

FOUNDATION FOR INTELLIGENT PHYSICAL AGENTS

FIPA Abstract Architecture Specification

Document title	FIPA Abstract Architecture Specification		
Document number	XC00001J	Document source	FIPA TC Architecture
Document status	Experimental	Date of this status	2002/25/02
Supersedes	None		
Contact	fab@fipa.org		
Change history			
2001/25/02	See <i>Informative Annex E — Change-Log</i>		

© 2002 Foundation for Intelligent Physical Agents - <http://www.fipa.org/>

Geneva, Switzerland

Notice

Use of the technologies described in this specification may infringe patents, copyrights or other intellectual property rights of FIPA Members and non-members. Nothing in this specification should be construed as granting permission to use any of the technologies described. Anyone planning to make use of technology covered by the intellectual property rights of others should first obtain permission from the holder(s) of the rights. FIPA strongly encourages anyone implementing any part of this specification to determine first whether part(s) sought to be implemented are covered by the intellectual property of others, and, if so, to obtain appropriate licenses or other permission from the holder(s) of such intellectual property prior to implementation. This specification is subject to change without notice. Neither FIPA nor any of its Members accept any responsibility whatsoever for damages or liability, direct or consequential, which may result from the use of this specification.

19 **Foreword**

20 The Foundation for Intelligent Physical Agents (FIPA) is an international organization that is dedicated to promoting the
21 industry of intelligent agents by openly developing specifications supporting interoperability among agents and agent-
22 based applications. This occurs through open collaboration among its member organizations, which are companies
23 and universities that are active in the field of agents. FIPA makes the results of its activities available to all interested
24 parties and intends to contribute its results to the appropriate formal standards bodies.

25 The members of FIPA are individually and collectively committed to open competition in the development of agent-
26 based applications, services and equipment. Membership in FIPA is open to any corporation and individual firm,
27 partnership, governmental body or international organization without restriction. In particular, members are not bound
28 to implement or use specific agent-based standards, recommendations and FIPA specifications by virtue of their
29 participation in FIPA.

30 The FIPA specifications are developed through direct involvement of the FIPA membership. The status of a
31 specification can be Preliminary, Experimental, Standard, Deprecated or Obsolete. More detail about the process of
32 specification may be found in the FIPA Procedures for Technical Work. A complete overview of the FIPA specifications
33 and their current status may be found in the FIPA List of Specifications. A list of terms and abbreviations used in the
34 FIPA specifications may be found in the FIPA Glossary.

35 FIPA is a non-profit association registered in Geneva, Switzerland. As of January 2000, the 56 members of FIPA
36 represented 17 countries worldwide. Further information about FIPA as an organization, membership information, FIPA
37 specifications and upcoming meetings may be found at <http://www.fipa.org/>.

38 **Contents**

39	1	Introduction.....	1
40	1.1	Contents.....	1
41	1.2	Audience.....	1
42	1.3	Acknowledgements.....	2
43	2	Scope and Methodology.....	3
44	2.1	Background.....	3
45	2.2	Why an Abstract Architecture?.....	4
46	2.3	Scope of the Abstract Architecture.....	4
47	2.3.1	Areas that are not Sufficiently Abstract.....	5
48	2.3.2	Areas for Future Consideration.....	5
49	2.4	Going From Abstract to Concrete Specifications.....	5
50	2.5	Methodology.....	7
51	2.6	Status of the Abstract Architecture.....	8
52	2.7	Evolution of the Abstract Architecture.....	8
53	3	Themes of the Abstract Architecture.....	9
54	3.1	Focus on Agent Interoperability.....	10
55	3.2	An Exemplar System.....	10
56	4	Architectural Overview.....	11
57	4.1	Agents and Services.....	11
58	4.2	Starting an Agent.....	11
59	4.3	Agent Directory Services.....	11
60	4.3.1	Registering an Agent.....	12
61	4.3.2	Discovering an Agent.....	12
62	4.4	Service Directory Services.....	13
63	4.5	Agent Messages.....	13
64	4.5.1	Message Structure.....	13
65	4.5.2	Message Transport.....	14
66	4.6	Agents Send Messages to Other Agents.....	14
67	4.7	Providing Message Validity and Encryption.....	16
68	4.8	Providing Interoperability.....	17
69	5	Architectural Elements.....	18
70	5.1	Introduction.....	18
71	5.1.1	Classification of Elements.....	18
72	5.1.2	Key-Value Tuples.....	18
73	5.1.3	Services.....	20
74	5.1.4	Format of Element Description.....	20
75	5.1.5	Abstract Elements.....	20
76	5.2	Agent.....	22
77	5.2.1	Summary.....	22
78	5.2.2	Relationships to Other Elements.....	22
79	5.2.3	Description.....	23
80	5.3	Agent Attribute.....	23
81	5.3.1	Summary.....	23
82	5.3.2	Relationships to Other Elements.....	23
83	5.3.3	Description.....	23
84	5.4	Agent Communication Language.....	23
85	5.4.1	Summary.....	23
86	5.4.2	Relationships to Other Elements.....	23
87	5.4.3	Description.....	23
88	5.5	Agent Directory Entry.....	24
89	5.5.1	Summary.....	24
90	5.5.2	Relationships to Other Elements.....	24

91	5.5.3	Description	24
92	5.6	Agent Directory Service	24
93	5.6.1	Summary	24
94	5.6.2	Relationships to Other Elements	24
95	5.6.3	Actions	24
96	5.6.4	Description	26
97	5.7	Agent Locator	27
98	5.7.1	Summary	27
99	5.7.2	Relationships to Other Elements	27
100	5.7.3	Description	27
101	5.8	Agent Name	27
102	5.8.1	Summary	27
103	5.8.2	Relationships to Other Elements	27
104	5.8.3	Description	28
105	5.9	Content	28
106	5.9.1	Summary	28
107	5.9.2	Relationships to Other Elements	28
108	5.9.3	Description	28
109	5.10	Content Language	28
110	5.10.1	Summary	28
111	5.10.2	Relationships to Other Elements	29
112	5.10.3	Description	29
113	5.11	Encoding Representation	29
114	5.11.1	Summary	29
115	5.11.2	Relationships to Other Elements	29
116	5.11.3	Description	29
117	5.12	Encoding Service	29
118	5.12.1	Summary	29
119	5.12.2	Relationships to Other Elements	29
120	5.12.3	Actions	30
121	5.12.4	Description	31
122	5.13	Envelope	31
123	5.13.1	Summary	31
124	5.13.2	Relationship to Other Elements	31
125	5.13.3	Description	31
126	5.14	Explanation	31
127	5.14.1	Summary	31
128	5.14.2	Relationship to Other Elements	31
129	5.14.3	Description	32
130	5.15	Message	32
131	5.15.1	Summary	32
132	5.15.2	Relationships to other elements	32
133	5.15.3	Description	32
134	5.16	Message Transport Service	32
135	5.16.1	Summary	32
136	5.16.2	Relationships to Other Elements	32
137	5.16.3	Actions	32
138	5.16.4	Description	34
139	5.17	Ontology	34
140	5.17.1	Summary	34
141	5.17.2	Relationships to Other Elements	34
142	5.17.3	Description	34
143	5.18	Payload	35
144	5.18.1	Summary	35
145	5.18.2	Relationships to Other Elements	35
146	5.18.3	Description	35

147	5.19	Service.....	35
148	5.19.1	Summary.....	35
149	5.19.2	Relationships to Other Elements	35
150	5.19.3	Description	35
151	5.20	Service Address	35
152	5.20.1	Summary.....	35
153	5.20.2	Relationships to Other Elements	36
154	5.20.3	Description	36
155	5.21	Service Attributes	36
156	5.21.1	Summary.....	36
157	5.21.2	Relationships to Other Elements	36
158	5.21.3	Description	36
159	5.22	Service Directory Entry.....	36
160	5.22.1	Summary.....	36
161	5.22.2	Relationships to Other Elements	36
162	5.22.3	Description	37
163	5.23	Services Directory Service	37
164	5.23.1	Summary.....	37
165	5.23.2	Relationships to Other Elements	37
166	5.23.3	Description	37
167	5.23.4	Actions	37
168	5.24	Service Id.....	39
169	5.24.1	Summary.....	39
170	5.24.2	Relationships to other elements	40
171	5.24.3	Description	40
172	5.25	Service Location Description.....	40
173	5.25.1	Summary.....	40
174	5.25.2	Relationships to Other Elements	40
175	5.25.3	Description	40
176	5.26	Service Locator.....	40
177	5.26.1	Summary.....	40
178	5.26.2	Relationships to Other Elements	40
179	5.26.3	Description	41
180	5.27	Service Root	41
181	5.27.1	Summary.....	41
182	5.27.2	Relationships to Other Elements	41
183	5.27.3	Description	41
184	5.28	Service Signature	41
185	5.28.1	Summary.....	41
186	5.28.2	Relationships to Other Elements	41
187	5.28.3	Description	41
188	5.29	Service Type.....	42
189	5.29.1	Summary.....	42
190	5.29.2	Relationships to Other Elements	42
191	5.29.3	Description	42
192	5.30	Signature Type	42
193	5.30.1	Summary.....	42
194	5.30.2	Relationships to Other Elements	42
195	5.30.3	Description	42
196	5.31	Transport	42
197	5.31.1	Summary.....	42
198	5.31.2	Relationships to Other Elements	43
199	5.31.3	Description	43
200	5.32	Transport Description	43
201	5.32.1	Summary.....	43
202	5.32.2	Relationships to Other Elements	43

203	5.32.3	Description	43
204	5.33	Transport Message.....	43
205	5.33.1	Summary.....	43
206	5.33.2	Relationships to Other Elements	43
207	5.33.3	Description	43
208	5.34	Transport Specific Address	44
209	5.34.1	Summary.....	44
210	5.34.2	Relationships to Other Elements	44
211	5.34.3	Description	44
212	5.35	Transport Specific Property.....	44
213	5.35.1	Summary.....	44
214	5.35.2	Relationships to Other Elements	44
215	5.35.3	Description	44
216	5.36	Transport Type	44
217	5.36.1	Summary.....	44
218	5.36.2	Relationships to Other Elements	44
219	5.36.3	Description	45
220	6	Agent and Agent Information Model.....	46
221	6.1	Agent Relationships	46
222	6.2	Transport Message Relationships	47
223	6.3	Agent Directory Entry Relationships	48
224	6.4	Service Directory Entry Relationships.....	49
225	6.5	Message Elements.....	50
226	6.6	Message Transport Elements	51
227	7	References	52
228	8	Informative Annex A — Goals of Service Model.....	53
229	8.1	Scope	53
230	8.2	Variety of Services	53
231	8.3	Bootstrapping	53
232	8.4	Dynamic services	53
233	8.5	Granularity.....	53
234	8.6	Example	53
235	9	Informative Annex B — Goals of Message Transport Service Abstractions.....	55
236	9.1	Scope	55
237	9.2	Variety of Transports.....	55
238	9.3	Support for Alternative Transports within a Single System.....	55
239	9.4	Desirability of Transport Agnosticism.....	56
240	9.5	Desirability of Selective Specificity.....	56
241	9.6	Connection-Based, Connectionless and Store-and-Forward Transports.....	56
242	9.7	Conversation Policies and Interaction Protocols	56
243	9.8	Point-to-Point and Multiparty Interactions	56
244	9.9	Durable Messaging	57
245	9.10	Quality of Service	57
246	9.11	Anonymity.....	57
247	9.12	Message Encoding.....	57
248	9.13	Interoperability and Gateways.....	58
249	9.14	Reasoning about Agent Communications.....	58
250	9.15	Testing, Debugging and Management	58
251	10	Informative Annex C — Goals of Directory Service Abstractions	59
252	10.1	Scope.....	59
253	10.2	Variety of Directory Services	59
254	10.3	Desirability of Directory Agnosticism	59
255	10.4	Desirability of Selective Specificity	60
256	10.5	Interoperability and Gateways.....	60
257	10.6	Reasoning about Agent Directory	60
258	10.7	Testing, Debugging and Management	60

259	11	Informative Annex D — Goals for Security and Identity Abstractions	61
260	11.1	Introduction	61
261	11.2	Overview.....	61
262	11.3	Areas to Apply Security	61
263	11.3.1	Content Validity and Privacy During Message Transport.....	61
264	11.3.2	Agent Identity	62
265	11.3.3	Agent Principal Validation	62
266	11.3.4	Code Signing Validation.....	63
267	11.4	Risks Not Addressed	63
268	11.4.1	Code or Data Peeping	63
269	11.4.2	Code or Data Alteration	63
270	11.4.3	Concerted Attacks.....	63
271	11.4.4	Copy and Replay	63
272	11.4.5	Denial of Service.....	63
273	11.4.6	Misinformation Campaigns	63
274	11.4.7	Repudiation.....	64
275	11.4.8	Spoofing and Masquerading.....	64
276	11.5	Glossary of Security Terms	64
277	12	Informative Annex E — Change-Log.....	66
278	12.1	2001/11/01 - change delta with respect to XC00001J	66

279 1 Introduction

280 This document, and the specifications that are derived from it, defines the FIPA Abstract Architecture. The parts of the
281 FIPA abstract architecture include:

282
283 A specification that defines architectural elements and their relationships (this document).
284

285 Guidelines for the specification of agent systems in terms of particular software and communications technologies
286 (Guidelines for Instantiation).
287

288 Specifications governing the interoperability and conformance of agents and agent systems (Interoperability
289 Guidelines).
290

291 Note that the latter two documents are not yet available.

292
293 See *Section 2, Scope and Methodology* for a fuller introduction to this document.
294

295 1.1 Contents

296 This document is organized into the following sections and a series of annexes.

297
298 This **Introduction**.

299
300 The **Scope and methodology** section explains the background of this work, its purpose, and the methodology that
301 has been followed. It describes the role of this work in the overall FIPA work program and discusses both the
302 current status of the work and way in which the document is expected to evolve.
303

304 The **Themes of the Abstract Architecture** section that explains the style and the themes of the Abstract
305 Architecture specification.
306

307 The **Architectural overview** presents an overview of the architecture with some examples. It is intended to
308 provide the appropriate context for understanding the subsequent sections.
309

310 The **Architectural Elements** section comprises the FIPA architecture components.

311
312 The **Agent and Agent Information Model** defines UML pattern relationships between **Architectural Elements**.
313

314 The annexes include:

315
316 **Goals of Service Model**

317
318 **Goals of Message Transport Service Abstractions**

319
320 **Goals of Directory Service Abstractions.**

321
322 **Goals for Security and Identity Abstractions.**
323

324 1.2 Audience

325 The primary audience for this document is developers of concrete specifications for agent systems – specifications
326 grounded in particularly technologies, representations, and programming models. It may also be read by the users of
327 these concrete specifications, including implementers of agent platforms, agent systems, and gateways between agent
328 systems.
329

330 This document describes an abstract architecture for creating intentional multi-agent systems. It assumes that the
331 reader has a good understanding about the basic principles of multi-agent systems. It does not provide the background
332 material to help the reader assess whether multi-agent systems are an appropriate model for their system design, nor
333 does it provide background material on topics such as Agent Communication Languages, BDI systems, or distributed
334 computing platforms.

335 The abstract architecture described in this document will guide the creation of concrete specifications of different
336 elements of the FIPA agent systems. The developers of the concrete specifications must ensure that their work
337 conform to the abstract architecture in order to provide specifications with appropriate levels of interoperability.
338 Similarly, those specifying applications that will run on FIPA compliant agent systems will need to understand what
339 services and features that they can use in the creation of their applications.
340

341 **1.3 Acknowledgements**

342 This document was developed by members of FIPA TC A, the Technical Committee of FIPA charged with this work.
343 Other FIPA Technical Committees also made substantial contributions to this effort, and we thank them for their effort
344 and assistance.
345

346 2 Scope and Methodology

347 This section provides a context for the Abstract Architecture, the scope of the work and methodology employed.
348

349 2.1 Background

350 FIPA's goal in creating agent standards is to promote inter-operable agent applications and agent systems. In 1997
351 and 1998, FIPA issued a series of agent system specifications that had as their goal inter-operable agent systems.
352 This work included specifications for agent infrastructure and agent applications. The infrastructure specifications
353 included an agent communication language, agent services, and supporting management ontologies. There were also
354 a number of application domains specified, such as personal travel assistance and network management and
355 provisioning.

356
357 At the heart FIPA's model for agent systems is agent communication, where agents can pass semantically meaningful
358 messages to one another in order to accomplish the tasks required by the application. In 1998 and 1999 it became
359 clear that it would be useful to support variations in those messages:

360 How those messages are transferred (that is, the transport).

361 How those messages are represented (e.g. s-expressions, bit-efficient binary objects, XML).

362 Optional attributes of those messages, such as how to authenticate or encrypt them.

363
364
365 It also became clear that to create agent systems, which could be deployed in commercial settings, it was important to
366 understand and to use existing software environments. These environments included elements such as:

367 Distributed computing platforms or programming languages,

368 Messaging platforms,

369 Security services,

370 Directory services, and,

371 Intermittent connectivity technologies.

372
373 FIPA was faced with two choices: to incrementally revise specifications to add various features such as intermittent
374 connectivity, or to take a more holistic approach. The holistic approach, which FIPA adopted in January of 1999, was
375 to create an architecture that could accommodate a wide range of commonly used mechanisms, such as various
376 message transports, directory services and other commonly, commercially available development platforms. For
377 detailed discussions of the goals of the architecture, see:

378 *Section 8, Informative Annex A — Goals of Service Model*

379 *Section 9, Informative Annex B — Goals of Message Transport Service Abstraction*

380 *Section 10, Informative Annex C — Goals of Directory Service Abstractions*

381 *Section 11, Informative Annex D — Goals for Security and Identity Abstractions*

382
383 These describe in greater detail the design considerations that were considered when creating this abstract
384 architecture. In addition, FIPA needed to consider the relationship between the existing FIPA 97, FIPA 98 and FIPA
385 2000 work and the abstract architecture. While more validation is required, the FIPA 2000 work is in part a concrete
386 realization of this abstract architecture. While one of the goals in creating this architecture was to maintain full
387
388
389
390
391
392
393
394
395
396
397

398 compatibility with the FIPA 97 and 98 specifications, this was not entirely feasible, especially when trying to support
399 multiple implementations.

400 Agent systems built according to FIPA 97 and 98 specifications will be able to inter-operate with agent systems built
401 according to the abstract architecture through transport gateways with some limitations. The FIPA 2000 architecture is
402 a closer match to the abstract architecture, and will be able to fully inter-operate via gateways. The overall goal in this
403 architectural approach is to permit the creation of systems that seamlessly integrate within their specific computing
404 environment while interoperating with agent systems residing in separate environments.
405

406 **2.2 Why an Abstract Architecture?**

407 The first purpose of this work is to foster interoperability and reusability. To achieve this, it is necessary to identify the
408 elements of the architecture that must be codified. Specifically, if two or more systems use different technologies to
409 achieve some functional purpose, it is necessary to identify the common characteristics of the various approaches.
410 This leads to the identification of *architectural abstractions*: abstract designs that can be formally related to every valid
411 implementation.
412

413 By describing systems abstractly, one can explore the relationships between fundamental elements of these agent
414 systems. By describing the relationships between these elements, it becomes clearer how agent systems can be
415 created so that they are interoperable. From this set of architectural elements and relations one can derive a broad set
416 of possible concrete architectures, which will interoperate because they share a common abstract design.
417

418 Because the abstract architecture permits the creation of multiple concrete realizations, it must provide mechanisms to
419 permit them to interoperate. This includes providing transformations for both transport and encodings, as well as
420 integrating these elements with the basic elements of the environment.
421

422 For example, one agent system may transmit ACL messages using the OMG IIOP protocol. A second may use IBM's
423 MQ-series enterprise messaging system. An analysis of these two systems – how senders and receivers are identified,
424 and how messages are encoded and transferred – allows us to arrive at a series of architectural abstractions involving
425 messages, encodings, and addresses.
426

427 **2.3 Scope of the Abstract Architecture**

428 The primary focus of this abstract architecture is to create semantically meaningful message exchange between
429 agents which may be using different messaging transports, different Agent Communication Languages, or different
430 content languages. This requires numerous points of potential interoperability. The scope of this architecture includes:

431 A model of services and discovery of services available to agents and other services.

432
433 Message transport interoperability.

434
435 Supporting various forms of ACL representations.

436
437 Supporting various forms of content language.

438
439 Supporting multiple directory services representations.
440

441 It must be possible to create implementations that vary in some of these attributes, but which can still interoperate.
442 Some aspects of potential standardization are outside of the scope of this architecture. There are three different
443 reasons why things are out of scope:
444

445
446 The area cannot be described abstractly.

447
448 The area is not yet ready for standardization, or there was not yet sufficient agreement about how to standardize it.
449

450 The area is sufficiently specialized that it does not currently need standardization.

451

452 Some of the key areas that are **not** included in this architecture are:

453

454 Agent lifecycle and management.

455

456 Agent mobility.

457

458 Domains.

459

460 Conversational policy.

461

462 Agent Identity.

463

464 The next sections describe the rationale for this in more detail. However, it is extremely important to understand that
465 the abstract architecture does not prohibit additional features – it merely addresses how interoperable features should
466 be implemented. It is anticipated that over time some of these areas will be part of the interoperability of agent
467 systems.

468

469 2.3.1 Areas that are not Sufficiently Abstract

470 An abstraction may not appear in the abstract architecture because is there is no clean abstraction for different models
471 of implementation. Two examples of this are agent lifecycle management and security related issues.

472

473 For example, in examining agent lifecycle, it seems clear there are a minimum set of features that are required:
474 Starting an agent, stopping an agent, "freezing" or "suspending" an agent, and "unfreezing" or "restarting" an agent. In
475 practice, when one examines how various software systems work, very little consistency is detected inside the
476 mechanisms, or in how to address and use those mechanisms. Although it is clear that concrete specifications will
477 have to address these issues, it is not clear how to provide a unifying abstraction for these features. Therefore there
478 are some architectural elements that can only appear at the concrete level, because the details of different
479 environments are so diverse.

480

481 Security has similar issues, especially when trying to provide security in the transport layer, or when trying to provide
482 security for attacks that can occur because a particular software environment has characteristics that permits that sort
483 of attack. Agent mobility is another implementation specific model that cannot easily be modelled abstractly.

