragments*) that will consist of executors, each in charge of realizing a part of the interaction. Each fragment will be co-localized to the component using them.

By default, for each port, a fragment (an executor) is produced. If several ports are always co-localized because it corresponds to the semantic of the connector, their behavior can be provided by the same fragment. This is an implementation choice for the connector developer.

The following figure shows the connector with its fragments at execution time:

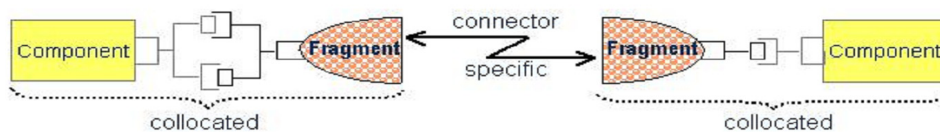


Figure 7.3 - Connector Representation at Execution Time (Fragments)

The communication mechanism between the fragments is connector specific and will be addressed only for DDS support in this specification.

The connector concept brings another way of seeing CCM: connectors are used to provide interaction (in particular communication support) between components, and are realized via fragments collocated with the concerned components. This contrasts with the classical approach, which entails CORBA servants for facets typically provided by code generation and encapsulating the component executors. An implementation compliant with the present connector specification is not required to provide CORBA servants and CORBA object references for the component facets.

7.2.3.1 IDL3+ Representation

The new IDL3 keyword **connector** allows defining connectors. A connector definition is very similar to a simplified component's one as a connector is just meant to gather ports (simple or extended). It thus cannot include the **support** keyword.

The following is an example of a connector definition:

```
connector Data_Cnx {
  mirrorport Data_ControlledConsumer cc;
  provides Data_Pusher p;
};
```

7.2.3.2 Connector Attributes

A connector can declare attributes in the same way as components. Attributes are declared at connector definition level and are reflected in each fragment at realization level. For instance in a DDS connector, the topic can be seen as an attribute and the value of the topic is reflected on each fragment that composes the connector: each fragment of the connector will work on the same topic.

7.2.3.3 Connector Inheritance

A connector can inherit from another connector. It means that the new connector is composed of all the ports and attributes of the inherited connector in addition to all the ones that are locally defined.

The syntax used to declare a connector inheritance is similar to the one used to declare a component inheritance.

7.2.3.4 Composite Connectors

A connector (type) can have multiple implementations. As it is the case for components, such an implementation may be an assembly of other components. For example, an implementation of a local FIFO queue can be provided by a monolithic implementation, but if this FIFO should enable distribution, an alternative implementation needs to provide multiple fragments co-localized with the components using them. These fragments can be considered as sub-components within an assembly (parts within UML composite structures), i.e., an implementation of a connector with multiple fragments is an assembly implementation. There is no restriction on the level of assembly implementations, for instance a fragment might itself be realized by an assembly implementation. The advantage of assembly implementations is twofold: first, they enable to express the fragmented implementation of connectors by concepts already existing in CCM. Second, assembly implementations enable the composition of connectors, which facilitates the development of new connectors.

Consider the example of remote a FIFO. One possible implementation is a FIFO on the consumer's site and a remote access. The structure of such a remote FIFO implementation is shown in Figure 7.4. It is composed of two fragments called respectively SocketClient and FIFO_Socket_f_pull.

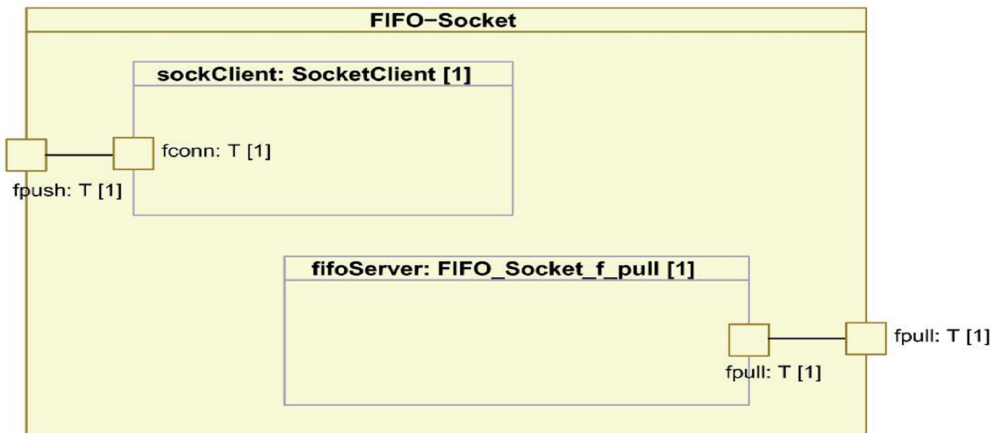


Figure 7.4 - Example of a Distributed FIFO Implementation

Figure 7.5 shows the detailed implementation of the second fragment (**FIFO_Socket_f_pull**) that is itself an assembly of 2 fragments: **SocketServer** and **ConnFIFO**.

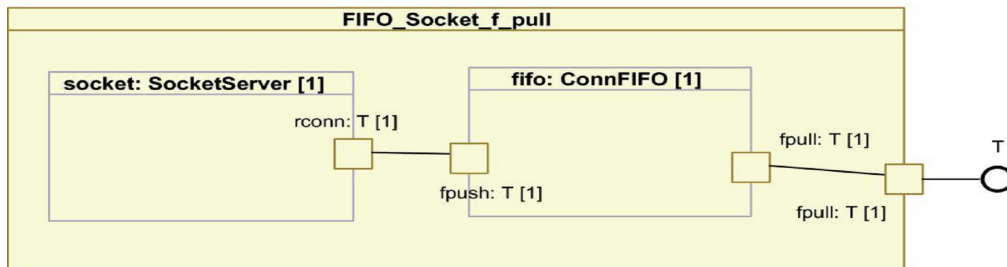


Figure 7.5 - Assembly Implementation of a Connector Fragment

It is thus possible to create new connector implementations by re-assembling existing connectors or fragment implementations. In case of the example, the socket could be replaced by another transport mechanism, for instance an inter process communication.

7.2.3.5 Translation to IDL3

The mapping of connector definition to standard IDL3 is trivial: The connector definitions are removed, but the information is used to provide type information at the assembly level. The information shall be described at assembly level to check whether the binding of two ports from a component and a connector, respectively binds identical provided and used types or vice versa.

In a CCM framework providing the support for connector extension, the connector definition in “IDL3+” can be used to generate partly the fragment executors where the connector implementation will be realized (see “Programming Model for Connectors” on page 63 for more details on connectors realization).

7.3 Generic Interaction Support with Templates

As extended ports and connectors are meant to capture interaction logics, their main benefit is obtained if they can be parameterized by types. In the above example, the port `Data_ControlledConsumer` and the connector `Data_Cnx` are only valid for manipulating elements of specific type `Data`. It would be very useful to define generic port type and connector that would provide similar interaction logics to any type of data.

For that purpose, an extension allowing parameterizing definitions of interfaces, ports and connectors has been added to IDL3+. Parameterized definitions can easily be resolved at IDL3+ compilation time.

7.3.1 Template Support Overview

The template support is aiming at integrating smoothly in the current IDL specification. It follows therefore the current rules that apply to existing predefined templates (such as sequences): the syntax is very lightweight and anonymous types are not allowed.

The template support allows defining all sensible parameterized interaction support.

Almost any sensible parameterized interaction support will associate at least one port type (itself comprising at least one parameterized interface) with one connector. *The components* that integrate the concrete port (resulting from the instantiation of the port type with a given parameter type) *and the related concrete connector* that provides the mirror port of the same port type *need eventually to use/provide exactly the same interface instantiation.*

If the port type on one hand and the connector on the other hand were placed in separate template definitions, this constraint would not be achievable due to anonymous types not being allowed. There is therefore a necessity to offer the mean to group several identically parameterized definitions in the same template scope. Modules are the only IDL grouping constructs. *Therefore the template support is introduced at the module level.*

Note that as parameterizing a module will result in de facto parameterization of all the embedded constructs, this support, offers a lot of possibilities despite its limited impact on the IDL grammar.

7.3.2 Template Modules

Using template modules is a two-stepped process:

- First of all the template module is declared.
- Secondly, its instantiation results in a concrete module that is usable as any module.

In addition, a template module can be referenced inside another similarly parameterized module, in order to reuse the related definitions.

7.3.2.1 Template Module Declaration

A template module is declared in adding to its declaration a list of comma-separated formal parameters embedded between angular brackets (< and >).

Formal parameters associate a type constraint and the formal parameter name. At instantiation time each formal parameter will be substituted by a concrete value. Only the concrete values that comply with their formal parameter type constraints will be accepted.

Type constraints can be:

- **typename**, meaning that any type will be acceptable;
- some more restricted type designators:
 - **interface**, meaning all interfaces;
 - **valuetype**, meaning all valuetype types;
 - **eventtype**, meaning all eventtypes;
 - **struct**, meaning all struct types;
 - **union**, meaning all union types;
 - **sequence**, meaning all sequence types;
 - **enum**, meaning all enum types;
 - a **const** primitive type, meaning that any constant of the required primitive type will be acceptable;
 - a sequence specification, with the constraints that its formal parameters must appear previously in the formal parameters list of the module. In this case, the passed parameter should be a sequence complying with the sequence specification⁵.

The following is a refactoring of the previous example, which has been generalized in order to be usable with any data type.

```

interface FlowControl {
    void suspend ();
    void resume();
    readonly attribute nb_waiting;
};

module Flow <typename T> {
    interface Pusher {
        void push(in T dat);
    };

    // Extended port definition
    porttype ControlledConsumer {
        provides Pusher consumer;
        uses FlowControl control;
    };

    // Connector
    connector Cnx {
        mirrorport ControlledConsumer cc;
        provides Pusher p;
    };
};

```

Note that all constructs that are not T-dependent (here the FlowControl interface) have been put outside the template module to avoid useless duplications. Note also that T-dependency of a construct may be direct, because the formal type is used in the definition (here the Pusher interface) or indirect when the definition makes use of a T-dependent construct (here the ControlledConsumer port type or the Cnx connector).

7.3.2.2 Template Module Instantiation

Once defined a template module has to be explicitly instantiated before being used. Instantiation consist in providing actual values to any formal parameters and a name to the resulting concrete module.

This is done by declaring the concrete module with a new form of the module declaration that inserts, between the keyword module and the module name, the template module instantiation with all values for formal parameters enclosed in angular brackets.

When the module is instantiated, all the embedded constructs are de facto instantiated with the proper parameters values.

The following is an example of instantiating the previously defined template module with the data type Data and using the port type in a component.

```

// module instantiation
module Flow<Data> Data_Flow;

```

-
5. This disposal is useful to pass, without duplication, an existing sequence type. Actually, the removal of anonymous types from IDL leads to each similar sequence instantiation be a different type. In case the interaction support needs to manipulate sequence<T> (T being the formal parameter of the template), then there is no means to use the same sequence as the rest of the application but to pass it as a formal parameter.

```

// component declaration
component C2 {
    port Data_Flow::ControlledConsumer p;
};

```

Applying then the IDL3+ to IDL3 translation rules will give the following result:

```

// Resulting IDL3 component definition
component C2 {
    provides Data_Flow::Pusher p_consumer;
    uses FlowControl p_control;
};

```

7.4 IDL3+ Grammar

Although they were originally defined here, the IDL3+ extensions are now maintained in a separate OMG specification. See [IDL4] sections 7.4.11 and 7.4.12 for the relevant Building Blocks and section 9.4.2 for the CCM with Generic Interaction Profile.

7.5 Programming Model for Connectors

This sub clause presents the rules a connector implementer has to follow. This is the counterpart of the CCM component model interfaces for connectors and connector ports. As presented in 7.5.6, 'Connector Implementation Interfaces' connectors' implementations consist in the collaboration of several objects, named fragments, or connector executors. They realize the implementation of the connector ports and are collocated with the components logically connected to the connector. The programming model is oriented towards the provision of objects corresponding to the ports of the connector under consideration. It is therefore composed of an API for fragments programming and fragments bootstrapping.

The following are the interfaces necessary to implement a connector's fragment:

```

module Components {
    interface CCMObject :Navigation,
                Receptacles,
                Events{
        CCMHome get_ccm_home();
        void configuration_complete() raises ( InvalidConfiguration );
        void remove () raises ( RemoveFailure );
    };

    interface KeylessCCMHome {
        CCMObject create_component() raises ( CreateFailure );
    };

    interface CCMHome {
        void remove_component() raises ( RemoveFailure );
    };

    interface HomeConfiguration : CCMHome {

```

```

        void set_configuration_values( in ConfigValues config );
    };
};

```

This presents the interfaces that need to be implemented by a connector provider. A **Components::CCMObject** interface has to be implemented for each identified fragment of the connector.

This set of interfaces is a subset of the component model coming from the Lightweight CCM specification. All the previous methods declared in interfaces have to be defined in the fragment implementation, in order to conform to all D&C deployment tools.

In a fragment's implementation, some of these interfaces could be left empty, others are mandatory, among them: **Navigation**, **Receptacles**, and **KeylessCCMHome**.

Note that the **Events** CCM interface is never used in connectors' executors. The reason is that, as the component and its connector's fragment are collocated, they only interact via synchronous calls (a potential asynchronous nature of the actual interaction between components would be provided by connector's fragments themselves).

7.5.1 Interface CCMObject

Given a **porttype**, the fragment inheriting **CCMObject** has to implement all necessary operations (**provide_facet**, **connect**, **disconnect**...) inherited from **Components::Navigation** and **Components:: Receptacles** interfaces in accordance with the **porttype**.

7.5.2 configuration_complete

This operation, similarly to components, will be called by the **Application::start** operation [D&C]. This operation is necessary to activate the handshake between connector fragments at deployment time, after the configuration of all components and connector fragments.

7.5.2.1 get_ccm_home

This operation, similarly to components, returns a **CCMHome** reference to the home that manages this component.

7.5.2.2 remove

This operation, similarly to components, is used to delete a fragment. Application failures during remove raise the **RemoveFailure** exception.

7.5.3 Interface KeylessCCMHome

This interface merely implements a bootstrapping facility to create connector fragment instances. The same interface is used by the components. As for components, an entry point allowing the container to create a connector's home instance is defined and is of type:

```

extern "C" { Components::HomeExecutorBase_ptr (*)(); } // in C++
Components_HomeExecutorBase* (*)(); // in C

```


7.5.3.1 create_component

This operation is called to create the connector fragment during deployment.

7.5.4 Interface HomeConfiguration

7.5.4.1 set_configuration_values

As for components, this operation establishes an attribute configuration for the target fragment object, as an instance of **Components::ConfigValues**. Factory operations on the home of fragment will apply this configurator to newly-created instances.

7.5.5 Equivalent IDL (w.r.t Equivalent IDL section in CCM)

The connector extension does not need to specify equivalent IDL interfaces deduced from ports since only generic operations inherited from **Navigation** and **Receptacles** are mandatory in the lightweight CCM profile, which is addressed by this specification.

If necessary for a connector, the rules to obtain equivalent interfaces are the same as for a component.

7.5.6 Connector Implementation Interfaces

This sub clause explains how can be implemented connector fragments.

The CCM provides a standardized Component Implementation Framework (CIF) defining the programming model for constructing component implementations.

The connector implementation (implementation of several fragments) is specific to the semantic it defines; it can be dependant of the underlying platform and is connector provider specific. For that reason, there is no need to standardize a counter part of the CIF for connectors.

As explained before, the implementation of a fragment inherits the **Components::CCMObject** interface and shall implements the specified operations of **Navigation, Receptacles**. This is mandatory to provide a connector that can be deployed and configured with lwCCM deployment framework (compliant to D&C specification, '[D&C]'). This implementation corresponds to a classical implementation of IDL interfaces using the standard language mapping.

As for components, the skeleton of connector fragments can be partly generated taking into account the transformation rules defined in the connector definition. This is fully the responsibility of the connector framework provider.

7.6 Connector Deployment and Configuration

This sub clause introduces all the modifications to the OMG D&C specification considered as necessary in order to deal with the packaging and deployment of connectors. The extensions are to be added in the PSM for CCM part of D&C reference in Clause 15, 'Deployment PSM for CCM.'

Note – This sub clause and the following are based on the D&C specification '[D&C]' - all conventions defined in this specification are applicable:

- In particular, standard attributes (e.g., label) have the semantics defined in the D&C specification.

- All classes that are not explicitly defined in this document are taken from the specification.
- In the UML diagrams, when no multiplicity is indicated on an association end, the multiplicity is one.

Note that extended ports and connectors (considered as CCM extensions) defined in the previous clauses, as an extension of IDL3 have no impact on the D&C PIM; it will only impact the PSM for CCM level [CCM].

7.6.1 Integration of Connectors in D&C

As said before, the objective of this specification is to provide new interaction modes for component-based applications. To achieve this goal, it shall not add complexity for the assembly of components. For this reason, the connectors in a component-based application design shall be seen as an interaction element that links 2 components and not as a new functional entity that will imply multiplication of connections at assembly level. Nevertheless it implies some modifications to the D&C Component Data Model at assembly level where connections will include connector information.

On the contrary, at Execution Data Model level, since the deployment plan aims to be [automatically] produced at planning phase by tools, and since it is a *flattened assembly*, the connector defined in the connection elements of the assembly will appear as artifacts that have to be deployed by the deployment tools. This implies that the fragment instances (artifacts) are described in the deployment plan with their configuration values; and that connections between components and their related fragment are basic connections (facet / receptacles).

7.6.2 Component Data Model

A connector is an entity very similar to a component. It is packaged, deployed, and owns implementation(s), as well as interfaces, etc. Therefore, it would not have been relevant to define a completely new data model for connectors.

7.6.2.1 Connector Description

Connectors may be packaged in the same way components are, thus most of the elements defined in the component data model are relevant in the case of connectors. However, component packages and connector packages shall be distinguishable; therefore a **ConnectorPackageDescription** class is defined.

Like a component package, a connector package owns descriptive information (interface description) and one or more implementation(s).

As far as the interface description is concerned, no differences exists between components and connectors, thus the **ComponentInterfaceDescription** class is used for connectors as well and is extended at PSM level to integrate extended port specificities.

In the following, all diagrams of the component data model impacted by the above statements are displayed.

The following figure displays the additions⁶ that are to be made to the Component Data Model at PSM for CCM level. Actually, two classes are added: **ConnectorPackageDescription**⁷ and **ConnectorImplementationDescription**.

6. Note that this diagram displays only the two classes that have to be added, along with their relations to already existing classes. All the classes originally defined in the specification are, even if not represented here, left intact, as well as their relations.

7. The association between **ComponentUsageDescription** and **ConnectorPackageConfiguration** is in mutual exclusion with those defined in the initial component data model between **ComponentUsageDescription** and **ComponentPackageDescription**.

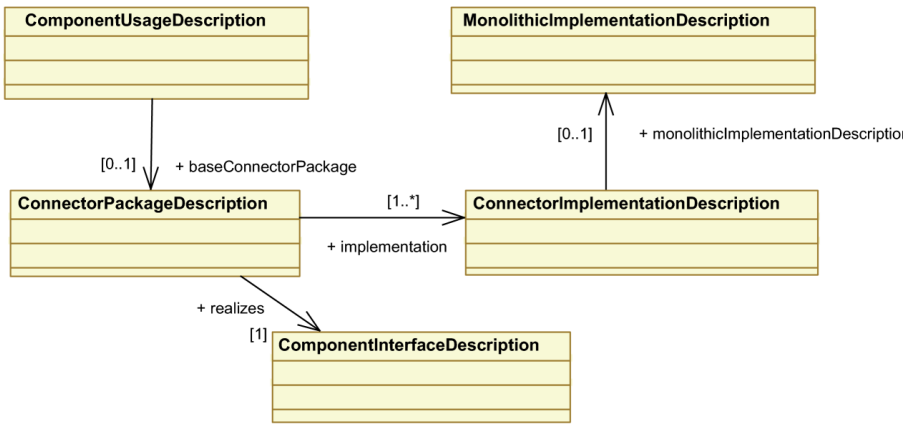


Figure 7.6 - Revised component data model overview

The two following figures give a detailed description of **ConnectorPackageDescription** and **ConnectorImplementationDescription** classes.

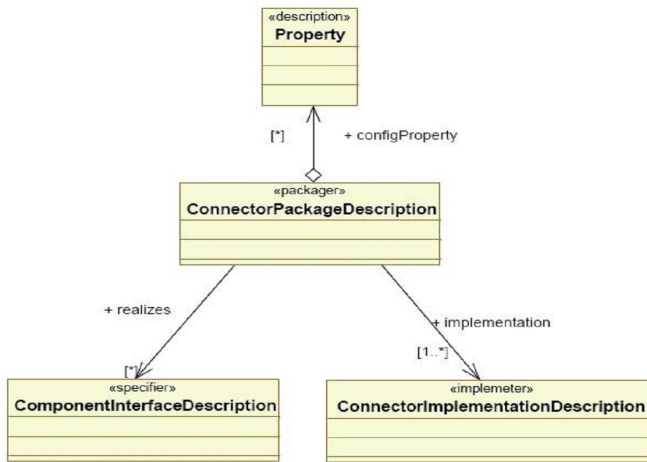


Figure 7.7 - ConnectorPackageDescription Class

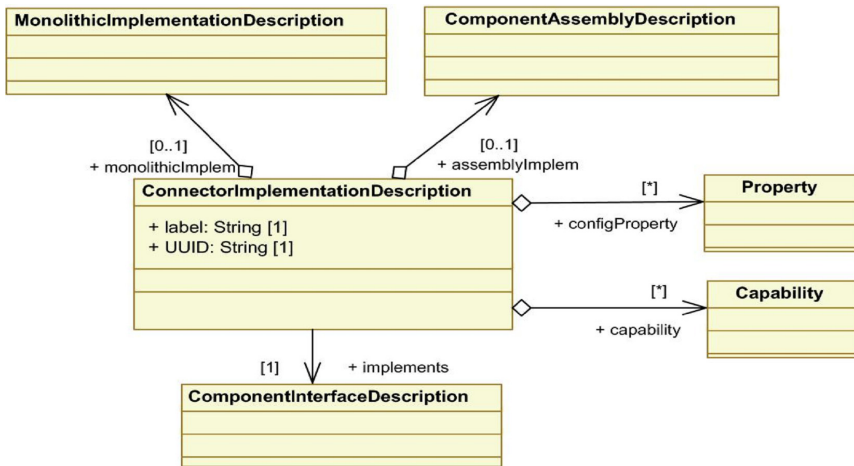


Figure 7.8 - ConnectorImplementationDescription Class

7.6.2.2 ConnectorPackageDescription

A **ConnectorPackageDescription** describes multiple alternative implementations of the same connector. It references the interface description for the connector and contains a number of configuration properties to configure the running connector (which may override implementation-defined properties and which may be overridden by a **PackageConfiguration**). These configuration properties enable the packager to define default values for a connector's properties regardless of which implementation for that component is chosen at deployment (planning) time.

7.6.2.3 ConnectorImplementationDescription

A **ConnectorImplementationDescription** describes a specific implementation of a connector. This implementation can be only monolithic. The **ConnectorImplementationDescription** may contain configuration properties that are used to configure each connector fragments instance ("default values"). Implementations may be tagged with user-defined capabilities. Administrators can then select among implementations using selection requirements in a **PackageConfiguration**.

The **ComponentInterfaceDescription** class is used to describe components and connectors. This description contains information on the ports of components and connectors.

ComponentPortDescription class shall be extended to support the extended ports. As explained in previous sections, extended ports are defined at least by their specific types (**specificType** member of **ComponentPortDescription**) but they can also be parameterized by several template parameters. The class is therefore extended with a **templateParam** member. The kind of port shall also support extended ports and inverse ports. The **CCMComponentPortKind** enumeration is extended with two values: **ExtendedPort**, **ExtendedMirrorPort**.

7.6.2.4 ComponentInterfaceDescription

The added **ComponentPortDescription::templateParam** (String [0..*]) contains all the template parameters types needed to parameterize the port (if extended). This member is null if the port is simple or if it is an extended port without template. If **templateParam** contains values, the **kind** attribute shall be **ExtendedPort** or **ExtendedMirrorPort**.

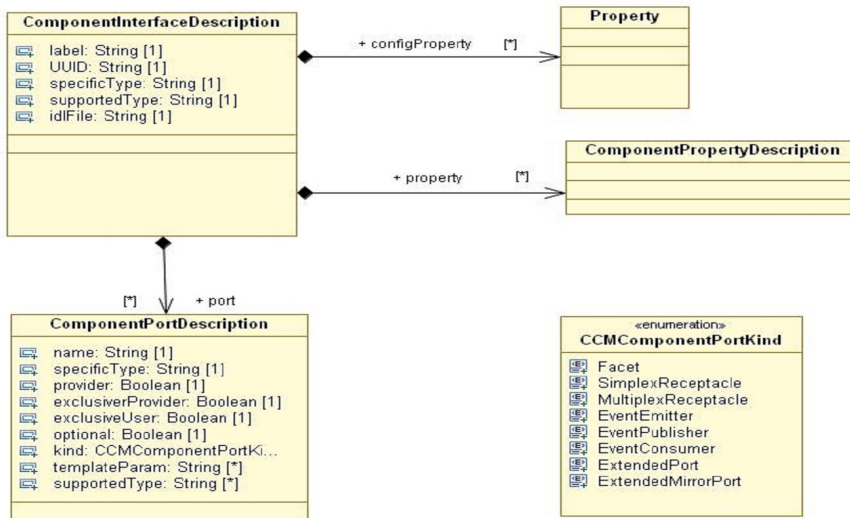


Figure 7.9 - Support for Extended Ports

The connector description will be part of .ccd files.

7.6.2.5 Component Assembly with Connectors

At D&C assembly level, using a connector shall result in a set of connections between components and shall not appear as a new component instance in the assembly.

In the D&C specification, the **ComponentAssemblyDescription** element contains information about subcomponent instances (**SubcomponentInstantiationDescription**), connections among ports (**AssemblyConnectionDescription**), and about the mapping of the assembly's properties (i.e., of the component that the assembly is implementing) to properties of its subcomponents.

Connectors at assembly level are considered as particular connections. It means that the **AssemblyConnectionDescription** need to be extended to support connector descriptions. At PSM for CCM level, the following extensions are specified:

The **AssemblyConnectionDescription** can be realized by a connector. Therefore, this class provides a direct association with **ConnectorPackageDescription**. The principle is similar to **SubComponentInstantiationDescription** that (by inheritance of **ComponentUsageDescription**) references **ComponentPackageDescription** itself referencing the connector definition (**ComponentInterfaceDescription**).

The association is 0..1. If the cardinality is 0 the connection is a basic CCM connection (facet □ receptacle and events), if it is 1 the connection is implemented by a connector.

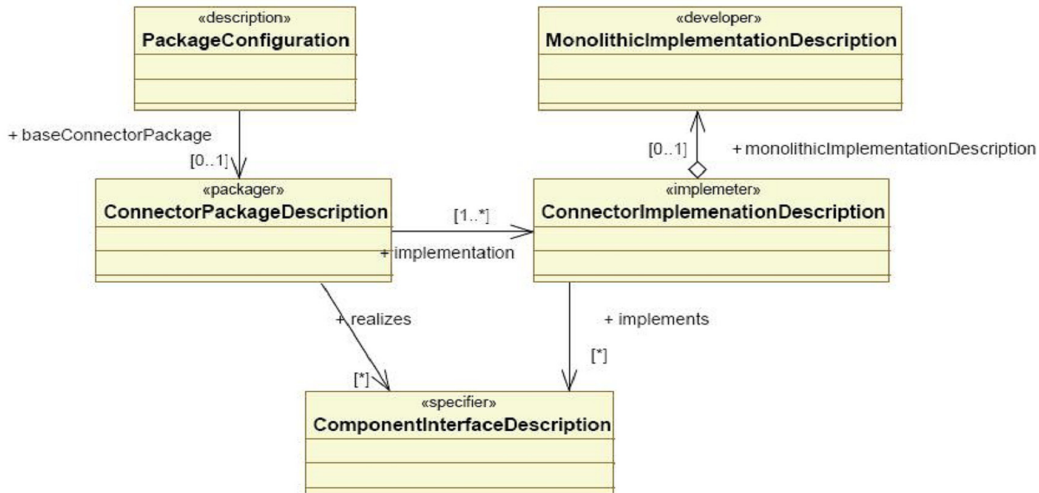


Figure 7.10 - AssemblyConnectionDescription Extension

7.6.3 Execution Data Model

At the Execution Data Model level, the deployment plan is produced from the assembly description and corresponds to the assembly fully flattened. All executable artifacts are part of the deployment plan.

Each connector fragment is described as artifacts (**ArtefactDeploymentDescription**), implementations (**MonolithicDeploymentDescription**), instances (**InstanceDeploymentDescriptions**). By definition a connector implementation is the result of its fragment implementations. Each fragment can be deployed on different a target, that's why at the execution model level, fragments are manipulated while at Component Data Model, connectors are manipulated.

The transformation from assembly level (designed in a modeling tool) to the resulting deployment plan can be easily generated since all parameterized typed are resolved when the assembly tool connects components with a connector. The resulting simple ports of components and connector fragments will be the endpoints to connect at deployment time.

This way of proceeding implies a very small impact on the existing deployment frameworks since they will deal with the same entities (artifacts, implementations, instances and connections). Nevertheless few extensions are necessary to allow the instantiation of connector fragments and their configuration.

7.6.3.1 Compliance with Entry Points

This refers to sub clause 10.6.1 of D&C [D&C] regarding the CCM entry points.

If the instance to be deployed is a connector, then the name of the execution parameter shall be “**home factory.**”

The parameter is of type String, and its name is the name of an entry point that has no parameters and that returns a pointer of type **Connectors::HomeExecutorComponents:: HomeExecutorBase**.

Thanks to this object, the deployment tool will call the **create_component()** operation on the **KeylessCCMHome** to instantiate a connector fragment.

7.6.4 Connector Configuration

The configuration of port type at assembly level produces the needed configuration values at deployment time.

All fragments of a given connector are in relation and have to be configured consistently. In some cases, this could require them to share configuration information that cannot be set statically. This dynamic initialization, if required, is connector implementation-specific and thus not specified. However it has to be completed by the end of the 'configuration_complete' phase of CCM deployment.

To configure the fragments the **Components::HomeConfiguration** IDL interface could be used. The method **set_configuration_values** is called in order to set the needed **ConfigValues** for the connector.

If two fragments need to exchange some configuration data (e.g., CORBA reference) the naming service could be used.

The configuration data are specified in the Component Deployment Plan file. Following is an example that shows how to configure fragments at deployment plan level.

```
<!-- ***** -->
<!-- ***** INSTANCES ***** -->
<!-- ***** -->
<!-- Instance for fragment_instance_1 -->
<instance id="fragment_instance_1">
  <name>fragment_instance_1</name>
  <node>node1</node>
  <implementation ref="fragment_impl_1"/>
  <configProperty>
    <name>mcast_addr</name>
    <type>string</type>
    <value>224.1.1.1</value>
  </configProperty>
  <configProperty>
    <name>mcast_port</name>
    <type>unsigned short</type>
    <value>31337</value>
  </configProperty>
  <configProperty>
    <name>msg_size</name>
    <type>unsigned long</type>
    <value>50</value>
  </configProperty>
</instance>
```

7.6.5 CCM Meta-model Extension to support Generic Interactions

In this sub clause, the basic concepts of the component model are summarized, based on the CCM meta-model [UML_CCM]. Central to it is the notion of Component definition (**ComponentDef**). It corresponds to the specification of a new component type, providing, using, and supporting possibly several interfaces, as well as consuming, emitting, or publishing event types. For configuration issues, attributes can be used as part of component definitions.

This part is based on the specification [UML_CCM] and it with new meta classes.

As an extension, the specification introduces the **ExtendedPortDef** as well as **ExtendedPortType** in the meta-model in order to allow definition of custom types of ports, the primary motivation being the reification at component level of interactions, which will be supported by the Connector concept.

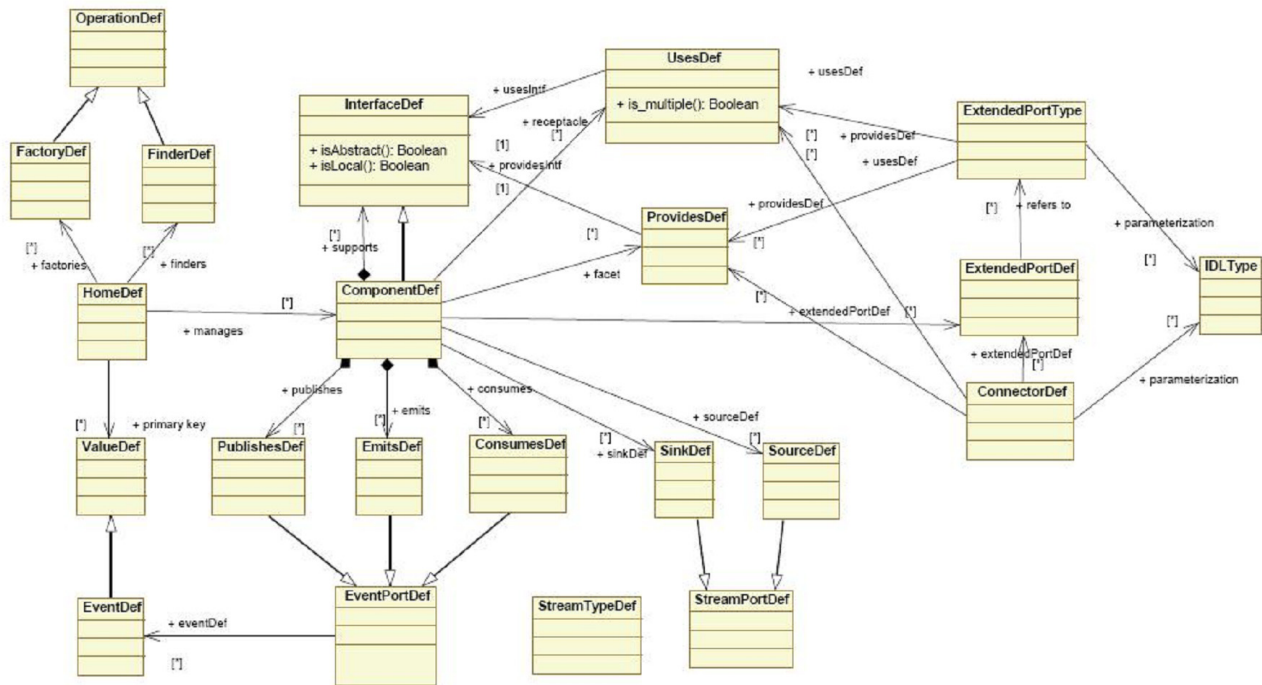


Figure 7.11 - Component Meta-Model – With Extended Ports and Connectors

ExtendedPortTypes are aggregation of zero or several provisions or needs of interfaces.

A matching relation for **ExtendedPortType** is defined as follows: two such types are compatible when they present one by one compatible **UsesDef** and **ProvidesDef**.

Finally, **ConnectorDef** is a new construct of the Component meta-model allowing modeling of connectors.

All those extensions are represented in Figure 7.11.

8 OMG CIDL Syntax and Semantics

8.1 Overview

This clause describes OMG Component Implementation Definition Language (CIDL) semantics and gives the syntax for OMG CIDL grammatical constructs.

The OMG Component Implementation Definition Language (CIDL) is a language used to describe the structure and state of component implementations. Component-enable ORB products generate implementation skeletons from CIDL definitions. Component builders extend these skeletons to create complete implementations.

OMG CIDL obeys the same lexical rules as OMG Persistent State Definition Language (PSDL) and OMG IDL, although new keywords are introduced to support concepts specific to component implementation descriptions.

The description of OMG CIDL's lexical conventions is presented in 8.2, Lexical Conventions. A description of OMG IDL preprocessing is presented in CORBA Core, IDL Syntax and Semantics, Preprocessing sub clause. The scope rules for identifiers in an OMG IDL specification are described in CORBA Core, IDL Syntax and Semantics, CORBA Module sub clause.

The OMG CIDL grammar is an extension of a combination of the OMG PSDL and OMG IDL grammars, with new constructs to define component implementations. OMG CIDL is a declarative language. The grammar is presented in “OMG CIDL Grammar” on page 82.

A source file containing specifications written in OMG CIDL must have a “.cdl” extension.

The description of OMG CIDL grammar uses the same syntax notation that is used to describe OMG IDL in Part 1 CORBA Interfaces, *IDL Syntax and Semantics* clause. For reference, Table 8.1 lists the symbols used in this format and their meaning.

Table 8.1 - IDL EBNF

Symbol	Meaning
::=	Is defined to be
	Alternatively
<text>	Nonterminal
“text”	Literal
*	The preceding syntactic unit can be repeated zero or more times
+	The preceding syntactic unit can be repeated one or more times
{ }	The enclosed syntactic units are grouped as a single syntactic unit
[]	The enclosed syntactic unit is optional—may occur zero or one time

8.2 Lexical Conventions

This sub clause presents the lexical conventions of OMG CIDL. In general OMG CIDL uses the same lexical conventions as OMG PSDL and OMG IDL. It does use additional keywords as described below.

8.2.1 Keywords

The identifiers listed in Table 8.2 are reserved for use as keywords in CIDL, and may not be used otherwise in CIDL, unless escaped with a leading underscore. These are in addition to the ones defined by PSDL and IDL, which may also not be used otherwise in CIDL, unless escaped with a leading underscore.

Table 8.2 - Keywords

bindsTo	delegatesTo	facet	proxy	session
composition	entity	implements	segment	storagehome
	executor	process	service	storedOn

8.3 OMG CIDL Grammar

The CIDL grammar is a combination of the PSDL and IDL grammars plus the following productions:

- (1) `<cidl_specification>` ::= `<import>*` `<cidl_definition>+`
- (2) `<cidl_definition>` ::= `<type_dcl>` `“,”`
| `<const_dcl>` `“,”`
| `<except_dcl>` `“,”`
| `<interface>` `“,”`
| `<cidl_module>` `“,”`
| `<storagehome>` `“,”`
| `<abstract_storagehome>` `“,”`
| `<storagetype>` `“,”`
| `<abstract_storagetype>` `“,”`
| `<value>` `“,”`
| `<type_id_dcl>` `“,”`
| `<type_prefix_dcl>` `“,”`
| `<event>` `“,”`
| `<component>` `“,”`
| `<home_dcl>` `“,”`
| `<composition>` `“,”`
- (3) `<cidl_module>` ::= `“module” <identifier>`
`“{” <cidl_definition>+ “}”`
- (4) `<composition>` ::= `“composition” <category> <identifier>`
`“{” <composition_body> “}”`
- (5) `<category>` ::= `“entity”`
| `“ ”`
| `“service”`
| `“session”`
- (6) `<composition_body>` ::= `<home_executor_def>`
`[<proxy_home_def>]`
- (7) `<home_executor_def>` ::= `“home” “executor” <identifier> “{”`
`<home_executor_body> “}” “;”`
- (8) `<home_executor_body>` ::= `<home_impl_dcl>`
`[<abstract_storagehome_binding>]`
`[<home_persistence_dcl>]`
`<executor_def>`

```

[<abstract_storagehome_delegation_spec>]
[ <executor_delegation_spec> ]
[ <abstract_spec> ]
(9) <home_impl_dcl> ::= "implements" <home_type_name> ";"
(10) <home_type_name> ::= <scoped_name>
(11) <abstract_storagehome_binding> ::= "bindsTo" <identifier> ";"
(12) <home_persistence_dcl> ::= "storedOn" <abstract_storagehome_name> ";"
(13) <executor_def> ::= "manages" <identifier>
    [ <executor_body> ] ";"
(14) <executor_body> ::= "{" <executor_member>+ "}"
(15) <executor_member> ::= <segment_def>
    | <feature_delegation_spec>
(16) <segment_def> ::= "segment" <identifier>
    "{" <segment_member>+ "}"
(17) <segment_member> ::= <segment_persistence_dcl> ";"
    | <facet_dcl> ";"
(18) <segment_persistence_dcl> ::= "storedOn" <abstract_storagehome_name> ";"
(19) <facet_dcl> ::= "provides" "facet" <identifier>
    { ";" <identifier> }*
(20) <feature_delegation_spec> ::= "delegatesTo" "storage"
    <feature_delegation_list>
(21) <feature_delegation_list> ::= "(" <feature_delegation> { ";" <feature_delegation> }* ")"
(22) <feature_delegation> ::= <feature_name> ":"
    <storage_member_name>
(23) <feature_name> ::= <identifier>
(24) <storage_member_name> ::= <identifier>
(25) <abstract_storagehome_delegation_spec> ::= "delegatesTo" "abstract"
    "storagehome" <delegation_list> ";"
(26) <executor_delegation_spec> ::= "delegatesTo" "executor"
    <delegation_list> ";"
(27) <delegation_list> ::= "(" <delegation> { ";" <delegation> }* ")"
(28) <delegation> ::= <operation_name> [ ":" <operation_name> ]
(29) <operation_name> ::= <identifier>
(30) <abstract_spec> ::= "abstract" <operation_list> ";"
(31) <operation_list> ::= "(" <operation_name>
    { ";" <operation_name> }* ")"
(32) <proxy_home_def> ::= "proxy" "home" <identifier>
    "{" <proxy_home_member>+ "}" ";"
(33) <proxy_home_member> ::= <home_delegation_spec> ";"
    | <abstract_spec>
(34) <home_delegation_spec> ::= "delegatesTo" "home" <delegation_list>

```

8.4 OMG CIDL Specification

A CIDL specification is like a PSDL and IDL specification that could also contain composition definitions. The syntax is:

```
(1) <cidl_specification> ::= <import>* <cidl_definition>+
(2) <cidl_definition> ::= <type_dcl> “;”
    | <const_dcl> “;”
    | <except_dcl> “;”
    | <interface> “;”
    | <cidl_module> “;”
    | <storagehome> “;”
    | <abstract_storagehome> “;”
    | <storagetype> “;”
    | <abstract_storagetype> “;”
    | <value> “;”
    | <type_id_dcl> “;”
    | <type_prefix_dcl> “;”
    | <event> “;”
    | <component> “;”
    | <home_dcl> “;”
    | <composition> “;”
(3) <cidl_module> ::= “module” <identifier>
    “{” <cidl_definition>+ “}”
```

8.5 Composition Definition

The syntax for composition definitions is as follows:

```
(4) <composition> ::= “composition” <category> <identifier> “{”
    <composition_body> “}”
(5) <category> ::= “entity”
    | “process”
    | “service”
    | “session”
(6) <composition_body> ::= <home_executor_def> [ <proxy_home_def> ]
```

A composition definition is a named scope that contains elements that constitute the composition. The elements of a composition definition are as follows:

- The keyword **composition**.
- The specification of the life cycle category, one of the keywords **service**, **session**, **process**, or **entity**. Subsequent definitions and declarations in the composition must be consistent with the declared category, as defined in “Life Cycle Category and Constraints” on page 85.
- An identifier that names the composition in the enclosing module scope.
- The composition body.

The composition body consists of the following elements:

- a mandatory home executor definition, and
- an optional proxy home definition.

8.5.1 Life Cycle Category and Constraints

Certain composition configurations are only valid for certain life cycle categories. The Container Programming Model sub clause describes the life cycle-related constraints from the perspective of the container. These constraints map onto corresponding constraints in component and composition definitions. The following lists define the CIDL constructs that are either mandatory or invalid for the designated life cycle category.

Note that these constraints supersede the conditionality of constructs based on CIDL syntax. If a construct is described below as mandatory for the category in question, it is mandatory regardless of its syntactic properties. All of the constructs described as invalid for a particular category are, of necessity, syntactically optional.

Table 8.3 - Constraints for service and session components

Service and Session	Mandatory	None
	Invalid	Abstract storage home bound to home executor: <abstract_storagehome_binding> in home executor body.
		Component home implemented by home executor specifies a primary key.
		Component home implemented by home executor specifies explicit finder operations.
		Segmented executor: <segment_def> in executor body.

Table 8.4 - Constraints for process components

Process	Mandatory	None
	Invalid	Component home implemented by home executor specifies a primary key.

Table 8.5 - Constraints for entity components

Entity	Mandatory	None
		Component home implemented by home executor specifies a primary key.
	Invalid	none

8.6 Home Executor Definition

The syntax for a home executor definition is as follows:

- (7) **<home_executor_def> ::= "home" "executor" <identifier>
"{" <home_executor_body> "}" ","**
- (8) **<home_executor_body> ::= <home_impl_dcl>
[<abstract_storagehome_binding>]
[<home_persistence_dcl>]
<executor_def>
[<abstract_storagehome_delegation_spec>]
[<executor_delegation_spec>]
[<abstract_spec>]**

A home executor definition consists of the following elements:

- the keywords **home** and **executor**,
- an identifier that names the home executor definition within the scope of the composition, and
- a home executor body.

The home executor body consists of the following elements:

- A home implementation declaration.
- An optional abstract storage home binding, specifying the storage home upon which the components managed by the home are stored.
- An optional home persistence declaration, identifying an abstract storage home upon which the state of the home executor itself is to be stored.
- An executor definition, describing the component executor managed by the home executor.
- An optional delegation specification describing the mapping of home operations to storage home operations.
- An optional delegation specification describing the mapping of home factory operations to the operations on the component executor.
- An optional abstract specification, declaring operations on the home executor that are to be left unimplemented, overriding default generated implementations.

The *<identifier>* in the header of the home executor definition is used as the basis for the name of the skeleton artifact generated by the CIF. The specific forms of the executors are defined in language mappings. The general requirements for language mappings of home executors are defined in this sub clause.

8.7 Home Implementation Declaration

The syntax of a home implementation declaration is as follows:

(9) `<home_impl_dcl> ::= "implements" <home_type_name> ";"`
(10) `<home_type_name> ::= <scoped_name>`

The home implementation declaration consists of the following elements:

- the keyword **implements**, and
- a scoped name denoting a component home imported from IDL.

The home implementation declaration specifies the component home that is to be implemented by the home executor being defined. The generated skeleton must support the home equivalent interface, as defined in "Equivalent Interfaces" on page 34. Implementations of orthodox home operations are generated if the life cycle category of the composition is either **entity** or **process** and the home executor specifies an abstract storage home binding, or if the life cycle category of the executor is either **session** or **service**.

The detailed semantics of generated implementations are described in this sub clause.

8.8 Storage Home Binding

The syntax for a storage home binding is as follows:

(11) <abstract_storagehome_binding> ::= “bindsTo” <abstract_storagehome_name> “;”

An abstract storage home binding declaration consists of the following elements:

- the keyword **bindsTo**, and
- an abstract storage home name.

8.9 Home Persistence Declaration

The syntax for a home persistence declaration is as follows:

(12) <home_persistence_dcl> ::= “storedOn” <abstract_storagehome_name> “;”

A home persistence declaration consists of the following elements:

- the keyword **storedOn**, and
- an abstract storage home name.

A home persistence declaration establishes that the home executor is itself persistent, and that its persistent state is managed by the container. The abstract storage type of the specified abstract storage home constitutes the state of the component home. The specific responsibilities of generated home executors related to home persistence are described in this sub clause.

8.10 Executor Definition

The syntax for an executor definition is as follows:

(13) <executor_def> ::= “manages” <identifier>
[<executor_body>] “;”

(14) <executor_body> ::= “{” <executor_member>+ “}”

(15) <executor_member> ::= <segment_def>
| <feature_delegation_spec>

An executor definition has the following elements:

- the keyword **manages**,
- an identifier that names the component executor being defined, and
- an executor body, containing one or more members enclosed in braces.

An executor member is either a *segment definition* or a *feature delegation specification*, as defined below.

The identifier in the executor definition forms the basis of the name of the programming artifact generated as the executor skeleton. The details of executor structure and responsibilities are defined in “Home Executor Definition” on page 85, and in CIDL language mappings.

8.11 Segment Definition

The syntax for a segment definition is as follows:

- (16) `<segment_def> ::= "segment" <identifier>
"{" <segment_member>+ "}"`
- (17) `<segment_member> ::= <segment_persistence_dcl> ";"
| <facet_dcl> ";"`

A segment definition consists of the following elements:

- the keyword **segment**,
- an identifier that names the segment in the scope of the executor definition, and
- one or more segment members enclosed in braces.

A segment member is either a *segment persistence declaration*, or a *facet declaration*, as described below.

If a segment definition occurs in an executor definition, the corresponding executor is said to be a segmented executor. If no segment definition occurs in an executor definition, the executor is said to be monolithic.

A separate skeleton is generated by the CIF for each segment of a segmented executor. Segments are independently activated. Each segment is assigned a segment identifier, which as a numeric value of type short, by the CIF implementation. The segment identifier is interpreted internally by the generated implementation during activation. Segment identifiers are also used in component identities, as described in “Component Identities” on page 164. There is no canonical mechanism for assigning segment identifier values (other than the component segment), as the values of segment identifiers does not affect portability or interoperability.

All executors have a distinguished segment, the component segment, that supports the component facet (i.e., the facet supporting the component equivalent interface). The segment identifier value of the component segment is always zero. If a component does not explicitly declare segments, the monolithic executor is still considered in some contexts to be the component segment executor.

The details of segment structure and implementation responsibilities are described in this sub clause.

8.12 Segment Persistence Declaration

The syntax for a segment persistence declaration is as follows:

- (18) `<segment_persistence_dcl> ::= "storedOn" <abstract_storagehome_name> ";"`

A segment persistence declaration has the following elements:

- the keyword **storedOn**, and
- an abstract storage home name.

A segment persistence declaration specifies the abstract storage home upon which the state of the segment will be stored. The abstract storage type of the storage home constitutes the state of the segment.

The detailed structure of segments, and implementation responsibilities with respect to segment persistence are described in this sub clause.

8.13 Facet Declaration

The syntax for a facet declaration is as follows:

(19) **<facet_dcl>** ::= “provides” “facet” <identifier>
{ “,” <identifier> }*

A facet declaration has the following elements:

- The keywords **provides** and **facet**.
- One or more identifiers separated by commas, where each identifier denotes a facet defined by the component type implemented by the composition (i.e., the component type managed by the home that is implemented by the home executor defined in the composition).

A facet declaration associates one or more component facets with the segment. The generated segment executor will provide the specified facets. A facet name may only appear in a single segment definition. Facets that are not explicitly declared in a segment definition are provided by the component segment.

The detailed structure of segments, and implementation responsibilities with respect to providing facets are described in this sub clause.

8.14 Feature Delegation Specification

The syntax for a feature delegation specification is as follows:

(20) **<feature_delegation_spec>** ::= “delegatesTo” “storage”
 <feature_delegation_list>
(21) **<feature_delegation_list>** ::= “(” **<feature_delegation>** { “,” **<feature_delegation>** }* “)”
(22) **<feature_delegation>** ::= **<feature_name>** “:”
 <storage_member_name>
(23) **<feature_name>** ::= **<identifier>**
(24) **<storage_member_name>** ::= **<identifier>**

A feature delegation specification has the following elements:

- the keywords **delegatesTo**, **abstract**, and **storagetype**, and
- a list of feature delegation specifications, enclosed in parentheses and separated by commas.

A feature delegation specification consists of the following elements:

- An identifier that denotes a stateful feature of the component implemented by the composition,
- A colon,
- An identifier that denotes a member of the abstract storage type of the abstract storage home specified in the abstract storage home binding in the home executor definition.

A feature delegation specification defines an association between a stateful feature of the component being implemented and a member of the abstract storage type that incarnates the component (or the component segment). The component executor skeleton generated by the CIF will provide implementations of feature management operations that store the feature’s state in the specified storage member. Stateful features include attributes, receptacles, and event sources.

The following constraints regarding feature delegation must be observed:

- Feature delegation specifications may only occur in an executor definition when the home executor specified an abstract storage home binding.
- The type of the storage member specified in a feature delegation must be compatible with the type of the feature. Compatibility, for the purposes of feature delegation is defined in Table 8.6.

Table 8.6 - Type compatibility for feature delegation purposes

Feature	Storage member type
attribute	Must be identical to feature for all types except object reference and valuetype; for object reference and valuetype storage member must be of identical type or base type (direct or indirect).
receptacle (simplex)	Must be identical to feature type or base interface (direct or indirect) of feature type.
receptacle (multiplex)	Sequence of type compatible with receptacle type as defined above.
emitter event source	Must be identical to feature type or base interface (direct or indirect) of feature type.
publisher event source	long*

* The persistent state maintained internally by the component is the **ChannelId** of the notification channel created by the container.

8.15 Abstract Storage Home Delegation Specification

The syntax for a storage home delegation specification is as follows:

```
(25) <abstract_storagehome_delegation_spec> ::= "delegatesTo" "abstract"
      "storagehome" <delegation_list> " ,"
(26)   <delegation_list> ::= "(" <delegation> { " ," <delegation> }* ")"
(27)   <delegation> ::= <operation_name> [ ":" <operation_name> ]
(28)   <operation_name> ::= <identifier>
```

An abstract storage home delegation specification has the following elements:

- The keywords **delegatesTo**, **abstract**, and **storagehome**.
- A list of delegation specifications enclosed in parentheses and separated by commas.

A delegation specification has the following elements:

- An identifier that denotes an operation on the home equivalent interface supported by the home executor.
- An optional delegation target, consisting of a colon, followed by identifier that denotes an operation on the abstract storage home to which the home is bound (i.e., the abstract storage home specified in the abstract storage home binding).

An abstract storage home delegation specification associates an operation on the home interface with an operation on the abstract storage home interface. The CIF shall generate an implementation of the specified home operation that delegates to the specified abstract storage home operation.

If the optional delegation target is omitted, the home operation is assumed to be delegated to an operation on the abstract storage home with the same name. If no such operation exists on the abstract storage home, the specification is not legal.

The signature of the abstract storage home operation must be compatible with the abstract storage home. Signature compatibility, from the perspective of abstract storage home delegation, has the following definition:

- If the home operation is an explicit **factory** operation, the abstract storage home operation must be an explicit **factory** operation.
- If the home operation is not a factory, the return type of the home operation must be identical to the return type of the abstract storage home operation, except when the return type is an object reference type or a value type. If the return type of the home operation is an object reference type or a value type, the return type of the storage home operation must be identical to, or more derived than, the return type of the home operation.
- For each exception explicitly raised by the storage home operation, an identical exception must appear in the **raises** clause of the home operation. The inverse is not true—the home operation may raise exceptions not raised by the abstract storage home operation.
- The number of parameters in the parameter lists of the home operation and the abstract storage home operation must be equal. Each parameter in the abstract storage home operation must be compatible with the parameter in the same position in the signature of the home operation, where compatibility is defined as follows:
 - If the parameter in the home operation is neither an object reference type nor a value type, the type of the corresponding parameter in the abstract storage home operation must be identical.
 - If the parameter type in the home operation is an object reference and the parameter is an **in** parameter, the corresponding parameter in the abstract storage home operation must be identical to, or a base type (direct or indirect) of, the parameter in the home operation.
 - If the parameter type in the home operation is an object reference and the parameter is an **out** parameter, the corresponding parameter in the abstract storage home operation must be identical to, or more derived than, the parameter in the home operation.
 - If the parameter type in the home operation is an object reference and the parameter is an **inout** parameter, the corresponding parameter in the abstract storage home operation must be identical to the parameter in the home operation.

The following additional constraints and rules apply to abstract storage home delegation:

- An operation on the home interface may delegate to at most one operation on the abstract storage home interface.
- An operation on the abstract storage home interface may be the target of at most one delegation from the home interface.
- Implicitly defined operations on the home (i.e., orthodox operations) delegate by default to cognate operations on the abstract storage home, as described by “Orthodox operations” on page 40. These default delegations may be overridden by explicit delegations. If an operation on the abstract storage home that is normally the default target of a delegation appears as the target of an explicit delegation, then the home operation that normally would have delegated to that target by default shall have no generated implementation (unless one is explicitly defined).

The detailed semantics and implementation responsibilities of delegated abstract storage home operations are described in this sub clause.

8.16 Executor Delegation Specification

The syntax for an executor delegation specification has the following form:

```
(29) <executor_delegation_spec> ::= “delegatesTo” “executor”
      <delegation_list> “;”
```

An executor delegation specification consists of the following elements:

- the keywords **delegatesTo** and **executor**, and
- a delegation list, identical structurally to the delegation list of the abstract storage home delegation specification.

An executor delegation specification defines an operation on the component executor, to which the specified home operation will be delegated. The following constraints apply to executor delegation specifications:

- Only factory operations may be delegated to the executor, including explicitly declared factories and implicit create operations.
- If no delegation target is explicitly specified, the operation defined on the executor shall have the same name as the delegating home operation.
- The signature of the defined operation on the executor shall be identical to the signature of the home operation, with the exception that the return type of the executor operation shall be void if the home does not specify a primary key, or the return type shall be the type of the primary key if the home specifies a primary key.

The CIF shall generate an implementation of the home operation that delegates to the defined operation on the executor. The detailed semantics and implementation responsibilities are described in this sub clause.

8.17 Abstract Spec Declaration

The syntax for an abstract spec has the following form:

```
(30)      <abstract_spec> ::= "abstract" <operation_list> ";"
(31)      <operation_list> ::= "(" <operation_name>
          { "," <operation_name> }* ")"
```

8.18 Proxy Home Declaration

The syntax for a proxy home declaration has the following form:

```
(32)      <proxy_home_def> ::= "proxy" "home" <identifier>
          "{" <proxy_home_member>+ "}" ";"
(33) <proxy_home_member> ::= <home_delegation_spec> ";"
          | <abstract_spec>
(34) <home_delegation_spec> ::= "delegatesTo" "home" <delegation_list>
```

9 CCM Implementation Framework

9.1 Introduction

The Component Implementation Framework (CIF) defines the programming model for constructing component implementations. Implementations of components and component homes are described in CIDL. See the “OMG CIDL Syntax and Semantics” clause for the definition and syntax. The CIF uses CIDL descriptions to generate programming skeletons that automate many of the basic behaviors of components, including navigation, identity inquiries, activation, state management, lifecycle management, and so on.

9.2 Component Implementation Framework (CIF) Architecture

As a programming abstraction, the CIF is designed to be compatible with the existing POA framework, but also to insulate programmers from its complexity. In particular, the CIF can be implemented using the existing POA framework, but it does not directly expose any elements of that framework.

9.2.1 Component Implementation Definition Language (CIDL)

The focal point of the CIF is Component Implementation Definition Language (CIDL), a declarative language for describing the structure and state of component implementations. Component-enabled ORB products generate implementation skeletons from CIDL definitions. Component builders extend these skeletons to create complete implementations.

9.2.2 Component persistence and behavior

CIDL is a superset of the Persistent State Definition Language, defined in the Persistent State Service specification (<http://www.omg.org/technology/documents/formal/persistent.htm>).

A CIDL implementation definition may optionally associate an abstract storage type with the component implementation, such that the abstract storage type defines the form of the internal state encapsulated by the component. When a component implementation declares an associated abstract storage type in this manner, the CIF and the run-time container environment cooperate to manage the persistence of the component state automatically.

This sub clause addresses the elements of the CIF that pertain to the implementation of a component’s behavior.

9.2.3 Implementing a CORBA Component

The remainder of this sub clause provides an overview of the concepts involved in building component implementations. It is intended to provide a high-level description that will serve as a framework for understanding the more formal descriptions that follow in subsequent sub clauses. While the information in this sub clause is normative (with the exception of italicized, indented rationale), it is not intended to be a complete or precise specification of the CIF, or all of the possible design options from which a component implementor may choose.

9.2.4 Behavioral elements: Executors

We coin the term *executor* to indicate the programming artifact that supplies the behavior of a component or a component home. In general, the terms *executor* or *component executor* refer to the artifact that implements the component type, and the term *home executor* refers to the artifact that implements the component home.

We chose to use the word executor rather than servant to avoid confusion with POA servants. POA servants, while conceptually similar to executors, are significantly different in detail, and map to different types in programming languages. Executor is pronounced with the accent on the second syllable (e.g., -ZEK-yoo-tor).

We have tried to avoid terminology that is specific to object-oriented programming languages, such as class, base class, derive, and so on, in an attempt to be precise and acknowledge that the CIF framework may be mapped to procedural programming languages. Hence, we typically use the word artifact or programming artifact to denote what may conveniently be thought of as a class, and likewise, the term skeleton to denote a generated abstract base class that is extended to form a complete implementation class. We hope this is not overly distracting to the reader.

9.2.5 Unit of implementation: Composition

An implementation of a component comprises a potentially complex set of artifacts that must exhibit specific relationships and behaviors in order to provide a proper implementation. The CIDL description of a component implementation is actually a description of this aggregate entity, of which the component itself may be a relatively small part. In order to enable more concise discussion, we coin the term *composition* to denote both the set of artifacts that constitute the unit of component implementation, and the definition itself. **composition** is the CIDL meta-type that corresponds to an implementation definition.

A composition definition specifies the following elements.

Component home

A composition definition specifies a component home type, imported from IDL. The specification of a component home implicitly identifies the component type for which the composition provides an implementation (i.e., the component type managed by the home, as specified in the IDL home definition).

Abstract Storage home binding

A composition optionally specifies an abstract storage home to which the component home is bound. The specification of an abstract storage home binding implicitly identifies the abstract storage type that incarnates the component. The relationship **between a home and the component it manages to isomorphic to the relationship** between an abstract storage home and the abstract storage type it manages. When a home binds to an abstract storage home, the component managed by the home is implicitly bound to the abstract storage type of this abstract storage home.

Home executor

A composition definition specifies a home executor definition. The name of the home executor definition is used as the name of the programming artifact (e.g., the class) generated by the CIF as the skeleton for the home executor. The contents of the home executor definition describe the relationships between the home executor and other elements of the composition, determining the characteristics of the generated home executor skeleton.

Component executor

A composition specifies an executor definition. The name of the executor definition is used as the name of the programming artifact generated by the CIF as the skeleton of the component executor. The body of the executor definition optionally specifies executor *segments*, which are physical partitions of the executor, encapsulating independent state and capable of being independently activated. Segments are described in “Segmented executors” on page 115. The executor body may also specify a mapping, or *delegation*, of certain component features (e.g., attributes) to storage members.

Delegation specification

A composition may optionally provide a specification of home operation delegation. This specification maps operations defined on the component home to isomorphic operations on either the abstract storage home or the component executor. The CIF uses this description to generate implementations of operations on the home executor, and to generate operation declarations on the component executor.

Proxy home

A composition may optionally specify a proxy home. The CIF supports the ability to define proxy home implementations, which are not required to be collocated with the container that executes the component implementation managed by the home. In some configurations, proxy homes can provide implementations of home operations without contacting the container that executes the actual home and component implementation. Support for proxy homes is intended to increase the scalability of the CORBA Component Model. The use of proxy homes is completely transparent to component clients and, to a great extent, transparent to component implementations. Proxy home behavior is described in “Proxy home delegation” on page 123.

9.2.6 Composition structure

A composition binds all of the previously-described elements together, and requires that the relationships between the bound entities define a consistent whole.

Note that a component home type necessarily implies a component type; that is, the managed component type specified in the home definition. Likewise, an abstract storage home implies an abstract storage type. It is unnecessary, therefore, for a composition to explicitly specify a component type or an abstract storage type. They are implicitly determined by the specification of a home and abstract storage home.

It may seem odd that the center of focus for compositions is the home rather than the component, but this works out to be reasonably intuitive in practice. The home is the primary point of contact for a client, and the home's interface and behavior have a major influence on the interaction between the client and the component.

A composition definition specifies a name that identifies the composition within the enclosing module scope, and which constitutes the name of a scope within which the contents of the composition are contained. The essential parts of a composition definition are the following:

- The name of the composition.
- The life cycle category of the component implementation, either **service**, **session**, **process**, or **entity**, as defined in “Component Categories” on page 137.
- The home type being implemented (which implicitly identifies the component type being implemented).
- The name of the home executor to be generated.
- The name of the component executor skeleton to be generated.

A composition definition has the following essential form:

```

composition <category> <composition_name> {
  home executor <home_executor_name> {
    implements <home_type> ;
    manages <executor_name>;
  };
};

```

where <composition_name> is the name of the composition, <category> identifies the life cycle category supported by the composition, <home_executor_name> is the name assigned to the generated home executor skeleton, <home_type> is the name of a component home type imported from IDL, and <executor_name> is the name assigned to the generated component executor skeleton.

This is a schematic representation of the minimal form of a composition, which specifies no state management. The structure of the composition specified by this schematic is illustrated in Figure 9.1. Note that the component type itself is not explicitly specified. It is unambiguously implied by the specification of the home type, as is the relationship between the executor and the component (i.e., that the executor *implements* the component).

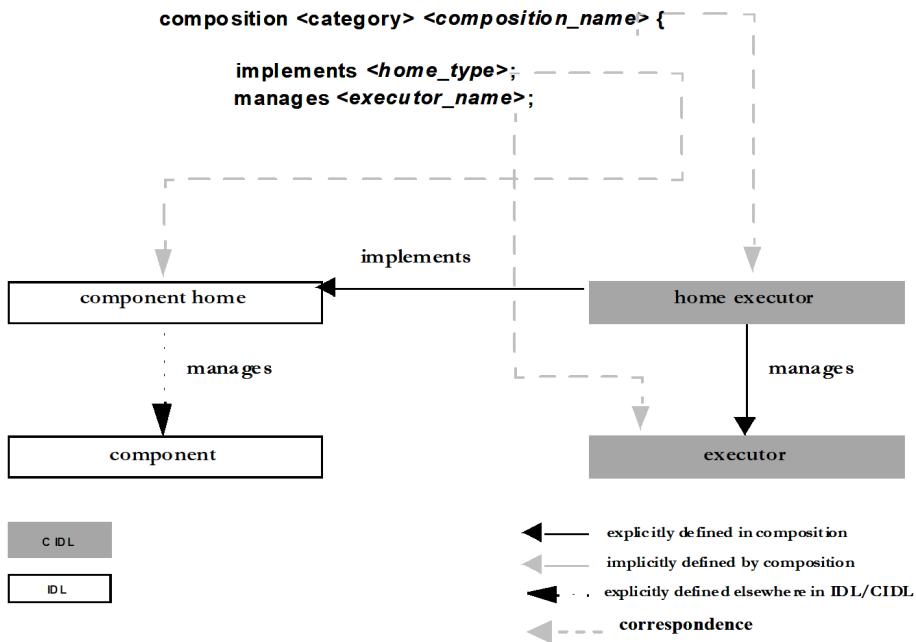


Figure 9.1- Minimal composition structure and relationships

General disclaimer and abdication of responsibility with regards to programming examples:

Before presenting programming examples, it should be noted that all examples are non-normative illustrations. In particular, the implementations provided in the examples of code that is to be generated by the CIF are merely schematic representations of the intended behaviors; they are by no means indicative of the actual content of a real implementation (e.g., they generally don't include exception handling, testing for validity, etc.).

Although the grammar for CIDL has not been presented yet, a simple example will help illustrate the concepts described in the previous sub clauses. Assume the following IDL component and home definitions:


```

-----
// Example 1
//
// USER-SPECIFIED IDL
//
module LooneyToons {
    interface Bird {
        void fly (in long how_long);
    };
    interface Cat {
        void eat (in Bird lunch);
    };
    component Toon {
        provides Bird tweety;
        provides Cat sylvester;
    };
home ToonTown manages Toon {};
};
-----

```

The following example shows a minimal CIDL definition that describes an implementation binding for those IDL definitions:

```

-----
// Example 1
//
// USER-SPECIFIED CIDL
//
import ::LooneyToons;

module MerryMelodies {

    // this is the composition:

    composition session ToonImpl {
        home executor ToonTownImpl {
            implements LooneyToons::ToonTown;
            manages ToonSessionImpl;
        };
    };
};
-----

```

In this example, *ToonImpl* is the name of the composition. It defines the name of the generated home executor to be *ToonTownImpl*, which implemented the *ToonTown* home interface imported from IDL. The home executor definition also specified the name of the component executor, *ToonSessionImpl*, which is managed by the home executor. Note that the component type (*Toon*) is not explicitly named—it is implied by the specification of the home *ToonTown*, which is known to manage the component type *Toon*. Thus, the declaration “**manages ToonSessionImpl**” implicitly defines the component executor *ToonSessionImpl* to be the implementation of the component type *Toon*.

This CIDL specification would cause the generation of the following artifacts:

- The skeleton for the component executor `ToonSessionImpl`
- The complete implementation of the home executor `ToonTownImpl`

We provide the following brief sketches of generated implementation skeletons in Java to help illustrate the programming model for component implementations.

Java **<interface>Operations** interfaces for all of the IDL interfaces are generated, precisely as currently specified by the current Java IDL language mapping:

```
-----  
// Example 1  
//  
// GENERATED FROM IDL SPECIFICATION:  
//  
package LooneyToons;  
  
import org.omg.Components.*;  
  
public interface BirdOperations {  
    public void fly (long how_long);  
}  
  
public interface CatOperations {  
    void eat(LooneyToons.Bird lunch);  
}  
  
public interface ToonOperations  
    extends CCMObjectOperations {  
    LooneyToons.Bird provide_tweety();  
    LooneyToons.Cat provide_sylvester();  
}  
  
public interface ToonTownExplicitOperations extends CCMHomeOperations { }  
  
public interface ToonTownImplicitOperations extends KeylessCCMHomeOperations  
{  
    Toon create();  
}  
  
public interface ToonTownOperations extends  
    ToonTownExplicitOperations,  
    ToonTownImplicitOperations { }  
-----
```

The `ToonImpl` executor skeleton class has the following form:

```
-----  
// Example 1  
//  
// GENERATED FROM CIDL SPECIFICATION:  
//  
package MerryMelodies;  
import LooneyToons;  
import org.omg.Components.*;  
  
abstract public class ToonSessionImpl  
    implements ToonOperations, SessionComponent,
```

```

ExecutorSegmentBase
{
    // Generated implementations of operations
    // inherited from SessionComponent and
    // ExecutorSegmentBase are omitted here.
    //

    protected ToonSessionImpl() {
        // generated implementation ...
    }

    // The following operations must be implemented
    // by the component developer:

    abstract public BirdOperations
        _get_facet_tweety();
    abstract public CatOperations
        _get_facet_sylvester();
}

```

*The generated executor abstract base class **ToonSessionImpl** implements all of the operations inherited by **ToonOperations**, including operations on **CCMObject** and its base interfaces. It also implements all of the operations inherited through **SessionComponent**, which are internal operations invoked by the container and the internals of the home implementation to manage executor instance lifecycle.*

*A complete implementation of the home executor **ToonTownImpl** is generated from the CIDL specification:*

```

-----
// Example 1
//
// GENERATED FROM CIDL SPECIFICATION:
//
package MerryMelodies;
import LooneyToons;
import org.omg.Components.*;

public class ToonTownImpl
implements LooneyToons, ToonTownOperations,
HomeExecutorBase, CCMHome
{
    // Implementations of operations inherited
    // from ExecutorBase and CCMHome
    // are omitted here.
    //
    // ToonHomeImpl also provides implementations
    // of operations inherited from the component
    // home interface ToonTown

    CCMObject create_component()
    {
        return create();
    }

    void remove_component(CCMObject comp)
    {

```

```

    }

    Toon create()
    {
    }
    // and so on...
}

```

The user-provided executor implementation must supply the following:

- Implementations of the operations `_get_tweety` and `_get_sylvester`, which must return implementations of the

BirdOperations and *CatOperations* interfaces

- said implementations of the behaviors of the facets **tweety** and **sylvester**, respectively

The following example shows one possible implementation strategy:

```

-----
// Example 1
//
// PROVIDED BY COMPONENT PROGRAMMER:
//
import LooneyToons.*;
import MerryMelodies.*;

public class myToonImpl extends ToonSessionImpl
implements BirdOperations, CatOperations {

    protected long timeFlown;
    protected Bird lastBirdEaten;

    public myToonImpl() {
        super();
        timeFlown = 0;
        lastBirdEaten = nil;
    }

    public void fly (long how_long) {
        timeFlown += how_long;
    }

    public void eat (Bird lunch) {
        lastBirdEaten = lunch;
    }

    public BirdOperations _get_facet_tweety() {
        return (BirdOperations) this;
    }

    public CatOperations _get_facet_sylvester() {
        return (CatOperations) this;
    }

}

```

This simple example implements all of the facets directly on the executor. This is not the only option; the programming objects that implement **BirdOperations** and **CatOperations** could be constructed separately and

managed by the executor class.

The final bit of implementation that the component programmer must provide is an extension of the home executor that acts as a component executor factory, by implementing the `create_executor_segment` method. This class must also provide an implementation of a static method called `create_home_executor` that returns a new instance of the home executor (as an `ExecutorSegmentBase`). This static method acts as an entry point for the entire composition.

```
-----  
// Example 1  
//  
// PROVIDED BY COMPONENT PROGRAMMER:  
//  
import LooneyToons.*;  
import MerryMelodies.*;  
  
public class myToonTownImpl extends ToonTownImpl  
{  
    protected myToonTownImpl() { super(); }  
  
    ExecutorSegmentBase  
    create_executor_segment (int segid) {  
        return new myToonImpl();  
    }  
  
    public static ExecutorSegmentBase  
    create_home_executor() {  
        return new myToonTownImpl();  
    }  
}  
-----
```

Note that these last two classes constitute the entirety of the code that must be supplied by the programmer. The implementations of operations for navigation, executor activation, object reference creation and management, and other mechanical functions are either generated or supplied by the container.

9.2.7 Compositions with Managed Storage

A composition definition may also contain a variety of optional specifications, most of which are related to state management. These include the following elements:

- An abstract storage home type to which the component home is bound (this implicitly identifies the abstract storage type to which the component itself is bound).
- The life cycle category of the composition must be either **entity** or **process** to support managed storage.

When state management is added to a composition definition, the definition takes the following general form, expressed as a schematic:

```
composition <category> <composition_name> {  
    home executor <home_executor_name> {  
        implements <home_type> ;  
        bindsTo <abstract_storage_home>;  
        manages <executor_name>;  
    };  
};
```

where the additional elements are as follows: *<abstract_storage_home>* denotes a particular abstract storage home provided by the catalog.

The structure of the resulting composition and the relationships between the elements is illustrated in Figure 9.2.

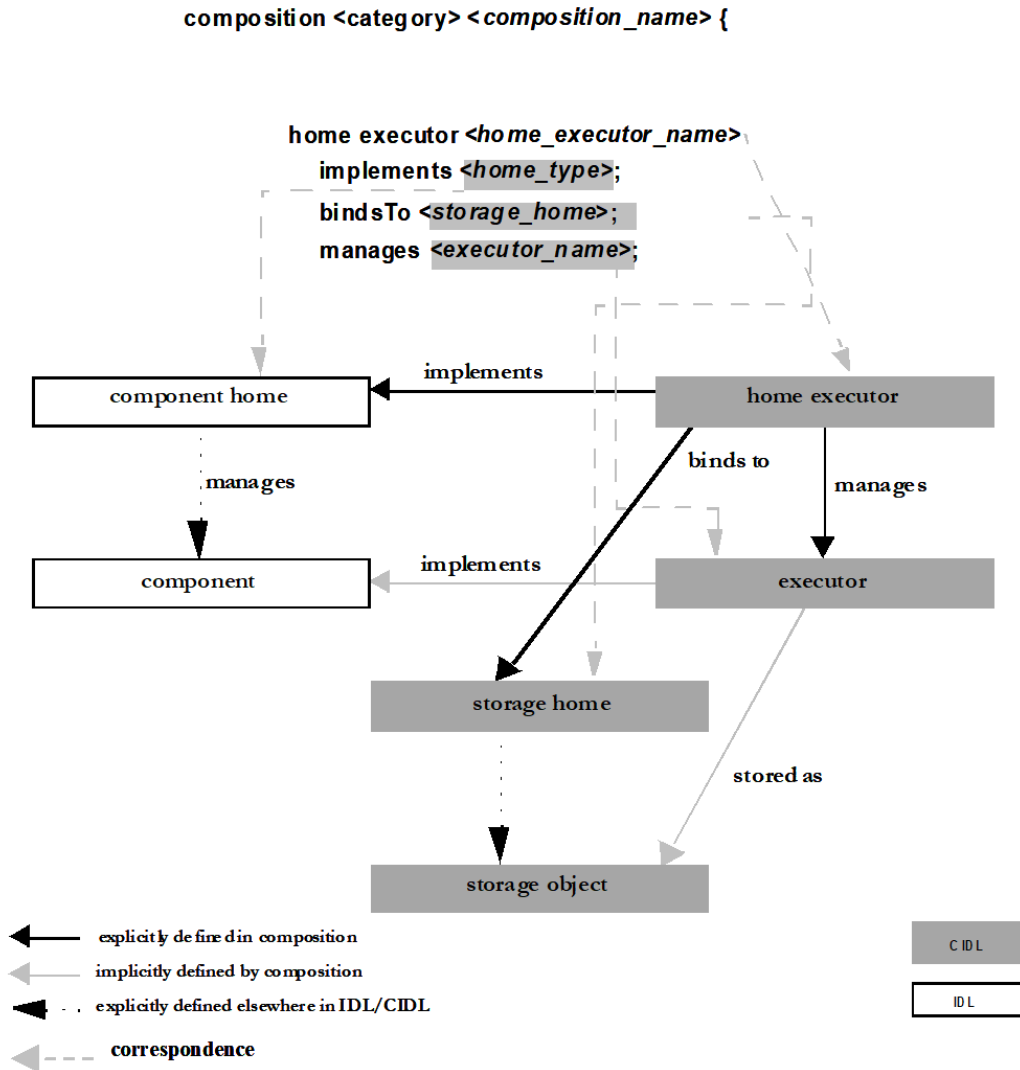


Figure 9.2- Structure of composition with managed storage

In many cases, it is expected that an abstract storage home will be intentionally designed to support a particular component home.

9.2.8 Relationship between Home Executor and Abstract Storage Home

When a composition specifies managed storage, the relationship between the home executor and the abstract storage home to which the home executor binds determines many of the characteristics of the implementation, including what implementation elements may be generated and how they will behave. This sub clause provides an overview of the basic concepts involved in home implementations and their relationships to abstract storage homes.

In general, operations on a home interface provide life cycle management. As described in “Homes” on page 34, when a home definition does not specify a primary key, the resulting equivalent home interface has the following operations:

- A generic **create_component** operation inherited from **KeylessCCMHome**,
- a **remove_component** operation inherited from **CCMHome**, and
- an implicitly-defined type-specific parameter-less **create** operation.

When a home definition specifies a primary key, the resulting equivalent home interface has the following operations:

- A **remove_component** operation inherited from **CCMHome**,
- an implicitly-defined type-specific **create** operation with a primary key parameter,
- an implicitly-defined type-specific **remove** operation with a primary key parameter, and
- an implicitly-defined type-specific **find_by_primary_key** operation.

