



Reference Architecture Foundation for Service Oriented Architecture Version 1.0

Draft 03

XX XXX 2010

Specification URIs:

This Version:

<http://docs.oasis-open.org/soa-rm/soa-ra/v1.0/soa-ra-cd-02.html>
<http://docs.oasis-open.org/soa-rm/soa-ra/v1.0/soa-ra-cd-02.doc>
<http://docs.oasis-open.org/soa-rm/soa-ra/v1.0/soa-ra-cd-02.pdf> (Authoritative)

Previous Version:

<http://docs.oasis-open.org/soa-rm/soa-ra/v1.0/soa-ra-pr-01.html>
<http://docs.oasis-open.org/soa-rm/soa-ra/v1.0/soa-ra-pr-01.doc>
<http://docs.oasis-open.org/soa-rm/soa-ra/v1.0/soa-ra-pr-01.pdf>

Latest Version:

<http://docs.oasis-open.org/soa-rm/soa-ra/v1.0/soa-ra.html>
<http://docs.oasis-open.org/soa-rm/soa-ra/v1.0/soa-ra.doc>
<http://docs.oasis-open.org/soa-rm/soa-ra/v1.0/soa-ra.pdf> (Authoritative)

Technical Committee:

OASIS Service Oriented Architecture Reference Model TC

Chair(s):

Ken Laskey, MITRE Corporation

Editor(s):

Jeff A. Estefan, Jet Propulsion Laboratory, jeffrey.a.estefan@jpl.nasa.gov
Ken Laskey, MITRE Corporation, klaskey@mitre.org
Francis G. McCabe, Individual, fmccabe@gmail.com
Danny Thornton, Northrop Grumman, danny.thornton@ngc.com

Related work:

This specification is related to:

[OASIS Reference Model for Service Oriented Architecture](#)

Abstract:

This document specifies the OASIS Reference Architecture Foundation for Service Oriented Architecture. It follows from the concepts and relationships defined in the OASIS Reference Model for Service Oriented Architecture. While

it remains abstract in nature, the current document describes the foundation upon which a SOA concrete architecture can be built.

Our focus in this architecture is on an approach to integrating business with the information technology needed to support it. The issues involved with integration are always present, but, we find, are thrown into clear focus when business integration involves crossing ownership boundaries.

This architecture follows the recommended practice of describing architecture in terms of models, views, and viewpoints, as prescribed in ANSI¹/IEEE² 1471-2000 and ISO³/IEC⁴ 42010-2007 Standards. This Reference Architecture is intended to be of value to Enterprise Architects, Business and IT Architects as well as CIOs and other senior executives involved in strategic business and IT planning.

The Reference Architecture has three main views: the Service Ecosystem view which focuses on the way that participants are part of a Service Oriented Architecture ecosystem; the Realizing Services view which addresses the requirements for constructing a Service Oriented Architecture; and the Owning Service Oriented Architecture view which focuses on the governance and management of SOA-based systems.

Status:

This document was last revised or approved by the SOA Reference Model TC on the above date. The level of approval is also listed above. Check the “Latest Version” or “Latest Approved Version” location noted above for possible later revisions of this document.

Technical Committee members should send comments on this specification to the Technical Committee’s email list. Others should send comments to the Technical Committee by using the “Send A Comment” button on the Technical Committee’s web page at http://www.oasis-open.org/committees/tc_home.php?wg_abbrev=soa-rm.

For information on whether any patents have been disclosed that may be essential to implementing this specification, and any offers of patent licensing terms, please refer to the Intellectual Property Rights section of the Technical Committee web page <http://www.oasis-open.org/committees/soa-rm/ipr.php>.

The non-normative errata page for this specification is located at http://www.oasis-open.org/committees/tc_home.php?wg_abbrev=soa-rm.

¹ American National Standards Institute

² Institute of Electrical and Electronics Engineers

³ International Organization for Standardization

⁴ International Electrotechnical Commission

Notices

Copyright © OASIS® 1993–2009. All Rights Reserved.

All capitalized terms in the following text have the meanings assigned to them in the OASIS Intellectual Property Rights Policy (the "OASIS IPR Policy"). The full Policy may be found at the OASIS website.

This document and translations of it may be copied and furnished to others, and derivative works that comment on or otherwise explain it or assist in its implementation may be prepared, copied, published, and distributed, in whole or in part, without restriction of any kind, provided that the above copyright notice and this section are included on all such copies and derivative works. However, this document itself may not be modified in any way, including by removing the copyright notice or references to OASIS, except as needed for the purpose of developing any document or deliverable produced by an OASIS Technical Committee (in which case the rules applicable to copyrights, as set forth in the OASIS IPR Policy, must be followed) or as required to translate it into languages other than English.

The limited permissions granted above are perpetual and will not be revoked by OASIS or its successors or assigns.

This document and the information contained herein is provided on an "AS IS" basis and OASIS DISCLAIMS ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY OWNERSHIP RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

OASIS requests that any OASIS Party or any other party that believes it has patent claims that would necessarily be infringed by implementations of this OASIS Committee Specification or OASIS Standard, to notify OASIS TC Administrator and provide an indication of its willingness to grant patent licenses to such patent claims in a manner consistent with the IPR Mode of the OASIS Technical Committee that produced this specification.

OASIS invites any party to contact the OASIS TC Administrator if it is aware of a claim of ownership of any patent claims that would necessarily be infringed by implementations of this specification by a patent holder that is not willing to provide a license to such patent claims in a manner consistent with the IPR Mode of the OASIS Technical Committee that produced this specification. OASIS may include such claims on its website, but disclaims any obligation to do so.

OASIS takes no position regarding the validity or scope of any intellectual property or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; neither does it represent that it has made any effort to identify any such rights. Information on OASIS' procedures with respect to rights in any document or deliverable produced by an OASIS Technical Committee can be found on the OASIS website. Copies of claims of rights made available for publication and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary

rights by implementers or users of this OASIS Committee Specification or OASIS Standard, can be obtained from the OASIS TC Administrator. OASIS makes no representation that any information or list of intellectual property rights will at any time be complete, or that any claims in such list are, in fact, Essential Claims.

The name "OASIS" is a trademark of OASIS, the owner and developer of this specification, and should be used only to refer to the organization and its official outputs. OASIS welcomes reference to, and implementation and use of, specifications, while reserving the right to enforce its marks against misleading uses. Please see <http://www.oasis-open.org/who/trademark.php> for above guidance.

Unified Modeling Language™, UML®, Object Management Group™, and OMG™ are trademarks of the Object Management Group.

10/26/2010

Table of Contents

1	Introduction	11
1.1	Context for Reference Architecture for SOA.....	11
1.1.1	What is a Reference Architecture?.....	11
1.1.2	What is this Reference Architecture?	12
1.1.3	Relationship to the OASIS Reference Model for SOA	12
1.1.4	Relationship to other Reference Architectures.....	13
1.1.5	Relationship to other SOA Open Standards.....	13
1.1.6	Expectations set by this Reference Architecture Foundation.....	14
1.2	Service Oriented Architecture – An Ecosystems Perspective	15
1.3	Viewpoints, Views and Models	15
1.3.1	ANSI/IEEE 1471-2000::ISO/IEC 42010-2007	15
1.3.2	UML Modeling Notation.....	17
1.4	Viewpoints of this Reference Architecture	18
1.4.1	Service Ecosystem Viewpoint	19
1.4.2	Realizing Service Oriented Architectures Viewpoint	20
1.4.3	Owning Service Oriented Architectures Viewpoint.....	20
1.5	Terminology	20
1.5.1	Usage of Terms.....	20
1.6	References.....	20
1.6.1	Normative References.....	20
1.6.2	Non-Normative References.....	21
2	Architectural Goals and Principles	23
2.1	Goals and Critical Success Factors of this Reference Architecture.....	23
2.1.1	Goals	24
2.1.2	Critical Success Factors.....	25
2.2	Principles of this Reference Architecture	26
3	Service Ecosystem View.....	29
3.1	Social Structure Model.....	30
3.1.1	Actors, Delegates and Participants	32
3.1.2	Roles in Social Structures	34
3.1.3	Shared State and Social Facts.....	36
3.1.4	Policies and Contracts.....	37
3.1.5	Resources and Ownership	39
3.1.6	Life-cycle of Social Structures	42
3.2	Acting in a SOA Ecosystem Model	42
3.2.1	Action and Joint Action.....	42
3.2.2	Goals and Capabilities	46

3.2.3 Trust and Risk	47
3.2.4 Social Actions	49
3.2.5 Communication for Action	51
3.3 Semantics in a SOA Ecosystem Model	54
3.3.1 Assertions	54
3.3.2 Semantic Engagement	58
3.4 Architectural Implications	60
4 Realizing Service Oriented Architectures View	62
4.1 Service Description Model	62
4.1.1 The Model for Service Description	64
4.1.2 Use Of Service Description	74
4.1.3 Relationship to Other Description Models	80
4.1.4 Architectural Implications	81
4.2 Service Visibility Model	83
4.2.1 Visibility to Business	83
4.2.2 Visibility	85
4.2.3 Architectural Implications	90
4.3 Interacting with Services Model	91
4.3.1 Interaction Dependencies	91
4.3.2 Actions and Events	92
4.3.3 Message Exchange	93
4.3.4 Composition of Services	96
4.3.5 Architectural Implications of Interacting with Services	100
4.4 Policies and Contracts Model	102
4.4.1 Policy and Contract Representation	102
4.4.2 Policy and Contract Enforcement	104
4.4.3 Architectural Implications	106
5 Owning Service Oriented Architectures View	107
5.1 Governance Model	107
5.1.1 Understanding Governance	108
5.1.2 A Generic Model for Governance	109
5.1.3 Governance Applied to SOA	114
5.1.4 Architectural Implications of SOA Governance	120
5.2 Security Model	121
5.2.1 Secure Interaction Concepts	122
5.2.2 Where SOA Security is Different	125
5.2.3 Security Threats	125
5.2.4 Security Responses	126
5.2.5 Architectural Implications of SOA Security	129

5.3 Management Model	129
5.3.1 Management and Governance	133
5.3.2 Management Contracts and Policies	133
5.3.3 Management Infrastructure	133
5.3.4 Service Life-cycle	134
5.4 SOA Testing Model	134
5.4.1 Traditional Software Testing as Basis for SOA Testing	135
5.4.2 Testing and the SOA Ecosystem	136
5.4.3 Elements of SOA Testing	138
5.4.4 Testing SOA Services	144
5.4.5 Architectural Implications for SOA Testing	148
6 Conformance	149
A. Acknowledgements	150
B. Critical Factors Analysis	151
B.1 Goals	151
Critical Success Factors	151
Requirements	151
CFA Diagrams	151

Table of Figures

Figure 1 Example UML class diagram—Resources.....	17
Figure 2 Critical Factors Analysis of the Reference Architecture	24
Figure 3 Model elements described in the Service Ecosystem view.....	30
Figure 4 Social Structure	30
Figure 5 Enterprise as a Social Structure.....	Error! Bookmark not defined.
Figure 6 Actors, Participants and Delegates	32
Figure 7 Service Participants.....	33
Figure 8 Roles, Rights and Responsibilities	34
Figure 9 Shared State and Social Facts.....	36
Figure 10 Policies and Contracts.....	38
Figure 11 Propositions.....	56
Figure 12 Assertions and Promises.....	58
Figure 13 Resources	40
Figure 14 Resource Ownership.....	41
Figure 15 Actions, Real World Effect and Events.....	42
Figure 16 Joint Action.....	44
Figure 17 Needs and Capabilities	46
<i>Figure 18 Trusting Actor and Willingness.....</i>	47
<i>Figure 19 Assessing Trust and Risk.....</i>	49
Figure 20 Acting within Social Structures	50
Figure 21 Business Transaction	50
Figure 22 Communication as Joint Action	52
Figure 23 Communicative actions as Service Actions.....	53
Figure 24 Communicating Agreement.....	Error! Bookmark not defined.
Figure 25 Model Elements Described in the Realizing Service Oriented Architectures View.....	62
<i>Figure 26 General Description.....</i>	65
Figure 27 Representation of a Description	67
<i>Figure 28 Service Description</i>	69
Figure 29 Service Interface	70
<i>Figure 30 Service Functionality</i>	71
Figure 31 Model for Policies and Contracts as related to Service Participants.....	72
Figure 32 Action-Level and Service-Level Policies	73
Figure 33 Policies and Contracts, Metrics, and Compliance Records.....	74
Figure 34 Relationship Between Action and Service Description Components	75

Figure 35 Execution Context	79
Figure 36 Interaction Description.....	80
Figure 37 Visibility to Business	84
Figure 38 Mediated Service Awareness.....	86
Figure 39 Awareness In a SOA Ecosystem	88
Figure 40 Business, Description and Willingness.....	89
Figure 41 Service Reachability	90
Figure 42 Interaction dependencies.	92
Figure 43 A "message" conveys either an action or an event.....	92
Figure 44 Fundamental SOA message exchange patterns (MEPs).....	94
Figure 45 Simple model of service composition.....	96
Figure 46 Abstract example of orchestration of service-oriented business process.	98
Figure 47 Abstract example of choreography of service-oriented business collaboration.	100
Figure 48 Policies and Contracts.....	103
Figure 49 Model Elements Described in the Owning Service Oriented Architectures View	107
Figure 50 Motivating governance model	110
Figure 51 Setting up governance model.....	111
Figure 52 Carrying out governance model	112
Figure 53 Ensuring governance compliance model	113
Figure 54 Relationship among types of governance	116
Figure 55 Confidentiality and Integrity	123
Figure 56 Authentication.....	123
Figure 57 Authorization	124
Figure 58 Managing resources in a SOA	130

1 Introduction

Service Oriented Architecture (SOA) is an architectural paradigm that has gained significant attention within the information technology (IT) and business communities. The SOA ecosystem described in this document occupies the boundary between business and IT. It is neither wholly IT nor wholly business, but is of both worlds. Neither business nor IT completely own, govern and manage this SOA ecosystem. Both sets of concerns must be accommodated for the SOA ecosystem to fulfill its purposes.⁵

The OASIS Reference Model for SOA [SOA-RM] provides a common language for understanding the important features of SOA but does not address the issues involved in constructing, using or owning a SOA-based system. This document focuses on these aspects of SOA.

The intended audiences of this document and expected benefits to be realized include non-exhaustively:

- Enterprise Architects - will gain a better understanding when planning and designing enterprise systems of the principles that underlie Service Oriented Architecture;
- Standards Architects and Analysts - will be able to better position specific specifications in relation to each other in order to support the goals of SOA;
- Decision Makers - will be better informed as to the technology and resource implications of commissioning and living with a SOA-based system; in particular, the implications following from multiple ownership domains; and
- Users - will gain a better understanding of what is involved in participating in a SOA-based system.

1.1 Context for Reference Architecture for SOA

1.1.1 What is a Reference Architecture?

A reference architecture models the abstract architectural elements in the domain independent of the technologies, protocols, and products that are used to implement the domain. It differs from a reference model in that a reference model describes the important concepts and relationships in the domain focusing on what distinguishes the elements of the domain; a reference architecture elaborates further on the model to show a more complete picture that includes showing what is involved in realizing the modeled entities.

It is possible to define reference architectures at many levels of detail or abstraction, and for many different purposes. A reference architecture need not be a concrete architecture; i.e., depending on the requirements being addressed by the reference

⁵ By *business* we refer to any activity that people are engaged in. We do not restrict the scope of SOA ecosystems to commercial applications.

36 architecture, it may not be necessary to completely specify all the technologies,
37 components and their relationships in sufficient detail to enable direct implementation.

38 **1.1.2 What is this Reference Architecture?**

39 This Reference Architecture is an abstract realization of SOA, focusing on the elements
40 and their relationships needed to enable SOA-based systems to be used, realized and
41 owned while avoiding reliance on specific concrete technologies. It is identified as a
42 *Reference Architecture Foundation* because it takes a first principles approach to
43 architectural modeling of SOA-based systems.

44 While requirements are addressed more fully in Section 2, the key assumptions that we
45 make in this Reference Architecture is that SOA-based systems involve:

- 46 • resources that are distributed across ownership boundaries;
- 47 • people and systems interacting with each other, also across ownership
48 boundaries;
- 49 • security, management and governance that are similarly distributed across
50 ownership boundaries; and
- 51 • interaction between people and systems that is primarily through the exchange of
52 messages with reliability that is appropriate for the intended uses and purposes.

53 Even in contexts that apparently have no ownership boundaries, such as within a single
54 organization, the reality is that different groups and departments often behave as
55 though they had ownership boundaries between them. This reflects organizational
56 practice; as well as reflecting the real motivations and desires of the people running
57 those organizations.

58 Below, we talk about such an environment as a *service ecosystem*. Informally, our goal
59 in this Reference Architecture is to show how Service Oriented Architecture fits into the
60 life of users and stakeholders, how such systems may be realized effectively, and what
61 is involved in owning and managing them. We believe that this approach will serve two
62 purposes: to ensure that service ecosystems can be realized using appropriate
63 technology, and to permit the audience to focus on the important issues without
64 becoming over-burdened with the details of a particular implementation technology.

65 **1.1.3 Relationship to the OASIS Reference Model for SOA**

66 The primary contribution of the OASIS Reference Model for Service Oriented
67 Architecture is that it identifies the key characteristics of SOA, and it defines many of the
68 important concepts needed to understand what SOA is and what makes it important.
69 This Reference Architecture Foundation takes the Reference Model as its starting point
70 in particular in relation to the vocabulary of important terms and concepts.

71 The Reference Architecture described herein goes a step further than the Reference
72 Model in that it shows how SOA-based systems can be realized – albeit in an abstract
73 way. As noted above, SOA-based systems are better thought of as ecosystems rather
74 than stand-alone software products. Consequently, how they are used and managed is
75 at least as important architecturally as how they are constructed.

76 In terms of approach, the primary difference between the Reference Model and this
77 document referred to as the Reference Architecture Foundation is that the former

78 focuses entirely on a common language of the distinguishing features of SOA. This
79 document introduces concepts and architectural elements as needed in order to fulfill
80 the core requirement of using, realizing and owning SOA-based systems.

81 **1.1.4 Relationship to other Reference Architectures**

82 It is fully recognized that other SOA reference architectures have emerged in the
83 industry, both from the analyst community and the vendor/solution provider community.
84 Some of these reference architectures are quite abstract in relation to specific
85 implementation technologies, while others are based on a solution or technology stack.
86 Still others use middleware technology such as an Enterprise Service Bus (ESB) as the
87 architectural foundation.

88 As with the Reference Model, this Reference Architecture is primarily focused on large-
89 scale distributed IT systems where the participants may be legally separate entities. It is
90 quite possible for many aspects of this Reference Architecture to be realized on quite
91 different platforms.

92 In addition, this Reference reference Architecture achitecture, as the title illustrates, is
93 intended to provide foundational concepts on which to build other reference
94 architectures and eventual concrete architectures. The relationship to other industry
95 reference architectures for SOA and related SOA open standards is described below in
96 Section 1.1.5

97 **1.1.5 Relationship to other SOA Open Standards**

98 The “Navigating the SOA Open Standards Landscape Around Architecture” joint white
99 paper from OASIS, OMG, and The Open Group [**SOA-NAV**] was written to help the
100 SOA community at large navigate the myriad of overlapping technical products
101 produced by these organizations with specific emphasis on the “A” in SOA, i.e.,
102 Architecture.