484

485 Both of these topics will be addressed in the *Instantiation Guidelines*, because they are an important part of how agent
486 systems are created. However, they cannot be modelled abstractly, and are therefore not included at the *abstract* level
487 of the architecture.

488

489 2.3.2 Areas for Future Consideration

490 FIPA may address a number of areas of agent standardization in the future. These include ontologies, domains,
491 conversational policies and mechanisms that are used to control systems (resource allocation and access control
492 lists), and agent identity. These all represent ideas requiring further development.

493

494 This architecture does not address application interoperability. The current model for application interoperability is that
495 agents that communicate using a shared set of semantics (such as represented by an ontology) can potentially
496 interoperate. This architecture does not extend this model any further.

497

498 2.4 Going From Abstract to Concrete Specifications

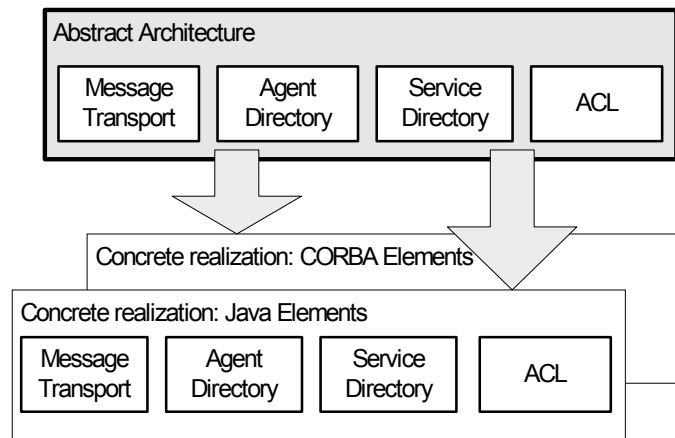
499 This document describes an abstract architecture. Such an architecture cannot be directly implemented, but instead
500 the forms the basis for the development of concrete architectural specifications. Such specifications describe in precise
501 detail how to construct an agent system, including the agents and the services that they rely upon, in terms of concrete

502 software artefacts, such as programming languages, applications programming interfaces, network protocols,
 503 operating system services, and so forth.

504
 505 In order for a concrete architectural specification to be FIPA compliant, it must have certain properties. First, the
 506 concrete architecture must include mechanisms for agent registration and agent discovery and inter-agent message
 507 transfer. These services must be explicitly described in terms of the corresponding elements of the FIPA abstract
 508 architecture. The definition of an abstract architectural element in terms of the concrete architecture is termed a
 509 *realization* of that element; more generally, a concrete architecture will be said to *realize* all or part of an abstraction.

510
 511 The designer of the concrete architecture has considerable latitude in how he or she chooses to realize the abstract
 512 elements. If the concrete architecture provides only one encoding for messages, or only one transport protocol, the
 513 realization may simplify the programmatic view of the system. Conversely, a realization may include additional options
 514 or features that require the developer to handle both abstract and platform-specific elements. That is to say that the
 515 existence of an abstract architecture does not *prohibit* the introduction of elements useful to make a good agent
 516 system, it merely sets out the *minimum* required elements.

517



518

519

520

521

Figure 1: Abstract Architecture Mapped to Various Concrete Realizations

522 The abstract architecture also describes *optional* elements. Although an element is optional at the abstract level, it may
 523 be *mandatory* in a particular realization. That is, a realization may require the existence of an entity that is optional at
 524 the abstract level (such as a **message-transport-service**), and further specify the features and interfaces that the
 525 element must have in that realization.

526

527 It is also important to note that a realization can be of the entire architecture, or just one element. For example, a
 528 series of concrete specifications could be created that describe how to represent the architecture in terms of particular
 529 programming language, coupled to a sockets-based message transport. Messages are handled as objects with that
 530 language, and so on.

531

532 On the other hand, there may be a single element that can be defined concretely, and then used in a number of
 533 different systems. For example, if a concrete specification were created for the **agent-directory-service** element that
 534 describes the schemas to use when implemented in LDAP, that particular element might appear in a number of
 535 different agent systems.

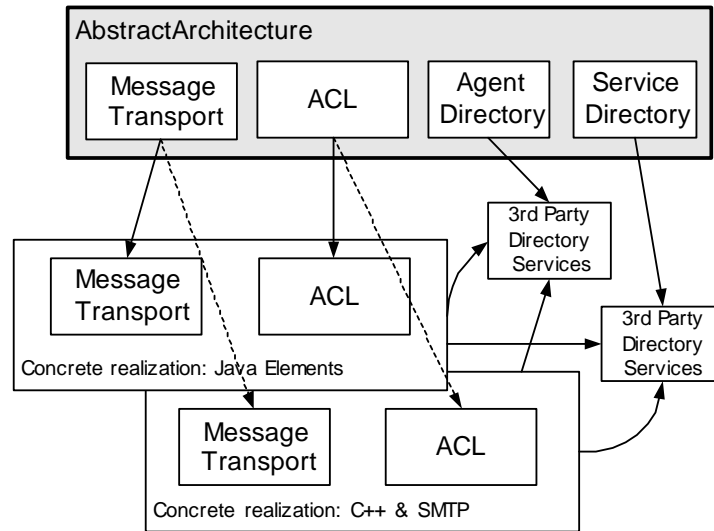


Figure 2: Concrete Realizations Using a Shared Element Realization

In this example, the concrete realization of directory is to implement the directory services in LDAP. Several realizations have chosen to use this directory service model.

2.5 Methodology

This abstract architecture was created by the use of UML modelling, combined with the notions of design patterns as described in [Gamma95]. Analysis was performed to consider a variety ways of structuring software and communications components in order to implement the features of an intelligent multi-agent system. This ideal agent system was to be capable of exhibiting execution autonomy and semantic interoperability based on an intentional stance. The analysis drew upon many sources:

- The abstract notions of agency and the design features that flow from this.

- Commercial software engineering principles, especially object-oriented techniques, design methodologies, development tools and distributed computing models.

- Requirements drawn from a variety of applications domains.

- Existing FIPA specifications and implementations.

- Agent systems and services, including FIPA and non-FIPA designs.

- Commercially important software systems and services, such as Java, CORBA, DCOM, LDAP, X.500 and MQ Series.

The primary purpose of this work is to foster interoperability and reusability. To achieve this, it is necessary to identify the elements of the architecture that must be codified. Specifically, if two or more systems use different technologies to achieve some functional purpose, it is necessary to identify the common characteristics of the various approaches. This leads to the identification of *architectural elements*: abstract designs that can be formally related to every valid implementation.

For example, one agent system may transmit ACL messages using the OMG IIOP protocol. A second may use IBM's MQ-series enterprise messaging system. An analysis of these two systems – how senders and receivers are identified, and how messages are encoded and transferred – allows us to arrive at a series of architectural abstractions involving messages, encodings, and addresses.

574

575 In some areas, the identification of common abstractions is essential for successful interoperation. This is particularly
576 true for agent-to-agent message transfer. The end-to-end support of a common agent communication language is at
577 the core of FIPA's work. These essential elements, which correspond to mandatory implementation specifications are
578 here described as *mandatory architectural elements*. Other areas are less straightforward. Different software systems,
579 particularly different types of commercial middleware systems, have specialized frameworks for software deployment,
580 configuration, and management, and it is hard to find common principles. For example, security and identity remain
581 tend to be highly dependent on implementation platforms. Such areas will eventually be the subjects of architectural
582 specification, but not all systems will support them. These architectural elements are *optional*.

583

584 This document models the elements and their relationships. In *Section 3, Themes of the Abstract Architecture* there is
585 an holistic overview of the architecture. In *Section 4, Architectural Overview* there is a structural overview of the
586 architecture. In *Section 5, Architectural Elements*, each of the architectural elements is described. In *Section 6, Agent
587 and Agent Information Model* there are diagrams in UML notation to describe the relationships between the elements.
588

589 **2.6 Status of the Abstract Architecture**

590 There are several steps in creating the abstract architecture:

591

592 1. Modelling of the abstract elements and their relationships.

593

594 2. Representing the other requirements on the architecture that cannot be modelled abstractly.

595

596 3. Describing interoperability points.

597

598 This document represents the first item in the list. It is nearing completion, and ready for review.

599

600 The second step is satisfied by *guidelines for instantiation*. This document will not be written until at least one
601 implementation based on the abstract architecture has been created, as it is desirable to base such a document on
602 actual implementation experience.

603

604 Interoperability points and conformance are defined by specific *interoperability profiles*. These profiles will be created
605 as required during the creation of concrete specifications.

606

607 **2.7 Evolution of the Abstract Architecture**

608 One of the challenges involved in creating this specification was drawing the line between elements that belong in the
609 abstract architecture and those which belong in concrete instantiations of the architecture. As FIPA creates several
610 concrete specifications, and explores the mechanisms required to properly manage interoperation of these
611 implementations, some features of the concrete architectures may be abstracted and incorporated in the FIPA abstract
612 architecture. Likewise, certain abstract architectural elements may eventually be dropped from the abstract
613 architecture, but may continue to exist in the form of concrete realizations.

614

615 The current placement of various elements as mandatory or optional is somewhat tentative. It is possible that some
616 elements that are currently optional will, upon further experience in the development of the architecture become
617 mandatory.

618

3 Themes of the Abstract Architecture

The overall approach of the abstract architecture is deeply rooted in object-oriented design, including the use of design patterns and UML modelling. As such, the natural way to envision the elements of the architecture is as a set of abstract object classes that can act as the input to the high level design of specific implementations.

Although the architecture explicitly avoids any specific model of composing its elements, its natural expression is a set of object classes comprising an agent platform that supports agents and services.

The following diagram depicts the hierarchical relationships between the abstraction defined by this document and the elements of a specific instantiation:

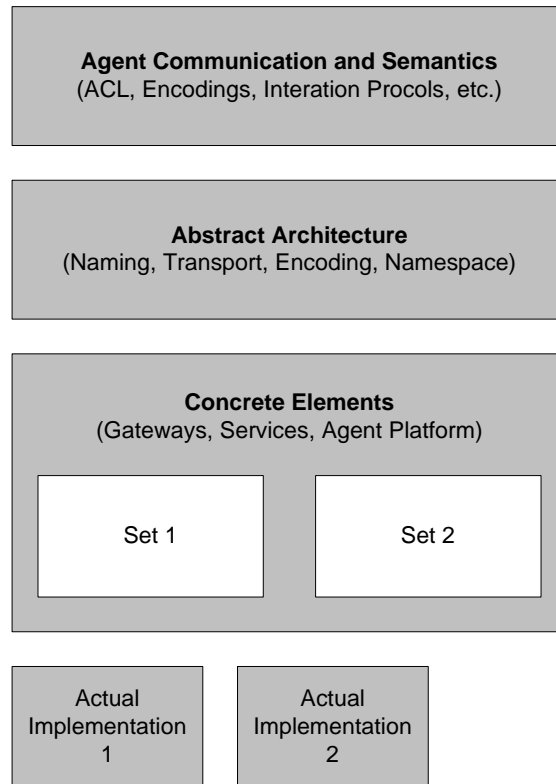


Figure 3: Relationship between Abstract and Concrete Architecture Elements

630

631

632

633

634

635

636

637

638

639

640

641

642

643

644

645

646

647

648

Several themes pervade the architecture; these capture the interaction between elements and their intended use.

The first theme is of opaque typed elements, which can be understood by specific implementations of a service. For example, the details of each transport description are opaque to other layers of the system. The transport descriptor provides a transport type, such as *fipa-tcpip-raw-socket* which acts to select the specific transport service that can interpret the transport-specific-address. Thus, a given address element, opaque to other portions of the system, might be *foo.bar.baz.com:1234* which would be readily understood by the above transport service. Opaque typed elements are used in both message encoding and directory services.

This theme leads to an elegant solution for extensibility. Additional implementations of a service may be dynamically added to an environment by defining a new opaque typed element and associating it with the new service. For example, it may be required that a transport mechanism such as the Simple Object Access Protocol (SOAP) be supported within the environment. The transport type ontology would be extended to include a new term, *fipa-soap-v1*. Note that this resembles a polymorphic type scheme.

649 A second repeated theme is the creation of an association (in the form of a contract) between an agent and a service,
650 such that the agent may then use the service through a returned handle. Note that this theme is intentionally well
651 suited for implementation through the factory design patterns.

652
653 For those familiar with the "design pattern" approach to describing system structure, these themes may be naturally
654 implemented using the factory pattern.
655

656 **3.1 Focus on Agent Interoperability**

657 The Abstract Architecture focuses on core interoperability between agents. These include:

658 Managing multiple message transport schemes,

659 Managing message encoding schemes, and,

660 Locating agents and services via directory services.
661
662

663
664
665 The Abstract Architecture explicitly avoids issues internal to the structure of an agent. It also largely defers details of
666 agent services to more concrete architecture documents.
667

668 After reading through the abstract architecture, many readers may feel that it lacks a number of elements they would
669 have expected to be included. Examples include the notion of an "agent-platform," "gateways" between agent systems,
670 bootstrapping of agent systems and agent configuration and coordination.
671

672 These elements are not included in the abstract architecture because they are inherently coupled with specific
673 implementations of the architecture, rather than across all possible implementations. The forthcoming document
674 "Concrete Architectural Elements" will describe many of these elements in terms of specific environments. Beyond this,
675 some elements will exist only in specific instantiations.
676

677 **3.2 An Exemplar System**

678 In order to further illuminate the intended use of the architectural elements, let us consider an agent platform,
679 implemented in an object oriented environment. The system uses the components of the abstract architecture to
680 implement two separate object factories; a transport factory and an encoding factory. A directory service is also
681 provided, with access through a static object.
682

683 Agents in the environment are constructed as objects, each running on a permanent thread. Each has access to the
684 two agent factories, as well as the directory service.
685

686 When an agent wants to send a message to another agent, it uses the directory service to obtain a set of transport-
687 descriptors for the agent. It then passes these transport-descriptors to the transport factory, which returns a transport-
688 handle. It should be noted that the transport factory and handle are not parts of the abstract architecture, but rather
689 artefacts of the specific implementation. The agent then uses an encoder provided by the encoding factory, to
690 transform the message into the desired encoding. Finally it transfers this encoded message to the recipient via the
691 selected transport.

692 4 Architectural Overview

693 The FIPA architecture defines at an abstract level how two agents can locate and communicate with each other by
 694 registering themselves and exchanging messages. To do this, a set of architectural elements and their relationships
 695 are described. In this section the basic relationships between the elements of the FIPA agent system are described. In
 696 *Section 5, Architectural Elements* and *Section 6, Agent and Agent Information Model*, there are descriptions of each
 697 element (including mandatory or optional status) and UML Models for the architecture, respectively.

698
 699 This section gives a relatively high level description of the notions of the architecture. It does not explain all of the
 700 aspects of the architecture. Use this material as an introduction, which can be combined with later sections to reach a
 701 fuller understanding of the abstract architecture.
 702

703 4.1 Agents and Services

704 **Agents** communicate by exchanging messages which represent speech acts, and which are encoded in an **agent-**
 705 **communication-language**.

706
 707 **Services** provide support services for **agents**. In addition to a number of standard services including **agent-directory-**
 708 **services** and **message-transport-services** this version of the Abstract Architecture defines a general service model
 709 that includes a **service-directory-service**.
 710

711 The Abstract architecture is explicitly neutral about how **services** are presented. They may be implemented either as
 712 **agents** or as software that is accessed via method invocation, using programming interfaces such as those provided in
 713 Java, C++, or IDL. An **agent** providing a **service** is more constrained in its behaviour than a general-purpose agent. In
 714 particular, these agents are required to preserve the semantics of the service. This implies that these agents do not
 715 have the degree of autonomy normally attributed to agents. They may not arbitrarily refuse to provide the service.
 716

717 It should be noted that if **services** are implemented as **agents** there are potential problems that may arise with
 718 discovering and communicating with these services. The resolution of these issues is beyond the scope of this
 719 document.
 720

721 4.2 Starting an Agent

722 On start-up an agent must be provided with a **service-root**. Typically the provider of the **service-root** will be a
 723 **service-directory-service** which will supply a set of **service-locators** for available agent lifecycle support services,
 724 such as **message-transport-services**, **agent-directory-services** and **service-directory-services**. In general, a
 725 **service-root** will provide sufficient entries to either describe all of the services available within the environment
 726 directly, or it will provide pointers to further services which will describe these services.
 727

728 4.3 Agent Directory Services

729 The basic role of the **agent-directory-service** is to provide a location where **agents** register their descriptions as
 730 **agent-directory-entries**. Other **agents** can search the **agent-directory-entries** to find **agents** with which they wish to
 731 interact.
 732

733 The **agent-directory-entry** is a **key-value-tuple** consisting of at least the following two **key-value-pairs**:
 734

Agent-name	A globally unique name for the agent
Agent-locator	One or more transport-descriptions , each of which is a self describing structure containing a transport-type , a transport-specific-address and zero or more transport-specific-properties used to communicate with the agent

735

In addition the **agent-directory-entry** may contain other descriptive attributes, such as the services offered by the **agent**, cost associated with using the **agent**, restrictions on using the **agent**, etc.

Note that the keys **agent-name** and **agent-locator** are short-form for the fully qualified names in the FIPA controlled namespace. See *Section 5.1.2, Key-Value Tuples* for further details.

4.3.1 Registering an Agent

Agent A wishes to advertise itself as a provider of some service. It first binds itself to one or more **transports**. In some implementations it will delegate this task to the **message-transport-service**; in others it will handle the details of, for example, contacting an ORB, or registering with an RMI registry, or establishing itself as a listener on a message queue. As a result of these actions, the agent is addressable via one or more **transports**.

Having established bindings to one or more **message-transport-services** the agent must advertise its presence. The agent realizes this by constructing an **agent-directory-entry** and registering it with the **agent-directory-service**. The **agent-directory-entry** includes the **agent-name**, its **agent-locator** and optional attributes that describe the service. For example, a stock service might advertise itself in abstract terms as {agent-service, "com.dowjones.stockticker"} and {ontology, org.fipa.ontology.stockquote}¹.


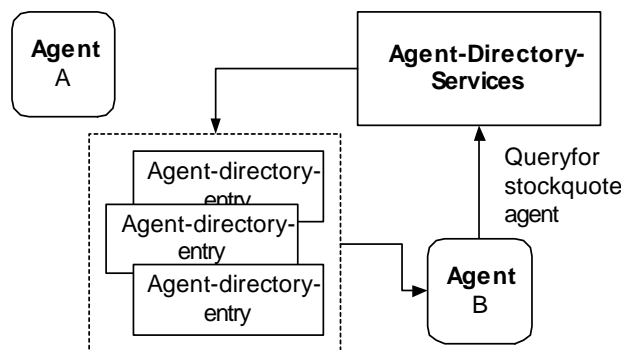


Figure 4: An Agent Registers with a Directory Service

4.3.2 Discovering an Agent

Agents can use the **agent-directory-service** to locate other agents with which to communicate. With reference to Figure 5, if agent B is seeking stock quotes, it may search for an agent that advertises use of the stockquote ontology. Technically, this would involve searching for an **agent-directory-entry** that includes the **key-value-pair** {ontology, {com, dowjones, ontology, stockquote}}. If it succeeds it will retrieve the **agent-directory-entry** for agent A. It might also retrieve other **agent-directory-entries** for agents that support that ontology.



¹ Note that the quoted string in the first example is a quoted value whereas the other elements are abstract names represented as tuples that may be encoded in a variety of different ways.

Figure 5: Directory Query

Agent B can then examine the returned **agent-directory-entries** to determine which agent best suits its needs. The **agent-directory-entries** include the **agent-name**, the **agent-locator**, which contains information related to how to communicate with the agent, and other optional attributes.

4.4 Service Directory Services

The basic role of the **service-directory-service** is to provide a consistent means by which agents and services can discover services. Operationally, the **service-directory-service** provides a location where **services** can register their service descriptions as **service-directory-entries**. Also, **agents** and **services** can search the **service-directory-service** to locate services appropriate to their needs.

The **service-directory-service** is analogous to but different to the **agent-directory-services**; the latter are oriented towards discovering **agents** whereas the former is oriented to discovering **services**. In practice also, the two kinds of directories may have radically different realizations. For example, on some systems a **service-directory-service** may be modelled simply as a fixed table of a small size whereas the **agent-directory-service** may be modelled using LDAP or other distributed directory technologies.

The entries in a **service-directory-service** are service descriptions consisting of a tuple containing a **service-id**, **service-type**, a **service-locator** and a set of optional **service-attributes**. The **service-locator** is a typed structure that may be used by **services** and **agents** to access the service.

The **service-directory-entry** is a **key-value-tuple** consisting of at least the following **key-value-pairs**:

Service-id	A globally unique name for the service
Service-type	The categorized <i>type</i> of the service
Service-locator	One or more key-value tuples containing a signature type , service signature and service address each

Additional **service-attributes** may be included that contain other descriptive properties of the **service**, such as the cost associated with using the **service**, restrictions on using the **service**, etc.

As a foundation for bootstrapping, each realization of the **service-directory-service** will provide agents with a **service-root**, which will take the form of a set of **service-locators** including at least one **service-directory-service**. (pointing to itself).

4.5 Agent Messages

In FIPA agent systems agents communicate with one another, by sending messages. Three fundamental aspects of message communication between agents are the message structure, message representation and message transport.

4.5.1 Message Structure

The structure of a **message** is a **key-value-tuple** (see *Section 5.1.2, Key-Value Tuples*) and is written in an **agent-communication-language**, such as FIPA ACL. The **content** of the **message** is expressed in a **content-language**, such as KIF or SL. **Content** expressions can be grounded by ontologies referenced within the **ontology key-value-tuple**. The messages also contain the **sender** and **receiver** names, expressed as **agent-names**. **Agent-names** are unique name identifiers for an agent. Every message has one sender and zero or more receivers. The case of zero receivers enables broadcasting of messages such as in ad-hoc wireless networks.

Messages can recursively contain other messages.