9.2.8.1 Primary Key Binding

A component home can define its primary key as a valuetype with a number of public data members, whereas abstract storage home defines keys as lists of attributes. A composition can only bind a component home with a primary key to an abstract storage home that defines a key on a state member whose type is this valuetype. If there is more than one key satisfying this condition, the first key is used.

For example:

```
valuetype SSN {
    public string social_security_number;
};

abstract storagetype Person {
    readonly state SSN social_security_number;
    state string name;
    state string address;
};

abstract storagehome PersonStore of Person {
    key social_security_number;
};
```

A home with primary key SSN can be bound to **PersonStore**. The key **social_security_number** is called the matching key.

9.2.8.2 Implicit delegation of home operations

When a composition specifies managed storage, finder operations can be implemented in terms of finder operations on the abstract storage home to which the home executor is bound.

Table 9.1 - Delegation of finder operations to finder operations on the bound abstract storagehome

home operation	abstract storagehome operation
<i>component</i> find_by_primary_key (<i>key</i>)	<i>ref<X></i> find_ref_by_matching_key_name (<i>matching_key</i>)

- The **find_by_primary_key** operation uses the **find_ref_by_matching_key_name** operation on the abstract storagehome. The returned storage reference is used to create an object reference for the component and returned to the invoking client.
- Destruction operations delegate to **destroy_object** operations on the reference.

The validity of these implementation semantics are predicated on the following assumptions:

- The initial state of the storage object created by the storage home constitutes a valid initial state for the component.
- All of the persistent state of the component is defined on (or reachable from) the storage object whose PID is associated with the component instance.
- The executor is monolithic, not segmented. Home operations can also be delegated to abstract storage homes when the executor is segmented, but the process is slightly more complex, and is discussed in full in “Segmented executors” on page 115.

If these assumptions do not hold (in particular, either of the first two), the component implementor can provide custom implementations of one or more home operations to accommodate the implementation requirements.

The following example extends the previous example to illustrate managed storage and storage home delegation. The example highlights differences from the previous, and does not repeat elements that are identical:

```

-----
// Example 2
//
// USER-SPECIFIED IDL
//
module LooneyToons { // IDL

    ... identical to previous example, except for the addition of the
        primary key:

    valuetype EpisodeName : Components::PrimaryKeyBase {
        public string name;
    };
    home ToonTown manages Toon primarykey EpisodeName {
    };
};
-----

```


The CIDL now defines abstract storage types, abstract storage homes, and a catalog. The composition binds :

```
-----  
// Example 2  
//  
// USER-SPECIFIED CIDL  
//  
import ::LooneyToons;  
  
module MerryMelodies {  
  
    abstract storagetype ToonState {  
        state LooneyToons::EpisodeName episode_name;  
        state string name;  
        state unsigned long time_flowin;  
        state LooneyToons::Bird last_bird_eaten;  
    };  
  
    abstract storagehome ToonStateHome of ToonState  
    {  
        key episode_name;  
        factory create(episode_name);  
    };  
  
    catalog ToonCatalog {  
        provides ToonStateHome TSHome;  
    };  
  
    // this is the composition:  
  
    composition entity ToonImpl {  
        uses catalog { ToonCatalog store; };  
        home executor ToonTownImpl {  
            implements LooneyToons::ToonTown {  
                bindsTo ToonStateHome;  
                manages ToonEntityImpl;  
            };  
        };  
    };  
};  
-----
```

*In this example, the composition binds the component home **ToonTown** to the abstract storage home **ToonStateHome**, and thus, implicitly binds the component type **Toon** to the abstract storage type **ToonState**. Note that the primary key (if any) in the home must match a key in the abstract storage home. As will be seen later in the CIDL grammar specification, the keyword **entity** in the implementation binding declaration specifies a particular lifecycle model for the resulting implementation.*

This CIDL specification would cause the generation of the following programming objects:

- *The skeleton for the component executor **ToonEntityImpl***

- The implementation of the home executor `ToonTownImpl`
- The incarnation interface for the abstract storage type `ToonState`
- The interface for the abstract storage home `ToonStateHome`
- The interface for the catalog `ToonCatalog`.

Note that the complete implementation of the home executor may not be able to be generated in some cases, e.g., when no abstract storage type is declared or when user-defined operations with arbitrary signatures appear on the component home definition.

Note also that the implementations of the storage-related interfaces `ToonState` and `ToonStateHome` are not necessarily provided by the same product that generates the component implementation skeletons. The CIF is specifically designed to decouple the executor implementation from the storage implementation, so that these capabilities may be provided by different products. A component-enabled ORB product is only required to generate the programming interfaces for the abstract storage type and homes through which the executor implementation will interact with one or more storage mechanisms. The implementations of these interfaces may be supplied separately, perhaps deferred until run-time.

The interfaces generated from the IDL are identical, with the exception of the addition of the primary key:

```

-----
// Example 2
//
// GENERATED FROM IDL SPECIFICATION:
//
package LooneyToons;

import org.omg.Components.*;

... same as previous except for the following:

public interface ToonTownImplicitOperations {
    Toon create(LooneyToons.EpisodeName key)
        throws DuplicateKey, InvalidKey;
    Toon find_by_primary_key
        (LooneyToons.EpisodeName key)
        throws UnknownKey, InvalidKey;
    void remove(LooneyToons.EpisodeName key)
        throws UnknownKey, InvalidKey;
    LooneyToons.EpisodeName
        get_primary_key(Toon comp);
}

public interface ToonTownOperations extends
ToonTownExplicitOperations,
ToonTownExplicitOperations {}
-----

```

The abstract storage type `ToonState` results in the generation of the following incarnation interfaces:

```

-----
// Example 2
//
// GENERATED FROM CIDL SPECIFICATION:
//
package MerryMelodies;
import org.omg.CosPersistentState.*;

```

```

import LooneyToons.*;
public interface ToonState extends StorageObject {
    public string name();
    public void name (String val);
    public long time_flown();
    public void time_flown (long val);
    public Bird last_bird_eaten();
    public void last_bird_eaten (Bird val);
}

```

The storage home **ToonStateHome** results in the generation of the following interface:

```

-----
// Example 2
//
// GENERATED FROM CIDL SPECIFICATION:
//
// no explicit operations
public interface ToonStateHome
    extends StorageHomeBase {

    public ToonState
    find_by_episode_name (EpisodeName k);

    public ToonStateRef
    find_ref_by_episode_name (EpisodeName k);
}

```

The **ToonImpl** executor skeleton class has the following form:

```

-----
// Example 2
//
// GENERATED FROM CIDL SPECIFICATION:
//
package MerryMelodies;
import LooneyToons;

abstract public class ToonImpl
implements LooneyToons.ToonOperations,
ExecutorSegmentBase, PersistentComponent
{
    // Generated implementations of operations
    // inherited from CCMObject and
    // ExecutorSegmentBase and PersistentComponent
    // are omitted here.
    //
    // ToonImpl also provides implementations of
    // operations inherited from ToonState, that
    // delegate to a separate incarnation object:

    protected ToonStateIncarnation _state;
}

```

```

protected ToonImpl() { _state = null; }

public void set_incarnation (ToonState state) {
    _state = state;
}

// The following operations must be implemented
// by the component developer:

abstract public BirdOperations
    _get_facet_tweety();
abstract public CatOperations
    _get_facet_sylvester();
}

```

An implementation of the home executor ToonHomeImpl is generated from the CIDL specification:

```

-----
// Example 2
//
// GENERATED FROM CIDL SPECIFICATION:
//
package MerryMelodies;
import LooneyToons;

public class ToonTownImpl
implements LooneyToons.ToonTownOperations,
PersistentComponent, ExecutorSegmentBase
{
    // Implementations of operations inherited
    // from PersistentComponent and
    // ExecutorSegmentBase
    // are omitted here.
    //
    // ToonHomeImpl also provides implementations
    // of operations inherited from the component
    // home interface ToonTown, that delegate
    // designated operations on the storage home
    //
    // values set during initialization
    // and activation:
    protected Entity2Context _origin;
    protected ToonStateHome _storageHome;
    ...

    Toon create(EpisodeName key)
    {
        // create a storage object with the key

        ToonState new_state = _storageHome.create(key);

        // REVISIT - Bernard Normier 7/27/1999
        // don't know how to complete this method
    }
}

```

```

    }

    Toon find(EpisodeName key)
    {
        ToonStateRef ref =
            _storageHome.find_ref_by_episode_name(key);
        // create reference from ref
        // and return , same as above...
    }

    // and so on...
}
-----

```

The user-provided executor uses the storage accessors and mutators on the incarnation:

```

-----
// Example 2
//
// PROVIDED BY COMPONENT PROGRAMMER:
//
import LooneyToons.*;
import MerryMelodies.*;

public class myToonImpl extends ToonImpl
implements BirdOperations, CatOperations {

    public myToonImpl() { super(); }

    void fly (long how_long) {
        _state.timeFlown
            ( _state.timeFlown() + how_long);
    }
    void eat (Bird lunch) {
        _state.last_bird_eaten(lunch);
    }
    BirdOperations get_facet_tweety() {
        return (BirdOperations) this;
    }
    CatOperations get_facet_sylvester() {
        return (CatOperations) this;
    }
}

```

9.2.8.3 Explicit delegation of home operations

The previous sub clause described the default home executor implementation generated by the CIF. Default delegation can only be implemented for home operations or the home base interfaces, and implicitly-defined home operations (i.e., *orthodox* home operations). The syntax for home definitions permits explicitly-defined factory operations, finder operations, and operations with arbitrary signatures to be declared on the home. The CIF makes no assumptions about the semantics of these operations (i.e., the *heterodox* operations), other than the assumptions that factory operations return references for newly-created components, and finder operations return references for existing components that were indirectly identified by the parameters of the finder operation. Implementations of these operations are not generated by

default. CIDL does, however, allow the component implementor to specify explicitly how heterodox home operations are implemented. A CIDL home executor definition may optionally include the declarations illustrated in the following schematic CIDL example:

```
composition <category> <composition_name> {
  ...
  home executor <home_executor_name> {
    ... // assume storage management specified

    delegatesTo abstract storagehome (
      <home_op0> : <storage_home_op0>,
      <home_op1> : <storage_home_op1>, ...
    );
    delegatesTo executor(
      <home_op2> : <executor_op2>, ...
    );
    abstract(<home_op3>, <home_op4>, ...);
  };
};
```

9.2.8.3.1 Delegation to abstract storage home

The **delegatesTo abstract storagehome** declaration specifies a sequence of operation mappings, where each operation mapping specifies the name of an operation on the home, and the name of an operation on the storage home. The signatures of the operations must be compatible, as defined in “Home inheritance” on page 38. Based on this declaration, the CIF generates implementations of the home operations on the home executor that delegate to the specified operations on the abstract storage home.

9.2.8.3.2 Delegation to executor

The **delegatesTo executor** declaration specifies a sequence of operation mappings, similar to the **delegatesTo abstract storagehome** declaration. The name on the left hand side of the mapping (i.e., to the left of the colon, ‘:’) must denote an explicitly-declared factory operation on the home, or the identifier “**create**,” denoting the implicitly-declared factory operation. The right hand side of each mapping specifies the name of an abstract operation that will be generated on the component executor. The component implementor provides the implementation of the executor operation, and the CIF provides an implementation of the operation on the home executor that delegates to the executor.

The delegation of home operations to executors is problematic, since home operations (other than factories) have no target component. For this reason, only factory operations may be delegated to the component executor. The CIF implements this delegation by defining an additional facet on the component executor, called a *factory facet*. A factory facet is only exposed to the home executor; clients cannot navigate to the factory facet, and the factory facet is not exposed in component meta-data, or described in the **FacetDescription** values returned from **Navigation::get_all_facets**.

The implementation of the factory operation on the home executor that delegates to the component executor must first create an object reference that denotes the factory facet. The home operation then invokes the mapped factory operation on the object reference, causing the activation of the component and ensuring that the execution of the operation on the component occurs in a proper invocation context.

If the factory operation being delegated is any operation other than the orthodox **create** operation, and the home definition includes a primary key specification, the operation generated on the factory facet of the component executor returns a value of the specified primary key type. The delegating operation on the home executor associates the primary key value returned from the component executor with the storage object (i.e., the storage object's PID) created to incarnate the component instance.

The use of PID values to create object references obviates the need to have two versions of a create method on the executor, as is the case in EJB with create and postCreate methods. An appropriate calling context can be created before the factory operation is invoked on the executor.

These precise semantics of and requirements for factory operations delegated to the executor are described in detail in “Factory operations” on page 37.

9.2.8.3.3 Suppressing generated implementation

The **abstract** specification overrides the generation of implementations for orthodox home operations. The name of any explicitly-defined operation on the home may be specified in the operation list of the abstract declaration. The CIF will not implement the specified operations, instead leaving unimplemented abstract operation declarations (on whatever appropriate equivalent exists for the particular language mapping).

The following example extends the previous example to illustrate delegation of home operations to the abstract storage home and the executor. The example highlights differences from the previous, and does not repeat elements that are identical:

```
-----  
// Example 3  
//  
// USER-SPECIFIED IDL  
//  
module LooneyToons { // IDL  
  
    ... identical to previous example, except for the home:  
  
    home ToonTown manages Toon primarykey EpisodeName {  
        factory createToon(  
            in string name, in long num, in Bird bref);  
        void arbitrary_operation();  
    };  
};  
-----
```

The CIDL now defines abstract storage types, abstract storage homes, and a catalog. The composition binds:

```
-----  
// Example 3  
//  
// USER-SPECIFIED CIDL  
//  
import ::LooneyToons;  
  
module MerryMelodies {
```

... identical to the previous example, except for:

```
abstract storagehome ToonStateHome of ToonState
{
    key episode_name;
    factory create();
    void do_something();
};

composition entity ToonImpl {
    uses catalog { ToonCatalog store; };
    home executor ToonTownImpl {
        implements LooneyToons::ToonTown;
        bindsTo store.TSHome;
        manages ToonEntityImpl;
        delegatesTo abstract storagehome
            (arbitrary_operation : do_something);
        delegatesTo executor ( createToon : createToon );
    };
};
```

In this example, the `arbitrary_operation` on the home interface `ToonTown` is delegated to the storage home operation `do_something`. Note that the operations have identical signatures. The `createToon` factory operation is delegated to an operation of the same name on the executor. This delegation causes the implicit definition of a factory facet on the component with the following interface:

```
interface ToonImplFactoryFacet {
    EpisodeName createToon(
        in string name, in long num, in Bird bref);
};
```

This interface is not part of the public interface of the component; its use is restricted to the home executor. In fact, the IDL need not be generated. All of the code that uses the factory facet is either generated by the CIF, or derived from CIF-generated skeletons, so the CIF can simply generate language mappings for the interface without actually providing any IDL for it. Note also that only a subset of the normal language mapping artifacts are required, including (in the case of Java) the abstract Operations interface, the POA tie class to be used internally by the executor, and a local stub to allow the home executor to make a delegating invocation. There is no need to generate a remote stub, as the facet is never exposed outside of the container.

The abstract storage home `ToonStateHome` interface has the added `do_something` operation on the explicit interface:

```
// Example 3
//
// GENERATED FROM CIDL SPECIFICATION:
//
public interface ToonStateHome
```



```

extends StorageHomeBase {
    public void do_something();
    // ...
}

```

The `ToonImpl` executor skeleton class supports an additional facet (the factory facet), which is returned by the `_get_factory_facet` operation:

```

-----
// Example 3
//
// GENERATED FROM CIDL SPECIFICATION:
//
package MerryMelodies;
import LooneyToons;

abstract public class ToonImpl
implements LooneyToons.ToonOperations,
ExecutorSegmentBase, PersistentComponent {
... same as previous
    // The following operations must be implemented
    // by the component developer:

    abstract public ToonImplFactoryFacetOperations
        _get_factory_facet();
    abstract public BirdOperations
        _get_facet_tweety();
    abstract public CatOperations
        _get_facet_sylvester();
}

```

The CIF generates implementations of the delegated operations on the home executor:

```

-----
// Example 3
//
// GENERATED FROM CIDL SPECIFICATION:
//
package MerryMelodies;
import LooneyToons;

public class ToonTownImpl
implements LooneyToons.ToonTownOperations,
CCMHome, ExecutorSegmentBase

    // values set during initialization

```

```

// and activation:
protected ToonStateHome _storageHome;
protected Entity2Context _origin;

...

Toon createToon(
    String name, long num, Bird bref)
{
    ToonState new_state=
        _storageHome.create();
    // etc.
}

void arbitrary_operation() {
    _storageHome.do_something();
}

...

}

```

The user-provide executor must implement the factory facet and operation:

```

-----
// Example 3
//
// PROVIDED BY COMPONENT PROGRAMMER:
//
import LooneyToons.*;
import MerryMelodies.*;

public class myToonImpl extends ToonImpl
implements BirdOperations, CatOperations,
ToonImplFactoryFacetOperations{
...
...

    EpisodeName
    createToon(String name, long num, Bird bref) {
        // presumably, the main reason for doing
        // this kind of delegation is to initialize
        // state in the context of the component:
        how_long(num);
        last_bird_eaten(bref);
        EpisodeNameDefaultFactory _keyFactory
            = new EpisodeNameDefaultFactory();
    }
}

```

```

        return _keyFactory.create(name);
    }

    ToonImplFactoryFacetOperations
    _get_factory_facet() {
        return
            (ToonImplFactoryFacetOperations) this;
    }
    ...
}

```

9.2.9 Executor Definition

The home executor definition must include an executor definition. An executor definition specifies the following characteristics of the component executor:

- The name of the executor, which is used as the name of the generated executor skeleton.
- Optionally, one or more distinct segments, or physical partitions of the executor. Each segment encapsulates independent state and is capable of being independently activated. Each segment also provides at least one facet.
- Optionally, the generation of operation implementations that manage the state of stateful component features (i.e., receptacles, attributes, and event sources) as members of the component incarnation.
- A delegation declaration that describes a correspondence between stateful component features and members of the abstract storage type that incarnates the component. The CIF uses this declaration to generate implementations of the feature-specific operations (e.g., **connect_** and **disconnect_** operations for receptacles, accessors, and mutators for attributes) that store the state associated with each specified feature in the storage member indicated on the right hand side of the delegation.

9.2.9.1 Segmented executors

A component executor may be *monolithic* or *segmented*. A monolithic executor is, from the container's perspective, a single artifact. A segmented executor is a set of physically distinct artifacts. Each segment may have a separate abstract state declaration. Each segment must provide at least one facet defined on the component definition. The life cycle category of the composition must be **entity** or **process** if the executor specifies segmentation.

The primary purpose for defining segmented executors is to allow requests on a subset of the component's facets to be serviced without requiring the entire component to be activated. Segments are independently activated. When the container receives a request whose target is a facet of a segmented executor, the container activates only the segment that provides the required facet.

The following schematic CIDL implicitly defines the normative form for the declaration of a segmented executor:

```

composition <category> <composition_name> {
    ...
    home executor <home_executor_name> {
        ... // assume storage management specified
        ...
        manages <executor_name> {
            segment <segment_name0> {

```

```

        storedOn <abstract_storage_home>;
        provides ( <facet_name0> , <facet_name1> , ... );
    };
    segment <segment_name1> { ... };
    ...
};

```

The abstract storage home specified in the segment's **storedOn** declaration implicitly specifies the abstract storage type that incarnates the segment. The home executor will use this abstract storage home to create and manage instances of the segment state (i.e., incarnations). If the component home specifies a primary key, then all of the abstract storage homes associated with executor segments must specify a matching key. The facets specified in the segment's **provides** declaration are implemented on the segment.

A segmented executor has a distinguished segment associated with the component. The component segment is implicitly declared, and supplies all of the facets not provided by separate segments, as well as all other component features and supported interfaces.

Figure 9.3, and Figure 9.4, illustrate the structure of monolithic and segmented executors, and the relationships between facets, storage objects, and segments. These figures also illustrate the identity information that is embedded in component and facet object references. Component identity information is described in more detail in “Component Identity” on page 11.

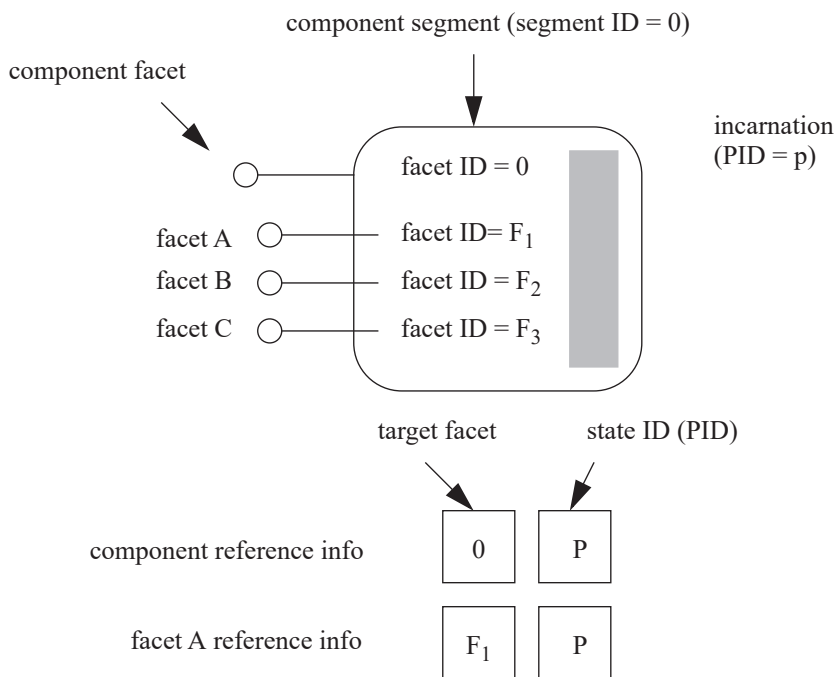


Figure 9.3- Monolithic executor and reference information structure

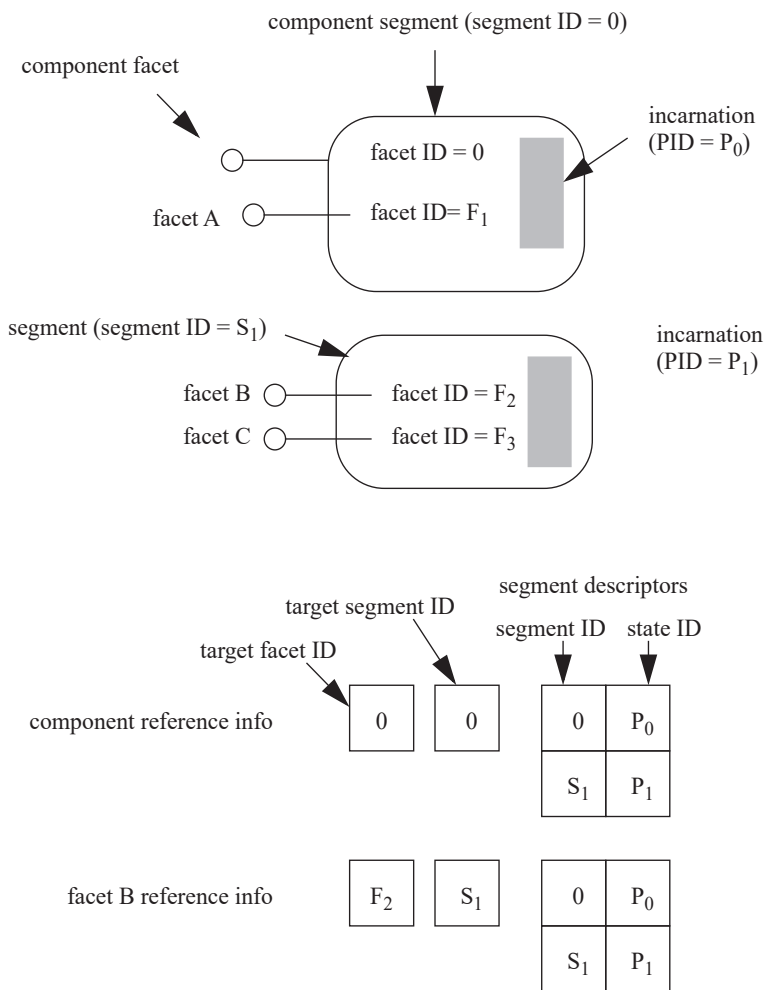


Figure 9.4- Segmented executor and reference information structure

The details of the structure and behavior of segments and requirements for their implementation are specified in “Segmented executors” on page 115.

The following example extends the previous example 2 to illustrate segmented executors. The example highlights differences from the previous, and does not repeat elements that are identical:

```

-----
//
// USER-SPECIFIED IDL
//
module LooneyToons { // IDL
    ... identical to previous example 2
};
-----

```

The CIDL now defines abstract storage types and abstract storage homes. The composition binds :

```
-----  
//  
// USER-SPECIFIED CIDL  
//  
import ::LooneyToons;  
  
module MerryMelodies {  
  
    ... identical to example 2 except for new storage, storage home  
        and executor definitions  
  
    abstract storagetype ToonState {  
        state LooneyToons::EpisodeName episode_name;  
        state string name;  
        state LooneyToons::Bird last_bird_eaten;  
    };  
  
    abstract storagehome ToonStateHome of ToonState {  
        key episode_name;  
    }; };  
  
    abstract storagetype BirdSegState {  
        state unsigned long time_flowin;  
    };  
  
    abstract storagehome BirdSegStateHome of BirdSegState {  
        key episode_name;  
    };  
  
    composition entity ToonImpl {  
        home executor ToonTownImpl {  
            implements LooneyToons::ToonTown {  
                bindsTo ToonStateHome;  
                manages ToonEntityImpl {  
                    segment BirdSegment {  
                        storedOn BirdSegStateHome;  
                        provides (tweety);  
                    };  
                };  
            };  
        };  
    };  
};  
-----
```

The storage home *BirdSegStateHome* is bound to the segment *BirdSegment*, which implicitly binds the segment executor for *BirdSegment* to the abstract storage type *BirdSegState*. This segment provides the facet *tweety*, leaving the remaining facet (*sylvester*) on the component segment.

The mappings of the CIDL abstract storage types and abstract storage homes are not presented, as they are not affected by the segmentation.

The generated component executor base class *ToonImpl* is also not presented, as the changes are trivial. The facet accessor *_get_facet_tweety* is no longer present on the component executor. There are other internal changes that are not visible to the component implementor. The executor for the new *BirdSegment* has the following form:

```
-----  
// Example 4  
//  
// GENERATED FROM CIDL SPECIFICATION:  
//  
package MerryMelodies;  
import LooneyToons;  
  
abstract public class BirdSegment  
implements ExecutorSegmentBase,  
PersistentComponent  
{  
    // Generated implementations of operations  
    // inherited from CCMObject and  
    // ExecutorSegmentBase and PersistentComponent  
    // are omitted here.  
    //  
  
    protected BirdSegState _state;  
  
    protected BirdSegment() { _state = null; }  
  
    public void set_incarnation (  
        BirdSegState state) {  
        _state = state;  
    }  
  
    // The following operations must be implemented  
    // by the component developer:  
  
    abstract public BirdOperations  
        _get_facet_tweety();  
}  
-----
```

Note that the *BirdSegment* executor does not implement any IDL interface directly, as does the component segment. It is remotely accessible only through a provided facet.

A generated implementation of the home executor *ToonHomeImpl* is considerably different from the previous example 2. The create method must create references for all of the segments and construct a *ComponentId* with the proper information::

```
-----  
//  
// GENERATED FROM CIDL SPECIFICATION:  
//  
package MerryMelodies;  
import LooneyToons;
```

```

public class ToonTownImpl
implements LooneyToons.ToonTownOperations,
CCMHome, ExecutorSegmentBase
{
    // Implementations of operations inherited
    // from CCMHome and ExecutorSegmentBase
    // are omitted here.
    //
    // ToonHomeImpl also provides implementations
    // of operations inherited from the component
    // home interface ToonTown, that delegate
    // designated operations on the storage home
    //
    // values set during initialization
    // and activation:
    protected Entity2Context _origin;
    protected ToonStateHome _toonStorageHome;
    protected BirdSegStateHome _birdStorageHome;
    ...

    Toon create(EpisodeName key)
    {
        ToonState new_toon =
            _toonStorageHome.create(key);
        // etc.
    }
}

```

There are now two segment executors to implement:

```

//
// PROVIDED BY COMPONENT PROGRAMMER:
//
import LooneyToons.*;
import MerryMelodies.*;

public class myToonImpl extends ToonImpl
implements CatOperations {

    public myToonImpl() { super(); }

    void fly (long how_long) {
        _state.timeFlown
            ( _state.timeFlown() + how_long);
    }
    void eat (Bird lunch) {
        _state.last_bird_eaten(lunch);
    }
    BirdOperations get_facet_tweety() {
        return (BirdOperations) this;
    }
    CatOperations get_facet_sylvester() {
        return (CatOperations) this;
    }
}

```



```

}

public class myBirdSegImpl extends BirdSegment
implements BirdOperations {

    public myBirdSegImpl() { super(); }

    void fly (long how_long) {
        _state.timeFlown
        ( _state.timeFlown() + how_long);
    }

    BirdOperations get_facet_tweety() {
        return (BirdOperations) this;
    }
}

```

The programmer must also supply a different implementation of the create_executor_segment operation on the home executor; that uses the segment ID value to determine which executor to create.

```

-----
// Example 4
//
// PROVIDED BY COMPONENT PROGRAMMER:
//
import LooneyToons.*;
import MerryMelodies.*;

public class myToonTownImpl extends ToonTownImpl
{
    protected myToonTownImpl() { super(); }

    ExecutorSegmentBase
    create_executor_segment (int segid) {

        // case discriminator values are constants
        // generated on the executor segment classes
        switch (segid) {
            case ToonImpl._segment_id_value :
                return new myToonImpl();
            case BirdSegment._segment_id_value :
                return new myBirdSegImpl();
            default
                ... raise an exception
        }
    }
    ...
}

```

9.2.9.2 Delegation of feature state

An executor may also optionally declare a correspondence between stateful component features (which include receptacles, attributes, and event sources) and members of the abstract storage type that incarnates the component (or the distinguished component segment, in the case of a segmented executor). The CIF uses this declaration to generate implementations of the feature-specific operations (e.g., `connect_` and `disconnect_` operations for receptacles, accessors, and mutators for attributes) that store the state associated with each specified feature in the storage member indicated on the right hand side of the delegation. The following schematic CIDL illustrates a feature delegation:

```
composition <category> <composition_name> {
  ...
  home executor <home_executor_name> {
    ... // assume storage management specified
    ...
    manages <executor_name> {
      delegatesTo abstract storagetype (
        <feature_name0> : <storage_member_name0> ,
        <feature_name1> : <storage_member_name1> , ...
      );
    };
  };
};
```

The type of the storage member must be compatible with the type associated with the feature, as defined in the *Component Model* sub clause. In the case of attributes, the CIF-generated implementations of accessors and mutators retrieve and store the attribute value in the specified storage member. The executor programming model allows implementors to intercept invocations of the generated accessor and mutator invocations and replace or extend their behaviors. In the case of receptacles and event sources, the implementations of the **`connect_<receptacle_name>`**, **`disconnect_<receptacle_name>`**, **`connect_<source_name>`**, **`disconnect_<source_name>`**, **`subscribe_<source_name>`**, and **`unsubscribe_<source_name>`** operations store the connected object references in the specified members of the storage object that incarnates the component.

This mechanism is only particularly useful if the connected object references are persistent references, capable of causing server and object activation if necessary.

9.2.10 Proxy Homes

A composition definition may include a *proxy home* declaration. A proxy home implements the component home interface specified by the composition definition, but the implementation is not required to be collocated with the container where the components managed by the home are activated.

Proxy homes are, in essence, remote projections of the actual home implementation, which is always collocated with the executing component implementation. A proxy home may be able to implement some subset (or potentially, all) of the operations defined on the component home without contacting the actual home implementation. Operations that cannot be locally implemented by the proxy home are delegated to the actual home. The run-time implementation of the CIF (including the supporting infrastructure of the container and the home finder) is responsible for maintaining the associations between proxy homes and the actual home they represent. The container provides an interface for registering proxy homes, described in “The ProxyHomeRegistration Interface” on page 162.

Proxy homes offer the capacity for considerably increased scalability over collocated homes, particularly when the home operations can be implemented locally by the proxy home implementation. The following schematic CIDL illustrates a proxy home definition:

```
composition <category> <composition_name> {
    ...
    home executor <home_executor_name> {
        implements <home_type> ;
        bindsTo <abstract_storage_home>;
        ...
    };
    proxy home <proxy_executor_name> {
        delegatesTo home ( <home_op0>, <home_op1>, ... );
        abstract (<home_op2>, <home_op3>, ...);
    };
};
```

The **<proxy_executor_name>** is used as the name of the generated skeleton artifact for the proxy home executor. The proxy home declaration implicitly acquires the characteristics of the actual home, as declared in the home executor definition (which must precede the proxy home definition in the composition scope). In particular, the proxy home implements the same home, and binds to the same abstract storage home. The operation delegations specified in the actual home executor definition are also acquired by the proxy home, but certain delegations are transformed according to rules specified in “Proxy home delegation” on page 123.

9.2.10.1 Proxy home delegation

For proxy homes in compositions that specify managed state, the CIF assumes that the proxy home has connectivity to the same persistent store as the actual home. Based on this assumption, the default implementations of orthodox operations on the proxy home executor are delegated directly to the storage home, precisely as they are in the actual home executor. In general, other operations are delegated to the actual home, by default, although the specific rules for determining the implementation of proxy home operations are somewhat more involved, and are described completely in “Implementing a CORBA Component” on page 93.

9.2.11 Component Object References

The CIF defines an information model for component object references. This information model is encapsulated within the **object_key** field of an IIOP profile, or an equivalent field in other profiles. The information model is an abstraction; no standard encoding within an **object_key** is specified. It is the responsibility of the container and the underlying ORB to encode this information for insertion into object references and to extract this information from the **object_key** in incoming requests, decode it, and use it to activate the appropriate component or segment and dispatch the request to the proper facet.

The **Entity2Context** interface, described in “The Entity2Context Interface” on page 169 is used by the component implementation to provide this information to the container, with which the container creates the object references for the component and its facets. The **ComponentId** interface encapsulates the component reference information. Examples 2, 3, and 4 in the previous sub sections illustrate the use of the **Entity2Context** and **ComponentId** interfaces to create object references. Figure 9.3 and Figure 9.4 illustrate the structure of the information encapsulated in **ComponentId**, and its relationship to executor structure.

9.2.11.1 Facet identifiers

The CIF implementation allocates numeric identifiers to facets. The facet ID values are interpreted by generated code in the component implementation, so the assignment of values does not need to be uniformly specified; a given CIF implementation's choice of facet ID values does not affect portability or interoperability.

9.2.11.2 Segment identifiers

The CIF implementation must also allocate numeric identifiers to segments. Similar to facet IDs, segment IDs are also interpreted by the component implementation, so no uniform allocation mechanism is specified. The implementation of **create_executor_segment** (on the home executor implementation) provided by the component implementor must interpret segment ID values in order to create and return the appropriate segment executor. The generated implementations of segment executor skeletons define symbolic constants to assist the component implementor in this mapping.

9.2.11.3 State identifiers

State identifier is an abstraction that generalizes different representations of state identifiers, the primary of which is the **pid** of the CORBA persistent state service. The generic representation of a state identifier is **StateIdValue**, an abstract valuetype from which specific, concrete state identity types are derived. Implementations of the concrete sub-types are responsible for converting their representations to byte sequences and back again.

9.2.11.4 Monolithic reference information

Monolithic references contain a facet identifier and a single state identifier. The facet identifier denotes the target facet of the reference (or, of requests made on the reference). The state identifier is interpreted by the component implementation and used to retrieve the component's state. In the case of automatically managed state, the CIF-generated implementation interprets the state identifier as a **pid**, using it to incarnate the component's storage object.

Note that navigation from one facet's reference to another consists of merely replacing the target facet identifier with the facet identifier of the desired facet. This can be accomplished without activating the component.

9.2.11.5 Segmented reference information

The reference information for segmented executors consists of the following:

- a target facet identifier,
- a target segment identifier,
- a sequence of segment descriptors, each of which contains:
 - the segment identifier of the segment being described,
 - the state identifier for the segment.

The target facet identifier denotes the target of requests made on the reference, and the target segment identifier denotes the segment on which that facet is implemented. The sequence of segment descriptors contains one element for each segment, including the component segment. This sequence is invariant for all references to a given component, over the lifetime of the component.

In the case of segmented executors, navigation is accomplished by replacing the facet and segment identifiers.

9.2.11.6 Component identity

The state identifier of the component segment (or the single state identifier in the case of monolithic executors) is interpreted as the unique identity of the component, within the scope of the home to which it belongs. Equivalence of component identity is defined as equivalence of state identifier values of the component segment.

9.3 Language Mapping

9.3.1 Overview

This part describes the language mapping for CORBA Components and defines interfaces that are used to implement components and homes. The language mapping, like the mapping for the client side, is based on equivalent IDL. For components and homes, local interfaces are defined. The user then implements these local interfaces using existing language mapping rules.

There are two strategies for implementing a component, coined *monolithic* and *locator*. In the monolithic strategy, the user implements all attributes, supported interfaces, and event consumers in a single executor interface. In the locator strategy, the user implements a locator, and the container uses this locator to retrieve references to executors for each port of a component. The decision which strategy is being used is made by the home, which can return a reference to either the monolithic or to the locator.

It is expected that the monolithic strategy is more simple to use and that it is sufficient for most use cases, while the locator strategy gives the user even more control over the life cycle of each executor.

Interfaces are designated internal or callback. Callback interfaces are implemented by the user and called by the container, while internal interfaces are provided by the container.

Some callback interfaces may be optionally implemented by the user. In order to optionally implement an interface, the user must define an interface, in IDL, that inherits both the base interface and the optional interface. For example, to inherit the optional **SessionSynchronization** interface in the implementation of a **Bank** component, the user would declare a new local interface, as shown below.

```
local interface MyBank :  
  Components::SessionSynchronization,  
  CCM_Bank };
```

Optional interfaces are used by services that require the component's cooperation (and therefore callback hooks). To determine whether an implementation supports an optional interface, the container narrows the object reference to that interface.

Internal interfaces are used by the container and various services to provide runtime information to the component. The component accesses internal interfaces through the context reference that it acquires through the **set_session_context** operation.

Details about existing internal and callback interfaces can be found in the Container Programming Model clause. Some of those interfaces are forward-referenced in this sub clause.

9.3.2 Common Interfaces

EnterpriseComponent is an empty callback interface that serves as common base for all component implementations, whether monolithic or locator-based.

```
module Components {
    local interface EnterpriseComponent {};
};
```

Note – The **EnterpriseComponent** interface is also defined in the Container Programming Model clause.

The **ExecutorLocator** interface is a callback interface that is used for the locator implementation strategy.

```
module Components {
    local interface ExecutorLocator : EnterpriseComponent {
        Object obtain_executor (in string name)
            raises (CCMException);
        void release_executor (in Object exc)
            raises (CCMException);
        void configuration_complete()
            raises (InvalidConfiguration);
    };
};
```

If a home, in creating a component, returns an **ExecutorLocator**, the container will invoke its **obtain_executor** operation prior to each invocation to retrieve the implementation for a port. The port name, given in the **name** parameter, is the same as used in the component's interface description in IDL, or the component's name for the "main" executor. The **obtain_executor** operation returns a local object reference of the expected type, as detailed below. The **CCMException** exception may be raised in case of a system error that prohibits locating the requested executor.

The **release_executor** operation is called by the container once the current invocation on an executor that was obtained through the **obtain_executor** operation has finished. The locator can thus release any resources that were acquired as part of the **obtain_executor** operation.

The **configuration_complete** operation is called to propagate the **configuration_complete** operation on the **CCMObject** interface to the component implementation.

Implementations of the **ExecutorLocator** interface for a service or session component must implement the **Components::SessionComponent** interface. Implementations of the **ExecutorLocator** interface for a process or entity component must implement the **Components::EntityComponent** interface.

Note – **Object** is used as the return type of the **obtain_executor** operation, because there is yet no IDL type for the common base of all local objects. Since local objects inherit from **Object**, this is not a problem.

The **HomeExecutorBase** interface is a common base for all home implementations.

```
module Components {
    local interface HomeExecutorBase {};
};
```

9.3.3 Mapping Rules

This sub clause defines equivalent interfaces that are generated for each interface, eventtype, component, and home.

9.3.3.1 Interfaces

For each non-abstract and non-local interface, a local *facet executor* interface is generated. This facet executor interface has the same name as the original interface with a “**CCM_**” prefix, and inherits the original interface. So for an interface of name *<interface name>*, the facet executor interface has the following form:

```
local interface CCM_<interface name> : <interface name> { };
```

If a component provides an interface as a facet, the user implements the facet executor interface rather than the original interface in order to achieve a local implementation.

Note – A container implementation may choose to limit generation of facet executor interfaces to only those interfaces that are actually used as a facet.

9.3.3.2 Eventtypes

For each eventtype, a local *consumer executor* interface is generated. For an eventtype *<eventtype name>*, a local interface with the same name, but with a “**CCM_**” prefix and a postfix of “**Consumer**” is generated. This interface has a single **push** operation with no result, and the eventtype as a single **in** parameter:

```
local interface CCM_<eventtype name>Consumer  
{  
    void push (in <eventtype name> ev);  
};
```

9.3.3.3 Components

A component maps to three local interfaces; two of them are callback interfaces, and one is an internal interface. The *monolithic executor* callback interface is for use in monolithic implementations, the *main executor* callback interface is for use in locator-based implementations. Both callback interfaces inherit the component’s base and supported interfaces. They also both expose the component’s attributes.

In addition, the monolithic executor callback interface also contains operations for acquiring references to facets, and for consuming events - in the locator approach, these jobs are mediated by the locator.

An internal *context* interface is defined for each component. It is implemented by the container and handed to the component as session or entity context. The context interface contains component-specific runtime information (e.g., for pushing events into event source ports).

Component Main Executor Interface

The *main executor* callback interface as used by the locator approach is defined by the following rules:

1. For each component *<component name>*, a local *main executor* interface with the same name as the component, but with a prefix of “**CCM_**” and a postfix of “**_Executor**” is defined.

2. The main executor interface contains all attributes declared by the component.
3. If the component has a base component with a name of *<base name>*, the main executor interface inherits **CCM_<base name>_Executor**. If the component does not have a base, the main executor interface inherits **Components::EnterpriseComponent**.
4. If the component has supported interfaces, they are inherited by the main executor interface.

If the container desires to acquire a reference to the main executor, it calls the **obtain_executor** operation of the **ExecutorLocator** with the name parameter set to *<component name>*.

Component Monolithic Executor Interface

The *monolithic executor* callback interface is defined by the following rules:

1. For each component *<component name>*, a local *monolithic executor* interface with the same name as the component and a prefix of “**CCM_**” is defined.
2. The monolithic executor interface contains all attributes declared by the component.
3. If the component has a base component with a name of *<base name>*, the monolithic executor interface inherits **CCM_<base name>**. If the component does not have a base, the monolithic executor interface inherits **Components::EnterpriseComponent**.
4. If the component has supported interfaces, they are inherited by the monolithic interface.
5. Additional operations are added to the monolithic interface for facets and event sinks.
6. Above rules can be satisfied by inheriting the main executor interface and adding operations for facets and event sinks. This is an optional design choice by the container implementation.

In a service and session component, the user may optionally inherit the **Components::SessionComponent** interface in the implementation of a monolithic executor in order to be notified by the container of activation and passivation. In a process or entity component, the user may optionally inherit the **Components::EntityComponent** interface in the implementation of a monolithic executor.

Component Context Interface

The *context* internal interface is defined by the following rules:

1. For each component *<component name>*, a local *context* interface with the same name as the component, but with a prefix of “**CCM_**” and a postfix of “**_Context**” is defined.
2. If the component has a base component with a name of *<base name>*, the context interface inherits **CCM_<base name>_Context**. If the component does not have a base, the context interface inherits **Components::CCMContext**.
3. Additional operations are added to the context interface for receptacles and event sources.

The container will implement an interface that inherits both the above context interface and either **Components::SessionContext** or **Components::EntityContext**, depending on the type of the component. The component implementation can narrow the **Components::SessionContext** or **Components::EntityContext** reference that it receives to the above component-specific context interface.

9.3.3.4 Example

For the following component declaration in IDL,

```
interface Hello {
    void sayHello ();
};

component HelloWorld supports Hello {
    attribute string message;
};
```

the following local interfaces are generated:

```
local interface CCM_Hello : Hello
{
}

local interface CCM_HelloWorld_Executor :
    Components::EnterpriseComponent, Hello
{
    attribute string message;
};

local interface CCM_HelloWorld :
    Components::EnterpriseComponent, Hello
{
    attribute string message;
};

local interface CCM_HelloWorld_Context :
    Components::CCMContext
{
};
```

Read on for further contents of these interfaces.

9.3.3.5 Ports

This sub clause defines equivalent operations that are added to either of the three interfaces for each port definition.

Facets

For each facet, an equivalent operation is defined in the monolithic executor interface. For a facet of name *<name>* and type *<type>*, an operation with the same name as the facet but with a “**get_**” prefix is generated. This operation has an empty parameter list and a reference of the interface’s facet executor type as return value:

```
CCM_<type> get_<name> ();
```

Users may optionally implement facet interfaces directly in the monolithic executor implementation by declaring a new local interface that inherits both the monolithic executor interface and the facet executor, and by then returning a reference to itself in the implementation of the above operation. Example:

```

// IDL
component MyComponent {
    provides MyInterface MyFacet;
};

// User IDL
local interface MyComponentImpl :
    CCM_MyComponent, CCM_MyInterface
{};

// C++
CCM_MyInterface_ptr
MyComponent_Impl::get_MyFacet ()
{
    return CCM_MyInterface::_duplicate (this);
}

```

If the locator strategy is used, the container calls the **obtain_executor** operation on the **ExecutorLocator** with the **name** parameter set to *<name>* in order to acquire a reference to the facet executor that matches this facet port.

Receptacles

For each receptacle, an equivalent operation is defined in the context interface. The signature of this operation depends on whether the receptacle is simplex or multiplex.

For a simplex receptacle of name *<name>* and type *<type>*, an operation of the same name as the receptacle but with a “**get_connection_**” prefix is generated. The operation has an empty parameter list, and an object reference of the interface’s type as return value:

```
<type> get_connection_<name> ();
```

If there is no connection, this operation returns a nil reference.

For a multiplex receptacle of name *<name>* and type *<type>*, an operation of the same name as the receptacle but with a “**get_connections_**” prefix is generated. The operation has an empty parameter list and a sequence of type *<name>Connections* as return value (this type is defined by the client-side equivalent IDL):

```
<name>Connections get_connections_<name> ();
```

9.3.3.5.1 Publisher and Emitter

For each publisher and emitter port, an equivalent operation is defined in the context interface. For a publisher or emitter port of name *<name>* and type *<type>*, an operation of the same name as the port but with a “**push_**” prefix is generated. This operation has no return value and a single **in** parameter containing the event.

```
void push_<name> (in <type> ev);
```

The component may call this operation in order to push an event to the consumer (for emitter ports) or to all subscribers (for publisher ports). The container is responsible for delivering the event.

9.3.3.5.2 Consumer

For each consumer port, an equivalent operation is defined in the monolithic executor interface. For a consumer port of name *<name>* and type *<type>*, an operation of the same name as the port but with a “**push_**” prefix is generated. This operation has no return value and a single **in** parameter containing the event.

```
void push_<name> (in <type> ev);
```

For component implementations that use the monolithic strategy, the container invokes this operation whenever a client sends an event to this sink.

For component implementations that use the locator strategy, the container calls the **obtain_executor** operation on the **ExecutorLocator** with the name parameter set to *<name>* in order to acquire a reference to an implementation of the eventtype’s consumer executor interface.

9.3.3.6 Home

For each home, three callback interfaces are generated, similar in structure to the interfaces defined on the client side. The three interfaces are named the *Implicit*, *Explicit*, and *Main* home executor.

9.3.3.6.1 Home Explicit Executor Interface

The *home explicit executor* callback interface is defined by the following rules:

1. For each home *<home name>*, a local *explicit executor* interface with the same name as the home, but with a prefix of “**CCM_**” and a postfix of “**Explicit**” is defined.
2. The explicit executor interface contains all attributes and operations declared by the home.
3. If the home has a base with a name of *<base name>*, the explicit executor interface inherits **CCM_<base name>Explicit**. If the home does not have a base, the explicit executor interface inherits **Components::HomeExecutorBase**.
4. If the home has supported interfaces, they are inherited by the explicit executor interface.
5. Additional operations are added to the explicit executor interface for factories and finders, see below.

9.3.3.6.2 Home Implicit Executor Interface

The contents of the *home implicit executor* callback interface depend on whether the home is keyless or keyed.

9.3.3.6.3 Implicit Executor Interface for Keyless Homes

For a *keyless* home *<home name>*, a local implicit executor interface with the same name as the home, but with a prefix of “**CCM_**” and a postfix of “**Implicit**” is defined. This interface contains a single **create** operation with the following signature:

```
local interface CCM_<home name>Implicit {  
    Components::EnterpriseComponent create ()  
    raises (Components::CCMException);  
};
```

The container calls the implicit **create** operation in order to create a new component instance. The operation can return either a reference to a monolithic executor or to an **ExecutorLocator**. In the former case, the container assumes that the monolithic strategy is used, otherwise it will use the locator strategy. The implementation may raise the **CCMException** exception in order to indicate a system-level error.

9.3.3.6.4 Implicit Executor Interface for Explicitly or Implicitly Keyed Homes

For a keyed home *<home name>* with a key of *<key type>* or a keyless home *<home name>* that derives from a keyed home with a key of *<key type>*, a local implicit executor interface with the same name as the home, but with a prefix of “**CCM_**” and a postfix of “**Implicit**” is defined. This interface contains the following operations:

```
local interface CCM_<home name>Implicit {
    Components::EnterpriseComponent
        create (in <key type> key)
            raises (Components::CCMException);
    Components::EnterpriseComponent
        find_by_primary_key (in <key type> key)
            raises (Components::CCMException);
    void remove (in <key_type> key)
        raises (Components::CCMException);
};
```

The container calls the **create** operation in order to create a new component associated with the specified primary key value. The operation can return either a reference to a monolithic executor or to an **ExecutorLocator**. In the former case, the container assumes that the monolithic strategy is used, otherwise it will use the locator strategy. The operation may raise the **CCMException** exception to indicate a system-level error.

The container calls the **find_by_primary_key** operation in order to find an existing component associated with the specified primary key value. The operation shall return the same reference to a monolithic executor or to an **ExecutorLocator** as it was previously returned from a **create** operation. The operation may raise the **CCMException** exception to indicate a system-level error.

The container calls the **remove** operation in order to remove the component identified by the specified primary key value. The operation may raise the **CCMException** exception to indicate a system-level error.

9.3.3.6.5 Home Main Executor Interface

For each home *<home name>*, a local *main executor* interface with the same name as the home and a prefix of “**CCM_**” is defined. The main executor interface inherits both the implicit and explicit executor interfaces, as shown below.

```
local interface CCM_<home name> :
    CCM_<home name>Explicit,
    CCM_<home name>Implicit
{
};
```

The main executor interface does not have any other contents.

In the implementation of a home main executor for a service and session component, the user may optionally inherit the **Components::SessionComponent** interface in order to be notified by the container of activation and passivation. In the implementation of a home main executor for a process or entity component, the user may optionally inherit the **Components::EntityComponent** interface.

Note – This structure allows implementation inheritance for the explicit interface without name clashes in the implicit interface.

9.3.3.6.6 Factories

For each factory in the home, an operation is defined in the explicit home executor interface. This operation has the same parameter list as the factory and the return type **EnterpriseComponent**. As with the home's **create** operation, factories can return either a reference to a monolithic executor or to an **ExecutorLocator**.

Factories are assumed to return a new component instance.

9.3.3.6.7 Finders

For each finder in the home, an operation is defined in the explicit home executor interface. This operation has the same parameter list as the finder and the return type **EnterpriseComponent**. As with the home's **create** operation, finders can return either a reference to a monolithic executor or to an **ExecutorLocator**.

Finders may return existing or new component instances. If a finder decides to return an existing component instance, it shall return the same reference to a monolithic executor or to an **ExecutorLocator** as it was previously returned from a factory or from the **create** operation.

9.3.3.6.8 Entry Points

Some programming languages require the existence of user-provided entry points, or *Home Factories*. These entry points are not part of the language mapping; they are dealt with in the *Packaging and Deployment* clause.

Home Factories, if required by a language mapping, shall return a reference to an instance of the home's main executor interface.

9.3.3.6.9 Example

The following example shows a **Bank** home that manages an **Account** component.

```
home Bank manages Account {  
    factory open (in string name);  
    void close (in string name);  
};
```

In this example, the following equivalent interfaces would be generated.

```
local interface CCM_BankExplicit :  
    Components::HomeExecutorBase  
{  
    Components::EnterpriseComponent open (in string name);  
    void close (in string name);  
};
```

```
local interface CCM_BankImplicit :  
{
```

```
    Components::EnterpriseComponent create ()  
        raises (Components::CCMException);  
};  
  
local interface CCM_Bank :  
    CCM_BankExplicit,  
    CCM_BankImplicit  
{  
};
```

The user would then implement the **CCM_Bank** interface and eventually provide an entry point that creates a **CCM_Bank** instance.

10 The Container Programming Model

10.1 General

The container is the server's runtime environment for a CORBA component implementation. This environment is implemented by a deployment platform such as an application server or a development platform like an IDE. A deployment platform typically provides a robust execution environment designed to support very large numbers of simultaneous users. A development platform would provide enough of a runtime to permit customization of CORBA components prior to deployment but perhaps support a limited number of concurrent users. From the point of view of the CORBA component implementation, such differences are "qualities of service" characteristics and have no effect on the set of interfaces the component implementor can rely on. This clause is organized as follows:

- "Introduction" on page 135 introduces the programming model and defines the elements that comprise it.
The container programming model is an API framework designed to simplify the task of building a CORBA application. Although the framework does not exclude the component developer from using any function currently defined in CORBA, it is intended to be complete enough in itself to support a broad spectrum of applications.
- "The Server Programming Environment" on page 138 describes the programming model the component implementor is to follow.
The programming model identifies the architectural choices which must be made to develop a CORBA component which can be deployed in a container.
- "Server Programming Interfaces - Basic Components" on page 150 describes the interfaces seen by the component developer.
These interfaces constitute the contract between the container provider and the component implementor. Together with the client programming interfaces defined in the *Component Model* clause, which can be used by servers as well as clients, they define the server programmer's API.
- "The Client Programming Model" on page 170 describes the client view of a CORBA component.
The **client programming model** has been described previously (see the *Component Model* clause). This sub clause describes the specific use of CORBA required by a client, which is **NOT** itself a CORBA component, to use a CORBA component written to the server programming model described in "Server Programming Interfaces - Basic Components" on page 150.

10.2 Introduction

The container programming model is made up of several elements:

- The **external API types** define the interfaces available to a component client.
- The **container API type** defines the API framework used by the component developer.
- The CORBA usage model defines the interactions between the container and the rest of CORBA (including the POA, the ORB, and the CORBA services).
- The component category is the combination of the container API type (i.e., the server view) and the external API types (i.e., the client view).

The overall architecture is depicted in Figure 10.1.

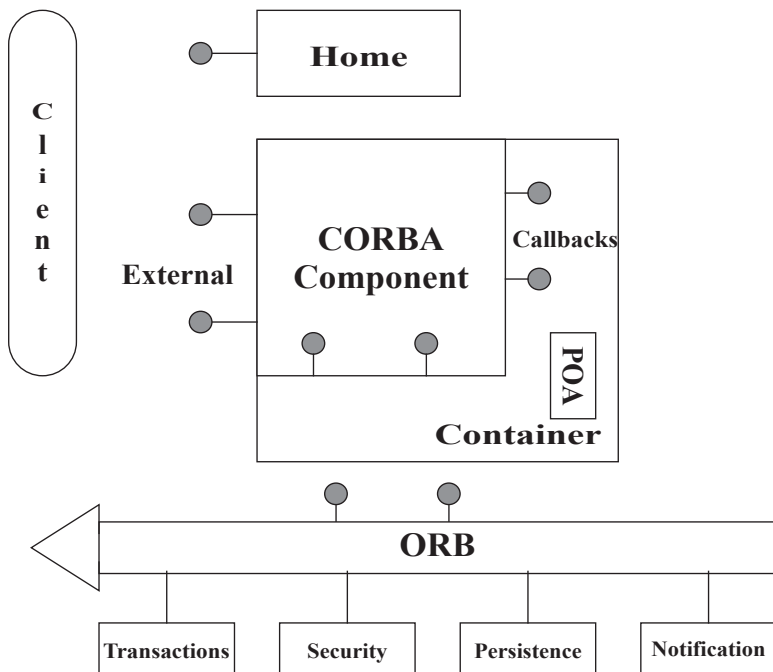


Figure 10.1- The Architecture of the Container Programming Model

The external API types are defined by the component IDL including the home specification. These interfaces are righteous CORBA objects and are stored in the Interface Repository for client use.

The **container API type** is a framework made up of internal interfaces and callback interfaces used by the component developer. These are defined using the new **local interface** declaration in IDL for specifying locality-constrained interfaces. The container API type is selected using CIDL, which describes component implementations.

The EJB session bean and entity bean can be viewed as two examples of container API type since they offer different sets of framework APIs to the EJB programmer. However, each of them also implies a client view (i.e., the external API types). EJB does not define a term for the two framework API sets it supports.

The CORBA usage model is controlled by policies that specify distinct interaction patterns with the POA and a set of CORBA services. These are defined by CIDL, augmented using XML, and used by the container factory to create a POA when the container is created.

The component category is a specific combination of external API types and container API type used to implement an application with the CORBA component technology.

10.2.1 External API Types

The external API types of a component are the contract between the component developer and the component client. We distinguish between two forms of external API types: the **home** interface and the **application** interfaces.

*These are analogous to the **EJBHome** and **EJBObject** interfaces of Enterprise JavaBeans.*

Home interfaces support operations that allow the client to obtain references to one of the application interfaces the component implements. From the client's perspective, two design patterns are supported - factories for creating new objects and finders for existing objects. These patterns are distinguished by the presence of a **primarykey** parameter in the home IDL declaration.

- A home interface with a **primarykey** declaration supports finders and its client is a **keyfull** client.
- A home interface without a **primarykey** declaration does not support finders and its client is a **keyless** client. All home types support factory operations.

10.2.2 Container API Type

The **container API type** defines an API framework; that is, the contract between a specific component and its container. This specification defines two base types that define the common APIs and a set of derived types that provide additional function. The **session** container API type defines a framework for components using transient object references. The **entity** container API type defines a framework for components using persistent object references.

10.2.3 CORBA Usage Model

A CORBA usage model specifies the required interaction pattern between the container, the POA, and the CORBA services. We define three **CORBA usage models** as part of this specification. Since all support the same set of CORBA services, they are distinguished only by their interaction with the POA.

- **stateless** - which uses transient object references in conjunction with a POA servant that can support any **Objectld**.
- **conversational** - which uses transient references in conjunction with a POA servant that is dedicated to a specific **Objectld**.
- **durable** - which uses persistent references in conjunction with a POA servant that is dedicated to a specific **Objectld**.

It should be obvious that the fourth possibility (persistent references with a POA servant that can support any Objectld) makes no sense and is therefore not included.

10.2.4 Component Categories

The component categories are defined as the valid combinations of external API types, container API type, and CORBA usage model. Table 10.1 summarizes the categories and identifies their EJB equivalent.

Table 10.1 - Definition of the Component Categories

CORBA Usage Model	Container API Type	Primary Key	Component Categories	EJB Bean Type
stateless	session	No	Service	-
conversational	session	No	Session	Session
durable	entity	No	Process	-
durable	entity	Yes	Entity	Entity

10.3 The Server Programming Environment

The component container provides interfaces to the component. These interfaces support access to CORBA services (transactions, security, notification, and persistence) and to other elements of the component model. This sub clause describes the features of the container that are selected by the deployment descriptor packaged with the component implementation. These features comprise the design decisions to be made in developing a CORBA component. Details of the interfaces provided by the container are provided in “Server Programming Interfaces - Basic Components” on page 150.

10.3.1 Component Containers

Containers provide the run-time execution environment for CORBA components. A container is a framework for integrating transactions, security, events, and persistence into a component’s behavior at runtime. A container provides the following functions for its component:

- All component instances are created and managed at runtime by its container.
- Containers provide a standard set of services to a component, enabling the same component to be hosted by different container implementations.

Components and homes are deployed into containers with the aid of container specific tools. These tools generate additional programming language and metadata artifacts needed by the container. The tools provide the following services:

- Editing the configuration metadata,
- editing the deployment metadata, and
- generating the implementations needed by the containers to support the component.

The container framework defines two forms of interfaces:

- **Internal interfaces** - These are locality-constrained interfaces defined as **local interface** types, which provide container functions to the CORBA component.
*These are similar to the **EJBContext** interface in Enterprise JavaBeans.*
- **Callback interfaces** - These are also **local interface** types invoked by the container and implemented by a CORBA component.

*These interfaces provide functions analogous to the **SessionBean** and **EntityBean** interfaces defined by Enterprise JavaBeans.*

This architecture is depicted in Figure 10.1 on page 136.

We define a small set of **container API types** to support a broad spectrum of component behavior with their associated **internal** and **callback** interfaces as part of this specification. These **container API types** are defined using local interfaces.

Additional component behavior is controlled by policies specified in the deployment descriptor. This specification defines policies that support POA interactions (CORBA usage model), servant lifetime management, transactions, security, events, and persistence.

CORBA containers are designed to be used as Enterprise JavaBeans containers. This allows a CORBA infrastructure to be the foundation of EJB, enabling a more robust implementation of the EJB specification. To support enterprise Beans natively within a CORBA container, the container must support the API frameworks defined by the EJB specification. This architecture is defined in the Integrating with Enterprise JavaBeans clause of this specification.

10.3.2 CORBA Usage Model

The CORBA Component Specification defines a set of **CORBA usage models** that create either **TRANSIENT** or **PERSISTENT** object references and use either a 1:1 or 1:N mapping of **Servant** to **Objectld**. These CORBA usage models are summarized in Table 10.2. A given component implementation shall support one and only one CORBA usage model.

Table 10.2 - CORBA Usage Model Definitions

CORBA Usage Model	Object Reference	Servant:OID Mapping
stateless	TRANSIENT	1:N
conversational	TRANSIENT	1:1
durable	PERSISTENT	1:1
(Invalid)	PERSISTENT	1:N

A CORBA usage model is specified using CIDL and is used to either create or select a component container at deployment time.

10.3.2.1 Component References

TRANSIENT objects support only the factory design pattern. They are created by operations on the home interface defined in the `component` declaration.

PERSISTENT objects support either the factory design pattern or the finder design pattern, depending on the component category. **PERSISTENT** objects support **self-managed** or **container-managed** persistence. **PERSISTENT** objects can be used with the CORBA persistent state service or any user-defined persistence mechanism. When the CORBA persistent state service is used, servant management is aligned with the **Persistentld** defined by the CORBA persistent state service and the container supports the transformation of an **Objectld** to and from a **Persistentld**. A **Persistentld** provides a persistent handle for a class of objects whose permanent state resides in a persistent store (e.g., a database).

Home references are exported for client use by registering them with a **HomeFinder** which the client subsequently interrogates or by binding them to the CORBA naming service in the form of externally visible names.

EJB clients find references to EJBHome using JNDI, the Java API for CosNaming. Placing home references is CosNaming supports both the CORBA component client and the EJB client programming models.

10.3.2.2 Servant to Objectld Mapping

Component implementations may use either the 1:1 or 1:N mapping of **Servant** to **Objectld** with **TRANSIENT** references (**stateless** and **conversational** CORBA usage model, respectively) but may use only the 1:1 mapping with **PERSISTENT** references.

- A 1:N mapping allows a **Servant** to be shared among all requests for the same interface and therefore requires the object to be stateless (i.e., it has no identity).

- A 1:1 mapping binds a **Servant** to a specific **Objectld** for an explicit servant lifetime policy (see “Servant Lifetime Management” on page 140) and therefore is stateful.

10.3.2.3 Threading Considerations

CORBA components support two threading models: **serialize** and **multithread**. A threading policy of **serialize** means that the component implementation is not thread safe and the container will prevent multiple threads from entering the component simultaneously. A threading policy of **multithread** means that the component is capable of mediating access to its state without container assistance and multiple threads will be allowed to enter the component simultaneously. Threading policy is specified in CIDL.

*A threading policy of **serialize** is required to support an enterprise Bean since they are defined to be single-threaded.*

10.3.3 Component Factories

A home is a component factory, responsible for creating instances of all interfaces exported by a component. Factory operations are defined on the home interface using the **factory** declaration. A default factory is automatically defined whose implementation may be generated by tools using the information provided in the **component** IDL. Specialized factories; for example, factories that accept user-defined input arguments must be implemented by the component developer. Factory operations are typically invoked by clients but may also be invoked as part of the implementation of the component. A CORBA component implementation can locate its home interface using an interface provided by the container.

10.3.4 Component Activation

CORBA components rely on the automatic activation features of the POA to tailor the behavior of the components using information present in the component’s deployment descriptor. Once references have been exported, clients make operation requests on the exported references. These requests are then routed by the ORB to the POA that created the reference and then the component container. This enables the container to control activation and passivation for components, apply policies defined in the component’s descriptor, and invoke callback interfaces on the component as necessary.

10.3.5 Servant Lifetime Management

Servants are programming language objects that the POA uses to dispatch operation requests based on the **Objectld** contained in the object key. The server programming model for CORBA components includes facilities to efficiently manage the memory associated with these programming objects. To implement this sophisticated memory management scheme, the server programmer makes several design choices:

- The container API type must be chosen.
- The CORBA usage model must be chosen.
- A servant lifetime policy is selected. CORBA components support four servant lifetime policies (**method**, **transaction**, **component**, and **container**).
- The designer is required to implement the callback interface associated with his choice.

The servant lifetime policies are defined as follows.

method

The **method** servant lifetime policy causes the container to activate the component on every operation request and to passivate the component when that operation has completed. This limits memory consumption to the duration of an operation request but incurs the cost of activation and passivation most frequently.

transaction

The **transaction** servant lifetime policy causes the container to activate the component on the first operation request within a transaction and leave it active until the transaction completes and which point the component will be passivated. Memory remains allocated for the duration of the transaction.

component

The **component** servant lifetime policy causes the container to activate the component on the first operation request and leave it active until the component implementation requests it to be passivated. After the operation that requests the passivation completes, the component will be passivated by the container. Memory remains allocated until explicit application request.

container

The **container** servant lifetime policy causes the container to activate the component on the first operation request and leave it active until the container determines it needs to be passivated. After the current operation completes, the component will be passivated by the container. Memory remains allocated until the container decides to reclaim it.

Table 10.3 shows the relationship between the CORBA usage model, the container API type, and the servant lifetime policies.

Table 10.3 - Servant Lifetime Policies by Container API Type

CORBA Usage Model	Container API Type	Valid Servant Lifetime Policies
stateless	session	method
conversational	session	method, transaction, component, container
durable	entity	method, transaction, component, container

Servant lifetime policies may be defined for each segment within a component.

10.3.6 Transactions

CORBA components may support either **self-managed transactions** (SMT) or **container-managed transactions** (CMT). A component using self-managed transactions will not have transaction policies defined with its deployment descriptor and is responsible for transaction demarcation using either the container's **UserTransaction** interface or the CORBA transaction service. A component using container-managed transactions defines transaction policies in its associated descriptor. The selection of container-managed transactions vs. self-managed transactions is a component-level specification.

When container-managed transactions are selected, additional transaction policies are defined in the component's deployment descriptor. The container uses these descriptions to make the proper calls to the CORBA transaction service. The transaction policy defined in the component's deployment descriptor is applied by the container prior to invoking the operation. Differing transaction policy declarations can be made for operations on any of the component's ports as well as for the component's home interface.

Table 10.4 summarizes the effects of the various transaction policy declarations and the presence or absence of a client transaction on the transaction that is used to invoke the requested operation on the component.

Table 10.4 - Effects of Transaction Policy Declaration

Transaction Attribute	Client Transaction	Component's Transaction
NOT_SUPPORTED	-	-
	T1	-
REQUIRED	-	T2
	T1	T1
SUPPORTS	-	-
	T1	T1
REQUIRES_NEW	-	T2
	T1	T2
MANDATORY	-	EXC (TRANSACTION_REQUIRED)
	T1	T1
NEVER	-	-
	T1	EXC (INVALID_TRANSACTION)

not_supported

This component does not support transactions. If the client does not provide a current transaction, the operation is invoked immediately. If the client provides a current transaction, it is suspended (**CosTransactions::Current::suspend**) before the operation is invoked and resumed (**CosTransactions::Current::resume**) when the operation completes.

required

This component requires a current transaction to execute successfully. If one is supplied by the client, it is used to invoke the operation. If one is not provided by the client, the container starts a transaction (**CosTransactions::Current::begin**) before invoking the operation and attempts to commit the transaction (**CosTransactions::Current::commit**) when the operation completes.

supports

This component will support transactions if one is available. If one is provided by the client, it is used to invoke the operation. If one is not provided by the client, the operation is invoked outside the scope of a transaction.

requires_new

This component requires its own transaction to execute successfully. If no transaction is provided by the client, the container starts one (**CosTransactions::Current::begin**) before invoking the operation and tries to commit it (**CosTransactions::Current::commit**) when the operation completes. If a transaction is provided by the client, it is

first suspended (**CosTransactions::Current::suspend**), a new transaction is started (**CosTransactions::Current::begin**), the operation invoked, the component's transaction attempts to commit (**CosTransactions::Current::commit**), and the client's transaction is resumed (**CosTransactions::Current::resume**).

mandatory

The component requires that the client be in a current transaction before this operation is invoked. If the client is in a current transaction, it is used to invoke the operation. If not, the **TRANSACTION_REQUIRED** exception shall be raised.

never

This component requires that the client not be in a current transaction to execute successfully. If no current transaction exists, the operation is invoked. If a current transaction exists, the **INVALID_TRANSACTION** exception shall be raised.

10.3.7 Security

Security policy is applied consistently to all categories of components. The container relies on CORBA security to consume the security policy declarations from the deployment descriptor and to check the active credentials for invoking operations. The security policy remains in effect until changed by a subsequent invocation on a different component having a different policy.

Access permissions are defined by the deployment descriptor associated with the component. The granularity of permissions must be aligned by the deployer with a set of rights recognized by the installed CORBA security mechanism since it will be used to check permissions at operation invocation time. Access permissions can be defined for any of the component's ports as well as the component's home interface.

Note – The security model used by EJB and being adopted by CORBA components requires the secure transportation of security credentials between systems. Today that is only possible if SECIOP is used as the CORBA transport.

10.3.8 Events

CORBA components use a simple subset of the CORBA notification service to emit and consume events. The subset can be characterized by the following attributes:

- Events are represented as **valuetypes** to the component implementor and the component client.
- The event data structure is mapped to an **any** in the body of a structured event presented to and received from CORBA notification.
- The fixed portion of the structured event is added to the event data structure by the container on sending and removed from the event data structure when receiving.
- Components support two forms of event generation using the push model:
 - A component may be an exclusive supplier of a given type of event.
 - A component may supply events to a shared channel that other CORBA notification users are also utilizing.
- A CORBA component consumes both forms of events using the push model.

- Events have transaction and security policies associated with the component's event ports as defined in the deployment descriptor.
- All channel management is implemented by the container, not the component.
- Filters are set administratively by the container, not the component.

Because events can be emitted and consumed by clients as well as component implementations, operations for emitting and consuming events are generated from the specifications in component IDL. The container is responsible for mapping these operations to the CORBA notification service to provide a robust event distribution network.

10.3.8.1 Transaction Policies for Events

Transaction policies are defined for component event ports, which include both events being generated and events being consumed. The possible values are as follows:

normal

A **normal** event policy indicates the event should be generated or consumed outside the scope of a transaction. If a current transaction is active, it is suspended before sending the event or invoking the operation on the proxy object provided by the component.

default

A **default** event policy indicates the event should be generated or consumed regardless of whether a current transaction exists. If a current transaction is active, the operation is transactional. If not, it is non-transactional.

transaction

A **transaction** event policy indicates the event should be generated or consumed within the scope of a transaction. If a current transaction is not active, a new one is initiated before sending the event or invoking the operation on the proxy object provided by the component. The new transaction is committed as soon as the operation is complete.

Transaction policy declarations can be defined in the deployment descriptor for each event port defined by the component.

10.3.8.2 Security Policies for Events

CORBA components permits access control policies based on roles to be associated with the generation and consumption of events. This is accomplished by associating ACLs with the component ports used to emit/publish and consume events and using CORBA security to restrict access. These policies provide access control based on role for both event generation and consumption.

10.3.9 Persistence

The **entity** container API type supports the use of a persistence mechanism for making component state durable; for example, storing it in a persistent store like a database. The **entity** container API type defines two forms of persistence support:

- **container-managed persistence (CMP)** - the component developer simply defines the state that is to be made persistent and the container (in conjunction with generated code) automatically saves and restores state as required.

Container-managed persistence is selected by defining the abstract state associated with a component segment using the state declaration language of the CORBA persistent state service and connecting that state declaration to a component segment using CIDL.

- **self-managed persistence (SMP)** - the component developer assumes the responsibility for saving and restoring state when requested to do so by the container.

Self-managed persistence is selected via CIDL declaration and triggered by the container invoking the callback interfaces (which the component must implement) defined later in this clause (“Server Programming Interfaces - Basic Components” on page 150).

Table 10.5 summarizes the choices and their required responsibilities.

Table 10.5 - Persistence Support for Entity Container API Type

Persistence Support	Persistence Mechanism	Responsibility	Persistence Classes	Callback Interfaces
Container Managed	CORBA	Container	Generated Code	Generated Code
Container Managed	User	Container	Component implements	Generated Code
Self-managed	CORBA	Component	Generated Code	Component implements
Self-managed	User	Component	Component implements	Component implements

Container-managed vs. self-managed persistence is selected via the deployment descriptor for each segment of the component.

10.3.9.1 Container-managed Persistence

Container-managed persistence may be accomplished using the CORBA persistent state service or any user-defined persistence mechanism. When the CORBA persistent state service is used, the container manages all interactions with the persistence provider and the component developer need not use the persistence interfaces offered by the container. With container-managed persistence using the CORBA persistent state service, it is possible to provide automatic code generation for the storage factories, finders, and some callback operations.

If container-managed persistence is to be accomplished with a user-defined persistence mechanism, the component developer must implement the various persistence classes defined in the persistence framework.

Container-managed persistence is selected using CIDL and tailored using XML at deployment time to specify connections to specific persistence providers and persistent stores.

10.3.9.2 Self-managed Persistence

Self-managed persistence is also supported by the **entity** container API type. Like container-managed persistence, the component developer has two choices: to use the CORBA persistent state service or some user-defined persistence mechanism. But since no declarations are available to support code generation, the component developer is responsible for implementing both the callback interfaces and the persistence classes. The container supports access to a component persistence abstraction provided by the CORBA persistent state service, which hides many of the details of the underlying persistence mechanism from the component developer.

Self-managed persistence is selected using CIDL and tailored using XML at deployment time to specify connections to specific persistence providers and persistent stores.

10.3.10 Application Operation Invocation

The application operations of a component can be specified on both the component's supported interfaces and the provided interfaces. These operations are normal CORBA object invocations.

Application operations may raise exceptions, both application exceptions (i.e., those defined as part of the IDL interface definition) and system exceptions (those that are not). Exceptions defined as part of the IDL interfaces defined for a component (that includes both provided interfaces and supported interfaces) are raised back to the client directly and do not affect the current transaction. All other exceptions raised by the application are intercepted by the container which then raises the TRANSACTION_ROLLEDBACK exception to the client, if a transaction is active. Otherwise they are reported back to the client directly.

10.3.11 Component Implementations

A component implementation consists of one or more **executors**. Each **executor** describes the implementation characteristics of a particular component segment. The session container API type consists of a single **executor** with a single segment that is activated in response to an operation request on any component facet. The entity container API type can be made up of multiple segments, each of which is associated with a different abstract state declaration. Each segment is independently activated when an operation request on a facet associated with that segment is received.

10.3.12 Component Levels

The CORBA component specification defines two levels of component function that can be used by component developers and supported by CORBA container providers:

- **basic** - The basic CORBA component supports a single interface (or multiple interfaces related by inheritance) and does not define any ports (provided interfaces or event source/sinks). The implementation of a basic component may use transaction, security, and simple persistence (i.e., a single segment) and relies on its container to manage the construction of CORBA object references.

The basic component is functionally equivalent to the EJB 1.1 Component Architecture.

- **extended** - The extended component is a basic component with multiple ports (supported interfaces, provided interfaces and/or event source/sinks). The implementation of the extended component may use all basic function, advanced persistence (multiple segments) plus the event model and participates in the construction of component object references.

The component interfaces defined in this specification have been structured into functional modules corresponding to the two levels of components defined above.

- Basic container APIs are defined in “Server Programming Interfaces - Basic Components” on page 150.
- Extended container APIs are defined in “Server Programming Interfaces - Extended Components” on page 160.

Partitioning the component function into two discrete packages permits the EJB 1.1 APIs to be used to implement basic CORBA components in Java. It also supports the construction of CORBA components in any supported CORBA language that can be accessed by EJB clients. This is described further in the “Integrating with Enterprise JavaBeans” clause.

10.3.13 Component Categories

As indicated in “Component Categories” on page 137, this specification defines four component categories whose behavior is specified by the two **container API types**. Additionally we reserve a component category to describe the empty container (i.e., a container API type that does not use one of the API frameworks defined in this specification). The four component categories are described briefly in the following sub clauses. The component categories are independent of the component levels defined in “Component Levels” on page 146.

10.3.13.1 The Service Component

The **service** component is a CORBA component with the following properties:

- no state
- no identity
- behavior

The lifespan of a **service** component is equivalent to the lifetime of a single operation request (i.e., **method**) so it is useful for functions such as command objects that have no duration beyond the lifetime of a single client interaction with them. A service component can also be compared to a traditional TP monitor program like a Tuxedo service or a CICS transaction. A service component provides a simple way of wrapping existing procedural applications.

A service component is equivalent to a stateless EJB session bean.

Table 10.6 summarizes the characteristics of a service component as seen by the server programmer.

Table 10.6 - Service Component Design Characteristics

Design Characteristic	Property
External Interfaces	As defined in the component IDL
Internal Interfaces	Base Set plus SessionContext (basic) Session2Context (extended)
Callback Interfaces	SessionComponent
CORBA Usage Model	stateless
External API Types	keyless
Client Design Pattern	Factory
Persistence	No
Servant Lifetime Policy	method
Transactions	May use, but not included in current transaction
Events	Transactional or Non-transactional
Executor	Single segment with a single servant and no managed storage

Because of its absence of state, any programming language servant can service any **ObjectId**, enabling such servants to be managed as a pool or dynamically created as required, depending on usage patterns. Because a service component has no identity, **ObjectIds** can be managed by the POA, not the component implementor, and the client sees only the factory design pattern. The service component can use either container-managed or self-managed transactions.