103 This joint white paper explains and positions standards for SOA reference models,
104 ontologies, reference architectures, maturity models, modeling languages, and
105 standards work on SOA governance. It outlines where the works are similar, highlights
106 the strengths of each body of work, and touches on how the work can be used together
107 in complementary ways. It is also meant as a guide to users for selecting those
108 specifications most appropriate for their needs.

109 While the understanding of SOA and SOA Governance concepts provided by these
110 works is similar, the evolving standards are written from different perspectives. Each
111 specification supports a similar range of opportunity, but has provided different depths
112 of detail for the perspectives on which they focus. Therefore, although the definitions
113 and expressions may differ somewhat, there is agreement on the fundamental concepts
114 of SOA and SOA Governance.

115 The following is a summary taken from [**SOA-NAV**] of the positioning and guidance on
116 the specifications:

- 117 • The OASIS Reference Model for SOA (SOA RM) is the most abstract of the
118 specifications positioned. It is used for understanding of core SOA concepts

- 119 • The Open Group SOA Ontology extends, refines, and formalizes some of the
120 core concepts of the SOA RM. It is used for understanding of core SOA
121 concepts and facilitate a model-driven approach to SOA development.
- 122 • The OASIS Reference Architecture Foundation for SOA (this document) is an
123 abstract, foundational reference architecture addressing the ecosystem viewpoint
124 for building and interacting within the SOA paradigm. It is used for understanding
125 different elements of SOA, the completeness of SOA architectures and
126 implementations, and considerations for reaching across ownership boundaries
127 where there is no single authoritative entity for SOA and SOA governance.
- 128 • The Open Group SOA Reference Architecture is a layered architecture from
129 consumer and provider perspective with cross cutting concerns describing these
130 architectural building blocks and principles that support the realizations of SOA. It
131 is used for understanding the different elements of SOA, deployment of SOA in
132 enterprise, basis for an industry or organizational reference architecture,
133 implication of architectural decisions, and positioning of vendor products in a
134 SOA context.
- 135 • The Open Group SOA Governance Framework is a governance domain
136 reference model and method. It is for understanding SOA governance in
137 organizations. The OASIS Reference Architecture for SOA Foundation contains
138 an abstract discussion of governance principles as applied to SOA across
139 boundaries
- 140 • The Open Group SOA Integration Maturity Model (OSIMM) is a means to assess
141 an organization's maturity within a broad SOA spectrum and define a roadmap
142 for incremental adoption. It is used for understanding the level of SOA maturity in
143 an organization
- 144 • The Object Management Group SoaML Specification supports services modeling
145 UML extensions. It can be seen as an instantiation of a subset of the Open
146 Group RA used for representing SOA artifacts in UML.

147 Fortunately, there is a great deal of agreement on the foundational core concepts
148 across the many independent specifications and standards for SOA. This could be best
149 explained by broad and common experience of users of SOA and its maturity in the
150 marketplace. It also provides assurance that investing in SOA-based business and IT
151 transformation initiatives that incorporate and use these specifications and standards
152 helps to mitigate risks that might compromise a successful SOA solution.

153 **1.1.6 Expectations set by this Reference Architecture Foundation**

154 This Reference Architecture Foundation is not a complete blueprint for realizing SOA-
155 based systems. Nor is it a technology map identifying all the technologies needed to
156 realize SOA-based systems. It does identify many of the key aspects and components
157 that will be present in any well designed SOA-based system. In order to actually use,
158 construct and manage SOA-based systems, many additional design decisions and
159 technology choices will need to be made.

160 1.2 Service Oriented Architecture – An Ecosystems Perspective

161 Many systems cannot be understood by a simple decomposition into parts and
162 subsystems – in particular when there are many interactions between the parts. For
163 example, a biological ecosystem is a self-sustaining association of plants, animals, and
164 the physical environment in which they live. Understanding an ecosystem often
165 requires a holistic perspective rather than one focusing on the system's individual parts.

166 This Reference Architecture views the SOA architectural paradigm from an ecosystems
167 perspective: a space that people, machines and services inhabit in order to further both
168 their own objectives and the objectives of the larger community.

169 Viewed as whole, a SOA-based system is a network of independent services,
170 machines, and people who operate, affect, use, and govern those services as well as
171 the suppliers of equipment and personnel to these people and services.

172 In a SOA ecosystem there may not be any single person or organization that is really "in
173 control" or "in charge" of the whole ecosystem; although there are definite stakeholders
174 involved, each of whom has some control and influence over the community.

175 The three key principles that inform our approach to a SOA ecosystem are:

- 176 • a SOA is a *medium for exchange of value* between independently acting
177 *participants*;
- 178 • participants (and **stakeholders** in general) have legitimate claims to *ownership* of
179 resources that are made available via the SOA; and
- 180 • the behavior and performance of the participants are subject to *rules of engagement*
181 which are captured in a series of policies and contracts.

182 1.3 Viewpoints, Views and Models

183 1.3.1 ANSI/IEEE 1471-2000::ISO/IEC 42010-2007

184 This Reference Architecture follows the ANSI⁶/IEEE⁷ 1471-2000 and ISO⁸/IEC⁹ 42010-
185 2007 standard. Recommended Practice for Architectural Description of Software-
186 Intensive Systems [**ANSI/IEEE 1471, ISO/IEC 42010**]. An architectural description
187 conforming to the ANSI/IEEE 1471-2000::ISO/IEC 42010-2007 recommended practice
188 is described by a clause that includes the following six (6) elements:

- 189 1. Architectural description identification, version, and overview information
- 190 2. Identification of the system **stakeholders** and their concerns judged to be
191 relevant to the architecture
- 192 3. Specifications of each viewpoint that has been selected to organize the
193 representation of the architecture and the rationale for those selections

⁶ American National Standards Institute

⁷ Institute of Electrical and Electronics Engineers

⁸ International Organization for Standardization

⁹ International Electrotechnical Commission

- 194 4. One or more architectural views
195 5. A record of all known inconsistencies among the architectural description's
196 required constituents
197 6. A rationale for selection of the architecture (in particular, showing how the
198 architecture supports the identified stakeholders' concerns).

199 The ANSI/IEEE 1471-2000::ISO/IEC 42010-2007 defines the following terms:

200 **Architecture**

201 The fundamental organization of a system embodied in its components, their
202 relationships to each other, and to the environment, and the principles guiding its
203 design and evolution.

204 **Architectural Description**

205 A collection of products that document the architecture.

206 **System**

207 A collection of components organized to accomplish a specific function or set of
208 functions.

209 **System Stakeholder**

210 A system stakeholder is an individual, team, or organization (or classes thereof)
211 with interests in, or concerns relative to, a system.

212 A stakeholder's concern should not be confused with a formal requirement. A concern is
213 an area or topic of interest. Within that concern, system stakeholders may have many
214 different requirements. In other words, something that is of interest or importance is not
215 the same as something that is obligatory or of necessity [TOGAF v9].

216 When describing architectures, it is important to identify stakeholder concerns and
217 associate them with viewpoints to insure that those concerns will be addressed in some
218 manner by the models that comprise the views on the architecture. The ANSI/IEEE
219 1471-2000::ISO/IEC 42010-2007 defines views and viewpoints as follows:

220 **View**

221 A representation of the whole system from the perspective of a related set of
222 concerns.

223 **Viewpoint**

224 A specification of the conventions for constructing and using a view. A pattern or
225 template from which to develop individual views by establishing the purposes and
226 audience for a view and the techniques for its creation and analysis.

227 In other words, a view is what the stakeholders see whereas the viewpoint defines the
228 perspective from which the view is taken.

229 It is important to note that viewpoints are independent of a particular system. In this
230 way, the architect can select a set of candidate viewpoints first, or create a set of
231 candidate viewpoints, and then use those viewpoints to construct specific views that will
232 be used to organize the architectural description. A view, on the other hand, is specific
233 to a particular system. Therefore, the practice of creating an architectural description
234 involves first selecting the viewpoints and then using those viewpoints to construct

235 specific views for a particular system or subsystem. Note that ANSI/IEEE 1471-
236 2000::ISO/IEC 42010-2007 requires that each view corresponds to exactly one
237 viewpoint. This helps maintain consistency among architectural views which is a
238 normative requirement of the standard.

239 A view is comprised of one or more architectural models, where model is defined as:

240 **Model**

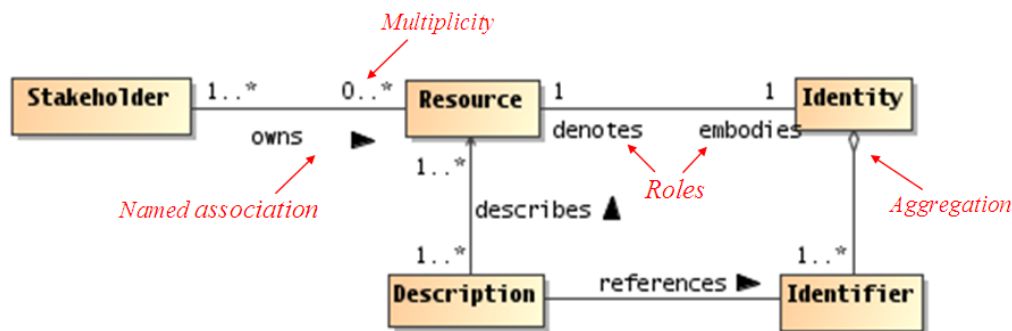
241 An abstraction or representation of some aspect of a thing (in this case, a
242 system)

243 Each architectural model is developed using the methods established by its associated
244 architectural viewpoint. An architectural model may participate in more than one view.

245 **1.3.2 UML Modeling Notation**

246 To help visualize structural and behavioral architectural concepts, it is useful to depict
247 them using an open standard visual modeling language. Although many architecture
248 description languages exist in practice, we have adopted the Unified Modeling
249 Language™ 2 (UML® 2) [UML 2] as the primary viewpoint modeling language. It
250 should be noted that while UML 2 is used in this Reference Architecture, formalization
251 and recommendation of a UML Profile for SOA is beyond the scope of this specification.
252 Every attempt is made to utilize normative UML unless otherwise noted.

253 Figure 1 illustrates an annotated example of a UML class diagram that is used to
254 represent a visual model depiction of the Resources Model in the Service Ecosystem
255 View (Section 3).



256

257 *Figure 1 Example UML class diagram—Resources.*

258 Lines connecting boxes (classifiers) represent associations between things. An
259 association has two roles (one in each direction). A role can have multiplicity, for
260 example, one or more (“1..*”) **stakeholders** own zero or more (“0..*”) **resources**. The
261 role from classifier A to B is labeled closest to B, and vice versa, for example, the role
262 between **resource** to **Identity** can be read as **resource** embodies **Identity**, and
263 **Identity** denotes a **resource**.

264 Mostly, we use named associations, which are denoted with a verb or verb phrase
265 associated with an arrowhead. A named association reads from classifier A to B, for
266 example, one or more **stakeholders** owns zero or more **resources**. Named
267 associations are a very effective way to model relationships between concepts.

268 An open diamond (at the end of an association line) denotes an aggregation, which is a
269 part-of relationship, for example, **Identifiers** are part of **Identity** (or conversely, **Identity**
270 is made up of **Identifiers**).

271 A stronger form of aggregation is known as composition, which involves using a filled-in
272 diamond at the end of an association line (not shown in above diagram). For example,
273 if the association between **Identity** and **Identifier** were a composition rather than an
274 aggregation as shown, deleting **Identity** would also delete any owned **Identifiers**.
275 There is also an element of exclusive ownership in a composition relationship between
276 classifiers, but this usually refers to specific instances of the owned classes (objects).

277 This is by no means a complete description of the semantics of all diagram elements
278 that comprise a UML class diagram, but rather is intended to serve as an illustrative
279 example for the reader. It should be noted that this Reference Architecture utilizes
280 additional class diagram elements as well as other UML diagram types such as
281 sequence diagrams and component diagrams. The reader who is unfamiliar with the
282 UML is encouraged to review one or more of the many useful online resources and
283 book publications available describing UML (see, for example, www.uml.org).
284

285 **1.4 Viewpoints of this Reference Architecture**

286 This Reference Architecture is partitioned into three views that conform to three primary
287 viewpoints: the Service Ecosystems View; the Realizing Service Oriented Architecture
288 view, and the Owning Service Oriented Architectures view. For this Reference
289 Architecture, there is a one-to-one correspondence between viewpoints and views (see
290 Table 1)

Viewpoint Element	Viewpoint		
	<i>Service Ecosystem</i>	<i>Realizing Service Oriented Architectures</i>	<i>Owning Service Oriented Architecture</i>
Main concepts	Captures what SOA means for people to participate in a service ecosystem.	Deals with the requirements for constructing a SOA.	Addresses issues involved in owning and managing a SOA.
Stakeholders	People using SOA, Decision Makers, Enterprise Architects, Standards Architects and Analysts.	Standards Architects, Enterprise Architects, Business Analysts, Decision Makers.	Service Providers, Service Consumers, Enterprise Architects, Decision Makers.
Concerns	Conduct business safely and effectively.	Effective construction of SOA-based systems.	Processes for engaging in a SOA are effective, equitable, and assured.
Modeling Techniques	UML class diagrams	UML class, sequence,, component, activity, communication, and composite structure diagrams	UML class and communication diagrams

291 Table 1 Viewpoint specifications for the OASIS Reference Architecture Foundation for SOA

292 **1.4.1 Service Ecosystem Viewpoint**

293 The Service Ecosystem Viewpoint is intended to capture what using a SOA-based
 294 system means as an environment for people to conduct their business. We do not limit
 295 the applicability of SOA-based systems to commercial and enterprise systems. We use
 296 the term **business** to include any activity of interest to a user; especially activities
 297 shared by multiple users.

298 The stakeholders who have key roles in or concerns addressed by this viewpoint are
 299 decision makers and *people*. The primary concern for people is to ensure that they can
 300 use a SOA to conduct their business in a safe and effective way. For decision makers,
 301 their primary concern revolves around the relationships between people and
 302 organizations using systems for which the decision makers are responsible.

303 Given the public nature of the Internet, and the intended use of SOA to allow people to
 304 access and provide services that cross **ownership boundaries**, it is necessary to be
 305 able to be somewhat explicit about those boundaries and what it means to cross an
 306 **ownership boundary**.

307 1.4.2 Realizing Service Oriented Architectures Viewpoint

308 The Realizing Service Oriented Architectures Viewpoint focuses on the infrastructure
309 elements that are needed to support the construction of SOA-based systems. From this
310 viewpoint, we are concerned with the application of well-understood technologies
311 available to system architects to realize the vision of a SOA that may cross **ownership**
312 **boundaries**.

313 The stakeholders are essentially anyone involved in designing, constructing and
314 deploying a SOA-based system.

315 1.4.3 Owning Service Oriented Architectures Viewpoint

316 The Owning Service Oriented Architectures Viewpoint addresses the issues involved in
317 owning a SOA as opposed to using one or building one. Many of these issues are not
318 easily addressed by automation; instead, they often involve people-oriented processes
319 such as governance bodies.

320 Owning a SOA-based system involves being able to manage an evolving system. In
321 our view, SOA-based systems are more like ecosystems than conventional applications;
322 the challenges of owning and managing SOA-based systems are the challenges of
323 managing an ecosystem. Thus, in this view, we are concerned with how systems are
324 managed effectively, how decisions are made and promulgated to the required end
325 points, and how to ensure that people may use the system effectively and that malicious
326 people cannot easily corrupt it for their own gain.

327 1.5 Terminology

328 The key words “MUST”, “MUST NOT”, “REQUIRED”, “SHALL”, “SHALL NOT”,
329 “SHOULD”, “SHOULD NOT”, “RECOMMENDED”, “MAY”, and “OPTIONAL” in this
330 document are to be interpreted as described in **[RFC2119]**.

331 References are surrounded with **[square brackets and are in bold text]**.

332 Terms such as this “Reference Architecture” and “Reference Architecture Foundation”
333 refer to this document, and “the Reference Model” refers to the OASIS Reference Model
334 for Service Oriented Architecture”. **[SOA-RM]**.

335 1.5.1 Usage of Terms

336 Certain terms are used throughout this document, such as model, action, and rule that
337 have formal definitions within the Reference Architecture. Where a reference to a
338 formally defined concept is intended, we use a bold font such as **actor**, **action** and **joint**
339 **action**. In addition, these words are hyperlinked to their definition within the document.
340 Where the more colloquial and informal meaning is intended, words are used without
341 special emphasis.

342 1.6 References

343 1.6.1 Normative References

344 **[ANSI/IEEE 1471]** *IEEE Recommended Practice for Architectural Description of*
345 *Software-Intensive Systems*, American National Standards

- 346 Institute/Institute for Electrical and Electronics Engineers,
347 September 21, 2000.
- 348 **[ISO/IEC 42010]** International Organization for Standardization and
349 International Electrotechnical Commission, *System and software*
350 *engineering — Recommended practice for architectural description*
351 *of software-intensive systems*, July 15, 2007.
- 352 **[RFC2119]** S. Bradner, *Key words for use in RFCs to Indicate Requirement*
353 *Levels*, <http://www.ietf.org/rfc/rfc2119.txt>, IETF RFC 2119, March
354 1997.
- 355 **[SOA-RM]** OASIS Standard, "Reference Model for Service Oriented
356 Architecture 1.0, 12 October 2006. [http://docs.oasis-open.org/soa-](http://docs.oasis-open.org/soa-rm/v1.0/soa-rm.pdf)
357 [rm/v1.0/soa-rm.pdf](http://docs.oasis-open.org/soa-rm/v1.0/soa-rm.pdf)
- 358 **[UML 2]** *Unified Modeling Language: Superstructure*, Ver. 2.1.1, OMG
359 Adopted Specification, OMG document formal/2007-02-05, Object
360 Management Group, Needham, MA, February 5, 2007.
- 361 **[WA]** Architecture of the World Wide Web, W3C, 2004.
362 <http://www.w3.org/TR/webarch>.
- 363 **[WSA]** David Booth, et al., "Web Services Architecture", W3C Working
364 Group Note, World Wide Web Consortium (W3C) (Massachusetts
365 Institute of Technology, European Research Consortium for
366 Informatics and Mathematics, Keio University), February, 2004.
367 <http://www.w3.org/TR/2004/NOTE-ws-arch-20040211/>

368 1.6.2 Non-Normative References

- 369 **[BLOOMBERG/SCHMELZER]** Jason Bloomberg and Ronald Schmelzer, *Service*
370 *Orient or Be Doomed!*, John Wiley & Sons: Hoboken, NJ, 2006.
- 371 **[COX]** D. E. Cox and H. Kreger, "Management of the service-oriented
372 architecture life cycle," "IBM Systems Journal" "44", No. 4, 709-
373 726, 2005
- 374 **[ERA]** A. Fattah, "Enterprise Reference Architecture," paper presented at
375 22nd Enterprise Architecture Practitioners Conference, London, UK,
376 April 2009.
- 377 **[ITU-T Rec. X.700 | ISO/IEC 10746-3:1996(E)]** Information processing systems—
378 Open Systems Interconnection—Basic Reference Model—Part 4:
379 Management Framework", International Telecommunication Union,
380 International Organization for Standardization and International
381 Electrotechnical Commission, Geneva, Switzerland, 1989.
- 382 **[NEWCOMER/LOMOW]** Eric Newcomer and Greg Lomow, *Understanding SOA with*
383 *Web Services*, Addison-Wesley: Upper Saddle River, NJ, 2005.
- 384 **[OECD]** Organization for Economic Cooperation and Development,
385 Directorate for Financial, Fiscal and Enterprise Affairs, OECD
386 Principles of Corporate Governance, SG/CG(99) 5 and 219, April
387 1999.