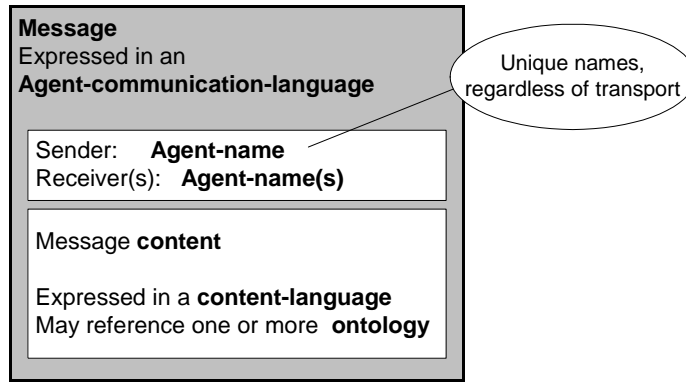


Figure 6: A Message

815
816
817
818

819 **4.5.2 Message Transport**

820 When a **message** is sent it is encoded into a **payload**, and included in a **transport-message**. The **payload** is
821 encoded using the **encoding-representation** appropriate for the transport. For example, if the **message** is going to be
822 sent over a low bandwidth transport (such a wireless connection) a bit efficient representation may used instead of a
823 string representation to allow more efficient transmission.

824
825 The **transport-message** itself is the **payload** plus the **envelope**. The **envelope** includes the sender and receiver
826 **transport-descriptions**. The **transport-descriptions** contain the information about how to send the message (via
827 what transport, to what address, with details about how to utilize the transport). The **envelope** can also contain
828 additional information, such as the **encoding-representation**, data related security, and other realization specific data
829 that needs be visible to the **transport** or recipient(s).

830

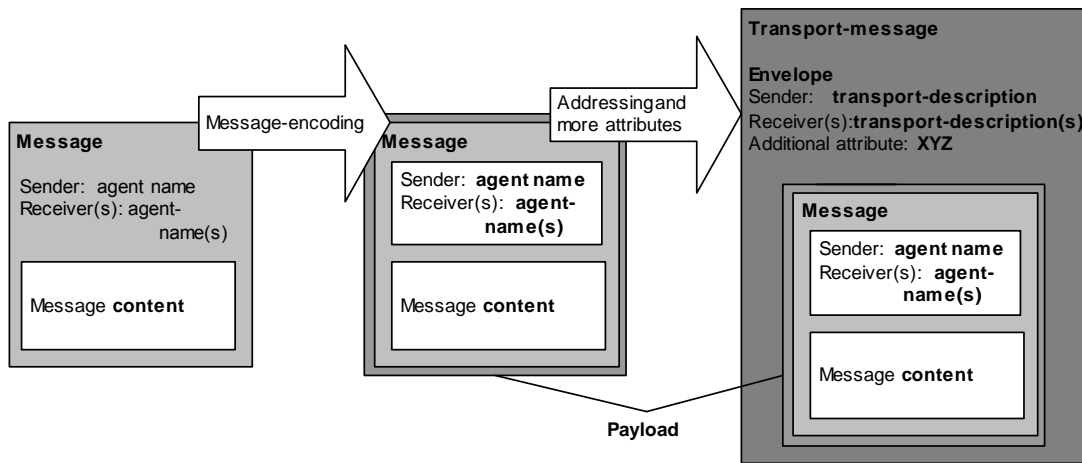


Figure 7: A Message Becomes a Transport-message

831
832
833
834
835
836
837
838
839
840

835 In the above diagram, a **message** is encoded into a **payload** suitable for transport over the selected **message-**
836 **transport**. It should be noted that **payload** adds nothing to the message, but only encodes it into another
837 representation. An appropriate **envelope** is created that has sender and receiver information that uses the **transport-**
838 **description** data appropriate to the transport selected. There may be additional envelope data also included. The
839 combination of the payload and envelope is termed as a **transport-message**.

841 **4.6 Agents Send Messages to Other Agents**

842 In FIPA agent systems agents are intended to communicate with one another. Hence, here are some of the basic
843 notions about agents and their communications:

844

845

Each **agent** has an **agent-name**. This **agent-name** is unique and unchangeable. Each agent also has one or more **transport-descriptions**, which are used by other agents to send a **transport-message**. Each **transport-description** correlates to a particular form of message **transport**, such as IOP, SMTP, or HTTP. A **transport** is a mechanism for transferring messages. A **transport-message** is a message that sent from one agent to another in a format (or encoding) that is appropriate to the **transport** being used. A set of **transport-descriptions** can be held in an **agent-locator**.

851

For example, there may be an **agent** with the **agent-name** "ABC". This agent is addressable through two different transports, HTTP and SMTP. Therefore, the agent has two **transport-descriptions**, which are held in the **agent-locator**. The transport descriptions are as follows:

855

Directory entry for ABC

857

Agent-name: ABC

858

Agent Locator:

Transport-type

HTTP

SMTP

Transport-specific-address

http://www.whiz.net/abc

Abc@lowcal.whiz.net

Transport-specific-property

(none)

(none)

860

Agent-attributes:

Attrib-1: yes

861

Attrib-2: yellow

862

Language: French, German, English

863

Preferred negotiation: contract-net

864

865

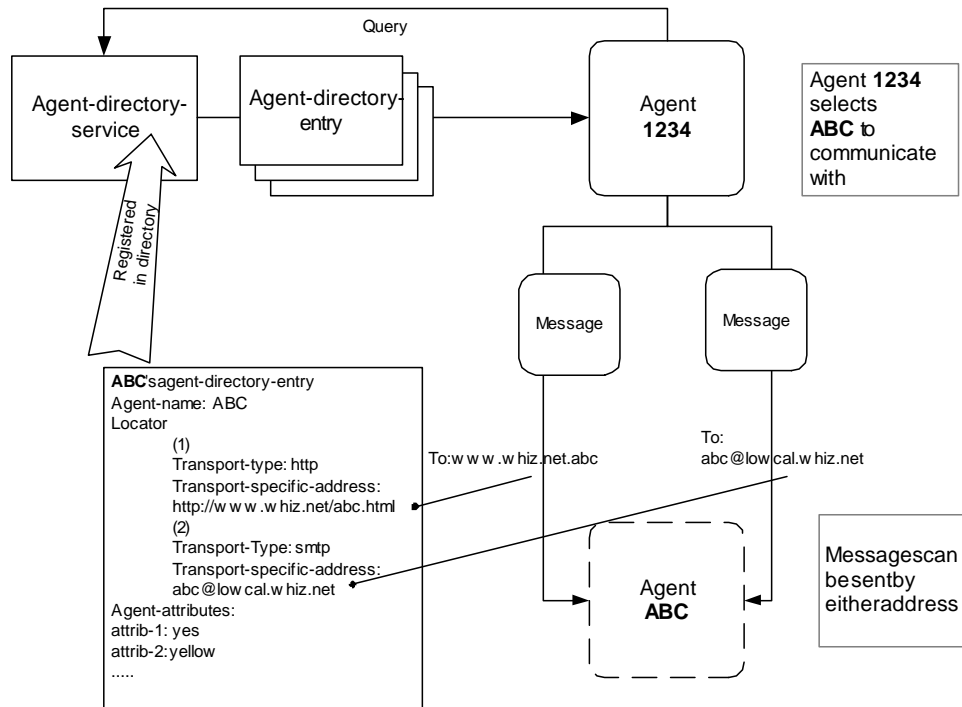
Note: in this example, the **agent-name** is used as part of the **transport-descriptions**. This is just to make these examples easier to read. There is *no* requirement to do this.

867

868

Another agent can communicate with agent "ABC" using either **transport-description**, and thereby know which agent it is communicating with. In fact, the second agent can even change transports and can continue its communication. Because the second agent knows the **agent-name**, it can retain any reasoning it may be doing about the other agent, without loss of continuity.

871



872

873

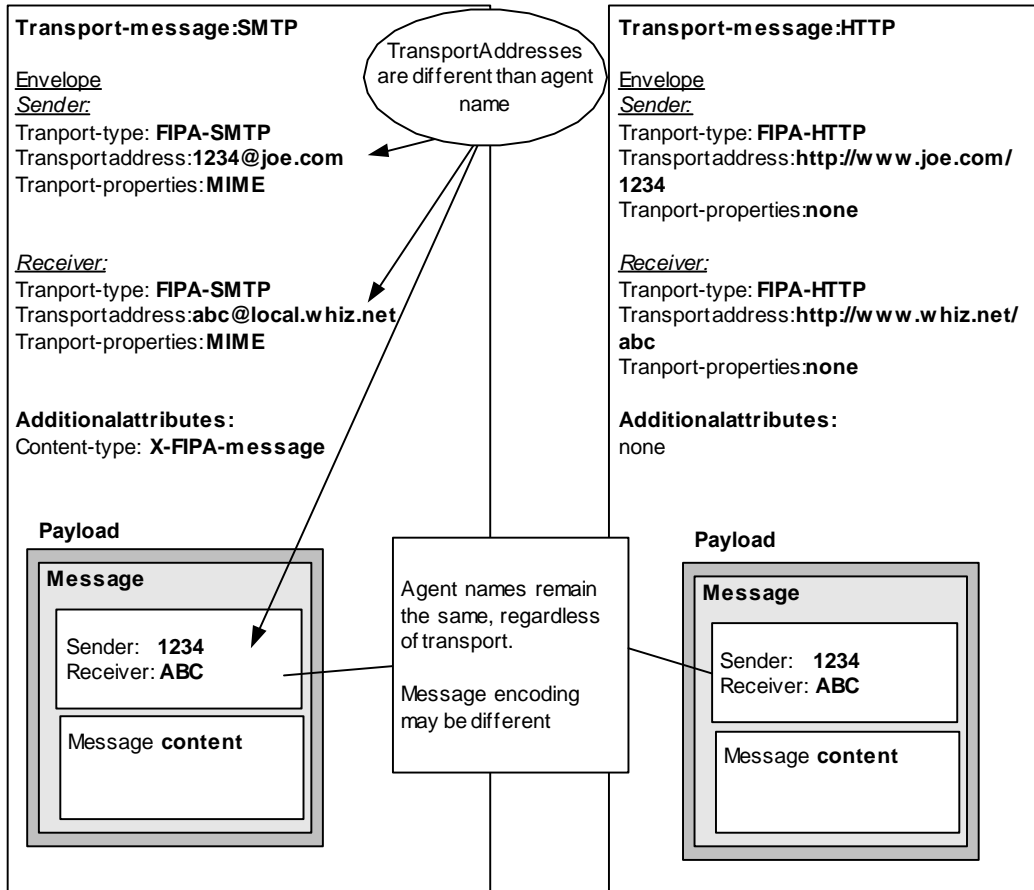
874

Figure 8: Communicating Using Any Transport

875
876
877
878
879
880
881
882

In the above diagram, Agent 1234 can communicate with Agent ABC using either an SMTP transport or an HTTP transport. In either case, if Agent 1234 is doing any reasoning about agents that it communicates with, it can use the **agent-name** "ABC" to record which agent it is communicating with, rather than the transport description. Thus, if it changes transports, it would still have continuity of reasoning.

Here's what the messages on the two different transports might look like:



883
884
885
886
887
888
889
890

Figure 9: Two Transport-Messages to the Same Agent

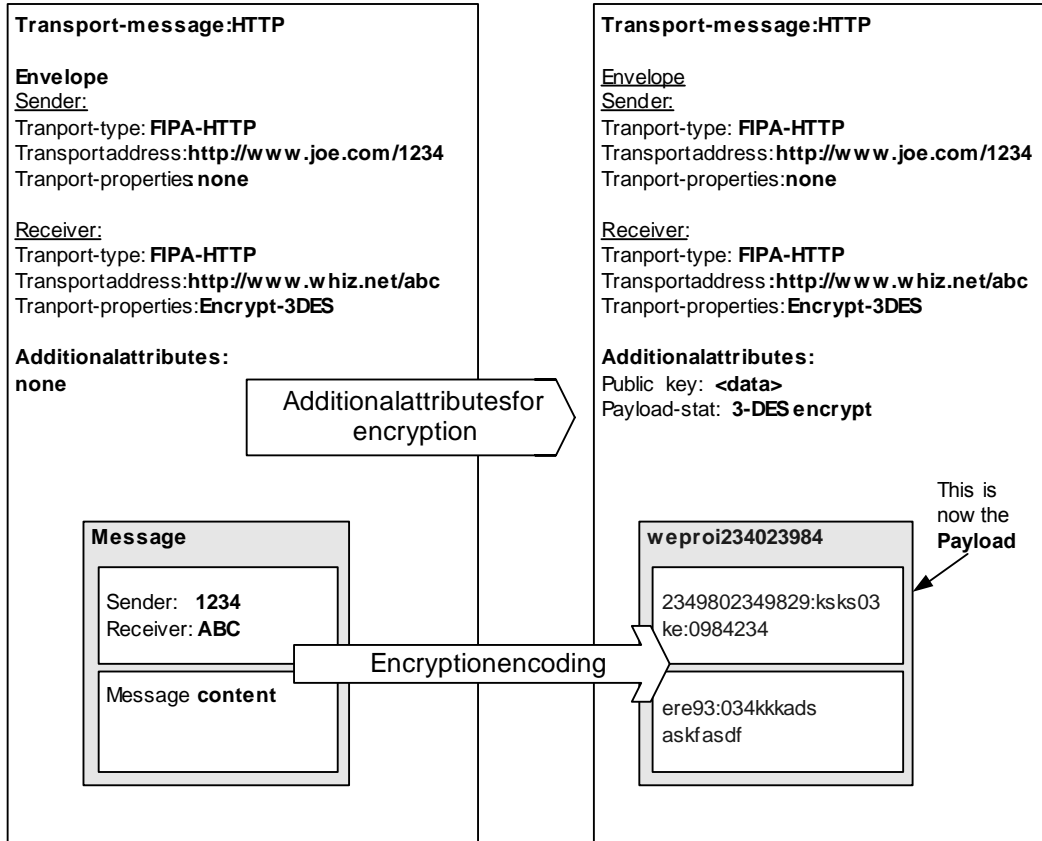
In the diagram above, the **transport-description** is different, depending on the transport that is going to be used. Similarly, the **message-encoding** of the **payload** may also be different. However, the **agent-names** remain consistent across the two message representations.

891 **4.7 Providing Message Validity and Encryption**

892 There are many aspects of security that can be provided in agent systems. See *Section 11, Informative Annex D —*
893 *Goals for Security and Identity Abstractions* for a discussion of possible security features. In this abstract architecture,
894 there is a simple form of security: message validity and message encryption. In message validity, messages can be
895 sent in such a way that any modification during transmission is identifiable. In message encryption, a message is sent
896 in encrypted form such that non-authorized entities cannot comprehend the message content.

897
898 In the abstract architecture these features are accommodated through **encoding-representations** and the use of
899 additional attributes in the **envelope**. For example, as the payload is encoded, one of the encodings could be to a
900 digitally encrypted set of data, using a public key and preferred encryption algorithm. Additional parameters are added
901 to the envelope to indicate these characteristics.

902



903
904
905
906
907
908
909

Figure 10: Encrypting a Message Payload

In the above diagram, the payload is encrypted, and additional attributes added to the envelope to support the encryption. These attributes must remain unencrypted in order that the receiving party is able to use them.

910 **4.8 Providing Interoperability**

911 There are two ways in which the abstract architecture makes provision for interoperability. The first is **transport**
912 interoperability. The second is **message** representation interoperability.

913
914 To provide interoperability, there are certain elements that must be included throughout the architecture to permit
915 multiple implementations. For example, earlier it was noted that an **agent** has both an **agent-name** and an **agent-**
916 **locator**. The **locator** contains **transport-descriptions**, each of which contains information necessary for a particular
917 transport to send a message to the corresponding agent. The semantics of agent communication require that an
918 agent's name be preserved throughout its lifetime, regardless of what transports may be used to communicate with it.
919

920 5 Architectural Elements

921 The elements of the abstract architecture are defined here. For each element, the semantics are described informally
 922 followed by the relationships between the element and others.

924 5.1 Introduction

925 5.1.1 Classification of Elements

926 The word **element** is used here to indicate an item or entity that is part of the architecture, and participates in
 927 relationships with other elements of the architecture.

928
 929 The architectural elements are classified as *mandatory* or *optional*. Mandatory elements must appear in all
 930 instantiations of the FIPA abstract architecture. They describe the fundamental services, such as agent registration
 931 and communications. These elements are the core aspects of the architecture. Optional elements are not mandatory;
 932 they represent architecturally useful features that may be shared by some, but not all, concrete instantiations. The
 933 abstract architecture only defines those optional elements that are highly likely to occur in multiple instantiations of the
 934 architecture.

935
 936 These descriptors and classifications are summarised in *Table 1*.

937

Word	Definition
Can, May	In relationship descriptions, the word can or may is used to indicate this is an optional relationship. For example, a service <i>may</i> provide an API invocation, but it is not required to do so.
Element, or architectural element	A member of this abstract architecture. The word element is used here to indicate an item or entity that is part of the architecture, and participates in relationships with other elements of the architecture.
Mandatory	Description of an element or relationship. Required in all fully functional implementations of the FIPA Abstract Architecture.
Must	In relationship descriptions, the word must is used to indicate this is a mandatory relationship. For example, an agent <i>must</i> have an agent-name means that an agent is required to have an agent-name .
Optional	Description of an element or relationship. May appear in any implementation of the FIPA Abstract Architecture, but is not required. Functionality that is common enough that it was included in model.
Realize, realization	To create a concrete specification or instantiation from the abstract architecture. For example, there may be a design to implement the abstract notion of agent-directory-services in LDAP. This could also be said that there is a <i>realization</i> of agent-directory-services .
Relationship	A connection between two elements in the architecture. The relationship between two elements is named (for example "is an instance of", "sends message to") and may have other attributes, such as whether it is required, optional, one-to-one, or one-to-many. The term as used in this document, is very much the way the term is used in UML or other system modelling techniques.

938

939

940

Table 1: Terminology

941 5.1.2 Key-Value Tuples

942 Many of the elements of the abstract architecture are defined to be **key-value-tuples**, or **KVTs**. For example, an ACL
 943 message, its envelope, and agent descriptions are all KVTs. The concept of a **KVT** is central to the notion of
 944 architectural extensibility, and so it is discussed in some length here.

945

946 A **KVT** consists of an unordered set of **key-value-pairs**. Each **key-value-pair** has two elements, as the term implies.
 947 The first element, the **key**, is a **pair-element** drawn from an administered name space. All keys defined by the Abstract
 948 Architecture are drawn from a name space managed by FIPA. This makes it possible for concrete architectures, or
 949 individual implementations, to add new architectural elements in a manner which is guaranteed not to conflict with the
 950 Abstract Architecture. The second element of the **key-value-pair** is the **value**. The type of value depends on the **key**.
 951 In many cases, the value is another **pair-element**, an identifier drawn from a name-space. In other cases, the **value** is
 952 a constant or expression of some specific type.

953
 954 The rest of this section describes the rules governing the names for **keys** and **values**.
 955

956 Traditionally, **pair-elements** have been treated as simple text strings. It is more useful to adopt a more abstract model
 957 in which abstract identifiers and keywords may be encoded in a variety of different ways.
 958

959 It is also important that the FIPA elements represented as **key-value-tuples** should be extensible. There are three
 960 types of extension that can be envisaged:

- 961 Official FIPA sanctioned standard extensions,
 962
 963 Durable vendor-specific extensions, and,
 964
 965 Temporary, probably private, extensions.
 966

967
 968 The last of these has traditionally been addressed by using a particular prefix string ("X-").
 969

970 Every **pair-element** is an ordered tuple of **tokens**. This tuple denotes a name within a hierarchical namespace, in
 971 which the first **token** in the tuple is at the highest level in the hierarchy and the rightmost is the leaf. Examples of tuples
 972 are:

- 973 {org, fipa, standard, ontology, foo}
 974 {com, sun, java, agent, performative, brainwash}
 975 {x, cc}
 976 {protocol}
 977

978
 979 A **pair-element** containing more than one **token** is a **qualified-element**. In a **qualified-element**, the left-most **token**
 980 must correspond to one of the top-level ICANN domain names, or to an **anonymous-token**. The latter is used to
 981 introduce temporary, experimental **qualified-elements**.
 982

983 With reference to the FQN (Fully Qualified Name) field in Table 2, if a **pair-element** contains only one **token**, it is an
 984 **unqualified-element**. An **unqualified-element** is interpreted according to Table 2, as though its **token** were
 985 appended to a tuple of tokens defining a FIPA standard name space, as follows:
 986

987 For example, the **pair-element**

988 { {ontology}, {foo} }

989 is equivalent to,

990 { {org, fipa, standard, message, ontology}, {org, fipa, standard, message, ontology, foo} }

991
 992 The natural encoding of a **pair-element** is as a sequence of text strings separated by dots. Thus the **pair-element**

993 { {org, fipa, standard, message, ontology}, {org, fipa, standard, message, ontology, foo} },
 994

995 will naturally be encoded as:

996
 997 org.fipa.standard.message.ontology org.fipa.standard.message.ontology.foo
 998
 999
 1000
 1001

1002

1003 **5.1.3 Services**

1004 A **service** is defined in terms of a set of **actions** that it supports. Each action defines an interaction between the
 1005 **service** and the **agent** using the service. The semantics of these actions are described informally, to minimize
 1006 assumptions about how they might be reified in a concrete specification.
 1007

1008 **5.1.4 Format of Element Description**

1009 The architectural elements are described below. The format of the description is:

1010

1011 **Summary.** A summary of the element.

1012

1013 **Relationship to other elements.** A complete description of the relationship of this element to the other
 1014 architectural elements.