10.3.13.2 The Session Component

The **session** component is a CORBA component with the following properties:

- transient state
- identity (which is not persistent)
- behavior

The lifespan of a **session** component is specified using the servant lifetime policies defined in “Servant Lifetime Management” on page 140. A session component (with a **transaction** lifetime policy) is similar to an MTS component and is useful for modeling things like iterators, which require transient state for the lifetime of a client interaction but no persistent store. A session component is equivalent to the stateful session bean found in EJB.

Table 10.7 summarizes the characteristics of a session component as seen by the server programmer.

Table 10.7 - Session Component Design Characteristics

Design Characteristic	Property
External Interfaces	As defined in the component IDL
Internal Interfaces	Base Set plus SessionContext (basic) Session2Context (extended)
Callback Interfaces	SessionComponent plus (optionally) SessionSynchronization
CORBA usage model	conversational
Client Design Pattern	Factory
External API Types	keyless
Persistence	No
Servant Lifetime Policy	Any
Transactions	May use, but not included in current transaction
Events	Transactional or Non-transactional
Executor	Single segment with a single servant and no managed storage

A programming language servant is allocated to an **ObjectId** for the duration of the servant lifetime policy specified. At that point, the servant can be returned to a pool and re-used for a different **ObjectId**. Alternatively, servants may be dynamically created as required, depending on usage patterns. Because a session component has no persistent identity, **ObjectIds** can be managed by the container, however extended components may choose to participate in creating references if desired, and the client sees only the factory design pattern. The session component shall use either container-managed or self-managed transactions.

10.3.13.3 The Process Component

The process component is a CORBA component with the following properties:

- Persistent state, which is not visible to the client and is managed by the **process** component implementation or the container.

- Persistent identity, which is managed by the **process** component and can be made visible to the client only through user-defined operations.
- Behavior, which may be transactional.

The process component is intended to model objects that represent business processes (e.g., applying for a loan, creating an order, etc.) rather than entities (e.g., customers, accounts, etc.). The major difference between **process** components and **entity** components is that the **process** component does not expose its persistent identity to the client (except through user-defined operations).

Table 10.8 summarizes the characteristics of process component as seen by the server programmer.

Table 10.8 - Process Component Design Characteristics

Design Characteristic	Property
External Interfaces	As defined in component IDL
Internal Interfaces	Base set plus EntityContext (basic) Entity2Context (extended)
Callback Interfaces	EntityComponent
CORBA usage model	durable
Client Design Pattern	Factory
External API Types	keyless
Persistence	Self-managed with or without PSS or Container-managed with or without PSS
Servant Lifetime Policy	Any
Transactions	May use, and can be included in current transaction
Events	Non-transactional or transactional events
Executor	Multiple segments with associated managed storage

A process component may have transactional behavior. The container will interact with the CORBA transaction service to participate in the commit process. The process component shall use container-managed transactions. This is identical to the EJB restriction for Entity Beans.

The process component can use **container-managed** or **self-managed** persistence using either the CORBA persistent state service or a user-defined persistence mechanism. The implications of the various choices are described in “Persistence” on page 144. The entity container uses callback interfaces, which enable the process component’s implementation to retrieve and save state data at activation and passivation respectively.

10.3.13.4 The Entity Component

The **entity** component is a CORBA component with the following properties:

- Persistent state, which is visible to the client and is managed by the **entity** component implementation or the container.
- Identity, which is architecturally visible to its clients through a **primarykey** declaration.
- Behavior, which may be transactional.

As a fundamental part of the architecture, **entity** components expose their persistent state to the client as a result of declaring a **primarykey** value on their home declaration. The entity component may be used to implement the entity bean in EJB.

Table 10.9 summarizes the characteristics of **entity** component as seen by the server programmer:

Table 10.9 - Entity Component Design Characteristics

Design Characteristic	Property
External Interfaces	As defined in the component IDL
Internal Interfaces	Base set plus EntityContext (basic) Entity2Context (extended)
Callback Interfaces	EntityComponent
CORBA usage model	durable
Client Design Pattern	Factory or Finder
External API Types	keyfull
Persistence	Self-managed with or without PSS or Container-managed with or without PSS
Servant Lifetime Policy	Any
Transactions	May use, and can be included in current transaction
Events	Non-transactional or transactional events
Executor	Multiple segments with associated managed storage

The entity component shall use container-managed transactions. The container shall interact with the CORBA transaction service to participate in the commit process. This is identical to the EJB restriction for Entity Beans.

The entity component can use **container-managed** or **self-managed** persistence using either the CORBA persistent state service or a user-defined persistence mechanism. The implications of the various choices are described in “Persistence” on page 144. The entity container uses callback interfaces that enable the entity component’s implementation to retrieve and save state data at activation and passivation, respectively.

10.4 Server Programming Interfaces - Basic Components

This sub clause defines the local interfaces used and provided by the component developer for basic components. These interfaces are then grouped as follows:

- Interfaces common to both container API types.
- Interfaces supported by the session container API type only.
- Interfaces supported by the entity container API type only.

Unless otherwise indicated, all of these interfaces are defined within the **Components** module.

10.4.1 Component Interfaces

All components deal with three sets of interfaces:

- **Internal** interfaces that are used by the component developer and provided by the container to assist in the implementation of the component's behavior.
- **External** interfaces that are used by the client and implemented by the component developer.
- **Callback** interfaces that are used by the container and implemented by the component, either in generated code or directly, in order for the component to be deployed in the container.

A container API type defines a base set of internal interfaces which the component developers use in their implementation. These interfaces are then augmented by others that are unique to the component category being developed.

- **CCMContext** - serves as a bootstrap and provides accessors to the other internal interfaces including access to the runtime services implemented by the container.

Each container API type has its own specialization of **CCMContext**, which we refer to as a context.

- **UserTransaction** - wraps the demarcation subset of the CORBA transaction service required by the application developer.
- **EnterpriseComponent** - the base class that all **callback** interfaces derive from.

All components implement a callback interface that is determined by the component category. It serves the same role as `EnterpriseBean` in EJB.

When a component instance is instantiated in a container, it is passed a reference to its context, a local interface used to invoke services. For basic components, these services include transactions and security. The component uses this reference to invoke operations required by the implementation at runtime beyond what is specified in its deployment descriptor.

10.4.2 Interfaces Common to both Container API Types

This sub clause describes the interfaces and operations provided by both **container API types** to support all categories of CORBA components.

10.4.2.1 The CCMContext Interface

The **CCMContext** is an **internal** interface that provides a component instance with access to the common container-provided runtime services applicable to both **container API types**. It serves as a “bootstrap” to the various services the container provides for the component.

The **CCMContext** provides the component access to the various services provided by the container. It enables the component to simply obtain all the references it may require to implement its behavior.

```
typedef SecurityLevel2::Credentials Principal; exception IllegalState { };
```

```
local interface CCMContext {
    Principal get_caller_principal();
    CCMHome get_CCM_home();
    boolean get_rollback_only() raises (IllegalState);
    Transaction::UserTransaction get_user_transaction()
        raises (IllegalState);
    boolean is_caller_in_role (in string role);
```

```
void set_rollback_only() raises (IllegalState);  
};
```

get_caller_principal

The **get_caller_principal** operation obtains the CORBA security credentials in effect for the caller. Security on the server is primarily controlled by the security policy in the deployment descriptor for this component. The component may use this operation to determine the credentials associated with the current client invocation.

get_CCM_home

The **get_CCM_home** operation is used to obtain a reference to the home interface. The home is the interface that supports factory and finder operations for the component and is defined by the **home** declaration in component IDL.

get_rollback_only

The **get_rollback_only** operation is used by a component to test if the current transaction has been marked for rollback. The **get_rollback_only** operation returns **TRUE** if the transaction has been marked for rollback, otherwise it returns **FALSE**. If no transaction is active, the **IllegalState** exception shall be raised. When **get_rollback_only** is issued by a component, it results in a **CosTransaction::Current::get_status** being issued to the CORBA transaction service and the **status** value returned being tested for the **MARKED_ROLLBACK** state.

get_user_transaction

The **get_user_transaction** operation is used to access the **Transaction::UserTransaction** interface. The **UserTransaction** interface is used to implement self-managed transactions. The **IllegalState** exception shall be raised if this component is using container-managed transactions.

is_caller_in_role

The **is_caller_in_role** operation is used by the CORBA component to compare the current credentials to the credentials defined by the role parameter. If they match, **TRUE** is returned. If not, **FALSE** is returned.

set_rollback_only

The **set_rollback_only** operation is used by a component to mark an existing transaction for abnormal termination. If no transaction is active, the **IllegalState** exception shall be raised. When **set_rollback_only** is issued by a component, it results in a **CosTransaction::Current::rollback_only** being issued to the CORBA transaction service. The rules for the use of this operation are equivalent to the rules of its corresponding CORBA transaction service operation.

10.4.2.2 The Home Interface

A home is an **external** interface that supports factory and finder operations for the component. These operations are generated from the **home** IDL declaration (see “Homes” on page 34). The context supports an operation (**get_CCM_home**) to obtain a reference to the component’s home interface.

10.4.2.3 The UserTransaction Interface

A CORBA component may use either container-managed or self-managed transactions, depending on the component category. With container-managed transactions, the component implementation relies on the transaction policy declarations packaged with the deployment descriptor and contains no transaction APIs in its implementation code.

This is identical to container-managed transactions in EJB or the default processing of an MTS component.

A component specifying self-managed transactions may use the CORBA transaction service directly to manipulate the current transaction or it may choose to use a simpler API, defined by this specification, which exposes only those transaction demarcation functions needed by the component implementation.

Manipulation of the current transaction shall be consistent between the client, the transaction policy specified in the deployment descriptor, and the component implementation.

For example, if the client or the container starts a transaction, the component may not end it (commit or rollback). The rules to be used are defined by the CORBA transaction service.

If the component uses the **CosTransactions::Current** interface, all operations defined for **Current** may be used as defined by the CORBA transaction service with the following exceptions:

- The **Control** object returned by **suspend** may only be used with **resume**.
- Operations on **Control** are not supported with CORBA components and may raise the **NO_IMPLEMENT** system exception.

*The **Control** interface in the CORBA transaction service supports accessors to the **Coordinator** and **Terminator** interfaces. The **Coordinator** is used to build object versions of XA resource managers. The **Terminator** is used to allow a transaction to be ended by someone other than the originator. Since neither function is within the scope of the demarcation subset of CORBA transactions used with CORBA components, we allow CORBA transaction services implementations used with CORBA components to raise the **NO_IMPLEMENT** exception. This provides the same level of function as the **bean-managed** transaction policy in Enterprise JavaBeans.*

The **UserTransaction** is an **internal** interface implemented by the container and is defined within its own module, **Transaction**, within the **Components** module (**Components::Transaction**). Because the **UserTransaction** is a wrapper over **CosTransactions::Current**, it is thread specific. The **UserTransaction** exposes a simple demarcation subset of the CORBA transaction service to the component. The context supports an operation (**get_user_transaction**) to obtain a reference to the **UserTransaction** interface. The **UserTransaction** interface is defined by the following IDL.

```
typedef sequence<octet> TranToken;
exception NoTransaction { };
exception NotSupported { };
exception SystemError { };
exception RollbackError { };
exception HeuristicMixed { };
exception HeuristicRollback { };
exception Security { };
exception InvalidToken { };
```

```
enum Status {
    ACTIVE,
    MARKED_ROLLBACK,
    PREPARED,
    COMMITTED,
    ROLLED_BACK,
    NO_TRANSACTION,
    PREPARING,
    COMMITTING,
```

ROLLING_BACK

};

```
local interface UserTransaction {
    void begin () raises (NotSupported, SystemError);
    void commit () raises (RollbackError , NoTransaction,
                          HeuristicMixed, HeuristicRollback,
                          Security, SystemError);
    void rollback () raises (NoTransaction, Security, SystemError);
    void set_rollback_only () raises (NoTransaction, SystemError);
    Status get_status() raises (SystemError);
    void set_timeout (in long to) raises (SystemError);
    TranToken suspend () raises (NoTransaction, SystemError);
    void resume (in TranToken txtoken)
                raises (InvalidToken, SystemError);
};
```

begin

The **begin** operation is used by a component to start a new transaction and associate it with the current thread. When **begin** is issued by a component, it results in a **CosTransaction::Current::begin** with **report_heuristics** set to **TRUE** being issued to the CORBA transaction service. The rules for the use of this operation are equivalent to the rules of its corresponding CORBA transaction service operation. The **NotSupported** exception is returned if it is received from the CORBA transaction service. Since nested transactions are not supported by CORBA component containers, this indicates an attempt to start a new transaction when an existing transaction is active. All other exceptions are converted to the **SystemError** exception.

commit

The **commit** operation is used by a component to terminate an existing transaction normally. When **commit** is issued by a component, it results in a **CosTransaction::Current::commit** being issued to the CORBA transaction service. The rules for the use of this operation are equivalent to the rules of its corresponding CORBA transaction service operation. If no transaction is active, the **NoTransaction** exception shall be raised. If the **TRANSACTION_ROLLEDBACK** system exception is returned, it is converted to the **RollbackError** exception. The **CosTransaction::HeuristicMixed** and **CosTransaction::HeuristicRollback** exceptions are reported as the **HeuristicMixed** and **HeuristicRollback** exceptions respectively. The **NO_PERMISSION** system exception is converted to the **Security** exception. All other exceptions are converted to the **SystemError** exception.

rollback

The **rollback** operation is used by a component to terminate an existing transaction abnormally. When **rollback** is issued by a component, it results in a **CosTransaction::Current::rollback** being issued to the CORBA transaction service. The rules for the use of this operation are equivalent to the rules of its corresponding CORBA transaction service operation. If no transaction is active, the **NoTransaction** exception shall be raised. The **NO_PERMISSION** system exception is converted to the **Security** exception. All other exceptions are converted to the **SystemError** exception.

set_rollback_only

The **set_rollback_only** operation is used by a component to mark an existing transaction for abnormal termination. When **set_rollback_only** is issued by a component, it results in a **CosTransaction::Current::rollback_only** being issued to the CORBA transaction service. The rules for the use of this operation are equivalent to the rules of its corresponding CORBA transaction service operation. If no transaction is active, the **NoTransaction** exception shall be raised. All other exceptions shall be converted to the **SystemError** exception.

get_status

The **get_status** operation is used by a component to determine the status of the current transaction. If no transaction is active, it returns the **NoTransaction** status value. Otherwise it returns the state of the current transaction. When **get_status** is issued by a component, it results in a **CosTransaction::Current::get_status** being issued to the CORBA transaction service. The status values returned by this operation are equivalent to the status values of its corresponding CORBA transaction service operation. All exceptions shall be converted to the **SystemError** exception.

set_timeout

The **set_timeout** operation is used by a component to associate a time-out value with the current transaction. The timeout value (**to**) is specified in seconds. When **set_timeout** is issued by a component, it results in a **CosTransaction::Current::set_timeout** being issued to the CORBA transaction service. The rules for the use of this operation are equivalent to the rules of its corresponding CORBA transaction service operation. All exceptions are converted to the **SystemError** exception.

suspend

The **suspend** operation is used by a component to disconnect an existing transaction from the current thread. The **suspend** operation returns a **TranToken**, which can only be used in a subsequent **resume** operation. When **suspend** is issued by a component, it results in a **CosTransaction::Current::suspend** being issued to the CORBA transaction service. The rules for the use of this operation are more restrictive than the rules of its corresponding CORBA transaction service operation:

- Only one transaction may be suspended.
- The suspended transaction is the only transaction that may be resumed.

If no transaction is active, the **NoTransaction** exception shall be raised. All other exceptions are converted to the **SystemError** exception.

resume

The **resume** operation is used by a component to reconnect a transaction previously suspended to the current thread. The **TranToken** identifies the suspended transaction that is to be resumed. If the transaction identified by **TranToken** has not been suspended, the **InvalidToken** exception shall be raised. When **resume** is issued by a component, it results in a **CosTransaction::Current::resume** being issued to the CORBA transaction service. The rules for the use of this operation are more restrictive than the rules of its corresponding CORBA transaction service operation since the single suspended transaction is the only transaction that may be resumed. All other exceptions are converted to the **SystemError** exception.

*The **UserTransaction** interface is equivalent to the **UserTransaction** interface (`javax.transaction.UserTransaction`) in EJB with the addition of the **suspend** and **resume** operations.*

10.4.2.4 The EnterpriseComponent Interface

All CORBA components must implement an interface derived from the **EnterpriseComponent** interface to be housed in a component container. **EnterpriseComponent** is a **callback** interface that defines no operations.

```
local interface EnterpriseComponent { };
```

10.4.3 Interfaces Supported by the Session Container API Type

This sub clause describes the interfaces supported by the session container API type. This includes both **internal** interfaces provided by the container and **callback** interfaces, which must be implemented by components deployed in this container API type.

10.4.3.1 The SessionContext Interface

The **SessionContext** is an **internal** interface that provides a component instance with access to the container-provided runtime services. It serves as a “bootstrap” to the various services the container provides for the component. The **SessionContext** enables the component to simply obtain all the references it may require to implement its behavior.

```
exception IllegalState { };
```

```
local interface SessionContext : CCMContext {  
    Object get_CCM_object() raises (IllegalState);  
};
```

get_CCM_object

The **get_CCM_object** operation is used to get the reference used to invoke the component. For basic components, this will always be the component reference. For extended components, this will be a specific facet reference. If this operation is issued outside of the scope of a **callback** operation, the **IllegalState** exception is returned.

10.4.3.2 The SessionComponent Interface

The **SessionComponent** is a **callback** interface implemented by a session CORBA component. It provides operations for disassociating a context with the component and to manage servant lifetimes for a session component.

```
enum CCMExceptionReason {  
    SYSTEM_ERROR,  
    CREATE_ERROR,  
    REMOVE_ERROR,  
    DUPLICATE_KEY,  
    FIND_ERROR,  
    OBJECT_NOT_FOUND,  
    NO_SUCH_ENTITY};
```

```
exception CCMException {CCMExceptionReason reason};
```

```
local interface SessionComponent : EnterpriseComponent {  
    void set_session_context ( in SessionContext ctx)  
        raises (CCMException);
```

```

    void ccm_activate() raises (CCMException);
    void ccm_passivate() raises (CCMException);
    void ccm_remove () raises (CCMException);
};

```

set_session_context

The **set_session_context** operation is used to set the **SessionContext** of the component. The container calls this operation after a component instance has been created. This operation is called outside the scope of an active transaction. The component may raise the **CCMException** with the **SYSTEM_ERROR** minor code to indicate a failure caused by a system level error.

ccm_activate

The **ccm_activate** operation is called by the container to notify a session component that it has been made active. The component instance should perform any initialization required prior to operation invocation. The component may raise the **CCMException** with the **SYSTEM_ERROR** minor code to indicate a failure caused by a system level error.

ccm_passivate

The **ccm_passivate** operation is called by the container to notify a session component that it has been made inactive. The component instance should release any resources it acquired at activation time. The component may raise the **CCMException** with the **SYSTEM_ERROR** minor code to indicate a failure caused by a system level error.

ccm_remove

The **ccm_remove** operation is called by the container when the servant is about to be destroyed. It informs the component that it is about to be destroyed. The component may raise the **CCMException** with the **SYSTEM_ERROR** minor code to indicate a failure caused by a system level error.

10.4.3.3 The SessionSynchronization Interface

The **SessionSynchronization** interface is a **callback** interface that may optionally be implemented by the session component. It permits the component to be notified of transaction boundaries by its container.

```

exception CCMException {CCMExceptionReason reason;};

```

```

local interface SessionSynchronization {
    void after_begin () raises (CCMException);
    void before_completion () raises (CCMException);
    void after_completion (
        in boolean committed) raises (CCMException);
};

```

after_begin

The **after_begin** operation is called by the container to notify a session component that a new transaction has started, and that the subsequent operations will be invoked in the context of the transaction. The component may raise the **CCMException** with the **SYSTEM_ERROR** minor code to indicate a failure caused by a system level error.

before_completion

The **before_completion** operation is called by the container just prior to the start of the two-phase commit protocol. The container implements the **CosTransactions::Synchronization** interface of the CORBA transaction service and invokes the **before_completion** operation on the component before starting its own processing. The component may raise the CCMException with the **SYSTEM_ERROR** minor code to indicate a failure caused by a system level error.

after_completion

The **after_completion** operation is called by the container after the completion of the two-phase commit protocol. If the transaction has committed, the **committed** value is set to **TRUE**. If the transaction has been rolled back, the **committed** value is set to **FALSE**. The container implements the **CosTransactions::Synchronization** interface of the CORBA transaction service and invokes the **after_completion** operation on the component after completing its own processing. The component may raise the CCMException with the **SYSTEM_ERROR** minor code to indicate a failure caused by a system level error.

10.4.4 Interfaces Supported by the Entity Container API Type

This sub clause describes the interfaces supported by the entity container API type. This includes both **internal** interfaces provided by the container and **callback** interfaces that must be implemented by components deployed in this container API type.

10.4.4.1 The EntityContext Interface

The **EntityContext** is an **internal** interface that provides a component instance with access to the container-provided runtime services. It serves as a “bootstrap” to the various services the container provides for the component.

The **EntityContext** enables the component to simply obtain all the references it may require to implement its behavior.

```
exception IllegalState { };
```

```
local interface EntityContext : CCMContext {  
    Object get_CCM_object () raises (IllegalState);  
    PrimaryKeyBase get_primary_key () raises (IllegalState);  
};
```

get_CCM_object

The **get_CCM_object** operation is used to obtain the reference used to invoke the component. For basic components, this will always be the component reference. For extended components, this will be a specific facet reference. If this operation is issued outside of the scope of a **callback** operation, the **IllegalState** exception is returned.

get_primary_key

The **get_primary_key** operation is used by an **entity** component to access the primary key value declared for this component’s home. This operation is equivalent to issuing the same operation on the component’s home interface. If this operation is issued outside of the scope of a **callback** operation, the **IllegalState** exception is returned.

10.4.4.2 The EntityComponent Interface

The **EntityComponent** is a **callback** interface implemented by both process and entity components. It contains operations to manage the persistent state of the component.

Note – As currently defined, any operation request will cause the container to activate the component segment, if required. Since the component reference is well-structured, we could consider the possibility of trapping navigation operations prior to activation and executing them without actually activating the component (or we could leave that to clever implementations).

```
exception CCMException {CCMExceptionReason reason;};
```

```
local interface EntityComponent : EnterpriseComponent {  
    void set_entity_context (in EntityContext ctx)  
        raises (CCMException);  
    void unset_entity_context ()raises (CCMException);  
    void ccm_activate () raises (CCMException);  
    void ccm_load ()raises (CCMException);  
    void ccm_store ()raises (CCMException);  
    void ccm_passivate ()raises (CCMException);  
    void ccm_remove ()raises (CCMException);  
};
```

set_entity_context

The **set_entity_context** operation is used to set the **EntityContext** of the component. The container calls this operation after a component instance has been created. This operation is called outside the scope of an active transaction. The component may raise the **CCMException** with the **SYSTEM_ERROR** minor code to indicate a failure caused by a system level error.

unset_entity_context

The **unset_entity_context** operation is used to remove the **EntityContext** of the component. The container calls this operation just before a component instance is destroyed. This operation is called outside the scope of an active transaction. The component may raise the **CCMException** with the **SYSTEM_ERROR** minor code to indicate a failure caused by a system level error.

ccm_activate

The **ccm_activate** operation is called by the container to notify the component that it has been made active. For most CORBA component implementations, no action is required. The component instance should perform any initialization (other than establishing its state) required prior to operation invocation. This operation is called within an unspecified transaction context. The component may raise the **CCMException** with the **SYSTEM_ERROR** minor code to indicate a failure caused by a system level error.

ccm_load

The **ccm_load** operation is called by the container to instruct the component to synchronize its state by loading it from its underlying persistent store. When container-managed persistence is implemented using the CORBA persistent state service, this operation can be implemented in generated code. If self-managed persistence is being used, the component is

responsible for locating its state in a persistent store. This operation executes within the scope of the current transaction. The component may raise the **CCMException** with the **SYSTEM_ERROR** minor code to indicate a failure caused by a system level error.

ccm_store

The **ccm_store** operation is called by the container to instruct the component to synchronize its state by saving it in its underlying persistent store. When container-managed persistence is implemented using the CORBA persistent state service, this operation can be implemented in generated code. If self-managed persistence is being used, the component is responsible for saving its state in the persistent store. This operation executes within the scope of the current transaction. The component may raise the **CCMException** with the **SYSTEM_ERROR** minor code to indicate a failure caused by a system level error.

ccm_passivate

The **ccm_passivate** operation is called by the container to notify the component that it has been made inactive. For most CORBA component implementations, no action is required. The component instance should perform any termination processing (other than saving its state) required prior to being passivated. This operation is called within an unspecified transaction context. The component may raise the **CCMException** with the **SYSTEM_ERROR** minor code to indicate a failure caused by a system level error.

ccm_remove

The **ccm_remove** operation is called by the container when the servant is about to be destroyed. It informs the component that it is about to be destroyed. This operation is always called outside the scope of a transaction. The component raises the **CCMException** with the **REMOVE_ERROR** minor code if it does not allow the destruction of the component. The component may raise the **CCMException** with the **SYSTEM_ERROR** minor code to indicate a failure caused by a system level error.

*The **EntityComponent** interface is equivalent to the **EntityBean** interface in Enterprise JavaBeans. Container-managed persistence with the CORBA persistent state service supports automatic code generation for **ccm_load** and **ccm_store**. For self-managed persistence, the component implementor provides the **ccm_load** and **ccm_store** methods. Since both process and entity components have persistent state and container-managed persistence, the same callback interfaces can be used.*

10.5 Server Programming Interfaces - Extended Components

This sub clause defines the local interfaces used and provided by the component developer for extended components. These interfaces are grouped as in “Server Programming Interfaces - Basic Components” on page 150. Unless otherwise indicated, all of these interfaces are defined within the **Components** module. Extended components add interfaces in the following areas:

- **CCM2Context** - adds functions unique to extended components.

Each container API type has its own specialization of **CCM2Context** that we refer to as a context. The context for extended components adds accessors to persistence services and supports operations for managing servant lifetime policy, and creating and managing object references in conjunction with the POA.

- **ComponentId** - encapsulates a component identifier, which is an abstract information model used to locate the component’s state.

Only the **entity container API type** supports the **ComponentId** interface.

10.5.1 Interfaces Common to both Container API Types

This sub clause describes the interfaces and operations provided for extended components by both **container API types** to support all categories of CORBA components.

10.5.1.1 The CCM2Context Interface

The **CCM2Context** is an **internal** interface that extends the **CCMContext** interface to provide the extended component instance with access to additional container-provided runtime services applicable to both **container API types**. These services include advanced persistence using the CORBA Persistent State service, and runtime management of component references and servants using the POA. The **CCM2Context** is defined by the following IDL:

```
typedef CosPersistentState::CatalogBase CatalogBase;
typedef CosPersistentState::Typeld Typeld;

exception PolicyMismatch { };
exception PersistenceNotAvailable { };
local interface CCM2Context : CCMContext {
    HomeRegistration get_home_registration ();
    void req_passivate () raises (PolicyMismatch);

    CatalogBase get_persistence (in Typeld catalog_type_id)
        raises (PersistenceNotAvailable);
};
```

get_home_registration

The **get_home_registration** operation is used to obtain a reference to the **HomeRegistration** interface. The **HomeRegistration** is used to register component homes so they may be located by the **HomeFinder**.

req_passivate

The **req_passivate** operation is used by the component to inform the container that it wishes to be passivated when its current operation completes. To be valid, the component must have a servant lifetime policy of **component** or **container**. If not, the **PolicyMismatch** exception shall be raised.

get_persistence

The **get_persistence** operation provides the component access to a persistence framework provided by an implementation of the CORBA Persistence State service. It returns a **CosPersistentState::CatalogBase**, which serves as an index to the available storage homes. The **CatalogBase** is identified by its **CosPersistentState::Typeld catalog_type_id**. If the **CatalogBase** identified by **catalog_type_id** is not available on this container, the **PersistenceNotAvailable** exception shall be raised.

10.5.1.2 The HomeRegistration Interface

The **HomeRegistration** is an **internal** interface that may be used by the CORBA component to register its home so it can be located by a **HomeFinder**.

*The **HomeRegistration** interface allows a component implementation to advertise a home instance that can be used to satisfy a client's **find_home** request. It may also be used by an administrator to do the same thing. It is likely*

*that the combination of **HomeRegistration** and **HomeFinder** interfaces will work within the domain of a single container provider unless multiple implementations use other shareable directory mechanisms (e.g., an LDAP global directory). Federating **HomeFinders** is a similar problem to federating CORBA security domains and we defer to the security people for an architecture for such federation rather than attempting to specify such an architecture in this specification.*

The **HomeRegistration** interface is defined by the following IDL:

```
local interface HomeRegistration {
    void register_home (
        in CCMHome home_ref,
        in string home_name);
    void unregister_home (in CCMHome home_ref);
};
```

register_home

The **register_home** operation is used to register a component home with the **HomeFinder** so it can be located by a component client. The **home_ref** parameter identifies the home being registered and can be used to obtain both the **CORBA::ComponentIR::ComponentDef** (**CCMHome::get_component_def**) and the **CORBA::InterfaceDef** (**CORBA::Object::get_interface_def**) to support both **HomeFinder::find_home_by_component_type** and **HomeFinder::find_home_by_home_type**. The **home_name** parameter identifies an Interoperable Naming Service (INS) name that can be used as input to the **HomeFinder::find_home_by_name** operation. If the **home_name** parameter is NULL, no name is associated with this home so this home cannot be retrieved by name.

unregister_home

The **unregister_home** operation is used to remove a component home from the **HomeFinder**. Once **unregister_home** completes, a client will never be returned a reference to the home specified as being unregistered. The **home_ref** parameter identifies the home being unregistered.

10.5.1.3 The ProxyHomeRegistration Interface

Because CORBA components exploit the dynamic activation features of the POA, it is possible for some component types to provide a home that is not collocated with the component instances it creates. This permits load balancing criteria to be applied in selecting the actual server and POA where this instance will be created. The **ProxyHomeRegistration** is an **internal** interface, derived from **HomeRegistration**, which can be used by the CORBA component to register a remote home (i.e., one that is not collocated with the component) so it can be returned by a **HomeFinder**. The **ProxyHomeRegistration** interface is defined by the following IDL:

```
exception UnknownActualHome { };
exception ProxyHomeNotSupported { };

local interface ProxyHomeRegistration : HomeRegistration {
    void register_proxy_home (
        in CCMHome rhome,
        in CCMHome ahome)
        raises (UnknownActualHome, ProxyHomeNotSupported);
};
```

register_proxy_home

The **register_proxy_home** operation is used to register a component home, not collocated with the instances that it can create, with the **HomeFinder** so the proxy home can be used by component clients. The **rhome** parameter identifies the proxy home being registered. The **ahome** parameter identifies the actual home that the **rhome** is associated with. If the actual home specified by **ahome** is not known, the **UnknownActualHome** exception shall be raised. If this component does not support proxy homes, the **ProxyHomeNotSupported** exception shall be raised. Support for proxy homes is a component implementation option.

10.5.2 Interfaces Supported by the Session Container API Type

This sub clause describes the interfaces supported for extended components by the session container API type. This includes both **internal** interfaces provided by the container and **callback** interfaces, which must be implemented by components deployed in this container API type.

10.5.2.1 The Session2Context Interface

The **Session2Context** is an **internal** interface that extends the **SessionContext** to provide a component instance with access to additional container-provided runtime services for the session container API type. It adds the ability to create references for components deployed in a **session** container API type. The **Session2Context** is defined by the following IDL:

```
enum BadComponentReferenceReason {
    NON_LOCAL_REFERENCE,
    NON_COMPONENT_REFERENCE,
    WRONG_CONTAINER,
};

exception BadComponentReference {
    BadComponentReferenceReason reason;
};

exception IllegalState { };

local interface Session2Context : SessionContext, CCM2Context {
    Object create_ref (in CORBA::RepositoryId repid);
    Object create_ref_from_oid (
        in CORBA::OctetSeq oid,
        in CORBA::RepositoryId repid);
    CORBA::OctetSeq get_oid_from_ref (in Object objref)
        raises (IllegalState, BadComponentReference);
};

create_ref
```

The **create_ref** operation is used to create a reference to be exported to clients to invoke operations. The **repid** parameter identifies the **RepositoryId** associated with the interface for which a reference is being created. Invocations on the new object reference are delivered to the appropriate segment of the component that invokes this operation. The **RepositoryId** must match the **RepositoryId** of the component itself, one of its bases, one of its supported interfaces, or one of its facets.

create_ref_from_oid

The **create_ref_from_oid** operation is used to create a reference to be exported to clients that includes information provided by the component which it can use on subsequent operation requests. The **oid** parameter identifies the **ObjectSeq** to be encapsulated in the reference and the **repid** parameter identifies the **RepositoryId** associated with the interface for which a reference is being created.

get_oid_from_ref

The **get_oid_from_ref** operation is used by the component to extract the **oid** encapsulated in the reference. The **objref** parameter specifies the reference that contains the **oid**. This operation must be called within an operation invocation. If not, the **IllegalState** exception shall be raised. If the reference was not created by this container, the **BadComponentReference** with the **WRONG_CONTAINER** minor code is raised.

10.5.3 Interfaces Supported by the Entity Container API Type

This sub clause describes the interfaces provided for extended components by the entity container API type. This includes both **internal** interfaces provided by the container and **callback** interfaces, which must be implemented by components deployed in this container API type.

10.5.3.1 Component Identifiers

The **ComponentId** interface is an **internal** interface provided by the entity container API type through which the component implementation and the container exchange identity information, referred to as *component identifiers*. The **ComponentId** interface encapsulates a component identifier, which is an abstract information model. The **ComponentId** interface is used in the following ways:

- Component implementations (usually home executor implementations) create component identifiers to describe new components, and to create object references that encapsulate the provided description. The **Entity2Context** interface acts as a factory for component identifiers and as the factory for object references.
- The container encodes the information encapsulated by the component identifier in the object identifier value it uses internally to create the object reference on the encapsulated POA. The encoding is not specified, since a container's choice of encoding does not affect interoperability or portability.
- While dispatching an incoming request, the container extracts and decodes the component identifier from the **ObjectId**. The extracted component identifier is made available to the component executor through the context before the request is dispatched to the component.
- When the container invokes **ccm_load** in the component executor, the implementation of **ccm_load** uses the contents of the component identifier to locate and incarnate the required component state.

In the following discussions, component identifiers, and component object references are sometimes used as though the terms were synonymous. Since there is a one-to-one relationship between a component identifier and an object reference created from the component identifier, this discussion occasionally uses the term “component reference” to mean “the component reference created from the component identifier in question,” for the sake of brevity.

The **ComponentId** interface does not explicitly specify the state representation it encapsulates. The abstract state is implied by the interface and reflects the structure of the executor it describes (see the *CCM Implementation Framework* clause for a complete discussion of executor structure).

A component identifier encapsulates the following information:

- A *facet identifier* value denoting the target facet of the component reference.
- A *segment identifier* value denoting the target segment of the component reference (i.e., the segment that supports the target facet).
- A sequence of *segment descriptors*.

A segment descriptor includes the following:

- A segment identifier denotes the segment being described, and
- a *state identifier* value that denotes the persistent state of the segment in some storage mechanism.

A monolithic executor is represented as a degenerate case of the generalized component identifier, where the target segment identifier is set to zero and the sequence of segment descriptors contains a single element, whose segment identifier is zero and whose state identifier denotes the persistent state of the component's single segment.

The facet identifier value zero is reserved to denote the component facet; that is, the facet that supports the component equivalent interface. The segment identifier value zero is reserved to denote the segment that supports the component facet. For monolithic executors, the segment identifier value is always zero.

State identifier is an abstraction that generalizes a variety of possible state identity schemes. This specification provides a mechanism for describing state identifiers that can be extended by component implementors, allowing customization for storage mechanisms that do not support the standard persistence interfaces.

The **ComponentId** local interface and supporting constructs are defined by the following IDL:

```

typedef short SegmentId;
const SegmentId COMPONENT_SEGMENT = 0;

typedef short FacetId;
const FacetId COMPONENT_FACET = 0;

typedef sequence<octet> IdData;
typedef CosPersistentState::Pid PersistentId;

exception InvalidStateIdData {};

typedef short StateIdType;
const StateIdType PERSISTENT_ID = 0;

abstract valuetype StateIdValue {
    StateIdType get_sid_type();
    IdData get_sid_data();
};

local interface StateIdFactory {
    StateIdValue create (in IdData data) raises (InvalidStateIdData);
};

valuetype PersistentIdValue : StateIdValue {
    private PersistentId pid;
    PersistentId get_pid();
};

```

```

    factory init (in PersistentId pid);
};

valuetype SegmentDescr {
    private StateIdValue sid;
    private SegmentId seg;
    StateIdValue get_sid();
    SegmentId get_seg_id();
    factory init (in StateIdValue sid, in SegmentId seg);
};

typedef sequence<SegmentDescr> SegmentDescrSeq;

local interface ComponentId {
    FacetId get_target_facet();
    SegmentId get_target_segment();
    StateIdValue get_target_state_id (in StateIdFactory sid_factory)
        raises (InvalidStateIdData);
    StateIdValue get_segment_state_id (
        in SegmentId seg,
        in StateIdFactory sid_factory)
        raises (InvalidStateIdData);
    ComponentId create_with_new_target (
        in FacetId new_target_facet,
        in SegmentId new_target_segment);
    SegmentDescrSeq get_segment_descrs (
        in StateIdFactory sid_factory)
        raises (InvalidStateIdData);
};

```

10.5.3.2 StateIdValue abstract valuetype

The **StateIdValue** type is the base valuetype for concrete, storage-specific state identity values. The container interacts with state identities completely in terms of this interface. A single pre-defined concrete value type derived from **StateIdValue** is provided for **PersistentId** state identities. Component implementors, or suppliers of storage mechanisms that do not support the CORBA component persistence model can provide their own state identity types by deriving from **StateIdValue** and implementing the required behaviors properly.

get_sid_type

The **get_sid_type** operation returns a discriminator (physically, a short) that identifies the type of the state identity encapsulated by the **StateIdValue**. This specification defines the value zero (0) to denote a **Components::Extended::PersistentId** state identifier.

get_sid_data

The **get_sid_data** operation returns the encapsulated state identity expressed in a canonical form, as a sequence of octets. The implementation of the derived concrete value type is responsible for converting its encapsulated data into this form, and for supplying a factory that can construct an instance of the concrete type from an **IdData** value (a sequence of octets).

10.5.3.3 StateIdFactory Interface

StateIdFactory is the abstract base interface for factories of state identity values derived from **StateIdValue**. An implementation of **StateIdFactory** must be supplied with the implementation of a concrete state identity type. If the **IdData** octet sequence provided in the **data** parameter cannot be decoded to create a proper instance of the expected state identity concrete type, the operation raises an **InvalidStateIdData** exception.

create

The **create** operation constructs an instance of a concrete state identifier from the octet sequence parameter. This operation performs the inverse of the transformation performed by the **get_sid_data**.

10.5.3.4 PersistentIdValue valuetype

The **PersistentIdValue** type is a specialization of **StateIdValue** that encapsulates a **PersistentId** value for inclusion in a component identifier.

get_pid

The **get_pid** operation returns the **PersistentId** value encapsulated by the value type.

init

The initializer for **PersistentIdValue** creates an instance of the valuetype that encapsulates the **PersistentId** value passed as a parameter.

get_sid_value

The implementation of **get_sid_value** for **PersistentIdValue** performs no transformation on the encapsulated **PersistentId** value. The sequence of octets returned by **get_sid_value** is identical to the encapsulated **PersistentId** value.

10.5.3.5 SegmentDescr valuetype

The **SegmentDescr** type describes an executor segment, encapsulating a segment identifier and a state identifier. A component identifier for a segmented executor encapsulates a sequence of **SegmentDescr** instances.

get_sid

The **get_sid** operation returns the state identity value of the segment being described.

get_seg_id

The **get_seg_id** operation returns the segment identifier of the segment being described.

init

This initializer sets the value of the encapsulated segment identifier and state identifier to the values of the respective parameters.

10.5.3.6 ComponentId Interface

The **ComponentId** interface encapsulates a complete component identity. Instances of **ComponentId** can only be created by the **Entity2Context** interface, which is supplied by the container, or by duplicating an existing component identifier with a new target value, with **ComponentId::create_with_new_target**. Instances of **ComponentId** are also provided by the **EntityContext** interface in the context of a CORBA invocation. The value of the component identifier provided by the **Entity2Context** shall be identical to the component identifier value used to create the object reference on which the invocation was made. The **ComponentId** interface is a read-only interface. Once a component identifier is constructed by the **create_component_id** operation or constructed internally and provided through the **Entity2Context** interface, the value of the component identifier cannot be altered.

get_target_facet

The **get_target_facet** operation returns the facet identifier of the facet, which is the target of the component reference; that is, the target of requests made on the component reference.

get_target_segment

The **get_target_segment** operation returns the segment identifier of the target segment; that is, the segments that provide the target facet.

get_target_state_id

The **get_target_state_id** operation returns the state identifier of the target segment. The **StateIdFactory** specified in the **sid_factory** parameter is used by the implementation of **get_target_state_id** to construct the proper state identifier from the octet sequence encapsulated by the component identifier. If the state identifier of the target segment is a **PersistentIdValue**, the **sid_factory** parameter may be nil. Container implementations shall provide a default implementation of **StateIdFactory** to be used when the encapsulated state identifier value is a **PersistentIdValue**. If provided (or default) factory cannot construct a correct state identifier of the expected type from the undecoded octet sequence encapsulated by the component identifier, the operation raises an **InvalidStateIdData** exception.

get_segment_state_id

The **get_segment_state_id** operation returns the state identifier of the segment specified by the **seg** parameter. The semantics are otherwise identical to **get_target_state_id**, with respect to the meaning and use of the **sid_factory** parameter.

get_segment_descrs

The **get_segment_descrs** operation returns a sequence containing all of the segment descriptors encapsulated by the component identifier. The sequence is a copy of the encapsulated sequence. The state identifier factory in the **sid_factory** parameter (or the default) is used by the implementation of **get_segment_descrs** to construct state identifiers of the appropriate concrete subtype of **StateIdValue**. If provided (or default) factory cannot construct a correct state identifier of the expected type from the undecoded octet sequence encapsulated by the component identifier, the operation raises an **InvalidStateIdData** exception.

create_with_new_target

The **create_with_new_target** operation creates a new component identifier that is identical to the target component identifier, except that the target facet and target segment values are replaced with the values of the **new_target_facet** and **new_target_segment** parameters, respectively.

This operation is intended primarily to be used in implementing navigation operations.

10.5.3.7 The Entity2Context Interface

The **Entity2Context** is an **internal** interface that extends the **EntityContext** interface to provide the extended component with access to additional container-provided runtime services for managing object references and advanced persistence. Object references for components deployed in an entity container API type can choose to use the CORBA Persistent State service or some user defined persistence mechanism. The **ComponentId** interface (defined in “ComponentId Interface” on page 168) encapsulates this distinction when a reference is to be used. The **Entity2Context** is defined by the following IDL.

```
exception BadComponentReference {
    BadComponentReferenceReason reason; };
exception IllegalState { };

local interface Entity2Context : EntityContext, CCM2Context {
    ComponentId get_component_id ()
        raises (IllegalState);
    ComponentId create_component_id (
        in FacetId target_facet,
        in SegmentId target_segment,
        in SegmentDescrSeq seq_descrs);
    ComponentId create_monolithic_component_id (
        in FacetId target_facet,
        in StateIdValue sid);
    Object create_ref_from_cid (
        in CORBA::RepositoryId repid,
        in ComponentId cid);
    ComponentId get_cid_from_ref (
        in Object objref) raises (BadComponentReference);
};
```

get_component_id

The **get_component_id** operation is used to obtain a reference to the **ComponentId** interface. The **ComponentId** interface encapsulates a persistence identifier that can be used to access the component’s persistence state. If this operation is issued outside of the scope of a **callback** operation, the **IllegalState** exception is returned.

create_component_id

The **create_component_id** operation creates a component identifier value, initializing it with the values specified in the parameters. The **target_facet** parameter contains the facet identifier of the target facet, the **target_segment** parameter contains the segment identifier of the target segment, and the **seq_descrs** parameter contains a sequence of segment descriptors describing all of the segments that constitute the component executor.

create_monolithic_component_id

The **create_monolithic_component_id** operation provides a simplified signature for creating a component identifier value for monolithic executors, which have a single segment. The **target_facet** parameter contains the facet identifier of the target facet, and the **sid** parameter contains the state identifier for the single executor segment. The target segment

identifier encapsulated by the component identifier is set to zero, and the sequence of segment descriptors encapsulated by the component identifier has a single element, initialized with segment identifier value zero, and state identifier value specified by the **sid** parameter.

create_ref_from_cid

The **create_ref_from_cid** operation is used by a component factory to create an object reference that can be exported to clients. The **cid** parameter specifies the **ComponentId** value to be placed in the object reference and made available (using the **get_component_id** operation on the context) when the **EntityComponent callback** operations are invoked. The **repid** parameter identifies the **RepositoryId** associated with the interface for which a reference is being created.

get_cid_from_ref

The **get_cid_from_ref** operation is used by a persistent component to retrieve the **ComponentId** encapsulated in the reference (**objref**). The **ComponentId** interface supports operations to locate the state in some persistent store. The **BadComponentReference** exception can be raised if the input reference is not local (**NON_LOCAL_REFERENCE**), not a component reference (**NON_COMPONENT_REFERENCE**), or created by some other container (**WRONG_CONTAINER**).

*The **ComponentId** structure is dependent on the home implementation and the container; in particular, its implementation of the **Entity2Context** interface. It is likely that a **ComponentId** created by one container will not be understandable by another, hence the possibility of the **WRONG_CONTAINER** exception.*

10.6 The Client Programming Model

This sub clause describes the architecture of the component programming model as seen by the client programmer. The client programming model as defined by IDL extensions has been described previously (see the Component Model clause). This sub clause focuses on the use of standard CORBA by the client who wishes to communicate with a CORBA component implemented in a **Component Server**. It enables a CORBA client, which is not itself a CORBA component, to communicate with a CORBA component.

The client interacts with a CORBA component through two forms of external interfaces - a **home** interface and one or more **application** interfaces. Home interfaces support operations that allow the client to obtain references to an application interface which the component implements.

From the client's perspective, the home supports two design patterns - factories for creating new objects and finders for existing objects. These are distinguished by the presence of a **primarykey** parameter in the home IDL.

- If a **primarykey** is defined, the home supports both factories and finders and the client may use both.
- If a **primarykey** is not defined, the home supports only the factory design pattern and the client must create new instances.

Two forms of clients are supported by the CORBA component model:

- Component-aware clients - These clients know they are making requests against a component (as opposed to an ordinary CORBA object) and can therefore avail themselves of unique component function; for example, navigation among multiple interfaces and component type factories.
- Component-unaware clients - These clients do not know that the interface they are making requests against is implemented by a CORBA component so they can only invoke functions supported by an ordinary CORBA object; for example, looking up a name in a Naming or Trader service, searching for a particular type of factory using a factory finder, etc.

10.6.1 Component-aware Clients

Clients that are defined using the IDL extensions in the Component Model clause are referred to as **component-aware** clients. Such clients can avail themselves of the unique features of CORBA components that are not supported by ordinary CORBA objects. The interaction between these clients and a CORBA component are outlined in the following sub clauses. A **component-aware** client interacts with a component through one or more CORBA interfaces:

- The equivalent interface implied by the **component** IDL declaration.
- Zero or more supported interfaces declared on the **component** specification.
- Zero or more interfaces defined by the **provides** clauses in the **component** definition.
- The home interface that supports factory and finder operations.

Furthermore a component-aware client locates those interfaces using the **Components::HomeFinder** or a naming service. The starting point for client interactions with the component is the **resolve_initial_references** operation on **CORBA::ORB** that provides the initial set of object references.

10.6.1.1 Initial References

Initial references for all services used by a component client are obtained using the **CORBA::ORB::resolve_initial_references** operation. This operation currently supports the following references required by a component client:

- Name Service (“**NameService**”)
- Transaction Current (“**TransactionCurrent**”)
- Security Current (“**SecurityCurrent**”)
- Notification Service (“**NotificationService**”)
- Interface Repository (“**InterfaceRepository**”) for DII clients
- Home Finder (“**ComponentHomeFinder**”)

The client uses ComponentHomeFinder (defined in “Home Finders” on page 42) to obtain a reference to the **HomeFinder** interface.

10.6.1.2 Factory Design Pattern

For factory operations, the client invokes a **create** operation on the home. Default create operations are defined for each category of CORBA components for which code can be automatically generated. These operations return an object of type **CORBA::Component** that must be narrowed to the specific type. Alternatively, the component designer may specify custom factories as part of the **component** definition to define a type-specific signature for the **create** operation. Because these operations are defined in IDL, operation names can be chosen by the component designer. All that is required is that the operations return an object of the appropriate type.

A client using the factory design pattern uses the **HomeFinder** to locate the component factory (**CCMHome**) by interface type. The **HomeFinder** returns a type-specific factory reference, which can then be used to create new instances of the component interface. Once created, the client makes operation requests on the reference representing the interface. This is illustrated by the following code fragment:

```

// Resolve HomeFinder
org.omg.CORBA.Object objref = orb.resolve_initial_references("ComponentHomeFinder");

ComponentHomeFinder ff = ComponentHomeFinderHelper.narrow(objref);

org.omg.CORBA.Object of = ff.find_home_by_type(AHomeHelper.id());

AHome F = AHomeHelper.narrow (of);
org.omg.Components.ComponentBase AInst = F.create();
A Areal = AHelper.narrow (AInst);

// Invoke Application Operation
answer = A.foo(input);

```

10.6.1.3 Finder Design Pattern

A component-aware client wishing to use an existing component instance (rather than create a new instance) uses a **finder** operation. Finders are supported for entity components only. Client's may use the **HomeFinder** as described in "Home Finders" on page 42 to locate the component's home or they may use CORBA naming to look up a specific instance of the home by symbolic name.

A client using the finder design pattern uses the **CosNaming::NamingContext** interface to look up a symbolic name. The naming service returns an object reference of the type previously bound. The client then makes operation requests on the reference representing the interface. This is illustrated by the following code fragment:

```

org.omg.CORBA.Object objref = orb.resolve_initial_references("NamingService");

NamingContext ncRef = NamingContextHelper.narrow(objref);

// Resolve the Object Reference in Naming
NameComponent nc = new NameComponent("A","");
NameComponent path[] = {nc};
A aRef = AHelper.narrow(ncRef.resolve(path));

// Invoke Application Operation
answer = A.foo(input);

```

10.6.1.4 Transactions

A component-aware client may optionally define the boundaries of the transaction to be used with CORBA components. If so, it uses the CORBA transaction service to ensure that the active transaction is associated with subsequent operations on the CORBA component.

The client obtains a reference to **CosTransactions::Current** by using the **CORBA::ORB::resolve_initial_references** operation specifying an **ObjectID** of “**TransactionCurrent**.” This permits the client to define the boundaries of the transaction; that is, how many operations will be invoked within the scope of the client’s transaction. All operations defined for **Current** may be used as defined by the CORBA Transaction service with the following exceptions:

- The **Control** object returned by **get_control** and **suspend** may only be used with **resume**.
- Operations on **Control** may raise the **NO_IMPLEMENT** exception with CORBA components.

*The **Control** interface in the CORBA transaction service supports accessors to the **Coordinator** and **Terminator** interfaces. The **Coordinator** is used to build object versions of XA resource managers. The **Terminator** is used to allow a transaction to be ended by someone other than the originator. Since neither function is within the scope of the demarcation subset of CORBA transactions used with CORBA components, we allow CORBA transaction services implementations used with CORBA components to raise the **NO_IMPLEMENT** exception.*

The following code fragment shows a typical usage:

```
org.omg.CORBA.Object objref = orb.resolve_initial_references("TransactionCurrent");

Current txRef = CurrentHelper.narrow(objRef);
txRef.begin();
// Invoke Application Operation
answer = A.foo(input);
txRef.commit();
```

10.6.1.5 Security

A component-aware client uses the existing CORBA security mechanism to manage security for a CORBA component. There are two scenarios possible:

- Use of SSL for establishing client credentials
CORBA security today does not define a standard API for clients to use with SSL to set the credentials that will be used to authorize subsequent requests. The credentials must be set in a way that is proprietary to the client ORB.
- Use of SECIOP by the client ORB.
In this case, CORBA security does define an API and it must be used by the client to establish the credentials to be used to authorize subsequent requests.

Security processing for CORBA components uses a subset of CORBA security. For SECIOP, the client sets the credentials to be used with subsequent operations on the component by using operations on the **SecurityLevel2::PrincipalAuthenticator**. The client obtains a reference to **SecurityLevel2::Current** by using the

CORBA::ORB::resolve_initial_references operation specifying an **ObjectID** of “**SecurityCurrent**.” This permits the client to access the **PrincipalAuthenticator** interface to associate security credentials with subsequent operations. The following code fragment shows a typical usage:

```
org.omg.CORBA.Object objref = orb.resolve_initial_references("SecurityCurrent");

org.omg.SecurityLevel2.PrincipalAuthenticator secRef = org.omg.SecurityLevel2.
PrincipalAuthenticatorHelper.narrow (objRef);

secRef.authenticate(...);

// Invoke Application Operation
answer = A.foo(input);
```

10.6.1.6 Events

Component-aware clients wishing to **emit** or **consume** events use the component APIs defined in the *Component Model* clause. Alternatively, they may use CORBA notification directly and conform to the subset supported by CORBA components (see “Events” on page 174 for details).

10.6.2 Component-unaware Clients

CORBA components can also be used by clients who are unaware that they are making requests against a component. Such clients can see only a single interface (the supported interface of a component) and do not support navigation.

10.6.2.1 Initial References

Component-unaware clients obtain initial references using existing CORBA mechanisms, viz.

CORBA::ORB::resolve_initial_references. It is unlikely, however, that this mechanism would be used to obtain a reference to the **HomeFinder**.

10.6.2.2 Factory Design Pattern

The factory design pattern can be used by component-unaware clients only if the supported interface has application operations defined. This permits existing CORBA objects to be easily converted to CORBA components, transparently to their existing clients. The following techniques can be used:

- The reference to a factory finder (typically the **CosLifeCycle::FactoryFinder**) can be stored in the Naming or Trader service and looked up by the client before creating the instance.
- A reference to the home interface can be obtained from the Naming service.
- The reference to the home interface can be obtained from a Trader service.
- After locating a factory finder, the factory can be located using the existing **find_factories** operation or by using the new **find_factory** operation on the **CosLifeCycle::FactoryFinder** interface.

*The current **CosLifeCycle find_factories** operation returns a sequence of factories to the client requiring the client to choose the one which will create the instance. To allow the server (i.e., the **FactoryFinder**) to make the*

*selection, we also add a new **find_factory** operation to **CosLifeCycle** which allows the server to choose the “best” factory for the client request based on its knowledge of workload, etc.*

A **FactoryFinder** will return an **Object**. A component-unaware client may expect to narrow this to **CosLifeCycle::GenericFactory** and use the generic create operation. For this reason, we allow the default creation operation on home to return a **GenericFactory** interface. This is fully described in “Homes” on page 34.

- A stringified object reference can be retrieved from a file known by the component-unaware client.

Once a reference to the home has been obtained, the client can create component instances and make operation requests on the component. Each component exports at least one IDL interface. A supported interface must be used by the client to invoke the component’s application operations. Provided interfaces cannot be located using the factory design pattern.

10.6.2.3 Finder Design Pattern

A component-unaware client can use CORBA naming to locate an existing **entity** component. Unlike the factory design pattern, the name to be looked up by the client can be either a supported interface or any of the provided interfaces. The following techniques can be used:

- A symbolic name associated with the component’s home can be looked up in a Naming service to make an invocation of the finder operations.
- Alternatively, the reference to the home interface can be obtained from a Trader service.
- The finder operation can be invoked on the **entity** component to return a reference to the client.

10.6.2.4 Transactions

This is the same as component-aware clients (See “Transactions” on page 173). However, the possibility of the **NO_IMPLEMENT** exception being raised for operations on **Control** may have a more serious impact, since the component-unaware client may not be expecting that to happen.

10.6.2.5 Security

This is the same as component-aware clients (See “Security” on page 173).

10.6.2.6 Events

Component-unaware clients wishing to **emit** or **consume** events must use the equivalent CORBA notification interfaces and stay within the subset supported by CORBA components (see “Events” on page 143 for details). This is illustrated by the following code fragment:

```
org.omg.CORBA.Object objref = orb.resolve_initial_references("NotificationService");

org.omg.CosNotifyChannelAdmin.EventChannelFactory evfRef = org.omg.EventChannel-
FactoryHelper.narrow(objRef);

// Create an Event Channel
org.omg.CosNotifyChannelAdmin.EventChannel evcRef = evfRef.create_channel(...);
```

```
// Obtain a SupplierAdmin
org.omg.CosNotifyChannelAdmin.SupplierAdmin publisher = evcRef.new_for_suppliers
(...);

// And a ConsumerProxy
org.omg.CosNotifyComm.ProxyConsumer proxy = pub-
lisher.obtain_notification_push_consumer (...);

// Publish a structured event
proxy.push_structured_event(...);
```


11 Integrating with Enterprise JavaBeans

11.1 Introduction

This clause describes how an Enterprise JavaBeans (EJB) component can be used by CORBA clients, including CORBA components. The EJB will have a CORBA component style remote interface that is described by CORBA IDL (including the component extensions).

This clause also describes how a CORBA component can be used by a Java client, including an Enterprise JavaBeans component. The CORBA component will have an EJB style remote interface that is defined following the Enterprise JavaBeans specification.

The concepts in this clause follow in the same prescription for interworking as laid out in Part 2 - CORBA Interoperability, *Interoperability Architecture* clause where it is discussed as follows.

How interworking can be practically achieved is illustrated in an Interworking Model, shown in Figure 11.1. It shows how an object in Object System B can be mapped and represented to a client in Object System A. From now on, this will be called a B/A mapping. For example, mapping a CORBA Component Model object to be visible to an EJB client is a CCM/EJB mapping.

On the left is a client in object system A, that wants to send a request to a target object in system B, on the right. We refer to the entire conceptual entity that provides the mapping as a bridge. The goal is to map and deliver any request from the client transparently to the target.

To do so, we first provide an object in system A called a View. The View is an object in system A that presents the identity and interface of the target in system B mapped to the vernacular of system A, and is described as an A View of a B target. The View exposes an interface, called the View Interface, which is isomorphic to the target's interface in system B. The methods of the View Interface convert requests from system A clients into requests on the target's interface in system B. The View is a component of the bridge. A bridge may be composed of many Views.

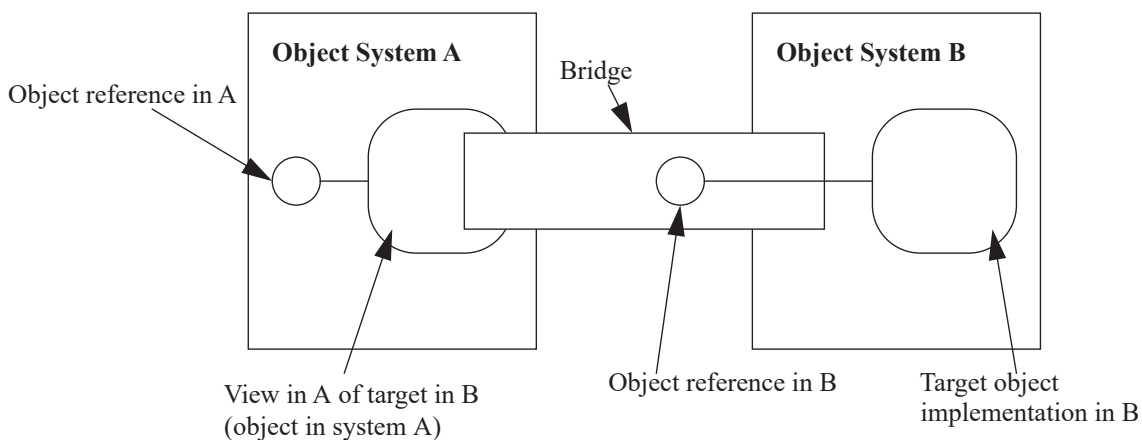


Figure 11.1- B/A Interworking Model

The bridge maps interface and identify forms between different object systems. Conceptually, the bridge holds a reference in B for the target (although this is not physically required). The bridge must provide a point of rendezvous between A and B, and may be implemented using any mechanism that permits communication between the two systems (IPC, RPC, network, shared memory, and so forth) sufficient to preserve all relevant object semantics.

The client treats the View as though it is the real object in system A, and makes the request in the vernacular request form of system A. The request is translated into the vernacular of object system B, and delivered to the target object. The net effect is that a request made on an interface in A is transparently delivered to the intended instance in B.

The Interworking Model works in either direction. For example, if system A is EJB, and system B is CCM, then the View is called the EJB View of the CCM target. The EJB View presents the target's interface to the EJB client. Similarly if system A is CCM and system B is EJB, then the View is called the CCM View of the EJB target. The CCM View presents the target's interface to the CCM client.

11.2 Enterprise JavaBeans Compatibility Objectives and Requirements

The objective is to allow the creation of distributed applications that mix CORBA components running in CORBA component servers with EJB components running in an EJB technology-based server. This objective allows a developer to create an application by reusing existing components of either kind.

This requires development time and runtime translations between the CORBA component and EJB domains provided by mediated bridges. It also requires that:

- A CORBA component view for an EJB complies with the EJB to CORBA mapping specification. In particular, this requires that:
 - An EJB definition be mapped to a CORBA component definition following the Java Language to IDL mapping plus the extensions to that mapping that are specified in this clause.
 - Value objects of one kind (e.g., Keys for EJB) have counterpart value objects of the other kind.
 - CORBA components accessible via **CosNaming** have their EJB views accessible via **JNDI**, and vice versa.
- An EJB view for a CORBA component complies with the EJB specification.

An application is to be built using both EJB and CORBA components deployed in their respective containers. At component development time, EJB components are originally defined in Java and CORBA components are originally defined in IDL. When applications are assembled using both, the application assembly environment will most commonly dictate which model these components must present to developers. During application assembly, developers construct clients (which themselves may be components) that make use of components in the way most natural to the particular environment. Thus in a CORBA environment clients will expect to make use of both the CCM model and the EJB model as CORBA components, and in an EJB environment, clients will expect to make use of both kinds as enterprise beans. All four combinations of clients and components are illustrated in Figure 11.2.

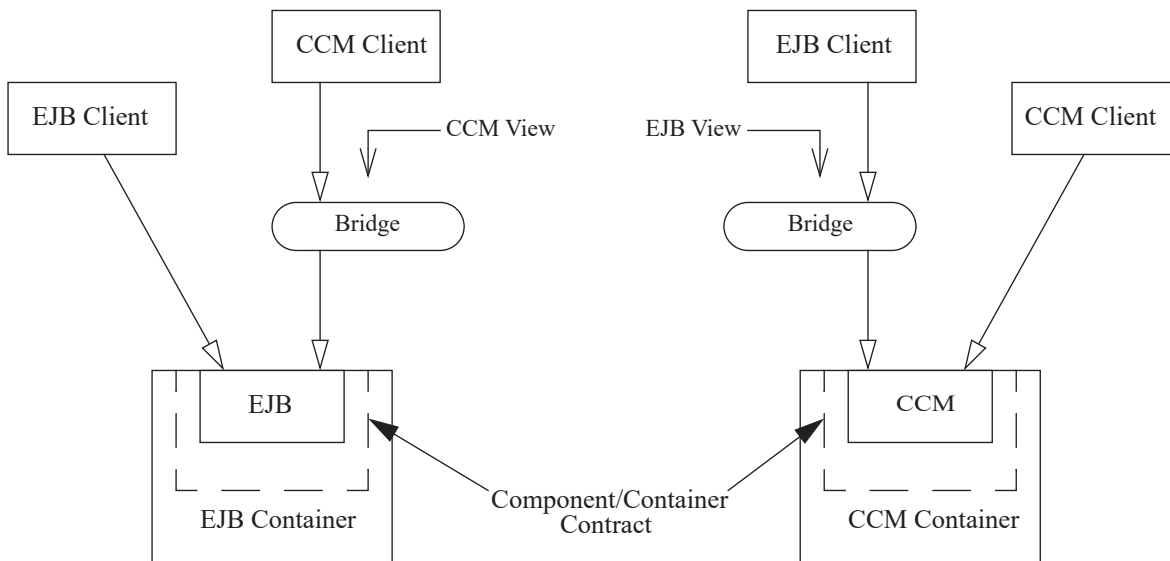


Figure 11.2- Interoperation in a mixed environment

In this scenario, components of one kind are made accessible to clients of another by way of two mechanisms: generation of bindings at development time and method translation at runtime. Thus, the containers provide an EJB view of a CORBA component and a CCM view of an EJB.

For application developers in a CORBA environment, EJBs specified in Java are mapped to CORBA IDL for use by CCM clients, and at runtime client calls on CCM methods are translated by a bridge into EJB methods. In effect, the EJBs *are* CORBA components.

For application developers in an EJB environment, CORBA components specified in IDL are mapped to Java interfaces for use by EJB clients, and at runtime client calls on EJB methods are translated by a bridge into CCM methods. In effect, the CORBA components *are* EJBs.

11.3 CORBA Component Views for EJBs

This kind of view allows a CORBA client — either a CORBA component or any piece of code that uses CORBA, and either component-aware or not — to access an EJB as a CORBA component. To do this, two things are needed:

1. A mapping of the definition of the existing EJB into the definition of a CORBA component. This mapping takes an EJB's RMI remote interface and home interface and produces an equivalent CORBA component definition.
2. A translation, at run-time, of CORBA component requests performed by a CORBA client into EJB requests. This translation can be performed in terms of either straight delegation, or as an interpretation of a CORBA client request in terms of EJB requests.

11.3.1 Mapping of EJB to Component IDL Definitions

An EJB definition includes the following EJB interfaces:

- EJB home interface - This interface extends the pre-defined **EJBHome** interface.
- EJB remote interface - This interface extends the pre-defined **EJBObject** interface.

Thus, for the purposes of this clause, at least these EJB interfaces must be mapped into IDL in order to obtain a CORBA component definition of a view that a CORBA client can use to make requests on an existing EJB. An EJB home interface definition maps into a CORBA component's home definition, whose implied IDL inherits from **CCMHome**. This means that **EJBHome** is mapped into **CCMHome**. Likewise, an EJB remote interface definition maps into a basic CORBA component definition, whose implied IDL inherits from **CCMObject**. This means that **EJBObject** is mapped into **CCMObject**.

In addition, **EJBHome** and **EJBObject** make use of the following pre-defined EJB interfaces:

- **HomeHandle**
- **Handle**
- **EJBMetaData**

Handles are an EJB concept that has no direct counterpart in CORBA components. Thus, **HomeHandle** and **Handle** are not directly mapped into equivalent IDL.

*Notice that although Interoperable Object References (IORs) and the ORB provided operations that manipulate them (**string_to_object** and **object_to_string**) are conceptually similar to Handles, there are enough differences between IORs and Handles to preclude a mapping from Handles to IORs.*

Meta data is available to a CORBA client but not in the same form as that provided by **EJBMetaData**. Given that an EJB maps into a CORBA component, whose definition produces the meta data that a CORBA client expects, mapping **EJBMetaData** into equivalent IDL is not required.

11.3.1.1 Java Language to IDL Mapping

The reader is assumed to be familiar with the specification for the Java to IDL mapping, whose major aspects are repeated here for convenience.

- A Java interface is an RMI/IDL remote interface if it at least extends **java.rmi.Remote** and all of its methods throw **java.rmi.RemoteException**.
- get- and set- name pattern names are translated to IDL attributes.
- IDL generated methods have only **in** parameters (but these can include object references to remote objects, allowing reference semantics normally obtained by using parameters of type **java.rmi.Remote**).
- Java objects that inherit from **java.io.Serializable** or **java.io.Externalizable** are mapped to a CORBA valuetype. All object types appearing in RMI remotable interfaces must inherit from these interfaces or from **java.rmi.Remote**. EJB **Key** and **Handle** types must inherit from **java.io.Serializable**.
 - However, the mapping does NOT require that methods on such objects or constructors be mapped to corresponding IDL operations on **valuetypes** and **init** specifications. The developer is expected to select those methods that should be mapped to IDL operations, and the method signatures must meet the requirements of the mapping.
 - Objects that inherit from **java.io.Externalizable** or that implement **writeObject** are understood to perform custom marshalling and the corresponding custom marshallers must be created for the CORBA valuetype.
- Arrays are mapped to “boxed” CORBA **valuetypes** containing sequences because Java arrays are dynamic.

- Java exceptions are subclassable; IDL exceptions are not. Consequently a name pattern is used to map to IDL exceptions. The Java exception object is mapped to a CORBA valuetype. The CORBA valuetype has an inheritance hierarchy like that of the corresponding Java exception object.
- Some additional programming is required to define Java classes (including EJB implementations) that are accessible via RMI/IIOP. This is to account for the fact that IIOP does not support distributed garbage collection.

11.3.1.2 EJB to IDL mapping

In general, the CORBA component that results from mapping an EJB will support an interface that is the Java to IDL map of the Remote interface of the EJB. The mapping rules are as follows.

11.3.1.2.1 Mapping the Remote Interface

- An EJB's remote interface maps to a definition of a basic CORBA component that supports the default interface. The form of the CORBA component definition is **component XXX supports XXXDefault**.
- An EJB's remote interface declaration is used to create a **supports** declaration and the corresponding IDL for the primary interface of the CORBA component that the EJB maps to. The identifier of this supported interface on the component is **XXXDefault**, where **XXX** is the name of the EJB remote interface. This generated interface is referred to as the *Default* interface of the component that the given EJB maps to.
- Each operation on the Remote interface is mapped under Java to IDL to an equivalent operation on the **XXXDefault** interface.
- Each pair of **getXXX** and **setXXX** methods in the EJB remote interface will be mapped to IDL attributes in the component definition itself. Any exceptions thrown by a **getXXX** method is mapped to an exception in the **getraises** clause of the mapped IDL attribute. Likewise, any exception thrown by a **setXXX** method is mapped to an exception in the **setraises** clause of the mapped IDL attribute. The actual definitions of the exceptions thrown are mapped following the Java to IDL rules.

11.3.1.2.2 Mapping the Home Interface

- An EJB's home interface maps to a definition of a CORBA component home. The form of the CORBA component home definition is **home YYY manages XXX**, where **YYY** is the name of the EJB home interface. Mapping an EJB home into a CORBA component home requires the existence of meta data that links the EJB home to the EJB that it hosts. These meta data are obtained from the EJB's deployment descriptor. Thus **XXX** is the name of the EJB that the EJB home hosts, as it is given in the EJB deployment descriptor.
- The EJB home methods called **create** are mapped into home **factory** declarations in IDL. The actual names of each of the **factory** operations are produced following the rules for mapping Java names to IDL names in the Java to IDL specification. The Java parameters of the operation are mapped to their corresponding IDL types and names as defined by Java to IDL.
- An EJB Primary Key class is mapped to a CORBA **valuetype** using the mapping rules in Java to IDL. This **valuetype** will be declared in the IDL for the CORBA component home as the primary key **valuetype** for the component. The key **valuetype** will inherit from **Components::PrimaryKeyBase**. If an EJB home uses a primary key, then the form of the CORBA component home definition is **home YYY manages XXX primarykey KKK**, where **KKK** is the name of the valuetype that the EJB primary key class maps to.
- The EJB home operation named **findByPrimaryKey** is mapped into the **find_by_primary_key(in <key-type> key)** operation on the component's implicit home interface.

- Finder and Creator EJB operations that return an RMI style object reference are mapped into Component IDL operations that return a CORBA Component Object Reference to **XXX**.

EJB home operations prefixed **find** whose return type is the type of the EJB hosted by the EJB home are mapped into component home **finder** operations in IDL. The actual names of each of the **finder** operations are produced following the rules for mapping Java names to IDL names in the Java to IDL specification. The Java parameters of the operation are mapped to their corresponding IDL types and names as defined by Java to IDL.

- Finder EJB operations that return a Java Enumeration are mapped into CORBA component operations that return a value of type **Enumeration**. This value type is declared as:

```

module Components {
  abstract valuetype Enumeration {
    boolean has_more_elements();
    CCMObject next_element();
  };
};

```

The Enumeration interface is just the RMI/IIOP image of the Java Enumeration class as defined in the JDK 1.1.6+. Sun has said that they intend to replace this with the JDK 1.2 (Java 2.0) Collections in a future version of the EJB specification. Subsequent to such a specification being issued, the CORBA components specification will be updated to correspond.

A concrete specialization of this abstract value type must be provided. This specialization has the form:

```

module Components {
  typedef sequence<CCMObject> CCMObjectSeq;
  valuetype DefaultEnumeration : Enumeration {
    private CCMObjectSeq objects;
  };
};

```

Any implementation of **DefaultEnumeration**, in any language, must provide implementations for the two **Enumeration** methods. Any client ORB that supports the interoperable bridge has to provide an implementation that knows how to read **DefaultEnumeration** from the wire and to use that information to provide a local implementation of these two methods. Any EJB container that supports the CCM-EJB bridge has to provide an implementation that knows how to construct itself from a **java.util.Enumeration** and then write itself to the wire as a **DefaultEnumeration**.

- In order for an EJB home definition that defines **findByPrimaryKey** to be successfully mapped onto a CORBA component home definition, it must define a **create** method that takes the primary key of the hosted EJB as its sole argument and returns an instance of the hosted EJB. This create method is mapped to **create(in <key-type> key)** on the CORBA component implicit home interface.

11.3.1.2.3 Mapping standard exceptions

The EJB exceptions **FinderException**, **CreateException**, **DuplicateKeyException**, and **RemoveException** thrown by methods to find, create, and remove an EJB are always mapped to the CCM exceptions **Components::FinderFailure**, **Components::CreateFailure**, **Components::DuplicateKeyValue**, and **Components::RemoveFailure**, respectively.

11.3.2 Translation of CORBA Component requests into EJB requests

A CORBA client that uses a CORBA component view on an EJB expects to be able to perform CORBA component requests on such a view. These requests need to be translated into EJB requests at run-time. This translation can be performed at the client-side, server-side, or a combination of the two. Table 11.1 lists the CORBA component operations that a CORBA client can perform requests on by interface, and it lists the corresponding EJB methods that these requests translate into, also by interface.

Table 11.1 - Translation of CCM operation requests into EJB method requests

CCM Interface	Operation called by client	EJB interface	Method invoked by bridge
CCMHome	ComponentDef get_component_def ();	EJBHome	EJBMetaData getEJBMetaData () throws RemoteException;
	CORBA::IObject get_home_def ();		EJBMetaData getEJBMetaData() throws RemoteException;
	void remove_component (in CCMObject comp) raises (RemoveFailure);		void remove (Handle handle) throws RemoveException, RemoteException;
<home-name>Explicit	<name> createXXX (<arg-list>) raises (CreateFailure, DuplicateKey Value, InvalidKey);	<home-name>	<name> create (<arg-list>) throws CreateException, DuplicateKey Exception;
	<name> findXXX (<arg-list>) raises (FinderFailure, <exceptions>);		<name> findXXX (<arg-list>) throws <exceptions>;
<home-name>Implicit	<name> create (in <key-type> key) raises (CreateFailure, DuplicateKey Value, InvalidKey);	<home-name>	<name> create (Object primaryKey) throws CreateException, DuplicateKeyException;
	<name> find_by_primary_key (in <key-type> key) raises (FinderFailure, UnknownKey Value, InvalidKey);		<name> findByPrimaryKey (<key-type> key) throws FinderException, ObjectNot FoundException;
	void remove (in <key-type> key) raises (RemoveFailure, Unknown KeyValue, InvalidKey);	EJBHome	void remove (Object primaryKey) throws RemoveException, RemoteException;
	<key_type> get_primary_key (in <name> comp);	EJBObject	Object getPrimaryKey () throws RemoteException;
CCMObject	ComponentDef get_component_def ();	EJBHome	EJBMetaData getEJBMetaData () throws RemoteException;
	CCMHome get_ccm_home ();	EJBObject	EJBHome getEJBHome() throws RemoteException;

Table 11.1 - Translation of CCM operation requests into EJB method requests

CCM Interface	Operation called by client	EJB interface	Method invoked by bridge
	PrimaryKeyBase get_primary_key ();	EJBObject	Object getPrimaryKey () throws RemoteException;
	void remove() raises (RemoveFailure);		void remove () throws RemoveException, RemoteException;
	void configuration_complete () raises (InvalidConfiguration);		Translation performed by bridge is to raise the NO_IMPLEMENT exception
<name>	<res-type> <operation> (<arg-list>) raises (<exceptions>);	<name>	<res-type> <operation> (<arg-list>) throws <exceptions>;
	<res-type> getXXX () throws <exceptions>;		<res-type> getXXX () throws <exceptions>;
	void setXXX (<arg-list>) throws <exceptions>;		void setXXX (<arg-list>) throws <exceptions>;

Notice that a CORBA client may use operations on object references such as **string_to_object** and **object_to_string** that may be considered as analogous to EJB **Handle** methods. However, these operations are not seen by the bridge since they are performed on the ORB and thus no translation for these operations on the part of the bridge is required.

The following restrictions apply:

- **create (in <key_type> key)** on the component implicit home interface can only be validly invoked by a CORBA client if the underlying EJB home declares the **findByPrimaryKey** operation.
- **remove (in <key_type> key)** on the component implicit home interface can only be validly invoked by a CORBA client if the underlying EJB home declares the **findByPrimaryKey** operation.
- **get_primary_key** on the component implicit home and on **CCMObject** can only be validly invoked by a CORBA client if the underlying EJB home declares the **findByPrimaryKey** operation.
- **configuration_complete** on **CCMObject** is not translated by the bridge, a request on this operation by a CORBA client raises the NO_IMPLEMENT exception.

11.3.3 Interoperability of the View

As stated in “Translation of CORBA Component requests into EJB requests” on page 183, translation of CORBA Component requests into EJB requests can happen at either the client-side, the server-side, or a combination of the two.

However, in order to provide interoperability of implementations of CORBA component views of EJBs, a minimal number of translation points must be performed and they must be performed at an explicitly defined location: either the client-side or the server-side. For the implementation of a CORBA component view of an EJB, and for an EJB home interface, the translation points are as follows.