- 388 [TOGAF v9] *The Open Group Architecture Framework (TOGAF) Version 9*
389 *Enterprise Edition*, The Open Group, Doc Number: G091, February
390 2009.
- 391 [WEILL] Harvard Business School Press, *IT Governance: How Top*
392 *Performers Manage IT Decision Rights for Superior Results*, Peter
393 Weill and Jeanne W. Ross, 2004
- 394 [DAMIANOU] Nicodemos C. Damianou , Thesis - A Policy Framework for
395 Management of Distributed Systems, University of London,
396 Department of Computing, 2002.
- 397 [LEVESON] Nancy G. Leveson, *Safeware: System Safety and Computers*,
398 Addison-Wesley Professional, Addison-Wesley Publishing
399 Company, Inc.: Boston, pg. 181, 1995.
- 400 [STEEL/NAGAPPAN/LAI] Christopher Steel and Ramesh Nagappan and Ray Lai,
401 *core Security Patterns:Best Practices and Strategies for J2EE, Web*
402 *Services and Identity Management*, Prentice Hall: 2005
- 403 [ISO/IEC 27002] International Organization for Standardization and
404 International Electrotechnical Commission, *Information technology*
405 *-- Security techniques – Code of practice for information security*
406 *management*, 2007
- 407 [SOA NAV] Heather Kreger and Jeff Estefan (Eds.), “Navigating the SOA Open
408 Standards Landscape Around Architecture,” Joint Paper, The Open
409 Group, OASIS, and OMG, July 2009.
410 <http://www.opengroup.org/projects/soa/uploads/40/20044/W096.pdf>

411 2 Architectural Goals and Principles

412 In this section, we identify both the goals of this Reference Architecture and the
413 architectural principles that underlie our approach to the Reference Architecture.

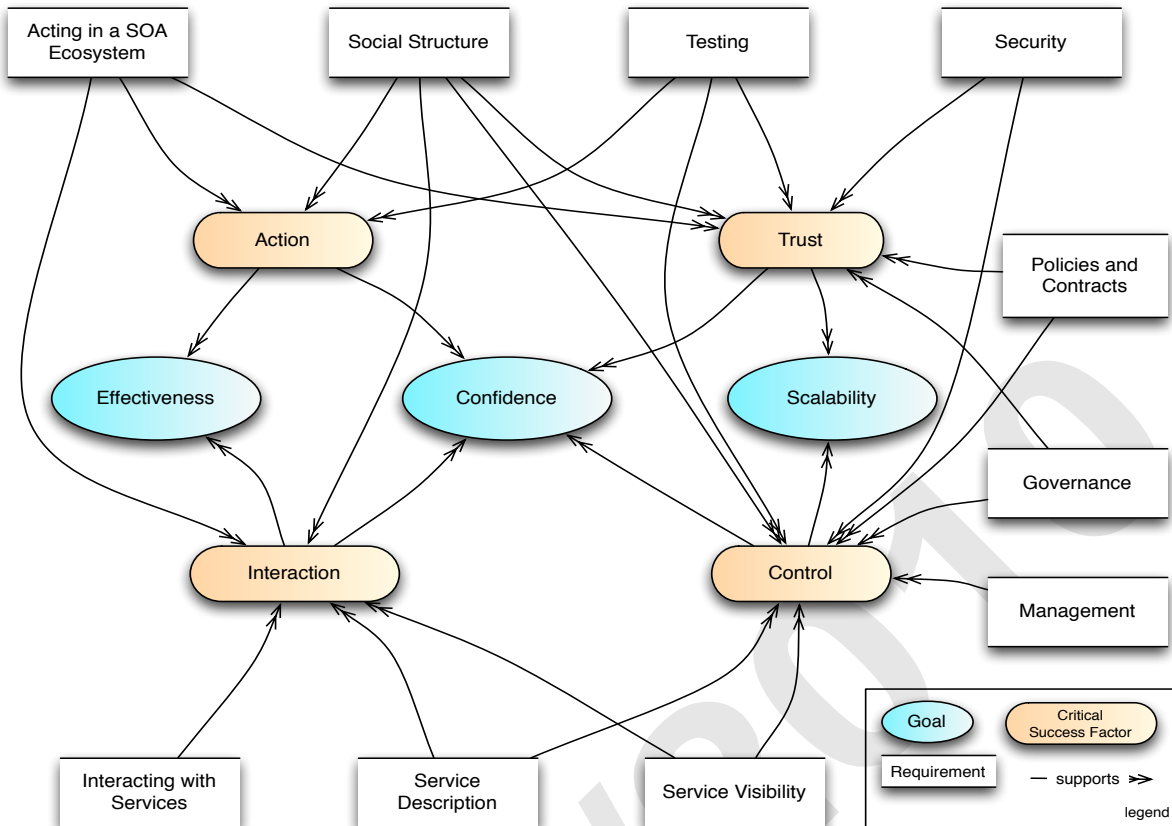
414 2.1 Goals and Critical Success Factors of this Reference Architecture

415 There are three principal goals of this Reference Architecture:

- 416 1. that it shows how SOA-based systems can effectively enable participants with
417 needs to interact with services with appropriate capabilities;
- 418 2. that participants can have a clearly understood level of confidence as they
419 interact using SOA-based systems; and
- 420 3. SOA-based systems can be scaled for small or large systems as needed.

421 There are four factors that are identified as being critical in the achievement of these
422 goals:

- 423 1. there must be a substantial account of participants' action within the ecosystem;
- 424 2. there must be an account of how participants' internal perceptions of the
425 reliability of others guides their behavior (i.e., the trust that participants may or
426 may not have in others)
- 427 3. there must be an account of how participants can interact with each other; and
- 428 4. there must be an account of how the management and governance of the entire
429 SOA ecosystem can be arranged.



430
431 *Figure 2 Critical Factors Analysis of the Reference Architecture*

432 Figure 2 represents a Critical Factors Analysis (CFA) diagram demonstrating the
433 relationship between the primary goals of this reference architecture, critical factors that
434 determine the success of the architecture and individual elements that need to be
435 modeled.

436 A CFA is a structured way of arriving at the requirements for a project, especially the
437 quality attribute (non-functional) requirements; as such, it forms a natural complement to
438 other requirements capture techniques such as use-case analysis, which are oriented
439 more toward functional requirements capture. The CFA requirement technique and the
440 diagram notation are summarized in Appendix B.

441 2.1.1 Goals

442 2.1.1.1 Effectiveness Goal

443 A primary purpose of this architecture is to show what is involved in SOA-based
444 systems to ensure that participants can use the facilities of the system to meet their
445 needs. This does not imply that every need has a SOA solution, but for those needs
446 that can benefit from a SOA approach, we look at what is needed to use the SOA
447 paradigm effectively.

448 The key factors that govern effectiveness from a participants' perspective are **action** –
449 especially action that crosses ownership boundaries – and **interaction** with other
450 participants in the ecosystem.

451 **2.1.1.2 Confidence Goal**

452 SOA-based systems should enable service providers and consumers to conduct their
453 business with the appropriate level of confidence in the interaction. Confidence is
454 especially important in situations that are high-risk; this includes situations involving
455 multiple ownership domains as well as situations involving the use of sensitive
456 resources.

457 Confidence has many dimensions: confidence in the successful interactions with other
458 participants, confidence in the assessment of trust, as well as confidence that the
459 ecosystem is properly managed.

460 **2.1.1.3 Scalability Goal**

461 The third goal of this Reference Architecture is scalability. In architectural terms, we
462 determine scalability in terms of the smooth growth of complexity of systems as the
463 number and complexity of services and interactions between participants increases.
464 Another measure of scalability is the ease with which interactions can cross ownership
465 boundaries.

466 **2.1.2 Critical Success Factors**

467 A critical success factor (CSF) is a property of the intended system, or a sub-goal that
468 directly supports a goal and there is strong belief that without it the goal is unattainable.
469 CSFs are not necessarily measurable in themselves. As illustrated in Figure 2, CSFs
470 can be associated with more than one goal.

471 In many cases critical success factors are often denoted by adjectives: reliability,
472 trustworthiness, and so on. In our analysis of the SOA paradigm however, it seems
473 more natural to identify four critical concepts (nouns) that characterize important
474 aspects of SOA:

475 **2.1.2.1 Action**

476 Participants' primary mode of participation in a SOA ecosystem is action; typically action
477 in the interest of achieving some desired **real world effect**. Understanding how action is
478 related to SOA is then critical to the paradigm.

479 Action is, of course, pervasive in the ecosystem; and many models in this Reference
480 Architecture address aspects of action. However, **action** is the central theme of the
481 models that labeled "Acting in a Social Context" and "Acting in a SOA Ecosystem".

482 **2.1.2.2 Trust**

483 The viability of a SOA ecosystem depends on participants being able to effectively
484 measure the trustworthiness of the system and of participants. Trust is a private
485 assessment of a participant's belief in the integrity and reliability of the SOA ecosystem
486 (see Section 3.2.3).

487 Trust can be analyzed in terms of trust in infrastructure facilities (otherwise known as
488 reliability), trust in the relationships and effects that are realized by interactions with
489 services, and trust in the integrity and confidentiality of those interactions particularly
490 with respect to external factors (otherwise known as security).

491 Note that there is a distinction between trust in a SOA-based system and trust in the
492 capabilities accessed via the SOA-based system. The former focuses on the role of
493 SOA-based systems as a *medium* for conducting business, the latter on the
494 trustworthiness of participants in such systems. This architecture focuses on the former,
495 while trying to encourage the latter.

496 2.1.2.3 Interaction

497 In order for a SOA ecosystem to function, it is essential that the means for participants
498 to interact with each other is available throughout the system. Interaction encompasses
499 not only the mechanics of communication but also the means for discovering and
500 offering communication.

501 2.1.2.4 Control

502 Given that a large-scale SOA-based system may be populated with many services, and
503 used by large numbers of people; managing SOA-based systems properly is a critical
504 factor for engendering confidence in them. This involves both managing the services
505 themselves and managing the relationships between people and the SOA-based
506 systems they are utilizing; the latter being more commonly identified with governance.

507 The governance of SOA-based systems requires an ability for decision makers to be
508 able to set policies about participants, services, and their relationships. It requires an
509 ability to ensure that policies are effectively described and enforced. It also requires an
510 effective means of measuring the historical and current performances of services and
511 participants.

512 The scope of management of SOA-based systems is constrained by the existence of
513 multiple ownership domains. Management may include setting policies such as
514 technology choices but may not, in some cases, include setting policies about the
515 services that are offered.

516 2.2 Principles of this Reference Architecture

517 The following principles serve as core tenets that guided the evolution of this Reference
518 Architecture. The ordered numbering of these principles does not imply priority order.

519 **Principle 1: Technology Neutrality**

520 **Statement:** Technology neutrality refers to independence from particular technologies.

521 **Rationale:** We view technology independence as important for three main reasons:
522 technology specific approach risks confusing issues that are technology
523 specific with those that are integrally involved with realizing SOA-based
524 systems; and we believe that the principles that underlie SOA-based
525 systems have the potential to outlive any specific technologies that are
526 used to deliver them. Finally, a great proportion of this architecture is
527 inherently concerned with people, their relationships to services on SOA-
528 based systems and to each other.

529 **Implications:** This Reference Architecture must be technology neutral, meaning that we
530 assume that technology will continue to evolve, and that over the lifetime
531 of this architecture that multiple, potentially competing technologies will
532 co-exist. Another immediate implication of technology independence is

533 that greater effort on the part of architects and other decision makers to
534 construct systems based on this architecture is needed.

535 **Principle 2: Parsimony**

536 Statement: Parsimony refers to economy of design, avoiding complexity where
537 possible and minimizing the number of components and relationships
538 needed.

539 Rationale: The hallmark of good design is parsimony, or “less is better.” It promotes
540 better understandability or comprehension of a domain of discourse by
541 avoiding gratuitous complexity, while being sufficiently rich to meet
542 requirements.

543 Implications: Occam’s (or Ockham’s) Razor applies, which states that the explanation
544 of any phenomenon should make as few assumptions as possible that
545 account for the observations. With respect to this Reference Architecture,
546 we aim to avoid the elaboration of unnecessary details.

547 The complement of a parsimonious design is a feature-rich design.
548 Parsimoniously designed systems tend to have fewer features.

549 **Principle 3: Separation of Concerns**

550 Statement: Separation of Concerns refers to the ability to cleanly delineate
551 architectural models in such a way that an individual stakeholder or a set
552 of stakeholders that share common concerns only see those models that
553 directly address their respective areas of interest. This principle could just
554 as easily be referred to as the Separation of Stakeholder Concerns
555 principle, but the focus here is predominantly on loose coupling of models.

556 Rationale: As SOA-based systems become more mainstream, and as they start to
557 become increasingly complex, it will be extremely important for the
558 architecture to be able to scale. Trying to maintain a single, monolithic
559 architecture that incorporates all models to address all possible system
560 stakeholders and their associated concerns will not only rapidly become
561 unmanageable with rising system complexity, but it will become unusable
562 as well.

563 Implications: This is a core tenet that drives this Reference Architecture to adopt the
564 notion of architectural viewpoints and corresponding views. A *viewpoint*
565 provides the formalization of the groupings of models representing one set
566 of concerns relative to an architecture, while a *view* is the actual
567 representation of a particular system. The ability to leverage an industry
568 standard that formalizes this notion of architectural viewpoints and views
569 helps us better ground these concepts for not only the developers of this
570 Reference Architecture but also for its readers. The IEEE Recommended
571 Practice for Architectural Description of Software-Intensive Systems
572 **[ANSI/IEEE 1471-2000::ISO/IEC 42010-2007]** is the standard that serves
573 as the basis for the structure and organization of this Reference
574 Architecture.

575 **Principle 4: Applicability**

576 Statement: Applicability refers to that which is relevant. Here, an architecture is
577 sought that is relevant to as many facets and applications of SOA-based
578 systems as possible; even those yet unforeseen.

579 Rationale: An architecture that is not relevant to its domain of discourse will not be
580 adopted and thus likely to languish.

581 Implications: This Reference Architecture needs to be relevant to the problem of
582 matching needs and capabilities under disparate domains of ownership; to
583 the concepts of “Intranet SOA” (SOA within the enterprise) as well as
584 “Internet SOA” (SOA outside the enterprise); to the concept of “Extranet
585 SOA” (SOA within the extended enterprise, i.e., SOA with suppliers and
586 trading partners); and finally, to “net-centric SOA” or “Internet-ready SOA.”

10/26/2010

587

3 Service Ecosystem View

588

589

590

591

592

593

594

595

596

597

598

No man is an island

*No man is an island entire of itself; every man
is a piece of the continent, a part of the main;
if a clod be washed away by the sea, Europe
is the less, as well as if a promontory were, as
well as any manner of thy friends or of thine
own were; any man's death diminishes me,
because I am involved in mankind.
And therefore never send to know for whom
the bell tolls; it tolls for thee.*

John Donne

599

The *Service Ecosystem View* focuses on what it means for people to conduct their business in the context of a SOA-based system. By business we mean to include any activity entered into whose objective is to satisfy some goal of a participant. By people we include organizations and also automated processes.

603

The people and organizations participating in a SOA-based ecosystem form a community. That community may be a single **enterprise** or a large peer-to-peer network of **enterprises** and individuals. Many of the activities that people engage in are themselves defined by the relationships between people and by the organizations to which they belong.

608

However, the primary motivation for participants to interact with each other is to achieve objectives – to get things done. Describing what it means to *act* in the SOA ecosystem when participants may be in different organizations, with different rules and expectations is one of the primary modeling objectives of this section.

612

Within a SOA ecosystem, the implication is that a dominant mode of communication is electronic, supported by IT resources and artifacts. Since there is inherent indirection involved when people and systems interact using electronic means, we lay the foundations for how *communication* can be used to represent **action**. However, it is important to understand that these are merely tools to an end and are usually not the primary interest of the **participants** of the ecosystem.

618

The Social Structure Model introduces the key elements that underlie the relationships between **participants**. The Acting in a SOA Ecosystem Model introduces the key concepts involved in actions, and shows how **ownership**, risk and transactions are key concepts in the SOA ecosystem.

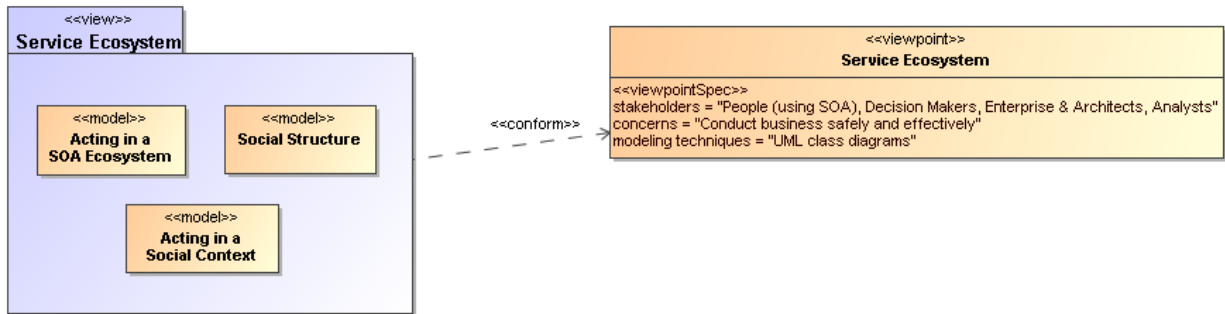
622

The Semantics in a SOA Ecosystem Model focuses on the concepts that underlie meaning within the SOA ecosystem. It introduces the fundamental concept of **proposition** and shows how **policies**, **facts** and communication are all dependent on semantics.

626

It is impossible to present the SOA ecosystem view in a strictly linear order. This is because the models and concepts involved are densely connected. To ease the reader's burden in following forward – and backward – references we make extensive use of intra-document hyperlinks such this link to **action**; a concept that is used

630 throughout. In fact *all* references to defined concepts are implemented as hyperlinks to
 631 their definition.

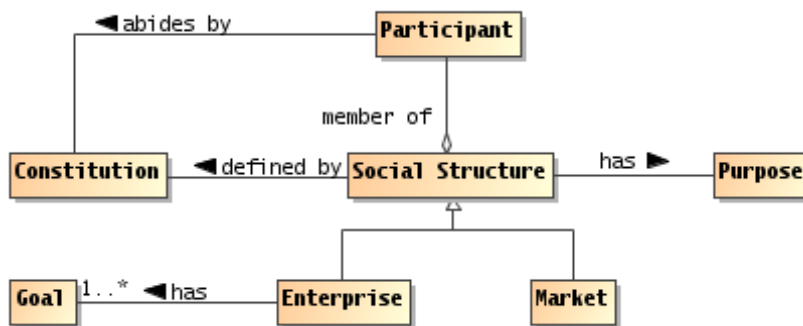


632
 633 Figure 3 Model elements described in the Service Ecosystem view

634 3.1 Social Structure Model

635 The actions undertaken by **participants** in a SOA ecosystem are performed in a *social*
 636 *context* that defines the relationships between the **participants**. We can formalize that
 637 context as a **social structure**.