1015 **Actions.** In the case of mandatory services, the actions that may be exerted by that service are described.
 1016

1017

1017 **Description.** Additional description and context for the element, along with explanatory notes and examples.
 1018

1018

1019 **5.1.5 Abstract Elements**

Element	Description	Fully Qualified Name (FQN)	Presence
Action-status	A status indication delivered by a service showing the success or failure of an action.	org.fipa.standard.service.action-status	Mandatory
Agent	A computational process that implements the autonomous, communicating functionality of an application.	org.fipa.standard.agent	Mandatory
Agent-attribute	A set of properties associated with an agent by inclusion in its agent-directory-entry .	org.fipa.standard.agent.agent-attribute	Optional
Agent-communication-language	A language with a precisely defined syntax semantics and pragmatics, which is the basis of communication between independently designed and developed agents .	org.fipa.standard.agent-communication-language	Mandatory
Agent-directory-entry	A composite entity containing the name , agent-locator , and agent-attributes of an agent .	org.fipa.standard.service.agent-directory-entry	Mandatory
Agent-directory-service	A service providing a shared information repository in which agent-directory-entries may be stored and queried	org.fipa.standard.service.agent-directory-service	Mandatory
Agent-locator	An agent-locator consists of the set of transport-descriptions used to communicate with an agent .	org.fipa.standard.service.message-transport-service.agent-locator	Mandatory
Agent-name	An opaque, non-forgeable token that uniquely identifies an agent .	org.fipa.standard.agent-name	Mandatory
Content	Content is that part of a message (communicative act) that represents the domain dependent component of the communication.	org.fipa.standard.message.content	Mandatory
Content-language	A language used to express the content of a communication between agents.	org.fipa.standard.message.content-language	Mandatory

Encoding-representation	A way of representing an abstract syntax in a particular concrete syntax. Examples of possible representations are XML, FIPA Strings, and serialized Java objects.	org.fipa.standard.encoding-service.encoding-representation	Mandatory
Encoding-service	A service that encodes a message to and from a payload .	org.fipa.standard.service.encoding-service	Mandatory
Envelope	That part of a transport-message containing information about how to send the message to the intended recipient(s). May also include additional information about the message encoding, encryption, etc.	org.fipa.standard.transport-message.envelope	Mandatory
Explanation	An encoding of the reason for a particular action-status .	org.fipa.standard.service.explanation	Optional
Message	A unit of communication between two agents. A message is expressed in an agent-communication-language , and encoded in an encoding-representation .	org.fipa.standard.message	Mandatory
Message-transport-service	A service that supports the sending and receiving of transport-messages between agents .	org.fipa.standard.service.message-transport-service	Mandatory
Ontology	A set of symbols together with an associated interpretation that may be shared by a community of agents or software. An ontology includes a vocabulary of symbols referring to objects in the subject domain, as well as symbols referring to relationships that may be evident in the domain.	org.fipa.standard.message.ontology	Optional
Payload	A message encoded in a manner suitable for inclusion in a transport-message .	org.fipa.standard.transport-message.payload	Mandatory
Service	A service provided for agents and other services .	org.fipa.standard.service	Mandatory
Service-address	A service-type specific string containing transport addressing information.	org.fipa.standard.service.service-address	Mandatory
Service-attributes	A set of properties associated with a service by inclusion in its service-directory-entry .	org.fipa.standard.service.service-attributes	Optional
Service-directory-entry	A composite entity containing the service-id , service-locator , and service-type of a service .	org.fipa.standard.service.service-directory-service.service-directory-entry	Mandatory
Service-directory-service	A directory service for registering and discovering services .	org.fipa.standard.service.service-directory-service	Mandatory
Service-id	A unique identifier of a particular service .	org.fipa.standard.service.service-id	Mandatory
Service-location-description	A key-value-tuple containing a signature-type a service-signature and service-address .	org.fipa.standard.service.service-location-description	Mandatory
Service-locator	A service-locator consists of the set of service-location-descriptions used to access a service .	org.fipa.standard.service.service-locator	Mandatory
Service-root	A set of service-directory-entries .	org.fipa.standard.service.service-root	Mandatory
Service-signature	A identifier that describes the binding signature for a service .	org.fipa.standard.service.service-type	Mandatory
Service-type	A key-value tuple describing the type of a	org.fipa.standard.service	Mandatory

	service.	e.service-type	
Signature-type	A key-value tuple describing the type of service-signature .	org.fipa.standard.servic e.signature-type	
Transport	A transport is a particular data delivery service supported by a given message-transport-service .	org.fipa.standard.servic e.message-transport- service.transport	Mandatory
Transport-description	A transport-description is a self describing structure containing a transport-type , a transport-specific-address and zero or more transport-specific-properties .	org.fipa.standard.servic e.message-transport- service.transport- description	Mandatory
Transport-message	The object conveyed from agent to agent . It contains the transport-description for the sender and receiver or receivers, together with a payload containing the message .	org.fipa.standard.transp ort-message	Mandatory
Transport-specific-address	A transport address specific to a given transport-type	og.fipa.standard.service .message-transport- service.transport- specific-address	Mandatory
Transport-specific-property	A transport-specific-property is a property associated with a transport-type .	org.fipa.standard.servic e.message-transport- service.transport- specific-property	Optional
Transport-type	A transport-type describes the type of transport associated with a transport-specific-address .	org.fipa.standard.servic e.message-transport- service.transport-type	Mandatory

Table 2: Abstract Elements

1020
1021
1022

5.2 Agent

1023

5.2.1 Summary

1024

An **agent** is a computational process that implements the autonomous, communicating functionality of an application. Typically, agents communicate using an **Agent Communication Language**. A concrete instantiation of **agent** is a mandatory element of every concrete instantiation of the abstract architecture.

1025
1026
1027
1028

5.2.2 Relationships to Other Elements

1029

Agent has an **agent-name**

1030

Agent may have **agent-attributes**

1031

Agent has an **agent-locator**, which lists the **transport-descriptions** for that agent

1032

Agent may be sent messages via a **transport-description**, using the **transport** corresponding to the **transport-description**

1033

1034

Agent may send a **transport-message** to one or more **agents**

1035

Agent may register with one or more **agent-directory-services**

1036

Agent may have an **agent-directory-entry**, which is registered with an **agent-directory-service**

1037

Agent may modify its **agent-directory-entry** as registered by an **agent-directory-service**

1038

Agent may delete its **agent-directory-entry** from an **agent-directory-service**.

1039

Agent may query for an **agent-directory-entry** registered within an **agent-directory-service**

1040

Agent is addressable by the mechanisms described in its **transport-descriptions** in its **agent-directory-entry**.

1041

1042

1043 5.2.3 Description

1044 In a concrete instantiation of the abstract architecture, an **agent** may be realized in a variety of ways, for example as a
 1045 Java™ component, a COM object, a self-contained Lisp program, or a TCL script. It may execute as a native process
 1046 on some physical computer under an operating system, or be supported by an interpreter such as a Java Virtual
 1047 Machine or a TCL system. The relationship between the **agent** and its computational context is specified by the agent
 1048 lifecycle. The abstract architecture does not address the lifecycle of agents as it is often handled differently in discrete
 1049 computational environments. Realizations of the abstract architecture *must* address these issues.
 1050

1051 5.3 Agent Attribute

1052 5.3.1 Summary

1053 An **agent-attribute** is one of a set of optional attributes that form part of the **agent-directory-entry** for an **agent**. They
 1054 are represented as **key-value-pairs** within the **key-value-tuple** that makes up the **agent-directory-entry**. The
 1055 purpose of the attributes is to allow searching for **agent-directory-entries** that match **agents** of interest. A concrete
 1056 instantiation of **agent-attribute** is an optional element of concrete instantiations of the abstract architecture.
 1057

1058 5.3.2 Relationships to Other Elements

1059 An **agent-directory-entry** may have zero or more **agent-attributes**
 1060 An **agent-attribute** describes aspects of an **agent**
 1061

1062 5.3.3 Description

1063 When an **agent** registers an **agent-directory-entry**, the **agent-directory-entry** may optionally contain **key-value-**
 1064 **pairs** that offer additional description of the **agent**. The values might include information about costs of using the
 1065 **agent** or **service**, features available, **ontologies** understood, names that the service is commonly known by, or any
 1066 other data that agents deem useful. This information can then be used to enhance search criteria exerted by **agents**
 1067 on the **agent-directory-service**.
 1068

1069 In practice, when defining realizations of this abstract architecture, domain specific specifications should exist
 1070 describing the **agent-attributes** to be used. This eases requirements for interoperation.
 1071

1072 5.4 Agent Communication Language

1073 5.4.1 Summary

1074 An **agent-communication-language** (ACL) is a language in which communicative acts can be expressed and hence
 1075 **messages** constructed. A concrete instantiation of **agent-communication-language** is a mandatory element of every
 1076 concrete instantiation of the abstract architecture.
 1077

1078 5.4.2 Relationships to Other Elements

1079 **Message** is written in an **agent-communication-language**

1080 5.4.3 Description

1081 FIPA ACL is described in detail in [FIPA00061] and FIPA communicative acts in [FIPA00037].
 1082

5.5 Agent Directory Entry

5.5.1 Summary

An **agent-directory-entry** is a **key-value tuple** consisting of the **agent-name**, an **agent-locator**, and zero or more **agent-attributes**. An **agent-directory-entry** refers to an **agent**; in some implementations this agent will provide a **service**. A concrete instantiation of **agent-directory-entry** is a mandatory element of every concrete instantiation of the abstract architecture.

5.5.2 Relationships to Other Elements

Agent-directory-entry contains the **agent-name** of the **agent** to which it refers

Agent-directory-entry contains one **agent-locator** of the **agent** to which it refers. The **agent-locator** contains one or more **transport-descriptions**

Agent-directory-entry is managed by and available from an **agent-directory-service**

Agent-directory-entry may contain **agent-attributes**

5.5.3 Description

Different realizations that use a common **agent-directory-service**, are strongly encouraged to adopt a common schema for storing **agent-directory-entries**. (This in turn implies the use of a common representation for **agent-locators**, **transport-descriptions**, **agent-names**, and so forth.)

Agents are not required to publish an **agent-directory-entry**. It is possible for agents to communicate with agents that have provided a **transport-description** through a private mechanism. For example, an agent involved in a negotiation may receive a **transport-description** directly from the party with which it is negotiating. This falls outside the scope of the using the **agent-directory-services** mechanisms.

5.6 Agent Directory Service

5.6.1 Summary

An **agent-directory-service** is a shared information repository in which **agents** may publish their **agent-directory-entries** and in which they may search for **agent-directory-entries** of interest. A concrete instantiation of **agent-directory-service** is a mandatory element of every concrete instantiation of the abstract architecture.

5.6.2 Relationships to Other Elements

Agent may register its **agent-directory-entry** with an **agent-directory-service**

Agent may modify its **agent-directory-entry** as registered by an **agent-directory-service**

Agent may delete its **agent-directory-entry** from an **agent-directory-service**

Agent may search for an **agent-directory-entry** registered within an **agent-directory-service**

An **agent-directory-service** must accept valid, authorized requests to register, de-register, delete, modify and identify agent descriptions

An **agent-directory-service** must accept valid, authorized requests for searching

5.6.3 Actions

An **agent-directory-service** supports the following actions.

5.6.3.1 Register

An **agent** may **register** an **agent-directory-entry** with an **agent-directory-service**. The semantics of this action are as follows:

The **agent** provides an **agent-directory-entry** that is to be registered. In initiating the action, the **agent** may control the scope of the action. Different reifications may handle this in different ways. The action may be addressed to a particular instance of an **agent-directory-service**, or the action may be qualified with some kind of scope parameter.

If the action is successful, the **agent-directory-service** will return an **action-status** indicating success. Following a successful **register**, the **agent-directory-service** will support legal **modify**, **delete**, and **query** actions with respect to the registered **agent-directory-entry**.

If the action is unsuccessful, the **agent-directory-service** will return an **action-status** indicating failure, together with an **explanation**. The range of possible explanations is, in general, specific to a particular reification. However a conforming reification must, where appropriate, distinguish between the following explanations:

Duplicate. The new entry "clashed" with some existing **agent-directory-entry**. Normally this would only occur if an existing **agent-directory-entry** had the same **agent-name**, but specific reifications may enforce additional requirements.

Access. The **agent** making the request is not authorized to perform the specified action.

Invalid. The **agent-directory-entry** is invalid in some way.

5.6.3.2 Modify

An **agent** may **modify** an **agent-directory-entry** that has been registered with an **agent-directory-service**. The semantics of this action are as follows:

The **agent** provides an **agent-directory-entry** which contains the same **agent-name** as the entry to be modified. In initiating the action, the **agent** may control the scope of the action. Different reifications may handle this in different ways. The action may be addressed to a particular instance of an **agent-directory-service**, or the action may be qualified with some kind of scope parameter.

The **agent-directory-service** verifies that the argument is a valid **agent-directory-entry**. It then searches for a registered **agent-directory-entry** with the same **agent-name**. If it does not find one, the action fails and an **explanation** provided. Otherwise it modifies the existing **agent-directory-entry** by examining each **key-value pair** in new **agent-directory-entry**. If the **value** is non-null, the **pair** is added to the new entry, replacing any existing **pair** with the same **key**. If the **value** is null, any existing **pair** with the same **key** is removed from the entry.

If the action is successful, the **agent-directory-service** will return an **action-status** indicating success, together with an **agent-directory-entry** corresponding to the new contents of the registered entry. Following a successful **register**, the **agent-directory-service** will support legal **modify**, **delete**, and **query** actions with respect to the modified **agent-directory-entry**.

If the action is unsuccessful, the **agent-directory-service** will return an **action-status** indicating failure, together with an **explanation**. The range of possible explanations is, in general, specific to a particular reification. However a conforming reification must, where appropriate, distinguish between the following explanations:

Not-found. The new entry did not match any existing **agent-directory-entry**. This would only occur if no existing **agent-directory-entry** had the same **agent-name**.

Access. The **agent** making the request is not authorized to perform the specified action.

Invalid. The new **agent-directory-entry** is invalid in some way.

5.6.3.3 Delete

An **agent** may **delete** an **agent-directory-entry** from an **agent-directory-service**. The semantics of this action are as follows:

1183
 1184 The **agent** provides an **agent-directory-entry** which has the same **agent-name** as that which is to be deleted. (The
 1185 rest of the **agent-directory-entry** is not significant.) In initiating the action, the **agent** may control the scope of the
 1186 action. Different reifications may handle this in different ways. The action may be addressed to a particular instance of
 1187 an **agent-directory-service**, or the action may be qualified with some kind of scope parameter.

1188
 1189 If the action is successful, the **agent-directory-service** will return an **action-status** indicating success. Following a
 1190 successful **delete**, the **agent-directory-service** will no longer support **modify**, **delete**, and **query** actions with respect
 1191 to the registered **agent-directory-entry**.

1192
 1193 If the action is unsuccessful, the **agent-directory-service** will return an **action-status** indicating failure, together with
 1194 an **explanation**. The range of possible explanations is, in general, specific to a particular reification. However a
 1195 conforming reification must, where appropriate, distinguish between the following explanations:

1196
 1197 *Not-found.* The new entry did not match any existing **agent-directory-entry**. This would only occur if no existing
 1198 **agent-directory-entry** had the same **agent-name**.

1199
 1200 *Access.* The **agent** making the request is not authorized to perform the specified action.

1201
 1202 *Invalid.* The **agent-directory-entry** is invalid in some way.
 1203

1204 5.6.3.4 Query

1205 An **agent** may **query** an **agent-directory-service** to locate **agent-directory-entries** of interest. The semantics of this
 1206 action are as follows:

1207
 1208 The **agent** provides an **agent-directory-entry** that is to be treated as a search pattern. In initiating the action, the
 1209 **agent** may control the scope of the action. Different reifications may handle this in different ways. The action may be
 1210 addressed to a particular instance of an **agent-directory-service**, or the action may be qualified with some kind of
 1211 scope parameter.

1212
 1213 The directory service verifies that the argument is a valid **agent-directory-entry**. It then searches for registered **agent-**
 1214 **directory-entries** that satisfy the search criteria. A registered entry satisfies the search criteria if there is a match
 1215 between each **key-value pair** in the submitted entry. The semantics of "matching" are likely to be reification-
 1216 dependent; at a minimum, there should be support for matching on the *same* value and on *any* value.

1217
 1218 If the action is successful, the **agent-directory-service** will return an **action-status** indicating success, together with a
 1219 set of **agent-directory-entries** that satisfy the search pattern. The mechanism by which multiple entries are returned,
 1220 and whether or not an agent may limit or terminate the delivery of results, is not defined in the abstract architecture and
 1221 is therefore reification dependent.

1222
 1223 If the action is unsuccessful, the **agent-directory-service** will return an **action-status** indicating failure, together with
 1224 an **explanation**. The range of possible explanations is, in general, specific to a particular reification. However a
 1225 conforming reification must, where appropriate, distinguish between the following explanations:

1226
 1227 *Not-found.* The search pattern did not match any existing **agent-directory-entry**.

1228
 1229 *Access.* The **agent** making the request is not authorized to perform the specified action.

1230
 1231 *Invalid.* The **agent-directory-entry** is invalid in some way.
 1232

1233 5.6.4 Description

1234 An **agent-directory-service** may be implemented in a variety of ways, using a general-purpose scheme such as
 1235 X.500 or some private agent-specific mechanism. Typically an **agent-directory-service** uses some hierarchical or

1236 federated scheme to support scalability. A concrete implementation may support such mechanisms automatically, or
 1237 may require each **agent** to manage its own directory usage.
 1238

1239 Different realizations that are based on the same underlying mechanism are strongly encouraged to adopt a common
 1240 schema for storing **agent-directory-entries**. This in turn implies the use of a common representation for **names**,
 1241 **locations**, and so forth. For example, considering multiple implementations of directory services in LDAP, it would be
 1242 useful for all of the implementations to interoperate because they are using the same schemas. Similarly, if there were
 1243 multiple implementations in NIS, they would need the same NIS data representation to interoperate.
 1244

1245 The **agent-directory-service** described here does not have the full flexibility found in the *directory-facilitator* (see
 1246 [FIPA00023]), of existing FIPA specifications. In practice, implementing the search capabilities of the existing *directory-*
 1247 *facilitator* is not feasible with most directory systems, that is, LDAP, X.500 and NIS. There seems to be a need for a
 1248 Lookup Service, which is here called the **agent-directory-service**, which allows an agent to identify and get the
 1249 **transport-description** for another agent, as well as a more complex search system, which can resolve complex
 1250 searches. The former system, which supports a single level of search on attributes, is the **agent-directory-service**.
 1251 The latter might be implemented as a broker, and might be implemented in systems that allow for arbitrary complexity
 1252 and nesting such as Prolog or LISP. This division of functionality reflects the experience of many implementations,
 1253 where there is a "quick" lookup service and a more robust, but slower complex search service.
 1254

1255 5.7 Agent Locator

1256 5.7.1 Summary

1257 An **agent-locator** consists of the set of **transport-descriptions**, which can be used to communicate with an **agent**. An
 1258 **agent-locator** may be used by a **message-transport-service** to select a **transport** for communicating with the **agent**,
 1259 such as an agent or a **service**. **Agent-locators** can also contain references to software interfaces. This can be used
 1260 when a **service** can be accessed programmatically, rather than via a messaging model. A concrete instantiation of
 1261 **agent-locator** is a mandatory element of every concrete instantiation of the abstract architecture.
 1262

1263 5.7.2 Relationships to Other Elements

1264 **Agent-locator** is a member of **agent-directory-entry**, which is registered with an **agent-directory-service**

1265 **Agent-locator** contains one or more **transport-descriptions**

1266 **Agent-locator** is used by **message-transport-service** to select a **transport**
 1267

1268 5.7.3 Description

1269 The **agent-locator** serves as a basic building block for managing address and transport resolution. An **agent-locator**
 1270 includes all of the **transport-descriptions** that may be used to contact the related **agent** or **service**.
 1271

1272 5.8 Agent Name

1273 5.8.1 Summary

1274 An **agent-name** is a means to identify an **agent** to other **agents** and **services**. It is expressed as a **key-value-pair**, is
 1275 unchanging (that is, it is immutable), and unique under normal circumstances of operation. A concrete instantiation of
 1276 **agent-name** is a mandatory element of every concrete instantiation of the abstract architecture.
 1277

1278 5.8.2 Relationships to Other Elements

1279 **Agent** has one **agent-name**

1280 **Message** must contain the **agent-names** of the sending and receiving **agents**

1281 **Agent-directory-entry** must contain the **agent-name** of the **agent** to which it refers
 1282

5.8.3 Description

An **agent-name** is an identifier (e.g., a GUID, Globally Unique Identifier) that is associated with the **agent** at creation time or initial registration. Name issuing should occur in a way that tends to ensure global uniqueness. This may be achieved, for example, through employing an algorithm that generates the name with a sufficient degree of stochastic complexity such as to induce a vanishingly small chance of a name collision.

The **agent-name** will typically be issued by another entity or service. Once issued, the unique identifier should not be dependent upon the continued existence of the third party that issued it. Ideally through, there will be some mechanism available that is capable of verifying name authenticity.

Beyond this durable relationship with the **agent** it denotes, the **agent-name** should have no semantics. It should not encode any actual properties of the agent itself, nor should it disclose related information such as agent **transport-description** or **location**. It should also not be used as a form of authentication of the agent. Authentication services must rely on the combination of a unique identifier plus additional information (for example, a certificate that makes the name tamper-proof and verifies its authenticity through a trusted third party).

A useful role of an **agent-name** is to support the use of BDI (belief/desire/intention) models within a multi-agent system. The **agent-name** can be used to correlate propositional attitudes with the particular **agents** that are believed to hold those attitudes.

Agents may also have "well-known" or "common" or "social" names, or "nicknames", or aliases by which they are popularly known. These names are often used to commonly identify an agent. For example, within an agent system, there may be a broker service for finding "air-fare" agents. The convention within this system is that this agent is nicknamed "Air-fare broker". In practice, this is implemented as an **agent-attribute**. The attribute could have the key "Nickname" with the value "Air-fare broker". However, the actual name of the agent providing the function is unique, to maintain the ability to distinguish between an agent providing that function in one cluster of agents, and another agent providing the same function in a different cluster of agents.

5.9 Content

5.9.1 Summary

Content is that part of a **message** (where a message is a communicative act) that represents the component of the communication that refers to a domain or topic area. **Content** is expressed using **content-languages**. Expressions contained within the content, or the entire content expression itself, can be put into context by one or more **ontologies**. A concrete instantiation of **content** is a mandatory element of every concrete instantiation of the abstract architecture.

5.9.2 Relationships to Other Elements

Content is expressed in a **content-language**

Content may reference one or more ontologies referenced in the **ontology** attribute of a **message**

Content is part of a **message**

5.9.3 Description

The **content** of a **message** is the propositional content of a speech act. It does not refer to everything within the message, including delimiters, as it does with some languages, but rather the domain specific component only.