11.3.3.1 Translation of specific method names

The following methods shall translate their names as indicated.

Table 11.2 - Translation of specific method names

CCM Interface	Method name	EJB Interface	Translation
CCMHome	get_component_def	EJBHome	getEJBMetaData
	remove_component		remove
<name>Implicit	find_by_primary_key	<name>	findByPrimaryKey
	remove		remove__java_lang_Object
	create		create__java_lang_Object
	get_primary_key	EJBObject	getPrimaryKey
CCMObject	get_ccm_home	EJBObject	getEJBHome
	get_primary_key		getPrimaryKey

11.3.3.2 Handling of standard exceptions

The following exceptions, caught by the indicated methods, shall be translated as indicated before raising them to their CORBA clients.

Table 11.3 - Handling of standard exceptions

CCM Interface	Method name	Exception caught	Translation
CCMHome	get_component_def	RemoteException	CORBA::UNKNOWN
	remove_component	RemoveException RemoteException	Components::RemoveFailure CORBA::UNKNOWN
<name>Implicit	create	DuplicateKeyException CreateException	Components::DuplicateKeyValue Components::CreateFailure
	find_by_primary_key	ObjectNotFoundException FinderException	Components::UnknownKeyValue Components::FinderFailure
	remove	RemoveException RemoteException	Components::RemoveFailure CORBA::UNKNOWN
	get_primary_key	RemoteException	CORBA::UNKNOWN
CCMObject	get_ccm_home	RemoteException	CORBA::UNKNOWN
	get_primary_key	RemoteException	CORBA::UNKNOWN
	remove	RemoveException RemoteException	Components::RemoveFailure CORBA::UNKNOWN

Note – RemoteException is translated into **CORBA::UNKNOWN** system exception according to rules defined in http://www.omg.org/technology/documents/formal/java_language_mapping_to_omg_idl.htm (Mapping RMI Exceptions to CORBA System Exceptions sub clause).

11.3.3.3 Handling of a primary key parameter

The methods **create**, **find_by_primary_key** and **remove**, defined by **<home>Implicit** shall translate the primary key valuetype they get as input parameter to a **CORBA::Any** equivalent. Likewise, the method **get_primary_key** defined by **<home>Implicit** shall translate the **CORBA::Any** value of the primary key it gets as a result from its request into an equivalent primary key valuetype before returning it.

The method **get_primary_key**, defined by **CCMObject**, shall translate the **CORBA::Any** value of the primary key it gets as a result from its request into an equivalent **Components::PrimaryKeyBase** valuetype before returning it.

11.3.4 CORBA Component view Example

In this sub clause we show a simple EJB together with the corresponding Component IDL. Note that the EJB deployment metadata is needed to generate the IDL; this is because the metadata binds together the Remote interface and the Home interface.

Below are the remote interfaces of the EJB.

```
package example;

class CustInfo implements java.io.Serializable
{
    public int custNo;
    public String custName;
    public String custAddr;
};

class CustBal implements java.io.Serializable
{
    public int custNo;
    public float acctBal;
};

interface CustomerInquiry extends javax.ejb.EJBObject
{
    CustInfo getCustInfo(int iCustNo)
        throws java.rmi.RemoteException;
    CustBal getCustBal(int iCustNo)
        throws java.rmi.RemoteException;
};
```

```

interface CustomerInquiryHome extends javax.ejb.EJBHome
{
    CustomerInquiry create()
        throws java.rmi.RemoteException;
};

```

Below are the contents of the descriptor classes as they might be expressed in an equivalent XML document.

```

<ejb-jar>
  <session>
    <description>
    </description>
    <ejb-name> CustomerInquiry </ejb-name>
    <home> example.CustomerInquiryHome </home>
    <remote> example.CustomerInquiry </remote>
    <ejb-class> example.CustomerInquiryBean </ejb-class>
    <session-type> Stateful </session-type>
  </session>
</ejb-jar>

```

The EJB is a session bean, and in this case, its **create** operation requires no parameters. The two operations take a key value and return values to the caller. The EJB implementation will use **JDBC** to retrieve the information to be returned by the operations on the **CustomerInquiry** EJB.

The serializable value classes are translated by RMI/IIOP into CORBA concrete **valuetypes** as follows:

```

valuetype CustInfo {
    public long custNo;
    public ::CORBA::WStringValue custName;
    public ::CORBA::WStringValue custAddr;
};

```

```

valuetype CustBal {
    public long custNo;
    public float custBal;
};

```

The information in the deployment descriptor and the home and remote interface declarations is introspected and used to generate the following IDL:

```

interface CustomerInquiryDefault {
    CustInfo getCustInfo(in long iCustNo);
    CustBal getCustBal(in long iCustNo);
};

```

```
component CustomerInquiry supports CustomerInquiryDefault {};
```

```
home CustomerInquiryHome manages CustomerInquiry {  
    factory create();  
};
```

11.4 EJB Views for CORBA Components

This kind of view allows a Java client — either an EJB or any other piece of Java code — to access a CORBA component as an EJB. To do this, two things are needed:

1. A mapping of the Component IDL definition of a CORBA component into an EJB definition. This mapping only considers that portion of the Component IDL language that has a counterpart in the EJB specification language and it ignores the rest. Notice that “The home and remote interfaces of the enterprise bean’s client view are defined as Java RMI interfaces. This allows the Container to implement the home and remote interfaces as distributed objects.” One implication of this is that the signatures on methods on an EJB’s remote interface can only include parameters with **in** semantics. That is, **out** and **inout** semantics for parameters is not allowed. As a consequence, the **out** and **inout** qualifiers for parameters in IDL interface method definitions are not included in the portion of Component IDL that can be mapped to an EJB definition.

Note however that a Java client does not have to use an EJB view in order to access a CCM. Any Java client can access a CCM directly via its IDL interface using a standard Java ORB, such as the one built into the JDK. This provides full access to all aspects of the CCM. Since the EJB view is derived using the IDL to Java mapping rules, the Java IDL interface is identical to the EJB view for all business operations. The only differences are in the operations mentioned in Table 11.4 on page 191 have slightly different names and signatures.

2. A translation, at run-time, of EJB requests performed by a Java client into CORBA component requests.

11.4.1 Mapping of Component IDL to Enterprise JavaBeans specifications

The portion of the Component extensions to the IDL language that can be mapped to the EJB specification language is denoted by the following subset of the Component extensions to IDL grammar.

```
<component_dcl> ::= <component_header> “{” <component_body> “}”  
<component_header> ::= “component” <identifier> [ <supported_interface_spec> ]  
<supported_interface_spec> ::= “supports” <scoped_name> { “,” <scoped_name> }*  
<component_body> ::= <component_export>*  
<component_export> ::= <attr_dcl> “;”  
<attr_dcl> ::= <readonly_attr_spec> | <attr_spec>  
<readonly_attr_spec> ::= “readonly” “attribute” <param_type_spec> <readonly_attr_declarator>  
<readonly_attr_declarator> ::= <simple_declarator> <raises_expr> | <simple_declarator> { “,”  
    <simple_declarator> }*  
<attr_spec> ::= “attribute” <param_type_spec> <attr_declarator>  
<attr_declarator> ::= <simple_declarator> <attr_raises_expr> | <simple_declarator> { “,”  
    <simple_declarator> }*
```

```

<attr_raises_expr> ::= <get_excep_expr> [ <set_excep_expr> ] | <set_excep_expr>
<get_excep_expr> ::= "getraises" <exception_list>
<set_excep_expr> ::= "setraises" <exception_list>
<exception_list> ::= "(" <scoped_name> { "," <scoped_name> } * ")"
<home_dcl> ::= <home_header> <home_body>
<home_header> ::= "home" <identifier> "manages" <scoped_name> [ <primary_key_spec> ]
<primary_key_spec> ::= "primarykey" <scoped_name>
<home_body> ::= "{" <home_export>* "}"
<home_export> ::= <factory_dcl> ";" | <finder_dcl> ";"
<factory_dcl> ::= "factory" <identifier> "(" [ <init_param_decls> ] ")" [ <raises_expr> ]
<finder_dcl> ::= "finder" <identifier> "(" [ <init_param_decls> ] ")" [ <raises_expr> ]

```

The rules for mapping a CORBA component definition into an EJB definition are defined in the following sub clauses. Where appropriate, these rules rely on the standard IDL to Java mapping.

11.4.1.1 Mapping the component definition

- A basic CORBA component definition is mapped to an EJB remote interface definition.
- The name of the EJB remote interface is the name of the basic CORBA component in the Component IDL definition.
- For each operation defined in each interface that the CORBA component **supports**, a method definition will be included in the EJB remote interface that the CORBA component maps to. That is, the EJB to which the basic CORBA component maps defines all the supported operations defined by the basic CORBA component.
- The signatures of the CORBA component operations are mapped to signatures of EJB remote interface methods following the IDL to Java mapping rules. Only signatures whose parameters have an **in** qualifier are allowed. Signatures that include parameters with **out** or **inout** qualifiers shall be signaled as an error.
- For each attribute **XXX** that the CORBA component defines, the corresponding EJB remote interface defines a pair of **getXXX** and **setXXX** methods, where **XXX** is the name of the given attribute. If the attribute definition includes a **getraises** exception clause, then the corresponding **getXXX** method definition in the EJB remote interface will include a throws exception clause. Likewise, if the attribute definition includes a **setraises** exception clause, then the corresponding **setXXX** method definition in the EJB remote interface will include a throws exception clause.
- Exceptions raised by CORBA component definition operations and attributes are mapped to exceptions thrown by EJB method definitions using the standard IDL to Java mapping rules.

11.4.1.2 Mapping the Component Home definition

- A CORBA component's home definition is mapped to an EJB home's remote interface definition. That is a definition of the form **home XXX manages YYY [primarykey KKK]** is mapped to an EJB home interface with name **xxx**.
- The methods defined by the EJB home remote interface include the implicit as well as the explicit methods of the CORBA component's home definition.
- Implicit CORBA component home operations are mapped to EJB home remote interface methods as follows:

- `<component_type> create (in <key_type> key)` raises `(Components::CreateFailure, Components::DuplicateKeyValue, Components::InvalidKey)`; maps to `<component_type> create (<key_type> key)` throws `DuplicateKeyException, CreateException`.
 - `<component_type> find_by_primary_key (in <key_type> key)` raises `(Components::FinderFailure, Components::UnknownKeyValue, Components::InvalidKey)`; maps to `<component_type> findByPrimaryKey(<key_type> key)` throws `ObjectNotFoundException, FinderException`.
 - `void remove (in <key_type> key)` raises `(Components::RemoveFailure, Components::UnknownKeyValue, Components::InvalidKey)`; maps to the `remove` by key method defined in `EJBHome`.
 - `<key_type> get_primary_key (in <component_type> comp)`; has no counterpart in an EJB home definition. Given that `EJBObject` already defines `getPrimaryKey`, it is not necessary to map `get_primary_key` on the implicit home to an EJB home operation.
- Explicit CORBA component basic home operations are mapped to EJB home remote interface methods as follows:
 - A **factory** operation maps to an overloaded **create** method with the corresponding arguments and exceptions.
 - An operation maps to a **find<identifier>** method with the corresponding arguments and exceptions, where **<identifier>** is the name of the **finder** operation.
 - The signatures of **factory** and **finder** operations are mapped to signatures of EJB home interface methods following the IDL to Java mapping rules.
 - A **valuetype** that is used to define the primary key of a CORBA component home is mapped to a Java class under the rules of the standard IDL to Java mapping. In addition, such a Java class is defined to extend `java.io.Serializable`.

11.4.1.3 Mapping standard exceptions

The CCM exceptions `Components::FinderFailure`, `Components::CreateFailure`, `Components::DuplicateKeyValue` and `Components::RemoveFailure` raised by methods to find, create and remove a CORBA component are always mapped to the EJB exceptions `FinderException`, `CreateException`, `DuplicateKeyException` and `RemoveException`, respectively.

11.4.2 Translation of EJB Requests into CORBA Component Requests

A Java client that uses an EJB view on a CORBA component expects to be able to perform EJB requests on such a view. These requests need to be translated into CORBA component requests at run-time. This translation can be performed at the client-side, the server-side, or a combination of the two. Table 11.4 lists the EJB methods that a Java client can perform requests on by interface, and it lists the corresponding CORBA component operations that these requests translate into, also by interface.

Table 11.4 - Translation of EJB method requests into CCM operation requests

EJB Interface	Method called by client	CCM interface	Operation called by bridge
EJBHome	EJBMetaData getEJBMetadata () throws RemoteException;	CCMHome	Translation performed by bridge does not call a CCM standard operation
	void remove (Handle handle) throws RemoveException, RemoteException;		void remove_component (in CCMObject comp) raises (RemoveFailure);
	void remove (Object primaryKey) throws RemoveException, RemoteException;	<home-name>Implicit	void remove (in <key-type> key) raises (RemoveFailure, UnknownKeyValue, InvalidKey);
	HomeHandle getHomeHandle () throws RemoteException;		Translation performed by bridge does not call a CCM standard operation
<home-name>	<name> create (<arg-list>) throws CreateException, DuplicateKeyException;	<home-name>Explicit	<name> createXXX (<arg-list>) raises (CreateFailure, DuplicateKeyValue, InvalidKey);
	<name> findByXXX (<arg-list>) throws <exceptions>;		<name> findXXX (<arg-list>) raises (FinderFailure, <exceptions>;
	<name> findByPrimaryKey (<key-type> key) throws FinderException, ObjectNotFoundException;	<home-name>Implicit	<name> find_by_primary_key (in <key-type> key) raises (FinderFailure, UnknownKeyValue, InvalidKey);
EJBObject	EJBHome getEJBHome () throws RemoteException;	CCMObject	CCMHome get_ccm_home ();
	Object getPrimaryKey () throws RemoteException;		PrimaryKeyBase get_primary_key ();
	void remove () throws RemoveException, RemoteException;		void remove () raises (RemoveFailure);
	boolean isIdentical (EJBObject object) throws RemoteException;	CORBA::Object	boolean is_equivalent ();
	Handle getHandle () throws RemoteException;		Translation performed by bridge does not call a CCM standard operation.
<name>	<res-type> <operation> (<arg-list>) throws <exceptions>;	<name>	<res-type> <operation> (<arg-list>) raises (<exceptions>;
	<res-type> getXXX () throws <exceptions>;		<res-type> get_XXX () raises (<exceptions>;
	void setXXX (<arg-list>) throws <exceptions>;		<res-type> set_XXX () raises (<exceptions>;

Table 11.4 - Translation of EJB method requests into CCM operation requests

EJB Interface	Method called by client	CCM interface	Operation called by bridge
EJBMetadata	EJBHome getEJBHome () throws RemoteException;		Translation performed by bridge on all these invocations does not call a CCM standard operation.
	Class getHomeInterfaceClass () throws RemoteException;		
	Class getRemoteInterfaceClass () throws RemoteException;		
	Class getPrimaryKeyClass () throws RemoteException;		
	boolean isSession () throws RemoteException;		
	boolean isStatelessSession() throws RemoteException		

In addition, the EJB programming model allows a Java client to:

- Locate EJB homes and distinguished EJB objects via **JNDI**.
- Demarcate transactions via a **UserTransaction** object, after locating this object via **JNDI**.

These requests are translated into similar requests provided by the CORBA component programming model, as follows:

- Location of home and EJB objects requires the definition of a mapping of JNDI to the COSNaming service. It also requires the mapping of a COSNaming name space into a JNDI name space.
- Transaction demarcation requires the definition of a mapping of JTA to the CORBA transaction service. It also requires that a **JNDI** name space location be populated with an object that implements **UserTransaction** and that maps to the corresponding CORBA transaction service object.

11.4.3 Interoperability of the View

As stated in “Translation of EJB Requests into CORBA Component Requests” on page 190 can happen at either the client-side, the server-side, or a combination of the two.

However, in order to provide interoperability of implementations of EJB views of CORBA components, a minimal number of translation points must be performed and they must be performed at an explicitly defined location: either the client-side or the server-side. For the implementation of an EJB view of a CORBA component, and for a CCM interface, the translation points are as follows.

11.4.3.1 Translation of specific method names

The following methods shall translate their names as indicated.

Table 11.5 - Translation of specific method names

EJB Interface	Method name	CCM Interface	Translation
EJBHome	remove	CCMHome	remove_component
<name>	findByPrimaryKey	<name>Implicit	find_by_primary_key
	remove__java_lang_Object		remove
	create__java_lang_Object		create
EJBObject	getEJBHome	CCMObject	get_ccm_home
	getPrimaryKey		get_primary_key
	isIdentical	CORBA::Object	is_equivalent

11.4.3.2 Handling of standard exceptions

The following exceptions, caught by the indicated methods, shall be translated as indicated before raising them to their EJB clients.

Table 11.6 - Handling of standard exceptions

EJB Interface	Method name	Exception caught	Translation
EJBHome	remove	Components::RemoveFailure CORBA system exceptions	RemoveException RemoteException
	remove__java_lang_Object	Components::RemoveFailure CORBA system exceptions	RemoveException RemoteException
<name>	create	Components::CreateFailure Components::DuplicateKeyValue	CreateException DuplicateKeyException
	findByPrimaryKey	Components::UnknownKeyValue Components::FinderFailure	ObjectNotFoundException FinderException
EJBObject	getEJBHome	CORBA system exceptions	RemoteException
	getPrimaryKey	CORBA system exceptions	RemoteException
	remove	Components::RemoveFailure CORBA system exceptions	RemoveException RemoteException
	isIdentical	CORBA system exceptions	RemoteException

Note – CORBA system exceptions are translated into **RemoteException** according to rules defined in http://www.omg.org/technology/documents/formal/java_language_mapping_to_omg_idl.htm (Mapping CORBA System Exceptions to RMI Exceptions sub clause).

11.4.3.3 Handling of a primary key parameter

The methods **create** and **findByPrimaryKey**, defined by <name>, and **remove__java_lang_Object**, defined by **EJBHome**, shall translate the primary key valuetype they get as input parameter to a **CORBA::Any** equivalent.

The method **getPrimaryKey**, defined by **EJBOject**, shall translate the **CORBA::Any** value of the primary key it gets as a result from its request into an equivalent Java Object valuetype before returning it.

11.4.4 Example

We show a simple CORBA component definition and its corresponding EJB mapping. The basic CORBA component **Account** is defined in terms of a regular IDL interface **AccountOps**. The home **AccountHome** is defined to manage **Account** and to use a primary key.

```
interface AccountOps {
    void debit( in double amt ) raises (NotEnoughFunds);
    void credit( in double amt );
};

component Account supports AccountOps {
    readonly attribute double balance;
};

valuetype AccountKey {
    public long acctNo;
};

home AccountHome manages Account primarykey AccountKey {
    finder largeAccount( double threshold );
};
```

The following EJB definition is derived from the definition of **Account** and its home.

```
public interface Account extends javax.ejb.EJBOject {
    public void debit( double amount )
        throws NotEnoughFunds, java.rmi.RemoteException;
    public void credit( double amount ) throws java.rmi.RemoteException;
    public double getBalance() throws java.rmi.RemoteException;
};

public class AccountKey implements java.io.Serializable {
    public long acctNo;
    public AccountKey( long k ) { acctNo = k; }
};

public interface AccountHome extends javax.ejb.EJBHome {
    public Account create( AccountKey key )
        throws DuplicateKeyException, CreateException,
        java.rmi.RemoteException;

    public Account findByPrimaryKey( Account key )
        throws ObjectNotFoundException, FinderException,
        java.rmi.RemoteException;
};
```

```

public Account findByLargeAccount( double threshold )
    throws java.rmi.RemoteException;
};

```

11.5 Compliance with the Interoperability of Integration Views

As stated in “Interoperability of the View” on page 184 and “Interoperability of the View” on page 192, request translations must happen at an explicitly defined location: either the client-side or the server side.

Rather than mandate one location arbitrarily, a number of levels of compliance with the interoperability of integration views are defined. Vendors shall clearly state what level of interoperability is supported by their implementations. These levels are:

- **NONE**: Integration view implementations that comply with this level actually perform no request translations. These implementations can still interoperate with other implementations that understand non-translated requests (e.g., implementations compliant with levels **SERVER-SIDE** and **FULL**).
- **CLIENT-SIDE**: Translation occurs either in the address space of a client stub or in a separate address space downstream from the client stub but before the resulting GIOP request gets sent to the server.
- **SERVER-SIDE**: Translation occurs either in the address space of a server skeleton or in a separate address space upstream from the server skeleton but after the GIOP request has been received from the client. The presence of a server-side view must not prevent native (i.e., non-translated) access to the component.
- **FULL**: Integration view implementations that comply with this level comply with both the **CLIENT-SIDE** and **SERVER-SIDE** levels. Note that a stand-alone bridge in a separate address space complies at this level since it is both upstream of the client (**SERVER-SIDE**) and downstream of the server (**CLIENT-SIDE**).
- **FULL**: Integration view implementations that comply with this level comply with both the **CLIENT** and the **SERVER** levels.

Table 11.7 illustrates the possible combinations of level compliance that are implied by the previous definitions. Rows in the table denote implementations compliant with a given level that send a request. Columns denote implementations compliant with a given level that receive a request. So, for example, a **SERVER-SIDE** implementation cannot interoperate with a **CLIENT-SIDE** implementation because the **SERVER-SIDE** implementation does not translate on send and the **CLIENT-SIDE** implementation does not translate on receive.

Table 11.7 - Compliance with the Interoperability of Integration Views

	NONE	CLIENT-SIDE	SERVER-SIDE	FULL
NONE	no	no	yes	yes
CLIENT-SIDE	no	yes	yes	yes
SERVER-SIDE	no	no	yes	yes
FULL	yes	yes	yes	yes

11.6 Comparing CCM and EJB

The following series of tables summarized the component APIs for Enterprise Java Beans (EJB 1.1) and Basic CORBA Components. The tables are organized as follows:

1. The home interfaces that define the remote access protocols for creating or finding EJBs or CORBA components (“The Home Interfaces” on page 196).
2. The component interfaces that define the remote access protocols for invoking business operations on EJBs or CORBA components (“The Component Interfaces” on page 197).
3. The callback interfaces that the CORBA component or EJB programmer must implement (“The Callback Interfaces” on page 199).
4. The Context interfaces that provide the component developer access to container-provided services (“The Context Interfaces” on page 200).
5. The Transaction interface that supports bean-managed or component-managed transactions (“The Transaction Interfaces” on page 201).
6. The metadata interfaces that support access to component metadata (“The Metadata Interfaces” on page 202).

11.6.1 The Home Interfaces

Table 11.8 compares the home interfaces and operations that make up the EJB and CORBA component models. In EJB, the **EJBHome** object is created by the EJB container provider’s tools and provides implementations for methods of the base class and delegates factory or finder methods on a derived class (**<name>Home**) to similarly named methods on the bean itself (**<name>Bean**).

In the CORBA component model, homes are defined as righteous CORBA objects and the associated factory or finder methods are generated as operations on the home and the component developer implements these directly so the container need not provide delegation support. The component developer may not even need to provide implementations for the default factory and finder operations if sufficient information is provided with the component’s definition.

For CORBA clients to use EJB implementations, the container provider must externalize **EJBHome** to the CORBA client as a CORBA component home. This is accomplished by extensions to the Java to IDL mapping defined in the Interface Repository Metamodel clause. For EJB clients to access CORBA component homes, the container provider must create an **EJBHome** object that serves as a bridge between equivalent operations on **EJBHome** and the CORBA component home. This bridge is also described in the Interface Repository Metamodel clause.

Table 11.8 - Comparing the home interfaces of EJB and CORBA components

Construct	EJB Form	CCM Form	Notes
Module	javax.ejb	Components	
Interface	EJBHome extends java.rmi.Remote	CCMHome	
Operation	public EJBMetaData get EJBMetaData () throws java.rmi.RemoteException	ComponentDef get_component_def ();	CORBA IR supports more metadata
	public HomeHandle getHomeHandle() throws java.rmi.RemoteException		CORBA::object_to_string provides same function
	public void remove (HomeHandle handle) throws java.rmi.RemoteException, RemoveException	void remove_component (in CCMObject component) raises (CCMException);	CORBA references instead of handles REMOVE_ERROR is minor code

Table 11.8 - Comparing the home interfaces of EJB and CORBA components

Construct	EJB Form	CCM Form	Notes
	public void remove (java.lang.Object primaryKey) throws java.rmi.RemoteException, RemoveException		similar operation is defined on <home>Implicit for Homes with primaryKey
Interface	HomeHandle extends java.io.Serializable		CORBA reference used for handle
	public EJBHome getEJBHome() throws java.rmi.RemoteException		CORBA::string_to_object
Module	<session-name>	<session-home>	
Interface	<session>home extends EJBHome	<session-home>::CCMHome, <session-home>Implicit, <session-home>Explicit	
Operation	public <session-name>Remote create (<arg-type> <arg-list>) throws CreateException	<session-component> create ();	Generated operation Inherited from <home>Implicit
Module	<entity-name>	<entity-home>	
Interface	<entity>home extends EJBHome	<entity-home>::CCMHome, <entity-home>Implicit, <entity-home>Explicit	
Operation	public <entity-name>Remote create (<arg-type> <arg-list>) throws CreateException, DuplicateKeyException	<entity-component> create () raises (InvalidKey, DuplicateKey);	Generated operation Inherited from <home>Implicit
	public <entity-name>Remote findByPrimaryKey (<arg-type> <arg-list>) throws FinderException, ObjectNotFoundException	<entity-component> find (in <key-type> key) raises (InvalidKey, UnknownKeyType);	Generated operation Inherited from <home>Implicit
	public <entity-name>Remote find<method> (<arg-type> <arg-list>) throws FinderException, ObjectNotFoundException	<entity-component> <find-method> (in <arg-type> <arg-list>) raises (<exceptions>);	Specified operation Inherited from <home>Explicit

11.6.2 The Component Interfaces

Table 11.9 compares the component interfaces and operations that make up the EJB and CORBA component models. In EJB, the **EJBObject** object is created by the EJB container provider's tools and provides implementations for methods of the base class and delegates business methods to a derived class (**<name>Remote**).

In the basic CORBA component model, components are defined as righteous CORBA objects and the associated business methods are defined as operations on a supported interface and the component developer implements these directly so the container need not provided delegation support.

For CORBA clients to use EJB implementations, the container provider must externalize **EJBObject** to the CORBA client as a CORBA component. This is accomplished by extensions to the Java to IDL mapping defined in the Interface Repository Metamodel clause. For EJB clients to access CORBA components, the container provider must create an **EJBObject** implementation that serves as a bridge between business methods on **EJBObject** and the basic CORBA component's supported interface. This bridge is also described in the *Interface Repository Metamodel* clause.

Table 11.9 - Comparing the remote interfaces of EJB and CORBA components

Construct	EJB Form	CCM Form	Notes
Module	javax.ejb	Components	
Interface	EJBObject extends java.rmi.Remote	CCMObject	
Operation	public EJBHome getEJBHome() throws java.rmi.RemoteException	CCMHome get_ccm_home();	
	public java.lang.Object primaryKey getPrimaryKey() throws java.rmi.RemoteException		operation defined on <entity>home
	public void remove (Handle handle) throws java.rmi.RemoteException, RemoveException	void remove() raises (CCMException);	CORBA references instead of handles; REMOVE_ERROR is minor code
	public Handle getHandle() throws java.rmi.RemoteException		CORBA::object_to_string
	public boolean isIdentical (EJBObject obj) throws java.rmi.RemoteException	boolean is_equivalent(in Object obj);	
Interface	Handle extends java.io.Serializable		CORBA reference used for handle
	public EJBObject getEJBObject() throws java.rmi.RemoteException		CORBA::string_to_object
Module	<session-bean>	<session-component>	
Interface	<session>Remote extends EJBObject	<session>::CCMObject	
	<res-type> <operation> (<arg-type> <arg-list>) throws <exceptions>	<res-type> <operation>(in <arg-type> <arg-list> raises (<exceptions>);	business methods
Module	<entity-bean>	<entity-component>	
Interface	<entity>Remote extends EJBObject	<entity>::CCMObject	
	<res-type> <operation> (<arg-type> <arg-list>) throws <exceptions>	<res-type> <operation> (in <arg-type> <arg-list> raises (<exceptions>);	business methods

11.6.3 The Callback Interfaces

Table 11.10 summarizes the callback interfaces the EJB programmer or basic CORBA component programmer must implement. The EJB interfaces are specified as Java interfaces in accordance with the EJB 1.1 specification dated June 28, 1999. The CCM interfaces are specified in IDL as defined in this specification.

Table 11.10 - Comparing EJB and CCM Callback Interfaces

Construct	EJB Form	CCM Form	Notes
Module	javax.ejb	Components::Basic	
Interface	EnterpriseBean	EnterpriseComponent	
Interface	SessionBean extends EnterpriseBean	SessionComponent::EnterpriseComponent	
Operation	public void setSessionContext (SessionContext ctx) throws EJBException	void set_session_context(in SessionContext ctx) raises (CCMException);	
	public void ejbActivate () throws EJBException	void ccm_activate () raises (CCMException);	
	public void ejbPassivate () throws EJBException	void ccm_passivate () raises (CCMException);	
	public void ejbRemove () throws EJBException	void ccm_remove () raises (CCMException);	
Interface	<name>Bean extends SessionBean		Home operations are not delegated in CCM.
Operation	public void ejbCreate (<Arg-type> <arg-list>) throws CreateException, EJBException)		Implemented on home, CREATE_ERROR is minor code
Interface	SessionSynchronization	SessionSynchronization	
Operation	public void afterBegin () throws EJBException	void after_begin () raises (CCMException);	
	public void beforeCompletion() throws EJBException	void before_completion () raises (CCMException);	
	public void afterCompletion (boolean committed) throws EJBException	void after_completion (in boolean committed) raises (CCMException);	
Interface	EntityBean extends EnterpriseBean	EntityComponent::EnterpriseComponent	
Operation	public void setEntityContext (EntityContext ctx) throws EJBException	void set_entity_context (in EntityContext ctx) raises CCMException;	
	public void unsetEntityContext () throws EJBException	void unset_entity_context () raises (CCMException);	
	public void ejbActivate () throws EJBException	void ccm_activate () raises (CCMException);	
	public void ejbLoad () throws EJBException	void ccm_load () raises (CCMException);	

Table 11.10 - Comparing EJB and CCM Callback Interfaces

Construct	EJB Form	CCM Form	Notes
	public void ejbStore () throws EJBException	void ccm_store() raises (CCMException);	
	public void ejbPassivate () throws EJBException	void ccm_passivate () raises (CCMException);	
	public void ejbRemove () throws RemoveException, EJBException	void ccm_remove () raises (CCMException);	REMOVE_ERROR is a minor code
Interface	<name>Bean extends EntityBean		Home operations are not delegated in CCM.
Operation	public <key-type> ejbcreate (<Arg-type> <arg-list>) throws CreateException, DuplicateKeyException, EJBException		Implemented on home, CREATE_ERROR and DUPLICATE_KEY are minor codes
	public void ejbPostCreate () throws CreateException, DuplicateKeyException, EJBException		post_create not required in CCM due to CORBA identity model
	public <key-type> findByPrimaryKey (<Arg-type> <arg-list>) throws FinderException, NoSuchEntityException, ObjectNotFoundException, EJBException		Implemented on home, FIND_ERROR, NO_SUCH_ENTITY and OBJECT_NOT_FOUND are minor codes
	public <key-type> find<method> (<Arg-type> <arg-list>) throws FinderException, NoSuchEntityException, ObjectNotFoundException, EJBException		Implemented on home, FIND_ERROR, NO_SUCH_ENTITY and OBJECT_NOT_FOUND are minor codes

11.6.4 The Context Interfaces

The context interfaces summarized in Table 11.11 provide accessors to services provided by the component container. They are used by the component developer when these services are required.

Table 11.11 - Comparing the EJB and CCM Context Interfaces

Construct	EJB Form	CCM Form	Notes
Module	javax.ejb	Components::Basic	
Interface	EJBContext	CCMContext	
Operation	public java.security.Principal getCallerPrincipal()	Principal get_caller_principal();	
	public EJBHome getEJBHome()	CCMHome get_ccm_home();	
	public boolean getRollbackOnly() throws java.lang.IllegalState	boolean get_rollback_only() raises (IllegalState);	

Table 11.11 - Comparing the EJB and CCM Context Interfaces

Construct	EJB Form	CCM Form	Notes
	public javax.transaction.UserTransaction getUserTransaction () throws java.lang.IllegalState	Transaction::UserTransaction get_user_transaction () raises (IllegalState);	
	public boolean isCallerInRole (java.lang.String (roleName)	boolean is_caller_in_role (in string role);	
	public void setRollbackOnly() throws java.lang.IllegalState	void set_rollback_only() raises IllegalState;	
Interface	SessionContext extends EJBContext	SessionContext::CCMContext	
Operation	public EJBObject getEJBObject() throws java.lang.IllegalState	CORBA::Object get_CCM_Object() raises (IllegalState);	this will be the component reference
Interface	EntityContext extends EJBContext	EntityContext::CCMContext	
Operation	public EJBObject getEJBObject() throws java.lang.IllegalState	CORBA::Object get_CCM_Object() raises (IllegalState);	this will be the component reference
	public java.lang.Object getPrimaryKey () throws java.lang.IllegalState	PrimaryKeyBase get_primary_key() raises (IllegalState);	

11.6.5 The Transaction Interfaces

Table 11.12 summarizes the transaction interfaces provided for bean-managed or component-managed transactions. Both EJB and CCM provide an accessor function in the context to obtain a reference to a transaction service. The transaction service supported for EJB is JTA, a subset of JTS which is equivalent to the CORBA transaction service (OTS). The transaction service supported for CORBA components is implemented by the component container as a wrapper over the CORBA transaction service. **Components::Transaction** is functionally equivalent to JTA (which is not a distinct compliance level for OTS) with the addition of **suspend** and **resume**.

Table 11.12 - Comparing the EJB Transaction service (JTA) with CORBA component transactions

Construct	EJB Form	CCM Form	Notes
Module	javax.transaction	Components::Transaction	
Interface	UserTransaction	UserTransaction	
Operation	public void begin() throws NotSupported, SystemException	void begin () raises (NotSupported, SystemError);	SystemError to avoid confusion with System Exception
	public void commit() throws RollbackException, HeuristicMixedException, HeuristicRollbackException, java.security.SecurityException, java.lang.IllegalStateException, SystemException	void commit() raises (RollbackError, HeuristicMixed, HeuristicRollback, Security, IllegalState, SystemError	map CORBA system exceptions TRANSACTION_ROLLED_BACK to ROLLBACK and NO_IMPLEMENT to SECURITY

Table 11.12 - Comparing the EJB Transaction service (JTA) with CORBA component transactions

Construct	EJB Form	CCM Form	Notes
	public void rollback() throws java.security.SecurityException, java.lang.IllegalStateException, SystemException	void rollback() raises (Security, IllegalState, SystemError);	
	public void setRollbackOnly() throws SystemException	void set_rollback_only() raises (SystemError);	
	public int getStatus() throws SystemException;	Status get_status() raises (SystemError);	
	public void setTransactionTimeout (int seconds) throws SystemException	void set_transaction_timeout(in long to) raises (SystemError);	
		TranToken suspend() raises (NoTransaction, SystemError);	CCM supports suspend/resume which JTA does not
		void resume(in TranToken) raises (invalidToken, SystemError);	CCM supports suspend/resume which JTA does not

11.6.6 The Metadata Interfaces

The EJB component model supports a limited set of metadata through the **EJBMetaData** interface. The CORBA component model extends the CORBA interface repository to add component-unique metadata for components. This meta-data is in addition to the metadata currently provided by the IR. When EJB clients access CORBA components, the container provider must provide an implementation of **EJBMetaData**, which supports the necessary metadata from the Interface Repository or the component descriptors. This is described further in Clause 12. When CORBA clients access EJB implementations, the Interface Repository is already populated for the **EJBHome** and **EJBObject** interfaces, enabling client requests to be satisfied. Table 11.13 compares the metadata supported by EJB and CORBA Components.

Table 11.13 - Comparing component metadata between EJB and CORBA components

Construct	EJB Form	CCM Form	Notes
Module	javax.ejb	IR	
Interface	EJBMetaData	ComponentDef	
	public EJBHome getEJBHome()		
	public java.lang.Class getHomeInterfaceClass()		
	public java.lang.Class getRemoteInterfaceClass()		
	public java.lang.Class getPrimaryKeyClass()		
	public boolean isSession()		
	public boolean isStatelessSession()		

12 Interface Repository Metamodel

12.1 Introduction

The first goal of the MOF-compliant metamodel is to express the extensions to IDL defined by the CORBA Component Model. Since these extensions are derived from the previously-existing IDL base, it is not possible to define a MOF-compliant metamodel for the extensions without defining a MOF-compliant metamodel for the IDL base.

Thus, the first MOF Package defined, entitled *BaseIDL*, is a MOF-compliant description of the pre-existing CORBA Interface Repository, while the second Package, entitled *ComponentIDL*, expresses the Component Model extensions. As shown by the following package diagram (Figure 12.1), the *ComponentIDL* Package is dependent upon the *BaseIDL* Package.

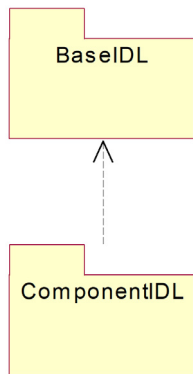


Figure 12.1- The Two Packages for the IDL Metamodel

12.1.1 BaseIDL Package

The base CORBA Interface Repository (IR) is described in Part 1 of this specification in the form of CORBA IDL. Because the MOF is more expressive than IDL, a range of legitimate MOF-compliant metamodels are equivalent to this IDL. For instance, multi-valued attributes and references expressed in IDL could be ordered or unordered, allow an instance to be contained in the collection only once or more than once. Further, specific multiplicity constraints could be specified; for example:

Can the sequence be empty?

Is there an upper bound?

As can be seen from an examination of the portion of the metamodel contained in the *BaseIDL* Package, many such questions are resolved via the more precise expression that the MOF enables.

12.1.1.1 A Structural Comparison of the BaseIDL Package with the Existing IR

Although the structure of the MOF-compliant CORBA IR is very similar to the existing CORBA IR, the authors have taken this opportunity to do some streamlining.

- In the existing CORBA IR, elements that are “typed,” (such as constants, attributes, etc.) hold an attribute of type *IDLType*. However, the same *IDLType* can be the type for many elements, so an attribute (with its composition semantics) is not appropriate. Instead, the MOF-compliant IR specifies the abstract *Typed* metaclass, and an Association between *Typed* and *IDLType*. This change eliminates the need for repeating the *type* attribute, which returns a *TypeCode*, in 6 different metaclasses.
- In the existing CORBA IR, *StructField*, *Parameter*, and *UnionField* are datatypes (structs). The MOF-compliant IR specifies them as full-blown metaclasses so that they can participate as derivations of the *Typed* metaclass.
- The MOF-compliant IR does not have to represent a repository since MOF-based servers inherently have such a construct. Thus, the MOF-compliant IR has no *Repository* metaclass and it specifies *Container* as a sub(meta)class of *Contained*, simplifying the hierarchy.
- The existing IR’s *IRObj* provides a *def_kind* readonly attribute. This information would be redundant in a MOF server, which inherently carries information describing the type of a metaobject. Thus, there is no *IRObj* metaclass in the MOF-compliant IR. However, it can be derived for a CORBA IR layer.
- In the existing IR, *UnionDef*, *StructDef*, *ExceptionDef*, and *OperationDef* inherit from *Container*. Since they each contain only a single type of object, it makes less sense for them to have a reference to a collection of *Contained* metaobjects. Instead, in the MOF-compliant IR they each hold their set of fields or parameters as attributes.
- As a simplification the two-stated enums *AttributeMode* and *OperationMode* have been eliminated. Attributes typed as *AttributeMode* or *OperationMode* have been turned into boolean-typed attributes.
- Basic CRUD operations for creating, reading, updating, and deleting metaobjects are generally not included in the metamodel, since these are generated automatically by the MOF-IDL mapping, which takes a MOF-compliant metamodel as input and deterministically derives the IDL for representing the metamodel in a repository.
- The existing IR duplicates many of the interfaces representing basic IR elements with structs representing the same elements. This duplication supports the ability to get a large collection of information required by a DII client without requiring the client to subsequently make repeated, possibly remote requests to objects in order to process the collection of information. Since the DII is optimized for the existing IR, this specification assumes that an IR layer will continue to service DII clients and thus does not attempt to provide this functionality in the MOF-compliant IR.

Figure 12.2 shows all of the metaclasses and relationships defined in the *BaseIDL* Package.

12.1.1.2 Typing

As mentioned earlier in this clause (“A Structural Comparison of the BaseIDL Package with the Existing IR” on page 203), the two critical elements of the BaseIDL Package supporting the typing of IR entities are the *Typed* and *IDLType* metaclasses. A *Typed* element references an *IDLType*, which has an attribute of type *TypeCode*.

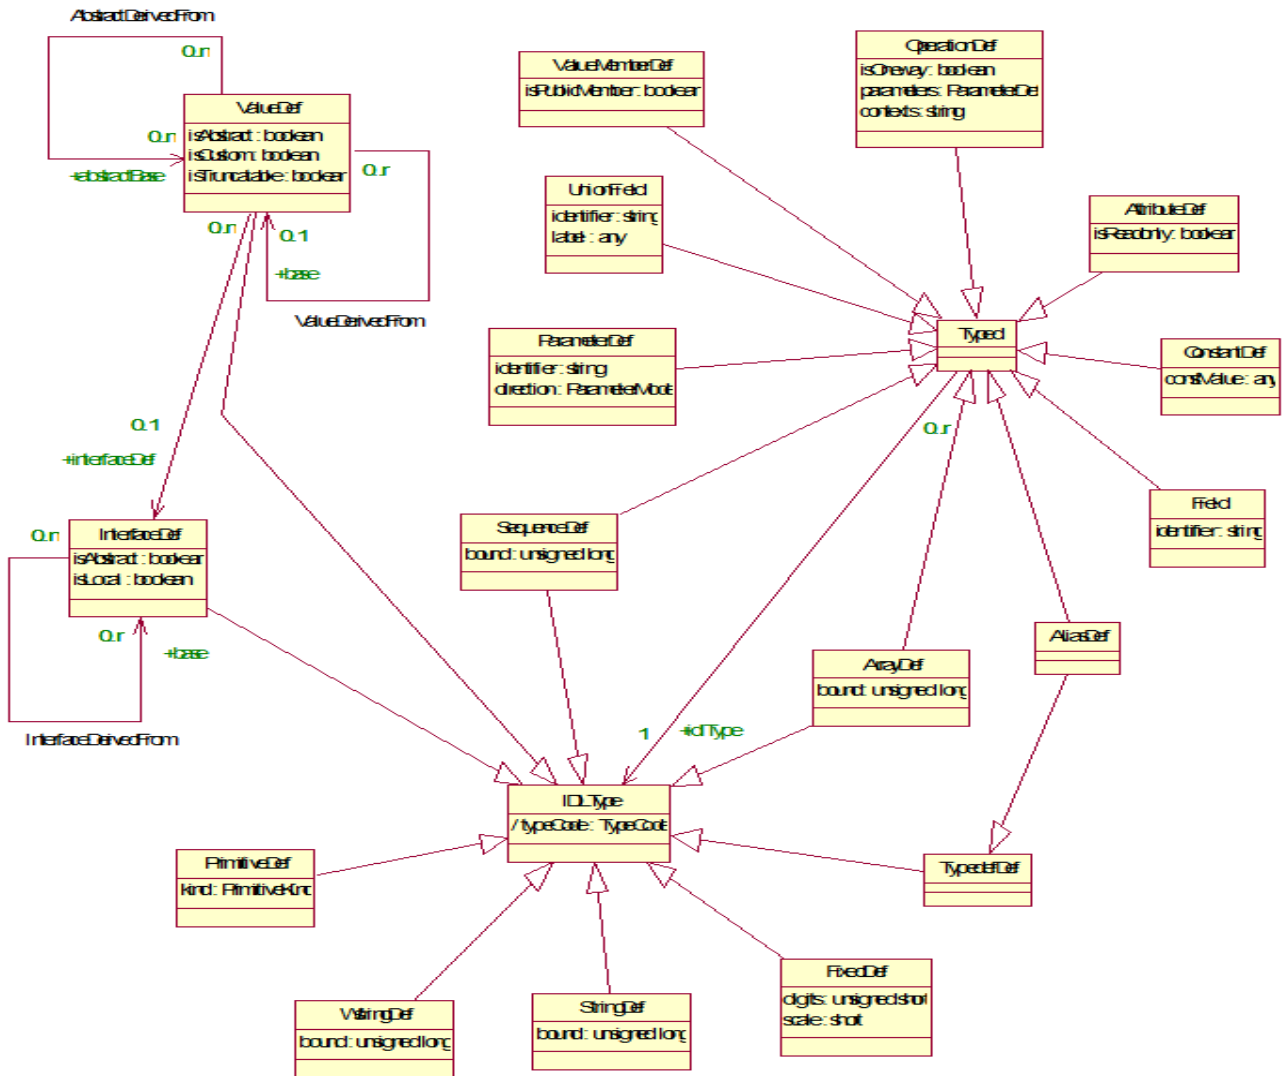


Figure 12.3- IDL Typing

12.1.1.3 Containment

Many elements in the metamodel descend from *Container* or *Contained*, in keeping with the structure of the original CORBA Interface Repository. As mentioned in the previous sub clause, the metamodel also derives *Container* from *Contained* so that an element that is logically a container and at the same time is defined in another container does not

have to inherit directly from both *Container* and *Contained*. However, this change requires that a constraint be written such that *ModuleDef* and only *ModuleDef* does not have to be defined in a *Container*. This constraint is included in the next sub clause on containment constraints. Figure 12.4 expresses the containment hierarchy.

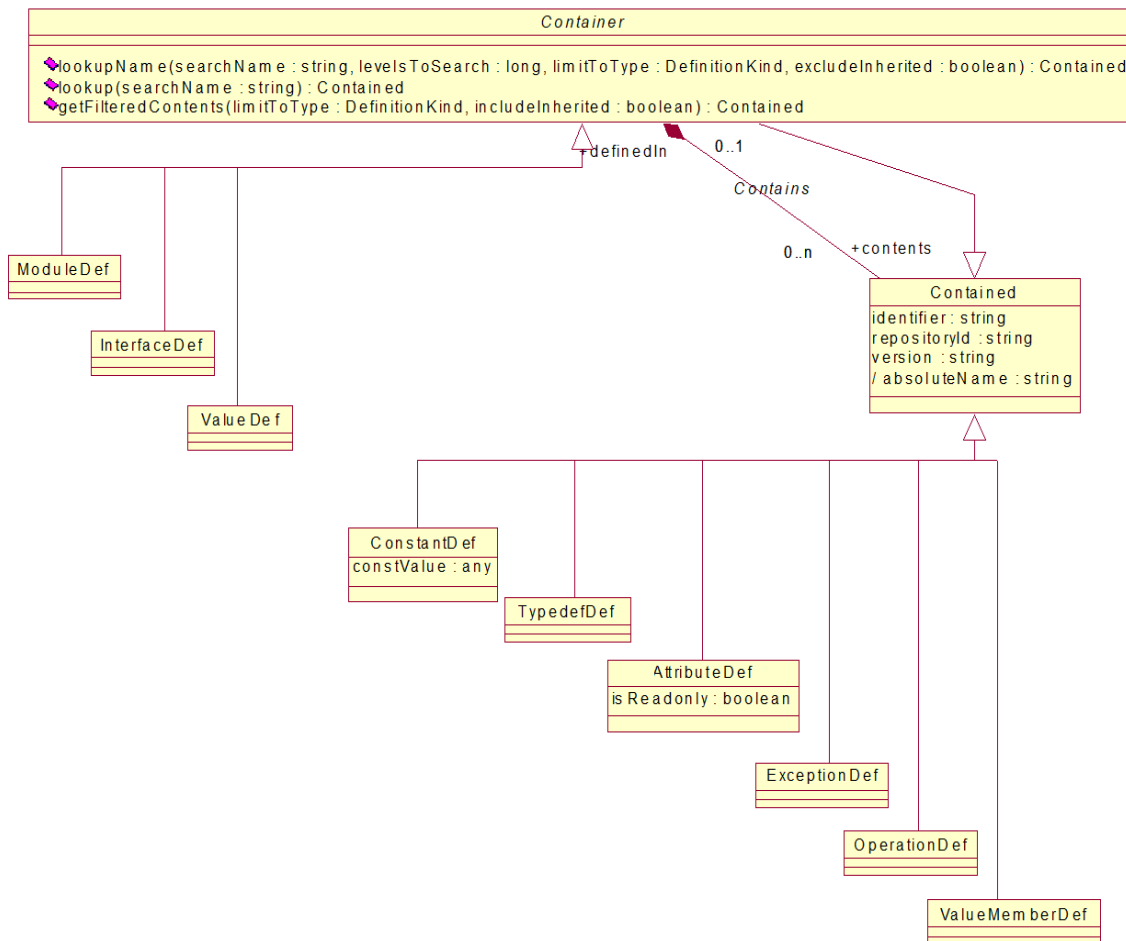


Figure 12.4- Containment Hierarchy

12.1.1.4 Containment Constraints

The Association between *Container* and *Contained* is named *Contains*. *Contains* is very general and is inherited by sub(meta)classes of *Container* and *Contained*. Unless further constrained, *Contains* would allow any *Container* to directly contain any *Contained* element. For example, a *ModuleDef* could contain an *OperationDef* and a *ValueDef* could contain an *InterfaceDef*. Clearly, the *Contains* Association must be constrained.

Figure 12.5 and Figure 12.6 express the containment constraints formally via the OMG’s Object Constraint Language (OCL). They also supplement the formal expressions with English natural language equivalents.

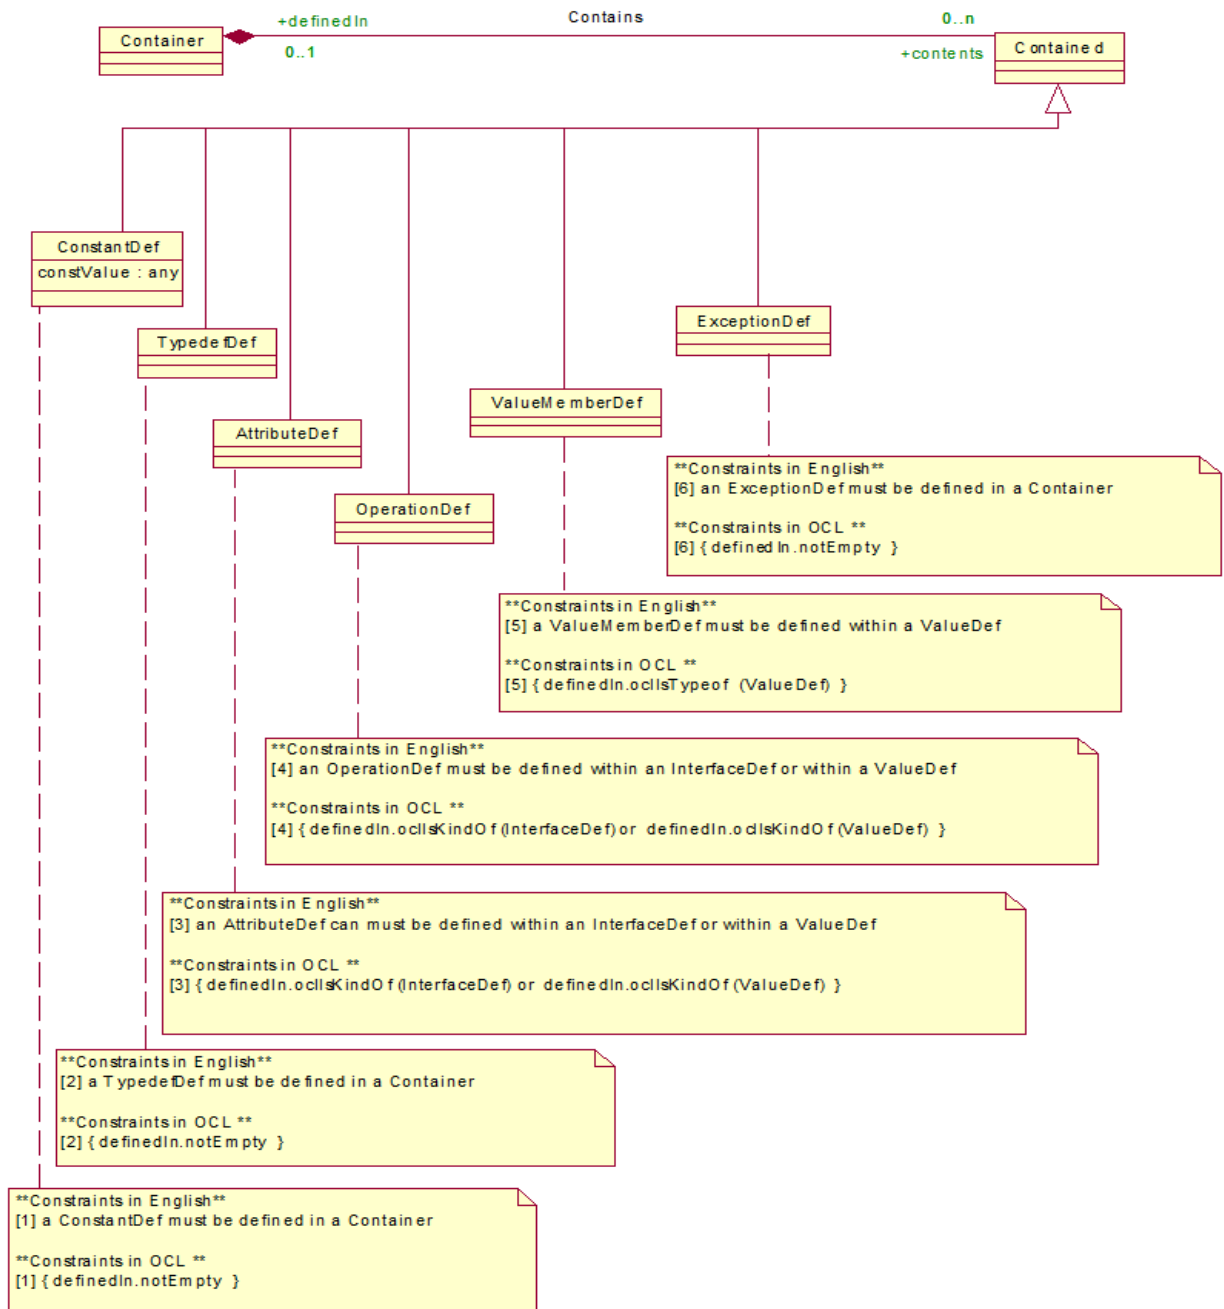


Figure 12.5- Containment Constraints--Subclasses of *Contained*

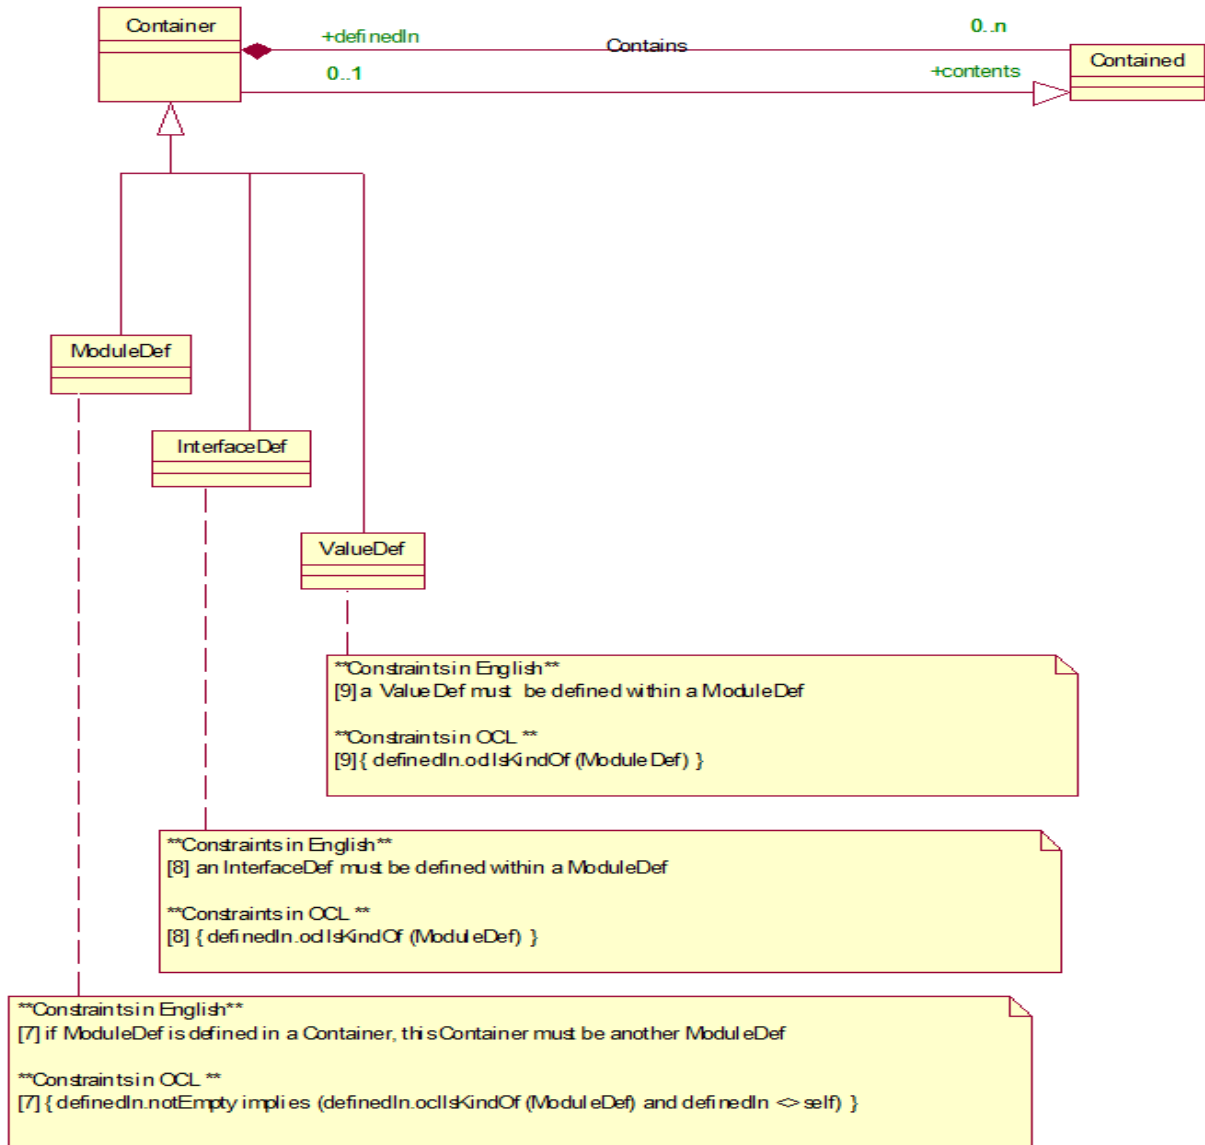


Figure 12.6- Containment Constraints--Subclasses of *Container*

12.1.1.5 Typedef and Type Derivations

Figure 12.7 expresses the hierarchy of derivatives of *Typedef* and *Typed*.

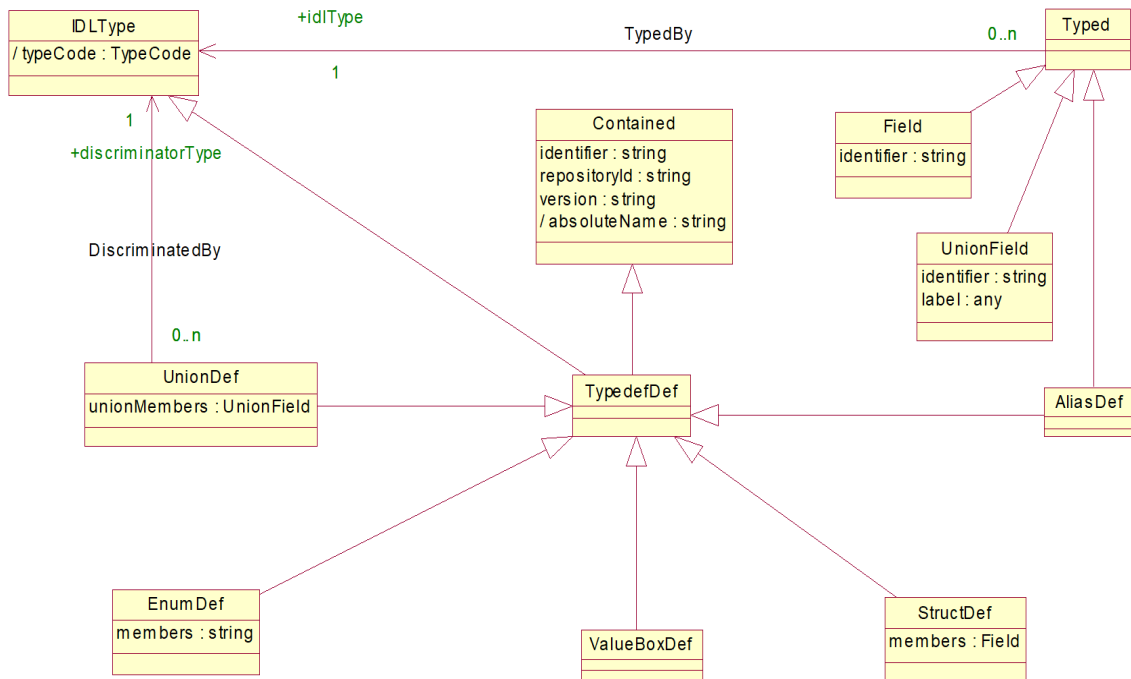


Figure 12.7- Derivations from *Typedef* and *Type*

12.1.1.6 Exceptions

Figure 12.8 shows the formal definition of the *ExceptionDef* metaclass. Note the inclusion of the newly-defined (in this specification) ability for attribute accessors and mutators to raise user-defined exceptions.

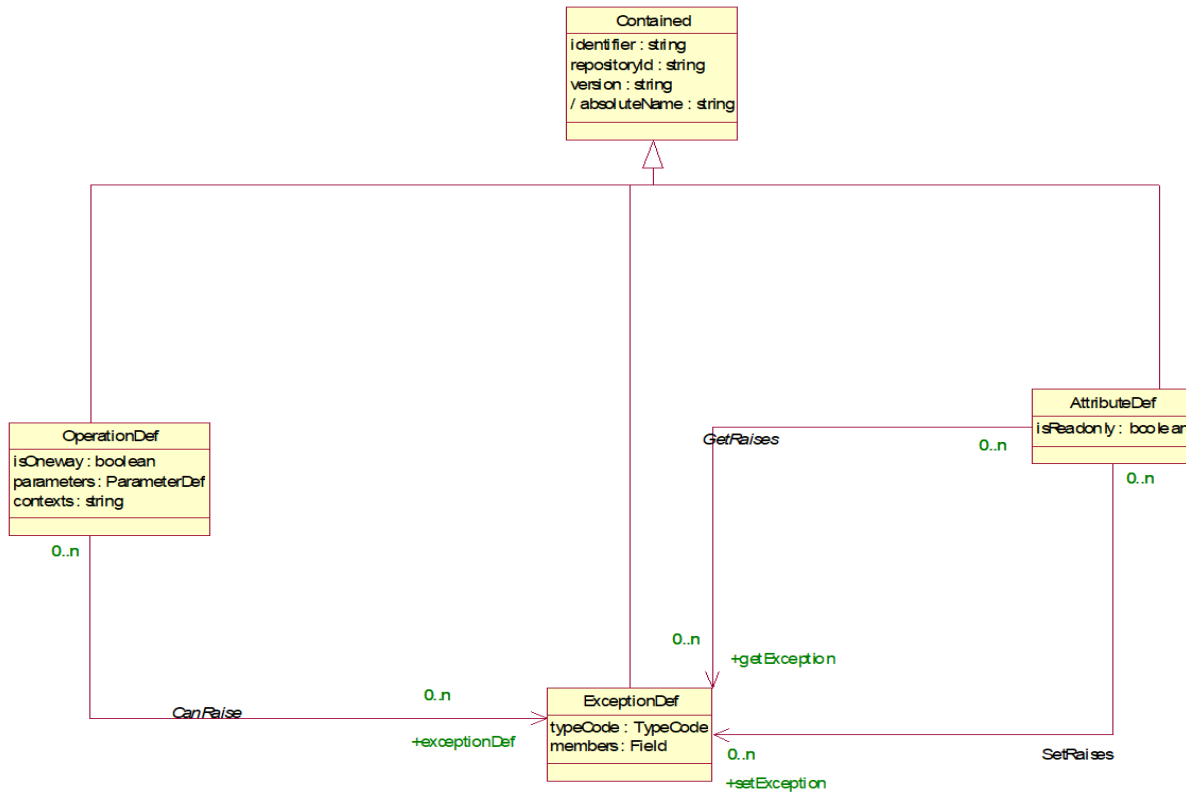


Figure 12.8- Exceptions

12.1.1.7 Value Types

CORBA 2.3 provided a model for types of objects that can be passed by value. The Objects By Value specification expanded the grammar of IDL and the structure of the Interface Repository to accommodate value types. Figure 12.9 focuses on the definition of value types in the MOF-compliant IR metamodel.

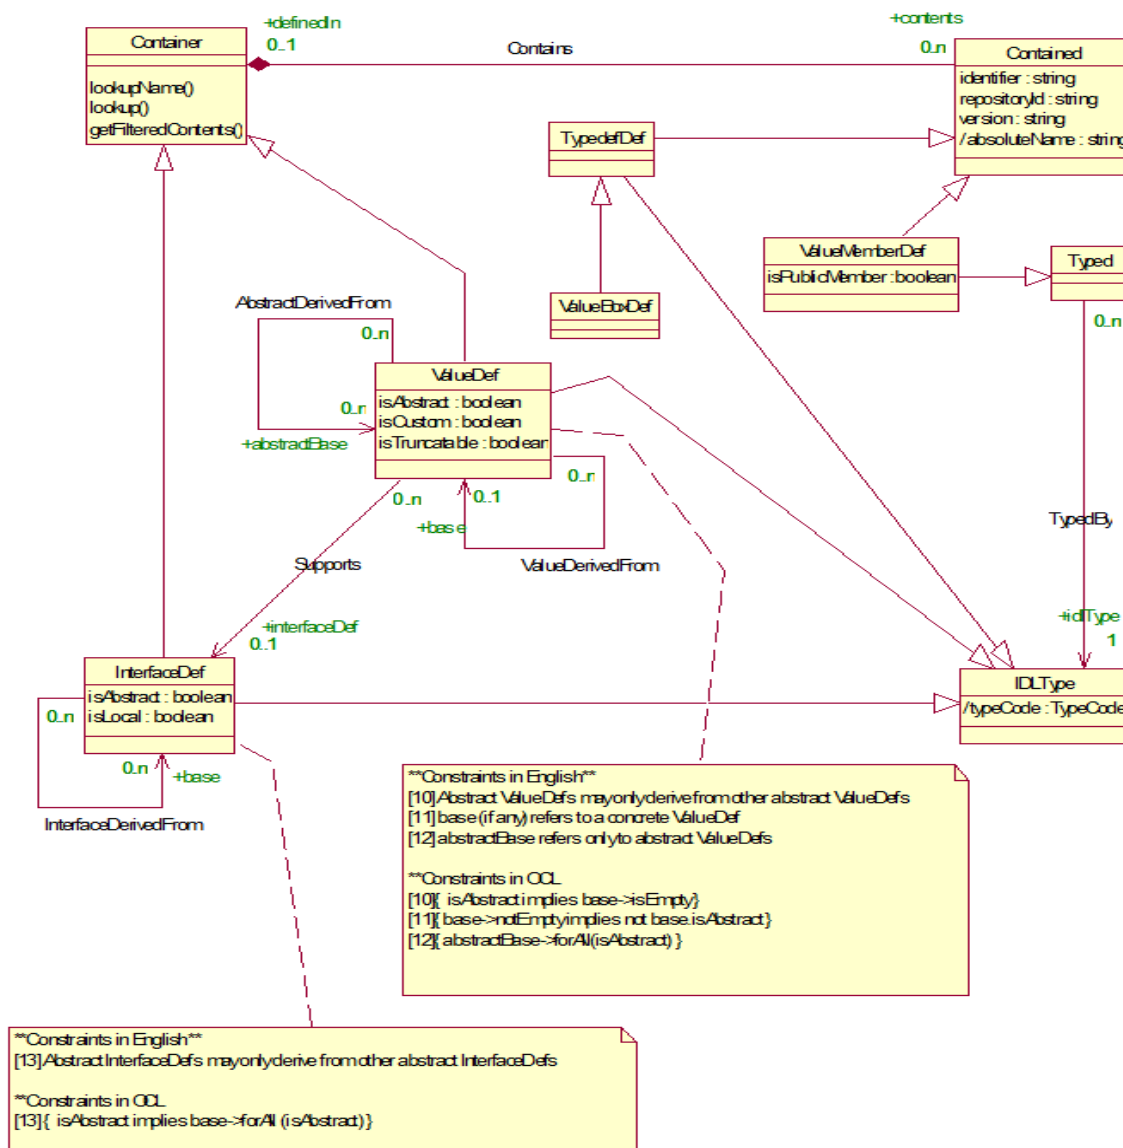


Figure 12.9- Value Types

12.1.1.8 Naming

Figure 12.10 focuses on the aspects of the metamodel that concern naming.

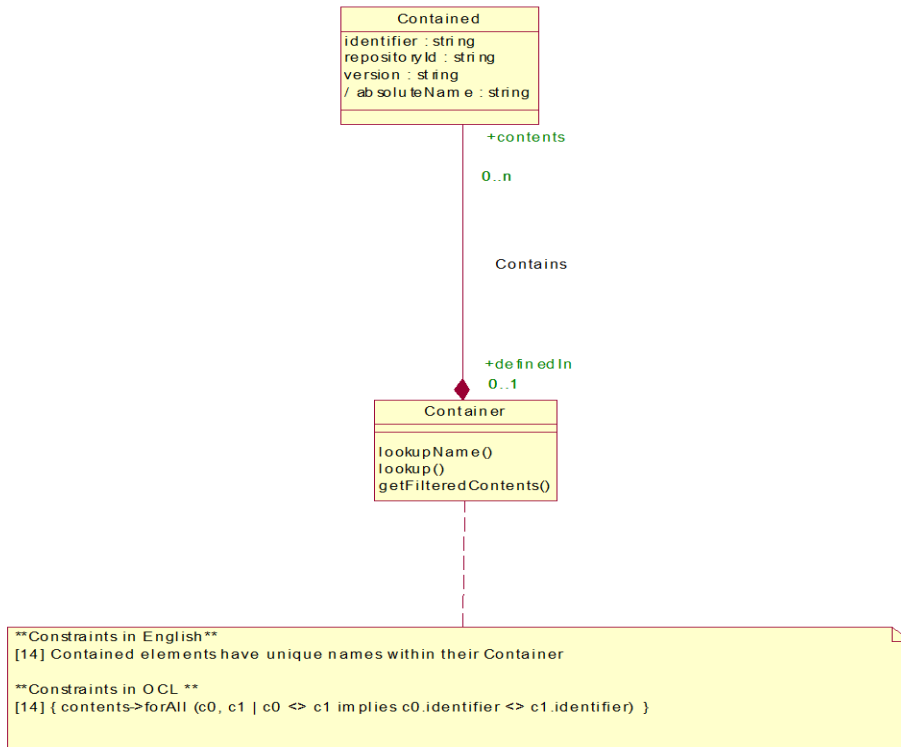


Figure 12.10- Naming

12.1.1.9 Operations

As mentioned earlier in this clause (“A Structural Comparison of the BaseIDL Package with the Existing IR” on page 203), the metamodel generally does not declare CRUD operations for the metaclasses, due to the fact that the MOF automatically generates such operations based on the structural metamodel. However, a few convenience operations are defined on the *Container* metaclass, as illustrated by Figure 12.11.

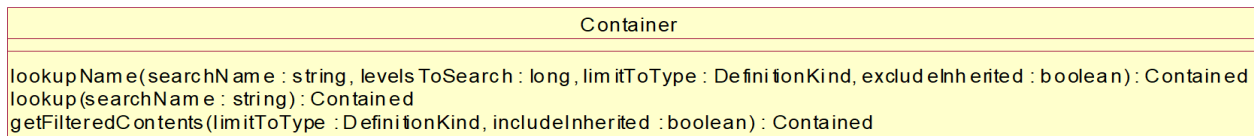


Figure 12.11- Convenience Operations

12.1.2 ComponentIDL Package

12.1.2.1 Overview

The following UML class diagram describes a metamodel representing the extensions to IDL defined by the CORBA Component Model. Just as these extensions are dependent on the base IDL defined in the CORBA Core, so is this metamodel dependent on a metamodel representing the base IDL.

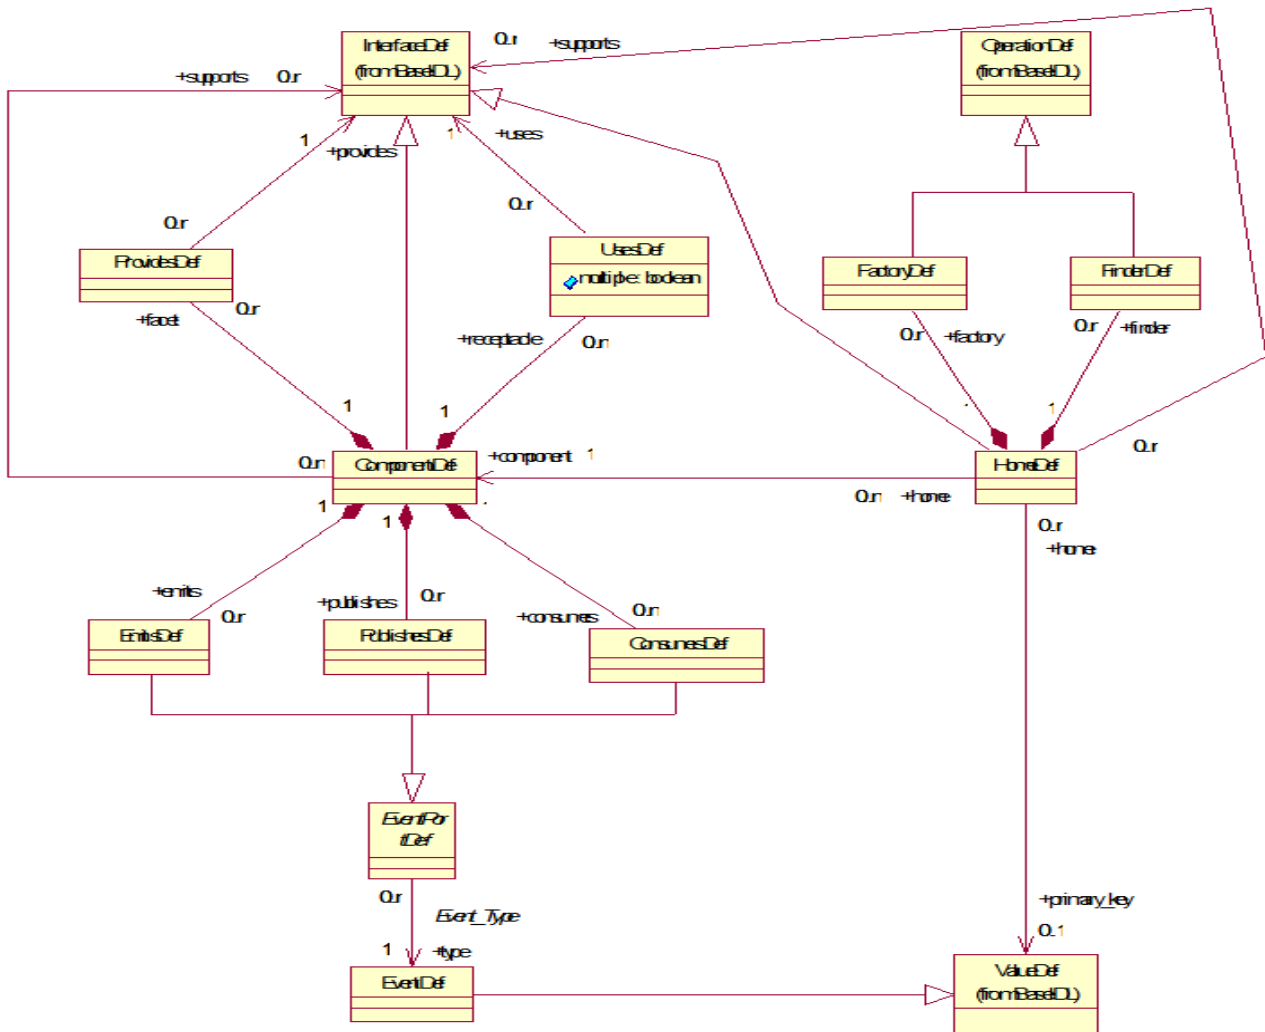


Figure 12.12- ComponentIDL Package - Main Diagram

12.1.2.2 Containers and Contained Elements

The following UML class diagram (Figure 12.13) describes the derivation of the metamodel elements from the BaseIDL *Container* and *Contained* elements:

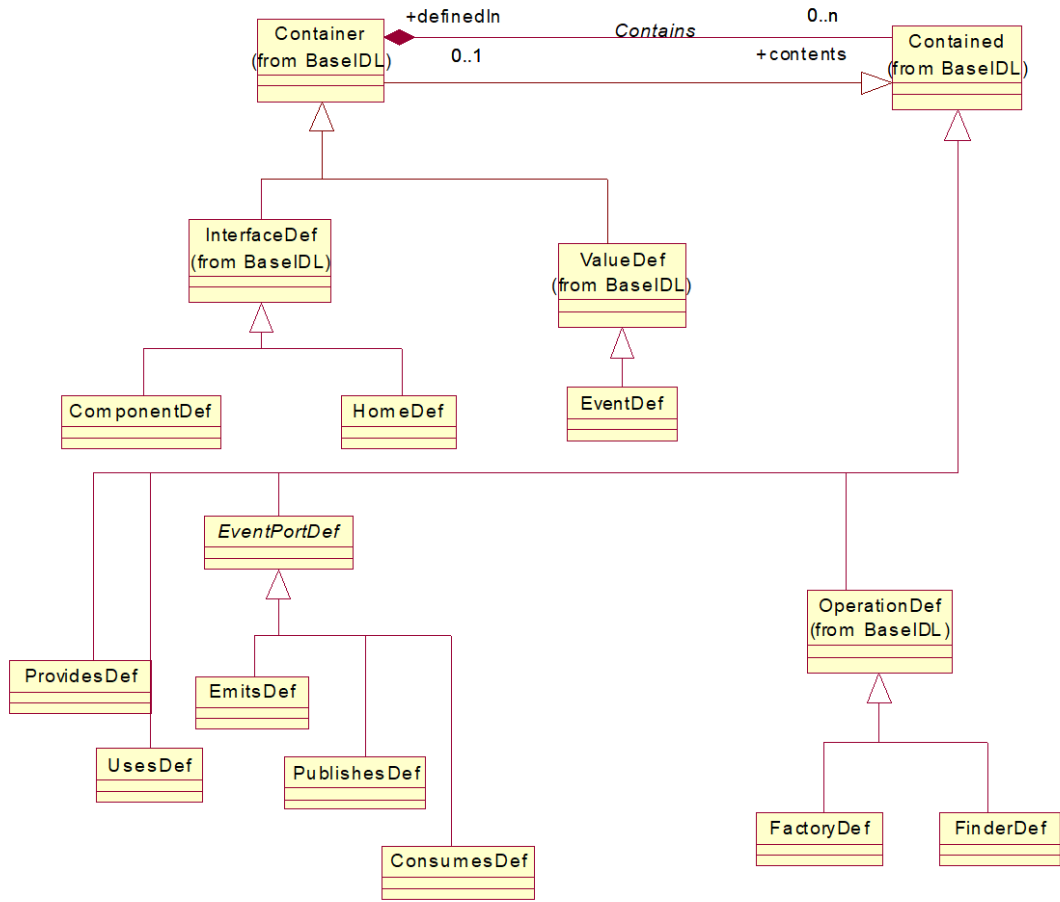


Figure 12.13- Containment Hierarchy

Each of the subtypes of *Contained* shown in Figure 12.13 can only be defined within certain subtypes of *Container*. Figure 12.14 formally specifies these constraints via the OMG’s Object Constraint Language (OCL), and supplements the OCL by expressing the constraints in natural language for the benefit of readers who are not familiar with OCL.