638 The Social Structure Model emphasizes the importance of defining and understanding
 639 the implications of crossing **ownership boundaries**. It is, for example, the foundation
 640 for understanding security, governance and management in the SOA ecosystem.
 641 However, the primary function of the Social Structure Model is to explicate the
 642 relationships between an individual **participant** and the social context of that
 643 **participant**.



644
 645 Figure 4 Social Structure

646 Social Structure

647 A **social structure**¹⁰ is a nexus of relationships amongst **participants** for a
 648 **purpose**

649 A **social structure** represents a collection of **participants**, but a collection that is
 650 brought together for a **purpose**. A nexus of relationships is a set of relationships. I.e.,
 651 there may be a large number of different kinds of relationships between **participants** in

¹⁰ Social structures are sometimes referred to as social institutions.

652 a **social structure**. The organizing principle for these relationships is the **social**
653 **structure**'s purpose.

654 A **social structure** may have any number of **participants**, and a given **participant** can
655 be a member of multiple **social structures**. Thus, there may be interaction among
656 **social structures**, sometimes resulting in disagreements when the premises of the
657 **social structures** do not align.

658 A **social structure** has a **purpose** – the reason for which it exists. All **social**
659 **structures** have a **purpose**, some **social structures** also have **goals**.

660 In this Reference Architecture Foundation, we are concerned primarily with **social**
661 **structures** that reflect the anticipated **participants** in SOA-based systems; these are
662 often embodied in legal and quasi-legal frameworks; i.e., they have some rules that are
663 commonly understood. For example, an **enterprise** is a common kind of **social**
664 **structure** that embodies a form of hierarchic organization; an online chat room
665 represents a **social structure** that is very loose and anarchic – modeled here as a
666 **market**. At the other extreme, the legal frameworks of entire countries and regions also
667 count as **social structures**.

668 It is not necessarily the case that the **social structures** involved in a particular
669 interaction are explicitly identified. For example, when a customer buys a book over the
670 Internet, the **social structure** that defines the validity of the transaction is often the
671 legal framework of the region associated with the book vendor. This legal jurisdiction
672 qualification is typically buried in the fine print of the service description.

673 **Constitution**

674 A **constitution** is a set of rules, written or unwritten, that defines a **social**
675 **structure**.

676 Every **social structure** defines the rules by which **participants** interact with each other
677 within the structure. A **social structure**'s rules are *abided to* by the **participants**. In
678 some cases, this is based on an explicit agreement, in other cases participants behave
679 as though they agree to the **constitution** without a formal agreement. In other cases,
680 participants abide by the rules with some degree of reluctance – this is an issue raised
681 later on when we discuss governance in SOA-based systems.

682 The SOA ecosystem is marked by two primary forms of **social structure** – the **peer**
683 **social structure** which is primarily oriented to the interrelationship between participants
684 within the ecosystem and the **enterprise** which represents a kind of *composite*
685 **participant** – an entity that has sufficient internal cohesiveness that allows us to
686 consider it as a potential **stakeholder** in its own right.

687 **Enterprise**

688 An **enterprise** is a **social structure** with internally established **goals** that reflect
689 its **purpose** and that can act as a **participant** within other **social structures**.

690 The **enterprise** is marked out as being associated with internal **goals** in a way that a
691 strict market type of **social structure** is not.

692 **Market**


693 A **market social structure** is the locus of interaction between participants who
694 are peers of one another.

695 If an **enterprise** is often the focus of the differing **roles** and **responsibilities** of
696 members, a **market** or meeting place is more concerned with the exchange of goods
697 and services for mutual benefit.

698 It is entirely possible for a given interaction between participants to take place within a
699 **social structure** that is an **enterprise** as well as being a **market** place. However,
700 interactions within a **peer social structure** are inherently across **ownership**
701 **boundaries**.

702 3.1.1 Actors, Delegates and Participants

703 As noted above, **social structures** have members. Some of these members are active,
704 others are not active within the **social structure** and others act on behalf of others.



705
706 *Figure 5 Actors, Participants and Delegates*

707 Stakeholder

708 A **stakeholder** in the SOA ecosystem is an individual entity, human or non-
709 human, or organization of entities that has an interest in the state of the
710 ecosystem.

711 Actor

712 An **actor** is an entity, human, non-human or organization of entities, that is
713 capable of **action**.

714 The concept of **actor** encompasses many kinds of entities, human and corporate
715 participants, even semi-autonomous computational agents. Two important kinds of
716 **actor** are **participants** and **delegates**.

717 Participant

718 A **participant** is a **stakeholder** that is an **actor** in a SOA ecosystem.

719 Non-Participant Stakeholder

720 A **non-participant stakeholder** is any **stakeholder** who is not an **actor** in the
721 ecosystem.

722 **Stakeholders** do not necessarily participate in the SOA ecosystem; indeed, the interest
723 of **non-participant stakeholders** may be in not acting directly but in still realizing the
724 benefits of a well-functioning ecosystem and not suffering unwanted consequences.

725 There are two main classes of such non-participatory **stakeholders**: third parties who
726 are affected by someone's use or provisioning of a service, and regulatory agencies
727 who wish to control the outcome of service interactions in some way.

728 An example of an affected third party may be someone using the service infrastructure
729 whose activities are impeded because an errant **participant** is consuming excessive
730 bandwidth in another interaction.

731 **Delegate**

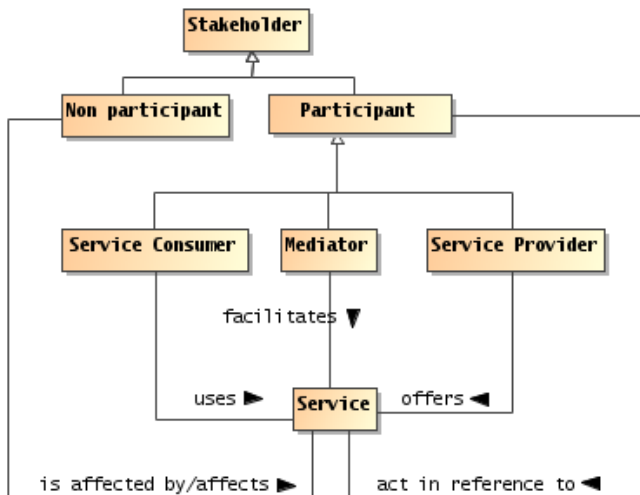
732 A **delegate** is an **actor** that is acting on behalf of a **stakeholder**.

733 There are many kinds of entities that may function in a SOA ecosystem. For example,
734 there may be software agents that permit people to offer and interact with services;
735 there may be **delegates** that represent the interests of other **stakeholders** – such as
736 security agents charged with managing the security of the ecosystem.

737 In the different models in this architecture we use the **actor** concept when it is not
738 important whether the entity involved is a **delegate**, **participant** or some other entity. If
739 the **actor** is acting on behalf of another, then we use the **delegate** concept. If the **actor**
740 is a **stakeholder** in the ecosystem then we use **participant**.

741 **3.1.1.1 Service Providers and Consumers**

742 Above, we distinguish between **participants** and nonparticipants. In a SOA
743 ecosystem, several types of **participants** play prominent **roles** – in particular, offering
744 and using **services**.



745
746 *Figure 6 Service Participants*

747 **Service Provider**

748 A **service provider** is a **participant** that offers a **service**.

749 **Service Consumer**

750 A **service consumer** is a **participant** that interacts with a **service** in order to
751 address a consumer need.

752 It is a common understanding that **service consumers** typically initiate service
753 interactions. Again, this is not necessarily true in all situations (for example, in publish-
754 and-subscribe scenarios, a **service consumer** may initiate an initial subscription, but
755 thereafter, the interactions are initiated by publishers). As with service providers, several
756 **stakeholders** may be involved in a service interaction supporting the consumer.

757 In many scenarios, **service providers** and **service consumers** do not represent truly
758 symmetric **roles**: each **participant** has different objectives and often has different

759 capabilities. However, the objectives and the conditions under which those objectives
760 align are critical for a successful interaction to proceed.

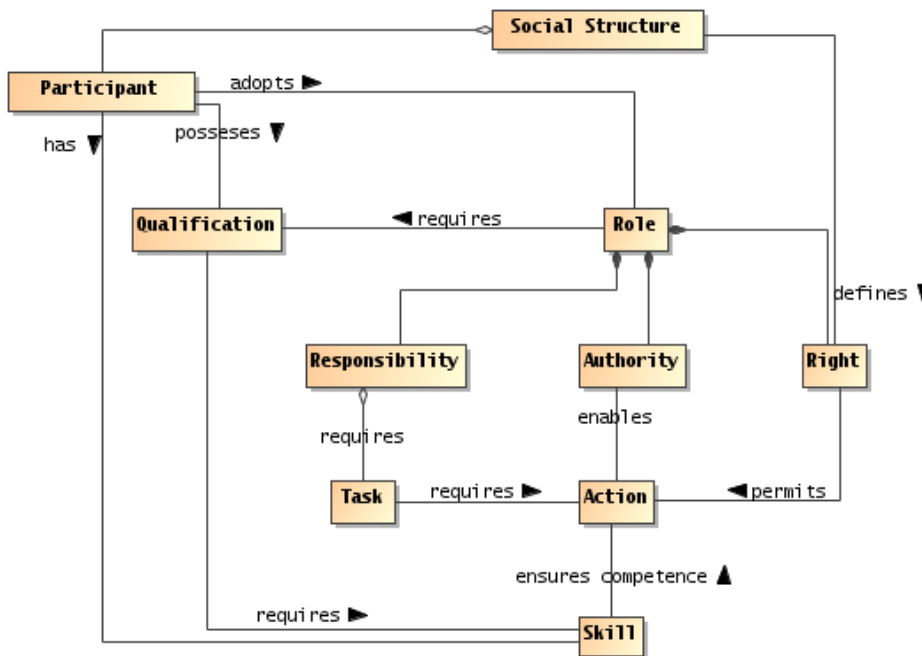
761 Service Mediator

762 A **service mediator** is a **participant** that facilitates the offering or use of services
763 in some way.

764 There are many kinds of **mediator**, for example a registry is a kind of mediator that
765 permits providers and consumers to find each other. Another example might be a filter
766 that enhances another service by translating messages between English and Japanese.

767 3.1.2 Roles in Social Structures

768 **Social structures** are abstractions: a **social structure** cannot directly perform **actions**
769 – only people or automated processes following the instructions of people can actually
770 do things. However, an **actor** may act on behalf of a **social structure** and certainly acts
771 within a **social structure** depending on the **roles** that the **actor** assumes.



772
773 *Figure 7 Roles, Rights and Responsibilities*

774 Role

775 A **role** is an identified relationship between a **participant** and a **social structure**
776 that defines the **rights**, **responsibilities**, **qualifications**, and **authorities** of that
777 **participant** within the context of the **social structure**.

778 A **participant** can be identified with one or more **roles**. Someone in authority in the
779 **social structure** may have formally designated the **participant** as assuming the **role**
780 with associated **rights** and **responsibilities**.

781 Conversely, someone who exhibits **qualification** and skill may by consensus assume
782 the **role** without any formal designation. For example, an office administrator who has
783 demonstrated facility with personal computers may be known as the 'goto' person for
784 people who need help with their computers.

785 Note that, while many **roles** are clearly identified, with appropriate names and
786 definitions of **responsibilities**, it is also entirely possible to separately bestow **rights**,
787 **responsibilities** and so on; usually in a temporary fashion. For example, when a
788 company president delegates the responsibility of ensuring that the company accounts
789 are correct to the chief engineer, this does not imply that the chief engineer is assuming
790 the full **role** of the company accountant.

791 **Role Player**

792 A **role player** is an **actor** that assumes a **role**. I.e., his actions and/or stance with
793 respect to other **participants** is consistent with the **role**.

794 In order for a person to act on behalf of some other person or on behalf of some legal
795 entity, it is required that they have the ability to do so and the **authority** to do so.

796 **Right**

797 A **right** is a predetermined **permission** that permits an **actor** to perform some
798 **action** or assumes a **role** in relation to the **social structure**.

799 **Rights** often are associated with additional constraints. For example, in most
800 circumstances, sellers have a right to refuse service to potential customers; but often
801 may only do so based on certain criteria.

802 **Authority**

803 **Authority** is the **right** to act on behalf of an organization or another person.

804 Usually, **authority** is constrained in terms of the kinds of actions that are authorized,
805 and in terms of the necessary skills and **qualifications** of the persons invoking the
806 **authority**.

807 An entity may authorize or be assigned another entity to act as its **delegate**. Often the
808 actions that are so authorized are restricted in some sense.

809 **Rights**, **authorities**, **responsibilities** and **roles** form the foundation for the security
810 architecture of the Reference Architecture. **Rights** and **responsibilities** have similar
811 structure to permissive and obligation policies; except that the focus is from the
812 perspective of the constrained **participant** rather than the constrained actions.

813 **Responsibility**

814 A **responsibility** is an **obligation** on a **role player** to perform some **action** or to
815 adopt a stance in relation to other **role players**.

816 **Skill**

817 A **skill** is a competence or capability to achieve some real world effect.

818 Skills are typically associated with **roles** in terms of requirements: a given role
819 description may require that the role player has a certain skill.

820 **Qualification**

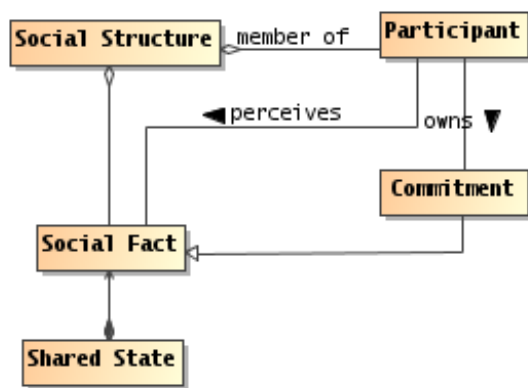
821 A **qualification** is a public recognition that an **actor** is capable of assuming a
822 **role**.

823 In most cases a **qualification** involves the recognition within a particular **social**
824 **structure** that the **actor** has the necessary set of **skills** that enable the **actor** to
825 assume some **role**.

826 There is a distinction between a **skill** – which is capability that a participant may have to
827 act – and a publicly accepted acknowledgement of the capability – i.e., a **qualification**.
828 For example, someone may have the skills to fly an airplane but not have a pilot's
829 license. Conversely, someone may have a pilot license, but because of some temporary
830 cause be incapable of flying a plane (they may be ill for example).

831 3.1.3 Shared State and Social Facts

832 Many of the actions performed by people and most of the important aspects of a
833 person's state are inherently social in nature. The social context of an **action** is what
834 gives it much of its meaning. We call actions in society **social actions** and, those facts
835 that are understood in a society, **social facts**. It is often the case that **social actions**
836 give rise to **social facts**.



837
838 *Figure 8 Shared State and Social Facts*

839 Social Fact

840 A **social fact** is a **fact** about a **social structure**.

841 **Social structures** provide a context in which **social facts** are given their meaning. For
842 example, the existence of a valid purchase order with a particular customer has a
843 meaning that is defined primarily by the company itself, together with the society that
844 the company is part of.

845 Compared to facts about the natural world, **social facts** are inherently abstract: they
846 only have meaning in the context of a **social structure**.

847 **Social facts** typically require some kind of ritual to establish the validity of the fact itself.
848 For example, the existence of an agreed contract typically requires both parties to sign
849 papers and to exchange those papers. If the signatures are not performed correctly, or if
850 the parties are not properly empowered to perform the ritual, then it is as though nothing
851 happened.

852 In the case of agreements reached by electronic means, this involves the exchange of
853 electronic messages; often with special tokens being exchanged in place of a hand-
854 written signature.

855 State

856 **State** is the condition that an entity is in at a particular time.

857 **State** is characterized by a set of **facts** that is true of the entity – in effect we are
858 concerned only with aspects of an entity that are potentially measurable. In principle,
859 the total **state** of an entity (or the world as a whole) is unbounded – potentially not even
860 countable. However, the principal means that an **actor** is aware of the **state** of an entity
861 is by what is known about the **state** – i.e., facts about the **state** – which is typically
862 finite.

863 For example, the total **state** of a lightbulb includes the temperature of the filament of the
864 bulb. It also includes a great deal of other **state** – the composition of the glass, the dirt
865 that is on the bulb’s surface and so on. However, an **actor** may be primarily interested
866 in whether the bulb is ‘on’ or ‘off’ and not on the amount of dirt accumulated. The
867 **actor’s** characterization of the **state** of the bulb reduces to the **fact**: ‘bulb is now on’
868 (say).

869 One of the key distinctions is between an **actor’s** private state that can only be
870 accessed by that **actor** and the public or **shared state** that is accessible to more than
871 one **actor**.

872 **Private State**

873 The **private state** of a **actor** is the **state** that is accessible to only that **actor**.

874 **Shared State**

875 **Shared state** is the **state** that is accessible by multiple **actors**.

876 Note that **shared state** *does not* imply the state *is* accessible to all **actors**. It simply
877 refers to that subset of state that *may* be accessed.

878 **Commitment**

879 A **commitment** is a **social fact** about the future that a **actor** is responsible for
880 ensuring is satisfied.

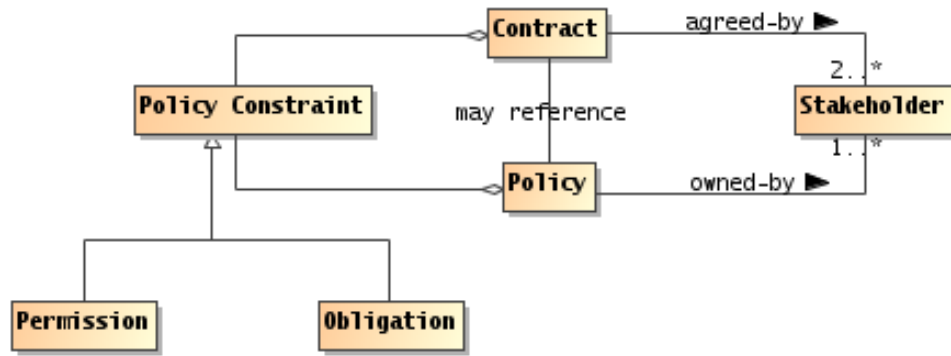
881 A **commitment** to deliver some good or service is a classic example of a fact about the
882 future. **Commitments** play an important role in **transactions** and **exchanges** between
883 **participants**.

884 Other important classes of **social facts** include the **policies** adopted by an organization,
885 any agreements that it is holding for **participants**, and the assignment of **participants**
886 to **roles** within the organization.

887 **3.1.4 Policies and Contracts**

888 As noted in the Reference Model, a **policy** represents some constraint that is
889 promulgated and enforced by a **stakeholder**. A **contract**, on the other hand,
890 represents an agreement by two or more **participants**. Enforcement of **contracts** may
891 or may not be the responsibility of the parties to the agreement.

892 In both cases, however, **policies** and **contracts** are naturally part of the state shared by
893 the **participants** in a **social structure**. In addition, some **policies** and **contracts** are
894 more integrally part of the **social structure** itself: forming part of the **social structure’s**
895 **constitution**.



896

897

Figure 9 Policies and Contracts

898

Policy

899

A **policy** is an **assertion** that is promulgated by a **stakeholder** in such a way as to enforce the **assertion's proposition**.