5.10 Content Language

5.10.1 Summary

A **content-language** is a language used to express the **content** of a communication between agents. FIPA allows considerable flexibility in the choice, form and encoding of a content language. However, content languages are

1331 required to be able to represent propositions, actions and terms (names of individual entities) if they are to make full
 1332 use of the standard FIPA performatives. A concrete instantiation of **content-language** is a mandatory element of every
 1333 concrete instantiation of the abstract architecture.
 1334

1335 5.10.2 Relationships to Other Elements

1336 **Content** is expressed in a **content-language**
 1337 **FIPA-SL** is an example of a **content-language**
 1338 **FIPA-RDF** is an example of a **content-language**
 1339 **FIPA-KIF** is an example of a **content-language**
 1340 **FIPA-CCL** is an example of a **content-language**
 1341

1342 5.10.3 Description

1343 The FIPA content language library is described in detail in [FIPA00007].
 1344

1345 5.11 Encoding Representation

1346 5.11.1 Summary

1347 An **encoding-representation** is a way of representing a **message** in a particular transport encoding. Examples of
 1348 possible representations are XML, Bit-efficient encoding and serialized Java objects. Typically an **encoding-**
 1349 **representation** is applied to the **payload** component of a **transport-message** to prepare it for transmission. A
 1350 concrete instantiation of **encoding-representation** is a mandatory element of every concrete instantiation of the
 1351 abstract architecture.
 1352

1353 5.11.2 Relationships to Other Elements

1354 **Payload** and the **message** and **content** contained within is encoded according to an **encoding-representation**
 1355 **Encoding-representation** is used by an **encoding-service**

1356 5.11.3 Description

1357 The way in which a message is encoded depends on the concrete architecture. If a particular architecture supports
 1358 only one form of encoding, no additional information is required. If multiple forms of encoding are supported, messages
 1359 may be made self-describing using techniques such as format tags, object introspection, and XML DTD references.
 1360

1361 5.12 Encoding Service

1362 5.12.1 Summary

1363 An **encoding-service** is a **service**. It provides the facility to encode a **message** or **content** into an **encoding-**
 1364 **representation** for use as a **transport-message payload**. This procedure must also function in reverse for decoding
 1365 **transport-messages**. A concrete instantiation of **encoding-service** is a mandatory element of every concrete
 1366 instantiation of the abstract architecture.
 1367

1368 5.12.2 Relationships to Other Elements

1369 **Encoding-service** converts a message into an **encoding-representation**
 1370 **Encoding-service** converts an **encoding-representation** into a **message**
 1371 **Encoding-service** can encode a **message** and message **content** as a **payload**
 1372 **Encoding-service** can decode a **payload** into a **message**
 1373 **Encoding-service** is a **service**
 1374

5.12.3 Actions

An **encoding-service** supports the following actions.

5.12.3.1 Transform Encoding/Decoding

An **agent** uses an **encoding-service** to convert a **message** to a **payload** and vice versa. That is, between **message** representation and a particular **encoding-representation**. It does this by invoking the **transform-encoding** action of the **encoding-service**. The semantics of this action are as follows:

To encode a message, the **agent** provides the **message** to the **encoding-service**, along with the type of encoding to be used. The encodings offered by the service may be queried using the **query-available-encodings** action described below. Encoding is context sensitive to ensure that appropriate **encoding-representations** are applied to specific message components. I.e. a **message** may be encoded in XML representation, but the **payload** that contains that **message** must be encoded for the transport to be used.

To decode a message, the encoded **payload** component of a **transport-message** is handed off to the **encoding-service** which decodes it into the **message**.

If the action is successful, the **encoding-service** will return an **action-status** indicating success, together with the encoded message component.

If the action is unsuccessful, the **encoding-service** will return an **action-status** indicating failure, together with an **explanation**. The range of possible explanations is, in general, specific to a particular reification. However a conforming reification must, where appropriate, distinguish between the following explanations:

Access. The **agent** making the request is not authorized to perform the specified action.

Invalid Message. The **message** to be encoded is invalid in some way.

Invalid Payload. The **payload** to be decoded is invalid in some way.

Invalid Encoding. The **encoding-representation** selected is unavailable.

5.12.3.2 Query Encoding Representation

An **agent** may query the **encoding-service** to resolve the **encoding-representation** with which the supplied **payload** has been encoded. It does this by invoking the **query-encoding-representation** action of the **encoding-transform-service**.

If the action is successful, the **encoding-service** will return an **action-status** indicating success. Additionally, the **encoding-representation** identity is returned.

If the action is unsuccessful, the **encoding-service** will return an **action-status** indicating failure, together with an **explanation**. The range of possible explanations is, in general, specific to a particular reification. However a conforming reification must, where appropriate, distinguish between the following explanations:

Access. The **agent** making the request is not authorized to perform the specified action.

Invalid. The encoded **payload** is invalid in some way.

Unidentifiable. The **encoding-representation** is unidentifiable by the **encoding-service**.

5.12.3.3 Query Available Encodings

An **agent** may query the **encoding-service** to provide a list of all **encoding-representations** known by the service. It does this by invoking the **query-available-encodings** action of the **encoding-service**.

1428

1429

If the action is successful, the **encoding-service** will return an **action-status** indicating success. Additionally, the available **encoding-representations** are supplied.

1430

1431

1432

1433

1434

1435

If the action is unsuccessful, the **encoding-service** will return an **action-status** indicating failure, together with an **explanation**. The range of possible explanations is, in general, specific to a particular reification. However a conforming reification must, where appropriate, distinguish between the following explanations:

1436

1437

Access. The **agent** making the request is not authorized to perform the specified action.

1438

5.12.4 Description

1439

1440

1441

1442

A concrete specification must realize a reification of the **encoding-service** in order that **agents** can encode and decode **encoding-representations** from and into a **message** representation, respectively. Every individual **encoding-representation** will require a specific codec for transforming to and from any **message** and **content** representation.

1443

5.13 Envelope

1444

5.13.1 Summary

1445

1446

1447

1448

1449

1450

An **envelope** is a **key-value tuple** that contains message delivery and encoding information. It is included in the **transport-message**, and includes elements such as the sender and receiver(s) **transport-descriptions**. It also contains the **encoding-representation** for the **message** and optionally, other message information such as validation and security data, or additional routing data. A concrete instantiation of **envelope** is a mandatory element of every concrete instantiation of the abstract architecture.

1451

5.13.2 Relationship to Other Elements

1452

1453

1454

1455

1456

1457

1458

Envelope contains **transport-descriptions**

Envelope optionally contains validity data (such as security keys for message validation)

Envelope optionally contains security data (such as security keys for message encryption or decryption)

Envelope optionally contains routing data

Envelope contains an **encoding-representation** for the **payload** being transported

Envelope is contained in **transport-message**

1459

5.13.3 Description

1460

1461

1462

1463

In the realization of the envelope data, the realization can specify envelope elements that are useful in the particular realization. These can include specialized routing data, security related data, or other data that can assist in the proper handling of the encoded **message**.

1464

5.14 Explanation

1465

5.14.1 Summary

1466

1467

1468

1469

An encoding of the reason for a particular **action-status**. When an action exerted by a service leads to a failure response, the **explanation** is an optional descriptor giving the reason why the particular action failed. A concrete instantiation of **explanation** is an optional element of every concrete instantiation of the abstract architecture.

1470

5.14.2 Relationship to Other Elements

1471

1472

Explanation qualifies an **action-status**.

1473 **5.14.3 Description**

1474 In terms of the three explicit services described by the abstract architecture, the **agent-directory-service**, **service-**
 1475 **directory-service** and **message-transport-service**, the relevant action **explanations** are listed in the appropriate
 1476 element subsections.

1477

1478 **5.15 Message**1479 **5.15.1 Summary**

1480 A **message** is an individual unit of communication between two or more **agents**. A **message** logically arises from and
 1481 programmatically corresponds to a communicative act, in the sense that a **message** encodes the communicative act.
 1482 Communicative acts can be recursively composed, so while the outermost act is directly encoded by the **message**,
 1483 taken as a whole a given **message** may represent multiple individual communicative acts. This is then encoded using
 1484 an **encoding-representation** and transmitted between **agents** over a **transport**. A **message** includes an indication of
 1485 the type of communicative act (for example, INFORM, REQUEST), the **agent-names** of the sender and receiver
 1486 **agents**, the **ontology** or **ontologies** to be used in interpreting the **content**, and the **content** of the **message** itself. A
 1487 **message** does not include any transport or addressing information. It is transmitted from sender to receiver(s) by
 1488 being encoded as the **payload** of a **transport-message**, which includes this information. A concrete instantiation of
 1489 **message** is a mandatory element of every concrete instantiation of the abstract architecture.

1490

1491 **5.15.2 Relationships to other elements**

1492 **Message** is written in an **agent-communication-language**

1493 **Message** contains **content**

1494 **Message** has an **ontology** attribute

1495 **Message** includes an **agent-name** corresponding to the sender of the message

1496 **Message** includes one or more **agent-name** corresponding to the receiver or receivers of the message

1497 **Message** is sent by an **agent**

1498 **Message** is received by one or more **agents**

1499 **Message** is transmitted as the **payload** of a **transport-message**

1500 **Message** is transformed to/from a **payload** by an **encoding-service**

1501

1502 **5.15.3 Description**

1503 The FIPA communicative acts library is described in detail in [FIPA00037].

1504

1505 **5.16 Message Transport Service**1506 **5.16.1 Summary**

1507 A **message-transport-service** is a **service**. It supports the sending and receiving of **transport-messages** between
 1508 **agents**. A concrete instantiation of **message-transport-service** is a mandatory element of every concrete instantiation
 1509 of the abstract architecture.

1510

1511 **5.16.2 Relationships to Other Elements**

1512 **Message-transport-service** may be invoked to send a **transport-message** to an **agent**

1513 **Message-transport-service** selects a **transport** based on the recipient's **transport-description**

1514 **Message-transport-service** is a **service**

1515

1516 **5.16.3 Actions**

1517 A **message-transport-service** supports the following actions.

1518

1519 5.16.3.1 Bind Transport

1520 An **agent** may form a contract with the **message-transport-service** to send and receive messages using a particular
 1521 **transport**. It does this by invoking the **bind-transport** action of the **message-transport-service**. The semantics of
 1522 this action are as follows:

1523

1524 The **agent** provides a **transport-description** corresponding to the **transport** to be used. (In initiating the action, the
 1525 **agent** may control the scope of the action. Different reifications may handle this in different ways. The action may be
 1526 addressed to a particular instance of a **agent-directory-service**, or the action may be qualified with some kind of
 1527 scope parameter.) Some or all of the elements of the **transport-description** may be missing, in which case the
 1528 **transport-service** may supply appropriate values. The **transport-service** attempts to create a usable transport-end-
 1529 point for the chosen **transport-type**, and constructs a **transport-specific-address** corresponding to this end-point.

1530

1531 If the action is successful, the **message-transport-service** will return an **action-status** indicating such, together with a
 1532 **transport-description** that has been completely filled in and is usable for message transport. The agent may use this
 1533 **transport-description** as part of its **agent-description**, and in constructing a **transport-message**.

1534

1535 Following a successful **bind-transport**, the **message-transport-service** will attempt to deliver any messages received
 1536 over the transport end-point to the **agent**.

1537

1538 If the action is unsuccessful, the **message-transport-service** will return an **action-status** indicating failure, together
 1539 with an **explanation**. The range of possible explanations is, in general, specific to a particular reification. However a
 1540 conforming reification must, where appropriate, distinguish between the following explanations:

1541

1542 *Access.* The **agent** making the request is not authorized to perform the specified action.

1543

1544 *Invalid.* The **transport-description** is invalid in some way.

1545

1546 5.16.3.2 Unbind Transport

1547 An **agent** may terminate a contract with the **message-transport-service** to send and receive messages using a
 1548 particular **transport**. It does this by invoking the **unbind-transport** action of the **message-transport-service**. The
 1549 semantics of this action are as follows:

1550

1551 The **agent** provides a **transport-description** returned by a previous **bind-transport** action. (In initiating the action, the
 1552 **agent** may control the scope of the action. Different reifications may handle this in different ways. The action may be
 1553 addressed to a particular instance of a **agent-directory-service**, or the action may be qualified with some kind of
 1554 scope parameter.) The **transport-service** identifies the corresponding transport-end-point and releases all transport
 1555 related resources.

1556

1557 If the action is successful, the **message-transport-service** will return an **action-status** indicating success.
 1558 Additionally, the **message-transport-service** will no longer attempt to deliver any messages to the **agents** associated
 1559 with the defunct transport binding.

1560

1561 If the action is unsuccessful, the **message-transport-service** will return an **action-status** indicating failure, together
 1562 with an **explanation**. The range of possible explanations is, in general, specific to a particular reification. However a
 1563 conforming reification must, where appropriate, distinguish between the following explanations:

1564

1565 *Not-found.* The **transport-description** does not correspond to a bound **transport**.

1566

1567 *Access.* The **agent** making the request is not authorized to perform the specified action.

1568

1569 *Invalid.* The **transport-description** is invalid in some way.

1570

5.16.3.3 Send Message

An **agent** may send a **transport-message** to another agent by invoking the **send-message** action of a **message-transport-service**. The semantics of this action are as follows:

The **agent** provides a **transport-message** to be sent. The **message-transport-service** examines the **envelope** of the message to determine how it should be handled.

If the action is successful, the **message-transport-service** will return an **action-status** indicating success. Following a successful **send-message**, the **message-transport-service** will make attempt to deliver the message to the recipient. However the successful completion of the **send-message** action should not be interpreted as indicating that delivery has been achieved.

If the action is unsuccessful, the **message-transport-service** will return an **action-status** indicating failure, together with an **explanation**. The range of possible explanations is, in general, specific to a particular reification. However a conforming reification must, where appropriate, distinguish between the following explanations:

Access. The **agent** making the request is not authorized to perform the specified action.

Invalid. The **transport-message** is invalid in some way.

5.16.3.4 Deliver Message

A **message-transport-service** may deliver a **transport-message** to an **agent** by invoking the **deliver-message** action of the **agent**. The semantics of this action are identical to those given for the **bind-transport** action.

5.16.4 Description

A concrete specification need not realize the notion of **message-transport-service** so long as the basic service provisions are satisfied. In the case of a concrete specification based on a single **transport**, the agent may use native operating system services or other mechanisms to achieve this service.

5.17 Ontology

5.17.1 Summary

An **Ontology** provides a vocabulary for representing and communicating knowledge about some topic and a set of relationships and properties that hold for the entities denoted by that vocabulary. A concrete instantiation of **ontology** is an optional element of concrete instantiations of the abstract architecture.

5.17.2 Relationships to Other Elements

Message has an **ontology** attribute that can contain references to one or more ontologies

Content is expressed in the context of one or more ontologies using the **ontology** message attribute

5.17.3 Description

An **ontology** is a set of symbols together with an associated interpretation that may be shared by a community of **agents** or **services**. An **ontology** includes a vocabulary of symbols referring to objects and relationships in the subject domain. An **ontology** also typically includes a set of logical statements expressing the constraints existing in the domain and restricting the interpretation of the vocabulary.

Ontologies must be nameable, discoverable and manageable.

1618 5.18 Payload

1619 5.18.1 Summary

1620 A **payload** is a **message** encoded in a manner suitable for inclusion in a **transport-message**. A concrete instantiation
1621 of **payload** is a mandatory element of every concrete instantiation of the abstract architecture.
1622

1623 5.18.2 Relationships to Other Elements

1624 **Payload** is an encoded **message**

1625 **Transport-message** contains a **payload**

1626 **Payload** is encoded according to an **encoding-representation**
1627

1628 5.18.3 Description

1629 See *Section 5.33.2, Relationships to Other Elements* which describes the **transport-message** element.
1630

1631 5.19 Service

1632 5.19.1 Summary

1633 A **service** is a functional coherent set of mechanisms that support the operation of **agents**, and other **services**. These
1634 are services used in the provisioning of *agent environments* and may be used as the basis for interoperation. A
1635 concrete instantiation of **service** is a mandatory element of every concrete instantiation of the abstract architecture.
1636

1637 Note: A service in this specification should not be confused with the service or services provided by agents
1638 implemented within instantiations of the architecture.
1639

1640 5.19.2 Relationships to Other Elements

1641 **Service** has a public set of behaviours and actions

1642 **Service** has a service description

1643 **Service** can be accessed by **agents**

1644 **Agent-directory-service** is an instance of **service**, and is mandatory

1645 **Message-transport-service** is an instance of **service**, and is mandatory

1646 **Service-directory-service** is an instance of **service**, and is mandatory

1647 A **service** has a **service-type**, a **service-id**, a **service-locator**

1648 A **service** can have a **service-directory-entry** in a **service-directory-service** containing the **service-id**, **service-type**
1649 and **service-locator**
1650

1651 5.19.3 Description

1652 FIPA will administer the name space of **services** according to the description given in Section 5.1.2. This is part of the
1653 concrete realization process. Having a clear naming scheme for the **services** will allow for optimised implementation
1654 and management of **services**.
1655

1656 5.20 Service Address

1657 5.20.1 Summary

1658 A **service-type** specific string that indicates how to bind to a particular **service**. A concrete instantiation of **service-**
1659 address is a mandatory element of every concrete instantiation of the abstract architecture.

1660 5.20.2 Relationships to Other Elements

1661 **Service-address** provides an address of a **service** that can be bound to by an **agent** or **service**

1662 **Services-locators** contain one or more **service-addresses**

1663 A **service-address** is qualified by a **signature-type**

1665 5.20.3 Description

1666 The **service address** is a **service-type** specific string that indicates how to bind to a **service**. The precise means by
 1667 which this binding is made is implementation and **service-type** specific; for example a **transport-service** that is bound
 1668 via RMI objects may give an RMI address of the Java object to bind to and thereby access the **transport-service**.
 1669 Alternatively, an **agent-directory-service** that is accessed via a TCP/IP socket may give a string containing the
 1670 hostname and port number.

1672 5.21 Service Attributes

1673 5.21.1 Summary

1674 **Service-attributes** are optional attributes that are part of the **service-directory-entry** for a **service**. They are
 1675 represented as **key-value-pairs** within the **key-value-tuple** that makes up the **service-directory-entry**. The purpose
 1676 of the attributes is to allow searching for **service-directory-entries** that match **services** of interest. A concrete
 1677 instantiation of **service-attributes** is an optional element of concrete instantiations of the abstract architecture.

1679 5.21.2 Relationships to Other Elements

1680 A **service-directory-entry** may have zero or more **service-attributes**

1681 **Service-attributes** describe aspects of a **service**

1683 5.21.3 Description

1684 When a **service** registers a **service-directory-entry**, the **service-directory-entry** may optionally contain **key-value-**
 1685 **pairs** that offer additional description of the **service**. The values might include information about costs of using the
 1686 **service**, features available, **ontologies** understood, names that the **service** is commonly known by, or any other
 1687 relevant data. This information can then be used to enhance the search criteria by which **services** are discovered in
 1688 the **service-directory-service**.

1689 In practice, when defining realizations of this abstract architecture, domain specific specifications should exist
 1690 describing the **service-attributes** to be used. This eases requirements for interoperation.

1693 5.22 Service Directory Entry

1694 5.22.1 Summary

1695 A **service-directory-entry** is a **key-value-tuple** consisting of a **service-id**, **service-type**, **service-locator** and zero or
 1696 more **service-attributes**. A concrete instantiation of **service-directory-entry** is a mandatory element of every
 1697 concrete instantiation of the abstract architecture.

1699 5.22.2 Relationships to Other Elements

1700 **Service-directory-entry** contains the **service-id** of the **service** to which it refers

1701 **Service-directory-entry** contains the **service-type** of the **service** to which it refers

1702 **Service-directory-entry** contains a **service-locator** of the **service** to which it refers

1703 **Service-directory-entry** may contain zero or more **service-attributes**

1704 **Service-directory-entry** is managed by and available from a **service-directory-service**

1705 **Services** are not required to publish a **service-directory-entry**

1707 5.22.3 Description

1708 A **service-directory-entry** is used to describe the identity, type, signature and address of a **service**, which is
 1709 accessed via programmatic means. A **service-directory-entry** also contains zero or more attribute value pairs, which
 1710 are used to distinguish on instance of a service from another. **Services** are registered to a **service-directory-service**
 1711 by adding a **service-directory-entry** to the directory.

1712
 1713 Different realizations that use a common **service-directory-service**, are strongly encouraged to adopt a common
 1714 schema for storing **service-directory-entries**.

1716 5.23 Services Directory Service

1717 5.23.1 Summary

1718 The **service-directory-service** is used to register and locate **services** within the FIPA infrastructure. Services
 1719 include, but are not limited to: **message-transport-services**, **agent-directory-services**, gateway services, and
 1720 message buffering services (note that the latter two services are not mandated by this specification). A **service-**
 1721 **directory-service** is also used to store the **service** descriptions of application oriented services, such as commercial
 1722 and business oriented services. A concrete instantiation of **service-directory-service** is a mandatory element of every
 1723 concrete instantiation of the abstract architecture.

1724
 1725 Note: Agents are not expected to register services in the **services-directory-service** which are not being used in
 1726 explicit provision of services for the platform. In addition, it would be expected that most services would not be register
 1727 by agents.

1728 5.23.2 Relationships to Other Elements

1729 **Service-directory-services** provides a directory of **service-directory-entries**

1730 **Services** may be registered within the **service-directory-service**.

1731 **Service-directory-service** is a **service**

1733 5.23.3 Description

1734 Each concrete implementation of this specification will provide a **service-directory-service**. The **service-directory-**
 1735 **service** will provide a simple registry for the **service** descriptions. Each realization of the **service-directory-service**
 1736 will provide agents with a **service-root**, which will take the form of a set of **service-locators** including at least one
 1737 **service-directory-service** (pointing to itself) In general, a **service-root** will provide sufficient entries to either describe
 1738 all of the services available within the environment directly, or it will provide pointers to further services which will
 1739 describe these services.

1740 5.23.4 Actions

1741 5.23.4.1 Register

1742 A service may **register** a **service** description in the form of a **service-directory-entry** with a **service-directory-**
 1743 **service**.

1744
 1745 The semantics of this action are as follows:

1746
 1747 The **service** provides a **service-directory-entry** that is to be registered. In initiating the action, the **service** may
 1748 control the scope of the action. Different reifications may handle this in different ways. The action may be addressed to
 1749 a particular instance of a **service-directory-service**, or the action may be qualified with some scope parameter.

1751 If the action is successful, the **service-directory-service** will return an **action-status** indicating success. Following a
 1752 successful **register**, the **service-directory-service** will support legal **delete**, and **query** actions with respect to the
 1753 registered **service-directory-entry**.