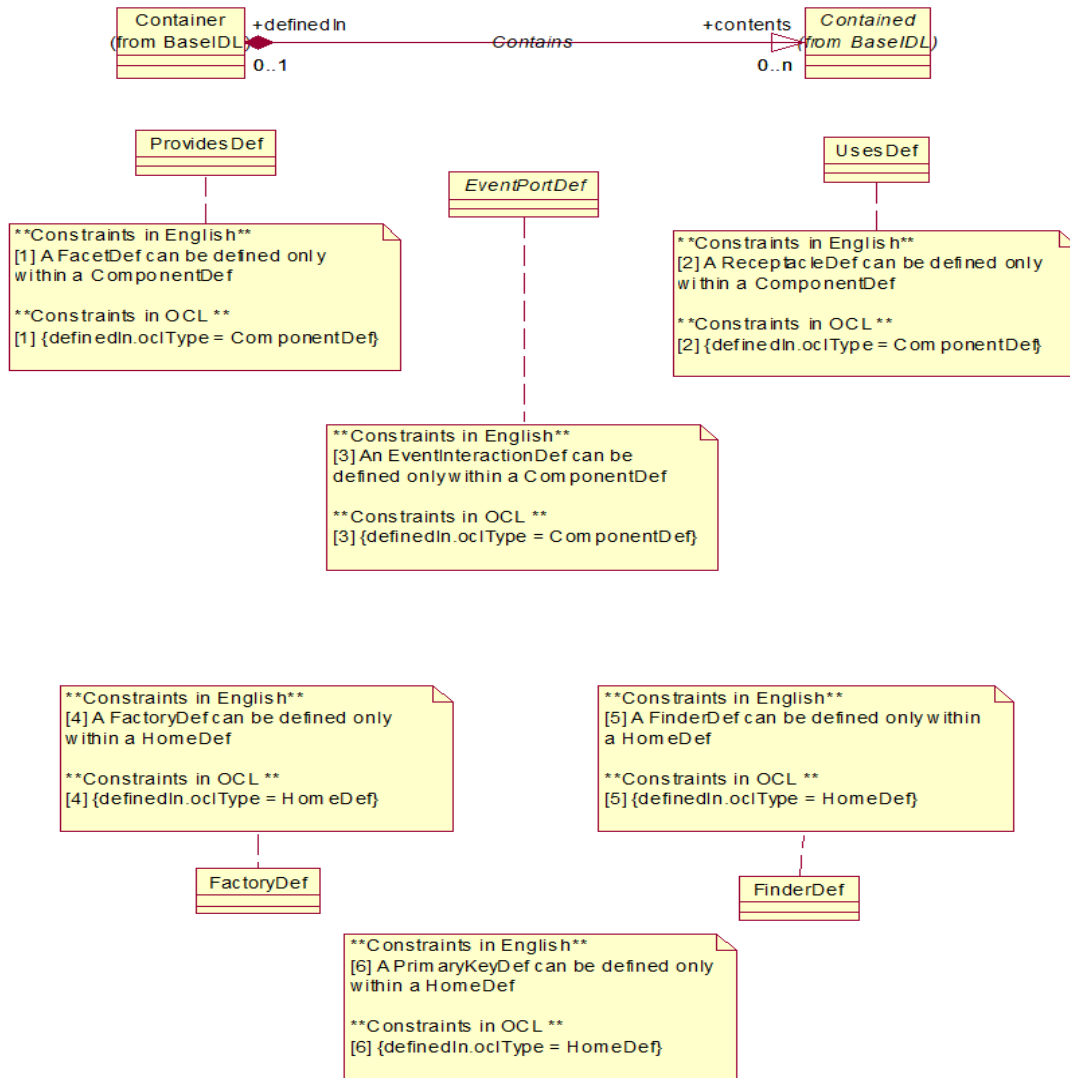


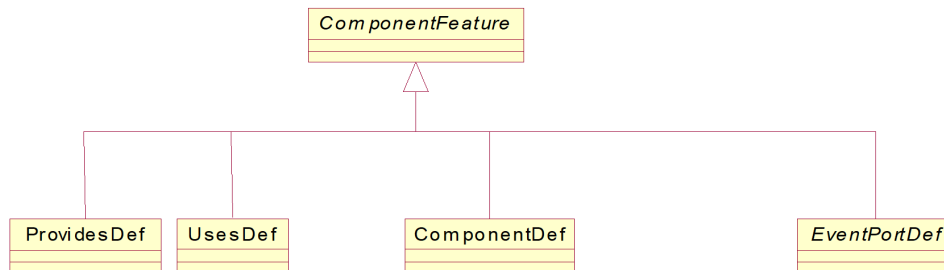
Figure 12.14- Constraints on Containment of Elements Defined In ComponentDef

An instance of *ComponentDef* describes a CORBA component in an abstract manner. The definition contains a description of all features of a component that are visible from the outside. More precise, it defines all interfaces including interfaces that are implicit or used for event communication. In detail, the features of component that are visible to the outside are:

- The component equivalent interface, containing all implicit operations, operations and attributes that are inherited by a component (also from supported interfaces), and attributes defined inside the component.
- The facets of a component; that is, all interfaces that are provided by the component to the outside.

- The receptacles of a component; that is, all interfaces that are used by a component.
- The events, which a component can emit, publish, or consume.

If a component is going to be implemented, all these features must be handled by the component implementation. To provide a common basis for defining the related implementation definitions (as part of CIF) the abstract metaclass *ComponentFeature* is defined. The metaclasses *ComponentDef*, *ProvidesDef*, *UsesDef*, and *EventPortDef* are defined as subclasses of the metaclass *ComponentFeature*.



All of *ComponentDef*'s composition Associations shown in the main diagram (Figure 12.12 on page 214) are derived from the *BaseIDL* metamodel's *Contains* Association between *Container* and *Contained*. As shown by Figure 12.13 on page 215, *ComponentDef* inherits that Association from *InterfaceDef*, which inherits it from *Container*.

The following class diagram (Figure 12.15 on page 218) details these derived Associations. A “/” prefix in an Association name denotes that the Association is derived, and sets the MOF's “isDerived” property for the Association. The constraints for each of the derived Associations are expressed in the OMG's Object Constraint Language and declare how the Associations are derived from the *Contains* Association.

The <<*implicit*>> stereotype is a standard UML stereotype that designates the Association as conceptual rather than manifest. An <<*implicit*>> Association is ignored when generating IDL for the metamodel via the MOF-IDL mapping. It is also ignored when deriving the XML DTD for the metamodel via the MOF-XML mapping specified by the XMI specification. The *Contains* association is sufficient for generating the accessor methods in the IDL allowing the containments to be traversed. If these Associations were not marked as <<*implicit*>>, then additional accessor methods would be generated to do the more focused traversals that they conceptualize. In the judgement of the submitters the generation of these additional accessor methods would expand the footprint of the IDL interfaces more than is warranted, given that the containments can be traversed by the single inherited *Contains* Association.

The fact that these <<*implicit*>> Associations are ignored when generating the IDL for the metamodel does not mean that they have no bearing on the contents of a repository. The “reflective” interfaces that all MOF metaobjects inherit have an operation called *metaObject* that returns a metaobject. This metaobject is part of the metamodel rather than part of a model; in other words, it is actually a meta-metaobject that is part of the description of the metamodel. The definitions of the <<*implicit*>> Associations in which a metaobject participates would be available via this meta-metaobject. The multiplicity constraints of these Associations would be available as well. Thus, for example, the fact that a *ComponentDef* aggregates zero or more *UsesDef* metaobjects is discoverable through such meta-metaobjects and thus serves as a formal constraint on the *Contains* Association from which the aggregation is derived.

Furthermore, when the state of the metamodel is streamed in conformance with the DTD for the MOF meta-metamodel, the state that specifies the <<*implicit*>> Associations are part of the stream. The DTD for the MOF meta-metamodel is contained in the XMI specification. XML streams conforming to that DTD and which contain the state of the IR metamodel are included in “MOF DTDs and IDL for the Interface Repository Metamodel” on page 223.

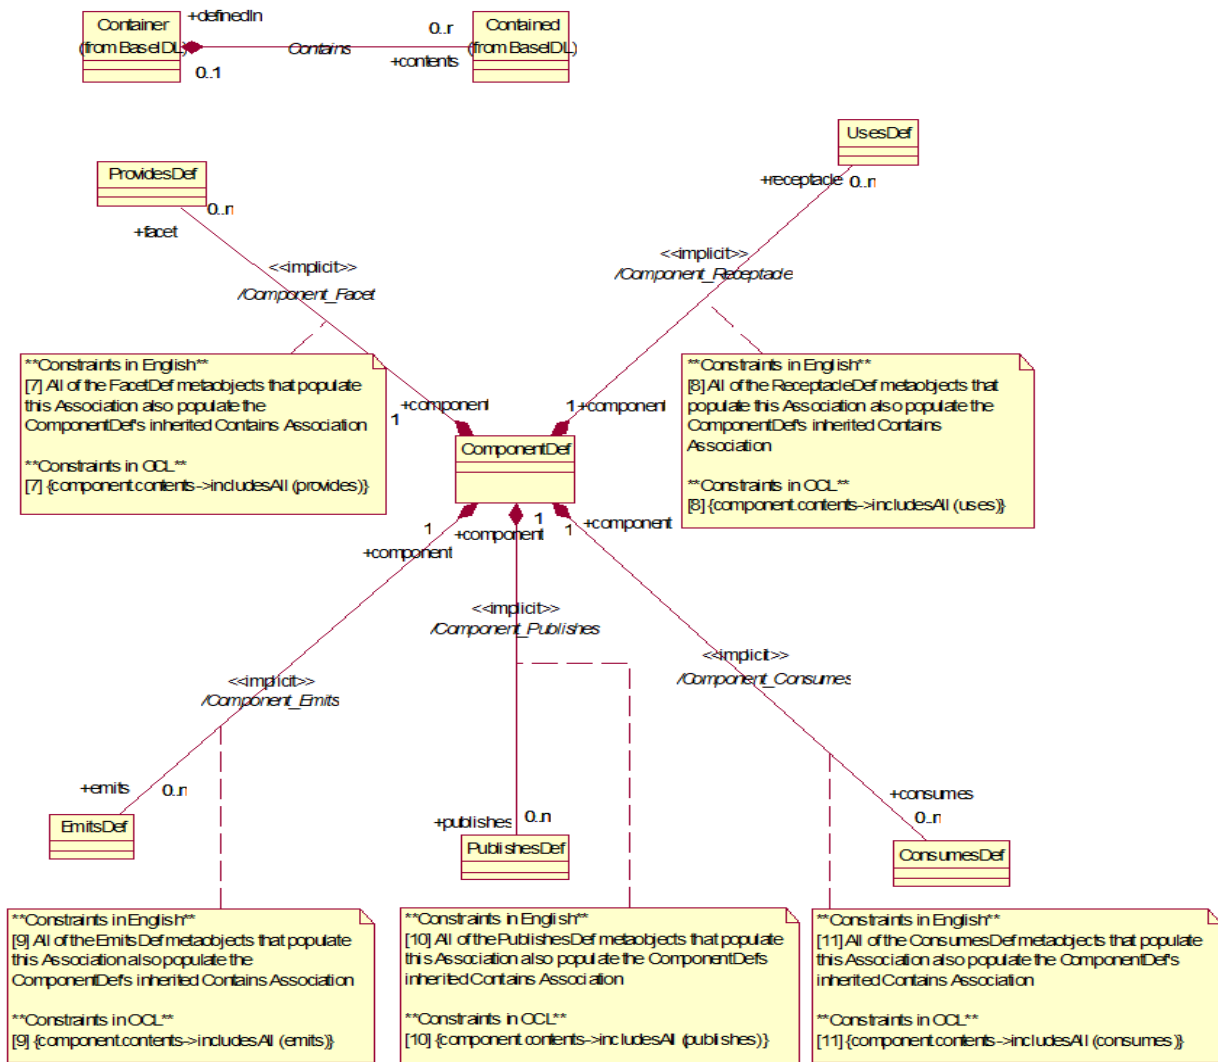


Figure 12.15- Implicit Derived Containments with ComponentDef as the Composite

HomeDef's composition Associations also are derived from the *Contains* Association. As shown in Figure 12.13 on page 215, *HomeDef* descends from *Container*. All of the components of its composition Associations descend from *Contained*. As with the derived Associations in which *ComponentDef* plays the composite role, the derived Associations in which *HomeDef* plays the composite role are marked as `<<implicit>>` to prevent excess IDL generation. Figure 12.16 formally defines the constraints that define the semantics of the derivations.

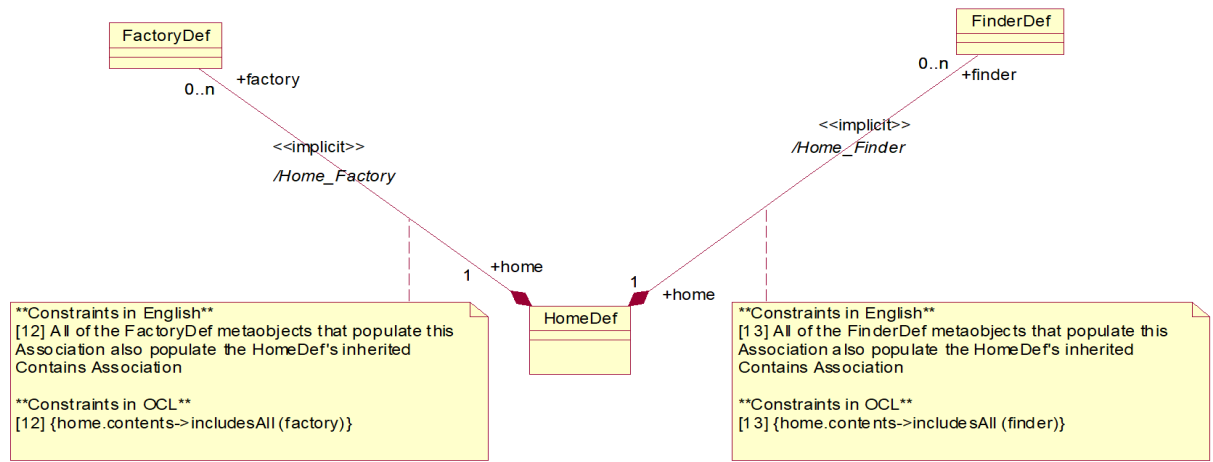


Figure 12.16- Implicit Derived Containments with HomeDef as the Composite

12.1.2.3 ValueDef Constraints EventsDef

The *ValueDef* metaclass, which is part of the *BaseIDL* Package, participates in a number of Associations defined by the *ComponentIDL* Package. The *emits*, *publishes*, and *consumes* declarations that are part of the component model IDL extensions all reference a *ValueDef*. Furthermore, the *primaryKey* declaration within *home* declarations references a *ValueDef*. However, the IDL type of the *ValueDef* is constrained, as explained in the *Component Model* clause.

Figure 12.17 expresses the *ValueDef* constraints formally. Note that it uses an OCL technique of defining a side-effect free operation to support recursion, which is required to traverse the transitive closure of a *ValueDef*'s inheritance hierarchy.

The Component IDL metamodel is changed to introduce the new metatype **eventtype**. This metatype is introduced as a metaclass *EventDef*, which is a specialization of *ValueDef*. Inheritance for instances of *EventDef* is allowed from instances of *ValueDef* and *EventDef*; however, instances of *ValueDef* are not allowed to inherit from instances of *EventDef*.



Figure 12.17- The new metatype eventtype

The former metaclass *EventDef* as contained in the metamodel for *ComponentIDL* in orbos/99-07-02 is renamed to *EventPortDef*. The Association between *EventPortDef* and *ValueDef* is removed. Instead, there is a similar Association defined between *EventPortDef* and *EventDef*. The metamodel for Component IDL is shown in Figure 12.12 on page 214.

12.1.2.4 Additional Type and Inheritance Constraints

Figure 12.18 and Figure 12.19 define additional constraints on the *ComponentDef*, *HomeDef*, *FactoryDef*, and *FinderDef* metaclasses.

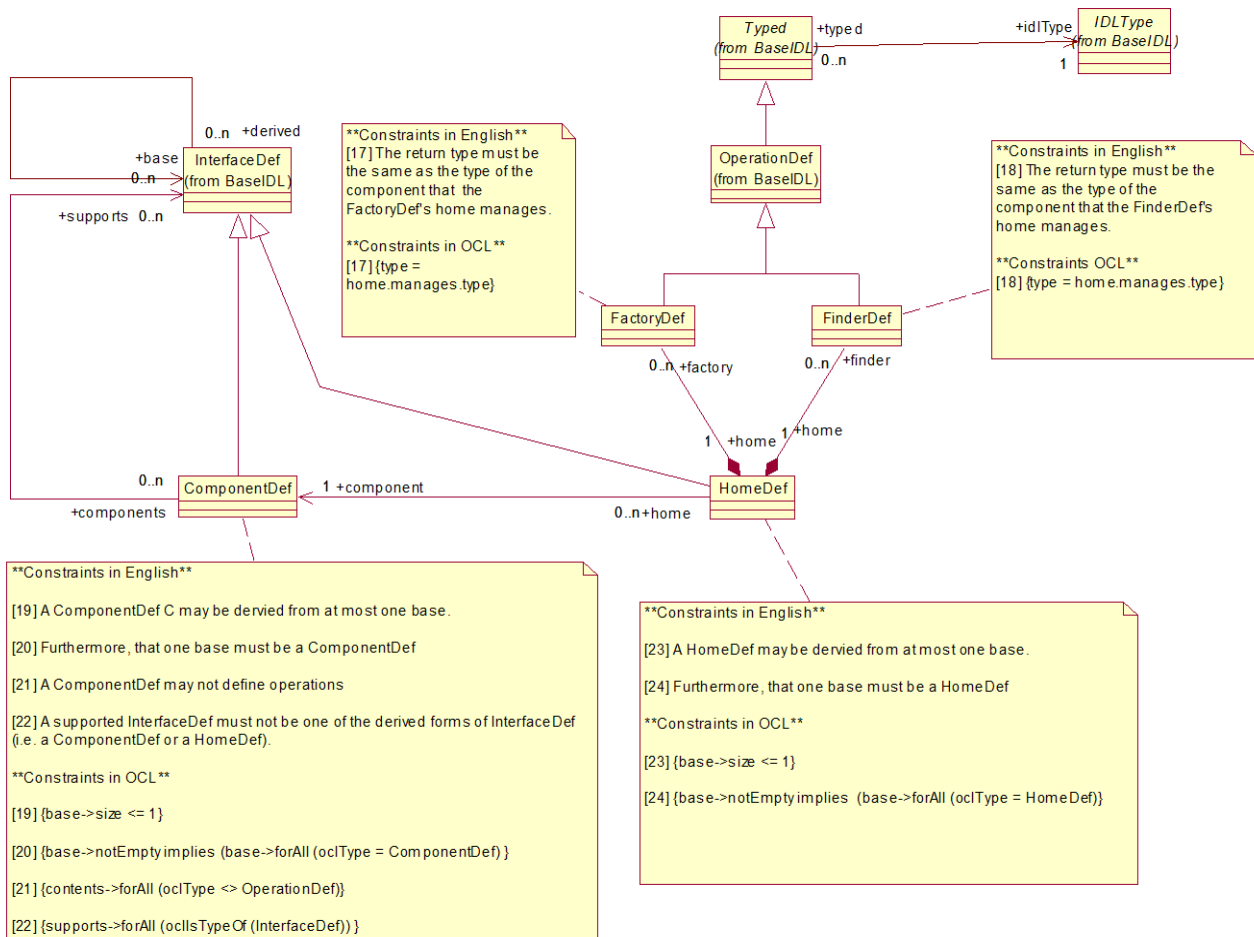


Figure 12.18- Additional Component, Home, Factory, and Finder Constraints

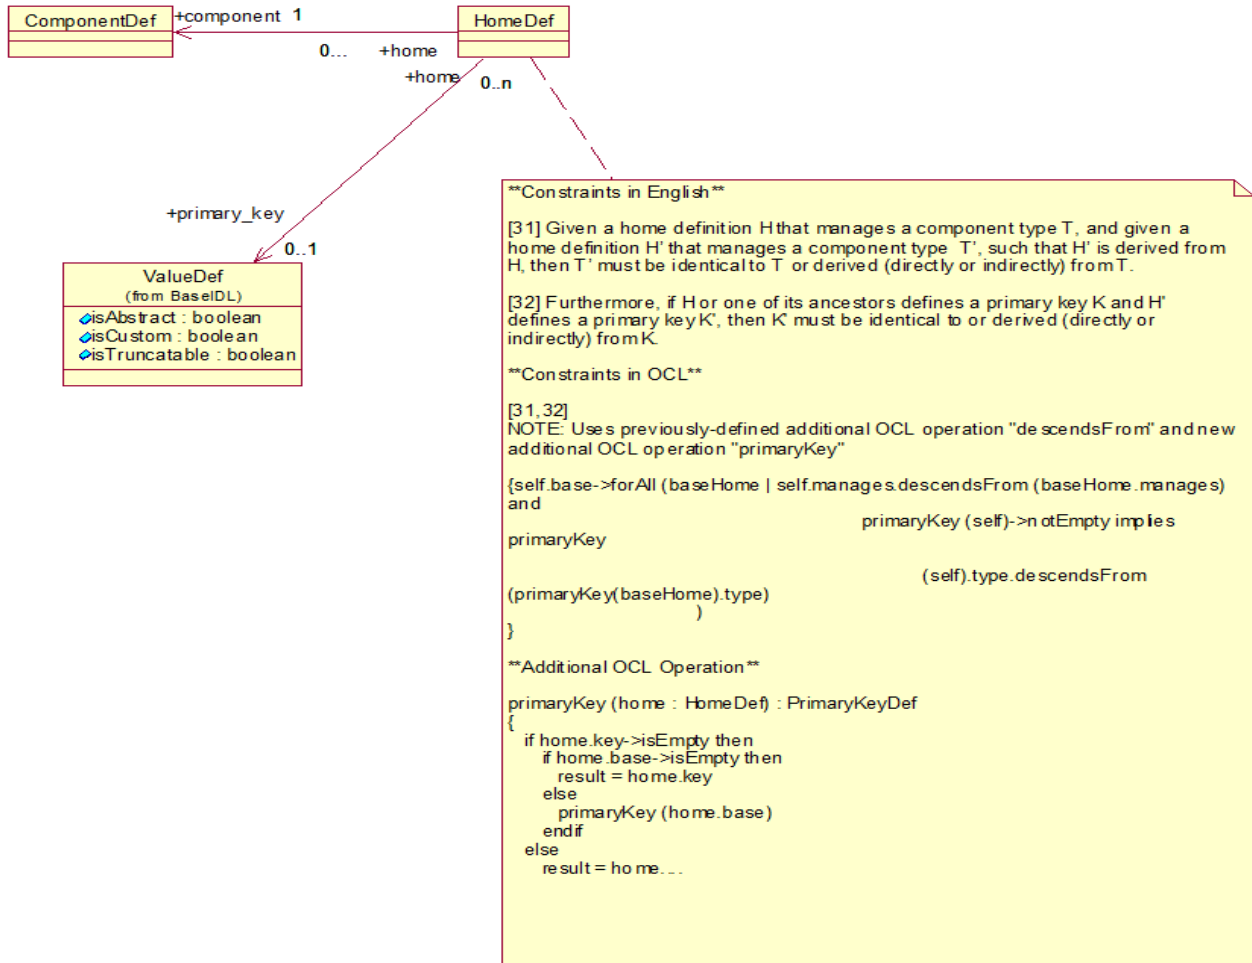


Figure 12.19- Home and Primary Key Constraints

12.1.2.5 Constraints on Basic Components

The CORBA Component Model defines the notion of *basic* components. The *ComponentDef* metaclass has an attribute named *isBasic*. The fact that a component is basic can actually be computed from the component declaration — if the component observes certain constraints, it is basic. Thus, strictly speaking, the *isBasic* attribute is not necessary. However, the attribute greatly simplifies the process of determining whether a component definition is basic.

Given the circumstances, it would seem appropriate to define the *isBasic* attribute as a derived one. In the MOF and UML, *isDerived* is an attribute of *Attribute* that indicates that the information can be computed from other information in the model. However, the XMI standard specifies that the state of derived attributes is not deposited in XMI/XML streams representing models. Thus, if the *isBasic* attribute were marked as derived, the state of the attribute would not appear in XMI streams representing CORBA-based object models. The authors have therefore decided not to mark the *isBasic* attribute of *ComponentDef* as derived.

The constraints on basic components are modeled formally as shown in Figure 12.20.

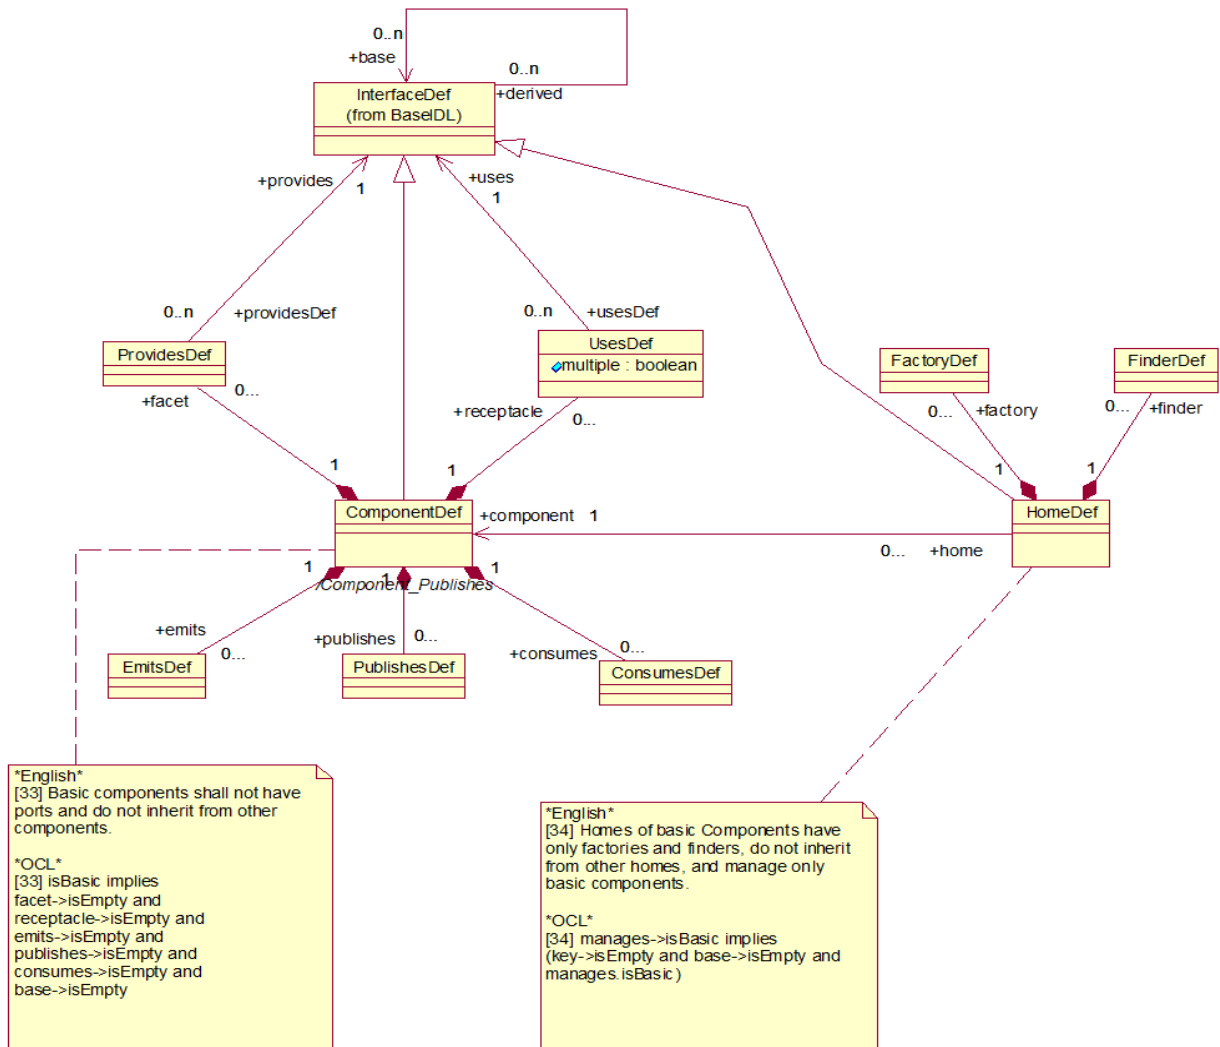


Figure 12.20- Constraints on Basic Components

12.2 Conformance Criteria

This clause identifies the conformance points required for compliant implementations of the interface repository metamodel architecture.

12.2.1 Conformance Points

In the previous sub clause, the MOF metamodel of the Interface Repository is defined. The following sub clause defines the XMI format for the exchange of Interface Repository metadata and the IDL for a MOF-compliant Interface Repository. Support for the generation and consumption of the XMI metadata and for the MOF-compliant IDL is optional.

12.3 MOF DTDs and IDL for the Interface Repository Metamodel

The XMI DTDs and IDL for the Interface Repository metamodel are presented in this sub clause. The DTDs are generated by applying the MOF-XML mapping defined by the XMI specification to the MOF-compliant metamodel described in “Introduction” on page 203. The IDL is generated by applying the MOF-IDL mapping defined in the MOF specification to the metamodels and was validated using the IDL compilers.

The IDL requires the inclusion of the reflective interfaces defined in the Meta Object Facility (MOF) specification (<http://www.omg.org/technology/documents/formal/mof.htm>).

12.3.1 XMI DTD

```
<!-- _____ -->
<!-- -->
<!-- XMI is the top-level XML element for XMI transfer text -->
<!-- _____ -->

<!ELEMENT XMI (XMI.header?, XMI.content?, XMI.difference*,
    XMI.extensions*)>
<!ATTLIST XMI
    xmi.version CDATA #FIXED "1.1"
    timestamp CDATA #IMPLIED
    verified (true|false) #IMPLIED>

<!-- _____ -->
<!-- -->
<!-- XMI.header contains documentation and identifies the model,
<!-- metamodel, and metamodel
<!-- _____ -->

<!ELEMENT XMI.header (XMI.documentation?, XMI.model*, XMI.metamodel*,
    XMI.metamodel*, XMI.import*)>

<!-- _____ -->
<!-- -->
<!-- documentation for transfer data -->
<!-- _____ -->

<!ELEMENT XMI.documentation (#PCDATA | XMI.owner | XMI.contact |
    XMI.longDescription | XMI.shortDescription |
    XMI.exporter | XMI.exporterVersion |
    XMI.notice)*>
<!ELEMENT XMI.owner ANY>
```

```

<ELEMENT XMI.contact ANY>
<ELEMENT XMI.longDescription ANY>
<ELEMENT XMI.shortDescription ANY>
<ELEMENT XMI.exporter ANY>
<ELEMENT XMI.exporterVersion ANY>
<ELEMENT XMI.exporterID ANY>
<ELEMENT XMI.notice ANY>

<!-- _____ -->
<!-- -->
<!-- XMI.element.att defines the attributes that each XML element -->
<!-- that corresponds to a metamodel class must have to conform to -->
<!-- the XMI specification. -->
<!-- _____ -->

<ENTITY % XMI.element.att
    'xmi.id ID #IMPLIED xmi.label CDATA #IMPLIED xmi.uuid
    CDATA #IMPLIED '>

<!-- _____ -->
<!-- -->
<!-- XMI.link.att defines the attributes that each XML element that -->
<!-- corresponds to a metamodel class must have to enable it to -->
<!-- function as a simple XLink as well as refer to model -->
<!-- constructs within the same XMI file. -->
<!-- _____ -->

<ENTITY % XMI.link.att
    'href CDATA #IMPLIED xmi.idref IDREF #IMPLIED xmi:link
    CDATA #IMPLIED xlink:inline (true|false) #IMPLIED
    xlink:actuate (show|user) #IMPLIED xlink:content-role
    CDATA #IMPLIED xlink:title CDATA #IMPLIED xlink:show
    (embed|replace|new) #IMPLIED xlink:behavior CDATA
    #IMPLIED '>

<!-- _____ -->
<!-- -->
<!-- XMI.model identifies the model(s) being transferred -->
<!-- _____ -->

<ELEMENT XMI.model ANY>
<!ATTLIST XMI.model %XMI.link.att;
    xmi.name CDATA #REQUIRED
    xmi.version CDATA #IMPLIED>

<!-- _____ -->
<!-- -->
<!-- XMI.metamodel identifies the metamodel(s) for the transferred -->
<!-- data -->
<!-- _____ -->

```



```

<ELEMENT XMI.metamodel ANY>
<!ATTLIST XMI.metamodel %XMI.link.att;
    xmi.name CDATA #REQUIRED
    xmi.version CDATA #IMPLIED>

<!-- _____ -->
<!-- -->
<!-- XMI.metamodel identifies the metamodel(s) for the -->
<!-- transferred data -->
<!-- _____ -->

<ELEMENT XMI.metamodel ANY>
<!ATTLIST XMI.metamodel %XMI.link.att;
    xmi.name CDATA #REQUIRED
    xmi.version CDATA #IMPLIED>

<!-- _____ -->
<!-- -->
<!-- XMI.import identifies imported metamodel(s) -->
<!-- _____ -->

<ELEMENT XMI.import ANY>
<!ATTLIST XMI.import %XMI.link.att;
    xmi.name CDATA #REQUIRED
    xmi.version CDATA #IMPLIED>

<!-- _____ -->
<!-- -->
<!-- XMI.content is the actual data being transferred -->
<!-- _____ -->

<ELEMENT XMI.content ANY>

<!-- _____ -->
<!-- -->
<!-- XMI.extensions contains data to transfer that does not conform -->
<!-- to the metamodel(s) in the header -->
<!-- _____ -->

<ELEMENT XMI.extensions ANY>
<!ATTLIST XMI.extensions
    xmi.extender CDATA #REQUIRED>

<!-- _____ -->
<!-- -->
<!-- extension contains information related to a specific model -->
<!-- construct that is not defined in the metamodel(s) in the header -->
<!-- _____ -->

<ELEMENT XMI.extension ANY>
<!ATTLIST XMI.extension %XMI.element.att; %XMI.link.att;

```

```
xmi.extender CDATA #REQUIRED
xmi.extenderID CDATA #IMPLIED>
```

```
<!-- _____ -->
<!-- -->
<!-- XMI.difference holds XML elements representing differences to a -->
<!-- base model -->
<!-- _____ -->
```

```
<ELEMENT XMI.difference (XMI.difference | XMI.delete | XMI.add |
    XMI.replace)*>
<!ATTLIST XMI.difference %XMI.element.att; %XMI.link.att;>
```

```
<!-- _____ -->
<!-- -->
<!-- XMI.delete represents a deletion from a base model -->
<!-- _____ -->
```

```
<ELEMENT XMI.delete EMPTY>
<!ATTLIST XMI.delete %XMI.element.att; %XMI.link.att;>
```

```
<!-- _____ -->
<!-- -->
<!-- XMI.add represents an addition to a base model -->
<!-- _____ -->
```

```
<ELEMENT XMI.add ANY>
<!ATTLIST XMI.add %XMI.element.att; %XMI.link.att;
    xmi.position CDATA "-1">
```

```
<!-- _____ -->
<!-- -->
<!-- XMI.replace represents the replacement of a model construct -->
<!-- with another model construct in a base model -->
<!-- _____ -->
```

```
<ELEMENT XMI.replace ANY>
<!ATTLIST XMI.replace %XMI.element.att; %XMI.link.att;
    xmi.position CDATA "-1">
```

```
<!-- _____ -->
<!-- -->
<!-- XMI.reference may be used to refer to data types not defined in -->
<!-- the metamodel -->
<!-- _____ -->
```

```
<ELEMENT XMI.reference ANY>
<!ATTLIST XMI.reference %XMI.link.att;>
```

```
<!-- _____ -->
```

```

<!--                                     -->
<!-- This sub clause contains the declaration of XML elements           -->
<!-- representing data types                                           -->
<!-- _____ -->

<IELEMENT XMI.TypeDefinitions ANY>
<IELEMENT XMI.field ANY>
<IELEMENT XMI.seqItem ANY>
<IELEMENT XMI.octetStream (#PCDATA)>
<IELEMENT XMI.unionDiscrim ANY>

<IELEMENT XMI.enum EMPTY>
<!ATTLIST XMI.enum xmi.value CDATA #REQUIRED>

<IELEMENT XMI.any ANY>
<!ATTLIST XMI.any %XMI.link.att;
          xmi.type CDATA #IMPLIED
          xmi.name CDATA #IMPLIED>

<IELEMENT XMI.CorbaTypeCode (XMI.CorbaTcAlias|XMI.CorbaTcStruct|
          XMI.CorbaTcSequence|XMI.CorbaTcArray|XMI.CorbaTcEnum|
          XMI.CorbaTcUnion|XMI.CorbaTcExcept|XMI.CorbaTcString|
          XMI.CorbaTcWstring|XMI.CorbaTcShort|XMI.CorbaTcLong|
          XMI.CorbaTcUshort|XMI.CorbaTcUlong|XMI.CorbaTcFloat|
          XMI.CorbaTcDouble|XMI.CorbaTcBoolean|XMI.CorbaTcChar|
          XMI.CorbaTcWchar|XMI.CorbaTcOctet|XMI.CorbaTcAny|
          XMI.CorbaTcTypeCode|XMI.CorbaTcPrincipal|XMI.CorbaTcNull|
          XMI.CorbaTcVoid|XMI.CorbaTcLongLong|XMI.CorbaTcUlongLong|
          XMI.CorbaTcObjRef|XMI.CorbaTcLongDouble)>
<!ATTLIST XMI.CorbaTypeCode %XMI.element.att;>

<IELEMENT XMI.CorbaTcAlias (XMI.CorbaTypeCode)>
<!ATTLIST XMI.CorbaTcAlias
          xmi.tcName CDATA #REQUIRED
          xmi.tcId CDATA #IMPLIED>

<IELEMENT XMI.CorbaTcStruct (XMI.CorbaTcField)*>
<!ATTLIST XMI.CorbaTcStruct
          xmi.tcName CDATA #REQUIRED
          xmi.tcId CDATA #IMPLIED>

<IELEMENT XMI.CorbaTcField (XMI.CorbaTypeCode)>
<!ATTLIST XMI.CorbaTcField
          xmi.tcName CDATA #REQUIRED>

<IELEMENT XMI.CorbaTcSequence (XMI.CorbaTypeCode|XMI.CorbaRecursiveType)>
<!ATTLIST XMI.CorbaTcSequence
          xmi.tcLength CDATA #REQUIRED>

<IELEMENT XMI.CorbaRecursiveType EMPTY>
<!ATTLIST XMI.CorbaRecursiveType

```

```

    xmi.offset CDATA #REQUIRED>

<ELEMENT XMI.CorbaTcArray (XMI.CorbaTypeCode)>
<!ATTLIST XMI.CorbaTcArray
    xmi.tcLength CDATA #REQUIRED>

<ELEMENT XMI.CorbaTcObjRef EMPTY>
<!ATTLIST XMI.CorbaTcObjRef
    xmi.tcName CDATA #REQUIRED
    xmi.tcId CDATA #IMPLIED>

<ELEMENT XMI.CorbaTcEnum (XMI.CorbaTcEnumLabel)*>
<!ATTLIST XMI.CorbaTcEnum
    xmi.tcName CDATA #REQUIRED
    xmi.tcId CDATA #IMPLIED>

<ELEMENT XMI.CorbaTcEnumLabel EMPTY>
<!ATTLIST XMI.CorbaTcEnumLabel
    xmi.tcName CDATA #REQUIRED>

<ELEMENT XMI.CorbaTcUnionMbr (XMI.CorbaTypeCode, XMI.any)>
<!ATTLIST XMI.CorbaTcUnionMbr
    xmi.tcName CDATA #REQUIRED>

<ELEMENT XMI.CorbaTcUnion (XMI.CorbaTypeCode, XMI.CorbaTcUnionMbr*)>
<!ATTLIST XMI.CorbaTcUnion
    xmi.tcName CDATA #REQUIRED
    xmi.tcId CDATA #IMPLIED>

<ELEMENT XMI.CorbaTcExcept (XMI.CorbaTcField)*>
<!ATTLIST XMI.CorbaTcExcept
    xmi.tcName CDATA #REQUIRED
    xmi.tcId CDATA #IMPLIED>

<ELEMENT XMI.CorbaTcString EMPTY>
<!ATTLIST XMI.CorbaTcString
    xmi.tcLength CDATA #REQUIRED>

<ELEMENT XMI.CorbaTcWstring EMPTY>
<!ATTLIST XMI.CorbaTcWstring
    xmi.tcLength CDATA #REQUIRED>

<ELEMENT XMI.CorbaTcFixed EMPTY>
<!ATTLIST XMI.CorbaTcFixed
    xmi.tcDigits CDATA #REQUIRED
    xmi.tcScale CDATA #REQUIRED>

<ELEMENT XMI.CorbaTcShort EMPTY>
<ELEMENT XMI.CorbaTcLong EMPTY>
<ELEMENT XMI.CorbaTcUshort EMPTY>
<ELEMENT XMI.CorbaTcUlong EMPTY>

```

```

<!ELEMENT XMI.CorbaTcFloat EMPTY>
<!ELEMENT XMI.CorbaTcDouble EMPTY>
<!ELEMENT XMI.CorbaTcBoolean EMPTY>
<!ELEMENT XMI.CorbaTcChar EMPTY>
<!ELEMENT XMI.CorbaTcWchar EMPTY>
<!ELEMENT XMI.CorbaTcOctet EMPTY>
<!ELEMENT XMI.CorbaTcAny EMPTY>
<!ELEMENT XMI.CorbaTcTypeCode EMPTY>
<!ELEMENT XMI.CorbaTcPrincipal EMPTY>
<!ELEMENT XMI.CorbaTcNull EMPTY>
<!ELEMENT XMI.CorbaTcVoid EMPTY>
<!ELEMENT XMI.CorbaTcLongLong EMPTY>
<!ELEMENT XMI.CorbaTcUlongLong EMPTY>
<!ELEMENT XMI.CorbaTcLongDouble EMPTY>

```

```

<!ATTLIST XMI xmlns:BaseIDL CDATA #IMPLIED>

```

```

<!-- _____ -->
<!-- -->
<!-- METAMODEL PACKAGE: BaseIDL -->
<!-- -->
<!-- _____ -->

```

```

<!ENTITY % BaseIDL:PrimitiveKind '(PK_NULL|PK_VOID|PK_SHORT|PK_LONG|
  PK_USHORT|PK_ULONG|PK_FLOAT|PK_DOUBLE|PK_BOOLEAN|PK_CHAR|PK_OCTET|PK_ANY|
  PK_LONGDOUBLE|PK_WSTRING|PK_TYPECODE|PK_WCHAR|PK_PRINCIPAL|PK_STRING|
  PK_ULONGLONG|PK_OBJREF|PK_LONGLONG)'>

```

```

<!ENTITY % BaseIDL:ParameterMode '(PARAM_IN|PARAM_OUT|PARAM_INOUT)'>

```

```

<!ENTITY % BaseIDL:DefinitionKind '(DK_NONE|DK_ALL|DK_ATTRIBUTE|
  DK_CONSTANT|DK_EXCEPTION|DK_INTERFACE|DK_MODULE|DK_OPERATION|DK_TYPEDEF|
  DK_ALIAS|DK_STRUCT|DK_UNION|DK_FIXED|DK_ENUM|DK_PRIMITIVE|DK_STRING|
  DK_SEQUENCE|DK_WSTRING|DK_ARRAY|DK_REPOSITORY)'>

```

```

<!-- _____ -->
<!-- -->
<!-- METAMODEL CLASS: BaseIDL.Typed -->
<!-- -->
<!-- _____ -->

```

```

<!ELEMENT BaseIDL.Typed.idlType (BaseIDL:IDLType)*>

```

```

<!ENTITY % BaseIDL:TypedFeatures 'XMI.extension |
  BaseIDL:Typed.idlType'>

```

```

<!ENTITY % BaseIDL:TypedAtts '%XMI.element.att; %XMI.link.att;
  idlType IDREFS #IMPLIED'>

```

```

<!ELEMENT BaseIDL.Typed (%BaseIDL:TypedFeatures;)*>

```

```

<!ATTLIST BaseIDL:Typed %BaseIDL:TypedAtts;>

<!-- _____ -->
<!-- -->
<!-- METAMODEL CLASS: BaseIDL:ParameterDef -->
<!-- -->
<!-- _____ -->

<IELEMENT BaseIDL:ParameterDef.identifier (#PCDATA|XML.reference)*>

<IELEMENT BaseIDL:ParameterDef.direction EMPTY>

<!ATTLIST BaseIDL:ParameterDef.direction xmi.value %BaseIDL:ParameterMode; #REQUIRED>

<ENTITY % BaseIDL:ParameterDefFeatures '%BaseIDL:TypedFeatures; |
BaseIDL:ParameterDef.identifier |
BaseIDL:ParameterDef.direction'>

<ENTITY % BaseIDL:ParameterDefAtts '%BaseIDL:TypedAtts;
identifier CDATA #IMPLIED
direction %BaseIDL:ParameterMode; #IMPLIED'>

<IELEMENT BaseIDL:ParameterDef (%BaseIDL:ParameterDefFeatures;)*>

<!ATTLIST BaseIDL:ParameterDef %BaseIDL:ParameterDefAtts;>

<!-- _____ -->
<!-- -->
<!-- METAMODEL CLASS: BaseIDL:Contained -->
<!-- -->
<!-- _____ -->

<IELEMENT BaseIDL:Contained.definedIn (BaseIDL:Container)*>

<IELEMENT BaseIDL:Contained.identifier (#PCDATA|XML.reference)*>

<IELEMENT BaseIDL:Contained.repositoryId (#PCDATA|XML.reference)*>

<IELEMENT BaseIDL:Contained.version (#PCDATA|XML.reference)*>

<ENTITY % BaseIDL:ContainedFeatures 'XML.extension |
BaseIDL:Contained.definedIn |
BaseIDL:Contained.identifier |
BaseIDL:Contained.repositoryId |
BaseIDL:Contained.version'>

<ENTITY % BaseIDL:ContainedAtts '%XML.element.att; %XML.link.att;
definedIn IDREFS #IMPLIED
identifier CDATA #IMPLIED
repositoryId CDATA #IMPLIED'>

```

```

    version CDATA #IMPLIED'>

<!ELEMENT BaseIDL:Contained (%BaseIDL:ContainedFeatures;)*>

<!ATTLIST BaseIDL:Contained %BaseIDL:ContainedAtts;>

<!-- _____ -->
<!-- -->
<!-- METAMODEL CLASS: BaseIDL.ConstantDef -->
<!-- -->
<!-- _____ -->

<!ELEMENT BaseIDL:ConstantDef.constValue (XML.any)>

<!ENTITY % BaseIDL:ConstantDefFeatures '%BaseIDL:TypedFeatures; |
BaseIDL:Contained.definedIn |
BaseIDL:Contained.identifier |
BaseIDL:Contained.repositoryId |
BaseIDL:Contained.version |
BaseIDL:ConstantDef.constValue'>

<!ENTITY % BaseIDL:ConstantDefAtts '%BaseIDL:TypedAtts;
definedIn IDREFS #IMPLIED
identifier CDATA #IMPLIED
repositoryId CDATA #IMPLIED
version CDATA #IMPLIED'>

<!ELEMENT BaseIDL:ConstantDef (%BaseIDL:ConstantDefFeatures;)*>

<!ATTLIST BaseIDL:ConstantDef %BaseIDL:ConstantDefAtts;>

<!-- _____ -->
<!-- -->
<!-- METAMODEL CLASS: BaseIDL.Container -->
<!-- -->
<!-- _____ -->

<!ELEMENT BaseIDL:Container.contents (BaseIDL:Contained|
BaseIDL:ConstantDef|BaseIDL:TypedDef|BaseIDL:StructDef|
BaseIDL:UnionDef|BaseIDL:EnumDef|BaseIDL:AliasDef|
BaseIDL:ValueBoxDef|BaseIDL:Container|BaseIDL:ModuleDef|
BaseIDL:InterfaceDef|ComponentIDL:ComponentDef|ComponentIDL:HomeDef|
BaseIDL:ValueDef|ComponentIDL:EventDef|CIF:ComponentImplDef|
CIF:HomeImplDef|BaseIDL:ValueMemberDef|
BaseIDL:OperationDef|ComponentIDL:FactoryDef|ComponentIDL:FinderDef|
BaseIDL:ExceptionDef|BaseIDL:AttributeDef|
CIF:ArtifactDef|CIF:SegmentDef|ComponentIDL:ProvidesDef|
ComponentIDL:UsesDef|ComponentIDL:EventPortDef|ComponentIDL:EmitsDef|
ComponentIDL:ConsumesDef|ComponentIDL:PublishesDef)*>

<!ENTITY % BaseIDL:ContainerFeatures '%BaseIDL:ContainedFeatures; |

```

```

BaseIDL:Container.contents'>

<!ENTITY % BaseIDL:ContainerAtts '%BaseIDL:ContainedAtts;'>

<!ELEMENT BaseIDL:Container (%BaseIDL:ContainerFeatures;)*>

<!ATTLIST BaseIDL:Container %BaseIDL:ContainerAtts;>

<!-- _____ -->
<!-- -->
<!-- METAMODEL CLASS: BaseIDL.ModuleDef -->
<!-- -->
<!-- _____ -->

<!ELEMENT BaseIDL:ModuleDef.prefix (#PCDATA|XMI.reference)*>

<!ENTITY % BaseIDL:ModuleDefFeatures '%BaseIDL:ContainerFeatures; |
BaseIDL:ModuleDef.prefix'>

<!ENTITY % BaseIDL:ModuleDefAtts '%BaseIDL:ContainerAtts;
prefix CDATA #IMPLIED'>

<!ELEMENT BaseIDL:ModuleDef (%BaseIDL:ModuleDefFeatures;)*>

<!ATTLIST BaseIDL:ModuleDef %BaseIDL:ModuleDefAtts;>

<!-- _____ -->
<!-- -->
<!-- METAMODEL CLASS: BaseIDL.IDLType -->
<!-- -->
<!-- _____ -->

<!ENTITY % BaseIDL:IDLTypeFeatures 'XMI.extension'>

<!ENTITY % BaseIDL:IDLTypeAtts '%XMI.element.att; %XMI.link.att;'>

<!ELEMENT BaseIDL:IDLType (%BaseIDL:IDLTypeFeatures;)*>

<!ATTLIST BaseIDL:IDLType %BaseIDL:IDLTypeAtts;>

<!-- _____ -->
<!-- -->
<!-- METAMODEL CLASS: BaseIDL.TypedefDef -->
<!-- -->
<!-- _____ -->

<!ENTITY % BaseIDL:TypedefDefFeatures '%BaseIDL:IDLTypeFeatures; |
BaseIDL:Contained.definedIn |
BaseIDL:Contained.identifier |
BaseIDL:Contained.repositoryId |
BaseIDL:Contained.version'>

```



```

<!ENTITY % BaseIDL:TypedefDefAtts '%BaseIDL:IDLTypeAtts;
  definedIn IDREFS #IMPLIED
  identifier CDATA #IMPLIED
  repositoryId CDATA #IMPLIED
  version CDATA #IMPLIED'>

<!ELEMENT BaseIDL:TypedefDef (%BaseIDL:TypedefDefFeatures;)*>

<!ATTLIST BaseIDL:TypedefDef %BaseIDL:TypedefDefAtts;>

<!-- _____ -->
<!-- -->
<!-- METAMODEL CLASS: BaseIDL:InterfaceDef -->
<!-- -->
<!-- _____ -->

<!ELEMENT BaseIDL:InterfaceDef.base (BaseIDL:InterfaceDef)*>

<!ELEMENT BaseIDL:InterfaceDef.isAbstract EMPTY>

<!ATTLIST BaseIDL:InterfaceDef.isAbstract xmi.value (true|false) #REQUIRED>

<!ELEMENT BaseIDL:InterfaceDef.isLocal EMPTY>

<!ATTLIST BaseIDL:InterfaceDef.isLocal xmi.value (true|false) #REQUIRED>

<!ENTITY % BaseIDL:InterfaceDefFeatures '%BaseIDL:IDLTypeFeatures; |
  BaseIDL:Contained.definedIn |
  BaseIDL:Contained.identifier |
  BaseIDL:Contained.repositoryId |
  BaseIDL:Contained.version |
  BaseIDL:Container.contents |
  BaseIDL:InterfaceDef.base |
  BaseIDL:InterfaceDef.isAbstract |
  BaseIDL:InterfaceDef.isLocal'>

<!ENTITY % BaseIDL:InterfaceDefAtts '%BaseIDL:IDLTypeAtts;
  definedIn IDREFS #IMPLIED
  identifier CDATA #IMPLIED
  repositoryId CDATA #IMPLIED
  version CDATA #IMPLIED
  base IDREFS #IMPLIED
  isAbstract (true|false) #IMPLIED
  isLocal (true|false) #IMPLIED'>

<!ELEMENT BaseIDL:InterfaceDef (%BaseIDL:InterfaceDefFeatures;)*>

<!ATTLIST BaseIDL:InterfaceDef %BaseIDL:InterfaceDefAtts;>

```

```

<!-- _____ -->
<!-- -->
<!-- METAMODEL CLASS: BaseIDL.Field -->
<!-- -->
<!-- _____ -->

<ELEMENT BaseIDL:Field.identifier (#PCDATA|XML.reference)*>

<ENTITY % BaseIDL:FieldFeatures '%BaseIDL:TypedFeatures; |
  BaseIDL:Field.identifier'>

<ENTITY % BaseIDL:FieldAtts '%BaseIDL:TypedAtts;
  identifier CDATA #IMPLIED'>

<ELEMENT BaseIDL:Field (%BaseIDL:FieldFeatures;)*>

<ATTLIST BaseIDL:Field %BaseIDL:FieldAtts;>

<!-- _____ -->
<!-- -->
<!-- METAMODEL CLASS: BaseIDL.StructDef -->
<!-- -->
<!-- _____ -->

<ELEMENT BaseIDL:StructDef.members (BaseIDL:Field)*>

<ENTITY % BaseIDL:StructDefFeatures '%BaseIDL:TypedefDefFeatures; |
  BaseIDL:StructDef.members'>

<ENTITY % BaseIDL:StructDefAtts '%BaseIDL:TypedefDefAtts;'>

<ELEMENT BaseIDL:StructDef (%BaseIDL:StructDefFeatures;)*>

<ATTLIST BaseIDL:StructDef %BaseIDL:StructDefAtts;>

<!-- _____ -->
<!-- -->
<!-- METAMODEL CLASS: BaseIDL.UnionDef -->
<!-- -->
<!-- _____ -->

<ELEMENT BaseIDL:UnionDef.discriminatorType (BaseIDL:IDLType)*>

<ELEMENT BaseIDL:UnionDef.unionMembers (BaseIDL:UnionField)*>

<ENTITY % BaseIDL:UnionDefFeatures '%BaseIDL:TypedefDefFeatures; |
  BaseIDL:UnionDef.discriminatorType |
  BaseIDL:UnionDef.unionMembers'>

<ENTITY % BaseIDL:UnionDefAtts '%BaseIDL:TypedefDefAtts;
  discriminatorType IDREFS #IMPLIED'>

```

```

<!ELEMENT BaseIDL:UnionDef (%BaseIDL:UnionDefFeatures;)*>

<!ATTLIST BaseIDL:UnionDef %BaseIDL:UnionDefAtts;>

<!-- _____ -->
<!-- -->
<!-- METAMODEL CLASS: BaseIDL.EnumDef -->
<!-- -->
<!-- _____ -->

<!ELEMENT BaseIDL:EnumDef.members (#PCDATA|XML.reference)*>

<!ENTITY % BaseIDL:EnumDefFeatures '%BaseIDL:TypedefDefFeatures; |
BaseIDL:EnumDef.members'>

<!ENTITY % BaseIDL:EnumDefAtts '%BaseIDL:TypedefDefAtts;
members CDATA #IMPLIED'>

<!ELEMENT BaseIDL:EnumDef (%BaseIDL:EnumDefFeatures;)*>

<!ATTLIST BaseIDL:EnumDef %BaseIDL:EnumDefAtts;>

<!-- _____ -->
<!-- -->
<!-- METAMODEL CLASS: BaseIDL.AliasDef -->
<!-- -->
<!-- _____ -->

<!ENTITY % BaseIDL:AliasDefFeatures '%BaseIDL:TypedefDefFeatures; |
BaseIDL:Typed.idlType'>

<!ENTITY % BaseIDL:AliasDefAtts '%BaseIDL:TypedefDefAtts;
idlType IDREFS #IMPLIED'>

<!ELEMENT BaseIDL:AliasDef (%BaseIDL:AliasDefFeatures;)*>

<!ATTLIST BaseIDL:AliasDef %BaseIDL:AliasDefAtts;>

<!-- _____ -->
<!-- -->
<!-- METAMODEL CLASS: BaseIDL.StringDef -->
<!-- -->
<!-- _____ -->

<!ELEMENT BaseIDL:StringDef.bound (#PCDATA|XML.reference)*>

<!ENTITY % BaseIDL:StringDefFeatures '%BaseIDL:IDLTypeFeatures; |
BaseIDL:StringDef.bound'>

<!ENTITY % BaseIDL:StringDefAtts '%BaseIDL:IDLTypeAtts;

```

```

bound CDATA #IMPLIED'>

<!ELEMENT BaseIDL:StringDef (%BaseIDL:StringDefFeatures;)*>

<!ATTLIST BaseIDL:StringDef %BaseIDL:StringDefAtts;>

<!-- _____ -->
<!-- -->
<!-- METAMODEL CLASS: BaseIDL.WStringDef -->
<!-- -->
<!-- _____ -->

<!ELEMENT BaseIDL:WStringDef.bound (#PCDATA|XML.reference)*>

<ENTITY % BaseIDL:WStringDefFeatures '%BaseIDL:IDLTypeFeatures; |
BaseIDL:WStringDef.bound'>

<ENTITY % BaseIDL:WStringDefAtts '%BaseIDL:IDLTypeAtts;
bound CDATA #IMPLIED'>

<!ELEMENT BaseIDL:WStringDef (%BaseIDL:WStringDefFeatures;)*>

<!ATTLIST BaseIDL:WStringDef %BaseIDL:WStringDefAtts;>

<!-- _____ -->
<!-- -->
<!-- METAMODEL CLASS: BaseIDL.FixedDef -->
<!-- -->
<!-- _____ -->

<!ELEMENT BaseIDL:FixedDef.digits (#PCDATA|XML.reference)*>

<!ELEMENT BaseIDL:FixedDef.scale (#PCDATA|XML.reference)*>

<ENTITY % BaseIDL:FixedDefFeatures '%BaseIDL:IDLTypeFeatures; |
BaseIDL:FixedDef.digits |
BaseIDL:FixedDef.scale'>

<ENTITY % BaseIDL:FixedDefAtts '%BaseIDL:IDLTypeAtts;
digits CDATA #IMPLIED
scale CDATA #IMPLIED'>

<!ELEMENT BaseIDL:FixedDef (%BaseIDL:FixedDefFeatures;)*>

<!ATTLIST BaseIDL:FixedDef %BaseIDL:FixedDefAtts;>

<!-- _____ -->
<!-- -->
<!-- METAMODEL CLASS: BaseIDL.SequenceDef -->
<!-- -->
<!-- _____ -->

```

```

<!ELEMENT BaseIDL:SequenceDef.bound (#PCDATA|XML.reference)*>

<!ENTITY % BaseIDL:SequenceDefFeatures '%BaseIDL:TypedFeatures; |
BaseIDL:SequenceDef.bound'>

<!ENTITY % BaseIDL:SequenceDefAtts '%BaseIDL:TypedAtts;
bound CDATA #IMPLIED'>

<!ELEMENT BaseIDL:SequenceDef (%BaseIDL:SequenceDefFeatures;)*>

<!ATTLIST BaseIDL:SequenceDef %BaseIDL:SequenceDefAtts;>

<!-- _____ -->
<!-- -->
<!-- METAMODEL CLASS: BaseIDL:ArrayDef -->
<!-- -->
<!-- _____ -->

<!ELEMENT BaseIDL:ArrayDef.bound (#PCDATA|XML.reference)*>

<!ENTITY % BaseIDL:ArrayDefFeatures '%BaseIDL:TypedFeatures; |
BaseIDL:ArrayDef.bound'>

<!ENTITY % BaseIDL:ArrayDefAtts '%BaseIDL:TypedAtts;
bound CDATA #IMPLIED'>

<!ELEMENT BaseIDL:ArrayDef (%BaseIDL:ArrayDefFeatures;)*>

<!ATTLIST BaseIDL:ArrayDef %BaseIDL:ArrayDefAtts;>

<!-- _____ -->
<!-- -->
<!-- METAMODEL CLASS: BaseIDL:PrimitiveDef -->
<!-- -->
<!-- _____ -->

<!ELEMENT BaseIDL:PrimitiveDef.kind EMPTY>

<!ATTLIST BaseIDL:PrimitiveDef.kind xmi.value %BaseIDL:PrimitiveKind; #REQUIRED>

<!ENTITY % BaseIDL:PrimitiveDefFeatures '%BaseIDL:IDLTypeFeatures; |
BaseIDL:PrimitiveDef.kind'>

<!ENTITY % BaseIDL:PrimitiveDefAtts '%BaseIDL:IDLTypeAtts;
kind %BaseIDL:PrimitiveKind; #IMPLIED'>

<!ELEMENT BaseIDL:PrimitiveDef (%BaseIDL:PrimitiveDefFeatures;)*>

<!ATTLIST BaseIDL:PrimitiveDef %BaseIDL:PrimitiveDefAtts;>

```

```

<!-- _____ -->
<!-- -->
<!-- METAMODEL CLASS: BaseIDL.UnionField -->
<!-- -->
<!-- _____ -->

<ELEMENT BaseIDL:UnionField.identifier (#PCDATA|XML.reference)*>

<ELEMENT BaseIDL:UnionField.label (XML.any)>

<ENTITY % BaseIDL:UnionFieldFeatures '%BaseIDL:TypedFeatures; |
BaseIDL:UnionField.identifier |
BaseIDL:UnionField.label'>

<ENTITY % BaseIDL:UnionFieldAtts '%BaseIDL:TypedAtts;
identifier CDATA #IMPLIED'>

<ELEMENT BaseIDL:UnionField (%BaseIDL:UnionFieldFeatures;)*>

<ATTLIST BaseIDL:UnionField %BaseIDL:UnionFieldAtts;>

<!-- _____ -->
<!-- -->
<!-- METAMODEL CLASS: BaseIDL.ValueMemberDef -->
<!-- -->
<!-- _____ -->

<ELEMENT BaseIDL:ValueMemberDef.isPublicMember EMPTY>

<ATTLIST BaseIDL:ValueMemberDef.isPublicMember xmi.value (true|false) #REQUIRED>

<ENTITY % BaseIDL:ValueMemberDefFeatures '%BaseIDL:TypedFeatures; |
BaseIDL:Contained.definedIn |
BaseIDL:Contained.identifier |
BaseIDL:Contained.repositoryId |
BaseIDL:Contained.version |
BaseIDL:ValueMemberDef.isPublicMember'>

<ENTITY % BaseIDL:ValueMemberDefAtts '%BaseIDL:TypedAtts;
definedIn IDREFS #IMPLIED
identifier CDATA #IMPLIED
repositoryId CDATA #IMPLIED
version CDATA #IMPLIED
isPublicMember (true|false) #IMPLIED'>

<ELEMENT BaseIDL:ValueMemberDef (%BaseIDL:ValueMemberDefFeatures;)*>

<ATTLIST BaseIDL:ValueMemberDef %BaseIDL:ValueMemberDefAtts;>

<!-- _____ -->
<!-- -->

```

```

<!-- METAMODEL CLASS: BaseIDL.ValueDef -->
<!-- -->
<!-- _____ -->

<!ELEMENT BaseIDL.ValueDef.interfaceDef (BaseIDL:InterfaceDef)*>

<!ELEMENT BaseIDL.ValueDef.base (BaseIDL:ValueDef)*>

<!ELEMENT BaseIDL.ValueDef.abstractBase (BaseIDL:ValueDef)*>

<!ELEMENT BaseIDL.ValueDef.isAbstract EMPTY>

<!ATTLIST BaseIDL.ValueDef.isAbstract xmi.value (true|false) #REQUIRED>

<!ELEMENT BaseIDL.ValueDef.isCustom EMPTY>

<!ATTLIST BaseIDL.ValueDef.isCustom xmi.value (true|false) #REQUIRED>

<!ELEMENT BaseIDL.ValueDef.isTruncatable EMPTY>

<!ATTLIST BaseIDL.ValueDef.isTruncatable xmi.value (true|false) #REQUIRED>

<!ENTITY % BaseIDL.ValueDefFeatures '%BaseIDL:ContainerFeatures; |
BaseIDL.ValueDef.interfaceDef |
BaseIDL.ValueDef.base |
BaseIDL.ValueDef.abstractBase |
BaseIDL.ValueDef.isAbstract |
BaseIDL.ValueDef.isCustom |
BaseIDL.ValueDef.isTruncatable'>

<!ENTITY % BaseIDL.ValueDefAtts '%BaseIDL:ContainerAtts;
interfaceDef IDREFS #IMPLIED
base IDREFS #IMPLIED
abstractBase IDREFS #IMPLIED
isAbstract (true|false) #IMPLIED
isCustom (true|false) #IMPLIED
isTruncatable (true|false) #IMPLIED'>

<!ELEMENT BaseIDL.ValueDef (%BaseIDL.ValueDefFeatures;)*>

<!ATTLIST BaseIDL.ValueDef %BaseIDL.ValueDefAtts;>

<!-- _____ -->
<!-- -->
<!-- METAMODEL CLASS: BaseIDL.ValueBoxDef -->
<!-- -->
<!-- _____ -->

<!ENTITY % BaseIDL.ValueBoxDefFeatures '%BaseIDL:TypedefDefFeatures;'>

<!ENTITY % BaseIDL.ValueBoxDefAtts '%BaseIDL:TypedefDefAtts;'>

```

```

<!ELEMENT BaseIDL:ValueBoxDef (%BaseIDL:ValueBoxDefFeatures;)*>

<!ATTLIST BaseIDL:ValueBoxDef %BaseIDL:ValueBoxDefAtts;>

<!-- _____ -->
<!-- -->
<!-- METAMODEL CLASS: BaseIDL:OperationDef -->
<!-- -->
<!-- _____ -->

<!ELEMENT BaseIDL:OperationDef.exceptionDef (BaseIDL:ExceptionDef)*>

<!ELEMENT BaseIDL:OperationDef.isOneway EMPTY>

<!ATTLIST BaseIDL:OperationDef.isOneway xmi.value (true|false) #REQUIRED>

<!ELEMENT BaseIDL:OperationDef.parameters (BaseIDL:ParameterDef)*>

<!ELEMENT BaseIDL:OperationDef.contexts (#PCDATA|XMI.reference)*>

<ENTITY % BaseIDL:OperationDefFeatures '%BaseIDL:TypedFeatures; |
BaseIDL:Contained.definedIn |
BaseIDL:Contained.identifier |
BaseIDL:Contained.repositoryId |
BaseIDL:Contained.version |
BaseIDL:OperationDef.exceptionDef |
BaseIDL:OperationDef.isOneway |
BaseIDL:OperationDef.parameters |
BaseIDL:OperationDef.contexts'>

<ENTITY % BaseIDL:OperationDefAtts '%BaseIDL:TypedAtts;
definedIn IDREFS #IMPLIED
identifier CDATA #IMPLIED
repositoryId CDATA #IMPLIED
version CDATA #IMPLIED
exceptionDef IDREFS #IMPLIED
isOneway (true|false) #IMPLIED
contexts CDATA #IMPLIED'>

<!ELEMENT BaseIDL:OperationDef (%BaseIDL:OperationDefFeatures;)*>

<!ATTLIST BaseIDL:OperationDef %BaseIDL:OperationDefAtts;>

<!-- _____ -->
<!-- -->
<!-- METAMODEL CLASS: BaseIDL:ExceptionDef -->
<!-- -->
<!-- _____ -->

<!ELEMENT BaseIDL:ExceptionDef.typeCode (XMI.CorbaTypeCode)>

```



```

<!ELEMENT BaseIDL:ExceptionDef.members (BaseIDL:Field)*>

<!ENTITY % BaseIDL:ExceptionDefFeatures '%BaseIDL:ContainedFeatures; |
BaseIDL:ExceptionDef.typeCode |
BaseIDL:ExceptionDef.members'>

<!ENTITY % BaseIDL:ExceptionDefAtts '%BaseIDL:ContainedAtts;'>

<!ELEMENT BaseIDL:ExceptionDef (%BaseIDL:ExceptionDefFeatures;)*>

<!ATTLIST BaseIDL:ExceptionDef %BaseIDL:ExceptionDefAtts;>

<!-- _____ -->
<!-- -->
<!-- METAMODEL CLASS: BaseIDL:AttributeDef -->
<!-- -->
<!-- _____ -->

<!ELEMENT BaseIDL:AttributeDef.setException (BaseIDL:ExceptionDef)*>

<!ELEMENT BaseIDL:AttributeDef.getException (BaseIDL:ExceptionDef)*>

<!ELEMENT BaseIDL:AttributeDef.isReadOnly EMPTY>

<!ATTLIST BaseIDL:AttributeDef.isReadOnly xmi.value (true|false) #REQUIRED>

<!ENTITY % BaseIDL:AttributeDefFeatures '%BaseIDL:TypedFeatures; |
BaseIDL:Contained.definedIn |
BaseIDL:Contained.identifier |
BaseIDL:Contained.repositoryId |
BaseIDL:Contained.version |
BaseIDL:AttributeDef.setException |
BaseIDL:AttributeDef.getException |
BaseIDL:AttributeDef.isReadOnly'>

<!ENTITY % BaseIDL:AttributeDefAtts '%BaseIDL:TypedAtts;
definedIn IDREFS #IMPLIED
identifier CDATA #IMPLIED
repositoryId CDATA #IMPLIED
version CDATA #IMPLIED
setException IDREFS #IMPLIED
getException IDREFS #IMPLIED
isReadOnly (true|false) #IMPLIED'>

<!ELEMENT BaseIDL:AttributeDef (%BaseIDL:AttributeDefFeatures;)*>

<!ATTLIST BaseIDL:AttributeDef %BaseIDL:AttributeDefAtts;>

<!-- _____ -->
<!-- -->

```

```

<!-- METAMODEL PACKAGE: CIF -->
<!-- -->
<!-- _____ -->

<ENTITY % CIF:ComponentCategory '(PROCESS|SESSION|ENTITY|SERVICE)'>

<!-- _____ -->
<!-- -->
<!-- METAMODEL CLASS: CIF:ArtifactDef -->
<!-- -->
<!-- _____ -->

<ENTITY % CIF:ArtifactDefFeatures '%BaseIDL:ContainedFeatures;'>

<ENTITY % CIF:ArtifactDefAtts '%BaseIDL:ContainedAtts;'>

<ELEMENT CIF:ArtifactDef (%CIF:ArtifactDefFeatures;)*>

<!ATTLIST CIF:ArtifactDef %CIF:ArtifactDefAtts;>

<!-- _____ -->
<!-- -->
<!-- METAMODEL CLASS: CIF:SegmentDef -->
<!-- -->
<!-- _____ -->

<ELEMENT CIF:SegmentDef.artifact (CIF:ArtifactDef)*>

<ELEMENT CIF:SegmentDef.features (ComponentIDL:ComponentFeature)*>

<ELEMENT CIF:SegmentDef.policies (CIF:Policy)*>

<ELEMENT CIF:SegmentDef.isSerialized EMPTY>

<!ATTLIST CIF:SegmentDef.isSerialized xmi.value (true|false) #REQUIRED>

<ENTITY % CIF:SegmentDefFeatures '%BaseIDL:ContainedFeatures; |
  CIF:SegmentDef.artifact |
  CIF:SegmentDef.features |
  CIF:SegmentDef.policies |
  CIF:SegmentDef.isSerialized'>

<ENTITY % CIF:SegmentDefAtts '%BaseIDL:ContainedAtts;
  artifact IDREFS #IMPLIED
  features IDREFS #IMPLIED
  policies IDREFS #IMPLIED
  isSerialized (true|false) #IMPLIED'>

<ELEMENT CIF:SegmentDef (%CIF:SegmentDefFeatures;)*>

<!ATTLIST CIF:SegmentDef %CIF:SegmentDefAtts;>

```

```

<!-- _____ -->
<!-- -->
<!-- METAMODEL CLASS: CIF:ComponentImplDef -->
<!-- -->
<!-- _____ -->

<!ELEMENT CIF:ComponentImplDef.component (ComponentIDL:ComponentDef)*>

<!ELEMENT CIF:ComponentImplDef.category EMPTY>

<!ATTLIST CIF:ComponentImplDef.category xmi.value %CIF:ComponentCategory; #REQUIRED>

<!ENTITY % CIF:ComponentImplDefFeatures '%BaseIDL:ContainerFeatures; |
  CIF:ComponentImplDef.component |
  CIF:ComponentImplDef.category'>

<!ENTITY % CIF:ComponentImplDefAtts '%BaseIDL:ContainerAtts;
  component IDREFS #IMPLIED
  category %CIF:ComponentCategory; #IMPLIED'>

<!ELEMENT CIF:ComponentImplDef (%CIF:ComponentImplDefFeatures;)*>

<!ATTLIST CIF:ComponentImplDef %CIF:ComponentImplDefAtts;>

<!-- _____ -->
<!-- -->
<!-- METAMODEL CLASS: CIF:Policy -->
<!-- -->
<!-- -->
<!-- _____ -->

<!ENTITY % CIF:PolicyFeatures 'XML.extension'>

<!ENTITY % CIF:PolicyAtts '%XML.element.att; %XML.link.att;'>

<!ELEMENT CIF:Policy (%CIF:PolicyFeatures;)*>

<!ATTLIST CIF:Policy %CIF:PolicyAtts;>

<!-- _____ -->
<!-- -->
<!-- METAMODEL CLASS: CIF:HomeImplDef -->
<!-- -->
<!-- _____ -->

<!ELEMENT CIF:HomeImplDef.home (ComponentIDL:HomeDef)*>

<!ELEMENT CIF:HomeImplDef.component_impl (CIF:ComponentImplDef)*>

<!ENTITY % CIF:HomeImplDefFeatures '%BaseIDL:ContainerFeatures; |
  CIF:HomeImplDef.home |

```

```

    CIF:HomeImplDef.component_impl'>

<!ENTITY % CIF:HomeImplDefAtts '%BaseIDL:ContainerAtts;
    home IDREFS #IMPLIED
    component_impl IDREFS #IMPLIED'>

<!ELEMENT CIF:HomeImplDef (%CIF:HomeImplDefFeatures;)*>

<!ATTLIST CIF:HomeImplDef %CIF:HomeImplDefAtts;>

<!-- _____ -->
<!-- -->
<!-- METAMODEL PACKAGE: ComponentIDL -->
<!-- -->
<!-- _____ -->

<!-- _____ -->
<!-- -->
<!-- METAMODEL CLASS: ComponentIDL:ComponentFeature -->
<!-- -->
<!-- _____ -->

<!ENTITY % ComponentIDL:ComponentFeatureFeatures 'XMI.extension'>

<!ENTITY % ComponentIDL:ComponentFeatureAtts '%XMI.element.att; %XMI.link.att;'>

<!ELEMENT ComponentIDL:ComponentFeature (%ComponentIDL:ComponentFeatureFeatures;)*>

<!ATTLIST ComponentIDL:ComponentFeature %ComponentIDL:ComponentFeatureAtts;>

<!-- _____ -->
<!-- -->
<!-- METAMODEL CLASS: ComponentIDL:ComponentDef -->
<!-- -->
<!-- _____ -->

<!ELEMENT ComponentIDL:ComponentDef.supports (BaseIDL:InterfaceDef)*>

<!ENTITY % ComponentIDL:ComponentDefFeatures '%BaseIDL:InterfaceDefFeatures; |
    ComponentIDL:ComponentDef.supports'>

<!ENTITY % ComponentIDL:ComponentDefAtts '%BaseIDL:InterfaceDefAtts;
    supports IDREFS #IMPLIED'>

<!ELEMENT ComponentIDL:ComponentDef (%ComponentIDL:ComponentDefFeatures;)*>

<!ATTLIST ComponentIDL:ComponentDef %ComponentIDL:ComponentDefAtts;>

<!-- _____ -->
<!-- -->
<!-- METAMODEL CLASS: ComponentIDL:ProvidesDef -->