900

901

Policies can often be said to be about something – they have an object. For example, there may be **policies** about the use of a **service**. In addition, and crucially for the purposes of this Reference Architecture, **policies** have a **owner** – the **stakeholder** that asserts the **policy** and is responsible for ensuring the enforcement of the **policy**.

902

903

904

905

Thirdly, **policies** represent constraints – some measurable limitation on the state or behavior of the object of the **policy**.

906

907

Policy Owner

A **policy owner** is a **stakeholder** that asserts and enforces the **policy**.

908

909

Policy Subject

A **policy subject** is an **actor** who is subject to the constraints of a **policy** or **contract**.

910

911

912

Policy Constraint

A **policy constraint** is a measurable **proposition** that characterizes the constraint that the **policy** is about.

913

914

915

Policy Topic

A **policy topic** is a category of **policy constraints** that are related by a area of concern.

916

917

918

For example, security represents an area of concern with many related **policy constraints**.

919

920

Policy Object

A **policy object** is an identifiable **state**, **action** or **resource** that is potentially constrained by the **policy**.

921

922

923

Contract

A **contract** represents an agreement by two or more **participants** to constrain their behavior and/or state.

924

925

926

Both **policies** and **contracts** embody a sense of enforcement: a constraint that is not enforced is a wish rather than a **policy**. **Policies** are **owned** by individual (or

927

928 aggregate) **stakeholders**; these **stakeholders** are responsible for ensuring that the
929 constraints in the **policy** are enforced – although, of course, the actual enforcement
930 may be delegated to a different mechanism.

931 However, where the **policy** or **contract** refers to a fundamental aspect of the **social**
932 **structure** itself, then such **policies** and **contracts** can be said to be part of the
933 **constitution** of the **social structure**.

934 **Permission**

935 A **permission** is a constraint that concerns allowed **actions** that an **actor** may
936 perform and/or the **states** the **actor** may be in.

937 Note that permissions are distinct from ability and from authority. Authority refers to the
938 legitimate nature of an **action** as performed by an **actor** on behalf of a **social**
939 **structure**, whereas **permission** does not always involve acting on behalf of anyone.

940 **Obligation**

941 An **obligation** constraint prescribes the **actions** that an **actor** must perform
942 and/or the **states** the **actor** must be in.

943 For example, once the **service consumer** and provider have entered into an
944 agreement to provide and consume a service, both **participants** incur **obligations**: the
945 consumer is obligated to pay for the service and the provider is obligated to provide the
946 service.

947 **Obligations** to maintain **state** may range from a requirement to maintain a minimum
948 balance on an account through a requirement that a service provider ‘remember’ that a
949 particular **service consumer** is logged in.

950 **Obligations** and **permissions** have a positive form and a negative form. A positive
951 **permission** refers to something that you may do, a negative **permission** refers to
952 something you should not do. A positive **obligation** refers to something you must do (or
953 maintain in the case of a **state** -oriented **obligation**. A negative **obligation** is similar to
954 a negative **permission**: you may be obliged to not perform some **action**.

955 **Obligations** and **permissions** are combinable, in the sense that you may have a
956 positive permission constraint (for example, you may use encryption in your messages),
957 whereas a negative permission constraint indicates that there is something you may not
958 do. Similarly, a positive **obligation** may be something like you must keep the balance of
959 your account positive; whereas an example of a negative **obligation** may be that the
960 bank will not cover a check for more than the balance in your account.

961 **Permissions** are often checkable a priori: before the intended **action** or before entering
962 the constrained **state**. However, **obligations** can normally only be verified a posteriori
963 through some form of auditing or verification process.


964 **3.1.5 Resources and Ownership**

965 Fundamentally, we view **ownership** as a relationship between a **stakeholder** and a
966 **resource**, where the **owner** has certain **rights** and **obligations** with respect to the
967 **resource**.

968 Typically, the **ownership** relationship is one of control: the **owner** of a **resource** can
969 control some aspect of the **resource**.

970 3.1.5.1 Resources

971 A **resource** is an entity that has value to someone. Key to the concept of **resources** is
972 that they are identifiable and may have an **owner**. We define **resource** as follows:



973
974 *Figure 10 Resources*

975 **Resource**

976 A **resource** is any entity **owned** by a **stakeholder** and that has **identity**.

977 A **resource** may have more than one identifier, but any well-formed identifier should
978 unambiguously resolve to the intended resource.

979 An important class of **resource** is the class of capabilities that underlie services. Other
980 examples of resources are services themselves, descriptions of entities (a kind of meta-
981 resource), IT infrastructure elements used to deliver services, contracts and **policies**,
982 and so on.

983 **Identity**

984 **Identity** is the collection of individual characteristics by which an entity, human or
985 nonhuman, is recognized or known.

986 The ability to unambiguously identify a resource in a SOA interaction is critical to
987 determine such things as authorizations, to understand what functions are being
988 performed and what the results mean, and to ensure repeatability or characterize
989 differences with future SOA interactions.

990 **Identifier**

991 An **identifier** is any sequence of characters that unambiguously connects a
992 resource with a particular identity.

993 **Identifiers** typically require a context in order to establish the connection between the
994 identifier and the resource. In a SOA ecosystem, it is good practice to use globally
995 unique identifiers; for example globally unique IRIs.

996 A given **resource** may have multiple identifiers, with different utility for different
997 contexts.

998 **Description**

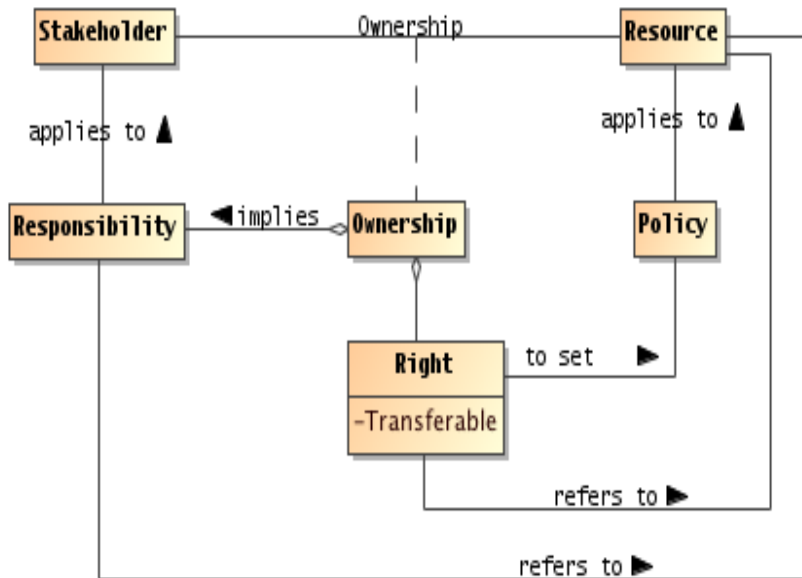
999 A **description** is a set of **assertions** about a **resource**.

1000 Description as related to the SOA ecosystem is discussed in detail in Section 4.1.

1001 3.1.5.2 Owning Resources

1002 Ownership

1003 **Ownership** is a set of **rights** and **responsibilities** that a **stakeholder** has in
1004 relation to a **resource**; including the **right** to transfer that **ownership** to another
1005 entity.



1006
1007 *Figure 11 Resource Ownership*

1008 To own a **resource** implies taking responsibility for creating, maintaining, and if it is to
1009 be available to others, provisioning the **resource**. More than one **stakeholder** may
1010 own different **rights** associated with a given **resource**, such as one **stakeholder** having
1011 the **right** to deploy a capability as a service, another owning the **rights** to the profits that
1012 result from using the capability, and yet another owning the **rights** to use the service.

1013 One who owns a **resource** may delegate **rights** and **responsibilities** to others, but
1014 typically retains some responsibility to see that the delegated **responsibilities** are met.
1015 There may also be joint **ownership** of a **resource**, where the **responsibility** is shared.

1016 A crucial property that distinguishes **ownership** from a more limited **right** to use is the
1017 **right** to transfer **ownership** to another person or organization. When a **resource** is
1018 being used without being owned, there is an implied requirement that at the end of a
1019 period of time the **rights** and **responsibilities** relating to the **resource** will be returned
1020 to the original owner of the **resource**.

1021 **Ownership** is defined in relation to the **social structure** relative to which **rights** and
1022 **responsibilities** are exercised. In particular, there may be constraints on how
1023 **ownership** may be transferred. For example, a government may not permit a
1024 corporation to transfer assets to a subsidiary in a different jurisdiction.

1025 Ownership Boundary

1026 An **ownership boundary** is the **social structure** within which the **rights** and
1027 **responsibilities** associated with a particular **ownership** may be recognized.

1028 Individual **participants** are *within* an **ownership boundary** in relation to a specific
1029 owned **resource** if they are members of the **social structure** that defines what
1030 **ownership** means in relation to the **resource**.

1031 3.1.6 Life-cycle of Social Structures

1032 Life Cycle

1033 A **social structure** has a **life cycle** associated with it.

1034 3.2 Acting in a SOA Ecosystem Model

1035 At the core of **participants**' interest in a SOA ecosystem is the concept of **action** –
1036 **participants** act in order to further their **goals**. Critically, **participants**' actions may
1037 involve systems and **resources** that do *not belong to them*.

1038 In this model we establish the key principles of **action** as an abstract concept. We
1039 elaborate on **action** in the context of a social context as **joint action**. Put simply, **joint**
1040 **actions** are simply coordinated **actions** that involve more than one **actor**.

1041 Given that **participants** must communicate with each other we also show the role of
1042 **communication** in **action** and **joint action**.

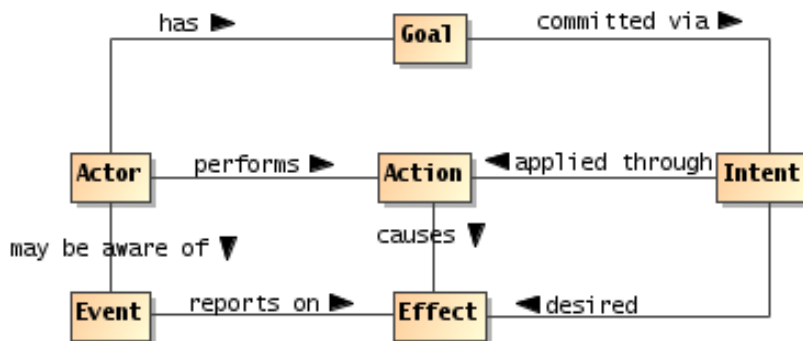
1043 A key aspect of **joint action** revolves around the **trust** that both parties must exhibit in
1044 order to participate in joint actions. This willingness to act and a mutual understanding
1045 of the information exchanged and the expected results is the particular focus of
1046 Section 3.2.3.

1047 3.2.1 Action and Joint Action

1048 Entities act in order to achieve their **goals**. The most basic form of **action** is that
1049 performed by a single **actor**. However, the form of **action** that is of most interest within
1050 a SOA ecosystem is that involving more than one **actor** – i.e., **joint action**.

1051 3.2.1.1 Action and Actors

1052 As modeled in Figure 12, **actions** are **purposeful** processes that **actors** engage in in
1053 order to achieve particular objectives.



1054
1055 *Figure 12 Actions, Real World Effect and Events*

1056 Action

1057 An **action** is the application of **intent** to achieve an **effect**.

1058 The aspect of **action** that distinguishes it from mere force or accident is that someone
1059 or something *intends* the **action** to occur. Of course, it should be pointed out that the
1060 actual effect of an **action** may not be the same as the intended **effect**.

1061 Whilst this definition of **action** is very general, we are mostly concerned in **actions** that
1062 have some impact on the SOA ecosystem.

1063 **Objective**

1064 An **objective** is a **real world effect** that an **actor** uses an **action** or set of
1065 **actions** to achieve.

1066 **Objectives** refer to **real world effects** that **actors** believe are achievable by a specific
1067 **action** or set of **actions**. Below, we define **goal** in terms of the state that an **actor** is
1068 attempting to achieve. In general, a **goal** is not linked with a specific **action** in the same
1069 way than an **objective** is.

1070 For example, someone may wish to have enough light to read a book. In order to satisfy
1071 that goal, the reader walks over to flip a light switch. The objective is to turn on the
1072 lamp, the goal is to be able to read.

1073 **Intent**

1074 **Intent** is the internal planning and orienting of an **actor** to achieve an **objective**.

1075 **Intent** is that internal process within an **actor** that leads it to perform an **action**. In order
1076 for an **actor** to perform an **action**, it must have internally planned to perform the **action**
1077 and it must engage the **action**.

1078 For example, for a person to pick up a cup with their hand, the brain has to plan and
1079 coordinate the movement of the muscles of the arm and the fingers. In addition, the
1080 brain must be oriented to the task of picking up the cup – i.e., the person must not only
1081 move muscles but must also intend to pick up the cup.

1082 In some situations it may be difficult for an observer in the SOA ecosystem to determine
1083 an **actor's** actual **intent**. This is because, **intent** is inherently part of an **actor's** private
1084 state.

1085 However, in most cases, **participants** in a SOA ecosystem make an assumption of
1086 **implied intent**.

1087 **Implied Intent**

1088 **Implied intent** is the **intent** that may be inferred by an observer when witnessing
1089 an **actor** perform an **action**.

1090 I.e., if an **actor** is seen performing an **action**, it is assumed that the **actor** also intended
1091 to perform the **action** – it was not an accident, nor was it the **action** of another **actor**.

1092 Much of the infrastructure of interaction is there to eliminate the potential for accidental
1093 or malicious actions.

1094 **3.2.1.2 Actions and Effects**

1095 **Effect**

1096 An **effect** is a measurable change in the state of the ecosystem resulting from an
1097 **action**.

1098 Note the normal **intent** of applying an **action** is to cause an **effect** that reflects the
1099 **actor's** objectives. However, there is often the possibility that the actual effects will
1100 include unintended consequences that fall outside of, and may even run counter to, the
1101 intent of the **actor**.

1102 Real World Effect

1103 A **real world effect** is a change to **shared state**.

1104 The intuition behind **real world effect** is that something actually happened as a result of
1105 some **action**. In particular, something changed that is relevant to a **participant**.

1106 Changes in the ecosystem may be *reported* by means of notifications of **events**:

1107 Event

1108 An **event** is an occurrence of which at least one **participant** has an interest in
1109 being aware.

1110 An **event** is often a corollary to **action**, when an **actor** performs an **action**, this causes
1111 an **effect** which is of interest to at least the **actor** performing the **action**; often to other
1112 **participants** also.

1113 However, there are other kinds of **events** which are not easily ascribed to the **actions** of
1114 **actors** – for example parts of a communications network becoming unavailable due to
1115 electrical storms. Such **events** are still important and measurable.

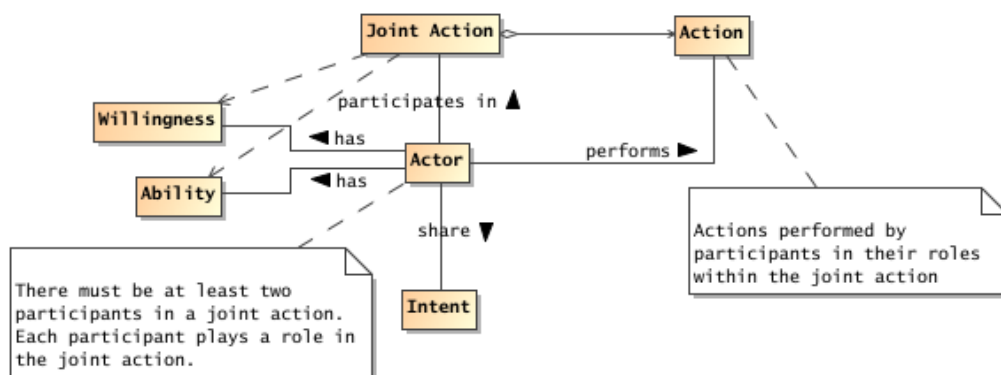
1116 Event Notification

1117 An **event notification** is the **action** of an **actor** becoming aware of an **event**.

1118 Note that, while performing an **action** may be an **event** that other **participants** have an
1119 interest in, an **event notification** that reports an **action** is not the same as the **action**
1120 itself.

1121 3.2.1.3 Joint Actions

1122 Joint actions are the foundation for understanding interaction between **participants** in a
1123 SOA ecosystem. In this Reference Architecture, we see joint actions at least two levels:
1124 as communication and as **participants** using and offering services.



1125
1126 *Figure 13 Joint Action*

1127 Joint Action

1128 A **joint action** is a coordinated set of **actions** involving the efforts of two or more
1129 **actors** to achieve an **effect**.

1130 In order for multiple **actors** to participate in a **joint action**, they must each act according
1131 to their **role** within the **joint action**. For example, a common example of a **joint action**
1132 is for one **actor** to speak to another.¹¹ A communication between **actors** cannot take
1133 place unless there is both a speaker and a listener – although it is not necessarily
1134 required that they both be active simultaneously.

1135 Note that the **effect** of a **joint action** is *not* always attributable to one or more **effects** of
1136 the individual **actions** of the participating **actors**.

1137 **Choreography**

1138 A **choreography** is a description of the individual **actions** to be performed by
1139 **actors** in order to successfully participate in one or more **joint actions**.

1140 A **choreography** defines how individual **actions** performed by **actors** can be
1141 aggregated together to denote **joint actions**.

1142 In any human context **joint actions** abound: people talking to each other, people buying
1143 and selling, people arranging their lives. In addition, **joint action** is at the heart of
1144 interactions within the context of a SOA ecosystem.

1145 There is another sense in which **joint actions** abound: even within a single incident of
1146 interaction there are typically several overlapping **joint actions** layered one on top of
1147 the other.

1148 For example, consider the ramifications of one **actor** exchanging a message with
1149 another **actor**. The act of communication that takes place between the **actors** is a **joint**
1150 **action** – specifically a **communicative action**. The **purpose** of the communication may
1151 be that the speaking **actor** wishes to invoke a **service action** involving the **service** that
1152 the listening **actor** offers. The **service action** itself may be a **joint action** such as
1153 making arrangements for the future **roles** of the **participants**. In this situation there are
1154 at least *three* senses of **joint action** – the **communicative action**, the **service action**
1155 and the **social action** establishing a **social fact**; there may be others.

1156 One might wish to insist that the highest-level interpretation of such a collection of **joint**
1157 **actions** was the only ‘real’ **joint action** and that the others (communication and service
1158 invocation) were merely subsidiary or implementing the **social action**.

1159 However, this is an interpretation that requires a strong viewpoint. Different viewpoints
1160 will lead to different **joint actions** being interpreted as most important. For example,
1161 from the viewpoint of ecosystem **governance**, the nature and fact of the established
1162 **agreement** may be dominant; from the viewpoint of ecosystem security, the
1163 **communicative action** may be dominant.

1164 In summary, the concept of **joint action** allows us to honor the fact that both parties in
1165 an interaction are required for there to be an actual effect; it allows us to separate out
1166 the different levels of the interaction into appropriate semantic layers; and it allows us to
1167 recombine those layers in potentially different ways whilst still achieving the intended
1168 **real world effects** of **action** in a SOA ecosystem.

¹¹ Where speaking and listening includes electronic message sending and receiving.