1754
 1755 If the action is unsuccessful, the **service-directory-service** will return an **action-status** indicating failure, together with
 1756 an **explanation**. The range of possible explanations is, in general, specific to a particular reification. However a
 1757 conforming reification must, where appropriate, distinguish between the following explanations:

1758
 1759 Duplicate – the new entry “clashed” with some existing **service-directory-entry**.

1760
 1761 Access – the **agent** or **service** making the request is not authorized to perform the specified action.

1762
 1763 Invalid – the **service-directory-entry** is invalid in some way.
 1764

1765 5.23.4.2 Delete

1766 A **service** may **delete** a **service-directory-entry** from a **service-directory-service**. The semantics of this action are
 1767 as follows:

1768
 1769 The **service** provides a **service-directory-entry** which has the same **service-id** as that which is to be deleted. (The
 1770 rest of the **service-directory-entry** is not significant.) In initiating the action, the **service** may control the scope of the
 1771 action. Different reifications may handle this in different ways. The action may be addressed to a particular instance of
 1772 a **service-directory-service**, or the action may be qualified with some scope parameter.

1773
 1774 If the action is successful, the **service-directory-service** will return an **action-status** indicating success. Following a
 1775 successful **delete**, the **service-directory-service** will no longer support **modify**, **delete**, and **query** actions with
 1776 respect to the deleted **service-directory-entry**.

1777
 1778 If the action is unsuccessful, the **service-directory-service** will return an **action-status** indicating failure, together with
 1779 an **explanation**. The range of possible explanations is, in general, specific to a particular reification. However a
 1780 conforming reification must, where appropriate, distinguish between the following explanations:

1781
 1782 Not-found – the new entry did not match any existing **service-directory-entry**. This would only occur if no existing
 1783 **service-directory-entry** had the same **service-id**

1784
 1785 Access – the **agent** or **service** making the request is not authorized to perform the specified action.

1786
 1787 Invalid – the **service-directory-entry** is invalid in some way.
 1788

1789 5.23.4.3 Query

1790 A **service** or **agent** may **query** a **service-directory-service** to locate **service-directory-entries** of interest. The
 1791 semantics of this action are as follows:

1792
 1793 The querying entity (**agent**) provides a **service-directory-entry** that is to be treated as a search pattern. In initiating
 1794 the action, the **agent** may control the scope of the action. Different reifications may handle this in different ways. The
 1795 action may be addressed to a particular instance of a **service-directory-service**, or the action may be qualified with
 1796 some scope parameter.

1797
 1798 The directory service verifies that the argument is a valid **service-directory-entry**. It then searches for registered
 1799 **service-directory-entries** that satisfy the search criteria. A registered entry satisfies the search criteria if there is a
 1800 match between each **key-value pair** in the submitted entry. The semantics of “matching” are likely to be reification-
 1801 dependent; at a minimum, there should be support for matching on the *same* value and on *any* value.

1802
 1803 If the action is successful, the **service-directory-service** will return an **action-status** indicating success, together with
 1804 a set of **service-directory-entries** that satisfy the search pattern. The mechanism by which multiple entries are

1805 returned, and whether or not an **agent** may limit or terminate the delivery of results, is not defined in the abstract
 1806 architecture and is therefore reification dependent.

1807
 1808 If the action is unsuccessful, the **service-directory-service** will return an **action-status** indicating failure, together with
 1809 an **explanation**. The range of possible explanations is, in general, specific to a particular reification. However a
 1810 conforming reification must, where appropriate, distinguish between the following explanations:

1811 Not-found – the search pattern did not match any existing **service-directory-entry**.

1812
 1813 Access – the **agent** or **service** making the request is not authorized to perform the specified action.

1814
 1815 Invalid – the **service-directory-entry** is invalid in some way.
 1816
 1817

1818 5.23.4.4 Modify

1819 A **service** may **modify** a **service-directory-entry** that has been registered with a **service-directory-service**. The
 1820 semantics of this action are as follows:

1821
 1822 The **service** provides a **service-directory-entry** which contains the same **service-id** as the entry to be modified. In
 1823 initiating the action, the **service** may control the scope of the action. Different reifications may handle this in different
 1824 ways. The action may be addressed to a particular instance of a **service-directory-service**, or the action may be
 1825 qualified with some scope parameter.
 1826

1827 The **service-directory-service** verifies that the argument is a valid **service-directory-entry**. It then searches for a
 1828 registered **service-directory-entry** with the same **service-id**. If it does not find one, the action fails and an
 1829 **explanation** provided. Otherwise it modifies the existing **service-directory-entry** by examining each **key-value-pair**
 1830 in new **service-directory-entry**. If the **value** is non-null, the **key-value-pair** is added to the new entry, replacing any
 1831 existing **key-value-pair** with the same **key** identity. If the **value** is null, any existing **key-value-pair** with the same **key**
 1832 identity is removed from the entry.
 1833

1834 If the action is successful, the **service-directory-service** will return an **action-status** indicating success, together with
 1835 a **service-directory-entry** corresponding to the new contents of the registered entry. Following a successful **modify**,
 1836 the **service-directory-service** will support legal **modify**, **delete**, and **query** actions with respect to the modified
 1837 **service-directory-entry**.

1838 If the action is unsuccessful, the **service-directory-service** will return an **action-status** indicating failure, together with
 1839 an **explanation**. The range of possible explanations is, in general, specific to a particular reification. However a
 1840 conforming reification must, where appropriate, distinguish between the following explanations:

1841 Not-found – the new entry did not match any existing **service-directory-entry**. This would only occur if no existing
 1842 **service-directory-entry** had the same **service-id**

1843
 1844 Access – the **agent** or **service** making the request is not authorized to perform the specified action.

1845
 1846 Invalid – the new **service-directory-entry** is invalid in some way.
 1847
 1848

1849 5.24 Service Id

1850 5.24.1 Summary

1851 The **service-id** provides uniqueness preservation within a given namespace. The **service-id** is used to test for
 1852 equivalence of a **service**, and for modifying, deleting and searching for **service-directory-entries** within a **service-**
 1853 **directory-service**. **Service-ids** are unique, and are intended only to be used to test for uniqueness and identity, not to
 1854 provide location or other extrinsic properties of the service. A concrete instantiation of **service-id** is a mandatory
 1855 element of every concrete instantiation of the abstract architecture.

1856 5.24.2 Relationships to other elements

1857 **Service-id** is used to identify a **service** within a **service-directory service**
 1858 **Service-id** is a component of a **service-directory entry**.

1859 5.24.3 Description

1860 A **service-id** is an immutable identifier (e.g. a GUID, Globally Unique Identifier) that is associated with the **service** at
 1861 creation time or initial registration. Name issuing should occur in a way that tends to ensure global uniqueness. This
 1862 may be achieved, for example, through employing an algorithm that generates the name with a sufficient degree of
 1863 stochastic complexity such as to induce a vanishingly small chance of a name collision.
 1864

1865 5.25 Service Location Description

1866 5.25.1 Summary

1867 A **service-location-description** is a set of one or more **key-value tuples**, each containing a **signature-type**, **service-**
 1868 **signature** and a **service-address**. In general, any **agent** or **service** wishing to use the **service** must 'already know'
 1869 how to operate the service. In particular, the **service-address** should be a data value of type known both to the agent
 1870 that it may use to invoke actions from the service. A concrete instantiation of **service-location-description** is a
 1871 mandatory element of every concrete instantiation of the abstract architecture.

1872 5.25.2 Relationships to Other Elements

1873 **Service-locator** contains one or more **service-location-descriptions**
 1874 **Service-location-description** contains **signature-type**
 1875 **Service-location-description** contains **service-signature**
 1876 **Service-location-description** contains **service-address**
 1877 **Service-location-description** is used by an **agent** to access a **service**
 1878

1879 5.25.3 Description

1880 A **service-location-description** is the parallel structure to a **transport-description** (which is a component of the
 1881 **agent-locator**), that describes how to access a **service**. Each **service-location-description** contains a **service-**
 1882 **signature** that defines how to call the service, a **signature-type** that type classifies the **service-signature** and a
 1883 **service-address** that identifies the addressable location of the **service**.
 1884

1885 5.26 Service Locator

1886 5.26.1 Summary

1887 A **service-locator** consists of the set of **service-location-descriptions**, which can be used to access and make use
 1888 of a **service**. In general, any **agent** or **service** wishing to use the **service** must 'already know' how to operate the
 1889 service. In particular, the **service-address** should be a data value of type known both to the agent that it may use to
 1890 invoke actions from the service. A concrete instantiation of **service-locator** is a mandatory element of every concrete
 1891 instantiation of the abstract architecture.
 1892

1893 5.26.2 Relationships to Other Elements

1894 **Service-locator** is a member of **service-directory-entry**, which is registered with a **service-directory-service**
 1895 **Service-locator** contains one or more **service-location-descriptions**
 1896 **Service-locator** is used by an **agent** to access a **service**
 1897

1898 5.26.3 Description

1899 A **service-locator** is the parallel structure to an **agent-locator**, which describes how to access a **service**. Each
 1900 **service-locator** includes all of the **service-location-descriptions** that may be used to access the associated **service**.
 1901

1902 5.27 Service Root

1903 5.27.1 Summary

1904 A **service-root** is a set of **service-directory-entries** made available to an **agent** at start-up. This is the mechanism by
 1905 which an **agent** can bootstrap lifecycle support services, such as **message-transport-services** and **agent-directory-**
 1906 **services**, to provide it with a connection to the outside environment. A concrete instantiation of **service-root** is a
 1907 mandatory element of every concrete instantiation of the abstract architecture.
 1908

1909 5.27.2 Relationships to Other Elements

1910 **Service-root** is used by an **agent** to bootstrap **services**

1911 **Service-root** is a set of **service-directory-entries**

1912 **Service-root** should contain a **service-directory-entry** for at least one **message-transport-service**

1913 **Service-root** should contain a **service-directory-entry** for at least one **agent-directory-service**

1914 **Service-root** should contain a **service-directory-entry** for at least one **service-directory-service**
 1915

1916 5.27.3 Description

1917 An **agent** must be provided with a **service-root** at initialization in order for it to be able to communicate with other
 1918 **agents** and **services**. Typically the provider of the **service-root** will be a **service-directory-service** which will supply
 1919 a set of service descriptions in the form of **service-directory-entries** for available agent lifecycle support services,
 1920 such as **message-transport-services**, **agent-directory-services** and **service-directory-services**. In general, a
 1921 **service-root** will provide sufficient entries to either describe all of the services available within the environment
 1922 directly, or it will provide pointers to further services which will describe these services.
 1923

1924 5.28 Service Signature

1925 5.28.1 Summary

1926 A **service-signature** is a Fully Qualified Name within an administered namespace that describes the binding signature
 1927 for a service. A concrete instantiation of **service-signature** is a mandatory element of every concrete instantiation of
 1928 the abstract architecture.

1929 5.28.2 Relationships to Other Elements

1930 **Service-signature** is a component of a **service-locator**

1931 **Service-signature** is qualified in terms of a **signature-type**
 1932

1933 5.28.3 Description

1934 Examples of **service-signatures** are:

1935
 1936 org.fipa.standard.service.java-rmi-binding

1937 org.omg.agent.idl-binding
 1938

1939 See **signature-type** for a description of these **service-signature** bindings.
 1940

1941 **5.29 Service Type**1942 **5.29.1 Summary**

1943 A **service-type** is a **key-value-tuple**, defining the *type* of a **service**. The set of possible values will be administered,
 1944 according to the process defined for **key-value-tuples** and by the appropriate namespace authority. A concrete
 1945 instantiation of **service-type** is a mandatory element of every concrete instantiation of the abstract architecture.
 1946

1947 **5.29.2 Relationships to Other Elements**

1948 **Service-type** is a component of a **service-directory-entry**

1949 **Service-type** qualifies the *type* of a **service**

1950

1951 **5.29.3 Description**

1952 **Service-type** is used to classify the **service** in terms of some administered namespace. The *type* provides a
 1953 contextual reference to **service** functionality. For example, the **service-address** component of the **service-locator**
 1954 uses **service-type** as a context for communication bindings.
 1955

1956 **5.30 Signature Type**1957 **5.30.1 Summary**

1958 A **signature-type** is a **key-value-tuple** describing the *type* of a **service-signature**. A **signature-type** allows the
 1959 interpretation of a **service-locator**, by associating it with a type of method signature binding. A concrete instantiation of
 1960 **signature-type** is an optional element of concrete instantiations of the abstract architecture.

1961 **5.30.2 Relationships to Other Elements**

1962 **Signature-type** is a component of a **service-locator**

1963 **Signature-type** qualifies the *type* of a **service-signature**

1964 **Signature-type** qualifies the *type* of a **service-address**

1965

1966 **5.30.3 Description**

1967 The **signature-type** keys access to the opaque portion of a **service-locator**. Examples of signatures are:

1968 5.30.3.1.1 org.fipa.standard.service.java-rmi -binding

1969 For this **signature-type**, the **service-signature** is the Java IDL of the Java method to be invoked and the **service-**
 1970 **address** is the URL for the target of the remote method invocation.
 1971

1972 5.30.3.1.2 org.omg.agent.idl-binding

1973 For this **signature-type**, the **service-signature** is the OMG CORBA IDL of the method to be invoked and the **service-**
 1974 **address** is the IOR of the remote object which is the target of the method invocation.
 1975

1976 **5.31 Transport**1977 **5.31.1 Summary**

1978 A **transport** is a particular **message** delivery service, such as a message transfer system, a datagram service, a byte
 1979 stream, or a shared scratchboard. Abstractly, a **transport** is a delivery system selected by virtue of the **transport-**
 1980 **description** used to deliver **messages** to an **agent**. A concrete instantiation of **transport** is a mandatory element of
 1981 every concrete instantiation of the abstract architecture.
 1982

1983 5.31.2 Relationships to Other Elements

1984 **Transport-description** can be mapped onto a **transport**
 1985 **Message-transport-service** may use one or more **transports** to effect message delivery
 1986 A **transport** may support one or more **transport-encodings**
 1987

1988 5.31.3 Description

1989 The mapping from **transport-description** to **transport** must be consistent across all realizations. FIPA will administer
 1990 ontology of transport names. Concrete specifications should define a concrete encoding for this ontology.
 1991

1992 5.32 Transport Description

1993 5.32.1 Summary

1994 A **transport-description** is a **key-value tuple** containing a **transport-type**, a **transport-specific-address** and zero or
 1995 more **transport-specific-properties**. A concrete instantiation of **transport-description** is a mandatory element of
 1996 every concrete instantiation of the abstract architecture.
 1997

1998 5.32.2 Relationships to Other Elements

1999 **Transport-description** has a **transport-type**
 2000 **Transport-description** has a set of **transport-specific-properties**
 2001 **Transport-description** has a **transport-specific-address**
 2002 **Agent-directory-entries** include one or more **transport-descriptions**
 2003 **Envelopes** contain one or more **transport-descriptions**
 2004

2005 5.32.3 Description

2006 **Transport-descriptions** are included in the **agent-directory-service**, describing where a registered agent may be
 2007 contacted. They can be included in the **envelope** for a **transport-message**, to describe how to deliver the message. In
 2008 addition, if a **message-transport-service** is implemented, **transport-descriptions** are used as input to the **message-**
 2009 **transport-service** to specify characteristics for additional delivery requirements for the delivery of **messages** to an
 2010 **agent**.

2011 5.33 Transport Message

2012 5.33.1 Summary

2013 A **transport-message** is the object conveyed from **agent** to **agent**. It contains the **envelope** containing **transport-**
 2014 **descriptions** for the sender and receiver(s) together with a **payload** containing the encoded **message**. A concrete
 2015 instantiation of **transport-message** is a mandatory element of every concrete instantiation of the abstract architecture.
 2016

2017 5.33.2 Relationships to Other Elements

2018 **Transport-message** contains a **payload**
 2019 **Transport-message** contains an **envelope**
 2020

2021 5.33.3 Description

2022 A concrete implementation may limit the number of receiving **transport-descriptions** in the **envelope** of a **transport-**
 2023 **message**. It may also establish particular relationships between the **agent-name** or **agent-names** for the receiver(s) in
 2024 the **payload**. For example, it may ensure that there is a one-to-one correspondence between **agent-names**. The
 2025 important thing to convey from **agent** to **agent** is the **payload**, together with sufficient **transport-message** context to

2026 send a reply. A gateway service or other transformation mechanism may unpack and reformat a **transport-message**
 2027 as part of its processing.
 2028

2029 **5.34 Transport Specific Address**

2030 **5.34.1 Summary**

2031 A **transport-specific-address** is an address specific to a particular **transport-type**. The format and description of the
 2032 address will be specific to this type. The address is used by a **transport-service** in conjunction with a **transport-type**
 2033 to construct transport connections. A concrete instantiation of **transport-specific-address** is an mandatory element of
 2034 every concrete instantiation of the abstract architecture.
 2035

2036 **5.34.2 Relationships to Other Elements**

2037 A **transport-specific-address** is a component of a **transport-description**.
 2038 A **transport-specific-address** is associated with a specific **transport-type**.
 2039

2040 **5.34.3 Description**

2041 The **transport-specific-address** provides a resolvable location descriptor, specific to a given **transport-type**, which
 2042 can be used by a **transport-service** to send and/or receive **messages**.
 2043

2044 **5.35 Transport Specific Property**

2045 **5.35.1 Summary**

2046 A **transport-specific-property** is property associated with a **transport-type**. These properties are used by the
 2047 **transport-service** to help it in constructing transport connections, based on the properties specified. A concrete
 2048 instantiation of **transport-specific-property** is an optional element of every concrete instantiation of the abstract
 2049 architecture.
 2050

2051 **5.35.2 Relationships to Other Elements**

2052 **Transport-description** includes zero or more **transport-specific-properties**
 2053

2054 **5.35.3 Description**

2055 The **transport-specific-properties** are not required for a particular **transport**. They may vary between **transports**.
 2056

2057 **5.36 Transport Type**

2058 **5.36.1 Summary**

2059 A **transport-type** describes the type of transport associated with a **transport-specific-address**. A concrete
 2060 instantiation of **transport-type** is a mandatory element of every concrete instantiation of the abstract architecture.
 2061

2062 **5.36.2 Relationships to Other Elements**

2063 **Transport-description** includes a **transport-type**
 2064

2065 **5.36.3 Description**

2066 FIPA will administer an **ontology** of **transport-types**. FIPA managed types will be flagged with the prefix of "FIPA-".
2067 Specific realizations can provide experimental types, which will be prefixed "X-"
2068

6 Agent and Agent Information Model

This section of the abstract architecture provides a series of UML class diagrams for key elements of the abstract architecture. In *Section 5, Architectural Elements* you can get a textual description of these elements and other aspects of the relationships between them.

Comment on notation: In UML, the notion of a 1 to many or 0 to many relationship is often noted as "1...*" or "0...*". However, the tool that was used to generate these diagrams used the convention "1...n" and "0...n" to express the concept of many.

6.1 Agent Relationships

Figure 11 outlines the basic relationships between an **agent** and other key elements of the FIPA abstract architecture. In other diagrams in this section are provided details on the **agent-locator**, and the **transport-message**.

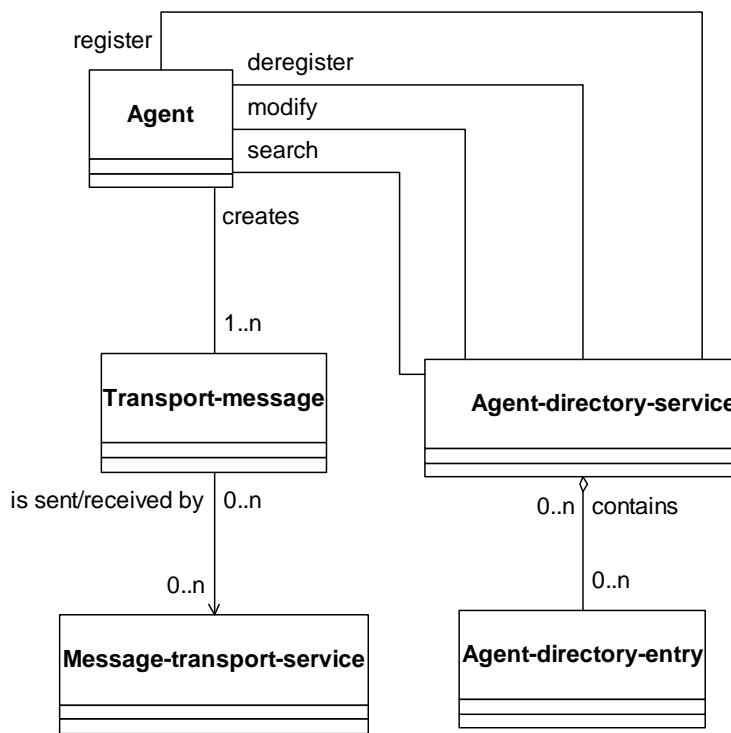
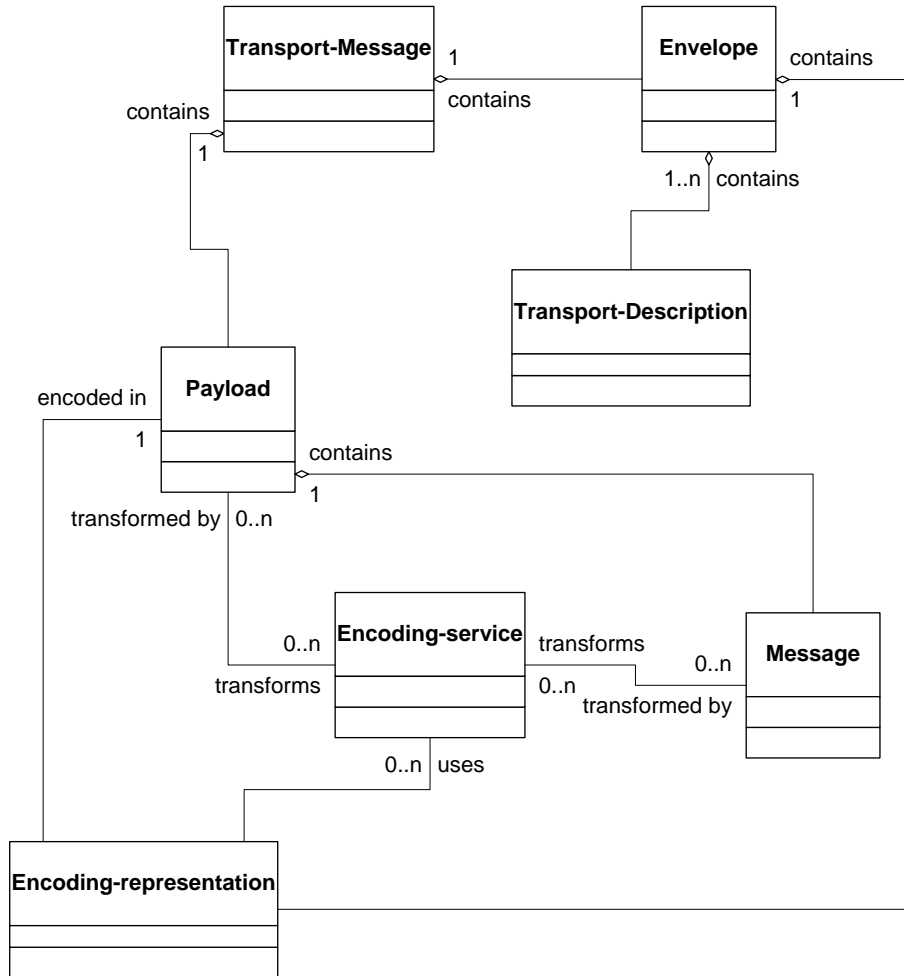


Figure 11: UML - Basic **Agent** Relationships

2085
2086
2087
2088

6.2 Transport Message Relationships

Transport-message is the object conveyed from **agent** to **agent**. It contains the **transport-description** for the sender and receiver or receivers, together with a **payload** containing the **message** (see *Figure 12*).