```

```

<!--                                     -->
<!-- _____ -->

<!ELEMENT ComponentIDL:ProvidesDef.provides (BaseIDL:InterfaceDef)*>

<!ENTITY % ComponentIDL:ProvidesDefFeatures '%BaseIDL:ContainedFeatures; |
  ComponentIDL:ProvidesDef.provides'>

<!ENTITY % ComponentIDL:ProvidesDefAtts '%BaseIDL:ContainedAtts;
  provides IDREFS #IMPLIED'>

<!ELEMENT ComponentIDL:ProvidesDef (%ComponentIDL:ProvidesDefFeatures;)*>

<!ATTLIST ComponentIDL:ProvidesDef %ComponentIDL:ProvidesDefAtts;>

<!--                                     -->
<!--                                     -->
<!-- METAMODEL CLASS: ComponentIDL:HomeDef -->
<!--                                     -->
<!-- _____ -->

<!ELEMENT ComponentIDL:HomeDef.component (ComponentIDL:ComponentDef)*>

<!ELEMENT ComponentIDL:HomeDef.primary_key (BaseIDL:ValueDef)*>

<!ELEMENT ComponentIDL:HomeDef.supports (BaseIDL:InterfaceDef)*>

<!ENTITY % ComponentIDL:HomeDefFeatures '%BaseIDL:InterfaceDefFeatures; |
  ComponentIDL:HomeDef.component |
  ComponentIDL:HomeDef.primary_key |
  ComponentIDL:HomeDef.supports'>

<!ENTITY % ComponentIDL:HomeDefAtts '%BaseIDL:InterfaceDefAtts;
  component IDREFS #IMPLIED
  primary_key IDREFS #IMPLIED
  supports IDREFS #IMPLIED'>

<!ELEMENT ComponentIDL:HomeDef (%ComponentIDL:HomeDefFeatures;)*>

<!ATTLIST ComponentIDL:HomeDef %ComponentIDL:HomeDefAtts;>

<!--                                     -->
<!--                                     -->
<!-- METAMODEL CLASS: ComponentIDL:FactoryDef -->
<!--                                     -->
<!-- _____ -->

<!ENTITY % ComponentIDL:FactoryDefFeatures '%BaseIDL:OperationDefFeatures;'>

<!ENTITY % ComponentIDL:FactoryDefAtts '%BaseIDL:OperationDefAtts;'>

```

```

<ELEMENT ComponentIDL:FactoryDef (%ComponentIDL:FactoryDefFeatures;)*>
<ATTLIST ComponentIDL:FactoryDef %ComponentIDL:FactoryDefAtts;>

<!-- _____ -->
<!-- -->
<!-- METAMODEL CLASS: ComponentIDL:FinderDef -->
<!-- -->
<!-- _____ -->

<ENTITY % ComponentIDL:FinderDefFeatures '%BaseIDL:OperationDefFeatures;'>
<ENTITY % ComponentIDL:FinderDefAtts '%BaseIDL:OperationDefAtts;'>
<ELEMENT ComponentIDL:FinderDef (%ComponentIDL:FinderDefFeatures;)*>
<ATTLIST ComponentIDL:FinderDef %ComponentIDL:FinderDefAtts;>

<!-- _____ -->
<!-- -->
<!-- METAMODEL CLASS: ComponentIDL:EventPortDef -->
<!-- -->
<!-- _____ -->

<ELEMENT ComponentIDL:EventPortDef.type (ComponentIDL:EventDef)*>
<ENTITY % ComponentIDL:EventPortDefFeatures '%BaseIDL:ContainedFeatures; |
ComponentIDL:EventPortDef.type'>
<ENTITY % ComponentIDL:EventPortDefAtts '%BaseIDL:ContainedAtts;
type IDREFS #IMPLIED'>
<ELEMENT ComponentIDL:EventPortDef (%ComponentIDL:EventPortDefFeatures;)*>
<ATTLIST ComponentIDL:EventPortDef %ComponentIDL:EventPortDefAtts;>

<!-- _____ -->
<!-- -->
<!-- METAMODEL CLASS: ComponentIDL:EmitsDef -->
<!-- -->
<!-- _____ -->

<ENTITY % ComponentIDL:EmitsDefFeatures '%ComponentIDL:EventPortDefFeatures;'>
<ENTITY % ComponentIDL:EmitsDefAtts '%ComponentIDL:EventPortDefAtts;'>
<ELEMENT ComponentIDL:EmitsDef (%ComponentIDL:EmitsDefFeatures;)*>
<ATTLIST ComponentIDL:EmitsDef %ComponentIDL:EmitsDefAtts;>

<!-- _____ -->

```

```

<!--                                     -->
<!-- METAMODEL CLASS: ComponentIDL:ConsumesDef -->
<!--                                     -->
<!-- _____ -->

<!ENTITY % ComponentIDL:ConsumesDefFeatures '%ComponentIDL:EventPortDefFeatures;*>
<!ENTITY % ComponentIDL:ConsumesDefAtts '%ComponentIDL:EventPortDefAtts;*>
<!ELEMENT ComponentIDL:ConsumesDef (%ComponentIDL:ConsumesDefFeatures;)*>
<!ATTLIST ComponentIDL:ConsumesDef %ComponentIDL:ConsumesDefAtts;>

<!-- _____ -->
<!--                                     -->
<!-- METAMODEL CLASS: ComponentIDL:UsesDef -->
<!--                                     -->
<!-- _____ -->

<!ELEMENT ComponentIDL:UsesDef.uses (BaseIDL:InterfaceDef)*>
<!ELEMENT ComponentIDL:UsesDef.multiple EMPTY>
<!ATTLIST ComponentIDL:UsesDef.multiple xmi.value (true|false) #REQUIRED>
<!ENTITY % ComponentIDL:UsesDefFeatures '%BaseIDL:ContainedFeatures; |
ComponentIDL:UsesDef.uses |
ComponentIDL:UsesDef.multiple'>
<!ENTITY % ComponentIDL:UsesDefAtts '%BaseIDL:ContainedAtts;
uses IDREFS #IMPLIED
multiple (true|false) #IMPLIED'>
<!ELEMENT ComponentIDL:UsesDef (%ComponentIDL:UsesDefFeatures;)*>
<!ATTLIST ComponentIDL:UsesDef %ComponentIDL:UsesDefAtts;>

<!-- _____ -->
<!--                                     -->
<!-- METAMODEL CLASS: ComponentIDL:PublishesDef -->
<!--                                     -->
<!-- _____ -->

<!ENTITY % ComponentIDL:PublishesDefFeatures '%ComponentIDL:EventPortDefFeatures;*>
<!ENTITY % ComponentIDL:PublishesDefAtts '%ComponentIDL:EventPortDefAtts;*>
<!ELEMENT ComponentIDL:PublishesDef (%ComponentIDL:PublishesDefFeatures;)*>
<!ATTLIST ComponentIDL:PublishesDef %ComponentIDL:PublishesDefAtts;>

```

```

<!-- _____ -->
<!-- -->
<!-- METAMODEL CLASS: ComponentIDL:EventDef -->
<!-- -->
<!-- _____ -->

<ENTITY % ComponentIDL:EventDefFeatures '%BaseIDL:ValueDefFeatures;'>

<ENTITY % ComponentIDL:EventDefAtts '%BaseIDL:ValueDefAtts;'>

<ELEMENT ComponentIDL:EventDef (%ComponentIDL:EventDefFeatures;)*>

<ATTLIST ComponentIDL:EventDef %ComponentIDL:EventDefAtts;>

```

12.3.2 IDL for the BaseIDL Package

```

#pragma prefix "ccm.omg.org"
#include "Reflective.idl"

module BaseIDL
{
  interface TypedClass;
  interface Typed;
  typedef sequence<Typed> TypedSet;
  interface ParameterDefClass;
  interface ParameterDef;
  typedef sequence<ParameterDef> ParameterDefSet;
  interface ContainedClass;
  interface Contained;
  typedef sequence<Contained> ContainedSet;
  interface ConstantDefClass;
  interface ConstantDef;
  typedef sequence<ConstantDef> ConstantDefSet;
  interface ContainerClass;
  interface Container;
  typedef sequence<Container> ContainerSet;
  interface ModuleDefClass;
  interface ModuleDef;
  typedef sequence<ModuleDef> ModuleDefSet;
  interface IDLTypeClass;
  interface IDLType;
  typedef sequence<IDLType> IDLTypeSet;
  interface TypedefDefClass;
  interface TypedefDef;
  typedef sequence<TypedefDef> TypedefDefSet;
  interface InterfaceDefClass;
  interface InterfaceDef;
  typedef sequence<InterfaceDef> InterfaceDefSet;
  interface FieldClass;

```



```

interface Field;
typedef sequence<Field> FieldSet;
interface StructDefClass;
interface StructDef;
typedef sequence<StructDef> StructDefSet;
interface UnionDefClass;
interface UnionDef;
typedef sequence<UnionDef> UnionDefSet;
interface EnumDefClass;
interface EnumDef;
typedef sequence<EnumDef> EnumDefSet;
interface AliasDefClass;
interface AliasDef;
typedef sequence<AliasDef> AliasDefSet;
interface StringDefClass;
interface StringDef;
typedef sequence<StringDef> StringDefSet;
interface WstringDefClass;
interface WstringDef;
typedef sequence<WstringDef> WstringDefSet;
interface FixedDefClass;
interface FixedDef;
typedef sequence<FixedDef> FixedDefSet;
interface SequenceDefClass;
interface SequenceDef;
typedef sequence<SequenceDef> SequenceDefSet;
interface ArrayDefClass;
interface ArrayDef;
typedef sequence<ArrayDef> ArrayDefSet;
interface PrimitiveDefClass;
interface PrimitiveDef;
typedef sequence<PrimitiveDef> PrimitiveDefSet;
interface UnionFieldClass;
interface UnionField;
typedef sequence<UnionField> UnionFieldSet;
interface ValueMemberDefClass;
interface ValueMemberDef;
typedef sequence<ValueMemberDef> ValueMemberDefSet;
interface ValueDefClass;
interface ValueDef;
typedef sequence<ValueDef> ValueDefSet;
interface ValueBoxDefClass;
interface ValueBoxDef;
typedef sequence<ValueBoxDef> ValueBoxDefSet;
interface OperationDefClass;
interface OperationDef;
typedef sequence<OperationDef> OperationDefSet;
interface ExceptionDefClass;
interface ExceptionDef;
typedef sequence<ExceptionDef> ExceptionDefSet;
interface AttributeDefClass;

```

```

interface AttributeDef;
typedef sequence<AttributeDef> AttributeDefSet;
interface BaseIDLPackage;
enum PrimitiveKind {PK_NULL, PK_VOID, PK_SHORT, PK_LONG, PK_USHORT, PK_ULONG,
PK_FLOAT, PK_DOUBLE, PK_BOOLEAN, PK_CHAR, PK_OCTET, PK_ANY, PK_LONGDOUBLE,
PK_WSTRING, PK_TYPECODE, PK_WCHAR, PK_PRINCIPAL, PK_STRING, PK_ULONGLONG, PK_OBJREF,
PK_LONGLONG};
enum ParameterMode {PARAM_IN, PARAM_OUT, PARAM_INOUT};
enum DefinitionKind {DK_NONE, DK_ALL, DK_ATTRIBUTE, DK_CONSTANT, DK_EXCEPTION,
DK_INTERFACE, DK_MODULE, DK_OPERATION, DK_TYPEDEF, DK_ALIAS, DK_STRUCT, DK_UNION,
DK_FIXED, DK_ENUM, DK_PRIMITIVE, DK_STRING, DK_SEQUENCE, DK_WSTRING, DK_ARRAY,
DK_REPOSITORY};

interface TypedClass : Reflective::RefObject
{
    readonly attribute TypedSet all_of_type_typed;
};

interface Typed : TypedClass
{
    IDLType idl_type ()
        raises (Reflective::MofError);
    void set_idl_type (in IDLType new_value)
        raises (Reflective::MofError);
}; // end of interface Typed

interface ParameterDefClass : TypedClass
{
    readonly attribute ParameterDefSet all_of_type_parameter_def;
    readonly attribute ParameterDefSet all_of_class_parameter_def;
    ParameterDef create_parameter_def (
        in string identifier,
        in ParameterMode direction)
        raises (Reflective::MofError);
};

interface ParameterDef : ParameterDefClass, Typed
{
    string identifier ()
        raises (Reflective::MofError);
    void set_identifier (in string new_value)
        raises (Reflective::MofError);
    ParameterMode direction ()
        raises (Reflective::MofError);
    void set_direction (in ParameterMode new_value)
        raises (Reflective::MofError);
}; // end of interface ParameterDef

interface ContainedClass : Reflective::RefObject
{
    readonly attribute ContainedSet all_of_type_contained;
};

```

```

};

interface Contained : ContainedClass
{
    Container defined_in ()
        raises (Reflective::NotSet, Reflective::MofError);
    void set_defined_in (in Container new_value)
        raises (Reflective::MofError);
    void unset_defined_in ()
        raises (Reflective::MofError);
    string identifier ()
        raises (Reflective::MofError);
    void set_identifier (in string new_value)
        raises (Reflective::MofError);
    string repository_id ()
        raises (Reflective::MofError);
    void set_repository_id (in string new_value)
        raises (Reflective::MofError);
    string version ()
        raises (Reflective::MofError);
    void set_version (in string new_value)
        raises (Reflective::MofError);
    string absolute_name ()
        raises (Reflective::MofError);
}; // end of interface Contained

interface ConstantDefClass : TypedClass, ContainedClass
{
    readonly attribute ConstantDefSet all_of_type_constant_def;
    readonly attribute ConstantDefSet all_of_class_constant_def;
    ConstantDef create_constant_def (
        in string identifier,
        in string repository_id,
        in string version,
        in any const_value)
        raises (Reflective::MofError);
};

interface ConstantDef : ConstantDefClass, Typed, Contained
{
    any const_value ()
        raises (Reflective::MofError);
    void set_const_value (in any new_value)
        raises (Reflective::MofError);
}; // end of interface ConstantDef

interface ContainerClass : ContainedClass
{
    readonly attribute ContainerSet all_of_type_container;
};

```

```

interface Container : ContainerClass, Contained
{
    ContainedSet contents ()
        raises (Reflective::MofError);
    void set_contents (in ContainedSet new_value)
        raises (Reflective::MofError);
    void unset_contents ()
        raises (Reflective::MofError);
    void add_contents (in Contained new_element)
        raises (Reflective::MofError);
    void modify_contents (
        in Contained old_element,
        in Contained new_element)
        raises (Reflective::NotFound, Reflective::MofError);
    void remove_contents (in Contained old_element)
        raises (Reflective::NotFound, Reflective::MofError);
    Contained lookup_name(
        in string search_name,
        in long levels_to_search,
        in DefinitionKind limit_to_type,
        in boolean exclude_inherited)
        raises (Reflective::MofError);
    Contained lookup(
        in string search_name)
        raises (Reflective::MofError);
    Contained get_filtered_contents(
        in DefinitionKind limit_to_type,
        in boolean include_inherited)
        raises (Reflective::MofError);
}; // end of interface Container

interface ModuleDefClass : ContainerClass
{
    readonly attribute ModuleDefSet all_of_type_module_def;
    readonly attribute ModuleDefSet all_of_class_module_def;
    ModuleDef create_module_def (
        in string identifier,
        in string repository_id,
        in string version,
        in string prefix)
        raises (Reflective::MofError);
};

interface ModuleDef : ModuleDefClass, Container
{
    string prefix ()
        raises (Reflective::MofError);
    void set_prefix (in string new_value)
        raises (Reflective::MofError);
}; // end of interface ModuleDef

```

```

interface IDLTypeClass : Reflective::RefObject
{
  readonly attribute IDLTypeSet all_of_type_idltype;
};

interface IDLType : IDLTypeClass
{
  CORBA::TypeCode type_code ()
  raises (Reflective::MofError);
}; // end of interface IDLType

interface TypedefDefClass : IDLTypeClass, ContainedClass
{
  readonly attribute TypedefDefSet all_of_type_typedef_def;
};

interface TypedefDef : TypedefDefClass, IDLType, Contained
{
}; // end of interface TypedefDef

interface InterfaceDefClass : IDLTypeClass, ContainerClass
{
  readonly attribute InterfaceDefSet all_of_type_interface_def;
  readonly attribute InterfaceDefSet all_of_class_interface_def;
  InterfaceDef create_interface_def (
    in string identifier,
    in string repository_id,
    in string version,
    in boolean is_abstract,
    in boolean is_local)
  raises (Reflective::MofError);
};

interface InterfaceDef : InterfaceDefClass, IDLType, Container
{
  InterfaceDefSet base ()
  raises (Reflective::MofError);
  void set_base (in InterfaceDefSet new_value)
  raises (Reflective::MofError);
  void unset_base ()
  raises (Reflective::MofError);
  void add_base (in InterfaceDef new_element)
  raises (Reflective::MofError);
  void modify_base (
    in InterfaceDef old_element,
    in InterfaceDef new_element)
  raises (Reflective::NotFound, Reflective::MofError);
  void remove_base (in InterfaceDef old_element)
  raises (Reflective::NotFound, Reflective::MofError);
  boolean is_abstract ()
  raises (Reflective::MofError);
};

```

```

void set_is_abstract (in boolean new_value)
    raises (Reflective::MofError);
boolean is_local ()
    raises (Reflective::MofError);
void set_is_local (in boolean new_value)
    raises (Reflective::MofError);
}; // end of interface InterfaceDef

interface FieldClass : TypedClass
{
    readonly attribute FieldSet all_of_type_field;
    readonly attribute FieldSet all_of_class_field;
    Field create_field (
        in string identifier)
        raises (Reflective::MofError);
};

interface Field : FieldClass, Typed
{
    string identifier ()
        raises (Reflective::MofError);
    void set_identifier (in string new_value)
        raises (Reflective::MofError);
}; // end of interface Field

interface StructDefClass : TypedefDefClass
{
    readonly attribute StructDefSet all_of_type_struct_def;
    readonly attribute StructDefSet all_of_class_struct_def;
    StructDef create_struct_def (
        in string identifier,
        in string repository_id,
        in string version,
        in Field members)
        raises (Reflective::MofError);
};

interface StructDef : StructDefClass, TypedefDef
{
    Field members ()
        raises (Reflective::MofError);
    void set_members (in Field new_value)
        raises (Reflective::MofError);
}; // end of interface StructDef

interface UnionDefClass : TypedefDefClass
{
    readonly attribute UnionDefSet all_of_type_union_def;
    readonly attribute UnionDefSet all_of_class_union_def;
    UnionDef create_union_def (
        in string identifier,

```

```

    in string repository_id,
    in string version,
    in UnionField union_members)
    raises (Reflective::MofError);
};

interface UnionDef : UnionDefClass, TypedefDef
{
    IDLType discriminator_type ()
        raises (Reflective::MofError);
    void set_discriminator_type (in IDLType new_value)
        raises (Reflective::MofError);
    UnionField union_members ()
        raises (Reflective::MofError);
    void set_union_members (in UnionField new_value)
        raises (Reflective::MofError);
}; // end of interface UnionDef

interface EnumDefClass : TypedefDefClass
{
    readonly attribute EnumDefSet all_of_type_enum_def;
    readonly attribute EnumDefSet all_of_class_enum_def;
    EnumDef create_enum_def (
        in string identifier,
        in string repository_id,
        in string version,
        in string members)
        raises (Reflective::MofError);
};

interface EnumDef : EnumDefClass, TypedefDef
{
    string members ()
        raises (Reflective::MofError);
    void set_members (in string new_value)
        raises (Reflective::MofError);
}; // end of interface EnumDef

interface AliasDefClass : TypedefDefClass, TypedClass
{
    readonly attribute AliasDefSet all_of_type_alias_def;
    readonly attribute AliasDefSet all_of_class_alias_def;
    AliasDef create_alias_def (
        in string identifier,
        in string repository_id,
        in string version)
        raises (Reflective::MofError);
};

interface AliasDef : AliasDefClass, TypedefDef, Typed
{

```

```

}; // end of interface AliasDef

interface StringDefClass : IDLTypeClass
{
  readonly attribute StringDefSet all_of_type_string_def;
  readonly attribute StringDefSet all_of_class_string_def;
  StringDef create_string_def (
    in unsigned long bound)
    raises (Reflective::MofError);
};

interface StringDef : StringDefClass, IDLType
{
  unsigned long bound ()
    raises (Reflective::MofError);
  void set_bound (in unsigned long new_value)
    raises (Reflective::MofError);
}; // end of interface StringDef

interface WstringDefClass : IDLTypeClass
{
  readonly attribute WstringDefSet all_of_type_wstring_def;
  readonly attribute WstringDefSet all_of_class_wstring_def;
  WstringDef create_wstring_def (
    in unsigned long bound)
    raises (Reflective::MofError);
};

interface WstringDef : WstringDefClass, IDLType
{
  unsigned long bound ()
    raises (Reflective::MofError);
  void set_bound (in unsigned long new_value)
    raises (Reflective::MofError);
}; // end of interface WstringDef

interface FixedDefClass : IDLTypeClass
{
  readonly attribute FixedDefSet all_of_type_fixed_def;
  readonly attribute FixedDefSet all_of_class_fixed_def;
  FixedDef create_fixed_def (
    in unsigned short digits,
    in short scale)
    raises (Reflective::MofError);
};

interface FixedDef : FixedDefClass, IDLType
{
  unsigned short digits ()
    raises (Reflective::MofError);
  void set_digits (in unsigned short new_value)

```



```

    raises (Reflective::MofError);
short scale ()
    raises (Reflective::MofError);
void set_scale (in short new_value)
    raises (Reflective::MofError);
}; // end of interface FixedDef

interface SequenceDefClass : TypedClass, IDLTypeClass
{
    readonly attribute SequenceDefSet all_of_type_sequence_def;
    readonly attribute SequenceDefSet all_of_class_sequence_def;
    SequenceDef create_sequence_def (
        in unsigned long bound)
        raises (Reflective::MofError);
};

interface SequenceDef : SequenceDefClass, Typed, IDLType
{
    unsigned long bound ()
        raises (Reflective::MofError);
    void set_bound (in unsigned long new_value)
        raises (Reflective::MofError);
}; // end of interface SequenceDef

interface ArrayDefClass : TypedClass, IDLTypeClass
{
    readonly attribute ArrayDefSet all_of_type_array_def;
    readonly attribute ArrayDefSet all_of_class_array_def;
    ArrayDef create_array_def (
        in unsigned long bound)
        raises (Reflective::MofError);
};

interface ArrayDef : ArrayDefClass, Typed, IDLType
{
    unsigned long bound ()
        raises (Reflective::MofError);
    void set_bound (in unsigned long new_value)
        raises (Reflective::MofError);
}; // end of interface ArrayDef

interface PrimitiveDefClass : IDLTypeClass
{
    readonly attribute PrimitiveDefSet all_of_type_primitive_def;
    readonly attribute PrimitiveDefSet all_of_class_primitive_def;
    PrimitiveDef create_primitive_def (
        in PrimitiveKind kind)
        raises (Reflective::MofError);
};

interface PrimitiveDef : PrimitiveDefClass, IDLType

```

```

{
  PrimitiveKind kind ()
    raises (Reflective::MofError);
  void set_kind (in PrimitiveKind new_value)
    raises (Reflective::MofError);
}; // end of interface PrimitiveDef

interface UnionFieldClass : TypedClass
{
  readonly attribute UnionFieldSet all_of_type_union_field;
  readonly attribute UnionFieldSet all_of_class_union_field;
  UnionField create_union_field (
    in string identifier,
    in any label)
    raises (Reflective::MofError);
};

interface UnionField : UnionFieldClass, Typed
{
  string identifier ()
    raises (Reflective::MofError);
  void set_identifier (in string new_value)
    raises (Reflective::MofError);
  any label ()
    raises (Reflective::MofError);
  void set_label (in any new_value)
    raises (Reflective::MofError);
}; // end of interface UnionField

interface ValueMemberDefClass : TypedClass, ContainedClass
{
  readonly attribute ValueMemberDefSet all_of_type_value_member_def;
  readonly attribute ValueMemberDefSet all_of_class_value_member_def;
  ValueMemberDef create_value_member_def (
    in string identifier,
    in string repository_id,
    in string version,
    in boolean is_public_member)
    raises (Reflective::MofError);
};

interface ValueMemberDef : ValueMemberDefClass, Typed, Contained
{
  boolean is_public_member ()
    raises (Reflective::MofError);
  void set_is_public_member (in boolean new_value)
    raises (Reflective::MofError);
}; // end of interface ValueMemberDef

interface ValueDefClass : ContainerClass, IDLTypeClass
{

```

```

readonly attribute ValueDefSet all_of_type_value_def;
readonly attribute ValueDefSet all_of_class_value_def;
ValueDef create_value_def (
    in string identifier,
    in string repository_id,
    in string version,
    in boolean is_abstract,
    in boolean is_custom,
    in boolean is_truncatable)
    raises (Reflective::MofError);
};

interface ValueDef : ValueDefClass, Container, IDLType
{
    InterfaceDef interface_def ()
        raises (Reflective::NotSet, Reflective::MofError);
    void set_interface_def (in InterfaceDef new_value)
        raises (Reflective::MofError);
    void unset_interface_def ()
        raises (Reflective::MofError);
    ValueDef base ()
        raises (Reflective::NotSet, Reflective::MofError);
    void set_base (in ValueDef new_value)
        raises (Reflective::MofError);
    void unset_base ()
        raises (Reflective::MofError);
    ValueDefSet abstract_base ()
        raises (Reflective::MofError);
    void set_abstract_base (in ValueDefSet new_value)
        raises (Reflective::MofError);
    void unset_abstract_base ()
        raises (Reflective::MofError);
    void add_abstract_base (in ValueDef new_element)
        raises (Reflective::MofError);
    void modify_abstract_base (
        in ValueDef old_element,
        in ValueDef new_element)
        raises (Reflective::NotFound, Reflective::MofError);
    void remove_abstract_base (in ValueDef old_element)
        raises (Reflective::NotFound, Reflective::MofError);
    boolean is_abstract ()
        raises (Reflective::MofError);
    void set_is_abstract (in boolean new_value)
        raises (Reflective::MofError);
    boolean is_custom ()
        raises (Reflective::MofError);
    void set_is_custom (in boolean new_value)
        raises (Reflective::MofError);
    boolean is_truncatable ()
        raises (Reflective::MofError);
    void set_is_truncatable (in boolean new_value)

```

```

    raises (Reflective::MofError);
}; // end of interface ValueDef

interface ValueBoxDefClass : TypedDefClass
{
    readonly attribute ValueBoxDefSet all_of_type_value_box_def;
    readonly attribute ValueBoxDefSet all_of_class_value_box_def;
    ValueBoxDef create_value_box_def (
        in string identifier,
        in string repository_id,
        in string version)
        raises (Reflective::MofError);
};

interface ValueBoxDef : ValueBoxDefClass, TypedDef
{
}; // end of interface ValueBoxDef

interface OperationDefClass : TypedClass, ContainedClass
{
    readonly attribute OperationDefSet all_of_type_operation_def;
    readonly attribute OperationDefSet all_of_class_operation_def;
    OperationDef create_operation_def (
        in string identifier,
        in string repository_id,
        in string version,
        in boolean is_oneway,
        in ParameterDef parameters,
        in string contexts)
        raises (Reflective::MofError);
};

interface OperationDef : OperationDefClass, Typed, Contained
{
    ExceptionDefSet exception_def ()
        raises (Reflective::MofError);
    void set_exception_def (in ExceptionDefSet new_value)
        raises (Reflective::MofError);
    void unset_exception_def ()
        raises (Reflective::MofError);
    void add_exception_def (in ExceptionDef new_element)
        raises (Reflective::MofError);
    void modify_exception_def (
        in ExceptionDef old_element,
        in ExceptionDef new_element)
        raises (Reflective::NotFound, Reflective::MofError);
    void remove_exception_def (in ExceptionDef old_element)
        raises (Reflective::NotFound, Reflective::MofError);
    boolean is_oneway ()
        raises (Reflective::MofError);
    void set_is_oneway (in boolean new_value)

```

```

    raises (Reflective::MofError);
ParameterDef parameters ()
    raises (Reflective::MofError);
void set_parameters (in ParameterDef new_value)
    raises (Reflective::MofError);
string contexts ()
    raises (Reflective::MofError);
void set_contexts (in string new_value)
    raises (Reflective::MofError);
}; // end of interface OperationDef

interface ExceptionDefClass : ContainedClass
{
    readonly attribute ExceptionDefSet all_of_type_exception_def;
    readonly attribute ExceptionDefSet all_of_class_exception_def;
    ExceptionDef create_exception_def (
        in string identifier,
        in string repository_id,
        in string version,
        in CORBA::TypeCode type_code,
        in Field members)
        raises (Reflective::MofError);
};

interface ExceptionDef : ExceptionDefClass, Contained
{
    CORBA::TypeCode type_code ()
        raises (Reflective::MofError);
    Field members ()
        raises (Reflective::MofError);
    void set_members (in Field new_value)
        raises (Reflective::MofError);
}; // end of interface ExceptionDef

interface AttributeDefClass : TypedClass, ContainedClass
{
    readonly attribute AttributeDefSet all_of_type_attribute_def;
    readonly attribute AttributeDefSet all_of_class_attribute_def;
    AttributeDef create_attribute_def (
        in string identifier,
        in string repository_id,
        in string version,
        in boolean is_readonly)
        raises (Reflective::MofError);
};

interface AttributeDef : AttributeDefClass, Typed, Contained
{
    ExceptionDefSet set_exception ()
        raises (Reflective::MofError);
    void set_set_exception (in ExceptionDefSet new_value)

```

```

    raises (Reflective::MofError);
void unset_set_exception ()
    raises (Reflective::MofError);
void add_set_exception (in ExceptionDef new_element)
    raises (Reflective::MofError);
void modify_set_exception (
    in ExceptionDef old_element,
    in ExceptionDef new_element)
    raises (Reflective::NotFound, Reflective::MofError);
void remove_set_exception (in ExceptionDef old_element)
    raises (Reflective::NotFound, Reflective::MofError);
ExceptionDefSet get_exception ()
    raises (Reflective::MofError);
void set_get_exception (in ExceptionDefSet new_value)
    raises (Reflective::MofError);
void unset_get_exception ()
    raises (Reflective::MofError);
void add_get_exception (in ExceptionDef new_element)
    raises (Reflective::MofError);
void modify_get_exception (
    in ExceptionDef old_element,
    in ExceptionDef new_element)
    raises (Reflective::NotFound, Reflective::MofError);
void remove_get_exception (in ExceptionDef old_element)
    raises (Reflective::NotFound, Reflective::MofError);
boolean is_readonly ()
    raises (Reflective::MofError);
void set_is_readonly (in boolean new_value)
    raises (Reflective::MofError);
}; // end of interface AttributeDef

struct InterfaceDerivedFromLink
{
    InterfaceDef base;
    InterfaceDef derived;
};
typedef sequence<InterfaceDerivedFromLink> InterfaceDerivedFromLinkSet;

interface InterfaceDerivedFrom : Reflective::RefAssociation
{
    InterfaceDerivedFromLinkSet all_interface_derived_from_links()
        raises (Reflective::MofError);
    boolean exists (
        in InterfaceDef base,
        in InterfaceDef derived)
        raises (Reflective::MofError);
    InterfaceDefSet base (in InterfaceDef derived)
        raises (Reflective::MofError);
    void add (
        in InterfaceDef base,
        in InterfaceDef derived)

```

```

    raises (Reflective::MofError);
void modify_base (
    in InterfaceDef base,
    in InterfaceDef derived,
    in InterfaceDef new_base)
    raises (Reflective::NotFound, Reflective::MofError);
void remove (
    in InterfaceDef base,
    in InterfaceDef derived)
    raises (Reflective::NotFound, Reflective::MofError);
}; // end of interface InterfaceDerivedFrom

struct DiscriminatedByLink
{
    IDLType discriminator_type;
    UnionDef union_def;
};
typedef sequence<DiscriminatedByLink> DiscriminatedByLinkSet;

interface DiscriminatedBy : Reflective::RefAssociation
{
    DiscriminatedByLinkSet all_discriminated_by_links()
        raises (Reflective::MofError);
    boolean exists (
        in IDLType discriminator_type,
        in UnionDef union_def)
        raises (Reflective::MofError);
    IDLType discriminator_type (in UnionDef union_def)
        raises (Reflective::MofError);
    void add (
        in IDLType discriminator_type,
        in UnionDef union_def)
        raises (Reflective::MofError);
    void modify_discriminator_type (
        in IDLType discriminator_type,
        in UnionDef union_def,
        in IDLType new_discriminator_type)
        raises (Reflective::NotFound, Reflective::MofError);
    void remove (
        in IDLType discriminator_type,
        in UnionDef union_def)
        raises (Reflective::NotFound, Reflective::MofError);
}; // end of interface DiscriminatedBy

struct TypedByLink
{
    IDLType idl_type;
    BaseIDL::Typed typed;
};
typedef sequence<TypedByLink> TypedByLinkSet;

```

```

interface TypedBy : Reflective::RefAssociation
{
  TypedByLinkSet all_typed_by_links()
    raises (Reflective::MofError);
  boolean exists (
    in IDLType idl_type,
    in BaseIDL::Typed typed)
    raises (Reflective::MofError);
  IDLType idl_type (in BaseIDL::Typed typed)
    raises (Reflective::MofError);
  void add (
    in IDLType idl_type,
    in BaseIDL::Typed typed)
    raises (Reflective::MofError);
  void modify_idl_type (
    in IDLType idl_type,
    in BaseIDL::Typed typed,
    in IDLType new_idl_type)
    raises (Reflective::NotFound, Reflective::MofError);
  void remove (
    in IDLType idl_type,
    in BaseIDL::Typed typed)
    raises (Reflective::NotFound, Reflective::MofError);
}; // end of interface TypedBy

```

```

struct SupportsLink
{
  InterfaceDef interface_def;
  ValueDef value_def;
};
typedef sequence<SupportsLink> SupportsLinkSet;

```

```

interface Supports : Reflective::RefAssociation
{
  SupportsLinkSet all_supports_links()
    raises (Reflective::MofError);
  boolean exists (
    in InterfaceDef interface_def,
    in ValueDef value_def)
    raises (Reflective::MofError);
  InterfaceDef interface_def (in ValueDef value_def)
    raises (Reflective::MofError);
  void add (
    in InterfaceDef interface_def,
    in ValueDef value_def)
    raises (Reflective::MofError);
  void modify_interface_def (
    in InterfaceDef interface_def,
    in ValueDef value_def,
    in InterfaceDef new_interface_def)
    raises (Reflective::NotFound, Reflective::MofError);

```



```

void remove (
    in InterfaceDef interface_def,
    in ValueDef value_def)
    raises (Reflective::NotFound, Reflective::MofError);
}; // end of interface Supports

struct ValueDerivedFromLink
{
    ValueDef base;
    ValueDef derived;
};
typedef sequence<ValueDerivedFromLink> ValueDerivedFromLinkSet;

interface ValueDerivedFrom : Reflective::RefAssociation
{
    ValueDerivedFromLinkSet all_value_derived_from_links()
        raises (Reflective::MofError);
    boolean exists (
        in ValueDef base,
        in ValueDef derived)
        raises (Reflective::MofError);
    ValueDef base (in ValueDef derived)
        raises (Reflective::MofError);
    void add (
        in ValueDef base,
        in ValueDef derived)
        raises (Reflective::MofError);
    void modify_base (
        in ValueDef base,
        in ValueDef derived,
        in ValueDef new_base)
        raises (Reflective::NotFound, Reflective::MofError);
    void remove (
        in ValueDef base,
        in ValueDef derived)
        raises (Reflective::NotFound, Reflective::MofError);
}; // end of interface ValueDerivedFrom

struct AbstractDerivedFromLink
{
    ValueDef abstract_derived;
    ValueDef abstract_base;
};
typedef sequence<AbstractDerivedFromLink> AbstractDerivedFromLinkSet;

interface AbstractDerivedFrom : Reflective::RefAssociation
{
    AbstractDerivedFromLinkSet all_abstract_derived_from_links()
        raises (Reflective::MofError);
    boolean exists (
        in ValueDef abstract_derived,

```

```

    in ValueDef abstract_base)
    raises (Reflective::MofError);
ValueDefSet abstract_base (in ValueDef abstract_derived)
    raises (Reflective::MofError);
void add (
    in ValueDef abstract_derived,
    in ValueDef abstract_base)
    raises (Reflective::MofError);
void modify_abstract_base (
    in ValueDef abstract_derived,
    in ValueDef abstract_base,
    in ValueDef new_abstract_base)
    raises (Reflective::NotFound, Reflective::MofError);
void remove (
    in ValueDef abstract_derived,
    in ValueDef abstract_base)
    raises (Reflective::NotFound, Reflective::MofError);
}; // end of interface AbstractDerivedFrom

struct SetRaisesLink
{
    ExceptionDef set_exception;
    AttributeDef set_attribute;
};
typedef sequence<SetRaisesLink> SetRaisesLinkSet;

interface SetRaises : Reflective::RefAssociation
{
    SetRaisesLinkSet all_set_raises_links()
        raises (Reflective::MofError);
    boolean exists (
        in ExceptionDef set_exception,
        in AttributeDef set_attribute)
        raises (Reflective::MofError);
    ExceptionDefSet set_exception (in AttributeDef set_attribute)
        raises (Reflective::MofError);
    void add (
        in ExceptionDef set_exception,
        in AttributeDef set_attribute)
        raises (Reflective::MofError);
    void modify_set_exception (
        in ExceptionDef set_exception,
        in AttributeDef set_attribute,
        in ExceptionDef new_set_exception)
        raises (Reflective::NotFound, Reflective::MofError);
    void remove (
        in ExceptionDef set_exception,
        in AttributeDef set_attribute)
        raises (Reflective::NotFound, Reflective::MofError);
}; // end of interface SetRaises

```

```

struct CanRaiseLink
{
    ExceptionDef exception_def;
    OperationDef operation_def;
};
typedef sequence<CanRaiseLink> CanRaiseLinkSet;

interface CanRaise : Reflective::RefAssociation
{
    CanRaiseLinkSet all_can_raise_links()
        raises (Reflective::MofError);
    boolean exists (
        in ExceptionDef exception_def,
        in OperationDef operation_def)
        raises (Reflective::MofError);
    ExceptionDefSet exception_def (in OperationDef operation_def)
        raises (Reflective::MofError);
    void add (
        in ExceptionDef exception_def,
        in OperationDef operation_def)
        raises (Reflective::MofError);
    void modify_exception_def (
        in ExceptionDef exception_def,
        in OperationDef operation_def,
        in ExceptionDef new_exception_def)
        raises (Reflective::NotFound, Reflective::MofError);
    void remove (
        in ExceptionDef exception_def,
        in OperationDef operation_def)
        raises (Reflective::NotFound, Reflective::MofError);
}; // end of interface CanRaise

struct GetRaisesLink
{
    ExceptionDef get_exception;
    AttributeDef get_attribute;
};
typedef sequence<GetRaisesLink> GetRaisesLinkSet;

interface GetRaises : Reflective::RefAssociation
{
    GetRaisesLinkSet all_get_raises_links()
        raises (Reflective::MofError);
    boolean exists (
        in ExceptionDef get_exception,
        in AttributeDef get_attribute)
        raises (Reflective::MofError);
    ExceptionDefSet get_exception (in AttributeDef get_attribute)
        raises (Reflective::MofError);
    void add (
        in ExceptionDef get_exception,

```

```

    in AttributeDef get_attribute)
    raises (Reflective::MofError);
void modify_get_exception (
    in ExceptionDef get_exception,
    in AttributeDef get_attribute,
    in ExceptionDef new_get_exception)
    raises (Reflective::NotFound, Reflective::MofError);
void remove (
    in ExceptionDef get_exception,
    in AttributeDef get_attribute)
    raises (Reflective::NotFound, Reflective::MofError);
}; // end of interface GetRaises

```

```

struct ContainsLink
{
    Container defined_in;
    Contained contents;
};
typedef sequence<ContainsLink> ContainsLinkSet;

```

```

interface Contains : Reflective::RefAssociation
{
    ContainsLinkSet all_contains_links()
        raises (Reflective::MofError);
    boolean exists (
        in Container defined_in,
        in Contained contents)
        raises (Reflective::MofError);
    Container defined_in (in Contained contents)
        raises (Reflective::MofError);
    ContainedSet contents (in Container defined_in)
        raises (Reflective::MofError);
    void add (
        in Container defined_in,
        in Contained contents)
        raises (Reflective::MofError);
    void modify_defined_in (
        in Container defined_in,
        in Contained contents,
        in Container new_defined_in)
        raises (Reflective::NotFound, Reflective::MofError);
    void modify_contents (
        in Container defined_in,
        in Contained contents,
        in Contained new_contents)
        raises (Reflective::NotFound, Reflective::MofError);
    void remove (
        in Container defined_in,
        in Contained contents)
        raises (Reflective::NotFound, Reflective::MofError);
}; // end of interface Contains

```

```

interface BaseIDLPackageFactory
{
    BaseIDLPackage create_base_idl_package ()
        raises (Reflective::MofError);
};

interface BaseIDLPackage : Reflective::RefPackage
{
    readonly attribute TypedClass typed_ref;
    readonly attribute ParameterDefClass parameter_def_ref;
    readonly attribute ContainedClass contained_ref;
    readonly attribute ConstantDefClass constant_def_ref;
    readonly attribute ContainerClass container_ref;
    readonly attribute ModuleDefClass module_def_ref;
    readonly attribute IDLTypeClass idltype_ref;
    readonly attribute TypedefDefClass typedef_def_ref;
    readonly attribute InterfaceDefClass interface_def_ref;
    readonly attribute FieldClass field_ref;
    readonly attribute StructDefClass struct_def_ref;
    readonly attribute UnionDefClass union_def_ref;
    readonly attribute EnumDefClass enum_def_ref;
    readonly attribute AliasDefClass alias_def_ref;
    readonly attribute StringDefClass string_def_ref;
    readonly attribute WstringDefClass wstring_def_ref;
    readonly attribute FixedDefClass fixed_def_ref;
    readonly attribute SequenceDefClass sequence_def_ref;
    readonly attribute ArrayDefClass array_def_ref;
    readonly attribute PrimitiveDefClass primitive_def_ref;
    readonly attribute UnionFieldClass union_field_ref;
    readonly attribute ValueMemberDefClass value_member_def_ref;
    readonly attribute ValueDefClass value_def_ref;
    readonly attribute ValueBoxDefClass value_box_def_ref;
    readonly attribute OperationDefClass operation_def_ref;
    readonly attribute ExceptionDefClass exception_def_ref;
    readonly attribute AttributeDefClass attribute_def_ref;
    readonly attribute InterfaceDerivedFrom interface_derived_from_ref;
    readonly attribute DiscriminatedBy discriminated_by_ref;
    readonly attribute TypedBy typed_by_ref;
    readonly attribute Supports supports_ref;
    readonly attribute ValueDerivedFrom value_derived_from_ref;
    readonly attribute AbstractDerivedFrom abstract_derived_from_ref;
    readonly attribute SetRaises set_raises_ref;
    readonly attribute CanRaise can_raise_ref;
    readonly attribute GetRaises get_raises_ref;
    readonly attribute Contains contains_ref;
};
}; // end of module BaseIDL

```

12.3.3 IDL for the ComponentIDL Package

```
#pragma prefix "ccm.omg.org"
#include "BaseIDL.idl"

module ComponentIDL
{
  interface ComponentFeatureClass;
  interface ComponentFeature;
  typedef sequence<ComponentFeature> ComponentFeatureSet;
  interface ComponentDefClass;
  interface ComponentDef;
  typedef sequence<ComponentDef> ComponentDefSet;
  interface ProvidesDefClass;
  interface ProvidesDef;
  typedef sequence<ProvidesDef> ProvidesDefSet;
  interface HomeDefClass;
  interface HomeDef;
  typedef sequence<HomeDef> HomeDefSet;
  interface FactoryDefClass;
  interface FactoryDef;
  typedef sequence<FactoryDef> FactoryDefSet;
  interface FinderDefClass;
  interface FinderDef;
  typedef sequence<FinderDef> FinderDefSet;
  interface EventPortDefClass;
  interface EventPortDef;
  typedef sequence<EventPortDef> EventPortDefSet;
  interface EmitsDefClass;
  interface EmitsDef;
  typedef sequence<EmitsDef> EmitsDefSet;
  interface ConsumesDefClass;
  interface ConsumesDef;
  typedef sequence<ConsumesDef> ConsumesDefSet;
  interface UsesDefClass;
  interface UsesDef;
  typedef sequence<UsesDef> UsesDefSet;
  interface PublishesDefClass;
  interface PublishesDef;
  typedef sequence<PublishesDef> PublishesDefSet;
  interface EventDefClass;
  interface EventDef;
  typedef sequence<EventDef> EventDefSet;
  interface ComponentIDLPackage;

  interface ComponentFeatureClass : Reflective::RefObject
  {
    readonly attribute ComponentFeatureSet all_of_type_component_feature;
  };
}
```

```

interface ComponentFeature : ComponentFeatureClass
{
}; // end of interface ComponentFeature

interface ComponentDefClass : BaseIDL::InterfaceDefClass, ComponentFeatureClass
{
  readonly attribute ComponentDefSet all_of_type_component_def;
  readonly attribute ComponentDefSet all_of_class_component_def;
  ComponentDef create_component_def (
    in string identifier,
    in string repository_id,
    in string version,
    in boolean is_abstract,
    in boolean is_local)
    raises (Reflective::MofError);
};

interface ComponentDef : ComponentDefClass, BaseIDL::InterfaceDef, ComponentFeature
{
  ProvidesDefSet facet ()
    raises (Reflective::MofError);
  void set_facet (in ProvidesDefSet new_value)
    raises (Reflective::MofError);
  void unset_facet ()
    raises (Reflective::MofError);
  void add_facet (in ProvidesDef new_element)
    raises (Reflective::MofError);
  void modify_facet (
    in ProvidesDef old_element,
    in ProvidesDef new_element)
    raises (Reflective::NotFound, Reflective::MofError);
  void remove_facet (in ProvidesDef old_element)
    raises (Reflective::NotFound, Reflective::MofError);
  EmitsDefSet emits ()
    raises (Reflective::MofError);
  void set_emits (in EmitsDefSet new_value)
    raises (Reflective::MofError);
  void unset_emits ()
    raises (Reflective::MofError);
  void add_emits (in EmitsDef new_element)
    raises (Reflective::MofError);
  void modify_emits (
    in EmitsDef old_element,
    in EmitsDef new_element)
    raises (Reflective::NotFound, Reflective::MofError);
  void remove_emits (in EmitsDef old_element)
    raises (Reflective::NotFound, Reflective::MofError);
  ConsumesDefSet consumes ()
    raises (Reflective::MofError);
  void set_consumes (in ConsumesDefSet new_value)
    raises (Reflective::MofError);
};

```

```

void unset_consumes ()
    raises (Reflective::MofError);
void add_consumes (in ConsumesDef new_element)
    raises (Reflective::MofError);
void modify_consumes (
    in ConsumesDef old_element,
    in ConsumesDef new_element)
    raises (Reflective::NotFound, Reflective::MofError);
void remove_consumes (in ConsumesDef old_element)
    raises (Reflective::NotFound, Reflective::MofError);
UsesDefSet receptacle ()
    raises (Reflective::MofError);
void set_receptacle (in UsesDefSet new_value)
    raises (Reflective::MofError);
void unset_receptacle ()
    raises (Reflective::MofError);
void add_receptacle (in UsesDef new_element)
    raises (Reflective::MofError);
void modify_receptacle (
    in UsesDef old_element,
    in UsesDef new_element)
    raises (Reflective::NotFound, Reflective::MofError);
void remove_receptacle (in UsesDef old_element)
    raises (Reflective::NotFound, Reflective::MofError);
BaseIDL::InterfaceDefSet supports ()
    raises (Reflective::MofError);
void set_supports (in BaseIDL::InterfaceDefSet new_value)
    raises (Reflective::MofError);
void unset_supports ()
    raises (Reflective::MofError);
void add_supports (in BaseIDL::InterfaceDef new_element)
    raises (Reflective::MofError);
void modify_supports (
    in BaseIDL::InterfaceDef old_element,
    in BaseIDL::InterfaceDef new_element)
    raises (Reflective::NotFound, Reflective::MofError);
void remove_supports (in BaseIDL::InterfaceDef old_element)
    raises (Reflective::NotFound, Reflective::MofError);
PublishesDefSet publishes ()
    raises (Reflective::MofError);
void set_publishes (in PublishesDefSet new_value)
    raises (Reflective::MofError);
void unset_publishes ()
    raises (Reflective::MofError);
void add_publishes (in PublishesDef new_element)
    raises (Reflective::MofError);
void modify_publishes (
    in PublishesDef old_element,
    in PublishesDef new_element)
    raises (Reflective::NotFound, Reflective::MofError);
void remove_publishes (in PublishesDef old_element)

```



```
    raises (Reflective::NotFound, Reflective::MofError);
}; // end of interface ComponentDef
```

```
interface ProvidesDefClass : BaseIDL::ContainedClass, ComponentFeatureClass
{
    readonly attribute ProvidesDefSet all_of_type_provides_def;
    readonly attribute ProvidesDefSet all_of_class_provides_def;
    ProvidesDef create_provides_def (
        in string identifier,
        in string repository_id,
        in string version)
        raises (Reflective::MofError);
};
```

```
interface ProvidesDef : ProvidesDefClass, BaseIDL::Contained, ComponentFeature
{
    BaseIDL::InterfaceDef provides ()
        raises (Reflective::MofError);
    void set_provides (in BaseIDL::InterfaceDef new_value)
        raises (Reflective::MofError);
    ComponentDef component ()
        raises (Reflective::MofError);
    void set_component (in ComponentDef new_value)
        raises (Reflective::MofError);
}; // end of interface ProvidesDef
```

```
interface HomeDefClass : BaseIDL::InterfaceDefClass
{
    readonly attribute HomeDefSet all_of_type_home_def;
    readonly attribute HomeDefSet all_of_class_home_def;
    HomeDef create_home_def (
        in string identifier,
        in string repository_id,
        in string version,
        in boolean is_abstract,
        in boolean is_local)
        raises (Reflective::MofError);
};
```

```
interface HomeDef : HomeDefClass, BaseIDL::InterfaceDef
{
    FinderDefSet finder ()
        raises (Reflective::MofError);
    void set_finder (in FinderDefSet new_value)
        raises (Reflective::MofError);
    void unset_finder ()
        raises (Reflective::MofError);
    void add_finder (in FinderDef new_element)
        raises (Reflective::MofError);
    void modify_finder (
        in FinderDef old_element,
```

```

    in FinderDef new_element)
    raises (Reflective::NotFound, Reflective::MofError);
void remove_finder (in FinderDef old_element)
    raises (Reflective::NotFound, Reflective::MofError);
ComponentDef component ()
    raises (Reflective::MofError);
void set_component (in ComponentDef new_value)
    raises (Reflective::MofError);
FactoryDefSet factory ()
    raises (Reflective::MofError);
void set_factory (in FactoryDefSet new_value)
    raises (Reflective::MofError);
void unset_factory ()
    raises (Reflective::MofError);
void add_factory (in FactoryDef new_element)
    raises (Reflective::MofError);
void modify_factory (
    in FactoryDef old_element,
    in FactoryDef new_element)
    raises (Reflective::NotFound, Reflective::MofError);
void remove_factory (in FactoryDef old_element)
    raises (Reflective::NotFound, Reflective::MofError);
BaseIDL::ValueDef primary_key ()
    raises (Reflective::NotSet, Reflective::MofError);
void set_primary_key (in BaseIDL::ValueDef new_value)
    raises (Reflective::MofError);
void unset_primary_key ()
    raises (Reflective::MofError);
BaseIDL::InterfaceDefSet supports ()
    raises (Reflective::MofError);
void set_supports (in BaseIDL::InterfaceDefSet new_value)
    raises (Reflective::MofError);
void unset_supports ()
    raises (Reflective::MofError);
void add_supports (in BaseIDL::InterfaceDef new_element)
    raises (Reflective::MofError);
void modify_supports (
    in BaseIDL::InterfaceDef old_element,
    in BaseIDL::InterfaceDef new_element)
    raises (Reflective::NotFound, Reflective::MofError);
void remove_supports (in BaseIDL::InterfaceDef old_element)
    raises (Reflective::NotFound, Reflective::MofError);
}; // end of interface HomeDef

```

```

interface FactoryDefClass : BaseIDL::OperationDefClass
{
    readonly attribute FactoryDefSet all_of_type_factory_def;
    readonly attribute FactoryDefSet all_of_class_factory_def;
    FactoryDef create_factory_def (
        in string identifier,
        in string repository_id,

```

```

    in string version,
    in boolean is_oneway,
    in BaseIDL::ParameterDef parameters,
    in string contexts)
    raises (Reflective::MofError);
};

```

```

interface FactoryDef : FactoryDefClass, BaseIDL::OperationDef
{
    HomeDef home ()
        raises (Reflective::MofError);
    void set_home (in HomeDef new_value)
        raises (Reflective::MofError);
}; // end of interface FactoryDef

```

```

interface FinderDefClass : BaseIDL::OperationDefClass
{
    readonly attribute FinderDefSet all_of_type_finder_def;
    readonly attribute FinderDefSet all_of_class_finder_def;
    FinderDef create_finder_def (
        in string identifier,
        in string repository_id,
        in string version,
        in boolean is_oneway,
        in BaseIDL::ParameterDef parameters,
        in string contexts)
        raises (Reflective::MofError);
};

```

```

interface FinderDef : FinderDefClass, BaseIDL::OperationDef
{
    HomeDef home ()
        raises (Reflective::MofError);
    void set_home (in HomeDef new_value)
        raises (Reflective::MofError);
}; // end of interface FinderDef

```

```

interface EventPortDefClass : BaseIDL::ContainedClass, ComponentFeatureClass
{
    readonly attribute EventPortDefSet all_of_type_event_port_def;
};

```

```

interface EventPortDef : EventPortDefClass, BaseIDL::Contained, ComponentFeature
{
    EventDef type ()
        raises (Reflective::MofError);
    void set_type (in EventDef new_value)
        raises (Reflective::MofError);
}; // end of interface EventPortDef

```

```

interface EmitsDefClass : EventPortDefClass, ComponentFeatureClass

```

```

{
  readonly attribute EmitsDefSet all_of_type_emits_def;
  readonly attribute EmitsDefSet all_of_class_emits_def;
  EmitsDef create_emits_def (
    in string identifier,
    in string repository_id,
    in string version)
    raises (Reflective::MofError);
};

interface EmitsDef : EmitsDefClass, EventPortDef, ComponentFeature
{
  ComponentDef component ()
    raises (Reflective::MofError);
  void set_component (in ComponentDef new_value)
    raises (Reflective::MofError);
}; // end of interface EmitsDef

interface ConsumesDefClass : EventPortDefClass, ComponentFeatureClass
{
  readonly attribute ConsumesDefSet all_of_type_consumes_def;
  readonly attribute ConsumesDefSet all_of_class_consumes_def;
  ConsumesDef create_consumes_def (
    in string identifier,
    in string repository_id,
    in string version)
    raises (Reflective::MofError);
};

interface ConsumesDef : ConsumesDefClass, EventPortDef, ComponentFeature
{
  ComponentDef component ()
    raises (Reflective::MofError);
  void set_component (in ComponentDef new_value)
    raises (Reflective::MofError);
}; // end of interface ConsumesDef

interface UsesDefClass : BaseIDL::ContainedClass, ComponentFeatureClass
{
  readonly attribute UsesDefSet all_of_type_uses_def;
  readonly attribute UsesDefSet all_of_class_uses_def;
  UsesDef create_uses_def (
    in string identifier,
    in string repository_id,
    in string version,
    in boolean multiple)
    raises (Reflective::MofError);
};

interface UsesDef : UsesDefClass, BaseIDL::Contained, ComponentFeature
{

```

```

ComponentDef component ()
    raises (Reflective::MofError);
void set_component (in ComponentDef new_value)
    raises (Reflective::MofError);
BaseIDL::InterfaceDef uses ()
    raises (Reflective::MofError);
void set_uses (in BaseIDL::InterfaceDef new_value)
    raises (Reflective::MofError);
boolean multiple ()
    raises (Reflective::MofError);
void set_multiple (in boolean new_value)
    raises (Reflective::MofError);
}; // end of interface UsesDef

interface PublishesDefClass : EventPortDefClass, ComponentFeatureClass
{
    readonly attribute PublishesDefSet all_of_type_publishes_def;
    readonly attribute PublishesDefSet all_of_class_publishes_def;
    PublishesDef create_publishes_def (
        in string identifier,
        in string repository_id,
        in string version)
        raises (Reflective::MofError);
};

interface PublishesDef : PublishesDefClass, EventPortDef, ComponentFeature
{
    ComponentDef component ()
        raises (Reflective::MofError);
    void set_component (in ComponentDef new_value)
        raises (Reflective::MofError);
}; // end of interface PublishesDef

interface EventDefClass : BaseIDL::ValueDefClass
{
    readonly attribute EventDefSet all_of_type_event_def;
    readonly attribute EventDefSet all_of_class_event_def;
    EventDef create_event_def (
        in string identifier,
        in string repository_id,
        in string version,
        in boolean is_abstract,
        in boolean is_custom,
        in boolean is_truncatable)
        raises (Reflective::MofError);
};

interface EventDef : EventDefClass, BaseIDL::ValueDef
{
}; // end of interface EventDef

```

```

struct ProvidesInterfaceLink
{
    BaseIDL::InterfaceDef provides;
    ProvidesDef provides_def;
};
typedef sequence<ProvidesInterfaceLink> ProvidesInterfaceLinkSet;

interface ProvidesInterface : Reflective::RefAssociation
{
    ProvidesInterfaceLinkSet all_provides_interface_links()
    raises (Reflective::MofError);
    boolean exists (
        in BaseIDL::InterfaceDef provides,
        in ProvidesDef provides_def)
    raises (Reflective::MofError);
    BaseIDL::InterfaceDef provides (in ProvidesDef provides_def)
    raises (Reflective::MofError);
    void add (
        in BaseIDL::InterfaceDef provides,
        in ProvidesDef provides_def)
    raises (Reflective::MofError);
    void modify_provides (
        in BaseIDL::InterfaceDef provides,
        in ProvidesDef provides_def,
        in BaseIDL::InterfaceDef new_provides)
    raises (Reflective::NotFound, Reflective::MofError);
    void remove (
        in BaseIDL::InterfaceDef provides,
        in ProvidesDef provides_def)
    raises (Reflective::NotFound, Reflective::MofError);
}; // end of interface ProvidesInterface

struct ComponentFacetLink
{
    ComponentDef component;
    ProvidesDef facet;
};
typedef sequence<ComponentFacetLink> ComponentFacetLinkSet;

interface ComponentFacet : Reflective::RefAssociation
{
    ComponentFacetLinkSet all_component_facet_links()
    raises (Reflective::MofError);
    boolean exists (
        in ComponentDef component,
        in ProvidesDef facet)
    raises (Reflective::MofError);
    ComponentDef component (in ProvidesDef facet)
    raises (Reflective::MofError);
    ProvidesDefSet facet (in ComponentDef component)
    raises (Reflective::MofError);
};

```

```

void add (
    in ComponentDef component,
    in ProvidesDef facet)
    raises (Reflective::MofError);
void modify_component (
    in ComponentDef component,
    in ProvidesDef facet,
    in ComponentDef new_component)
    raises (Reflective::NotFound, Reflective::MofError);
void modify_facet (
    in ComponentDef component,
    in ProvidesDef facet,
    in ProvidesDef new_facet)
    raises (Reflective::NotFound, Reflective::MofError);
void remove (
    in ComponentDef component,
    in ProvidesDef facet)
    raises (Reflective::NotFound, Reflective::MofError);
}; // end of interface ComponentFacet

struct HomeFinderLink
{
    HomeDef home;
    FinderDef finder;
};
typedef sequence<HomeFinderLink> HomeFinderLinkSet;

interface HomeFinder : Reflective::RefAssociation
{
    HomeFinderLinkSet all_home_finder_links()
        raises (Reflective::MofError);
    boolean exists (
        in HomeDef home,
        in FinderDef finder)
        raises (Reflective::MofError);
    HomeDef home (in FinderDef finder)
        raises (Reflective::MofError);
    FinderDefSet finder (in HomeDef home)
        raises (Reflective::MofError);
    void add (
        in HomeDef home,
        in FinderDef finder)
        raises (Reflective::MofError);
    void modify_home (
        in HomeDef home,
        in FinderDef finder,
        in HomeDef new_home)
        raises (Reflective::NotFound, Reflective::MofError);
    void modify_finder (
        in HomeDef home,
        in FinderDef finder,

```

```

    in FinderDef new_finder)
    raises (Reflective::NotFound, Reflective::MofError);
void remove (
    in HomeDef home,
    in FinderDef finder)
    raises (Reflective::NotFound, Reflective::MofError);
}; // end of interface HomeFinder

struct ComponentEmitsLink
{
    ComponentDef component;
    EmitsDef emits;
};
typedef sequence<ComponentEmitsLink> ComponentEmitsLinkSet;

interface ComponentEmits : Reflective::RefAssociation
{
    ComponentEmitsLinkSet all_component_emits_links()
        raises (Reflective::MofError);
    boolean exists (
        in ComponentDef component,
        in EmitsDef emits)
        raises (Reflective::MofError);
    ComponentDef component (in EmitsDef emits)
        raises (Reflective::MofError);
    EmitsDefSet emits (in ComponentDef component)
        raises (Reflective::MofError);
    void add (
        in ComponentDef component,
        in EmitsDef emits)
        raises (Reflective::MofError);
    void modify_component (
        in ComponentDef component,
        in EmitsDef emits,
        in ComponentDef new_component)
        raises (Reflective::NotFound, Reflective::MofError);
    void modify_emits (
        in ComponentDef component,
        in EmitsDef emits,
        in EmitsDef new_emits)
        raises (Reflective::NotFound, Reflective::MofError);
    void remove (
        in ComponentDef component,
        in EmitsDef emits)
        raises (Reflective::NotFound, Reflective::MofError);
}; // end of interface ComponentEmits

struct ComponentConsumesLink
{
    ComponentDef component;
    ConsumesDef consumes;
};

```



```

};
typedef sequence<ComponentConsumesLink> ComponentConsumesLinkSet;

interface ComponentConsumes : Reflective::RefAssociation
{
    ComponentConsumesLinkSet all_component_consumes_links()
        raises (Reflective::MofError);
    boolean exists (
        in ComponentDef component,
        in ConsumesDef consumes)
        raises (Reflective::MofError);
    ComponentDef component (in ConsumesDef consumes)
        raises (Reflective::MofError);
    ConsumesDefSet consumes (in ComponentDef component)
        raises (Reflective::MofError);
    void add (
        in ComponentDef component,
        in ConsumesDef consumes)
        raises (Reflective::MofError);
    void modify_component (
        in ComponentDef component,
        in ConsumesDef consumes,
        in ComponentDef new_component)
        raises (Reflective::NotFound, Reflective::MofError);
    void modify_consumes (
        in ComponentDef component,
        in ConsumesDef consumes,
        in ConsumesDef new_consumes)
        raises (Reflective::NotFound, Reflective::MofError);
    void remove (
        in ComponentDef component,
        in ConsumesDef consumes)
        raises (Reflective::NotFound, Reflective::MofError);
}; // end of interface ComponentConsumes

struct ComponentReceptacleLink
{
    ComponentDef component;
    UsesDef receptacle;
};
typedef sequence<ComponentReceptacleLink> ComponentReceptacleLinkSet;

interface ComponentReceptacle : Reflective::RefAssociation
{
    ComponentReceptacleLinkSet all_component_receptacle_links()
        raises (Reflective::MofError);
    boolean exists (
        in ComponentDef component,
        in UsesDef receptacle)
        raises (Reflective::MofError);
    ComponentDef component (in UsesDef receptacle)

```

```

    raises (Reflective::MofError);
UsesDefSet receptacle (in ComponentDef component)
    raises (Reflective::MofError);
void add (
    in ComponentDef component,
    in UsesDef receptacle)
    raises (Reflective::MofError);
void modify_component (
    in ComponentDef component,
    in UsesDef receptacle,
    in ComponentDef new_component)
    raises (Reflective::NotFound, Reflective::MofError);
void modify_receptacle (
    in ComponentDef component,
    in UsesDef receptacle,
    in UsesDef new_receptacle)
    raises (Reflective::NotFound, Reflective::MofError);
void remove (
    in ComponentDef component,
    in UsesDef receptacle)
    raises (Reflective::NotFound, Reflective::MofError);
}; // end of interface ComponentReceptacle

struct UsesInterfaceLink
{
    BaseIDL::InterfaceDef uses;
    UsesDef uses_def;
};
typedef sequence<UsesInterfaceLink> UsesInterfaceLinkSet;

interface UsesInterface : Reflective::RefAssociation
{
    UsesInterfaceLinkSet all_uses_interface_links()
        raises (Reflective::MofError);
    boolean exists (
        in BaseIDL::InterfaceDef uses,
        in UsesDef uses_def)
        raises (Reflective::MofError);
    BaseIDL::InterfaceDef uses (in UsesDef uses_def)
        raises (Reflective::MofError);
    void add (
        in BaseIDL::InterfaceDef uses,
        in UsesDef uses_def)
        raises (Reflective::MofError);
    void modify_uses (
        in BaseIDL::InterfaceDef uses,
        in UsesDef uses_def,
        in BaseIDL::InterfaceDef new_uses)
        raises (Reflective::NotFound, Reflective::MofError);
    void remove (
        in BaseIDL::InterfaceDef uses,

```

```

    in UsesDef uses_def)
    raises (Reflective::NotFound, Reflective::MofError);
}; // end of interface UsesInterface

struct ComponentHomeLink
{
    ComponentDef component;
    HomeDef home;
};
typedef sequence<ComponentHomeLink> ComponentHomeLinkSet;

interface ComponentHome : Reflective::RefAssociation
{
    ComponentHomeLinkSet all_component_home_links()
    raises (Reflective::MofError);
    boolean exists (
        in ComponentDef component,
        in HomeDef home)
    raises (Reflective::MofError);
    ComponentDef component (in HomeDef home)
    raises (Reflective::MofError);
    void add (
        in ComponentDef component,
        in HomeDef home)
    raises (Reflective::MofError);
    void modify_component (
        in ComponentDef component,
        in HomeDef home,
        in ComponentDef new_component)
    raises (Reflective::NotFound, Reflective::MofError);
    void remove (
        in ComponentDef component,
        in HomeDef home)
    raises (Reflective::NotFound, Reflective::MofError);
}; // end of interface ComponentHome

struct ComponentSupportsLink
{
    BaseIDL::InterfaceDef supports;
    ComponentDef components;
};
typedef sequence<ComponentSupportsLink> ComponentSupportsLinkSet;

interface ComponentSupports : Reflective::RefAssociation
{
    ComponentSupportsLinkSet all_component_supports_links()
    raises (Reflective::MofError);
    boolean exists (
        in BaseIDL::InterfaceDef supports,
        in ComponentDef components)
    raises (Reflective::MofError);
};

```

```

BaselDL::InterfaceDefSet supports (in ComponentDef components)
  raises (Reflective::MofError);
void add (
  in BaselDL::InterfaceDef supports,
  in ComponentDef components)
  raises (Reflective::MofError);
void modify_supports (
  in BaselDL::InterfaceDef supports,
  in ComponentDef components,
  in BaselDL::InterfaceDef new_supports)
  raises (Reflective::NotFound, Reflective::MofError);
void remove (
  in BaselDL::InterfaceDef supports,
  in ComponentDef components)
  raises (Reflective::NotFound, Reflective::MofError);
}; // end of interface ComponentSupports

```

```

struct HomeFactoryLink
{
  HomeDef home;
  FactoryDef factory;
};
typedef sequence<HomeFactoryLink> HomeFactoryLinkSet;

```

```

interface HomeFactory : Reflective::RefAssociation
{
  HomeFactoryLinkSet all_home_factory_links()
    raises (Reflective::MofError);
  boolean exists (
    in HomeDef home,
    in FactoryDef factory)
    raises (Reflective::MofError);
  HomeDef home (in FactoryDef factory)
    raises (Reflective::MofError);
  FactoryDefSet factory (in HomeDef home)
    raises (Reflective::MofError);
  void add (
    in HomeDef home,
    in FactoryDef factory)
    raises (Reflective::MofError);
  void modify_home (
    in HomeDef home,
    in FactoryDef factory,
    in HomeDef new_home)
    raises (Reflective::NotFound, Reflective::MofError);
  void modify_factory (
    in HomeDef home,
    in FactoryDef factory,
    in FactoryDef new_factory)
    raises (Reflective::NotFound, Reflective::MofError);
  void remove (

```

```

    in HomeDef home,
    in FactoryDef factory)
    raises (Reflective::NotFound, Reflective::MofError);
}; // end of interface HomeFactory

struct ComponentPublishesLink
{
    ComponentDef component;
    PublishesDef publishes;
};
typedef sequence<ComponentPublishesLink> ComponentPublishesLinkSet;

interface ComponentPublishes : Reflective::RefAssociation
{
    ComponentPublishesLinkSet all_component_publishes_links()
    raises (Reflective::MofError);
    boolean exists (
        in ComponentDef component,
        in PublishesDef publishes)
    raises (Reflective::MofError);
    ComponentDef component (in PublishesDef publishes)
    raises (Reflective::MofError);
    PublishesDefSet publishes (in ComponentDef component)
    raises (Reflective::MofError);
    void add (
        in ComponentDef component,
        in PublishesDef publishes)
    raises (Reflective::MofError);
    void modify_component (
        in ComponentDef component,
        in PublishesDef publishes,
        in ComponentDef new_component)
    raises (Reflective::NotFound, Reflective::MofError);
    void modify_publishes (
        in ComponentDef component,
        in PublishesDef publishes,
        in PublishesDef new_publishes)
    raises (Reflective::NotFound, Reflective::MofError);
    void remove (
        in ComponentDef component,
        in PublishesDef publishes)
    raises (Reflective::NotFound, Reflective::MofError);
}; // end of interface ComponentPublishes

struct EventTypeLink
{
    EventDef type;
    EventPortDef event;
};
typedef sequence<EventTypeLink> EventTypeLinkSet;

```

```

interface EventType : Reflective::RefAssociation
{
  EventTypeLinkSet all_event_type_links()
    raises (Reflective::MofError);
  boolean exists (
    in EventDef type,
    in EventPortDef event)
    raises (Reflective::MofError);
  EventDef type (in EventPortDef event)
    raises (Reflective::MofError);
  void add (
    in EventDef type,
    in EventPortDef event)
    raises (Reflective::MofError);
  void modify_type (
    in EventDef type,
    in EventPortDef event,
    in EventDef new_type)
    raises (Reflective::NotFound, Reflective::MofError);
  void remove (
    in EventDef type,
    in EventPortDef event)
    raises (Reflective::NotFound, Reflective::MofError);
}; // end of interface EventType

struct APrimaryKeyHomeLink
{
  BaseIDL::ValueDef primary_key;
  HomeDef home;
};
typedef sequence<APrimaryKeyHomeLink> APrimaryKeyHomeLinkSet;

interface APrimaryKeyHome : Reflective::RefAssociation
{
  APrimaryKeyHomeLinkSet all_a_primary_key_home_links()
    raises (Reflective::MofError);
  boolean exists (
    in BaseIDL::ValueDef primary_key,
    in HomeDef home)
    raises (Reflective::MofError);
  BaseIDL::ValueDef primary_key (in HomeDef home)
    raises (Reflective::MofError);
  void add (
    in BaseIDL::ValueDef primary_key,
    in HomeDef home)
    raises (Reflective::MofError);
  void modify_primary_key (
    in BaseIDL::ValueDef primary_key,
    in HomeDef home,
    in BaseIDL::ValueDef new_primary_key)
    raises (Reflective::NotFound, Reflective::MofError);
};

```

```

void remove (
    in BaseIDL::ValueDef primary_key,
    in HomeDef home)
    raises (Reflective::NotFound, Reflective::MofError);
}; // end of interface APrimaryKeyHome

struct ASupportsHomeDefLink
{
    BaseIDL::InterfaceDef supports;
    HomeDef home_def;
};
typedef sequence<ASupportsHomeDefLink> ASupportsHomeDefLinkSet;

interface ASupportsHomeDef : Reflective::RefAssociation
{
    ASupportsHomeDefLinkSet all_a_supports_home_def_links()
        raises (Reflective::MofError);
    boolean exists (
        in BaseIDL::InterfaceDef supports,
        in HomeDef home_def)
        raises (Reflective::MofError);
    BaseIDL::InterfaceDefSet supports (in HomeDef home_def)
        raises (Reflective::MofError);
    void add (
        in BaseIDL::InterfaceDef supports,
        in HomeDef home_def)
        raises (Reflective::MofError);
    void modify_supports (
        in BaseIDL::InterfaceDef supports,
        in HomeDef home_def,
        in BaseIDL::InterfaceDef new_supports)
        raises (Reflective::NotFound, Reflective::MofError);
    void remove (
        in BaseIDL::InterfaceDef supports,
        in HomeDef home_def)
        raises (Reflective::NotFound, Reflective::MofError);
}; // end of interface ASupportsHomeDef

interface ComponentIDLPackageFactory
{
    ComponentIDLPackage create_component_idl_package ()
        raises (Reflective::MofError);
};

interface ComponentIDLPackage : Reflective::RefPackage
{
    readonly attribute ComponentFeatureClass component_feature_ref;
    readonly attribute ComponentDefClass component_def_ref;
    readonly attribute ProvidesDefClass provides_def_ref;
    readonly attribute HomeDefClass home_def_ref;
    readonly attribute FactoryDefClass factory_def_ref;
};

```

```
readonly attribute FinderDefClass finder_def_ref;  
readonly attribute EventPortDefClass event_port_def_ref;  
readonly attribute EmitsDefClass emits_def_ref;  
readonly attribute ConsumesDefClass consumes_def_ref;  
readonly attribute UsesDefClass uses_def_ref;  
readonly attribute PublishesDefClass publishes_def_ref;  
readonly attribute EventDefClass event_def_ref;  
readonly attribute ProvidesInterface provides_interface_ref;  
readonly attribute ComponentFacet component_facet_ref;  
readonly attribute HomeFinder home_finder_ref;  
readonly attribute ComponentEmits component_emits_ref;  
readonly attribute ComponentConsumes component_consumes_ref;  
readonly attribute ComponentReceptacle component_receptacle_ref;  
readonly attribute UsesInterface uses_interface_ref;  
readonly attribute ComponentHome component_home_ref;  
readonly attribute ComponentSupports component_supports_ref;  
readonly attribute HomeFactory home_factory_ref;  
readonly attribute ComponentPublishes component_publishes_ref;  
readonly attribute EventType event_type_ref;  
readonly attribute APrimaryKeyHome a_primary_key_home_ref;  
readonly attribute ASupportsHomeDef a_supports_home_def_ref;  
};  
}; // end of module ComponentIDL
```


13 CIF Metamodel

13.1 CIF Package

In addition to the packages *BaseIDL* and *ComponentIDL*, a new package *CIF* is introduced that contains the metamodel for the Component Implementation Framework. This package obviously depends on the *ComponentIDL* package since its main purpose is to enable the modeling of implementations for components specified using the *ComponentIDL* definitions. This situation is depicted in the following diagram.

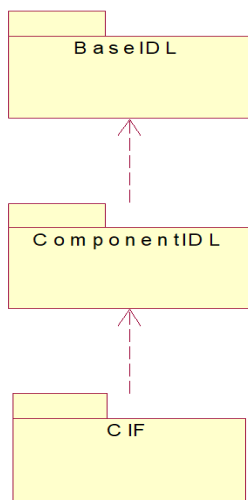


Figure 13.1- Package structure of CCM metamodels

13.2 Classes and Associations

The CIF metamodel defines additional metaclasses and associations. An overview on these is to be seen in Figure 13.2. Their meaning is explained on the following pages.

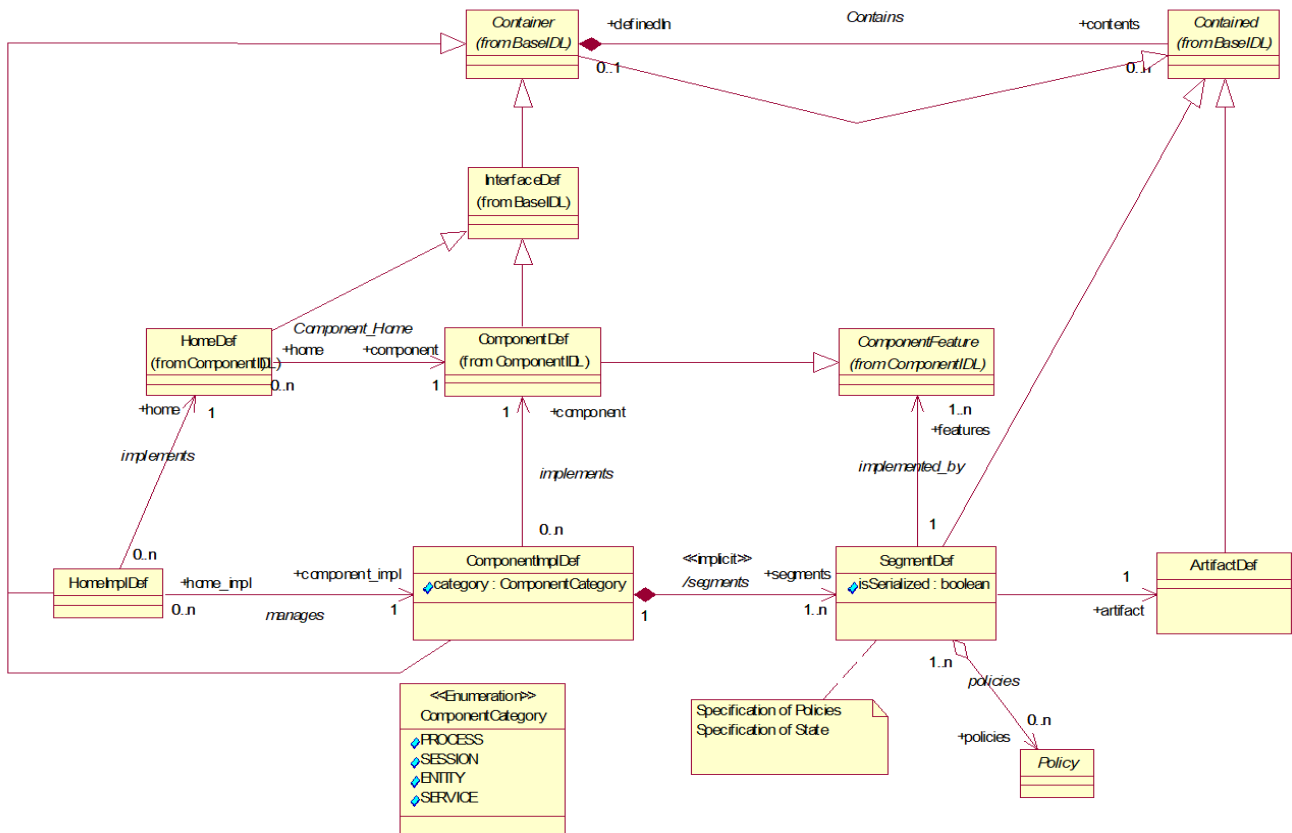


Figure 13.2- CIF metamodel (overview)

13.2.1 ComponentImplDef

This metaclass is used to model an implementation definition for a given component definition. It specifies an association to *ComponentDef* to allow instances to point exactly to the component the instance is going to implement. A *ComponentImplDef* always has exactly one *ComponentDef* associated while each *ComponentDef* might be implemented by different *ComponentImplDefs*.