1169 **3.2.2 Goals and Capabilities**

1170 **Actors** participating in a SOA ecosystem are often attempting to get other **actors** to *do*
1171 something. For example, a customer trying to buy a book has to convince the book
1172 selling **service** to deliver the book. Conversely the book selling service has to convince
1173 the customer to pay for it.

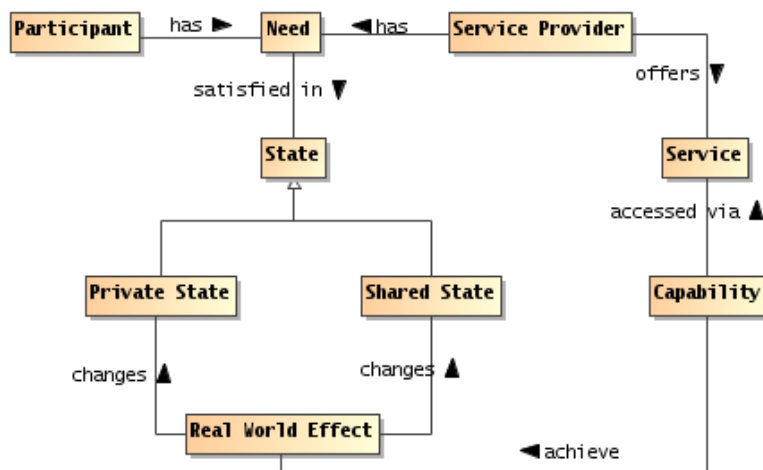
1174 There is a reason that **actors** are engaged in this **choreography**: different **actors** have
1175 different **needs** and have different **capabilities** for satisfying those needs.

1176 **Goal**

1177 A **goal** is an **assertion** about the **state** that an **actor** is seeking to establish or
1178 maintain.

1179 In the **Reference Model** this is known as **need**.

1180 In general, there is a *subsumption* relationship between **actors' goals** and their
1181 **objectives**: an **objective** can be considered to be *consistent* with one of more **goals**.
1182 Generally, a **goal** is a long term state of the world that may be, in practice, difficult to
1183 measure. On the other hand, an **objective** is a directly measurable and preferably
1184 predictable outcome of a particular **action** or set of **actions**.



1185
1186 *Figure 14 Needs and Capabilities*

1187 **Capability**

1188 A **capability** is an ability to achieve a **real world effect**.

1189 Both **goals** and the effects of using **capabilities** are expressed in terms of **state**: a
1190 need is expressed as a condition on the desired state and the **real world effect** of using
1191 capabilities is a change in the state of the world.

1192 In any interaction between **actors** it is reasonable to assume that they all have **goals**.
1193 For example, in the case of a **service provider**, the service's owners aim to address
1194 their **goals** as well as the needs of **participants** who use the **service**.

1195 When one **actor** agrees to a course of **action** as a result of its interactions with other
1196 **actors** – for example by employing a **capability** to achieve an **effect** – it is **adopting** an
1197 **objective**.

1198 The process of adopting objectives is required in order for **actors** to get their needs met
1199 by means of other **actors'** actions.

1200 Objective Adoption

1201 An **actor** may adopt an **objective** as a result of interacting with another **actor**.

1202 The **objective** that an **actor** adopts need not be identical with the **actor** that originated
1203 the **action**; however, it should be *consistent* with that **actor's** **goals**. This discrepancy
1204 may be for several reasons: the responding **actor** may not have the appropriate
1205 **semantic engagement** to adopt the originating **actor's objective**; or the **adopted**
1206 **objective** may lead to partial fulfillment of the originating **actor's** goals.

1207 One consequence of an **actor** adopting an **objective** on behalf of another **actor** is that
1208 the **actor** becomes **accountable** to the latter for the successful satisfaction of the
1209 **objective**.

1210 Accountability

1211 An **actor** is **accountable** to another **actor** when the former consents to achieve
1212 an identified **objective**.

1213 It is possible to characterize an **actor's** **accountability** in terms of **obligation policies**
1214 that are in force in relation to that **actor**.

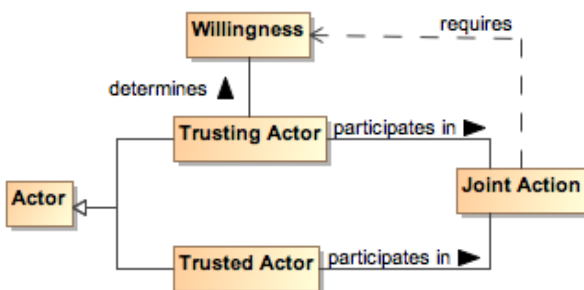
1215 3.2.3 Trust and Risk

1216 For interaction to occur between **actors**, they must be able to interact but especially
1217 they must be **willing** to participate in the various **joint actions** that make up the
1218 interaction.

1219 Willingness

1220 **Willingness** is the internal commitment of an **actor** to carry out its part of a **joint**
1221 **action**.

1222 An important prerequisite for **willingness** is **trust**. Without **trust** the required
1223 **willingness** will be significantly reduced or even non-existent.



1224
1225 Figure 15 Trusting Actor and Willingness

1226 Trust

1227 **Trust** is a private assessment or internal perception that some entity will perform
1228 actions that will lead to an identifiable set of **real world effects**.

1229 Implied in this definition of **trust** is that it is intrinsic to an **actor's** private perception; as
1230 opposed to an extrinsic property of **actors**. In addition, *measurement* is integral to **trust**
1231 – it is something that is evaluated.

1232 Note that while normally **trust** is associated with positive attributes, it is not strictly
1233 necessary that that is so. The complement of **trust** is the concept of **risk** – the
1234 assessment of undesirable potential:

1235 **Risk**

1236 **Risk** is a private assessment or internal perception that certain undesirable **real**
1237 **world effects** may come into being.

1238 Both **trust** and **risk** involve an active **actor** making the assessment and another **actor**
1239 whose merits are being evaluated:

1240 **Trusting Actor**

1241 A **trusting actor** is an **actor** who establishes and maintains **willingness** to
1242 proceed with an interaction based on its trust of other **actors**.

1243 **Trusted Actor**

1244 A **trusted actor** is an **actor** with which a **trusting actor** has sufficient **trust** to
1245 proceed with an interaction.

1246 **Trust** is involved in all interactions – it is necessary for *all* the **actors** involved to **trust**
1247 each other at least to the extent required for continuance of the interaction. The degree
1248 and nature of that **trust** will likely be different for each **actor**. Like **actors** in a stage
1249 play, each **actor** in an interaction must assess the reliability of its partners.

1250 An **actor** perceiving **risk** may take actions to mitigate the **risk**. At one extreme this will
1251 result in a refusal to interact with the suspect **actor**. Alternately, it may involve adding
1252 protection – for example by using encrypted communication and/or anonymization – to
1253 reduce the perception of **risk**. Often standard procedures are put in place to increase
1254 **trust** and to mitigate **risk**.

1255 **3.2.3.1 Assessing Trust and Risk**

1256 The assessments of **trust** and **risk** are based on **evidence** available to the **trusting**
1257 **actor**. In general, **actors** will seek **evidence** from their private knowledge of the **trusted**
1258 **actor** as well as **evidence** of the **reputation** of the **trusted actor**.

1259 **Evidence**

1260 **Evidence** is the accumulation of **facts** by which a **trusting actor** can assess
1261 **trust** and **risk**.

1262 The **evidence** that an **actor** uses to assess **trust** and **risk** may include a history of
1263 previous interaction between the **trusting** and **trusted actors** or previous interactions
1264 of the **trusted actor** with other **actors** for which the **facts** of their interactions are public;
1265 i.e., the **trusted actor's reputation**.

1266 **Reputation**

1267 **Reputation** is the set of **social facts** that form publicly available **evidence** by
1268 which a **trusting actor** can assess **trust** and **risk** of an **actor**.

1269 **Reputation** is the **evidence** that is publicly available within a **social structure** that
1270 members of that **social structure** have access to in order to help measure the
1271 trustworthiness or otherwise of a **trusted actor**.

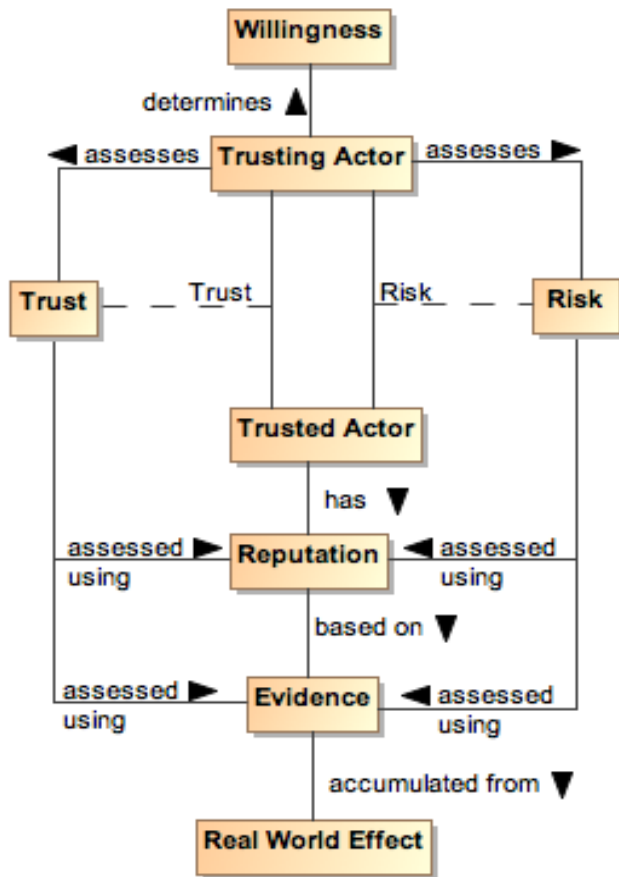


Figure 16 Assessing Trust and Risk

1272
1273

1274 **Trust** is based on the confidence the **trusting actor** has in the accuracy and sufficiency
1275 of the gathered **evidence** and the degree to which any assessment is appropriate for
1276 the situation for which trust is being assessed.

1277 In most situations, assessment of **trust** is not binary, i.e. an **actor** is not completely
1278 **trusted** or **untrusted**, because there is typically some degree of uncertainty in the
1279 accuracy or completeness of the evidence or the assessment. Similarly, there is
1280 uncertainty in the amount and consequences of potential **risk**.

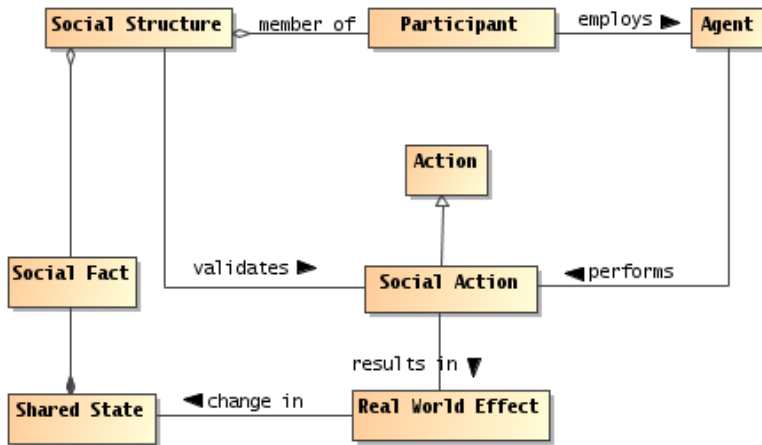
1281 The relevance of **trust** depends on the assessment of **risk**. If there is little or no
1282 perceived **risk**, then the degree of **trust** may not be relevant in assessing possible
1283 actions. For example, most people consider there to be an acceptable level of risk to
1284 privacy when using search engines, and submit queries without any sense of trust being
1285 considered.

1286 As perceived **risk** increases, the issue of **trust** becomes more of a consideration. For
1287 interactions with a high degree of **risk**, the **trusting actor** will typically require stronger
1288 or additional **evidence** when evaluating the balance between **risk** and **trust**.

1289 3.2.4 Social Actions

1290 In the context of SOA ecosystems, **actions** are often social in nature — one **participant**
1291 is asking another to do something that is directly related to the organization(s) that they

1292 are part of — and goal oriented — the **purpose** of interacting with a service is to satisfy
 1293 a need by attempting to ensure that a remote entity applies its capabilities to the need.



1294
 1295 *Figure 17 Acting within Social Structures*

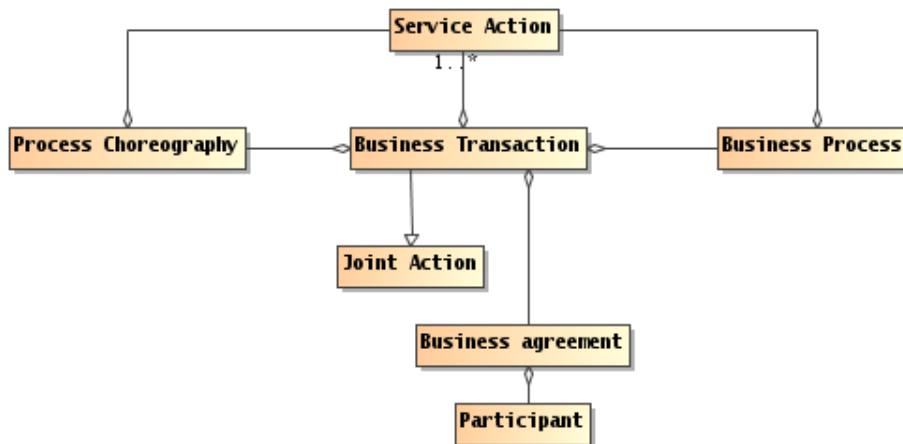
1296 **Social Action**

1297 **Social actions** are **joint actions** that are performed in order to affect the **social**
 1298 **structure** itself.

1299 A **social action** is defined primarily by the effect it has on the relationship between
 1300 **participants** and state of a **social structure** by establishing one or more new **social**
 1301 **facts**.

1302 **3.2.4.1 Transactions and Exchanges**

1303 An important class of **social action** is the **business transaction**, or **contract**
 1304 **exchange**. Many interactions between **participants** in a SOA ecosystem are based
 1305 around business transactions.



1306
 1307 *Figure 18 Business Transaction*

1308 **Business Transaction**

1309 A **business transaction** is a **social action** engaged in by two or more
 1310 **participants** in which the **ownership** of one of more **resources** is exchanged.

1311 A classic **business transaction** is buying some good or service, but there is a huge
1312 variety of kinds of possible business transactions.

1313 Key to the concept of business transaction is the contract or agreement to exchange.
1314 The form of the contract can vary from a simple handshake to an elaborately drawn
1315 contract with lawyers giving advice from all sides.

1316 A completed transaction establishes a set of **social facts** relating to the exchange;
1317 typically to the changes of **ownerships** of the **resources** being exchanged.

1318 **Business Agreement**

1319 A **business agreement** is an **agreement** entered into by two or more
1320 **participants** that constrains their future behaviors and permitted states.

1321 A **business agreement** is typically associated with **business transactions**: the
1322 transaction is guided by the agreement and an agreement can be the result of a
1323 transaction.

1324 **Business transactions** often have a well defined life-cycle: a negotiation phase in
1325 which the terms of the transaction are discussed, an agreement **joint action** which
1326 establishes the **commitment** to the transaction, an action phase in which the agreed-
1327 upon items are exchanged (they may need to be manufactured before they can be
1328 exchanged), and a termination phase in which there may be long-term **commitments**
1329 by both parties but no particular actions required (e.g., if the exchanged goods are
1330 found to be defective, then there is likely a **commitment** to repair or replace them).

1331 Within the SOA ecosystem, a **business transaction** often represents the top-most
1332 mode of interpretation of **service interactions**. When **participants** interact in a
1333 **service**, they exchange information and perform actions that have an effect in the
1334 world. These exchanges can be interpreted as realizing part of, and in support of,
1335 **business transactions**.

1336 **3.2.5 Communication for Action**

1337 Because there is inherently some separation between **actors** in a SOA ecosystem, they
1338 are effectively driven to use communication techniques to 'get their business done'.
1339 Communication and the interpretation of communicated content is the foundation of all
1340 interaction within the SOA ecosystem.

1341 When an **actor** sends a message to another **actor** there are at least two senses in
1342 which the **actor** can be said to be acting: by communicating with other **actors**; and by
1343 using that communication to communicate a **service action**.

1344 We define the **communicative action** as the **action** of message exchange:

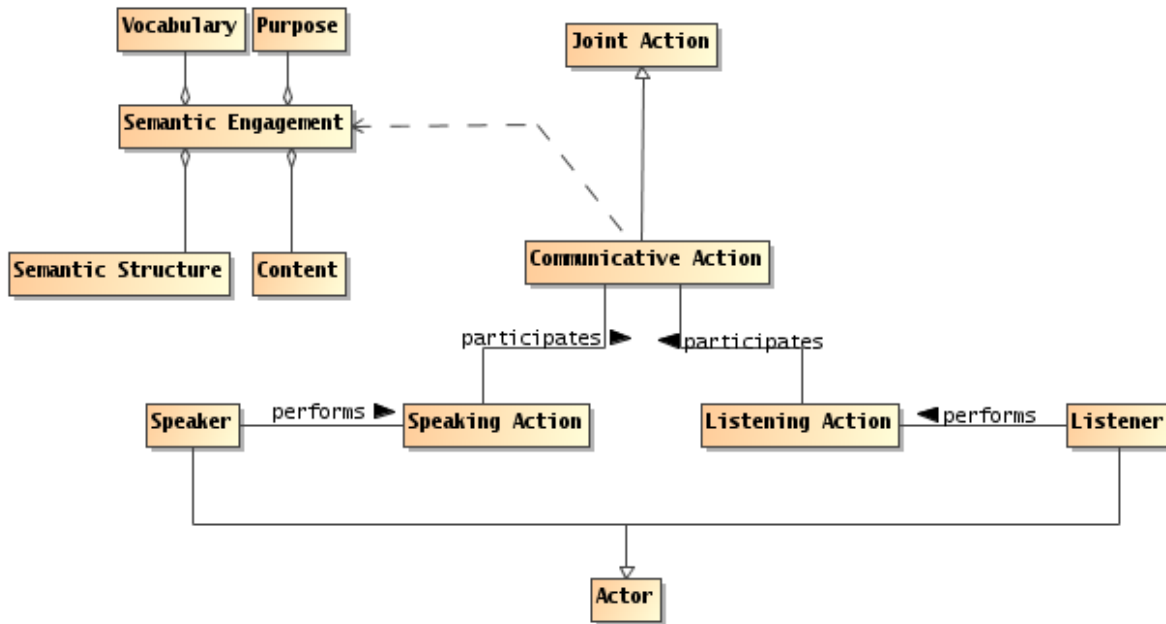


Figure 19 Communication as Joint Action

Communicative Action

A **communicative action** is a **joint action** in which an **actor** communicates with one or more other **actors**.

A **communicative action** has a **speaker** and at least one **listener**; each of whom must perform their part for the **communicative action** to occur. In addition, associated with the **communicative action** is the **content** of the communication – what is being communicated.

Speaker

A **speaker** is an **actor** who originates a **communicative action** by performing those **actions** needed to communicate.

Listener

A **listener** is an **actor** who performs **actions** needed to acquire a communication.

Typically, a **communicative action** involves one **participant** speaking and the other listening simultaneously; although there are many potential important variations, such as broadcast, writing and so on.

In an interaction between **participants** it is highly likely that the **speaking** and **listening roles** alternate throughout the interaction.