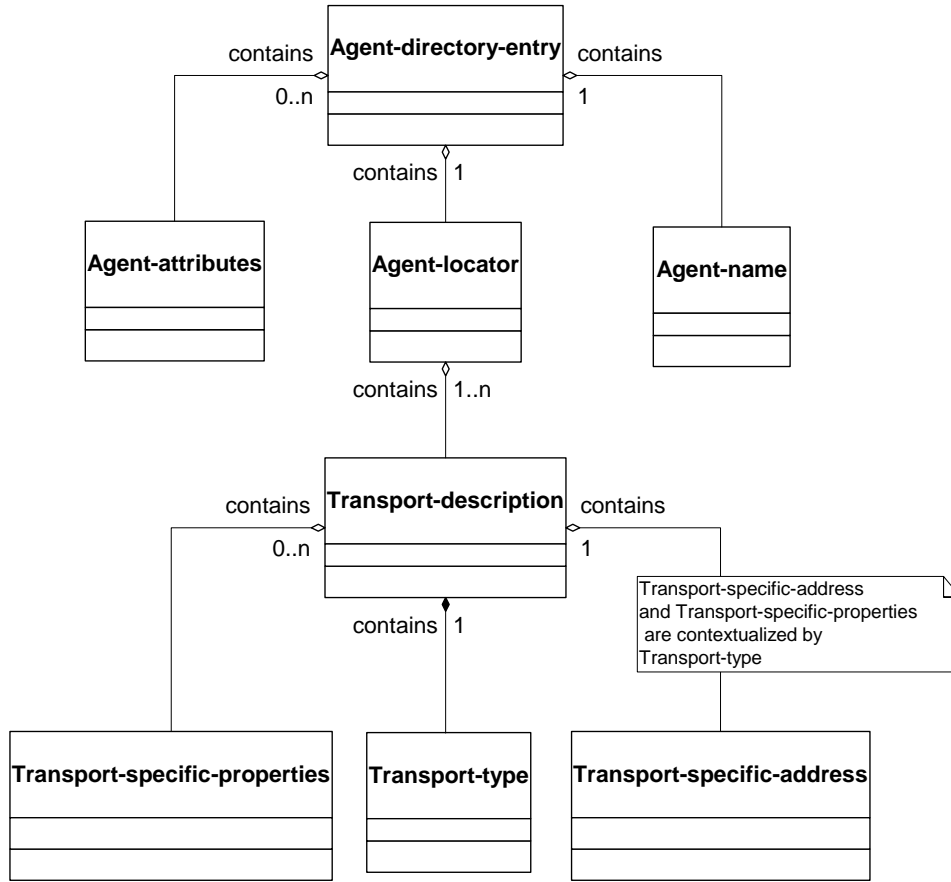
2089
2090
2091
2092

Figure 12: UML - Transport-Message Relationships

2093
2094
2095
2096

6.3 Agent Directory Entry Relationships

The **agent-directory-entry** contains the **agent-name**, **agent-locator** and **agent-attributes**. The **agent-locator** provides ways to address **messages** to an **agent**. It is also used in modifying **transport** requests (see *Figure 13*).



2097
2098
2099
2100
2101

Figure 13: UML - Agent-directory-entry and Agent-locator Relationships

6.4 Service Directory Entry Relationships

Figure 14 shows the hierarchical relationships within a **service-directory-entry** which contains the **service-id**, **service-type** and **service-locator**. The **service-locator** provides the means to contact and make use of a **service** and contains one or more **service-location-descriptions** which in turn each contain a **service-signature**, the **signature-type** and the **service-address**.

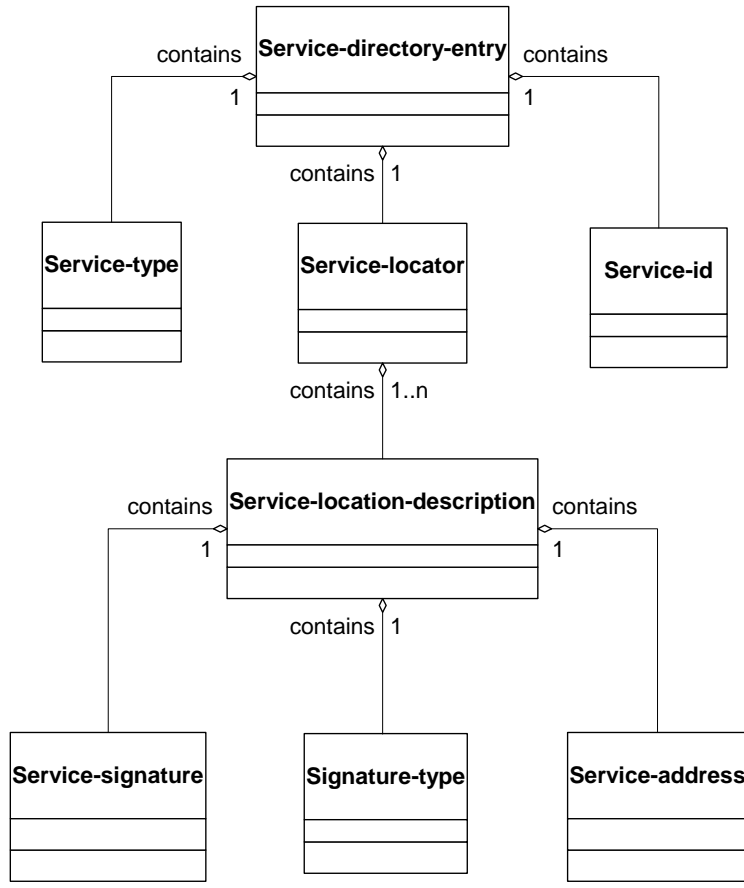


Figure 14: UML - Service-directory-entry and Service-locator Relationships

6.5 Message Elements

Figure 15 shows the elements in a **message**. A **message** is contained in a **transport-message** when messages are sent. Note that in Figure 14, the elements 'Communicative Act' and 'Performative' are not explicit architectural elements defined within this specification; they are informative entities relating to the semantics of the message as defined by the FIPA specification [FIPA00037]. Also, the multiplicity of the 'Ontologies' element refers to the fact more than one ontology may be referred to by the **ontology** architectural element which corresponds to the ACL message attribute 'Ontology' [FIPA00061].

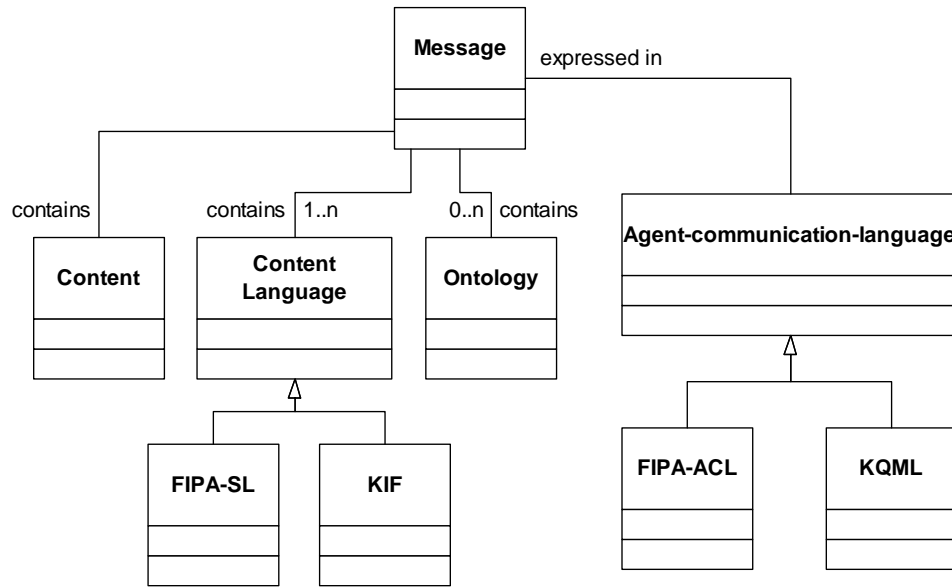
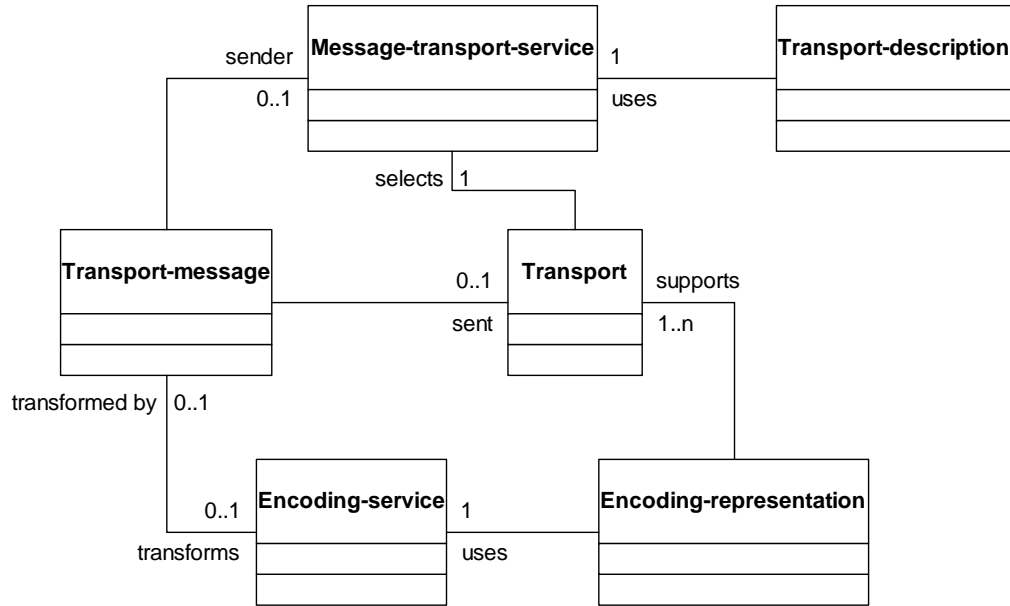


Figure 15: UML - Message Elements

2123
2124
2125
2126

6.6 Message Transport Elements

The **message-transport-service** is an option service that can send **transport-messages** between **agents**. These elements may participate in other relationships as well (see *Figure 16*).



2127
2128
2129
2130
2131

Figure 16: UML - Message-Transport Elements

7 References

2132

2133 [FIPA00007] FIPA Content Language Library Specification. Foundation for Intelligent Physical Agents, 2000.
2134 <http://www.fipa.org/specs/fipa00007/>

2135 [FIPA00023] FIPA Agent Management Specification. Foundation for Intelligent Physical Agents, 2000.
2136 <http://www.fipa.org/specs/fipa00023/>

2137 [FIPA00037] FIPA Communicative Act Library Specification. Foundation for Intelligent Physical Agents, 2000.
2138 <http://www.fipa.org/specs/fipa00037/>

2139 [FIPA00061] FIPA ACL Message Structure Specification. Foundation for Intelligent Physical Agents, 2000.
2140 <http://www.fipa.org/specs/fipa00061/>

2141 [Gamma95] Gamma, Helm, Johnson and Vlissides, Design Patterns. Addison-Wesley, 1995.

2142 [Searle69] Searle, J. L., Speech Acts. Cambridge University Press, 1969.

2143

2144 8 Informative Annex A — Goals of Service Model

2145 8.1 Scope

2146 Agents require the use of many services in order to interoperate with other agents. In order to create the essential
2147 abstractions for the various kinds of services that are essential to this mission, and to permit the straightforward
2148 incorporation of other services in a consistent framework we require a model of services themselves.

2149 8.2 Variety of Services

2150 Although there are a number of essential services required by the abstract architecture, a fully built out platform may
2151 include a wide variety of services not referenced in this document -- for example a platform may provide various kinds
2152 of buffering services. Since the actual services may vary dynamically it is desirable for agents and services to have a
2153 common model for discovering other services.

2154 8.3 Bootstrapping

2155 While the concrete realizations of the Abstract Architecture may have very different forms a common requirement
2156 exists for many systems for a clear and reliable method of bootstrapping services, agents and agent systems.
2157 Supporting bootstrapping is a clear aim of the service model
2158

2159 8.4 Dynamic services

2160 The set of services available to an agent may on some systems be quite fixed: they are made available on start-up and
2161 exist unchanged for the lifetime of the agent. However, on many – if not most – large scale systems, the set of services
2162 available to agents is in fact dynamic. Both the number, type and instantiations of services are all is often subject to
2163 change; for example, the message transport services available to an agent may vary depending on the circumstances.
2164

2165 It is an objective of the service model to provide a consistent framework permitting services themselves to be made
2166 dynamically available: services need to be able to dynamically register themselves, and agents and services may need
2167 to be able to dynamically discover the appropriate services.
2168

2169 8.5 Granularity

2170 An important – if informal – property of the service model is *granularity of services*. For example, it would may be
2171 possible to `break apart` a message transport service into a collection of transports each of which is registered
2172 independently with a service directory service. However, to do so would impose a significant burden on programmers
2173 wishing to make use of message transport: a key benefit of supporting an integrated message transport service is that
2174 it permits high-level convenience operations such as `reply to this message with this new message` or `send a
2175 message to this agent` without requiring a `manual` search of the service directory service each time.
2176

2177 In general the appropriate granularity of services depends on whether a range of related services is best viewed as
2178 instantiations of a single high-level service or whether they are interdependent but distinct kinds of service.
2179

2180 8.6 Example

2181 The following example illustrates how an entry in a service directory service can be formulated.
2182

2183 For our example, we consider locating a prototype buffering service, implemented as Java object. The service, being
2184 experimental, is contained within the name space, “org.fipa.experimental” and has the signature type “fipa-
2185 experimental.buffer-prototype”.
2186

2187 The Java object is accessed via the service address URL: `rmi://testbox.fipa.org/buffertest`
2188

2189 The method signature is:
2190 `public void setBuffer (BufferSessionContext ctx) throws java.rmi.RemoteException`

2191
2192 So, we register the object by generating a service directory entry containing:

```
2193 (service-id, "org.BT.experimental.buffer-prototype.test-1")
2194 (service-type, "org.fipa.experimental.buffer-prototype")
2195 (service-locator, ((signature-type, " org.fipa.service-signature-ontology java2.rmi"),
2196                   (service-signature, "fipa.agentpackages.experimentalbufferpackage"),
2197                   (service-address, "rmi://testbox.Norwich.bt.co.uk/1066/buffertest")))
```

2199
2200 The service-locator contains the signature-type which tells us that we use Java2 RMI to access the service. This tells
2201 us how to understand the next two elements of the locator, the service-signature and service-address. The service-
2202 signature is the Java package which you need to use to get at the methods provided by the buffering object. Finally,
2203 the service-address is the resolvable location at which the appropriate method can be found.

2204
2205
2206

9 Informative Annex B — Goals of Message Transport Service Abstraction

9.1 Scope

In order to create abstractions for the various architectural elements, it is necessary to examine the kinds of systems and infrastructures that are likely targets of real implementations of the abstract architecture. In this section, we examine some of the ways in which concrete messaging and messaging transports may differ. Authors of concrete architectural specifications must take these issues into account when considering what end-to-end assumptions they can safely make. The goals describe below give the reader an understanding of the objectives the authors of the abstract architecture had in mind when creating this architecture.

9.2 Variety of Transports

There are a wide variety of transport services that may be used to convey a message from one agent to another. The abstract architecture is neutral with respect to this variety. For any instantiation of the architecture, one must specify the set of transports that are supported, how new transports are added, and how interoperability is to be achieved. It is permissible for a particular concrete architecture to require that implementations of that architecture must support particular transports.

Different transports use a variety of different address representations. Instantiations of the message transport architecture may support mechanisms for validating addresses, and for selecting appropriate transport services based upon the form of address used. It is extremely undesirable for an agent to be required to parse, decode, or otherwise rely upon the format of an address.

The following are examples of transport services that may be used to instantiate this abstract architecture:

- Enterprise message systems such as those from IBM and Tibco.

- A Java Messaging System (JMS) service provider, such as Fiorano.

- CORBA IIOP used as a simple byte stream.

- Remote method invocation, using Java RMI or a CORBA-based interface.

- SMTP email using MIME encoding.

- XML over HTTP.

- Wireless Access Protocol.

- Microsoft Named Pipes.

9.3 Support for Alternative Transports within a Single System

Many application programming environments offer developers a variety of network protocols and higher-level constructs from which to implement inter-process communications, and it is becoming increasingly common for services to be made available over several different communications frameworks. It is expected that some instantiations of the FIPA architecture will allow the developer or deployer of agent systems to advertise the availability of their services over more than one message transport.

For this reason, the notion of transport address is here generalized to that of *destination*. A destination is an object containing one or more transport addresses. Each address is represented in a format that describes (explicitly or

2255 implicitly) the set of transports for which it is usable. (The precise mapping from address to transport is left to the
 2256 concrete specification, although in practice the mapping is likely to be one-to-one.)
 2257

2258 In its simplest form, a destination may be a single address that unambiguously defines the transport for which it can be
 2259 used.
 2260

2261 **9.4 Desirability of Transport Agnosticism**

2262 The abstract architecture is consistent with concrete architectures which provide "transport agnostic" services. Such
 2263 architectures will provide a programming model in which agents may be more or less aware of the details of transports,
 2264 addressing, and many other communications-related mechanisms. For example, one agent may be able to address
 2265 another in terms of some "social name", or in terms of service attributes advertised through the agent directory service
 2266 without being aware of addressing format, transport mechanism, required level of privacy, audit logging, and so forth.
 2267

2268 Transport agnosticism may apply to both senders and recipients of messages. A concrete architecture may provide
 2269 mechanisms whereby an agent may delegate some or all of the tasks of assigning transport addresses, binding
 2270 addresses to transport end-points, and registering addresses in white-pages or yellow-pages directories to the agent
 2271 platform.
 2272

2273 **9.5 Desirability of Selective Specificity**

2274 While transport agnosticism simplifies the development of agents, there are times when explicit control of specific
 2275 aspects of the message transport mechanism is required. A concrete architecture may provide programmatic access
 2276 to various elements in the message transport subsystem.
 2277

2278 **9.6 Connection-Based, Connectionless and Store-and-Forward Transports**

2279 The abstract architecture is compatible with connection-based, connectionless, and store-and-forward transports. For
 2280 connection-based transports, an instantiation may support the automatic reestablishment of broken connections. It is
 2281 desirable than instantiations that implement several of these modes of operation should support transport-agnostic
 2282 agents.
 2283

2284 **9.7 Conversation Policies and Interaction Protocols**

2285 The abstract architecture specifies a set of abstract objects that allows for the explicit representation of "a
 2286 conversation", i.e. a related set of messages between interlocutors that are logically related by some interaction
 2287 pattern. It is desirable that this property be achieved by the minimum of overhead at the infrastructure or message
 2288 level; in particular, it is important that interoperability remain un-compromised. For example, an implementation may
 2289 deliver messages to conversation-specific queues based on an interpretation of the message envelope. To achieve
 2290 interoperability with an agent that does not support explicit conversations (i.e. which does not allow individual
 2291 messages to be automatically associated with a particular higher-level interaction pattern), it is necessary to specify
 2292 the way in which the message envelope must be processed in order to preserve conversational semantics.
 2293

2294 *Note:* in the practice, we were not able to fully meet this goal. It remains a topic of future work.
 2295

2296 **9.8 Point-to-Point and Multiparty Interactions**

2297 The abstract architecture supports both point-to-point and multiparty message transport. For point-to-point interactions,
 2298 an agent sends a message to an address that identifies a single receiving agent. (An instantiation may support implicit
 2299 addressing, in which the destination is derived from the name of the intended recipient agent without the explicit
 2300 involvement of the sender.) For multiparty message transport, the address must identify a group of recipients. The
 2301 most common model for such message transport is termed "publish and subscribe", in which the address is a "topic" to
 2302 which recipients may subscribe. Other models, for example, "address lists", are possible.

2303
2304
2305
2306
2307
2308

Not all transport mechanisms support multiparty communications, and concrete architectures are not required to provide multiparty messaging services. Concrete architectures that do provide such services may support proxy mechanisms, so that agents and agent systems that only use point-to-point communications may be included in multiparty interactions.

2309 **9.9 Durable Messaging**

2310
2311
2312
2313

Some commercial messaging systems support the notion of durable messages, which are stored by the messaging infrastructure and may be delivered at some later point in time. It is desirable that a message transport architecture should take advantage of such services.

2314 **9.10 Quality of Service**

2315
2316
2317

The term quality of service refers to a collection of service attributes that control the way in which message transport is provided. These attributes fall into a number of categories:

2318
2319
2320
2321
2322
2323
2324
2325
2326
2327
2328
2329
2330
2331

Performance,

Security,

Delivery semantics,

Resource consumption,

Data integrity,

Logging and auditing, and,

Alternate delivery.