ComponentImplDef is specified as being a *Container*, by doing so, instances are able to contain other definitions. The only definitions that are allowed to be contained by a *ComponentImplDef* are instances of *SegmentDef*.

Currently, there is no inheritance specification for instances of *ComponentImplDef*.

13.2.2 SegmentDef

Instances of *SegmentDef* are used to model a segmented implementation structure for a component implementation. This means that the behavior for each component feature can be provided by a separate segment of the component implementation (most likely a separate programming language class in the code generated by the CIF tools) if necessary. It is also possible to specify that a segment provides the behavior for a number of component features including the extreme that only one segment implements all component features.

Instances of *SegmentDef* are always contained in instances of *ComponentImplDef* and therefore are derived from *Contained*. *SegmentDef* has an association to *ComponentFeature* so that instances can point to all features of a component which the segment is going to implement. *SegmentDef* has in addition an association to *ArtifactDef*, which are models of programming language constructs (classes) used to actually implement the behavior for component features. There is always exactly one artifact for each segment. However, artefacts may be shared between component implementations whereas segments cannot. That's why the distinction between artifacts and segments has been modeled in the CIF.

The attribute *isSerialized* is used to indicate that the access to segment is required to be serialized or not.

13.2.3 ArtifactDef

ArtifactDef is used to model abstractions from programming language constructs like classes. Instances from *ArtifactDef* in a model represent the elements that provide the behavior for features of a component. Since these can be shared across component implementations the distinctions between artifacts and segments have been made in the metamodel.

ArtifactDef is a specialization of the metaclass *Contained*, which means that artifacts are identifiable and contained in other definitions. The only allowed *Container* for *ArtifactDef* is *ModuleDef*.

13.2.4 Policy

Segment definitions modeled as instances of the metaclass *SegmentDef* may contain a set of policies that have to be applied to realizations of the segment in the implementation code. These policies include for example activation policies for the artifact associated to a segment. The complete set of required policies is not known yet and the metamodel is designed to be flexible. *Policy* is introduced as an abstract metaclass and concrete policies are expected to be defined as specializations of that class. *SegmentDef* aggregates a set of policies.

The metamodel for component implementations is shown in Figure 13.3.

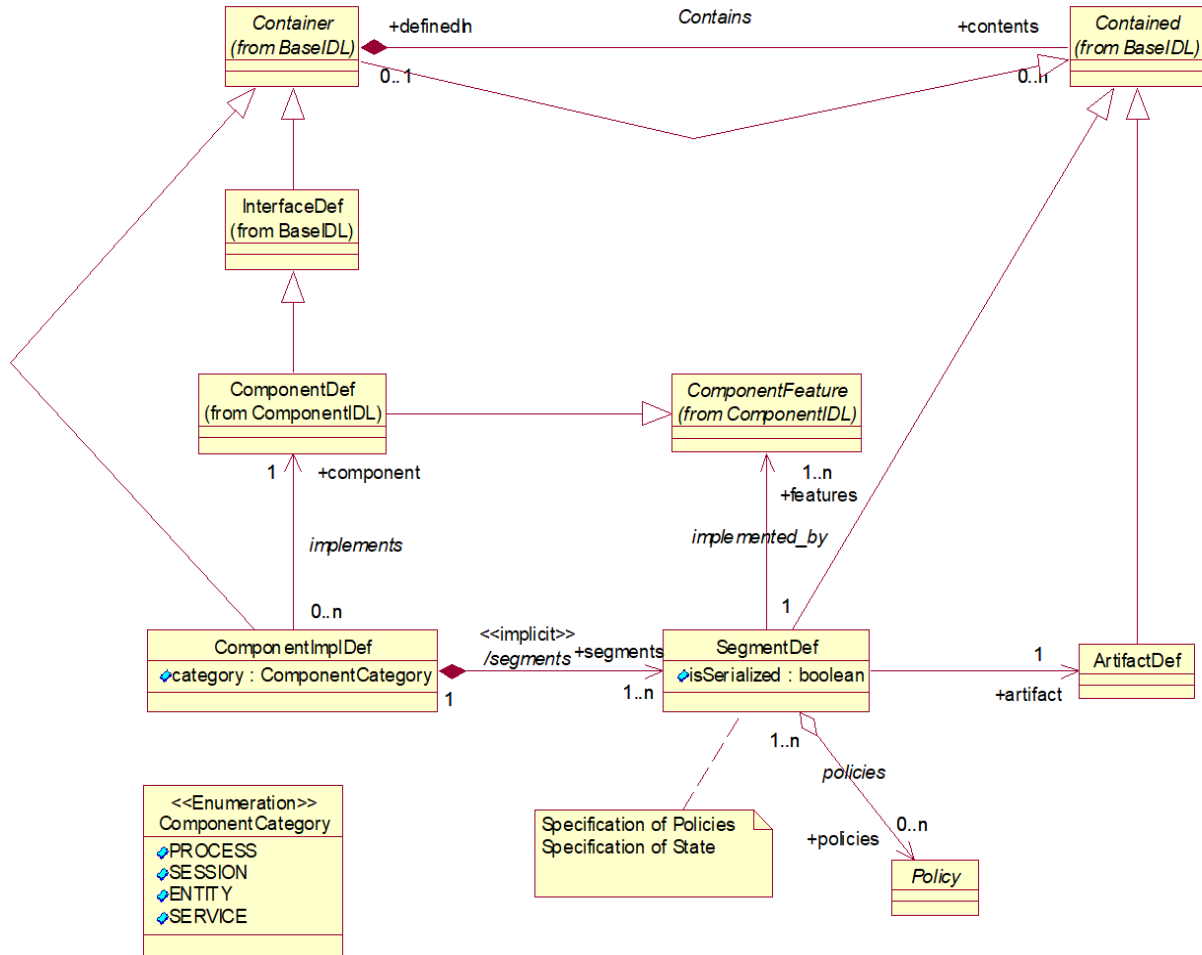


Figure 13.3- Metamodel for component implementations

13.2.5 HomeImplDef

Home interfaces have to be implemented and their implementation is not part of a component implementation. However, for a home implementation it has to be specified, which component implementation the home implementation manages. For that reason, the CIF metamodel contains a metaclass *HomeImplDef*. Each instance of *HomeImplDef* in a model implements exactly one instance of *HomeDef*. This relation is modeled by the association `implements` between both metaclasses. *HomeImplDef* inherits from the abstract metaclass *Container*. This is to allow a home implementation to be identifiable within a model and to contain other model elements. (For the time being, no such elements have been identified). Each home implementation manages exactly one component implementation, this relation is modeled by the association `manages`. It is required, that for each instance *x* of *HomeImplDef* the instance of *ComponentDef*, which is associated to the instance of *HomeDef* associated to *x* is the same instance as the instance of *ComponentDef* associated to the instance of *ComponentImplDef* which is associated to *x*.

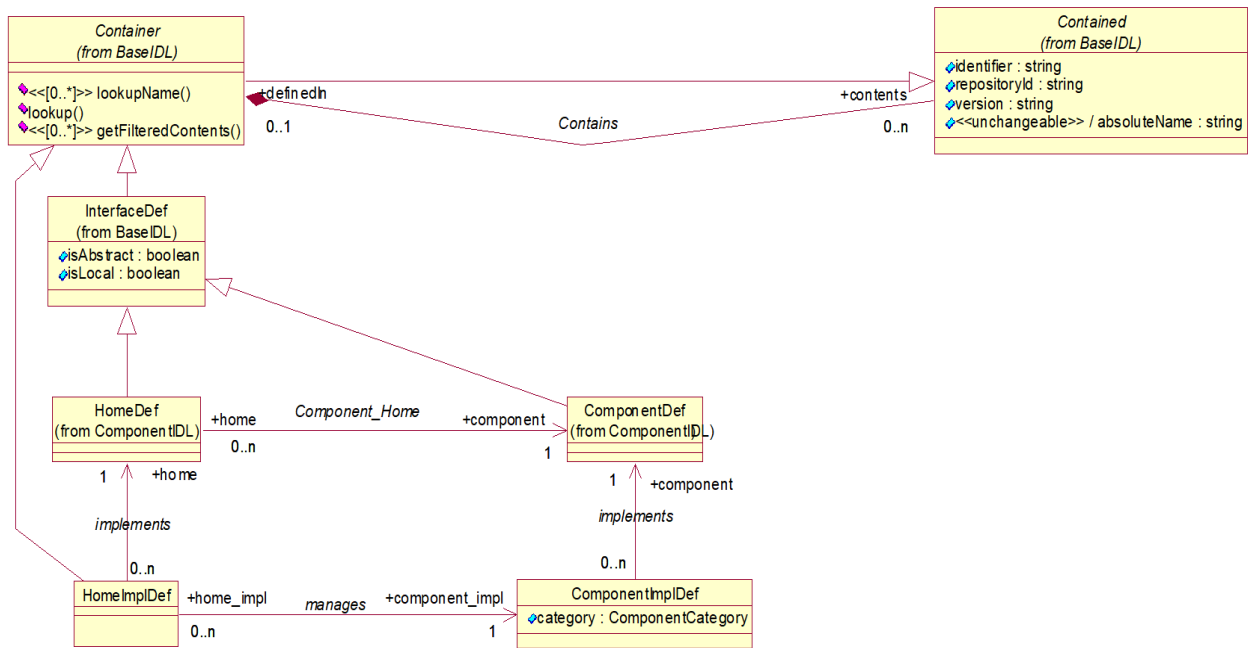


Figure 13.4- Metamodel for home implementations

13.3 Conformance Criteria

This sub clause identifies the conformance points required for compliant implementations of the Component Implementation Framework (CIF) metamodel architecture.

13.3.1 Conformance Points

In the previous sub clause, the MOF metamodel of the Component Implementation Framework (CIF) was defined. The following sub clause defines the XMI format for the exchange of CIF metadata and the IDL for a MOF-compliant CIF repository. Support for the generation and consumption of the XMI metadata and for the MOF-compliant IDL is optional.

13.4 MOF DTDs and IDL for the CIF Metamodel

The XMI DTDs and IDL for the Interface Repository metamodel are presented in this sub clause. The DTDs are generated by applying the MOF-XML mapping defined by the XMI specification to the MOF-compliant metamodel described in “CIF Package” on page 289. The IDL is generated by applying the MOF-IDL mapping defined in the MOF specification to the metamodels and was validated using the IDL compilers.

The IDL requires the inclusion of the reflective interfaces defined in <http://www.omg.org/spec/MOF/2.0> as part of the MOF specification.

13.4.1 XMI DTD

See “XMI DTD” on page 294.

13.4.2 IDL for the CIF Package

```
#pragma prefix "ccm.omg.org"
#include "ComponentIDL.idl"

module CIF
{
  interface ArtifactDefClass;
  interface ArtifactDef;
  typedef sequence<ArtifactDef> ArtifactDefSet;
  interface SegmentDefClass;
  interface SegmentDef;
  typedef sequence<SegmentDef> SegmentDefSet;
  interface ComponentImplDefClass;
  interface ComponentImplDef;
  typedef sequence<ComponentImplDef> ComponentImplDefSet;
  interface PolicyClass;
  interface Policy;
  typedef sequence<Policy> PolicySet;
  interface HomeImplDefClass;
  interface HomeImplDef;
  typedef sequence<HomeImplDef> HomeImplDefSet;
  interface CIFPackage;
  enum ComponentCategory {PROCESS, SESSION, ENTITY, SERVICE};

  interface ArtifactDefClass : BaseIDL::ContainedClass
  {
    readonly attribute ArtifactDefSet all_of_type_artifact_def;
    readonly attribute ArtifactDefSet all_of_class_artifact_def;
    ArtifactDef create_artifact_def (
      in string identifier,
      in string repository_id,
      in string version)
      raises (Reflective::MofError);
  };

  interface ArtifactDef : ArtifactDefClass, BaseIDL::Contained
  {
  }; // end of interface ArtifactDef

  interface SegmentDefClass : BaseIDL::ContainedClass
  {
    readonly attribute SegmentDefSet all_of_type_segment_def;
    readonly attribute SegmentDefSet all_of_class_segment_def;
    SegmentDef create_segment_def (
      in string identifier,
```

```

    in string repository_id,
    in string version,
    in boolean is_serialized)
    raises (Reflective::MofError);
};

```

```

interface SegmentDef : SegmentDefClass, BaseIDL::Contained
{
    ArtifactDef artifact ()
        raises (Reflective::MofError);
    void set_artifact (in ArtifactDef new_value)
        raises (Reflective::MofError);
    ComponentIDL::ComponentFeatureSet features ()
        raises (Reflective::MofError);
    void set_features (in ComponentIDL::ComponentFeatureSet new_value)
        raises (Reflective::MofError);
    void add_features (in ComponentIDL::ComponentFeature new_element)
        raises (Reflective::MofError);
    void modify_features (
        in ComponentIDL::ComponentFeature old_element,
        in ComponentIDL::ComponentFeature new_element)
        raises (Reflective::NotFound, Reflective::MofError);
    void remove_features (in ComponentIDL::ComponentFeature old_element)
        raises (Reflective::NotFound, Reflective::MofError);
    PolicySet policies ()
        raises (Reflective::MofError);
    void set_policies (in PolicySet new_value)
        raises (Reflective::MofError);
    void unset_policies ()
        raises (Reflective::MofError);
    void add_policies (in Policy new_element)
        raises (Reflective::MofError);
    void modify_policies (
        in Policy old_element,
        in Policy new_element)
        raises (Reflective::NotFound, Reflective::MofError);
    void remove_policies (in Policy old_element)
        raises (Reflective::NotFound, Reflective::MofError);
    boolean is_serialized ()
        raises (Reflective::MofError);
    void set_is_serialized (in boolean new_value)
        raises (Reflective::MofError);
}; // end of interface SegmentDef

```

```

interface ComponentImplDefClass : BaseIDL::ContainerClass
{
    readonly attribute ComponentImplDefSet all_of_type_component_impl_def;
    readonly attribute ComponentImplDefSet all_of_class_component_impl_def;
    ComponentImplDef create_component_impl_def (
        in string identifier,
        in string repository_id,

```

```

    in string version,
    in ComponentCategory category)
    raises (Reflective::MofError);
};

interface ComponentImplDef : ComponentImplDefClass, BaseIDL::Container
{
    ComponentIDL::ComponentDef component ()
        raises (Reflective::MofError);
    void set_component (in ComponentIDL::ComponentDef new_value)
        raises (Reflective::MofError);
    SegmentDefSet segments ()
        raises (Reflective::MofError);
    void set_segments (in SegmentDefSet new_value)
        raises (Reflective::MofError);
    void add_segments (in SegmentDef new_element)
        raises (Reflective::MofError);
    void modify_segments (
        in SegmentDef old_element,
        in SegmentDef new_element)
        raises (Reflective::NotFound, Reflective::MofError);
    void remove_segments (in SegmentDef old_element)
        raises (Reflective::NotFound, Reflective::MofError);
    ComponentCategory category ()
        raises (Reflective::MofError);
    void set_category (in ComponentCategory new_value)
        raises (Reflective::MofError);
}; // end of interface ComponentImplDef

interface PolicyClass : Reflective::RefObject
{
    readonly attribute PolicySet all_of_type_policy;
};

interface Policy : PolicyClass
{
}; // end of interface Policy

interface HomeImplDefClass : BaseIDL::ContainerClass
{
    readonly attribute HomeImplDefSet all_of_type_home_impl_def;
    readonly attribute HomeImplDefSet all_of_class_home_impl_def;
    HomeImplDef create_home_impl_def (
        in string identifier,
        in string repository_id,
        in string version)
        raises (Reflective::MofError);
};

interface HomeImplDef : HomeImplDefClass, BaseIDL::Container
{

```



```

ComponentIDL::HomeDef home ()
    raises (Reflective::MofError);
void set_home (in ComponentIDL::HomeDef new_value)
    raises (Reflective::MofError);
ComponentImplDef component_impl ()
    raises (Reflective::MofError);
void set_component_impl (in ComponentImplDef new_value)
    raises (Reflective::MofError);
}; // end of interface HomeImplDef

struct AArtifactSegmentDefLink
{
    ArtifactDef artifact;
    SegmentDef segment_def;
};
typedef sequence<AArtifactSegmentDefLink> AArtifactSegmentDefLinkSet;

interface AArtifactSegmentDef : Reflective::RefAssociation
{
    AArtifactSegmentDefLinkSet all_a_artifact_segment_def_links()
        raises (Reflective::MofError);
    boolean exists (
        in ArtifactDef artifact,
        in SegmentDef segment_def)
        raises (Reflective::MofError);
    ArtifactDef artifact (in SegmentDef segment_def)
        raises (Reflective::MofError);
    void add (
        in ArtifactDef artifact,
        in SegmentDef segment_def)
        raises (Reflective::MofError);
    void modify_artifact (
        in ArtifactDef artifact,
        in SegmentDef segment_def,
        in ArtifactDef new_artifact)
        raises (Reflective::NotFound, Reflective::MofError);
    void remove (
        in ArtifactDef artifact,
        in SegmentDef segment_def)
        raises (Reflective::NotFound, Reflective::MofError);
}; // end of interface AArtifactSegmentDef

struct SegmentsLink
{
    SegmentDef segments;
    ComponentImplDef component_impl_def;
};
typedef sequence<SegmentsLink> SegmentsLinkSet;

interface Segments : Reflective::RefAssociation
{

```

```

SegmentsLinkSet all_segments_links()
    raises (Reflective::MofError);
boolean exists (
    in SegmentDef segments,
    in ComponentImplDef component_impl_def)
    raises (Reflective::MofError);
SegmentDefSet segments (in ComponentImplDef component_impl_def)
    raises (Reflective::MofError);
void add (
    in SegmentDef segments,
    in ComponentImplDef component_impl_def)
    raises (Reflective::MofError);
void modify_segments (
    in SegmentDef segments,
    in ComponentImplDef component_impl_def,
    in SegmentDef new_segments)
    raises (Reflective::NotFound, Reflective::MofError);
void remove (
    in SegmentDef segments,
    in ComponentImplDef component_impl_def)
    raises (Reflective::NotFound, Reflective::MofError);
}; // end of interface Segments

struct ImplementedByLink
{
    ComponentIDL::ComponentFeature features;
    SegmentDef segment_def;
};
typedef sequence<ImplementedByLink> ImplementedByLinkSet;

interface ImplementedBy : Reflective::RefAssociation
{
    ImplementedByLinkSet all_implemented_by_links()
        raises (Reflective::MofError);
boolean exists (
    in ComponentIDL::ComponentFeature features,
    in SegmentDef segment_def)
    raises (Reflective::MofError);
ComponentIDL::ComponentFeatureSet features (in SegmentDef segment_def)
    raises (Reflective::MofError);
void add (
    in ComponentIDL::ComponentFeature features,
    in SegmentDef segment_def)
    raises (Reflective::MofError);
void modify_features (
    in ComponentIDL::ComponentFeature features,
    in SegmentDef segment_def,
    in ComponentIDL::ComponentFeature new_features)
    raises (Reflective::NotFound, Reflective::MofError);
void remove (
    in ComponentIDL::ComponentFeature features,

```

```

    in SegmentDef segment_def)
    raises (Reflective::NotFound, Reflective::MofError);
}; // end of interface ImplementedBy

struct PoliciesLink
{
    Policy policies;
    SegmentDef segment_def;
};
typedef sequence<PoliciesLink> PoliciesLinkSet;

interface Policies : Reflective::RefAssociation
{
    PoliciesLinkSet all_policies_links()
        raises (Reflective::MofError);
    boolean exists (
        in Policy policies,
        in SegmentDef segment_def)
        raises (Reflective::MofError);
    PolicySet policies (in SegmentDef segment_def)
        raises (Reflective::MofError);
    void add (
        in Policy policies,
        in SegmentDef segment_def)
        raises (Reflective::MofError);
    void modify_policies (
        in Policy policies,
        in SegmentDef segment_def,
        in Policy new_policies)
        raises (Reflective::NotFound, Reflective::MofError);
    void remove (
        in Policy policies,
        in SegmentDef segment_def)
        raises (Reflective::NotFound, Reflective::MofError);
}; // end of interface Policies

struct ManagesLink
{
    ComponentImplDef component_impl;
    HomeImplDef home_impl;
};
typedef sequence<ManagesLink> ManagesLinkSet;

interface Manages : Reflective::RefAssociation
{
    ManagesLinkSet all_manages_links()
        raises (Reflective::MofError);
    boolean exists (
        in ComponentImplDef component_impl,
        in HomeImplDef home_impl)
        raises (Reflective::MofError);
};

```

```

ComponentImplDef component_impl (in HomeImplDef home_impl)
  raises (Reflective::MofError);
void add (
  in ComponentImplDef component_impl,
  in HomeImplDef home_impl)
  raises (Reflective::MofError);
void modify_component_impl (
  in ComponentImplDef component_impl,
  in HomeImplDef home_impl,
  in ComponentImplDef new_component_impl)
  raises (Reflective::NotFound, Reflective::MofError);
void remove (
  in ComponentImplDef component_impl,
  in HomeImplDef home_impl)
  raises (Reflective::NotFound, Reflective::MofError);
}; // end of interface Manages

interface CIFPackageFactory
{
  CIFPackage create_cif_package ()
    raises (Reflective::MofError);
};

interface CIFPackage : Reflective::RefPackage
{
  readonly attribute ArtifactDefClass artifact_def_ref;
  readonly attribute SegmentDefClass segment_def_ref;
  readonly attribute ComponentImplDefClass component_impl_def_ref;
  readonly attribute PolicyClass policy_ref;
  readonly attribute HomeImplDefClass home_impl_def_ref;
  readonly attribute AArtifactSegmentDef a_artifact_segment_def_ref;
  readonly attribute Segments segments_ref;
  readonly attribute ImplementedBy implemented_by_ref;
  readonly attribute Policies policies_ref;
  readonly attribute Manages manages_ref;
};
}; // end of module CIF

```

14 Lightweight CCM Profile

14.1 Summary

This Lightweight CCM profile is a conformance point based on the extended model as defined in “Conformance and Compliance” on page 1. The table below defines the specific parts of this CCM specification that are impacted and the normative specific subsetting of CCM. In summary, the following general capabilities (and associated machinery) are excluded from the extended model to define this conformance point:

- Persistence (only session and service components are supported)
- Introspection
- Navigation
- Redundancies, preferring generic over specific.
- Segmentation (not allowed for session or service components)
- Transactions
- Security
- Configurators
- Proxy homes
- Home finders
- CIDL
- POA related mandates

Operations that are termed “disabled” in this conformance point are still part of the associated interface, but implementations may raise either `BAD_OPERATION` or `NO_IMPLEMENT` exceptions when they are invoked. This flexibility avoids the overhead required for the lightweight component skeletons to have explicit tests for all disabled operations. In the specific case of the disabled “**same_component**” operation, on the Navigation interface, throwing an exception of `NO_IMPLEMENT` is required, with a minor code of `IS_LIGHTWEIGHT_COMPONENT`. Clients that wish to clearly determine whether an object has a Lightweight CCM implementation can make this determination by performing a **get_component** on the object, and then perform the **same_component** operation to specifically test for a component meeting this conformance point.

Document impacts below are those that would make the CCM specification document describe this conformance point without reference to the unsupported features and disabled operations. However, the disabled operations are considered present in the interface to preserve the meaning of the interface repository IDs.

The following tables list the specification and document impacts. The normative description column contains all normative issues, without redundancy, and thus those columns contain all the functional “exclusions” to the specification and can be used as a quick reference list of detailed changes. The text in the “Document Impact” column describes the editing changes that would turn the CCM specification into a document describing only the Lightweight CCM profile.

14.2 Changes associated with excluding support for persistence

Normative Exclusion	Document Impact	Comment(s)
Exclude support for primary keys	6.1.4, paragraph 2: remove 6.1.5, paragraph 1: remove reference to associations between instances and keys 6.7, paragraph 3: remove 6.7.1, paragraph 1 and 3: remove references to primary key 6.7.1.2: remove 6.7.2: remove 6.7.4: remove references to primary key 6.7.5.1, paragraph 1 and 2: remove references to primary key 6.7.5.2, paragraph 1 and 2: remove references to primary key 9.1.1: remove	Only components with persistence can have primary keys (Table 7.4, sub clause 7.5.1).
Exclude support for CIDL	6.2.2.1, paragraph 1: remove references to CIDL Clause 7: remove 8.1, paragraph 1: remove references to CIDL 8.2.1: remove 8.2.2: remove	Note excluding CIDL is a consequence of excluding support for <i>both</i> persistence and segmentation.
Exclude support for abstract storage homes	6.7.4: remove references to storage home 8.2.5, paragraph 4: remove references to abstract storage home 8.2.6: remove references to storage home 8.2.7: remove 8.2.8: remove	Table 7-4 in sub clause 7.5.1 disallows such abstract storage homes on session or service components.
Exclude support for abstract storage types	8.2.2, paragraph 2: remove references to abstract storage type 8.2.6: remove references to abstract storage type 8.2.9, point 4: remove references to abstract storage type 8.2.9.2: remove 8.2.11.3: remove references to using the state identifier for persistence. 8.2.11.4: remove references to using the state identifier for persistence.	
Exclude the Implicit Executor Interface for Keyed Homes	8.3.3.6: remove paragraphs about the Implicit Executor Interface for keyed homes.	Keyed homes already excluded.

Remove support for finders in the explicit home executor	8.3.3.6: remove paragraphs about finders 9.3.2.2: remove reference to finders	Table 7-4 in sub clause 7.5.1 disallows such finder operations on session or service components.
Remove support for entity container API types. Only session containers exist.	8.2.11, paragraph 2: remove reference to Entity2Context 8.3.3.3: remove references to entity components 9.1.2: remove references to entity container API types 9.2.9: remove 9.2.10, 9.3: remove references to entity container API type 9.2.13.3, 9.2.13.4, 9.3.4: remove 9.4: remove reference to ComponentId 9.4.1.1: remove reference to persistence 9.4.3: remove 9.5.1.3, paragraph 1: remove references to entity components 9.5.2.3, point 3: remove references to entity components 9.1, 12.2, 12.4: remove Process and Entity component categories from the ComponentType enumeration.	
Remove durable CORBA usage model	9.1.3: remove description of durable usage model 9.1.4: remove table entries associated with durable usage model 9.2.2, 9.2.2.1, 9.2.2.2, 9.2.5: remove table entries and references associated with PERSISTENT object references	
Disable get_persistence operation from CCM2Context interface	9.4.1.1: remove reference to and description of this operation	
Exclude support for persistence	9.2, paragraph 1: remove reference to persistence 9.2.1, paragraph 1: remove reference to persistence 9.2.12: remove reference to persistence	

14.3 Changes associated with excluding support for introspection, navigation and type-specific operations redundant with generic operations

Normative Exclusion	Document Impact	Comment(s)
Disable provide_<name> operation in the component's equivalent interface	6.4.1: remove 6.4.3.2: remove 6.4.2, 6.4.3: replace provide_<name> with provide_facet	The generic version, provide_facet, in the Navigation interface, is retained as it is required for containers for deployment.
Disable get_all_facets, get_named_facets, same_component operations in Navigation interface.	6.4.3.3: remove these operations from the Navigation interface. Also remove the PortDescription, FacetDescription, and FacetDescriptions types. 6.4.3.4: remove 6.4.4: remove	Only the generic provide_facet remains enabled in the Navigation interface, and only for the equivalent interface.
Disable connect_<name>, disconnect_<name>, get_connection_<name> operations, get_connections_<name> for receptacles, in the component's equivalent interface	6.5.1: remove 6.5.2: replace references to type-specific connect and disconnect with generic equivalent. 6.5.2.3: remove	Introspection, or redundant with generic connect/disconnect operations in the Receptacles interface.
Disable get_connections, get_all_receptacles, get_named_receptacles operations in the Receptacles interface	6.5.3: remove these operations from the Receptacles interface. Also remove the ConnectionDescription and ReceptacleDescription types.	Introspection
Disable subscribe_<name>, unsubscribe_<name> operations for event publishers, in the equivalent interface	6.6.5: remove	These are redundant with the generic connect/disconnect operations in the Receptacles interface
Disable connect_<name>, disconnect_<name> operations for event emitters, in the equivalent interface	6.6.6: remove	These are redundant with the generic connect/disconnect operations in the Receptacles interface.
Disable get_consumer_<name> operations in the equivalent interface	6.6.7: remove	These operation are redundant with the generic provide_facet operation in the Navigation interface inherited by the equivalent interface.

Disable <code>get_all_consumers</code> , <code>get_named_consumers</code> , <code>get_all_emitters</code> , <code>get_named_emitters</code> , <code>get_all_publishers</code> , <code>get_named_publishers</code> in the Events interface inherited by the equivalent interface	6.6.8: remove these operations from the Events interface. Also remove the <code>ConsumerDescription</code> , <code>EmitterDescription</code> , <code>SubscriberDescription</code> , and <code>PublisherDescription</code> types.	These operations are for introspection.
Disable <code>get_component_def</code> , <code>get_home_def</code> operations on CCMHome interface	6.7.6: remove references to these operations	Introspection excluded
Disable <code>get_component_def</code> , <code>get_primary_key</code> , <code>get_all_ports</code> operations in the CCMObject interface	6.11.1: remove references to these operations	Introspection excluded

14.4 Changes associated with excluding support for segmentation

Normative Exclusion	Document Impact	Comment(s)
Exclude support for CIDL	Chapter 7: remove 8.2.1: remove	Note excluding CIDL is a consequence of excluding support for <i>both</i> persistence and segmentation. In some sense segmentation is only valid with persistence (see 9.2.11), but this is not the consistent message of the specification.
Exclude composition	8.2.5: remove 8.2.6: remove 8.2.7: remove 8.2.9.1: remove 8.2.11.2: remove 8.2.11.5: remove 9.2.5, 9.2.11: remove reference to segments	Sub clause 8.2.7 is also removed since it deals with persistence
Remove locator support	8.3.1, 8.3.2: remove references to locators, locator-based implementations and the <code>ExecutorLocator</code> interface 8.3.3.3, 8.3.3.4: remove references to main executors and locators 8.3.3.5, paragraph 4 and last paragraph: remove references to locator 8.3.3.6: remove reference to factories returning <code>ExecutorLocators</code>	
Restrict CIF metamodel to a single segment per implementation	12.1, 12.2, 12.4: restrict the multiplicity of the association between <code>ComponentImplDef</code> and <code>SegmentDef</code> to 1:1	

Remove segmentation support	8.2.11, 8.2.9, 9.2.12: remove reference to segmentation	
-----------------------------	---	--

14.5 Changes associated with excluding support for transactions

Normative Exclusion	Document Impact	Comment(s)
Remove support for the “transaction” servant lifetime policy	9.2.5: remove references to the transaction servant lifetime policy.	
Remove support for transaction policies	9.2.6: remove 9.2.8: remove reference to transaction policies 9.2.8.1: remove 9.2.10, 9.2.13.1, 9.2.13.2, 9.3.1, 9.3.3.2: remove reference to transactions	
Disable the <code>get_rollback_only</code> , <code>get_user_transaction</code> , and <code>set_rollback_only</code> operations in the <code>CCMContext</code> interface	9.3.2.1: remove references to these operations	
Disable the <code>UserTransaction</code> interface	9.3.2.3: remove	
Remove support for the <code>SessionSynchronization</code> interface	8.3.1, 9.2.12.2: remove reference to this interface 9.3.3.3: remove	
Exclude support for transactions	9.2, paragraph 1: remove reference to transaction 9.2.1, paragraph 1: remove reference to transaction 9.2.12: remove reference to transaction 9.5.1.1, point 2: remove reference to transaction 9.5.1.4: remove reference to transaction 9.5.2.4: remove reference to transaction	

14.6 Changes associated with excluding support for security

Normative Exclusion	Document Impact	Comment(s)
Remove support for the security policy declarations	9.2.7: remove 9.2.8, 0.3.1: remove reference to security policies 9.2.8.2: remove	
Disable the <code>get_caller_principal</code> and <code>is_caller_in_role</code> operations in the <code>CCMContext</code> interface	9.3.2.1: remove references to these operations	

14.7 Changes associated with excluding support for configurators

Normative Exclusion	Document Impact	Comment(s)
Exclude support for Configurator and StandardConfigurator interfaces	6.10.1: remove	Note that the set_configuration_values operation in the HomeConfiguration interface is retained. Support for Configurators is functionally redundant.
Disable the set_configurator operation in the HomeConfiguration interface	6.10.2.1: remove references to the set_configurator operation (3 places)	Note that the set_configuration_values operation in the HomeConfiguration interface is retained
Remove references to configurators	6.10.2, 6.10.2.1, 6.11.1: remove references to configurators	

14.8 Changes associated with excluding support for proxy homes

Normative Exclusion	Document Impact	Comment(s)
Exclude support for proxy homes	8.2.5 last paragraph: remove 8.2.10: remove 9.4.1.3: remove	
Exclude support for security	9.2, paragraph 1: remove reference to security 9.2.1, paragraph 1: remove reference to security 9.2.12: remove reference to security 9.5.1.1: remove reference to security 9.5.1.5: remove reference to security 9.5.2.5: remove reference to security	

14.9 Changes associated with excluding support for home finders

Normative Exclusion	Document Impact	Comment(s)
Disable support for the HomeFinder interface and the ComponentHomeFinder parameter value to the ORB::resolve_initial_references operation	6.8: remove 9.2.2.1: remove sentence about HomeFinders	Although 9.1.1 implies that homes without primary keys do not support finders, 6.8 implies that they do, thus this table assumes that home finders are not necessarily tied to primary keys and thus we explicitly exclude them in their own right.
Disable the get_home_registration operation in the CCM2Context interface	9.4.1.1: Remove references to and description of this operation 9.4.1.2: remove	

Exclude support for home finders and finder operations	6.7.1, paragraph 2, 6.7.1.1, 6.7.3: remove references to home finders and finder operations 6.7.3.2: remove 6.7.3.3, 6.7.4, paragraph 1, 6.7.5 (heterodox), 8.3.6, point 5, 9.3.2.1 (get_CCM_home), 9.5, 9.5.1, point 4 and last paragraph, 9.5.1.1, last point, 9.5.1.2: remove references to home finders and finder operations 9.5.1.3: remove 9.5.2.3: remove	
--	---	--

14.10 Changes adding additional restrictions to the extended model not represented by exclusions above

Normative restriction	Document Impact	Comment(s)
CIDL is removed because segmentation and persistence is removed, and thus the “API type” does not require specification.	9.1, 9.2.2: remove references to CIDL as means to specifying API type. 9.2.2.3: remove references to CIDL as a means to specify threading policies	Note that the distinction between service components and session components and threading policy is still reflected in the XML descriptors.
Access to the POA is excluded to provide proper containment and enable container implementations that integrate object adapter capabilities.	9.1: remove references to using implementation information to program POA policies. 9.2.4: remove 9.4, 4.4.1.1: remove reference to persistence: remove POA reference	Like all other container services, object adapter services can be satisfied by the standard service (e.g., POA), but, as with the other services, using the standard one is not mandated. Normative usage of the POA is not necessary and precludes simpler and more efficient container implementations.
Exclude self management of components’ references	9.4.2: remove	Such capabilities generally violate the deployment and configuration role of the container.

15 Deployment PSM for CCM

15.1 Introduction

This clause describes the mapping of the platform-independent model for Deployment and Configuration [D+C] to the CORBA Component Model platform.

15.2 Overview

The D&C data models are used in two different ways, first for persistent storage and distribution of information, and second for representing data at runtime. For persistent storage and distribution, the data models are mapped to XML schemas [XSD], so that information can be stored in XML files [XML] according to the model. We frequently use the term (and stereotype) *description* for the classes that define the data model. We use the term “*descriptor*” to refer to the XML file that contains the data. For runtime, the data models are mapped to IDL data structures.

The management classes are runtime entities and mapped to IDL interfaces only.

This sub clause does not include XML schema and IDL files, since both are generated according to rules. However, these files are supplied with this specification to show the results of this rule-based file generation. The rules that will be used to auto-generate these files from the platform independent model use stereotype classes and associations appropriately and then use rules set forth in the UML profile for CORBA.

This clause defines three transformations and two mappings.

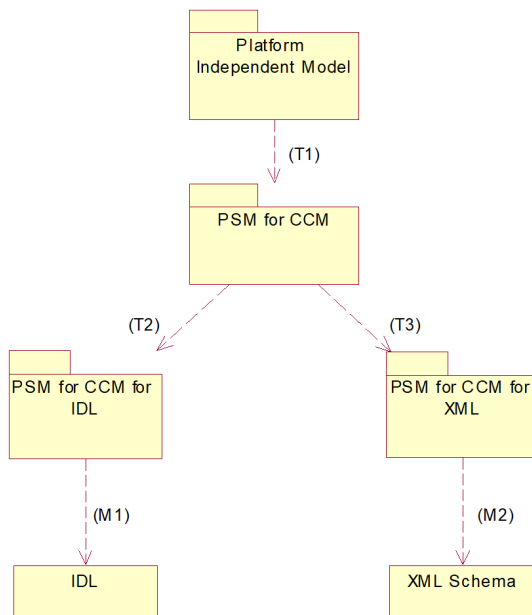


Figure 15.1- Model Transformations and Mappings for CCM

The first transformation, T1 (*PIM to PSM for CCM*), takes the platform-independent model, and refines it into a platform specific model for CCM. In this *PSM for CCM*, the abstract meta-concepts are concretized, and also some other classes are aligned with the CORBA Component Model.

The second transformation T2 (*PSM for CCM to PSM for IDL*) takes the *PSM for CCM* and transforms it into a *PSM for CCM for IDL* that can be used to generate concrete IDL from the model.

The third transformation T3 (*PSM for CCM to PSM for CCM for XML*) creates a *PSM for CCM for XML* that can be used to generate concrete XML schemas.

The motivation for transformations T2 and T3 is to transform the PIM into PSMs so that generic, rule-based mappings M1 and M2 can be used. (Note that some classes have different representations in IDL and XML, for example the **Any** class, prohibiting IDL and XML schema generation from the same model.) The motivation for transformation T1 is that some CCM specific transformations are necessary that are independent of the mapping to IDL or XML.

The M1 mapping is realized using the UML Profile for CORBA [UPC], the M2 mapping is realized using the XML Metadata Interchange (XMI) Version 2 [XMI] specification, “2, XML Schema Production.”

15.3 Definition of Meta-Concepts

This sub clause provides a concrete definition for the classes that are abstract in the PIM. This sub clause is unrelated to the transformations, which will be described in the following sub clauses.

15.3.1 Component

The abstraction of **Component** in the PIM is mapped to both components and homes for the CCM platform. Components in CCM have an interface, attributes, and ports. Homes do not have ports, but an interface and attributes. Both components and homes have explicitly “supported” interfaces in addition to the “equivalent” interface, that inherits all supported interfaces, and includes attributes and explicit operations in the component and home interface definitions.

Viewing homes as a kind of component allows this specification’s model to deploy homes (by themselves or as part of an assembly). Applications or other components in an assembly can then use the home to create component instances at runtime. This supports the full feature set of CCM, without requiring explicit home implementations.

If a CCM home or component supports an interface, their **ComponentInterfaceDescription** has a special port named “supports” that can be used in connections for any of the “supported” interfaces. If, in an assembly, a connection is to be provided by any of the component’s or home’s supported interfaces, then the port name of the **ComponentExternalPortEndpoint** or **SubcomponentPortEndpoint** class is “supports.” For CCM homes, this port also provides their equivalent interface. The “supports” port for CCM components does *not* provide the equivalent interface, since this would be problematic for assembly implementations of components. Home implementations are always monolithic. (Note that in CCM 3.0, assemblies did not allow connections to a component’s equivalent interface either.)

Configuration properties of components, as described by the **ComponentPropertyDescription** class, are attributes in the component or home interface or any inherited component or home interface, but not in any supported interface.

Note – The “supports” magic name has been chosen because it reflects the supported interface. Because it is an IDL keyword, it has little likelihood of conflicting with other port names.

15.3.2 ImplementationArtifact

The meta-concept of **ImplementationArtifact** is mapped to a file accessible by URL. This PSM still treats files as opaque. Agreement between the author of an implementation and the **NodeManager** over the contents of an implementation artifact is assumed. This agreement, or “contract,” is expressed in terms of execution parameters and an implementation’s dependencies on resources provided by the node.

15.3.3 Package

The meta-concept of a package is mapped to a zip file accessible by URI [URI], that includes implementation artifacts and descriptors. Packages have the “.cpk” extension and must contain a single Toplevel Package Descriptor containing a **ToplevelPackageDescription** element with the magic name “package.tpd.”

15.4 PIM to PSM for CCM Transformation

This sub clause defines transformation T1 (as described in the introduction for this clause). It takes the platform-independent model from Clause 6 of the Deployment and Configuration [D+C] specification and aligns classes with the CORBA Component Model. This involves changes to attributes, associations, and semantics of some classes. All classes from the PIM that are not refined here are imported into the PSM for CCM without change.

15.4.1 ComponentInterfaceDescription

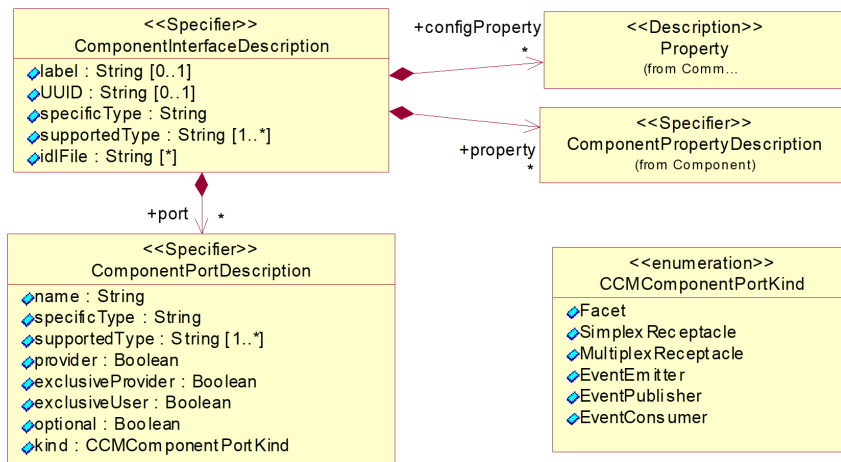


Figure 15.2- ComponentInterfaceDescription and ComponentPortDescription

The **ComponentInterfaceDescription** and **ComponentPortDescription** classes are augmented to support CCM.

The `idlFile` attribute is added to the **ComponentInterfaceDescription**. The `idlFile` attribute, if present, contains alternative URIs that reference an IDL file containing the component’s (or home’s) interface definition. The IDL file is not used within the deployment infrastructure; it may be included in a package for convenience. Since deployable applications have a component interface, some tools that deploy and execute such applications might need the IDL to interact with the ports of the application’s component interface.

The `kind` attribute is added to the **ComponentPortDescription** class and specifies the concrete port kind that is used. This information is required by the **NodeManager** and by assembly tools. In CCM, `EventConsumer` and `Facet` ports are considered providers, the other ports are users.

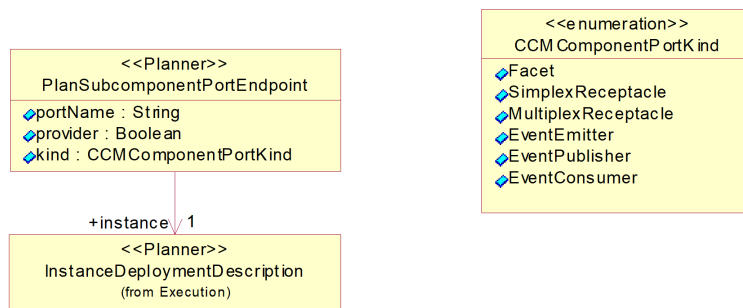
Repository Id strings are used to identify interface types, i.e., for the `specificType` and `supportedType` attributes.

For `Facet` ports, `supportedType` lists the Repository Id of the provided interface and any of its base interfaces that the developer (or tool) chooses to expose for connections. For receptacles, `supportedType` lists the Repository Id of the accepted interface. For `EventEmitter` and `EventPublisher` ports, `supportedType` lists the Repository Id of the accepted consumer interface. For `EventConsumer` ports, `supportedType` lists the Repository Id of the consumer interface and any of its base interfaces that the developer (or tool) chooses to expose for connections.

If the component or home supports one or more interfaces, this will be reflected by a **ComponentPortDescription** element of kind `Facet` with the magic name “`supports.`” The `specificType` attribute is left empty, the `supportedType` attribute lists the Repository Id of any of its supported interfaces and base interfaces that the developer wants to expose for connections.

Initially, a **ComponentInterfaceDescription** can be generated from a component’s or home’s IDL description with a defined set of configuration properties (from attributes) and default values for the `exclusiveProvider`, `exclusiveUser`, and `optional` attributes. If desired, a user can then adjust these three attributes for each port and also add configuration property default values to the **ComponentInterfaceDescription** by adding **Property** elements to the `configProperties` list.

15.4.2 PlanSubcomponentPortEndpoint



The `kind` attribute augments the `provider` attribute in the **PlanSubcomponentPortEndpoint** class and specifies the concrete port kind that is used. This information is required by the various managers in the Execution Management Model. The `provider` attribute still indicates a port that provides an object reference.

15.4.3 Application

The **start** operation on the **Application** class performs the `configuration_complete` operation in all component instances that are part of the application.

15.4.4 RepositoryManager

When artifact files are included in the package (as opposed to referenced via URL outside the package), the **RepositoryManager** must make its own copy of these artifacts during the `installPackage` operation. It must substitute a URL that references this copy of the artifact in the `location` attribute of **ImplementationArtifactDescription** elements delivered via its interface.

15.4.5 SatisfierProperty

This PSM has to define concrete types that are implied on the value of a **SatisfierProperty** by the **SatisfierPropertyKind**, and on the value of the **Property** that is matched against the satisfier.

- For the `Quantity` kind, the value of the **SatisfierProperty** is of type `unsigned long`. The value of the **Property** is ignored.
- For the `Capacity` kind, the value of the **SatisfierProperty** is of type `unsigned long` or `double`. The value of the **Property** must be of the same type.
- For the `Maximum` and `Minimum` kinds, the value of the **SatisfierProperty** is of type `long` or `double`. The value of the **Property** must be of the same type.
- For the `Attribute` kind, the value of the **SatisfierProperty** is of type `long`, `double`, `string`, or an enumeration type. In the case of `long`, `double`, or `string`, the value of the **Property** must be of the same type. If the value of the **SatisfierProperty** is of enumeration type, the value of the **Property** is of type `string`, containing the enumeration value that must compare equal to the **SatisfierProperty** value.
- For the `Selection` kind, the value of the **SatisfierProperty** is a sequence of type `long`, `double`, `string`, or an enumeration type. The same rules as for the `Attribute` kind apply.

15.5 PSM for CCM to PSM for CCM for IDL Transformation

This sub clause defines transformation T2 (as described in the introduction). It transforms the *PSM for CCM* into a *PSM for CCM for IDL* that can be used to generate concrete IDL using a rule-based mapping. Classes from the *PSM for CCM* are transformed to match the UML Profile for CORBA. Its rules are then used to generate concrete IDL.

The first sub clause describes generic mapping rules that are applied to all classes that are part of the *PSM for CCM*. The second sub clause defines special transformation rules for the classes that are abstract in the PIM.

All classes in the *PSM for CCM for IDL* are placed in the **Deployment** package, so that all resulting IDL structures and interfaces will be part of the `Deployment` IDL module.

15.5.1 Generic Transformation Rules

The mapping to IDL is accomplished using the rules set forth in the UML Profile for CORBA. To apply these rules, the stereotypes used in the platform-independent model are mapped to stereotypes for which a mapping is defined in the profile. The **«Description»** stereotype and all that inherit from it are mapped to the **«CORBAStruct»** stereotype; these classes are therefore mapped to CORBA structures. The **«Exception»** stereotype is mapped to the **«CORBAException»** stereotype; such classes become CORBA exceptions. The **«Enumeration»** stereotype is mapped to the **«CORBAEnum»** stereotype in order to become enum types in IDL. The **«Manager»** stereotype is mapped to the **«CORBAInterface»** stereotype so that these classes become CORBA interfaces.

To avoid redundancy and circular graphs, non-composite associations between classes with a common owner are expressed by an ordinal attribute at the source (navigating) end, with the name of the attribute being the role name plus the suffix “Ref,” and the type “unsigned long.” The value of this attribute is the index of the target element in its container, with the index of the first element being 0 (zero). To enable the usage of an index, the composition of the target element in its container is qualified with the “ordered” constraint.

Wherever the multiplicity of an attribute, parameter or return value is not exactly one (but 0..1, 1..* or *), a new class is introduced to represent a sequence of the type of the attribute, parameter, or return value. The sequence class has the «**CORBASequence**» stereotype, and its name is the english plural of the name of the type. The sequence class has a composition association with the element class that is navigable from the sequence to the element. The composition is qualified with the index of the sequence. The attribute, parameter or return value is then replaced with an attribute, parameter or return value, respectively, with the same name as before, but with the type being the newly introduced sequence class and the exactly one (1..1) multiplicity.

A similar rule is applied to all navigable association or composition ends whose multiplicity is not exactly one (but 0..1, 1..* or *): a new class is introduced to represent a sequence of the class at the navigable end; this sequence class is defined as describe above. The original association or composition end is then replaced with a navigable association or composition end, with the same role name as before, at the new sequence class, with a multiplicity of exactly one (1..1). According to the rules in the UML Profile for CORBA, these associations and compositions will then map to a structure member in IDL, its type being a named sequence of the referenced type.

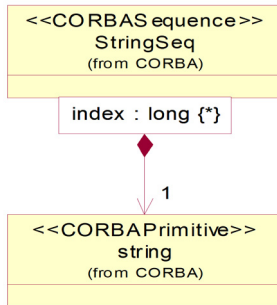
Excepted from the two rules above are attributes, parameters, return values or navigable association or composition ends where the type is **String**, `unsigned long` or **Endpoint**. Instead of defining new sequence types, the existing types in the CORBA package are being used; see below. Note that in combination, these rules map non-composite associations between classes with a common owner and a multiplicity other than 1 to sequence of “unsigned long” type.

Another exception from the rule above are attributes of type **String** with the 0..1 (zero or one) multiplicity. In this case, the multiplicity is updated to 1..1 (exactly one). If the value is missing in an XML representation of the model, the empty string is used as default value.

The inheritance relationships of classes with the «**Description**» stereotype (**SharedResource**, **Resource**, and **Capability**) classes are removed; all attributes and associations of the base class are attached to the derived class. Associations of classes with the «**Manager**» stereotype are removed from the *PSM for CCM for IDL*.

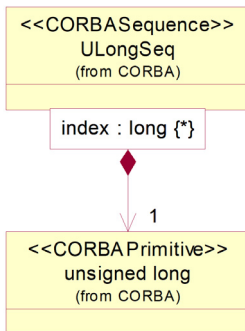
15.5.2 Special Transformation Rules

15.5.2.1 Sequence of String



A type representing a sequence of strings already exists in the **CORBA** package and can be re-used. Wherever the **String** type is used with a multiplicity other than exactly one, it is mapped to the **StringSeq** class from the CORBA package as shown above. It then maps to the `CORBA::StringSeq` type in IDL (from the `orb.idl` file).

15.5.2.2 Sequence of unsigned long



A type representing a sequence of the `unsigned long` type already exists in the **CORBA** package and can be re-used. Wherever the `unsigned long` type is used with a multiplicity other than exactly one, it is mapped to the **ULongSeq** class from the CORBA package as shown above. It then maps to the `CORBA::ULongSeq` type in IDL (from the `orb.idl` file). Sequences of the unsigned long type occur when a non-composite association between classes with a common owner with a multiplicity other than one occurs, according to the generic rule above.

15.5.2.3 Endpoint



The abstract **Endpoint** class is mapped to the **Object** class from the **CORBA** package. It will therefore map to the `Object` type in IDL.

15.5.2.4 DataType

```
<<CORBAPrimitive>>  
    typecode  
    (from CORBAProfile)
```

The abstract **DataType** class is mapped to the **typecode** class from the **CORBAProfile** package. It then maps to the `TypeCode` type in IDL.

15.5.2.5 Any

```
<<CORBAPrimitive>>  
    any  
    (from CORBAProfile)
```

The abstract **Any** class is mapped to the **any** class from the **CORBAProfile** package. It will then map to the `any` type in IDL.

15.5.2.6 Primitive Types

The UML data types **String**, **Integer** and **Boolean** are mapped to the classes **string**, **long** and **boolean** in the **CORBAProfile** package, respectively. They will then map to the `string`, `long`, and `boolean` types in IDL, respectively.

15.5.3 Mapping to IDL

After applying the transformations defined in this sub clause, IDL is generated by applying the rules set forth in the UML Profile for CORBA specification [UPC].

15.6 PSM for CCM to PSM for CCM for XML Transformation

This sub clause defines transformation T3 (as described in the introduction). It transforms the *PSM for CCM* into a *PSM for CCM for XML* that can be used to generate a concrete XML schema using the mapping rules described in “2, XML Schema Production” of the XML Metadata Interchange (XMI) Version 2 [XMI] specification.

15.6.1 Generic Transformation Rules

Data model elements, annotated with the «**Description**» or «**enumeration**» stereotype (or a stereotype that inherits from it), are used to generate an XML schema for representing metadata in XML documents for distribution, interchange or persistence. The only normative use of such XML-based metadata in this specification is for installing component packages using the **RepositoryManager**'s `installPackage` operation.

Management model elements, annotated with the «**Manager**» or «**Exception**» stereotype, are not part of the PSM for CCM for XML, they are mapped to IDL only.

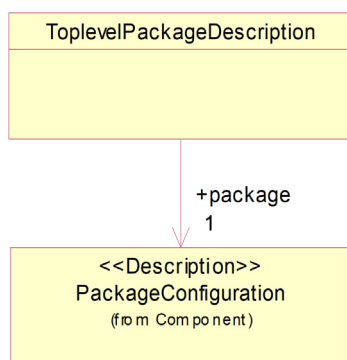
All classes in the PSM for CCM for XML are annotated with the “org.omg.xmi.contentType” tag set to the value “complex.”

All attributes are annotated with the “org.omg.xmi.element” tag set to “true.”

All packages are annotated with the “org.omg.xmi.nsURI” tag set to “http://www.omg.org/Deployment” and the “org.omg.xmi.nsPrefix” tag set to the value “Deployment.”

15.6.2 Special Transformation Rules

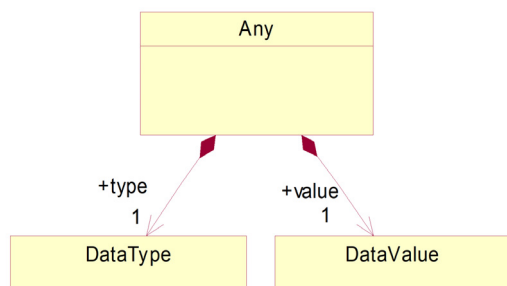
15.6.2.1 ToplevelPackageDescription



The **ToplevelPackageDescription** is introduced to point to the **PackageConfiguration** element for the top-level component package in a package.

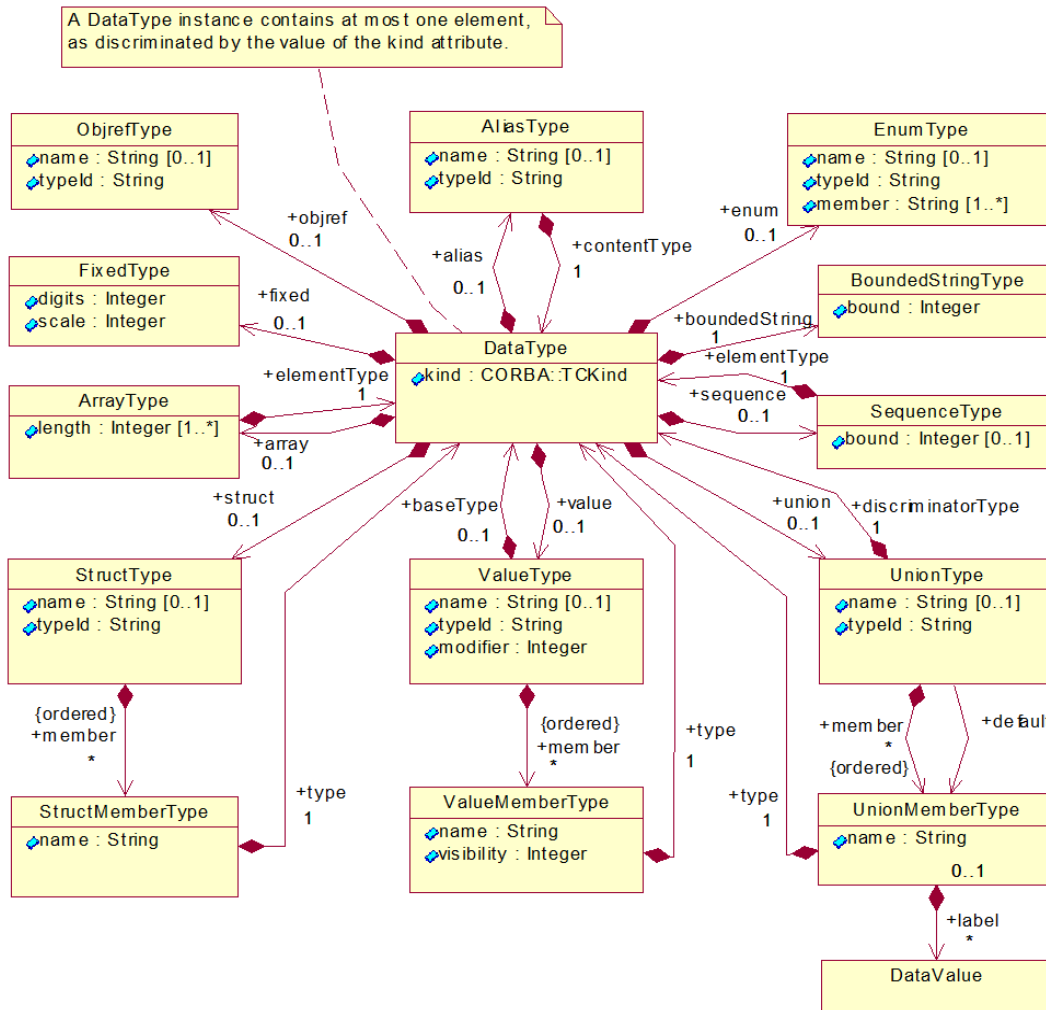
The motivation for this element is that a package may include component packages for sub-components. A selection mechanism is necessary to distinguish the top-level component package. This is accomplished by including a single Toplevel Package Descriptor with the magic name “package.tpd” into the package.

15.6.2.2 Any



An **Any** instance describes a typed value. It is mapped to a class that contains a **DataType** and a **DataValue**, which are elaborated below.

15.6.2.3 DataType



A **DataType** instance describes a type. It is mapped to a hierarchical structure as shown above, describing available types in IDL.

The **DataType** class contains a `kind` field that indicates the IDL type described by a **DataType** instance. The `kind` is of the enumeration type `CORBA::TCKind`, as defined in sub clause 4.11 of the CORBA specification.

If the `kind` is `tk_null`, `tk_void`, `tk_short`, `tk_long`, `tk_ushort`, `tk_ulong`, `tk_float`, `tk_double`, `tk_boolean`, `tk_char`, `tk_octet`, `tk_any`, `tk_TypeCode`, `tk_longlong`, `tk_ulonglong`, `tk_longdouble` or `tk_wchar`, the **DataType** element does not contain any other elements.

If the `kind` is `tk_string` or `tk_wstring`, then the **DataType** may optionally contain a **BoundedStringType** element indicating the upper bound for the string length. If the **DataType** does not contain a **BoundedStringType**, an unbounded string is assumed.

If the kind is `tk_objref`, `tk_component` or, `tk_home`, then the **DataType** may optionally contain an **ObjrefType** element describing the object reference's type (using its Repository Id). If the **DataType** does not contain an **ObjrefType** element, then an untyped object reference (Repository Id "IDL:omg.org/CORBA/Object:1.0") is assumed.

If the kind is `tk_struct` or `tk_except`, then the **DataType** contains a **StructType** element, which in turn describes a list of struct members.

If the kind is `tk_union`, then the **DataType** contains a **UnionType** element. **UnionType** contains the type of the discriminator and a number of typed elements, one of which may be the default member. Each member may be identified with multiple case labels. No label is associated with the default member.

If the kind is `tk_enum`, then the **DataType** contains an **EnumType** element describing the enumeration values.

If the kind is `tk_sequence`, then the **DataType** contains a **SequenceType** element. Its optional `bound` attribute indicates the sequence's upper bound. If the `bound` attribute is absent, the sequence is unbounded.

If the kind is `tk_array`, then the **DataType** contains an **ArrayType** element. Its `length` attribute indicates the length of the array. For multi-dimensional arrays, the multiplicity of the length attribute is greater than one, and the most significant dimension is listed first ("left to right" in IDL).

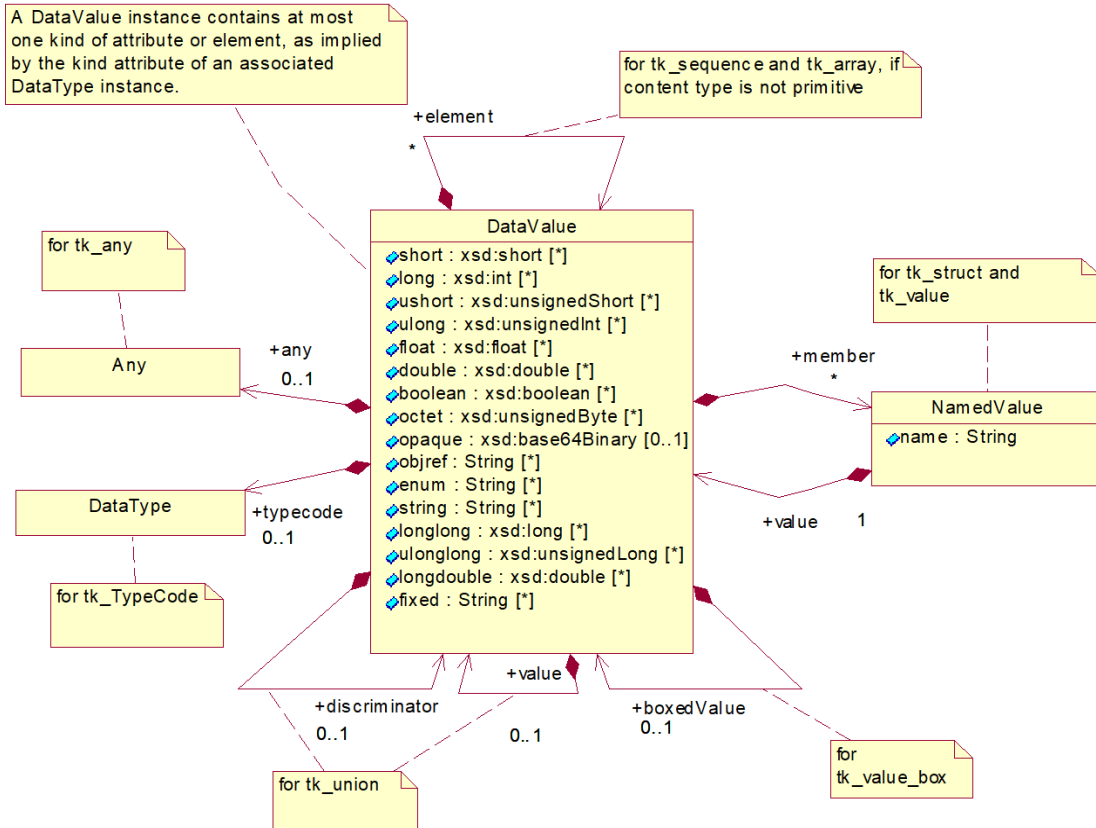
If the kind is `tk_alias` or `tk_value_box`, then the **DataType** contains an **AliasType** element.

If the kind is `tk_fixed`, then the **DataType** contains a **FixedType** element.

If the kind is `tk_value`, then the **DataType** contains a **ValueType** element. **ValueType** contains the type code of the concrete base type, if any, a type modifier (with values as defined by `CORBA::ValueModifier`), and a number of members. Each member has a name, type and visibility (with values as defined by `CORBA::Visibility`).

In **StructType**, **ValueType** and **EnumType**, the `name` attribute contains the name of the struct, valuetype or enum IDL type, and the `typeId` attribute contains its Repository Id.

15.6.2.4 DataValue



The **DataValue** class describes a value. It is mapped to a hierarchical structure as above, fully describing a value that can be described by an IDL type. A **DataValue** cannot exist by itself, it needs a matching **DataType** to describe its structure (see the **Any** class).

If the type's kind is `tk_null` or `tk_void`, **DataValue** is empty.

If the type's kind is `tk_short`, `tk_long`, `tk_ushort`, `tk_ulong`, `tk_float`, `tk_double`, `tk_boolean`, `tk_octet`, `tk_string`, `tk_longlong`, `tk_ulonglong`, or `tk_longdouble`, **DataValue** contains a single short, long, ushort, ulong, float, double, boolean, octet, string, longlong, ulonglong, or longdouble attribute, respectively. If the type's kind is `tk_wstring`, then **DataValue** also contains a string element.

If the type's kind is `tk_char` or `tk_wchar`, the **DataValue** contains a string attribute containing a string of length 1.

If the type's kind is `tk_enum`, the **DataValue** contains the enumeration value in the `enum` attribute.

If the type's kind is `tk_objref`, `tk_component` or `tk_home`, the **DataValue** contains a stringified object reference in the `objref` attribute.

If the type's kind is `tk_fixed`, the **DataValue** contains a `fixed` attribute holding a fixed-point decimal literal.

If the type's kind is `tk_sequence` or `tk_array`, and the sequence's or array's element type is equivalent to (i.e., not considering aliased types) `tk_short`, `tk_long`, `tk_ushort`, `tk_ulong`, `tk_float`, `tk_double`, `tk_boolean`, `tk_octet`, `tk_objref`, `tk_enum`, `tk_string`, `tk_longlong`, `tk_ulonglong`, `tk_longdouble`, `tk_wstring`, `tk_fixed`, `tk_component`, or `tk_home`, then the respective attribute has a multiplicity equal to the length of the sequence or array. In the case of multi-dimensional arrays, the least significant dimension is enumerated first.

If the type's kind is `tk_sequence` or `tk_array`, and the sequence's or array's element type is equivalent to `tk_char` or `tk_wchar`, then the **DataValue** contains a single `string` attribute. Each character in this string is used as an element of the sequence or array.

If the type's kind is `tk_sequence` or `tk_array`, and the sequence's or array's element type is equivalent to `tk_octet`, then the **DataValue** contains a single `opaque` attribute.

If the type's kind is `tk_sequence` or `tk_array`, and the sequence's or array's element is not of the types enumerated above, then the **DataValue** contains the elements of the sequence or array as **DataType** elements, using the `element` association.

If the type's kind is `tk_TypeCode` or `tk_any`, the **DataValue** contains a **DataType** or **Any** element, respectively.

If the type's kind is `tk_struct` or `tk_value`, the **DataValue** contains a **NamedValue** for each member of the structure or valuetype.

If the type's kind is `tk_union`, the **DataValue** contains a single **DataValue** as the union's discriminator, and zero or one **DataValue** elements, using the `value` association, as the member of the union.

If the type's kind is `tk_value_box`, the **DataValue** contains zero or one **DataValue** elements using the `boxedValue` association. If the `boxedValue` element is missing, a null value is implied.

15.6.2.5 Others

The **PackageConfiguration**, **DomainUpdateKind**, **Connection**, and **Endpoint** classes are used by the runtime models only and are not part of the PSM for XML.

15.6.3 Transformation Exceptions and Extensions

Metadata for a component package is usually spread out across several XML files, which are called descriptors. Certain associations are expected to be expressed by links within the same document, others are expected to link across documents. XMI takes care of both patterns by way of "proxies," which do not contain nested elements but a link attribute (either "`href`" or "`xlink:href`") referencing the target element by URI. A schema produced using the XMI rules for schema production allows proxies to appear anywhere.

Composition associations in UML express that the class at the composite end (the containing class) owns and contains the class at the part end (the contained class). It is typical, in XML documents, for instances of contained classes to be embedded within the instance of the containing class. However, it is also possible to store contained instances by themselves in a separate file by using a proxy (using "`href`" or "`xlink:href`") to reference the contained instance in a separate file. Since the multiplicity on the composite end of a composite association is always one to one in this specification, contained instances can only have a single such proxy reference.

For non-composite associations between classes with a common owner (composite end of composition), the definition of the class at the source end of the association must contain a proxy linking to the element at the target end of the association. The definition of the class at the source end cannot contain the definition of the element at the target end, because it is owned by the common owner, and its identity cannot be duplicated.

Non-composite associations between classes that do not have a common owner are usually expressed by the element defining the class at the source end containing a proxy that links to an external document. Direct containment is possible but may result in duplicated data.

While tools can decide to either combine information into a single XML document or to place information into arbitrary files, using XMI proxies to link to that information, it is expected that some model elements usually appear as the outermost document element of a standalone XML file. These commonly used descriptors are assigned descriptive terms and standard file extensions.

- A Component Package Descriptor contains a **ComponentPackageDescription** document element; it has the “.cpd” file extension.
- A Component Implementation Descriptor contains a **ComponentImplementationDescription** document element; it has the “.cid” file extension.
- An Implementation Artifact Descriptor contains an **ImplementationArtifactDescription** document element; it has the “.iad” file extension.
- A Component Interface Descriptor contains a **ComponentInterfaceDescription** document element; it has the “.ccd” (CORBA Component Descriptor) file extension.
- A Domain Descriptor contains a **Domain** document element; it has the “.cdd” (Component Domain Descriptor) file extension.
- A Deployment Plan Descriptor contains a **DeploymentPlan** document element; it has the “.cdp” (Component Deployment Plan) file extension.
- A Package Configuration Descriptor contains a **PackageConfiguration** document element; it has the “.pcd” file extension.
- A Toplevel Package Descriptor contains a **ToplevelPackageDescription** document element; it has the “package.tpd” file name.
- Package files use the “.cpk” extension.

Spreading information across files according to these patterns allow better reuse, for example by extracting an implementation from a package.

Proxies follow the linking semantics specified by XMI [XMI]. If a URI reference [URI] does not contain a fragment identifier (the “#id_value” part), then the target of the reference is the outermost document element of an descriptor file.

15.6.4 Interpretation of Relative References

URI references [URI] are used by proxies and appear in the `location` attribute of the **ImplementationArtifactDescription** and **ArtifactDeploymentDescription** classes and the `idlFile` attribute of the **ComponentInterfaceDescription** class.

XML documents that are part of a Component Package can use relative-path references (i.e., URIs that do not begin with a scheme name or a slash character) to reference documents and other artifacts within the same package.

The interpretation of relative URIs that are not relative-path references (i.e., network-path references that start with two slash characters, or absolute-path references that start with a single slash character), the interpretation of relative-path references that reference documents outside the package (by way of “. . .” path segments), and the interpretation of relative-path references in documents that are not contained in a Component Package (e.g., a Deployment Plan Descriptor) is implementation-specific. (Note: this allows XML processors to supply arbitrary Base URIs that do not necessarily relate to any file system but that must expose the Component Package’s hierarchical structure.)

15.6.5 Mapping to XML

After applying the transformations defined in this sub clause, an XML schema is generated by applying the rules set forth in the XML Metadata Interchange specification, “2, XML Schema Production.” [XMI]

15.7 Miscellaneous

15.7.1 Entry Points

This specification defines the interaction between an implementation artifact and the node manager as implementation-dependent, in order to not restrict the forms that an implementation artifact might have – executable files, loadable libraries, source files or scripts, for example.

However, to ensure source code compatibility in the common case without restricting implementation choice, entry points are defined here if the language is C++ and the implementation artifact is a shared library, or if the language is Java and the implementation artifact is a class file. In these two cases, there must be a specific execution parameter associated with the Monolithic Implementation Description.

If the instance to be deployed is a component, then the name of the execution parameter shall be “component factory.” The parameter is of type **String**, and its name is the name of an entry point that has no parameters and that returns a pointer of type `Components::EnterpriseComponent`.

If the instance to be deployed is a home, then the name of the execution parameter shall be “home factory.” The parameter is of type **String**, and its name is the name of an entry point that has no parameters and that returns a pointer of type `Components::HomeExecutorBase`.

For backwards compatibility, it is recommended that the name of the entry point should be the name of the component or home, prefixed with “create_” (e.g., “create_Account” for an Account component).

If the language is C++, then the entry points shall be qualified as ‘extern “C”’.

These well-defined entry points ensure that the user code for the entry point does not need to be changed when building components for different target environments. These definitions do not enable interoperability between containers and DLLs (even assuming the same compiler and ORB), thus additional interfaces are still required that are specific to container implementations. This implies that, as in CCM 3.0, component and home implementation DLLs are specific to the container implementation (and the code generation tools). Since there was and is no normative interoperability interfaces within a node, thus further implies that there is no vendor segmentation boundary within a node at all.

15.7.2 Homes

Note that this specification does not depend on the existence of homes; using the entry points defined above, a container is able to create component instances directly, without the need of creating a home first, and then using it as a factory for the component instance.

This is no loss in comparison to the Packaging and Deployment clause of CCM in CORBA 3.0. If a component instance is to be deployed as part of an assembly, the container has no way of providing a user-defined home with any parameters, and is therefore limited to keyless homes. However, a factory operation for the component instance as defined above can do its job as well as the parameter-challenged `create` operation that is part of a keyless home.

In contrast to the Packaging and Deployment clause, this specification recognizes homes as instances that can be deployed, and therefore enables the full range of home features.

15.7.3 Valuetype Factories

If an **ImplementationArtifact** contains valuetype factories, then its list of execution parameters shall include an element with the name “valuetype factories” and of type **ValuetypeFactoryList**, which is defined as

```
module Deployment {
    struct ValuetypeFactory {
        string repid;
        string valueentrypoint;
        string factoryentrypoint;
    };
    typedef sequence<ValuetypeFactory>
        ValuetypeFactoryList;
};
```

Each element of that sequence describes a valuetype factory that needs to be registered with the ORB in order to demarshal user-defined valuetypes. The `repid` field specifies the Repository Id of the valuetype created by the valuetype factory. The `factoryentrypoint` field specifies the name of an entry point that can be used to create an instance of the valuetype factory. If `valueentrypoint` is not the empty string, it specifies an entry point that can be used to create an instance of the valuetype.

If the language is C++, then the entry points shall be qualified as ‘extern “C”’.

15.7.4 Discovery and Initialization

The **ExecutionManager** must be able to find the **NodeManager** instances for all nodes in the **Domain**, so that it is able to deploy applications according to deployment plans that are based on the current contents of the Target Data Model. This is accomplished using the Naming Service.

- The user of the deployment system creates a naming context for a domain. Note that a naming context is expressible by a URL representation (e.g., a “`corbaname:`” reference).
- Implementations of the **ExecutionManager** interface must accept the address of the naming context as a configuration parameter, and use it to publish its own reference with the name “`ExecutionManager`” and the empty string as the id in that context.

- Implementations of the **TargetManager** interface must accept the address of the naming context as a configuration parameter, and use it to publish its own reference with the name “TargetManager” and the empty string as the id in that context.
- Implementations of the **NodeManager** interface must accept the address of the naming context as a configuration parameter, and use it to publish their own reference with the node’s name as the name and the id “NodeManager.” The node’s name must match the name attribute of the node in the Target Data Model.

Upon startup, the **ExecutionManager** finds the **TargetManager** in the Naming Service, and accesses the current **Domain** information. Based on the **Node** elements that are contained in the **Domain**, the **ExecutionManager** then calls the `joinDomain` operation of each **NodeManager**.

An **ExecutionManager** may offer functionality to “add” new nodes to the domain, or to remove nodes from the domain. In that case, the **ExecutionManager** looks up a **NodeManager** with a user-provided name in the Naming Service and then calls its `joinDomain` or `leaveDomain` operation, respectively. In addition, an **ExecutionManager** may offer to scan the Naming Service context for previously unregistered nodes, calling the `joinDomain` operation on each associated **NodeManager**.

Note that there is no direct relationship between domains and repositories. Therefore, implementations of the **RepositoryManager** interface are not registered in the Naming Service.

15.7.5 Location

URI references [URI] are handled by the **RepositoryManager** and **NodeManager** interfaces: the **RepositoryManager** receives URLs to packages as a parameter to the `installPackage` operation and must generate URLs pointing to itself in **ImplementationArtifactDescription** elements. The **NodeManager** receives URLs as attributes of the **ArtifactDeploymentDescription** elements that are part of the **DeploymentPlan**.

Both **RepositoryManager** and **NodeManager** shall be able to interpret URLs according to the `http` scheme. Additional schemes may optionally be supported.

Note – This requires **RepositoryManager** implementations to include both an `http` server and an `http` client [HTTP]. **NodeManager** need to implement `http` clients only, in order to download implementation artifacts from the repository.

The **RepositoryManager** must supply a “`http`” URI as part of the location attribute in the **ImplementationArtifactDescription** elements. A **RepositoryManager** may optionally include other alternative locations to provide **NodeManager** implementations with a choice of transports to use for downloading artifacts.

15.7.6 Segmentation

This specification obsoletes CCM’s idea of component segmentation. In the original CCM, assemblies provided just a single level of decomposition. Segments then offered a second level to split the implementation of a component into several independent pieces of code. This specification allows composition and decomposition on any level, and therefore the ability to add another level of decomposition on the lowest level is redundant. However, no parts of this specification inhibit a component author from using this feature of the CCM Implementation Framework.

15.8 Migration Issues

This sub clause deals with the issues of migrating from the Packaging and Deployment model that existed in version 3.0 of the CCM specification to the deployment model defined in this specification.

Note – The Packaging and Deployment clause of CCM 3.0, in its Component Deployment sub clause, defined interfaces that are involved in the deployment of components onto nodes. Similar interfaces might be useful in implementing the NodeManager, however, this specification does not prescribe any such node-level interfaces.