A given **communicative action** may have any number of **listeners**. Indeed, in some situations, it may not be possible for the **speaker** to be aware of the **listener** in a **communicative action**; however, this does not change the fundamentals of communication: without both a **speaker** and a **listener** there is no communication.

Content

Content is the information passed from the **speaker** to the **listener** in a **communicative action**.

1372 Note that an **actor** listening to a message not only acquires the message but must also
1373 be able to understand it. The extent of that understanding will depend on the **role** of the
1374 **actor** and its **purpose** of the communication.

1375 Even though communication is effected through **action**, it is not actually effective if the
1376 **listener** cannot understand the content of the communication. However, understanding
1377 can itself be characterized in terms of **semantic engagement**: informally the
1378 relationship an **actor** has with the world; communication in this context.

1379 We can characterize the necessary modes of understanding in terms of a shared
1380 **vocabulary** and a shared understanding of the **purpose** of the communication. More
1381 formally, we can say that a communication has a combination of syntax, **public**
1382 **semantics** and **illocutionary force**.

1383 Illocutionary Force

1384 The **illocutionary force** of a **communicative act** is the proximate **purpose** of
1385 the communication.


1386 For example, a **communicative action** may be a *request*, or it may *inform* the listener
1387 of some fact.

1388 Of course, the *ultimate purpose* for a communication may not be closely related to the
1389 proximate **purpose**. For example, a bank service may *inform* a customer that their
1390 account balance is too low; the ultimate **purpose** being to persuade the customer to
1391 augment the account.

1392 Note that, while it is often easier to visualize the semantics of communication in terms
1393 that reflect human experience, it is not required for interactions between **service**
1394 **consumers** and providers to particularly look like human speech. Machine-machine
1395 communication is typically highly stylized in form, it may have particular forms and it
1396 may involve particular terms not found in everyday human interaction.

1397 3.2.5.1 Using Communication for Service Action

1398 Like **communicative actions**, **service actions**, or **actions** involving a **service**, are
1399 inherently **joint actions** – there can be no **service action** without both the **service** and
1400 the **actor** originating the **action**. However, because there is a gap between the
1401 **participant** performing a service action and the service being acted upon, there must
1402 be a bridge across that gap; bridging this gap relies on the *counts as* relationship.



1403
1404 Figure 20 Communicative actions as Service Actions

1405 Service Action

1406 A **service action** is an element of the action model of a **service**.

1407 **Service actions** are inherently **joint actions**; they require both the entity performing the
1408 action and the service itself to participate in the **action**.

1409 When we state that a **communicative action counts as a service action**, we are
1410 relating a system of communication to a system of **action** against services.¹² Since a
1411 **participant** cannot (normally) act directly on a **service** it must use some means of
1412 mediating the **action**. However, from the perspective of all the **participants** involved,
1413 when a **participant** uses a **communicative action** appropriately, the **participants** are
1414 *expected* to understand the communication *as though* a **service action** were actually
1415 performed.

1416 **Counts as**

1417 **Counts as** is a relationship between two logical systems in which an **action**,
1418 **event** or **concept** in one system can be understood as another **action**, **event** or
1419 **concept** in another system.

1420 The two systems involved in SOA-based systems are the system of communication on
1421 the one hand and the system of services on the other.

1422 For example, suppose that **actor Alpha** wishes to communicate with **actor Beta** that it
1423 agrees to a 'deal' – maybe it is a **contract** that *Alpha* is agreeing to. I.e., *Alpha* must
1424 perform a **communicative action**. The fact of communication is not itself sufficient to
1425 ensure agreement: *Beta* must be able to **interpret** the **content** of the communication as
1426 agreement. In effect there are two coincident actions taking place: the act of
1427 communication and the act of agreeing.

1428 Taken together, the syntax, **semantics**, **vocabulary**, and the **illocutionary force** of
1429 communicated **content** is the basis of all interaction in a SOA ecosystem.

1430 **3.3 Semantics in a SOA Ecosystem Model**

1431 Semantics is important to the SOA ecosystem because it is pervasive: it is explicitly
1432 important in the communication between **actors**, but is also a driver for **policies**, and
1433 many other aspects of the ecosystem.

1434 This model undertakes to establish two primary concepts: that of **assertion** and that of
1435 **semantic engagement**. An **assertion** is a form of utterance, or more generally any
1436 assertion or conclusion that an **actor** may make in the context of the ecosystem;
1437 together with the role or purpose of the utterance. **Semantic engagement** refers to how
1438 an **actor** engages with, or understands, any **assertions** that it may encounter.

1439 **3.3.1 Assertions**

1440 We are concerned with several forms of **assertion** in this Reference Architecture:
1441 including, but not limited to, **facts** which denote **propositions** that are knowable by
1442 **actors**; **goals** and **objectives** which denote **propositions** that **actors** are trying to
1443 satisfy; **policies** which denote **propositions** that **actors** are enforcing and abiding by;
1444 and **contracts** which denote agreements between **actors**.

1445 In general we can characterize **assertions** in terms of a content **proposition** and a
1446 **stance**.

¹² Acting against a service should not be understood to mean acting to foil the effectiveness of the service; but simply as an action involving the normal operation of the service.

1447 **Assertion**

1448 An **assertion** is a **proposition** associated with a particular **stance**.

1449 A **proposition** is a testable predicate about the state of the ecosystem (including
1450 internal state of **actors**). The **stance** denotes the *relationship(s)* between the
1451 **proposition** and some **actor** or **actors**. Informally, we can say that the **proposition**
1452 denotes the ‘what’ of the **assertion** and the **stance** denotes the ‘why’ or ‘what for’ of the
1453 **assertion**.

1454 Typically we also give specific names to particular combinations of **proposition** and
1455 **stance**; for example a **fact** is a **proposition** that an **actor** can ‘know’ – *knowing* a **fact**
1456 is a **stance** that an **actor** can adopt in relation to the **proposition**.

1457 **Proposition**

1458 A **proposition** is an expression, normally in a language that has a well-defined
1459 written form, that denotes some property of a **domain** from the perspective of a
1460 **stakeholder**.

1461 The key properties of **propositions** are that they are expressions – i.e., they have a
1462 particular ‘form’ – and that the truth of a **proposition** is verifiable – using a **decision**
1463 **procedure**. Minimally, verification of a **proposition** is achieved by checking that the
1464 **proposition** and the **domain** are consistent with each other.¹³

1465 There is a distinction between *tautological* and *contingent* **propositions**. A tautology is
1466 something that is universally and necessarily true. A contingent **proposition** is one that
1467 depends on a **domain** for its truth – i.e., it is only true if it can be verified to be true. In
1468 general, contingent **propositions** are often more interesting to **actors** as they represent
1469 knowledge about the world that is potentially subject to change.

1470 The requirements for the written form of **propositions** will vary with the application.
1471 Some highly structured and formalized systems of include various forms of logic.

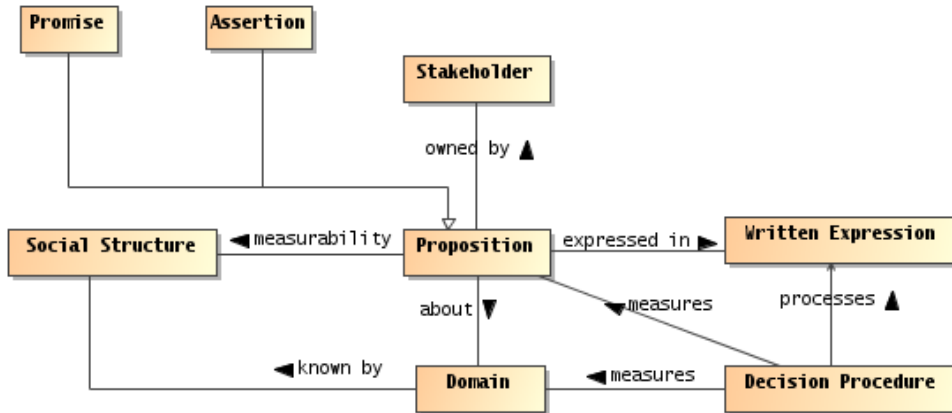
1472 **Written Expression**

1473 The **written expression** of a **proposition** is a formula written in a systematic
1474 system of marks.

1475 A key characteristic of such ‘systems of marks’ is that there are clear syntactic rules for
1476 what constitutes legal expressions. While there may be additional rules for determining
1477 how to interpret **written expressions**, it is generally not desirable to completely
1478 constrain the *interpretation* of **written expressions**.

1479 By **written expressions** we mean a combination of expressions built up from a
1480 **vocabulary** and a set of operators. There is no specific requirement that **written**
1481 **expressions** are actually written down on paper. Typically, **written expressions** take
1482 the form of highly regular tree structures. For example, an invoice will often follow pre-
1483 established standards for communicating invoices.

¹³ We exclude here the special case of proposition known as a tautology. Tautologies are important in the study of logic; the kinds of propositions that we are primarily interested in are those which pertain to the world; and as such are only *contingently* true.



1484
1485 *Figure 21 Propositions*

1486 **Vocabulary**

1487 A **vocabulary** is a set of terms, together with an **interpretation**, which may be
1488 referenced by a **written expression**.

1489 Typically, in a formally expressed **proposition**, the operators are ‘part of the language’
1490 and the lexicon of symbols – or **vocabulary** – is domain-specific referring to entities in
1491 the world.

1492 Particularly in the context of communication between **actors**, the **actors** must be able to
1493 understand the terms appearing in the communication. This is sometimes expressed in
1494 terms of ‘sharing a common **vocabulary**’; in any case, there must be *sufficient* shared
1495 understanding of the elements of the **written expression** for communication to be
1496 successful.

1497 **Vocabularies** are important in other contexts also: for example a **policy** statement may
1498 reference specific entities as the **policy subject**, **policy constraint** and so on. In order
1499 for other **actors** to be able to interpret **policies** correctly, they too must share the
1500 appropriate **vocabulary**.

1501 A shared **vocabulary** may range from a simple understanding of particular strings as
1502 commands to a sophisticated collection of terms that are formalized in shared
1503 ontologies.

1504 **Domain**

1505 A **domain** is a ‘world’ that is used as the basis for the truth of a **proposition**.

1506 When we say ‘world’, we are not restricted to the physical world. The criterion is an
1507 ability to discover facts about it. In our case governmental, commercial and **social**
1508 **structures** that form the backdrop for SOA-based systems are important examples of
1509 modeled worlds.

1510 **Interpretation**

1511 An **interpretation** is a mapping from an element of a **vocabulary** to an element
1512 of a **domain**.

1513 An **interpretation** is a way that elements of the **vocabulary** are understood in terms of
1514 the **domain**.

1515 **Semantics**

1516 The **semantics** of a **proposition** is the set of permissible **interpretations** of the
1517 **proposition**.

1518 In effect, **semantics** is rooted in the world – the meaning of an utterance depends on its
1519 relationship with the world.

1520 **Decision Procedure**

1521 A **decision procedure** is a process for determining whether a **proposition** is
1522 true, or is satisfied, in a **domain**.

1523 Decision procedures are algorithms, programs that can measure the world against a
1524 **proposition**'s **written expression** and answer the question whether the **domain**
1525 corresponds to the description.

1526 If the truth of a **proposition** is indeterminable, then a **decision procedure** is not
1527 complete, and the logic is un-decidable.

1528 Note that not all 'systems of marks' have a **decision procedure**. However, for the uses
1529 to which we put the concepts of **fact**, **policies**, **service descriptions**, and so on, we
1530 require that the language used to write **propositions** MUST have a **decision**
1531 **procedure**.

1532 Each system of *logic* has at least one **decision procedure** – by definition. Much of the
1533 art in designing a system of expressions and semantics is arranging for there to be a
1534 **decision procedure** and to ensure that there is at least one tractable **decision**
1535 **procedure**. This issue is especially important in designing **policy frameworks**.

1536 A critically important characteristic of a **proposition** is its meaning to the **actors** and
1537 **stakeholders** in the SOA ecosystem. What a **proposition** means to an **actor** depends
1538 on that **actor's stance** to the **proposition**.

1539 **Stance**

1540 **Stance** is the relationship that an **actor** (or group of **actors**) has to a
1541 **proposition**.

1542 The primary kinds of **stance** that are possible reflect the primary ways that an **actor** can
1543 relate to a **proposition**: the **proposition** may be something that the **actor** knows (or
1544 believes), the **actor** may desire that the **proposition** is satisfied, the **actor** may be
1545 actively engaged in satisfying the **proposition** with some planned **action**; or the **actor**
1546 may view the **proposition** as a **policy** that is to be enforced.

1547 In this section we highlight some of the basic examples of **stance**, leaving others in
1548 sections specially devoted to their exposition.

1549 **Fact**

1550 A **fact** is a **proposition** that can be known by an **actor**.

1551 It is, in principle, impossible for an external observer to determine what a given **actor**
1552 knows. However, when one **actor** 'tells' another **actor** something, then the first **actor**
1553 can infer that the second **actor** now also knows it.

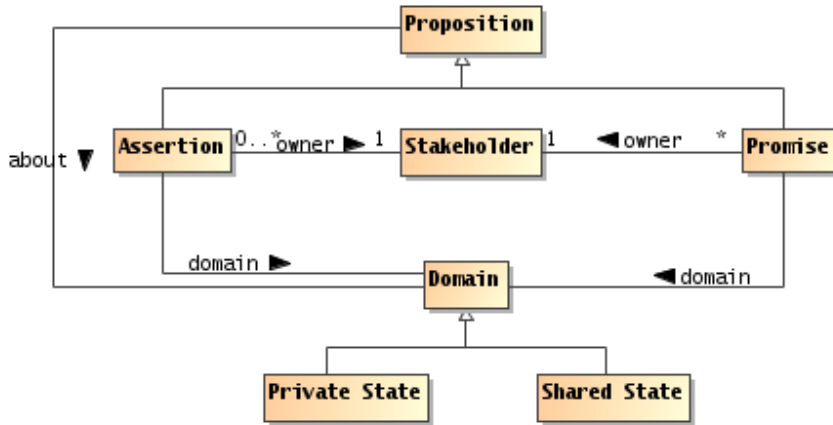


Figure 22 Assertions and Promises

Promise

A **promise** is a **proposition** regarding the future state of the world by a **stakeholder**.

In particular, a **promise** represents a **commitment** by the **stakeholder** to ensure that the **proposition** will be true at some point.

For example, an airline may report its record in on-time departures for its various flights. This is a claim made by the airline which is, in principle, verifiable. The same airline may promise that some percentage of its flights depart within 5 minutes of their scheduled departure. The truth of this promise depends on the effectiveness of the airline in meeting its **commitments**.

Other forms of assertion that are used in this Reference Architecture include **goals** which is something that an **actor** is seeking to establish or maintain; **objectives** which are **propositions** that **actors** seek to establish or maintain by means of specific **actions**; **purposes** which are externally asserted about entities; and **policies**.

Purpose

A **purpose** is a measurable **condition** ascribed to a thing or an **action** relating it to a **goal**.

By their nature, **purposes** are *external* to the purposed entities, whereas goals are *internal* to the entity.

3.3.2 Semantic Engagement

Any utterance can only be fully understood in terms of specific contexts; contexts that necessarily include the **actors** that are involved. For example, a **policy** statement that governs the **actions** relating to a particular **resource** has a different force for the **participant** that **owns** the **resource** to that for the **actor** that is trying to access it: the former interprets the **policy** as something that it is trying to enforce and the latter interprets the **policy** as something that constrains it.

1582 In addition, the ability of an **actor** to comprehend **assertions** is also very variable and
1583 context sensitive. For any given **actor**, for any given **action** – whether the **action** is
1584 private or is a **joint action** – the engagement of the **actor** can be informally understood
1585 as the **actor's** understanding of what it is doing.¹⁴

1586 Even within a single **joint action** the different **actors** involved may have very different
1587 means of understanding the activity. For example, an **actor** requesting a particular
1588 record from a database may understand the request in terms of accessing customer
1589 data. The database **actor** (assuming that the database has been personalized)
1590 interprets the same request as at best an SQL script to execute.

1591 **Semantic Engagement**

1592 **Semantic engagement** is the relationship between an **actor** and the possible
1593 **interpretations** of an **assertion**.

1594 Different **actors** will have differing capabilities and requirements for understanding
1595 **assertions**. This is true for human **actors**; but is especially true for automated **actors**.
1596 For example, a 'message forwarding agent' only needs to be able to understand
1597 communication sufficiently to determine how to forward the communication. On the
1598 other hand, a order processing **actor** needs to be able to determine if a communication
1599 is a purchase order in the correct form, whether there is sufficient stock, whether the
1600 customer has appropriate credit and so on.

1601 **3.3.3 Private and Public semantics**

1602 A SOA ecosystem can be viewed as a space in which **actors** share understanding as
1603 well as sharing **actions**. As such, we need to be able to distinguish the shared meaning
1604 of utterances from the full or private meanings of those utterances. An important
1605 distinction here is that of public versus private semantics:

1606 **Public Semantics**

1607 The **public semantics** of an **assertion** is that subset of the possible
1608 **interpretations** of the **assertion** that is available to any observer by virtue of the
1609 observer's situation in a **social structure**.

1610 The **public semantics** of an utterance is effectively the 'ecosystem view' of the
1611 utterance.

1612 Of course, the most obvious observer of a communication is the intended recipient of
1613 the communication. However, in general, the **public semantics** of a communication
1614 would enable *any* observer to make the same inferences.

1615 For example, a standard purchase order denotes a **commitment** to buy some goods or
1616 services. Any observer of such a standard purchase order would be entitled to interpret
1617 it as a purchase order (whether or not the purchase order was targeted at the observer).

¹⁴ We adopt a Turing test-based approach to understanding: if an actor behaves as though it understands an utterance then we assume that it does understand it.

1618 3.4 Architectural Implications

1619 3.4.1.1 Social structures

1620 A SOA ecosystem's **participants** are organized into various forms of **social structure**.
1621 Not all **social structures** are hierarchical: a SOA ecosystem should be able to
1622 incorporate peer-to-peer forms of organization as well as hierarchic structures. In
1623 addition, it should be possible to identify and manage any constitutional agreements
1624 that define the **social structures** present in a SOA ecosystem.

- 1625 • Different **social structures** have different rules of engagement
 - 1626 ○ Techniques for expressing constitutions are important
- 1627 • **social structures** have **roles** and members
 - 1628 ○ Techniques for identifying, managing members of **social structures**
 - 1629 ○ Techniques for describing **roles** and **role** adoption
- 1630 • **social structures** may be complex
 - 1631 ○ Child **social structures**' constitutions depend on their parent constitutions
- 1632 • Social structures overlap and interact
 - 1633 ○ A given **actor** may be member of multiple **social structures**
 - 1634 ○ Social structures may be associated with different jurisdictions
 - 1635 ○ Social structures may involved in disputes with one another
 - 1636 ■ Requiring conflict resolution

1637 3.4.1.2 The Importance of Action

1638 **participants** participate in a SOA ecosystem in order to get their needs met. This
1639 involves **action**; both individual actions and **joint actions**.

1640 Any architectural realization of a SOA ecosystem should address:

- 1641 • How actions are modeled:
 - 1642 ○ Identifying the performer or agent of the **action**;
 - 1643 ○ the target of the **action**; and the
 - 1644 ○ verb of the **action**.