2332
2333
2334
2335
2336
2337

Some of these attributes apply to a single message; others may apply to conversations or to particular types of message transport. Architecturally it is important to be able to determine what elements of quality of service are supported, to express (or negotiate) the desired quality of service, to manage the service features which are controlled via the quality of service, to relate the specified quality of service to a service performance guarantee, and to relate quality of service to interoperability specifications.

2338 **9.11 Anonymity**

2339
2340
2341
2342
2343
2344
2345

The abstract transport architecture supports the notion of anonymous interaction. Multiparty message transport may support access by anonymous recipients. An agent may be able to associate a transient address with a conversation, such that the address is not publicly registered with any agent management system or directory service; this may extend to guarantees by the message transport service to withhold certain information about the principal associated with an address. If anonymous interaction is supported, an agent should be able to determine whether or not its interlocutor is anonymous.

2346 **9.12 Message Encoding**

2347
2348
2349
2350

It is anticipated that FIPA will define multiple message encodings together with rules governing the translation of messages from one encoding to another. The message transport architecture allows for the development of instantiations that use one or more message encodings.

2351 9.13 Interoperability and Gateways

2352 The abstract agent transport architecture supports the development of instantiations that use transports, encodings,
2353 and infrastructure elements appropriate to the application domain. To ensure that heterogeneity does not preclude
2354 interoperability, the developers of a concrete architecture must consider the modes of interoperability that are feasible
2355 with other instantiations. Where direct end-to-end interoperability is impossible, impractical or undesirable, it is
2356 important that consideration be given to the specification of gateways that can provide full or limited interoperability.
2357 Such gateways may relay messages between incompatible transports, may translate messages from one encoding to
2358 another, and may provide quality-of-service features supported by one party but not another.
2359

2360 9.14 Reasoning about Agent Communications

2361 The agent transport architecture supports the notion of agents communicating and reasoning about the message
2362 transport process itself. It does not, however, define the ontology or conversation patterns necessary to do this, nor are
2363 concrete architectures required to provide or accept information in a form convenient for such reasoning.
2364

2365 9.15 Testing, Debugging and Management

2366 In general, issues of testing, debugging, and management are implementation-specific and will not be addressed in an
2367 abstract architecture. Individual instantiations may include specific interfaces, actions, and ontologies that relate to
2368 these issues, and may specify that these features are optional or normative for implementations of the instantiation.
2369

2370 **10 Informative Annex C — Goals of Directory Service Abstractions**

2371 This section describes the requirements and architectural elements of the abstract Directory Service. The directory
 2372 service is that part of the FIPA architecture which allows agents to register information about themselves in one or
 2373 more repositories, for those same agents to modify and delete this information, and for agents to search the
 2374 repositories for information of interest to them. The information that is stored is referred to a directory entry, and the
 2375 repository is an agent directory.
 2376

2377 **10.1 Scope**

2378 The purpose of the abstract architecture is to identify the key abstractions that will form the basis of all concrete
 2379 architectures. As such, it is necessarily both limited and non-specific. In this section, we examine some of the ways in
 2380 which concrete directory services may differ.
 2381

2382 **10.2 Variety of Directory Services**

2383 There are several directory services that may be used to store agent descriptions. The abstract architecture is neutral
 2384 with respect to this variety. For any instantiation of the architecture, one must specify the set of directory services that
 2385 are supported, how new directory services are added, and how interoperability is to be achieved. It is permissible for a
 2386 particular concrete architecture to require that implementations of that architecture must support particular directory
 2387 services.
 2388

2389 Different directory services use a variety of different representations for schemas and contents. Instantiations of the
 2390 agent directory architecture may support mechanisms for hiding these differences behind a common API and
 2391 encoding, such as the Java JNDI model or hyper-directory schemes. It is extremely undesirable for an agent to be
 2392 required to parse, decode, or otherwise rely upon different information encodings and schemas.
 2393

2394 The following are examples of directory systems that may be used to instantiate the abstract directory service:

2395 LDAP,

2396 NIS or NIS+,

2400 COS Naming,

2402 Novell NDS,

2404 Microsoft Active Directory,

2406 The Jini lookup service, and,

2408 A name service federation layer, such as JNDI.
 2409

2410 **10.3 Desirability of Directory Agnosticism**

2411 The abstract architecture is consistent with concrete architectures which provide "directory agnostic" services. Such a
 2412 model will support agents that are more or less completely unaware of the details of directory services. A concrete
 2413 architecture may provide mechanisms whereby an agent may delegate some or all of the tasks of assigning transport
 2414 addresses, binding addresses to transport end-points, and registering addresses in all available directories to the
 2415 agent platform.
 2416

2417 10.4 Desirability of Selective Specificity

2418 While directory agnosticism simplifies the development of agents, there are times when explicit control of specific
2419 aspects of the directory mechanism is required. A concrete architecture may provide programmatic access to various
2420 elements in the directory subsystem.
2421

2422 10.5 Interoperability and Gateways

2423 The abstract directory architecture supports the development of instantiations that use directory services appropriate to
2424 the application domain. To ensure that heterogeneity does not preclude interoperability, the developers of a concrete
2425 architecture must consider the modes of interoperability that are feasible with other instantiations. Where direct end-to-
2426 end interoperability is impossible, impractical or undesirable, it is important that consideration be given to the
2427 specification of gateways that can provide full or limited interoperability. Such gateways may extract agent descriptions
2428 from one directory service, transform the information if necessary, and publish it through another directory service.
2429

2430 10.6 Reasoning about Agent Directory

2431 The abstract directory architecture supports the notion of agents communicating and reasoning about the directory
2432 service itself. It does not, however, define the ontology or conversation patterns necessary to do this, nor are concrete
2433 architectures required to provide or accept information in a form convenient for such reasoning.
2434

2435 10.7 Testing, Debugging and Management

2436 In general, issues of testing, debugging, and management are implementation-specific and will not be addressed in an
2437 abstract architecture. Individual instantiations may include specific interfaces, actions, and ontologies that relate to
2438 these issues, and may specify that these features are optional or normative for implementations of the instantiation.
2439
2440

11 Informative Annex D — Goals for Security and Identity Abstractions

11.1 Introduction

In order to create abstractions for the various architectural elements, it is necessary to examine the kinds of systems and infrastructures that are likely targets of real implementations of the abstract architecture. In this section, we examine some of the ways in which security related issues may differ. Authors of concrete architectural specifications must take these issues into account when considering what end-to-end assumptions they can safely make. The goals describe below give the reader an understanding of the objectives the authors of the abstract architecture had in mind when creating this architecture.

In practice, only a very minor part of the security issues can be addressed in the abstract architecture, as most security issues are tightly coupled to their implementation.

In general, the amount of security required is highly dependent on the target deployment environment.

A glossary of security terms is located at the end of this section.

11.2 Overview

There are several aspects to security, which must permeate the FIPA architecture. They are:

Identity. The ability to determine the identity of the various entities in the system. By identifying an entity, another entity interacting with it can determine what policies are relevant to interactions with that entity. Identity is based on credentials, which are verified by a Credential Authority.

Access Permissions. Based on the identity of an entity, determine what policies apply to the entity. These policies might govern resource consumption, types of file access allowed, types of queries that can be performed, or other controlling policies.

Content Validity. The ability to determine whether a piece of software, a message, or other data has been modified since being dispatched by its originating source. Digitally signing data and then having the recipient verify the contents are unchanged often accomplish this. Other mechanisms such as hash algorithms can also be applied.

Content Privacy. The ability to ensure that only designated identities can examine software, a message or other data. To all others the information is obscured. This is often accomplished by encrypting the data, but can also be accomplished by transporting the data over channels that are encrypted.

Identity, or the use of credentials, is needed to supply the ability to control access, to provide content validity, and create content privacy. Each of these is discussed below.

11.3 Areas to Apply Security

This section describes the areas in which security can be applied within agent systems. In each case, the security related risks that are being guarded against are described. The assumption is that any agent or other entity in the system may have credentials that can be used to perform various forms of validation.

11.3.1 Content Validity and Privacy During Message Transport

There are two basic potential security risks when sending a message from one agent to another.

2488 The primary risk is that a message is intercepted, and modified in some way. For example, the interceptor software
 2489 inserts several extra numbers into a payment amount, and modifies the name of the check payee. After modification, it
 2490 is sent on to the original recipient. The other agent acts on the incorrect data. In a case like this, the *content* validity of
 2491 the message is broken.

2492
 2493 The secondary risk is that the message is read by another entity, and the data in it is used by that entity. The message
 2494 does reach its original destination intact. If this occurs, the privacy of the message is violated.

2495
 2496 Digital signing and encryption can address these risks, respectively. These two techniques can be abstractly presented
 2497 at two different layers of the architecture. The messages themselves (or probably just the **payload** part) can be signed
 2498 or encrypted. There are a number of techniques for this, PGP signing and encryption, Public Key signing and
 2499 encryption, one time transmission keys, and other cryptographic techniques. This approach is most effective when the
 2500 nature of underlying message transport is unknown or unreliable from a security perspective.

2501
 2502 The message transport itself can also provide the digital signing or encryption. There are a number of transports that
 2503 can provide such features: SKIP, IPSEC and CORBA Common Secure Interoperability Services. It seems prudent to
 2504 include both models within the architecture, since different applications and software environments will have very
 2505 different capabilities.

2506
 2507 There is another aspect of message transport privacy that comes from agents that misrepresent themselves. In this
 2508 scenario, an agent can register with directory services indicating that is a provider of some service, but in fact uses the
 2509 data it receives for some other purpose. To put it differently, how do you know *who* you are talking to? This topic is
 2510 covered under agent identity below.

2512 **11.3.2 Agent Identity**

2513 If agents and agent services have a digital identity, then agents can validate that:

2514 Agents they are exchanging messages with can be accurately identified, and,

2515 Services they are using are from a known, safe source.

2516
 2517 Similarly, services can determine whether the agent:

2518 Use identity to determine code access or access control decisions, or,

2519 Use agent identity for non-repudiation of transactions.

2520 **11.3.3 Agent Principal Validation**

2521 The Agent can contain a principal (for example a user), on whose behalf this code is running. The principal has one or
 2522 more credentials, and the credentials may have one or more roles that represent the principal.

2523 If an agent has a principal, the other agents can:

2524 Determine whether they want to interoperate with that agent,

2525 Determine what policy and access control to permit to that user, and,

2526 Use the identity to perform transactions.

2527 Services could perform similar actions.

2539
2540
2541
2542
2543
2544
2545
2546
2547
2548
2549
2550
2551
2552

11.3.4 Code Signing Validation

An agent can be code signed. This involves digitally signing the code with one or more credentials. If an agent is code signed, the platform could:

Validate the credential(s) used to sign the agent software. Credentials are validated with a credential authority,

If the credentials are valid, use policy to determine what access this code will have, or,

If the credentials are valid, verify that the code is not modified.

In addition, the Agent Platform can use the lack of digital signature to determine whether to allow the code to run, and policy to determine what access the code will have. In other words, some platforms may have the policy that will not permit code to run, or will restrict Access Permissions unless it is digitally signed.

2553
2554
2555
2556
2557

11.4 Risks Not Addressed

There are a number of other possible security risks that are not addressed, because they are general software issues, rather than unique or special to agents. However, designers of agent systems should keep these issues in mind when designing their agent systems.

2558
2559
2560

11.4.1 Code or Data Peeping

An entity can probe the running agent and extract useful information.

2561
2562
2563
2564

11.4.2 Code or Data Alteration

The unauthorized modification or corruption of an agent, its state, or data. This is somewhat addressed by the code signing, which does not cover all cases.

2565
2566
2567
2568
2569

11.4.3 Concerted Attacks

When a group of agents conspire to reach a set of goals that are not desired by other entities. These are particularly hard to guard against, because several agents may co-operate to create a denial of service attack in a feint to allow another agent to undertake the undesirable action.

2570
2571
2572
2573

11.4.4 Copy and Replay

An attempt to copy an agent or a message and clone or retransmit it. For example, a malicious platform creates an illegal copy, or a clone, of an agent, or a message from an agent is illegally copied and retransmitted.

2574
2575
2576
2577

11.4.5 Denial of Service

In a denial-of-service the attackers try to deny resources to the platform or an agent. For example, an agent floods another agent with requests and the receiving agent is unable to provide its services to other agents.

2578
2579
2580
2581

11.4.6 Misinformation Campaigns

The agent, platform, or service misrepresents information. This includes lying during negotiation, deliberately representing another agent, service or platform as being untrustworthy, costly, or undesirable.

2582 **11.4.7 Repudiation**

2583 An agent or agent platform denies that it has received/sent a message or taken a specific action. For example, a
 2584 commitment between two agents as the result of a contract negotiation is later ignored by one of the agents, denying
 2585 the negotiation has ever taken place and refusing to honour its part of the commitment.
 2586

2587 **11.4.8 Spoofing and Masquerading**

2588 An unauthorized agent or service claims the identity of another agent or piece of software. For example, an agent
 2589 registers as a Directory Service and therefore receives information from other registering agents.
 2590

2591 **11.5 Glossary of Security Terms**

2592 **Access permission** – Based on a credential model, the ability to allow or disallow software from taking an action. For
 2593 example, software with certain credentials may be allowed read a particular file, a group with different credentials may
 2594 be allowed to write to the file.

2595 *Examples: OS file system permissions, Java Security Profiles (check name), Database access controls.*
 2596

2597 **Authentication** – Using some credential model, ability to verify that the entity offering the credentials is who/what it
 2598 says it is.
 2599

2600 **Credential** – An item offered to prove that a user, a group, a software entity, a company, or other entities is who or
 2601 what it claims to be.

2602 *Examples: X.509 certificate, a user login and password pair, a PGP key, a response/challenge key, a fingerprint, a
 2603 retinal scan, a photo id. (Obviously, some of these are better suited to software than others!)*
 2604

2605 **Credential Authority** – An entity that determines whether the credential offered is valid, and that the credential
 2606 accurately identifies the individual offering it.

2607 *Examples: An X.509 certificate can be validated by a certificate authority. At a bar, the bartender is the credential
 2608 authority who determines whether your photo id represents you (he may then determine your access permissions to
 2609 available beverages!).*
 2610

2611 **Credential model** – The particular mechanism(s) being used to provide and authenticate credentials.
 2612

2613 **Code signing** – A particular case of digital signature (see below), where code is signed by the credentials of some
 2614 entity. The purpose of code signing is to identify the source of the code, and to verify that the code has not been
 2615 changed by another entity.

2616 *Examples: Java code signing, DCOM object signing, checksum verification.*
 2617

2618 **Digital signature** – Using a credential model to indicate the source of some data, and to ensure that the data is
 2619 unchanged since it was signed. Note: the word data is used very broadly here – it could a string, software, voice
 2620 stream, etc.

2621 *Examples: S/MIME mail, PGP digital signing, IPSEC (authentication modes)*
 2622

2623 **Encryption** – The ability to transform data into a format that can only be restored by the holder of a particular
 2624 credential. Used to prevent data from being observed by others.

2625 *Examples: SSL, S/MIME mail, PGP digital signing, IPSEC (encryption modes)*
 2626

2627 **Identity** – A person, server, group, company, software program that can be uniquely identified. Identities can have
 2628 credentials that identify them.
 2629

2630 **Lease** – An interval of time that some element, such as an identity or a credential is good for. Leases are very useful
 2631 when you want to restrict the length of commitment. For example, you may issue a temporary credential to an agent
 2632 that gives it 20 minutes in a given system, at which time the credential expires.
 2633

2634 **Policy** – Some set of actions that should be performed when a set of conditions is met. In the context of security, allow
2635 access permissions based on a valid credential that establishes an identity.

2636 *Examples: If a credential for a particular user is presented, allow him to access a file. If a credential for a particular role*
2637 *is presented, allow the agent to run with a low priority.*

2638

2639 **Role** – An identity that has an "group" quality. That is, the role does not uniquely identify an individual, or machine, or
2640 an agent, but instead identifies the identity in a particular context: as a system manager, as a member of the entry
2641 order group, as a high-performance calculation server, etc.

2642 *Examples: In various operating system groups, as applied to file system access. In Lotus Notes, the "role" concept.*
2643 *X.509 certificate role attributes.*

2644

2645 **Principal** – In the agent domain, the identity on whose behalf the agent is running. This may be a user, a group, a role
2646 or another software entity.

2647 *Examples: A shopping agent's principal is the user who launched it. An commodity trader agent's principal is a*
2648 *financial company. A network management agent's principal is the role of system admin, or super-user. In a small*
2649 *"worker bee" agent, the principal may be the delegated authority of the parent agent.*

2650

2651 12 Informative Annex E — Change-Log

2652 12.1 2001/11/01 - change delta with respect to XC00001J

2653	All document	directory-service becomes agent-directory-service .
2654	All document	directory-entry becomes agent-directory-entry .
2655	All document	locator becomes agent-locator .
2656	All document	Encoding-transform-service becomes encoding-service .
2657		
2658	Section 1, Paragraph 5	Note added concerning availability of documents.
2659	Section 1.1	Annexes updated to include new ones.
2660	Section 2.1	Changed text of second bullet point.
2661	Section 2.1	Section descriptions updated to include new annexes.
2662	Section 2.3, Paragraph 2	Added complete paragraph.
2663	Section 4.1, Paragraph 1	Changed 2nd sentence changed to include service-directory-service .
2664	Section 4.1, Paragraph 2	First sentence added.
2665	Section 4.2	Added complete section.
2666	Section 4.3	Table updated to correct agent-locator description.
2667	Section 4.3.1	Changed section heading.
2668	Section 4.3.2	Changed section heading.
2669	Section 4.4	Added complete section.
2670	Section 4.5, Paragraph 1	Changed "fundamental aspects" to include message representation.
2671	Section 4.5.1, Paragraph 1	Replaced 3rd sentence.
2672	Section 4.5.1, Figure 6	Receiver (and agent-name for receiver) made plural.
2673	Section 4.5.2	Added complete section.
2674	Section 4.5.3, Figure 7	Receiver (and agent-name for receiver) made plural.
2675	Section 5.1.5, Table 2	Included Fully Qualified Name column for each element
2676		Changed description of encoding-service .
2677		Changed service presence to be mandatory.
2678		Added service-address .
2679		Added service-attributes .
2680		Added service-directory-service .
2681		Added service-directory-entry .
2682		Added service-id .
2683		Added service-location-description .
2684		Added service-locator .
2685		Added service-root .
2686		Added service-signature .
2687		Added service-type .
2688		Added signature-type .
2689		Added transport-specific-address .
2690	Section 5.2	Added complete section.
2691	Section 5.3	Added complete section.
2692	Section 5.4.2	Removed first point.
2693	Section 5.6.1, Paragraph 1	Removed 2nd and 3rd sentence. Added new 2nd sentence.
2694	Section 5.6.1, Paragraph 2	Removed.
2695	Section 5.6.2	Added new relationship.
2696	Section 5.10.3	Changed 1st sentence so that GUID now an example.
2697	Section 5.11.1	Changed 1st sentence to include message reference.
2698		Moved 2nd and 3rd sentences to Section 5.11.3
2699		Added new 2nd sentence.
2700	Section 5.11.2	Changed 2nd relationship to be more accurate.
2701	Section 5.11.3	Added complete section.
2702	Section 5.13.1, Paragraph 1	Changed 2nd sentence to include Bit-efficient encoding.

2703		Added 3rd sentence.
2704	Section 5.13.1, Paragraph 2	Removed.
2705	Section 5.13.2	Changed 1st relationship.
2706		Removed 2nd, 3rd and 4th relationships.
2707		Added new 2nd relationship.
2708	Section 5.14.1	Added 3rd sentence.
2709	Section 5.14.2	Changed 2nd, 3rd and 4th relationship.
2710		Removed 5th relationship.
2711	Section 5.14.3.1	Changed section heading.
2712	Section 5.14.3.1. Paragraph 1	Changed 1st and 2nd sentences.
2713	Section 5.14.3.1. Paragraph 2	Changed 1st sentence.
2714	Section 5.14.3.1. Paragraph 3	Added complete paragraph.
2715	Section 5.14.3.1	Added 'invalid payload' explanation.
2716	Section 5.14.3	Added new 2nd sentence.
2717	Section 5.14.3	Deleted last 2 sentences.
2718	Section 5.16.1	Added last sentence.
2719	Section 5.16.3	Changed 1st to include service-directory-service .
2720	Section 5.17.1	Added new 4th and last sentences.
2721	Section 5.17.1	Added 'and ontologies' to 6th sentence.
2722	Section 5.17.3	Updated final two relationships.
2723	Section 5.19.2	Updated both relationships with respect to ontologies .
2724	Section 5.21.2	Added three new relationships related to service model.
2725	Section 5.22	Added complete section.
2726	Section 5.23	Added complete section.
2727	Section 5.24	Added complete section.
2728	Section 5.25	Added complete section.
2729	Section 5.26	Added complete section.
2730	Section 5.27	Added complete section.
2731	Section 5.28	Added complete section.
2732	Section 5.29	Added complete section.
2733	Section 5.30	Added complete section.
2734	Section 5.31	Added complete section.
2735	Section 5.32	Added complete section.
2736	Section 5.36	Added complete section.
2737	Section 6.2, Figure 12	Changed message-encoding-representation to encoding-representation .
2738		Changed transform-service to encoding-service .
2739		Changed role linking payload and message .
2740		Removed role linking transport-message and encoding-representation .
2741		Removed role linking transport-message and encoding-service .
2742		Removed payload-external-attributes .
2743		Added role linking envelope and encoding-representation .
2744	Section 6.3, Figure 13	Changed role linking agent-directory-service and agent-locator from 'contains 1..n' to 'contain 1'.
2745		Changed role linking agent-locator and transport-description from 'contains 1' to 'contain 1..n'.
2746		Changed role linking transport-description and transport-type from "has a" to "contains 1".
2747		
2748		
2749		
2750	Section 6.4	Added complete section.
2751	Section 6.5, Paragraph 1	Added final two sentences.
2752	Section 6.5, Figure 15	Changed role linking message and "communicative act" from 'contains 1..n' to 'is a'.
2753		Changed role linking "communicative act" and content from 'contains 1..n' to 'contains 1'.
2754		
2755	Section 7	Added reference for FIPA00095.
2756	Section 8	Added complete section.
2757	Section 9	Added complete section.
2758	Section 10	Added word 'service' into section heading.

2759 Section 13

Added complete section.