15.8.1 Component Implementations

The portable parts of CCM component implementation source code remains untouched. The generated code to enable interactions with the containers may change, requiring recompilation and linking. The non-portable hand written code in some implementations which was written assuming a particular container implementation would likely have to change — similar to porting the component to a different CCM system.

15.8.2 Component and Assembly Packages and Metadata

The metadata is changed to be based on XML schemas, and the basic models are different. Many lower level elements are not different, and it is expected that meta-data transformation (forward migration) will be able to be automated in the common cases where all the features used are supported.

This specification is kept simple in anticipation that broad (and necessarily complex) software packaging and distribution standards do not exist, and the W3C OSD specification (by Microsoft and Marinba in 1997) referenced by the original CCM specification did not become a standard. Future RFPs may want to consider mappings from such comprehensive standards into this simpler model that focuses on CCM applications.

The component data model stays within the scope of deployment and configuration and does not bring forward all the metadata aspects in the previous CCM specification that were not relevant to deployment and configuration. Furthermore, much of the metadata for informing containers of the requirements of component instances was not defined as part of an intervendedor boundary. Thus this specification assumes the use of two “private” channels of information between the development tools (and code generation) and the runtime environment (NodeManager). These are the resource requirements of the MonolithicImplementationDescription and the execParameters of the ImplementationArtifactDescription. The authors believe standardizing this metadata should be part of a true effort at vendor segmentation between CCM development tools and CCM runtime environments (assuming the same compiler and ORB), which does not exist and was not the mandate of this specification.

Beyond the necessity of validating configuration and connection among components, the one other metadata interoperability issue is to standardize the vocabulary for selection criteria, which is interoperation between users and implementers of component software. This is currently deferred due to the concurrence of the other specification for this language with this specification (see below).

15.8.3 Component Deployment Systems

Deployment systems need to be changed to support this specification. Most aspects of container implementations should be reusable.

15.9 Metadata Vocabulary

15.9.1 Implementation Selection Requirements

Selection requirements, part of both the **PackageConfiguration** and **SubcomponentInstantiationDescription** classes, express requirements that are meant to drive the selection among alternative implementations. The user of an implementation (creator of a package configuration or an assembly) is requesting services to be satisfied by a component implementation. The mechanism defined in this specification requires agreement of the vocabulary of these services on both sides, but there is no interoperable vocabulary defined. The currently active specification entitled “UML Profile for Modeling Quality of Service and Fault Tolerance Characteristics and Mechanisms” resulted in, among other things, “a Definition of Individual QoS Characteristics,” which provides an appropriate vocabulary to drive this mechanism.

When this QoS-driven vocabulary is connected to the CCM PSM, some other component metadata requirements, such as “humanlanguage” may also be added to the selection criteria language.

15.9.2 Monolithic Implementation Resource Requirements

As mentioned above, this vocabulary is a private communication channel between development tools and the NodeManager, since no other interoperability boundary exists between these two. Obviously some standardization could be easily done, based on previous CCM-defined metadata such as container supported persistence, transactions, and POA policies. If this limited scoping is not accepted by the Task Force, data model classes containing this type of information can easily be added to support both a defined resource vocabulary and even a separate container-services vocabulary for information that would never be part of a “resource finding” matching process with the target nodes, but needs to be conveyed to the runtime environment for component instances.

16 Deployment IDL for CCM

16.1 Overview

This clause describes interfaces and meta-data to facilitate the deployment of CORBA component based applications, and a process of generating concrete IDL for the interfaces and meta-data, by using the rules defined by the UML Profile for CORBA [UPC] specification on the Deployment PSM for CCM, which is derived from the PIM in the Deployment and Configuration [D+C] specification. With these rules, this clause contains the normative definition.

This clause contains IDL that has been produced by using these rules. It is non-normative, so in the case of discrepancies, this clause is relevant.

```
#include <orb.idl>
```

```
module Deployment {
```

```
    enum SatisfierPropertyKind {
```

```
        Quantity,  
        Capacity,  
        Minimum,  
        Maximum,  
        _Attribute,  
        Selection
```

```
    };
```

```
    struct SatisfierProperty {
```

```
        string name;  
        SatisfierPropertyKind kind;  
        boolean dynamic;  
        any value;
```

```
    };
```

```
    typedef sequence < SatisfierProperty > SatisfierProperties;
```

```
    struct SharedResource {
```

```
        string name;  
        ::CORBA::StringSeq resourceType;  
        ::CORBA::ULongSeq nodeRef;  
        SatisfierProperties property;
```

```
    };
```

```
    typedef sequence < SharedResource > SharedResources;
```

```
    struct Resource {
```

```
        string name;  
        ::CORBA::StringSeq resourceType;  
        SatisfierProperties property;
```

```
    };
```

```

typedef sequence < Resource > Resources;

struct Node {
    string name;
    string label;
    ::CORBA::ULongSeq sharedResourceRef;
    ::CORBA::ULongSeq connectionRef;
    Resources resource;
};

typedef sequence < Node > Nodes;

struct Interconnect {
    string name;
    string label;
    ::CORBA::ULongSeq connectionRef;
    ::CORBA::ULongSeq connectRef;
    Resources resource;
};

typedef sequence < Interconnect > Interconnects;

struct Bridge {
    string name;
    string label;
    ::CORBA::ULongSeq connectRef;
    Resources resource;
};

typedef sequence < Bridge > Bridges;

struct Property {
    string name;
    any value;
};

typedef sequence < Property > Properties;

struct Domain {
    string UUID;
    string label;
    SharedResources sharedResource;
    Nodes node;
    Interconnects interconnect;
    Bridges bridge;
    Properties infoProperty;
};

struct ResourceAllocation {
    string elementName;
    string resourceName;

```

```

    Properties property;
};

typedef sequence < ResourceAllocation > ResourceAllocations;

exception ResourceCommitmentFailure {
    string reason;
    long index;
    string propertyName;
    ::CORBA::AnySeq propertyValue;
};

interface ResourceCommitmentManager {
    void commitResources (in ResourceAllocations resources)
        raises (ResourceCommitmentFailure);
    void releaseResources (in ResourceAllocations resources)
        raises (ResourceCommitmentFailure);
};

enum DomainUpdateKind {
    Add,
    Delete,
    UpdateAll,
    UpdateDynamic
};

interface TargetManager {
    Domain getAllResources ();
    Domain getAvailableResources ();
    ResourceCommitmentManager commitResources (in ResourceAllocations resources)
        raises (ResourceCommitmentFailure);
    void destroyResourceCommitment (in ResourceCommitmentManager manager);
    void updateDomain (in ::CORBA::StringSeq elements, in Domain domainSubset, in DomainUpdateKind
updateKind);
};

typedef sequence < Object > Endpoints;

struct Connection {
    string name;
    Endpoints endpoint;
};

typedef sequence < Connection > Connections;

exception StartError {
    string name;
    string reason;
};

exception InvalidConnection {

```

```

    string name;
    string reason;
};

interface Application {
    void finishLaunch (in Connections providedReference, in boolean start)
        raises (StartError, InvalidConnection);
    void start ()
        raises (StartError);
};

exception ResourceNotAvailable {
    string name;
    string resourceType;
    string propertyName;
    string elementName;
    string resourceName;
};

exception InvalidProperty {
    string name;
    string reason;
};

exception InvalidNodeExecParameter {
    string name;
    string reason;
};

exception InvalidComponentExecParameter {
    string name;
    string reason;
};

exception StopError {
    string name;
    string reason;
};

interface ApplicationManager {
    Application startLaunch (in Properties configProperty, out Connections providedReference)
        raises (ResourceNotAvailable, StartError, InvalidProperty, InvalidNodeExecParameter, InvalidComponentExecParameter);
    void destroyApplication (in Application app)
        raises (StopError);
};

typedef sequence < Application > Applications;

enum CCMComponentPortKind {
    Facet,

```

```
SimplexReceptacle,  
MultiplexReceptacle,  
EventEmitter,  
EventPublisher,  
EventConsumer  
};
```

```
struct ComponentPortDescription {  
    string name;  
    string specificType;  
    ::CORBA::StringSeq supportedType;  
    boolean provider;  
    boolean exclusiveProvider;  
    boolean exclusiveUser;  
    boolean optional;  
    CCMComponentPortKind kind;  
};
```

```
typedef sequence < ComponentPortDescription > ComponentPortDescriptions;
```

```
struct ComponentPropertyDescription {  
    string name;  
    CORBA::TypeCode type;  
};
```

```
typedef sequence < ComponentPropertyDescription > ComponentPropertyDescriptions;
```

```
struct ComponentInterfaceDescription {  
    string label;  
    string UUID;  
    string specificType;  
    ::CORBA::StringSeq supportedType;  
    ::CORBA::StringSeq idlFile;  
    Properties configProperty;  
    ComponentPortDescriptions port;  
    ComponentPropertyDescriptions property;  
    Properties infoProperty;  
};
```

```
struct Requirement {  
    string name;  
    string resourceType;  
    Properties property;  
};
```

```
typedef sequence < Requirement > Requirements;
```

```
struct MonolithicDeploymentDescription {  
    string name;  
    ::CORBA::StringSeq source;  
    ::CORBA::ULongSeq artifactRef;
```

```

    Properties execParameter;
    Requirements deployRequirement;
};

typedef sequence < MonolithicDeploymentDescription > MonolithicDeploymentDescriptions;

enum ResourceUsageKind {
    None,
    InstanceUsesResource,
    ResourceUsesInstance,
    PortUsesResource,
    ResourceUsesPort
};

struct InstanceResourceDeploymentDescription {
    ResourceUsageKind resourceUsage;
    string requirementName;
    string resourceName;
    Properties property;
};

typedef sequence < InstanceResourceDeploymentDescription > InstanceResourceDeploymentDescriptions;

struct InstanceDeploymentDescription {
    string name;
    string node;
    ::CORBA::StringSeq source;
    unsigned long implementationRef;
    Properties configProperty;
    InstanceResourceDeploymentDescriptions deployedResource;
    InstanceResourceDeploymentDescriptions deployedSharedResource;
};

typedef sequence < InstanceDeploymentDescription > InstanceDeploymentDescriptions;

struct ComponentExternalPortEndpoint {
    string portName;
};

typedef sequence < ComponentExternalPortEndpoint > ComponentExternalPortEndpoints;

struct PlanSubcomponentPortEndpoint {
    string portName;
    boolean provider;
    CCMComponentPortKind kind;
    unsigned long instanceRef;
};

typedef sequence < PlanSubcomponentPortEndpoint > PlanSubcomponentPortEndpoints;

```

```

struct ExternalReferenceEndpoint {
    string location;
    boolean provider;
    string portName;
    ::CORBA::StringSeq supportedType;
};

typedef sequence < ExternalReferenceEndpoint > ExternalReferenceEndpoints;

struct ConnectionResourceDeploymentDescription {
    string targetName;
    string requirementName;
    string resourceName;
    Properties property;
};

typedef sequence < ConnectionResourceDeploymentDescription > ConnectionResourceDeploymentDe-
scriptions;

struct PlanConnectionDescription {
    string name;
    ::CORBA::StringSeq source;
    Requirements deployRequirement;
    ComponentExternalPortEndpoints externalEndpoint;
    PlanSubcomponentPortEndpoints internalEndpoint;
    ExternalReferenceEndpoints externalReference;
    ConnectionResourceDeploymentDescriptions deployedResource;
};

typedef sequence < PlanConnectionDescription > PlanConnectionDescriptions;

struct PlanSubcomponentPropertyReference {
    string propertyName;
    unsigned long instanceRef;
};

typedef sequence < PlanSubcomponentPropertyReference > PlanSubcomponentPropertyReferences;

struct PlanPropertyMapping {
    string name;
    ::CORBA::StringSeq source;
    string externalName;
    PlanSubcomponentPropertyReferences delegatesTo;
};

typedef sequence < PlanPropertyMapping > PlanPropertyMappings;

struct ImplementationDependency {
    string requiredType;
};

```

```

typedef sequence < ImplementationDependency > ImplementationDependencies;

struct ResourceDeploymentDescription {
    string requirementName;
    string resourceName;
    Properties property;
};

typedef sequence < ResourceDeploymentDescription > ResourceDeploymentDescriptions;

struct ArtifactDeploymentDescription {
    string name;
    ::CORBA::StringSeq location;
    string node;
    ::CORBA::StringSeq source;
    Properties execParameter;
    Requirements deployRequirement;
    ResourceDeploymentDescriptions deployedResource;
};

typedef sequence < ArtifactDeploymentDescription > ArtifactDeploymentDescriptions;

enum PlanLocalityKind {
    SameProcess,
    DifferentProcess,
    NoConstraint
};

struct PlanLocality {
    PlanLocalityKind constraint;
    ::CORBA::ULongSeq constrainedInstanceRef;
};

typedef sequence < PlanLocality > PlanLocalities;

struct DeploymentPlan {
    string label;
    string UUID;
    ComponentInterfaceDescription realizes;
    MonolithicDeploymentDescriptions implementation;
    InstanceDeploymentDescriptions instance;
    PlanConnectionDescriptions connection;
    PlanPropertyMappings externalProperty;
    ImplementationDependencies dependsOn;
    ArtifactDeploymentDescriptions artifact;
    Properties infoProperty;
    PlanLocalities localityConstraint;
};

interface DomainApplicationManager :
    ApplicationManager

```



```

{
    Applications getApplications ();
    DeploymentPlan getPlan ();
};

exception PlanError {
    string name;
    string reason;
};

typedef sequence < DomainApplicationManager > DomainApplicationManagers;

interface ExecutionManager {
    DomainApplicationManager preparePlan (in DeploymentPlan plan, in ResourceCommitmentManager
resourceCommitment)
    raises (ResourceNotAvailable, PlanError, StartError);
    DomainApplicationManagers getManagers ();
    void destroyManager (in DomainApplicationManager manager)
    raises (StopError);
};

interface Logger {
};

interface NodeApplicationManager :
    ApplicationManager
{
};

interface NodeManager {
    void joinDomain (in Domain theDomain, in TargetManager manager, in Logger log, in long updateInter-
val);
    void leaveDomain ();
    NodeApplicationManager preparePlan (in DeploymentPlan plan, in ResourceCommitmentManager
resourceCommitment)
    raises (StartError, PlanError);
    void destroyManager (in NodeApplicationManager appManager)
    raises (StopError);
    Resources getDynamicResources ();
};

interface NodeApplication :
    Application
{
};

interface DomainApplication :
    Application
{
};

```

```

exception NameExists {
};

exception PackageError {
    string source;
    string reason;
};

exception NoSuchName {
};

exception LastConfiguration {
};

exception InvalidReference {
};

struct ComponentPackageDescription;

typedef sequence < ComponentPackageDescription > ComponentPackageDescriptions;

struct PackageConfiguration;

typedef sequence < PackageConfiguration > PackageConfigurations;

struct ComponentPackageReference {
    string requiredUUID;
    string requiredName;
    ComponentInterfaceDescription requiredType;
};

typedef sequence < ComponentPackageReference > ComponentPackageReferences;

struct ComponentPackageImport {
    ::CORBA::StringSeq location;
};

typedef sequence < ComponentPackageImport > ComponentPackageImports;

struct SubcomponentInstantiationDescription {
    string name;
    ComponentPackageDescriptions basePackage;
    PackageConfigurations specializedConfig;
    Requirements selectRequirement;
    Properties configProperty;
    ComponentPackageReferences referencedPackage;
    ComponentPackageImports importedPackage;
};

typedef sequence < SubcomponentInstantiationDescription > SubcomponentInstantiationDescriptions;

```

```

struct SubcomponentPortEndpoint {
    string portName;
    unsigned long instanceRef;
};

typedef sequence < SubcomponentPortEndpoint > SubcomponentPortEndpoints;

struct AssemblyConnectionDescription {
    string name;
    Requirements deployRequirement;
    ComponentExternalPortEndpoints externalEndpoint;
    SubcomponentPortEndpoints internalEndpoint;
    ExternalReferenceEndpoints externalReference;
};

typedef sequence < AssemblyConnectionDescription > AssemblyConnectionDescriptions;

struct SubcomponentPropertyReference {
    string propertyName;
    unsigned long instanceRef;
};

typedef sequence < SubcomponentPropertyReference > SubcomponentPropertyReferences;

struct AssemblyPropertyMapping {
    string name;
    string externalName;
    SubcomponentPropertyReferences delegatesTo;
};

typedef sequence < AssemblyPropertyMapping > AssemblyPropertyMappings;

enum LocalityKind {
    SameNodeAnyProcess,
    SameNodeSameProcess,
    SameNodeDifferentProcess,
    DifferentNode,
    DifferentProcess,
    NoConstraint
};

struct Locality {
    LocalityKind constraint;
    ::CORBA::ULongSeq constrainedInstanceRef;
};

typedef sequence < Locality > Localities;

struct ComponentAssemblyDescription {
    SubcomponentInstantiationDescriptions instance;
    AssemblyConnectionDescriptions connection;
};

```

```

    AssemblyPropertyMappings externalProperty;
    Localities localityConstraint;
};

typedef sequence < ComponentAssemblyDescription > ComponentAssemblyDescriptions;

struct NamedImplementationArtifact;

typedef sequence < NamedImplementationArtifact > NamedImplementationArtifacts;

struct ImplementationArtifactDescription {
    string label;
    string UUID;
    ::CORBA::StringSeq location;
    Properties execParameter;
    Requirements deployRequirement;
    NamedImplementationArtifacts dependsOn;
    Properties infoProperty;
};

struct NamedImplementationArtifact {
    string name;
    ImplementationArtifactDescription referencedArtifact;
};

typedef sequence < ResourceUsageKind > ResourceUsageKinds;

struct ImplementationRequirement {
    ResourceUsageKinds resourceUsage;
    string resourcePort;
    string componentPort;
    string name;
    string resourceType;
    Properties property;
};

typedef sequence < ImplementationRequirement > ImplementationRequirements;

struct MonolithicImplementationDescription {
    Properties nodeExecParameter;
    NamedImplementationArtifacts primaryArtifact;
    ImplementationRequirements deployRequirement;
    Properties componentExecParameter;
};

typedef sequence < MonolithicImplementationDescription > MonolithicImplementationDescriptions;

struct Capability {
    string name;
    ::CORBA::StringSeq resourceType;
    SatisfierProperties property;
};

```

```

};

typedef sequence < Capability > Capabilities;

struct ComponentImplementationDescription {
    string label;
    string UUID;
    ComponentInterfaceDescription implements;
    ComponentAssemblyDescriptions assemblyImpl;
    MonolithicImplementationDescriptions monolithicImpl;
    Properties configProperty;
    Capabilities capability;
    ImplementationDependencies dependsOn;
    Properties infoProperty;
};

struct PackagedComponentImplementation {
    string name;
    ComponentImplementationDescription referencedImplementation;
};

typedef sequence < PackagedComponentImplementation > PackagedComponentImplementations;

struct ComponentPackageDescription {
    string label;
    string UUID;
    ComponentInterfaceDescription realizes;
    Properties configProperty;
    PackagedComponentImplementations implementation;
    Properties infoProperty;
};

struct PackageConfiguration {
    string label;
    string UUID;
    ComponentPackageDescriptions basePackage;
    PackageConfigurations specializedConfig;
    Requirements selectRequirement;
    Properties configProperty;
    ComponentPackageReferences referencedPackage;
    ComponentPackageImports importedPackage;
};

interface RepositoryManager {
    void installPackage (in string installationName, in string location, in boolean replace)
        raises (NameExists, PackageError);
    void createPackage (in string installationName, in PackageConfiguration package, in string baseLocation, in boolean replace)
        raises (NameExists, PackageError);
    PackageConfiguration findPackageByName (in string name)
        raises (NoSuchName);
};

```

```

PackageConfiguration findPackageByUUID (in string UUID)
  raises (NoSuchName);
::CORBA::StringSeq findNamesByType (in string type);
::CORBA::StringSeq getAllNames ();
::CORBA::StringSeq getAllTypes ();
void deletePackage (in string installationName)
  raises (NoSuchName);
};

struct ComponentUsageDescription {
  ComponentPackageDescriptions basePackage;
  PackageConfigurations specializedConfig;
  Requirements selectRequirement;
  Properties configProperty;
  ComponentPackageReferences referencedPackage;
  ComponentPackageImports importedPackage;
};

struct RequirementSatisfier {
  string name;
  ::CORBA::StringSeq resourceType;
  SatisfierProperties property;
};
};

```

17 XML Schema for CCM

17.1 Introduction

This clause describes interfaces and meta-data to facilitate the deployment of CORBA component based applications, and a process of generating a concrete XML schema for the meta-data, by using the rules defined by the XML Metadata Interchange (XMI) version 2 [XMI] specification on the Deployment PSM for CCM, which is derived from the PIM in the Deployment and Configuration [D+C] specification. With these rules, this clause contains the normative definition.

This clause contains an XML schema that has been produced by using these rules. It is non-normative, so in the case of discrepancies, this clause is relevant.

17.2 Schema

```
<xsd:schema
  xmlns:xsd="http://www.w3.org/2001/XMLSchema"
  xmlns:xmi="http://www.omg.org/XMI"
  xmlns:Deployment="http://www.omg.org/Deployment"
  targetNamespace="http://www.omg.org/Deployment"
  >
<xsd:import namespace="http://www.omg.org/XMI"/>
<xsd:complexType name="Any">
  <xsd:choice minOccurs="0" maxOccurs="unbounded">
    <xsd:element name="type" type="Deployment:DataType"/>
    <xsd:element name="value" type="Deployment:DataValue"/>
    <xsd:element ref="xmi:Extension"/>
  </xsd:choice>
  <xsd:attribute ref="xmi:id" use="optional"/>
  <xsd:attributeGroup ref="xmi:ObjectAttribs"/>
</xsd:complexType>
<xsd:element name="Any" type="Deployment:Any"/>
<xsd:complexType name="DataType">
  <xsd:choice minOccurs="0" maxOccurs="unbounded">
    <xsd:element name="kind" type="Deployment:TCKind"/>
    <xsd:element name="enum" type="Deployment:EnumType"/>
    <xsd:element name="objref" type="Deployment:ObjrefType"/>
    <xsd:element name="boundedString" type="Deployment:BoundedStringType"/>
    <xsd:element name="fixed" type="Deployment:FixedType"/>
    <xsd:element name="array" type="Deployment:ArrayType"/>
    <xsd:element name="sequence" type="Deployment:SequenceType"/>
    <xsd:element name="alias" type="Deployment:AliasType"/>
    <xsd:element name="struct" type="Deployment:StructType"/>
    <xsd:element name="value" type="Deployment:ValueType"/>
    <xsd:element name="union" type="Deployment:UnionType"/>
    <xsd:element ref="xmi:Extension"/>
  </xsd:choice>
  <xsd:attribute ref="xmi:id" use="optional"/>
  <xsd:attributeGroup ref="xmi:ObjectAttribs"/>
</xsd:complexType>
```

```

</xsd:complexType>
<xsd:element name="DataType" type="Deployment:DataType"/>
<xsd:complexType name="DataValue">
  <xsd:choice minOccurs="0" maxOccurs="unbounded">
    <xsd:element name="short" type="xsd:short"/>
    <xsd:element name="long" type="xsd:int"/>
    <xsd:element name="ushort" type="xsd:unsignedShort"/>
    <xsd:element name="ulong" type="xsd:unsignedInt"/>
    <xsd:element name="float" type="xsd:float"/>
    <xsd:element name="double" type="xsd:double"/>
    <xsd:element name="boolean" type="xsd:boolean"/>
    <xsd:element name="octet" type="xsd:unsignedByte"/>
    <xsd:element name="opaque" type="xsd:base64Binary"/>
    <xsd:element name="objref" type="xsd:string"/>
    <xsd:element name="enum" type="xsd:string"/>
    <xsd:element name="string" type="xsd:string"/>
    <xsd:element name="longlong" type="xsd:long"/>
    <xsd:element name="ulonglong" type="xsd:unsignedLong"/>
    <xsd:element name="longdouble" type="xsd:double"/>
    <xsd:element name="fixed" type="xsd:string"/>
    <xsd:element name="any" type="Deployment:Any"/>
    <xsd:element name="typecode" type="Deployment:DataType"/>
    <xsd:element name="element" type="Deployment:DataValue"/>
    <xsd:element name="discriminator" type="Deployment:DataValue"/>
    <xsd:element name="value" type="Deployment:DataValue"/>
    <xsd:element name="boxedValue" type="Deployment:DataValue"/>
    <xsd:element name="member" type="Deployment:NamedValue"/>
    <xsd:element ref="xmi:Extension"/>
  </xsd:choice>
  <xsd:attribute ref="xmi:id" use="optional"/>
  <xsd:attributeGroup ref="xmi:ObjectAttribs"/>
</xsd:complexType>
<xsd:element name="DataValue" type="Deployment:DataValue"/>
<xsd:complexType name="EnumType">
  <xsd:choice minOccurs="0" maxOccurs="unbounded">
    <xsd:element name="name" type="xsd:string"/>
    <xsd:element name="typeld" type="xsd:string"/>
    <xsd:element name="member" type="xsd:string"/>
    <xsd:element ref="xmi:Extension"/>
  </xsd:choice>
  <xsd:attribute ref="xmi:id" use="optional"/>
  <xsd:attributeGroup ref="xmi:ObjectAttribs"/>
</xsd:complexType>
<xsd:element name="EnumType" type="Deployment:EnumType"/>
<xsd:complexType name="ObjrefType">
  <xsd:choice minOccurs="0" maxOccurs="unbounded">
    <xsd:element name="name" type="xsd:string"/>
    <xsd:element name="typeld" type="xsd:string"/>
    <xsd:element ref="xmi:Extension"/>
  </xsd:choice>
  <xsd:attribute ref="xmi:id" use="optional"/>

```



```

<xsd:attributeGroup ref="xmi:ObjectAttribs"/>
</xsd:complexType>
<xsd:element name="ObjrefType" type="Deployment:ObjrefType"/>
<xsd:complexType name="BoundedStringType">
  <xsd:choice minOccurs="0" maxOccurs="unbounded">
    <xsd:element name="bound" type="xsd:string"/>
    <xsd:element ref="xmi:Extension"/>
  </xsd:choice>
  <xsd:attribute ref="xmi:id" use="optional"/>
  <xsd:attributeGroup ref="xmi:ObjectAttribs"/>
</xsd:complexType>
<xsd:element name="BoundedStringType" type="Deployment:BoundedStringType"/>
<xsd:complexType name="FixedType">
  <xsd:choice minOccurs="0" maxOccurs="unbounded">
    <xsd:element name="digits" type="xsd:string"/>
    <xsd:element name="scale" type="xsd:string"/>
    <xsd:element ref="xmi:Extension"/>
  </xsd:choice>
  <xsd:attribute ref="xmi:id" use="optional"/>
  <xsd:attributeGroup ref="xmi:ObjectAttribs"/>
</xsd:complexType>
<xsd:element name="FixedType" type="Deployment:FixedType"/>
<xsd:complexType name="ArrayType">
  <xsd:choice minOccurs="0" maxOccurs="unbounded">
    <xsd:element name="length" type="xsd:string"/>
    <xsd:element name="elementType" type="Deployment:DataType"/>
    <xsd:element ref="xmi:Extension"/>
  </xsd:choice>
  <xsd:attribute ref="xmi:id" use="optional"/>
  <xsd:attributeGroup ref="xmi:ObjectAttribs"/>
</xsd:complexType>
<xsd:element name="ArrayType" type="Deployment:ArrayType"/>
<xsd:complexType name="SequenceType">
  <xsd:choice minOccurs="0" maxOccurs="unbounded">
    <xsd:element name="bound" type="xsd:string"/>
    <xsd:element name="elementType" type="Deployment:DataType"/>
    <xsd:element ref="xmi:Extension"/>
  </xsd:choice>
  <xsd:attribute ref="xmi:id" use="optional"/>
  <xsd:attributeGroup ref="xmi:ObjectAttribs"/>
</xsd:complexType>
<xsd:element name="SequenceType" type="Deployment:SequenceType"/>
<xsd:complexType name="AliasType">
  <xsd:choice minOccurs="0" maxOccurs="unbounded">
    <xsd:element name="name" type="xsd:string"/>
    <xsd:element name="typeld" type="xsd:string"/>
    <xsd:element name="elementType" type="Deployment:DataType"/>
    <xsd:element ref="xmi:Extension"/>
  </xsd:choice>
  <xsd:attribute ref="xmi:id" use="optional"/>
  <xsd:attributeGroup ref="xmi:ObjectAttribs"/>

```

```

</xsd:complexType>
<xsd:element name="AliasType" type="Deployment:AliasType"/>
<xsd:complexType name="StructType">
  <xsd:choice minOccurs="0" maxOccurs="unbounded">
    <xsd:element name="name" type="xsd:string"/>
    <xsd:element name="typeld" type="xsd:string"/>
    <xsd:element name="member" type="Deployment:StructMemberType"/>
    <xsd:element ref="xmi:Extension"/>
  </xsd:choice>
  <xsd:attribute ref="xmi:id" use="optional"/>
  <xsd:attributeGroup ref="xmi:ObjectAttribs"/>
</xsd:complexType>
<xsd:element name="StructType" type="Deployment:StructType"/>
<xsd:complexType name="StructMemberType">
  <xsd:choice minOccurs="0" maxOccurs="unbounded">
    <xsd:element name="name" type="xsd:string"/>
    <xsd:element name="type" type="Deployment:DataType"/>
    <xsd:element ref="xmi:Extension"/>
  </xsd:choice>
  <xsd:attribute ref="xmi:id" use="optional"/>
  <xsd:attributeGroup ref="xmi:ObjectAttribs"/>
</xsd:complexType>
<xsd:element name="StructMemberType" type="Deployment:StructMemberType"/>
<xsd:complexType name="ValueType">
  <xsd:choice minOccurs="0" maxOccurs="unbounded">
    <xsd:element name="name" type="xsd:string"/>
    <xsd:element name="typeld" type="xsd:string"/>
    <xsd:element name="modifier" type="xsd:string"/>
    <xsd:element name="baseType" type="Deployment:DataType"/>
    <xsd:element name="member" type="Deployment:ValueMemberType"/>
    <xsd:element ref="xmi:Extension"/>
  </xsd:choice>
  <xsd:attribute ref="xmi:id" use="optional"/>
  <xsd:attributeGroup ref="xmi:ObjectAttribs"/>
</xsd:complexType>
<xsd:element name="ValueType" type="Deployment:ValueType"/>
<xsd:complexType name="ValueMemberType">
  <xsd:choice minOccurs="0" maxOccurs="unbounded">
    <xsd:element name="name" type="xsd:string"/>
    <xsd:element name="visibility" type="xsd:string"/>
    <xsd:element name="type" type="Deployment:DataType"/>
    <xsd:element ref="xmi:Extension"/>
  </xsd:choice>
  <xsd:attribute ref="xmi:id" use="optional"/>
  <xsd:attributeGroup ref="xmi:ObjectAttribs"/>
</xsd:complexType>
<xsd:element name="ValueMemberType" type="Deployment:ValueMemberType"/>
<xsd:complexType name="UnionType">
  <xsd:choice minOccurs="0" maxOccurs="unbounded">
    <xsd:element name="name" type="xsd:string"/>
    <xsd:element name="typeld" type="xsd:string"/>
  </xsd:choice>

```

```

    <xsd:element name="default" type="Deployment:UnionMemberType"/>
    <xsd:element name="discriminatorType" type="Deployment:DataType"/>
    <xsd:element name="member" type="Deployment:UnionMemberType"/>
    <xsd:element ref="xmi:Extension"/>
  </xsd:choice>
  <xsd:attribute ref="xmi:id" use="optional"/>
  <xsd:attributeGroup ref="xmi:ObjectAttribs"/>
</xsd:complexType>
<xsd:element name="UnionType" type="Deployment:UnionType"/>
<xsd:complexType name="UnionMemberType">
  <xsd:choice minOccurs="0" maxOccurs="unbounded">
    <xsd:element name="name" type="xsd:string"/>
    <xsd:element name="type" type="Deployment:DataType"/>
    <xsd:element name="label" type="Deployment:DataValue"/>
    <xsd:element ref="xmi:Extension"/>
  </xsd:choice>
  <xsd:attribute ref="xmi:id" use="optional"/>
  <xsd:attributeGroup ref="xmi:ObjectAttribs"/>
</xsd:complexType>
<xsd:element name="UnionMemberType" type="Deployment:UnionMemberType"/>
<xsd:complexType name="NamedValue">
  <xsd:choice minOccurs="0" maxOccurs="unbounded">
    <xsd:element name="name" type="xsd:string"/>
    <xsd:element name="value" type="Deployment:DataValue"/>
    <xsd:element ref="xmi:Extension"/>
  </xsd:choice>
  <xsd:attribute ref="xmi:id" use="optional"/>
  <xsd:attributeGroup ref="xmi:ObjectAttribs"/>
</xsd:complexType>
<xsd:element name="NamedValue" type="Deployment:NamedValue"/>
<xsd:complexType name="Bridge">
  <xsd:choice minOccurs="0" maxOccurs="unbounded">
    <xsd:element name="name" type="xsd:string"/>
    <xsd:element name="label" type="xsd:string"/>
    <xsd:element name="connect" type="Deployment:Interconnect"/>
    <xsd:element name="resource" type="Deployment:Resource"/>
    <xsd:element ref="xmi:Extension"/>
  </xsd:choice>
  <xsd:attribute ref="xmi:id" use="optional"/>
  <xsd:attributeGroup ref="xmi:ObjectAttribs"/>
</xsd:complexType>
<xsd:element name="Bridge" type="Deployment:Bridge"/>
<xsd:complexType name="Interconnect">
  <xsd:choice minOccurs="0" maxOccurs="unbounded">
    <xsd:element name="name" type="xsd:string"/>
    <xsd:element name="label" type="xsd:string"/>
    <xsd:element name="connection" type="Deployment:Bridge"/>
    <xsd:element name="connect" type="Deployment:Node"/>
    <xsd:element name="resource" type="Deployment:Resource"/>
    <xsd:element ref="xmi:Extension"/>
  </xsd:choice>

```

```

<xsd:attribute ref="xmi:id" use="optional"/>
<xsd:attributeGroup ref="xmi:ObjectAttribs"/>
</xsd:complexType>
<xsd:element name="Interconnect" type="Deployment:Interconnect"/>
<xsd:complexType name="Node">
  <xsd:choice minOccurs="0" maxOccurs="unbounded">
    <xsd:element name="name" type="xsd:string"/>
    <xsd:element name="label" type="xsd:string"/>
    <xsd:element name="connection" type="Deployment:Interconnect"/>
    <xsd:element name="sharedResource" type="Deployment:SharedResource"/>
    <xsd:element name="resource" type="Deployment:Resource"/>
    <xsd:element ref="xmi:Extension"/>
  </xsd:choice>
  <xsd:attribute ref="xmi:id" use="optional"/>
  <xsd:attributeGroup ref="xmi:ObjectAttribs"/>
</xsd:complexType>
<xsd:element name="Node" type="Deployment:Node"/>
<xsd:complexType name="Resource">
  <xsd:choice minOccurs="0" maxOccurs="unbounded">
    <xsd:element name="name" type="xsd:string"/>
    <xsd:element name="resourceType" type="xsd:string"/>
    <xsd:element name="property" type="Deployment:SatisfierProperty"/>
    <xsd:element ref="xmi:Extension"/>
  </xsd:choice>
  <xsd:attribute ref="xmi:id" use="optional"/>
  <xsd:attributeGroup ref="xmi:ObjectAttribs"/>
</xsd:complexType>
<xsd:element name="Resource" type="Deployment:Resource"/>
<xsd:complexType name="SharedResource">
  <xsd:choice minOccurs="0" maxOccurs="unbounded">
    <xsd:element name="name" type="xsd:string"/>
    <xsd:element name="resourceType" type="xsd:string"/>
    <xsd:element name="node" type="Deployment:Node"/>
    <xsd:element name="property" type="Deployment:SatisfierProperty"/>
    <xsd:element ref="xmi:Extension"/>
  </xsd:choice>
  <xsd:attribute ref="xmi:id" use="optional"/>
  <xsd:attributeGroup ref="xmi:ObjectAttribs"/>
</xsd:complexType>
<xsd:element name="SharedResource" type="Deployment:SharedResource"/>
<xsd:complexType name="Domain">
  <xsd:choice minOccurs="0" maxOccurs="unbounded">
    <xsd:element name="UUID" type="xsd:string"/>
    <xsd:element name="label" type="xsd:string"/>
    <xsd:element name="sharedResource" type="Deployment:SharedResource"/>
    <xsd:element name="node" type="Deployment:Node"/>
    <xsd:element name="interconnect" type="Deployment:Interconnect"/>
    <xsd:element name="bridge" type="Deployment:Bridge"/>
    <xsd:element name="infoProperty" type="Deployment:Property"/>
    <xsd:element ref="xmi:Extension"/>
  </xsd:choice>

```

```

<xsd:attribute ref="xmi:id" use="optional"/>
<xsd:attributeGroup ref="xmi:ObjectAttribs"/>
</xsd:complexType>
<xsd:element name="Domain" type="Deployment:Domain"/>
<xsd:complexType name="ResourceAllocation">
  <xsd:choice minOccurs="0" maxOccurs="unbounded">
    <xsd:element name="elementName" type="xsd:string"/>
    <xsd:element name="resourceName" type="xsd:string"/>
    <xsd:element name="property" type="Deployment:Property"/>
    <xsd:element ref="xmi:Extension"/>
  </xsd:choice>
  <xsd:attribute ref="xmi:id" use="optional"/>
  <xsd:attributeGroup ref="xmi:ObjectAttribs"/>
</xsd:complexType>
<xsd:element name="ResourceAllocation" type="Deployment:ResourceAllocation"/>
<xsd:complexType name="PlanPropertyMapping">
  <xsd:choice minOccurs="0" maxOccurs="unbounded">
    <xsd:element name="name" type="xsd:string"/>
    <xsd:element name="source" type="xsd:string"/>
    <xsd:element name="externalName" type="xsd:string"/>
    <xsd:element name="delegatesTo" type="Deployment:PlanSubcomponentPropertyReference"/>
    <xsd:element ref="xmi:Extension"/>
  </xsd:choice>
  <xsd:attribute ref="xmi:id" use="optional"/>
  <xsd:attributeGroup ref="xmi:ObjectAttribs"/>
</xsd:complexType>
<xsd:element name="PlanPropertyMapping" type="Deployment:PlanPropertyMapping"/>
<xsd:complexType name="PlanSubcomponentPropertyReference">
  <xsd:choice minOccurs="0" maxOccurs="unbounded">
    <xsd:element name="propertyName" type="xsd:string"/>
    <xsd:element name="instance" type="Deployment:InstanceDeploymentDescription"/>
    <xsd:element ref="xmi:Extension"/>
  </xsd:choice>
  <xsd:attribute ref="xmi:id" use="optional"/>
  <xsd:attributeGroup ref="xmi:ObjectAttribs"/>
</xsd:complexType>
<xsd:element name="PlanSubcomponentPropertyReference" type="Deployment:PlanSubcomponent-
PropertyReference"/>
<xsd:complexType name="PlanSubcomponentPortEndpoint">
  <xsd:choice minOccurs="0" maxOccurs="unbounded">
    <xsd:element name="portName" type="xsd:string"/>
    <xsd:element name="provider" type="xsd:string"/>
    <xsd:element name="kind" type="Deployment:CCMComponentPortKind"/>
    <xsd:element name="instance" type="Deployment:InstanceDeploymentDescription"/>
    <xsd:element ref="xmi:Extension"/>
  </xsd:choice>
  <xsd:attribute ref="xmi:id" use="optional"/>
  <xsd:attributeGroup ref="xmi:ObjectAttribs"/>
</xsd:complexType>
<xsd:element name="PlanSubcomponentPortEndpoint" type="Deployment:PlanSubcomponentPortEnd-
point"/>

```

```

<xsd:complexType name="PlanConnectionDescription">
  <xsd:choice minOccurs="0" maxOccurs="unbounded">
    <xsd:element name="name" type="xsd:string"/>
    <xsd:element name="source" type="xsd:string"/>
    <xsd:element name="deployRequirement" type="Deployment:Requirement"/>
    <xsd:element name="externalEndpoint" type="Deployment:ComponentExternalPortEndpoint"/>
    <xsd:element name="internalEndpoint" type="Deployment:PlanSubcomponentPortEndpoint"/>
    <xsd:element name="externalReference" type="Deployment:ExternalReferenceEndpoint"/>
    <xsd:element name="deployedResource" type="Deployment:ConnectionResourceDeploymentDe-
scription"/>
    <xsd:element ref="xmi:Extension"/>
  </xsd:choice>
  <xsd:attribute ref="xmi:id" use="optional"/>
  <xsd:attributeGroup ref="xmi:ObjectAttribs"/>
</xsd:complexType>
<xsd:element name="PlanConnectionDescription" type="Deployment:PlanConnectionDescription"/>
<xsd:complexType name="InstanceDeploymentDescription">
  <xsd:choice minOccurs="0" maxOccurs="unbounded">
    <xsd:element name="name" type="xsd:string"/>
    <xsd:element name="node" type="xsd:string"/>
    <xsd:element name="source" type="xsd:string"/>
    <xsd:element name="implementation" type="Deployment:MonolithicDeploymentDescription"/>
    <xsd:element name="configProperty" type="Deployment:Property"/>
    <xsd:element name="deployedResource" type="Deployment:InstanceResourceDeploymentDescrip-
tion"/>
    <xsd:element name="deployedSharedResource" type="Deployment:InstanceResourceDeploymentDe-
scription"/>
    <xsd:element ref="xmi:Extension"/>
  </xsd:choice>
  <xsd:attribute ref="xmi:id" use="optional"/>
  <xsd:attributeGroup ref="xmi:ObjectAttribs"/>
</xsd:complexType>
<xsd:element name="InstanceDeploymentDescription" type="Deployment:InstanceDeploymentDescrip-
tion"/>
<xsd:complexType name="MonolithicDeploymentDescription">
  <xsd:choice minOccurs="0" maxOccurs="unbounded">
    <xsd:element name="name" type="xsd:string"/>
    <xsd:element name="source" type="xsd:string"/>
    <xsd:element name="artifact" type="Deployment:ArtifactDeploymentDescription"/>
    <xsd:element name="execParameter" type="Deployment:Property"/>
    <xsd:element name="deployRequirement" type="Deployment:Requirement"/>
    <xsd:element ref="xmi:Extension"/>
  </xsd:choice>
  <xsd:attribute ref="xmi:id" use="optional"/>
  <xsd:attributeGroup ref="xmi:ObjectAttribs"/>
</xsd:complexType>
<xsd:element name="MonolithicDeploymentDescription" type="Deployment:MonolithicDeploymentDe-
scription"/>
<xsd:complexType name="ArtifactDeploymentDescription">
  <xsd:choice minOccurs="0" maxOccurs="unbounded">
    <xsd:element name="name" type="xsd:string"/>

```



```

<xsd:element name="location" type="xsd:string"/>
<xsd:element name="node" type="xsd:string"/>
<xsd:element name="source" type="xsd:string"/>
<xsd:element name="execParameter" type="Deployment:Property"/>
<xsd:element name="deployRequirement" type="Deployment:Requirement"/>
<xsd:element name="deployedResource" type="Deployment:ResourceDeploymentDescription"/>
<xsd:element ref="xmi:Extension"/>
</xsd:choice>
<xsd:attribute ref="xmi:id" use="optional"/>
<xsd:attributeGroup ref="xmi:ObjectAttribs"/>
</xsd:complexType>
<xsd:element name="ArtifactDeploymentDescription" type="Deployment:ArtifactDeploymentDescription"/>
<xsd:complexType name="DeploymentPlan">
<xsd:choice minOccurs="0" maxOccurs="unbounded">
<xsd:element name="label" type="xsd:string"/>
<xsd:element name="UUID" type="xsd:string"/>
<xsd:element name="realizes" type="Deployment:ComponentInterfaceDescription"/>
<xsd:element name="implementation" type="Deployment:MonolithicDeploymentDescription"/>
<xsd:element name="instance" type="Deployment:InstanceDeploymentDescription"/>
<xsd:element name="connection" type="Deployment:PlanConnectionDescription"/>
<xsd:element name="externalProperty" type="Deployment:PlanPropertyMapping"/>
<xsd:element name="dependsOn" type="Deployment:ImplementationDependency"/>
<xsd:element name="artifact" type="Deployment:ArtifactDeploymentDescription"/>
<xsd:element name="infoProperty" type="Deployment:Property"/>
<xsd:element name="localityConstraint" type="Deployment:PlanLocality"/>
<xsd:element ref="xmi:Extension"/>
</xsd:choice>
<xsd:attribute ref="xmi:id" use="optional"/>
<xsd:attributeGroup ref="xmi:ObjectAttribs"/>
</xsd:complexType>
<xsd:element name="DeploymentPlan" type="Deployment:DeploymentPlan"/>
<xsd:complexType name="ResourceDeploymentDescription">
<xsd:choice minOccurs="0" maxOccurs="unbounded">
<xsd:element name="requirementName" type="xsd:string"/>
<xsd:element name="resourceName" type="xsd:string"/>
<xsd:element name="property" type="Deployment:Property"/>
<xsd:element ref="xmi:Extension"/>
</xsd:choice>
<xsd:attribute ref="xmi:id" use="optional"/>
<xsd:attributeGroup ref="xmi:ObjectAttribs"/>
</xsd:complexType>
<xsd:element name="ResourceDeploymentDescription" type="Deployment:ResourceDeploymentDescription"/>
<xsd:complexType name="InstanceResourceDeploymentDescription">
<xsd:choice minOccurs="0" maxOccurs="unbounded">
<xsd:element name="resourceUsage" type="Deployment:ResourceUsageKind"/>
<xsd:element name="requirementName" type="xsd:string"/>
<xsd:element name="resourceName" type="xsd:string"/>
<xsd:element name="property" type="Deployment:Property"/>
<xsd:element ref="xmi:Extension"/>

```

```

</xsd:choice>
<xsd:attribute ref="xmi:id" use="optional"/>
<xsd:attributeGroup ref="xmi:ObjectAttribs"/>
</xsd:complexType>
<xsd:element name="InstanceResourceDeploymentDescription" type="Deployment:InstanceResourceDeploymentDescription"/>
<xsd:complexType name="ConnectionResourceDeploymentDescription">
<xsd:choice minOccurs="0" maxOccurs="unbounded">
<xsd:element name="targetName" type="xsd:string"/>
<xsd:element name="requirementName" type="xsd:string"/>
<xsd:element name="resourceName" type="xsd:string"/>
<xsd:element name="property" type="Deployment:Property"/>
<xsd:element ref="xmi:Extension"/>
</xsd:choice>
<xsd:attribute ref="xmi:id" use="optional"/>
<xsd:attributeGroup ref="xmi:ObjectAttribs"/>
</xsd:complexType>
<xsd:element name="ConnectionResourceDeploymentDescription" type="Deployment:ConnectionResourceDeploymentDescription"/>
<xsd:complexType name="PlanLocality">
<xsd:choice minOccurs="0" maxOccurs="unbounded">
<xsd:element name="constraint" type="Deployment:PlanLocalityKind"/>
<xsd:element name="constrainedInstance" type="Deployment:InstanceDeploymentDescription"/>
<xsd:element ref="xmi:Extension"/>
</xsd:choice>
<xsd:attribute ref="xmi:id" use="optional"/>
<xsd:attributeGroup ref="xmi:ObjectAttribs"/>
</xsd:complexType>
<xsd:element name="PlanLocality" type="Deployment:PlanLocality"/>
<xsd:complexType name="Capability">
<xsd:choice minOccurs="0" maxOccurs="unbounded">
<xsd:element name="name" type="xsd:string"/>
<xsd:element name="resourceType" type="xsd:string"/>
<xsd:element name="property" type="Deployment:SatisfierProperty"/>
<xsd:element ref="xmi:Extension"/>
</xsd:choice>
<xsd:attribute ref="xmi:id" use="optional"/>
<xsd:attributeGroup ref="xmi:ObjectAttribs"/>
</xsd:complexType>
<xsd:element name="Capability" type="Deployment:Capability"/>
<xsd:complexType name="ComponentPropertyDescription">
<xsd:choice minOccurs="0" maxOccurs="unbounded">
<xsd:element name="name" type="xsd:string"/>
<xsd:element name="type" type="Deployment:DataType"/>
<xsd:element ref="xmi:Extension"/>
</xsd:choice>
<xsd:attribute ref="xmi:id" use="optional"/>
<xsd:attributeGroup ref="xmi:ObjectAttribs"/>
</xsd:complexType>
<xsd:element name="ComponentPropertyDescription" type="Deployment:ComponentPropertyDescription"/>

```



```

<xsd:complexType name="ComponentPortDescription">
  <xsd:choice minOccurs="0" maxOccurs="unbounded">
    <xsd:element name="name" type="xsd:string"/>
    <xsd:element name="specificType" type="xsd:string"/>
    <xsd:element name="supportedType" type="xsd:string"/>
    <xsd:element name="provider" type="xsd:string"/>
    <xsd:element name="exclusiveProvider" type="xsd:string"/>
    <xsd:element name="exclusiveUser" type="xsd:string"/>
    <xsd:element name="optional" type="xsd:string"/>
    <xsd:element name="kind" type="Deployment:CCMComponentPortKind"/>
    <xsd:element ref="xmi:Extension"/>
  </xsd:choice>
  <xsd:attribute ref="xmi:id" use="optional"/>
  <xsd:attributeGroup ref="xmi:ObjectAttribs"/>
</xsd:complexType>
<xsd:element name="ComponentPortDescription" type="Deployment:ComponentPortDescription"/>
<xsd:complexType name="ComponentInterfaceDescription">
  <xsd:choice minOccurs="0" maxOccurs="unbounded">
    <xsd:element name="label" type="xsd:string"/>
    <xsd:element name="UUID" type="xsd:string"/>
    <xsd:element name="specificType" type="xsd:string"/>
    <xsd:element name="supportedType" type="xsd:string"/>
    <xsd:element name="idOfFile" type="xsd:string"/>
    <xsd:element name="configProperty" type="Deployment:Property"/>
    <xsd:element name="port" type="Deployment:ComponentPortDescription"/>
    <xsd:element name="property" type="Deployment:ComponentPropertyDescription"/>
    <xsd:element name="infoProperty" type="Deployment:Property"/>
    <xsd:element ref="xmi:Extension"/>
  </xsd:choice>
  <xsd:attribute ref="xmi:id" use="optional"/>
  <xsd:attributeGroup ref="xmi:ObjectAttribs"/>
</xsd:complexType>
<xsd:element name="ComponentInterfaceDescription" type="Deployment:ComponentInterfaceDescription"/>
<xsd:complexType name="ImplementationArtifactDescription">
  <xsd:choice minOccurs="0" maxOccurs="unbounded">
    <xsd:element name="label" type="xsd:string"/>
    <xsd:element name="UUID" type="xsd:string"/>
    <xsd:element name="location" type="xsd:string"/>
    <xsd:element name="execParameter" type="Deployment:Property"/>
    <xsd:element name="deployRequirement" type="Deployment:Requirement"/>
    <xsd:element name="dependsOn" type="Deployment:NamedImplementationArtifact"/>
    <xsd:element name="infoProperty" type="Deployment:Property"/>
    <xsd:element ref="xmi:Extension"/>
  </xsd:choice>
  <xsd:attribute ref="xmi:id" use="optional"/>
  <xsd:attributeGroup ref="xmi:ObjectAttribs"/>
</xsd:complexType>
<xsd:element name="ImplementationArtifactDescription" type="Deployment:ImplementationArtifactDescription"/>
<xsd:complexType name="MonolithicImplementationDescription">

```

```

<xsd:choice minOccurs="0" maxOccurs="unbounded">
  <xsd:element name="nodeExecParameter" type="Deployment:Property"/>
  <xsd:element name="primaryArtifact" type="Deployment:NamedImplementationArtifact"/>
  <xsd:element name="deployRequirement" type="Deployment:ImplementationRequirement"/>
  <xsd:element name="componentExecParameter" type="Deployment:Property"/>
  <xsd:element ref="xmi:Extension"/>
</xsd:choice>
<xsd:attribute ref="xmi:id" use="optional"/>
<xsd:attributeGroup ref="xmi:ObjectAttribs"/>
</xsd:complexType>
<xsd:element name="MonolithicImplementationDescription" type="Deployment:MonolithicImplementationDescription"/>
<xsd:complexType name="AssemblyPropertyMapping">
  <xsd:choice minOccurs="0" maxOccurs="unbounded">
    <xsd:element name="name" type="xsd:string"/>
    <xsd:element name="externalName" type="xsd:string"/>
    <xsd:element name="delegatesTo" type="Deployment:SubcomponentPropertyReference"/>
    <xsd:element ref="xmi:Extension"/>
  </xsd:choice>
  <xsd:attribute ref="xmi:id" use="optional"/>
  <xsd:attributeGroup ref="xmi:ObjectAttribs"/>
</xsd:complexType>
<xsd:element name="AssemblyPropertyMapping" type="Deployment:AssemblyPropertyMapping"/>
<xsd:complexType name="SubcomponentPropertyReference">
  <xsd:choice minOccurs="0" maxOccurs="unbounded">
    <xsd:element name="propertyName" type="xsd:string"/>
    <xsd:element name="instance" type="Deployment:SubcomponentInstantiationDescription"/>
    <xsd:element ref="xmi:Extension"/>
  </xsd:choice>
  <xsd:attribute ref="xmi:id" use="optional"/>
  <xsd:attributeGroup ref="xmi:ObjectAttribs"/>
</xsd:complexType>
<xsd:element name="SubcomponentPropertyReference" type="Deployment:SubcomponentPropertyReference"/>
<xsd:complexType name="SubcomponentPortEndpoint">
  <xsd:choice minOccurs="0" maxOccurs="unbounded">
    <xsd:element name="portName" type="xsd:string"/>
    <xsd:element name="instance" type="Deployment:SubcomponentInstantiationDescription"/>
    <xsd:element ref="xmi:Extension"/>
  </xsd:choice>
  <xsd:attribute ref="xmi:id" use="optional"/>
  <xsd:attributeGroup ref="xmi:ObjectAttribs"/>
</xsd:complexType>
<xsd:element name="SubcomponentPortEndpoint" type="Deployment:SubcomponentPortEndpoint"/>
<xsd:complexType name="AssemblyConnectionDescription">
  <xsd:choice minOccurs="0" maxOccurs="unbounded">
    <xsd:element name="name" type="xsd:string"/>
    <xsd:element name="deployRequirement" type="Deployment:Requirement"/>
    <xsd:element name="externalEndpoint" type="Deployment:ComponentExternalPortEndpoint"/>
    <xsd:element name="internalEndpoint" type="Deployment:SubcomponentPortEndpoint"/>
    <xsd:element name="externalReference" type="Deployment:ExternalReferenceEndpoint"/>
  </xsd:choice>
  <xsd:attribute ref="xmi:id" use="optional"/>
  <xsd:attributeGroup ref="xmi:ObjectAttribs"/>
</xsd:complexType>

```

```

    <xsd:element ref="xmi:Extension"/>
</xsd:choice>
<xsd:attribute ref="xmi:id" use="optional"/>
<xsd:attributeGroup ref="xmi:ObjectAttribs"/>
</xsd:complexType>
<xsd:element name="AssemblyConnectionDescription" type="Deployment:AssemblyConnectionDe-
scription"/>
<xsd:complexType name="SubcomponentInstantiationDescription">
<xsd:choice minOccurs="0" maxOccurs="unbounded">
<xsd:element name="name" type="xsd:string"/>
<xsd:element name="basePackage" type="Deployment:ComponentPackageDescription"/>
<xsd:element name="specializedConfig" type="Deployment:PackageConfiguration"/>
<xsd:element name="selectRequirement" type="Deployment:Requirement"/>
<xsd:element name="configProperty" type="Deployment:Property"/>
<xsd:element name="referencedPackage" type="Deployment:ComponentPackageReference"/>
<xsd:element name="importedPackage" type="Deployment:ComponentPackageImport"/>
<xsd:element ref="xmi:Extension"/>
</xsd:choice>
<xsd:attribute ref="xmi:id" use="optional"/>
<xsd:attributeGroup ref="xmi:ObjectAttribs"/>
</xsd:complexType>
<xsd:element name="SubcomponentInstantiationDescription" type="Deployment:SubcomponentInstan-
tiationDescription"/>
<xsd:complexType name="ComponentAssemblyDescription">
<xsd:choice minOccurs="0" maxOccurs="unbounded">
<xsd:element name="instance" type="Deployment:SubcomponentInstantiationDescription"/>
<xsd:element name="connection" type="Deployment:AssemblyConnectionDescription"/>
<xsd:element name="externalProperty" type="Deployment:AssemblyPropertyMapping"/>
<xsd:element name="localityConstraint" type="Deployment:Locality"/>
<xsd:element ref="xmi:Extension"/>
</xsd:choice>
<xsd:attribute ref="xmi:id" use="optional"/>
<xsd:attributeGroup ref="xmi:ObjectAttribs"/>
</xsd:complexType>
<xsd:element name="ComponentAssemblyDescription" type="Deployment:ComponentAssembly-
Description"/>
<xsd:complexType name="ComponentImplementationDescription">
<xsd:choice minOccurs="0" maxOccurs="unbounded">
<xsd:element name="label" type="xsd:string"/>
<xsd:element name="UUID" type="xsd:string"/>
<xsd:element name="implements" type="Deployment:ComponentInterfaceDescription"/>
<xsd:element name="assemblyImpl" type="Deployment:ComponentAssemblyDescription"/>
<xsd:element name="monolithicImpl" type="Deployment:MonolithicImplementationDescription"/>
<xsd:element name="configProperty" type="Deployment:Property"/>
<xsd:element name="capability" type="Deployment:Capability"/>
<xsd:element name="dependsOn" type="Deployment:ImplementationDependency"/>
<xsd:element name="infoProperty" type="Deployment:Property"/>
<xsd:element ref="xmi:Extension"/>
</xsd:choice>
<xsd:attribute ref="xmi:id" use="optional"/>
<xsd:attributeGroup ref="xmi:ObjectAttribs"/>

```

```

</xsd:complexType>
<xsd:element name="ComponentImplementationDescription" type="Deployment:ComponentImplement-
ationDescription"/>
<xsd:complexType name="ComponentPackageReference">
  <xsd:choice minOccurs="0" maxOccurs="unbounded">
    <xsd:element name="requiredUUID" type="xsd:string"/>
    <xsd:element name="requiredName" type="xsd:string"/>
    <xsd:element name="requiredType" type="Deployment:ComponentInterfaceDescription"/>
    <xsd:element ref="xmi:Extension"/>
  </xsd:choice>
  <xsd:attribute ref="xmi:id" use="optional"/>
  <xsd:attributeGroup ref="xmi:ObjectAttribs"/>
</xsd:complexType>
<xsd:element name="ComponentPackageReference" type="Deployment:ComponentPackageReference"/>
>
<xsd:complexType name="ComponentPackageDescription">
  <xsd:choice minOccurs="0" maxOccurs="unbounded">
    <xsd:element name="label" type="xsd:string"/>
    <xsd:element name="UUID" type="xsd:string"/>
    <xsd:element name="realizes" type="Deployment:ComponentInterfaceDescription"/>
    <xsd:element name="configProperty" type="Deployment:Property"/>
    <xsd:element name="implementation" type="Deployment:PackagedComponentImplementation"/>
    <xsd:element name="infoProperty" type="Deployment:Property"/>
    <xsd:element ref="xmi:Extension"/>
  </xsd:choice>
  <xsd:attribute ref="xmi:id" use="optional"/>
  <xsd:attributeGroup ref="xmi:ObjectAttribs"/>
</xsd:complexType>
<xsd:element name="ComponentPackageDescription" type="Deployment:ComponentPackageDescrip-
tion"/>
<xsd:complexType name="ComponentUsageDescription">
  <xsd:choice minOccurs="0" maxOccurs="unbounded">
    <xsd:element name="basePackage" type="Deployment:ComponentPackageDescription"/>
    <xsd:element name="specializedConfig" type="Deployment:PackageConfiguration"/>
    <xsd:element name="selectRequirement" type="Deployment:Requirement"/>
    <xsd:element name="configProperty" type="Deployment:Property"/>
    <xsd:element name="referencedPackage" type="Deployment:ComponentPackageReference"/>
    <xsd:element name="importedPackage" type="Deployment:ComponentPackageImport"/>
    <xsd:element ref="xmi:Extension"/>
  </xsd:choice>
  <xsd:attribute ref="xmi:id" use="optional"/>
  <xsd:attributeGroup ref="xmi:ObjectAttribs"/>
</xsd:complexType>
<xsd:element name="ComponentUsageDescription" type="Deployment:ComponentUsageDescription"/>
<xsd:complexType name="PackagedComponentImplementation">
  <xsd:choice minOccurs="0" maxOccurs="unbounded">
    <xsd:element name="name" type="xsd:string"/>
    <xsd:element name="referencedImplementation" type="Deployment:ComponentImplementationDe-
scription"/>
    <xsd:element ref="xmi:Extension"/>
  </xsd:choice>

```

```

    <xsd:attribute ref="xmi:id" use="optional"/>
    <xsd:attributeGroup ref="xmi:ObjectAttribs"/>
  </xsd:complexType>
  <xsd:element name="PackagedComponentImplementation" type="Deployment:PackagedComponentImplementation"/>
  <xsd:complexType name="NamedImplementationArtifact">
    <xsd:choice minOccurs="0" maxOccurs="unbounded">
      <xsd:element name="name" type="xsd:string"/>
      <xsd:element name="referencedArtifact" type="Deployment:ImplementationArtifactDescription"/>
      <xsd:element ref="xmi:Extension"/>
    </xsd:choice>
    <xsd:attribute ref="xmi:id" use="optional"/>
    <xsd:attributeGroup ref="xmi:ObjectAttribs"/>
  </xsd:complexType>
  <xsd:element name="NamedImplementationArtifact" type="Deployment:NamedImplementationArtifact"/>
  >
  <xsd:complexType name="ImplementationRequirement">
    <xsd:choice minOccurs="0" maxOccurs="unbounded">
      <xsd:element name="resourceUsage" type="Deployment:ResourceUsageKind"/>
      <xsd:element name="resourcePort" type="xsd:string"/>
      <xsd:element name="componentPort" type="xsd:string"/>
      <xsd:element name="name" type="xsd:string"/>
      <xsd:element name="resourceType" type="xsd:string"/>
      <xsd:element name="property" type="Deployment:Property"/>
      <xsd:element ref="xmi:Extension"/>
    </xsd:choice>
    <xsd:attribute ref="xmi:id" use="optional"/>
    <xsd:attributeGroup ref="xmi:ObjectAttribs"/>
  </xsd:complexType>
  <xsd:element name="ImplementationRequirement" type="Deployment:ImplementationRequirement"/>
  <xsd:complexType name="ComponentPackageImport">
    <xsd:choice minOccurs="0" maxOccurs="unbounded">
      <xsd:element name="location" type="xsd:string"/>
      <xsd:element ref="xmi:Extension"/>
    </xsd:choice>
    <xsd:attribute ref="xmi:id" use="optional"/>
    <xsd:attributeGroup ref="xmi:ObjectAttribs"/>
  </xsd:complexType>
  <xsd:element name="ComponentPackageImport" type="Deployment:ComponentPackageImport"/>
  <xsd:complexType name="PackageConfiguration">
    <xsd:choice minOccurs="0" maxOccurs="unbounded">
      <xsd:element name="label" type="xsd:string"/>
      <xsd:element name="UUID" type="xsd:string"/>
      <xsd:element name="basePackage" type="Deployment:ComponentPackageDescription"/>
      <xsd:element name="specializedConfig" type="Deployment:PackageConfiguration"/>
      <xsd:element name="selectRequirement" type="Deployment:Requirement"/>
      <xsd:element name="configProperty" type="Deployment:Property"/>
      <xsd:element name="referencedPackage" type="Deployment:ComponentPackageReference"/>
      <xsd:element name="importedPackage" type="Deployment:ComponentPackageImport"/>
      <xsd:element ref="xmi:Extension"/>
    </xsd:choice>

```

```

<xsd:attribute ref="xmi:id" use="optional"/>
<xsd:attributeGroup ref="xmi:ObjectAttribs"/>
</xsd:complexType>
<xsd:element name="PackageConfiguration" type="Deployment:PackageConfiguration"/>
<xsd:complexType name="Locality">
  <xsd:choice minOccurs="0" maxOccurs="unbounded">
    <xsd:element name="constraint" type="Deployment:LocalityKind"/>
    <xsd:element name="constrainedInstance" type="Deployment:SubcomponentInstantiationDescription"/>
    <xsd:element ref="xmi:Extension"/>
  </xsd:choice>
  <xsd:attribute ref="xmi:id" use="optional"/>
  <xsd:attributeGroup ref="xmi:ObjectAttribs"/>
</xsd:complexType>
<xsd:element name="Locality" type="Deployment:Locality"/>
<xsd:complexType name="RequirementSatisfier">
  <xsd:choice minOccurs="0" maxOccurs="unbounded">
    <xsd:element name="name" type="xsd:string"/>
    <xsd:element name="resourceType" type="xsd:string"/>
    <xsd:element name="property" type="Deployment:SatisfierProperty"/>
    <xsd:element ref="xmi:Extension"/>
  </xsd:choice>
  <xsd:attribute ref="xmi:id" use="optional"/>
  <xsd:attributeGroup ref="xmi:ObjectAttribs"/>
</xsd:complexType>
<xsd:element name="RequirementSatisfier" type="Deployment:RequirementSatisfier"/>
<xsd:complexType name="SatisfierProperty">
  <xsd:choice minOccurs="0" maxOccurs="unbounded">
    <xsd:element name="name" type="xsd:string"/>
    <xsd:element name="kind" type="Deployment:SatisfierPropertyKind"/>
    <xsd:element name="dynamic" type="xsd:string"/>
    <xsd:element name="value" type="Deployment:Any"/>
    <xsd:element ref="xmi:Extension"/>
  </xsd:choice>
  <xsd:attribute ref="xmi:id" use="optional"/>
  <xsd:attributeGroup ref="xmi:ObjectAttribs"/>
</xsd:complexType>
<xsd:element name="SatisfierProperty" type="Deployment:SatisfierProperty"/>
<xsd:complexType name="Requirement">
  <xsd:choice minOccurs="0" maxOccurs="unbounded">
    <xsd:element name="name" type="xsd:string"/>
    <xsd:element name="resourceType" type="xsd:string"/>
    <xsd:element name="property" type="Deployment:Property"/>
    <xsd:element ref="xmi:Extension"/>
  </xsd:choice>
  <xsd:attribute ref="xmi:id" use="optional"/>
  <xsd:attributeGroup ref="xmi:ObjectAttribs"/>
</xsd:complexType>
<xsd:element name="Requirement" type="Deployment:Requirement"/>
<xsd:complexType name="Property">
  <xsd:choice minOccurs="0" maxOccurs="unbounded">

```



```

    <xsd:element name="name" type="xsd:string"/>
    <xsd:element name="value" type="Deployment:Any"/>
    <xsd:element ref="xmi:Extension"/>
  </xsd:choice>
  <xsd:attribute ref="xmi:id" use="optional"/>
  <xsd:attributeGroup ref="xmi:ObjectAttribs"/>
</xsd:complexType>
<xsd:element name="Property" type="Deployment:Property"/>
<xsd:complexType name="ExternalReferenceEndpoint">
  <xsd:choice minOccurs="0" maxOccurs="unbounded">
    <xsd:element name="location" type="xsd:string"/>
    <xsd:element name="provider" type="xsd:string"/>
    <xsd:element name="portName" type="xsd:string"/>
    <xsd:element name="supportedType" type="xsd:string"/>
    <xsd:element ref="xmi:Extension"/>
  </xsd:choice>
  <xsd:attribute ref="xmi:id" use="optional"/>
  <xsd:attributeGroup ref="xmi:ObjectAttribs"/>
</xsd:complexType>
<xsd:element name="ExternalReferenceEndpoint" type="Deployment:ExternalReferenceEndpoint"/>
<xsd:complexType name="ComponentExternalPortEndpoint">
  <xsd:choice minOccurs="0" maxOccurs="unbounded">
    <xsd:element name="portName" type="xsd:string"/>
    <xsd:element ref="xmi:Extension"/>
  </xsd:choice>
  <xsd:attribute ref="xmi:id" use="optional"/>
  <xsd:attributeGroup ref="xmi:ObjectAttribs"/>
</xsd:complexType>
<xsd:element name="ComponentExternalPortEndpoint" type="Deployment:ComponentExternalPortEnd-
point"/>
<xsd:complexType name="ImplementationDependency">
  <xsd:choice minOccurs="0" maxOccurs="unbounded">
    <xsd:element name="requiredType" type="xsd:string"/>
    <xsd:element ref="xmi:Extension"/>
  </xsd:choice>
  <xsd:attribute ref="xmi:id" use="optional"/>
  <xsd:attributeGroup ref="xmi:ObjectAttribs"/>
</xsd:complexType>
<xsd:element name="ImplementationDependency" type="Deployment:ImplementationDependency"/>
<xsd:complexType name="ToplevelPackageDescription">
  <xsd:choice minOccurs="0" maxOccurs="unbounded">
    <xsd:element name="package" type="Deployment:PackageConfiguration"/>
    <xsd:element ref="xmi:Extension"/>
  </xsd:choice>
  <xsd:attribute ref="xmi:id" use="optional"/>
  <xsd:attributeGroup ref="xmi:ObjectAttribs"/>
</xsd:complexType>
<xsd:element name="ToplevelPackageDescription" type="Deployment:ToplevelPackageDescription"/>
<xsd:simpleType name="TCKind">
  <xsd:restriction base="xsd:string">
    <xsd:enumeration value="tk_null"/>
  </xsd:restriction>
</xsd:simpleType>

```

```

<xsd:enumeration value="tk_void"/>
<xsd:enumeration value="tk_short"/>
<xsd:enumeration value="tk_long"/>
<xsd:enumeration value="tk_ushort"/>
<xsd:enumeration value="tk_ulong"/>
<xsd:enumeration value="tk_float"/>
<xsd:enumeration value="tk_double"/>
<xsd:enumeration value="tk_boolean"/>
<xsd:enumeration value="tk_char"/>
<xsd:enumeration value="tk_octet"/>
<xsd:enumeration value="tk_any"/>
<xsd:enumeration value="tk_TypeCode"/>
<xsd:enumeration value="tk_Principal"/>
<xsd:enumeration value="tk_objref"/>
<xsd:enumeration value="tk_struct"/>
<xsd:enumeration value="tk_union"/>
<xsd:enumeration value="tk_enum"/>
<xsd:enumeration value="tk_string"/>
<xsd:enumeration value="tk_sequence"/>
<xsd:enumeration value="tk_array"/>
<xsd:enumeration value="tk_alias"/>
<xsd:enumeration value="tk_except"/>
<xsd:enumeration value="tk_longlong"/>
<xsd:enumeration value="tk_ulonglong"/>
<xsd:enumeration value="tk_longdouble"/>
<xsd:enumeration value="tk_wchar"/>
<xsd:enumeration value="tk_wstring"/>
<xsd:enumeration value="tk_fixed"/>
<xsd:enumeration value="tk_value"/>
<xsd:enumeration value="tk_value_box"/>
<xsd:enumeration value="tk_native"/>
<xsd:enumeration value="tk_abstract_interface"/>
<xsd:enumeration value="tk_local_interface"/>
<xsd:enumeration value="tk_component"/>
<xsd:enumeration value="tk_home"/>
<xsd:enumeration value="tk_event"/>
</xsd:restriction>
</xsd:simpleType>
<xsd:simpleType name="PlanLocalityKind">
  <xsd:restriction base="xsd:string">
    <xsd:enumeration value="SameProcess"/>
    <xsd:enumeration value="DifferentProcess"/>
    <xsd:enumeration value="NoConstraint"/>
  </xsd:restriction>
</xsd:simpleType>
<xsd:simpleType name="ResourceUsageKind">
  <xsd:restriction base="xsd:string">
    <xsd:enumeration value="None"/>
    <xsd:enumeration value="InstanceUsesResource"/>
    <xsd:enumeration value="ResourceUsesInstance"/>
    <xsd:enumeration value="PortUsesResource"/>
  </xsd:restriction>
</xsd:simpleType>

```



```

    <xsd:enumeration value="ResourceUsesPort"/>
  </xsd:restriction>
</xsd:simpleType>
<xsd:simpleType name="LocalityKind">
  <xsd:restriction base="xsd:string">
    <xsd:enumeration value="SameNodeAnyProcess"/>
    <xsd:enumeration value="SameNodeSameProcess"/>
    <xsd:enumeration value="SameNodeDifferentProcess"/>
    <xsd:enumeration value="DifferentNode"/>
    <xsd:enumeration value="DifferentProcess"/>
    <xsd:enumeration value="NoConstraint"/>
  </xsd:restriction>
</xsd:simpleType>
<xsd:simpleType name="CCMComponentPortKind">
  <xsd:restriction base="xsd:string">
    <xsd:enumeration value="Facet"/>
    <xsd:enumeration value="SimplexReceptacle"/>
    <xsd:enumeration value="MultiplexReceptacle"/>
    <xsd:enumeration value="EventEmitter"/>
    <xsd:enumeration value="EventPublisher"/>
    <xsd:enumeration value="EventConsumer"/>
  </xsd:restriction>
</xsd:simpleType>
<xsd:simpleType name="SatisfierPropertyKind">
  <xsd:restriction base="xsd:string">
    <xsd:enumeration value="Quantity"/>
    <xsd:enumeration value="Capacity"/>
    <xsd:enumeration value="Minimum"/>
    <xsd:enumeration value="Maximum"/>
    <xsd:enumeration value="Attribute"/>
    <xsd:enumeration value="Selection"/>
  </xsd:restriction>
</xsd:simpleType>
<xsd:element name="Deployment">
<xsd:complexType>
  <xsd:choice minOccurs="0" maxOccurs="unbounded">
    <xsd:element ref="Deployment:Any"/>
    <xsd:element ref="Deployment:DataType"/>
    <xsd:element ref="Deployment:DataValue"/>
    <xsd:element ref="Deployment:EnumType"/>
    <xsd:element ref="Deployment:ObjrefType"/>
    <xsd:element ref="Deployment:BoundedStringType"/>
    <xsd:element ref="Deployment:FixedType"/>
    <xsd:element ref="Deployment:ArrayType"/>
    <xsd:element ref="Deployment:SequenceType"/>
    <xsd:element ref="Deployment:AliasType"/>
    <xsd:element ref="Deployment:StructType"/>
    <xsd:element ref="Deployment:StructMemberType"/>
    <xsd:element ref="Deployment:ValueType"/>
    <xsd:element ref="Deployment:ValueMemberType"/>
    <xsd:element ref="Deployment:UnionType"/>
  </xsd:choice>
</xsd:complexType>

```

```

<xsd:element ref="Deployment:UnionMemberType"/>
<xsd:element ref="Deployment:NamedValue"/>
<xsd:element ref="Deployment:Bridge"/>
<xsd:element ref="Deployment:Interconnect"/>
<xsd:element ref="Deployment:Node"/>
<xsd:element ref="Deployment:Resource"/>
<xsd:element ref="Deployment:SharedResource"/>
<xsd:element ref="Deployment:Domain"/>
<xsd:element ref="Deployment:ResourceAllocation"/>
<xsd:element ref="Deployment:PlanPropertyMapping"/>
<xsd:element ref="Deployment:PlanSubcomponentPropertyReference"/>
<xsd:element ref="Deployment:PlanSubcomponentPortEndpoint"/>
<xsd:element ref="Deployment:PlanConnectionDescription"/>
<xsd:element ref="Deployment:InstanceDeploymentDescription"/>
<xsd:element ref="Deployment:MonolithicDeploymentDescription"/>
<xsd:element ref="Deployment:ArtifactDeploymentDescription"/>
<xsd:element ref="Deployment:DeploymentPlan"/>
<xsd:element ref="Deployment:ResourceDeploymentDescription"/>
<xsd:element ref="Deployment:InstanceResourceDeploymentDescription"/>
<xsd:element ref="Deployment:ConnectionResourceDeploymentDescription"/>
<xsd:element ref="Deployment:PlanLocality"/>
<xsd:element ref="Deployment:Capability"/>
<xsd:element ref="Deployment:ComponentPropertyDescription"/>
<xsd:element ref="Deployment:ComponentPortDescription"/>
<xsd:element ref="Deployment:ComponentInterfaceDescription"/>
<xsd:element ref="Deployment:ImplementationArtifactDescription"/>
<xsd:element ref="Deployment:MonolithicImplementationDescription"/>
<xsd:element ref="Deployment:AssemblyPropertyMapping"/>
<xsd:element ref="Deployment:SubcomponentPropertyReference"/>
<xsd:element ref="Deployment:SubcomponentPortEndpoint"/>
<xsd:element ref="Deployment:AssemblyConnectionDescription"/>
<xsd:element ref="Deployment:SubcomponentInstantiationDescription"/>
<xsd:element ref="Deployment:ComponentAssemblyDescription"/>
<xsd:element ref="Deployment:ComponentImplementationDescription"/>
<xsd:element ref="Deployment:ComponentPackageReference"/>
<xsd:element ref="Deployment:ComponentPackageDescription"/>
<xsd:element ref="Deployment:ComponentUsageDescription"/>
<xsd:element ref="Deployment:PackagedComponentImplementation"/>
<xsd:element ref="Deployment:NamedImplementationArtifact"/>
<xsd:element ref="Deployment:ImplementationRequirement"/>
<xsd:element ref="Deployment:ComponentPackageImport"/>
<xsd:element ref="Deployment:PackageConfiguration"/>
<xsd:element ref="Deployment:Locality"/>
<xsd:element ref="Deployment:RequirementSatisfier"/>
<xsd:element ref="Deployment:SatisfierProperty"/>
<xsd:element ref="Deployment:Requirement"/>
<xsd:element ref="Deployment:Property"/>
<xsd:element ref="Deployment:ExternalReferenceEndpoint"/>
<xsd:element ref="Deployment:ComponentExternalPortEndpoint"/>
<xsd:element ref="Deployment:ImplementationDependency"/>
<xsd:element ref="Deployment:ToplevelPackageDescription"/>

```

```
<xsd:element ref="xmi:Extension"/>
</xsd:choice>
<xsd:attribute ref="xmi:id" use="optional"/>
<xsd:attributeGroup ref="xmi:ObjectAttribs"/>
</xsd:complexType>
</xsd:element>
</xsd:schema>
```


Annex A - Acknowledgements

(normative)

A.1 List of Submitters/Supporters

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 Computing, 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
- THALES
- TIBCO, Inc.
- Tivoli Systems, Inc.
- Tri-Pacific Software, Inc.
- University of California, Santa Barbara
- University of Rhode Island
- Visual Edge Software, Ltd.
- Washington University