1645 Joint actions are actions involving multiple **actors**. Any explicit models of **joint action**
1646 should take into account

- 1647 • The choreography that defines the **joint action**.
- 1648 • The potential for **joint actions** to be multiply layered on top of each other

1649 3.4.1.3 Communications as a Means of Mediating Action

1650 Using message exchange for mediating **action** implies

- 1651 • Ensuring correct identification of the structure of messages:
 - 1652 ○ Identifying the syntax of the message;
 - 1653 ○ Identifying the vocabularies used in the communication
 - 1654 ○ Identifying the higher-level structure such as the illocutionary form of the
1655 communication
- 1656 • A principal objective of communication is to mediate **action**
 - 1657 ○ Messages convey actions and **events**
 - 1658 ○ Receiving a message is an **action**, but is not the same **action** as the
1659 action conveyed by the message

- 1660 ○ Actions are associated with objectives of the **actors** involved
1661 ▪ Explicit representation of objectives may facilitate automated
1662 processing of messages
1663 ○ An **actor** agreeing to adopt an objective becomes responsible for that
1664 objective

1665 **3.4.1.4 Semantics**

1666 Semantics is pervasive in a SOA ecosystem. There are many forms of utterance that
1667 are relevant to the ecosystem: apart from communicated content there are policy
1668 statements, goals, purposes, descriptions, and agreements which are all forms of
1669 utterance.

1670 The operation of the SOA ecosystem is significantly enhanced if

- 1671 • A careful distinction is made between **public semantics** and private semantics.
1672 In particular, it MUST be possible for **actors** to process content such as
1673 communications, descriptions and policies solely on the basis of the **public**
1674 **semantics** of those utterances.
- 1675 • A well founded semantics ensures that any **assertions** that are essential to the
1676 operator of the ecosystem (such as **policy statements**, and **descriptions**) have
1677 carefully chosen **written expressions** and associated **decision procedures**.
- 1678 • The role of **vocabularies** as a focal point for multiple **actors** to be able to
1679 understand each other is critical. While no two **actors** can fully share their
1680 interpretation of elements of **vocabularies**, ensuring that they do understand the
1681 **public** meaning of **vocabularies'** elements is essential.

1682 **3.4.1.5 Trust and Risk**

1683 In traditional systems, the balance between **trust** and **risk** is achieved by severely
1684 restricting interactions and by controlling the **participants** of a system.

1685 It is important that **actors** are able to explicitly reason about both **trust** and **risk** in order
1686 to effectively participate in a SOA ecosystem. The more open and public the SOA
1687 ecosystem is, the more important it is for **actors** to be able to reason about their
1688 participation.

1689 **3.4.1.6 Policies and Contracts**

- 1690 • Policies are constraints
 - 1691 ○ It is necessary to be able to express required policies
 - 1692 ○ It is necessary to be able to enforce the constraints
 - 1693 ○ It is necessary to manage potentially large numbers of policies
- 1694 • Policies have owners
 - 1695 ○ The right to establish policies is an aspect of the social structure.
- 1696 • Policies may not be consistent with one another
 - 1697 ○ **Policy** conflict resolution techniques
- 1698 • Agreements are constraints agreed to
 - 1699 ○ Contracts often need to be enforced by mechanisms of the social structure

1700

4 Realizing Service Oriented Architectures View

Make everything as simple as possible but no simpler.

Albert Einstein

The *Realizing Service Oriented Architectures View* focuses on the infrastructure elements that are needed in order to support the discovery and interaction with services. The key questions asked are "What are services, what support is needed and how are they realized?"

The models in this view include the Service Description Model, the Service Visibility Model, the Interacting with Services Model, and the Policies and Contracts Model.

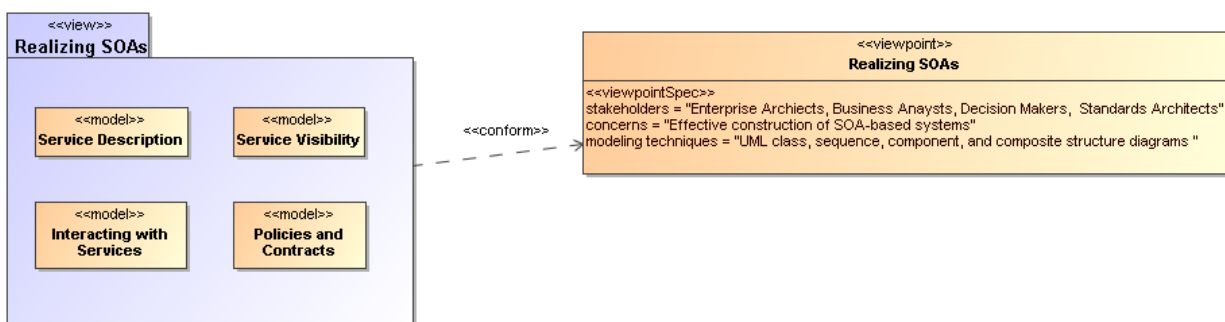


Figure 23 Model Elements Described in the Realizing Service Oriented Architectures View

The Service Description Model informs the **participants** of what services exist and the conditions under which these can be used. Some of those conditions follow from policies and agreements on policy that flow from the Policies and Contracts Model. The information in the service description as augmented by details of **policy** provides the basis for visibility as defined in the SOA Reference Model and captured in the Service Visibility Model. Finally, the process by which services as described are used under the defined conditions and agreements is described in the Interacting with Services Model.

4.1 Service Description Model

A service description is an artifact, usually document-based, that defines or references the information needed to use, deploy, manage and otherwise control a service. This includes not only the information and behavior models associated with a service to define the service interface but also includes information needed to decide whether the service is appropriate for the current needs of the **service consumer**. Thus, the service description will also include information such as service reachability, service functionality, and the policies and contracts associated with a service.

A service description artifact may be a single document or it may be an interlinked set of documents. For the purposes of this model, differences in representation are to be ignored, but the implications of a "web of documents" is discussed later in this section.

There are several points to note regarding the following discussion of service description:

- 1733 • The Reference Model states that one of the hallmarks of SOA is the large amount of
1734 associated description. The model presented below focuses on the description of
1735 services but it is equally important to consider the descriptions of the consumer,
1736 other **participants**, and needed **resources** other than services.
- 1737 • Descriptions are inherently incomplete but may be determined as *sufficient* when it is
1738 possible for the **participants** to access and use the described services based only
1739 on the descriptions provided. This means that, at one end of the spectrum, a
1740 description along the lines of “*That service on that machine*” may be sufficient for the
1741 intended audience. On the other extreme, a service description with a machine-
1742 process-able description of the semantics of its operations and **real world effects**
1743 may be required for services accessed via automated service discovery and
1744 planning systems.
- 1745 • Descriptions come with context, i.e. a given description comprises information
1746 needed to adequately support the context. For example, a list of items can define a
1747 version of a service, but for many contexts an indicated version number is sufficient
1748 without the detailed list. The current model focuses on the description needed by a
1749 **service consumer** to understand what the service does, under what conditions will
1750 the service do it, how well does the service do it, and what steps are needed by the
1751 consumer to initiate and complete a service interaction. Such information also
1752 enables the service provider to clearly specify what is being provided and the
1753 intended conditions of use.
- 1754 • Descriptions will change over time as, for example, the ingredients and nutrition
1755 information for food labeling continues to evolve. A requirement for transparency of
1756 transactions may require additional description for those associated contexts.
- 1757 • Description always proceeds from a basis of what is considered “common
1758 knowledge”. This may be social conventions that are commonly expected or possibly
1759 codified in law. It is impossible to describe everything and it can be expected that a
1760 mechanism as far reaching as SOA will also connect entities where there is
1761 inconsistent “common” knowledge.
- 1762 • Descriptions will become the collection point of information related to a service or
1763 any other **resource**, but it will not necessarily be the originating point or the
1764 motivation for generating this information. In particular, given a SOA service as the
1765 access to an underlying capability, the service may point to some of the capability’s
1766 previously generated description, e.g. a service providing access to a data store may
1767 reference update records that indicate the freshness of the data.
- 1768 • Descriptions of the provider and consumer are the essential building blocks for
1769 establishing the execution context of an interaction.

1770 These points emphasize that there is no one “right” description for all contexts and for
1771 all time. Several descriptions for the same subject may exist at the same time, and this
1772 emphasizes the importance of the description referencing source material maintained
1773 by that material’s owner rather than having multiple copies that become out of synch
1774 and inconsistent.

1775 It may also prove useful for a description assembled for one context to cross-reference
1776 description assembled for another context as a way of referencing ancillary information

1777 without overburdening any single description. Rather than a single artifact, description
1778 can be thought of as a web of documents that enhance the total available description.

1779 This Reference Architecture uses the term service description for consistency with the
1780 concept defined in the Reference Model. Some SOA literature treats the idea of a
1781 “service contract” as equivalent to service description. In this Reference Architecture,
1782 the term service description is preferred. Replacing service description with service
1783 contract implies just one side of the interaction is governing and misses the point that a
1784 single set of policies identified by a service description may lead to numerous contracts,
1785 i.e. service level agreements, leveraging the same description.

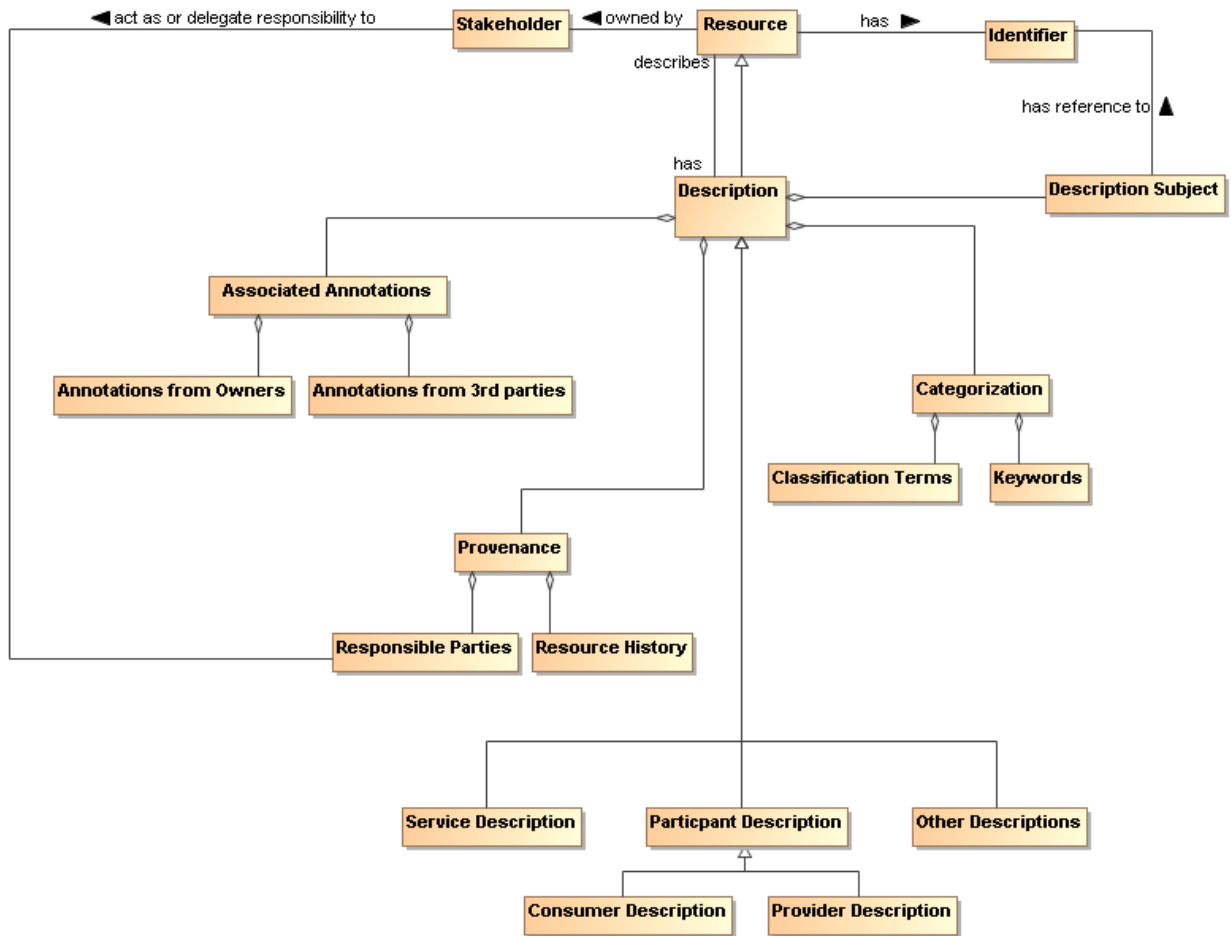
1786 **4.1.1 The Model for Service Description**

1787 *Figure 24* shows Service Description as a subclass of the general Description class,
1788 where Description is a subclass of the **resource** class as defined in Section **Error!**
1789 **Reference source not found.** In addition, each **resource** is assumed to have a
1790 description. The following section discusses the relationships among elements of
1791 general description and the subsequent sections focus on service description itself.
1792 Note, other descriptions, such as those of **participants**, are important to SOA but are
1793 not individually elaborated in this document.

1794

1795 **4.1.1.1 Elements Common to General Description**

1796 The general Description class is composed of a number of elements that are expected
1797 to be common among all specialized descriptions supporting a service oriented
1798 architecture. A registry often contains a subset of the description instance, where the
1799 chosen subset is identified as that which facilitates mediated discovery. Additional
1800 information contained in a more complete description may be needed to initiate and
1801 continue interaction.



1802
1803 *Figure 24 General Description*

1804 **4.1.1.1.1 Description Subject**

1805 The subject of a description is a **resource**. The value assigned to the Description
1806 Subject class may be of any form that provides understanding of what constitutes the
1807 **resource**, but it is often in human-readable text. The Description Subject **MUST** also
1808 reference the Identifier of the **resource** it describes so it can unambiguously identify the
1809 subject of each description instance.

1810 As a **resource**, Description also has an identifier with a unique value for each
1811 description instance. The description instance provides vital information needed to both
1812 establish visibility of the **resource** and to support its use in the execution context for the
1813 associated interaction. The identifier of the description instance allows the description
1814 itself to be referenced for discussion, access, or reuse of its content.

1815

1816 **4.1.1.1.2 Provenance**

1817 While the **resource** Identifier provides the means to know which subject and subject
1818 description are being considered, Provenance as related to the Description class
1819 provides information that reflects on the quality or usability of the subject. Provenance
1820 specifically identifies the entity (human, defined **role**, organization, ...) that assumes

1821 responsibility for the **resource** being described and tracks historic information that
1822 establishes a context for understanding what the **resource** provides and how it has
1823 changed over time. **Responsibilities** may be directly assumed by the **stakeholder** who
1824 owns a **resource** or the Owner may designate Responsible Parties for the various
1825 aspects of maintaining the **resource** and provisioning it for use by others. There may be
1826 more than one entity identified under Responsible Parties; for example, one entity may
1827 be responsible for code maintenance while another is responsible for provisioning of the
1828 executable code. The historical aspects may also have multiple entries, such as when
1829 and how data was collected and when and how it was subsequently processed, and as
1830 with other elements of description, may provide links to other assets maintained by the
1831 **resource** owner.

1832 **4.1.1.1.3 Keywords and Classification Terms**

1833 A traditional element of description has been to associate the **resource** being described
1834 with predefined keywords or classification taxonomies that derive from referenceable
1835 formal definitions and vocabularies. This Reference Architecture does not prescribe
1836 which vocabularies or taxonomies may be referenced, nor does it limit the number of
1837 keywords or classifications that may be associated with the **resource**. It does,
1838 however, state that a normative definition SHOULD be referenced, whether that be a
1839 representation in a formal ontology language, a pointer to an online dictionary, or any
1840 other accessible source. See Section 4.1.1.2 for further discussion on associating
1841 semantics with assigned values.

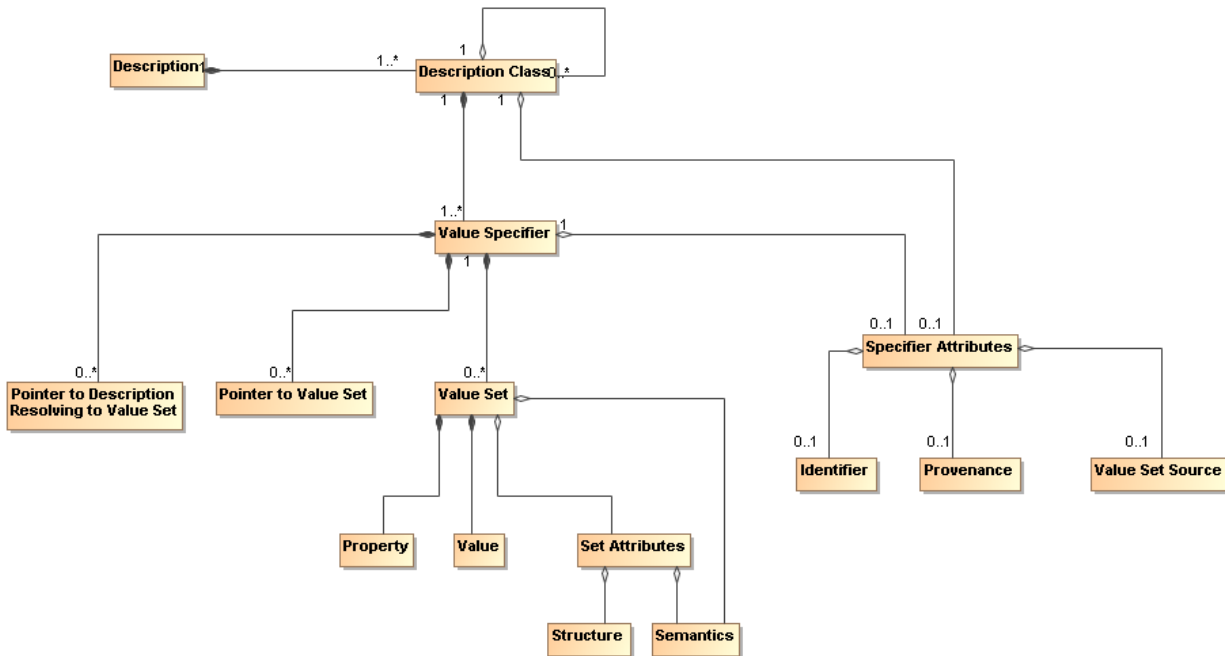
1842 **4.1.1.1.4 Associated Annotations**

1843 The general description instance may also reference associated documentation that is
1844 in addition to that considered necessary in this model. For example, the owner of a
1845 service may have documentation on best practices for using the service. Alternately, a
1846 third party may certify a service based on their own criteria and certification process;
1847 this may be vital information to other prospective consumers if they were willing to
1848 accept the certification in lieu of having to perform another certification themselves.
1849 Note, while the examples of Associated Documentation presented here are related to
1850 services, the concept applies equally to description of other entities.

1851

1852 **4.1.1.2 Assigning Values to Description Instances**

1853



1854

1855 *Figure 25 Representation of a Description*

1856 Figure 24 shows the template for a general description but individual description
1857 instances depend on the ability to associate meaningful values with the identified
1858 elements. Figure 25 shows a model for a collection of information that provides for value
1859 assignment and traceability for both the value meaning and the source of a value. The
1860 model is not meant to replace existing or future schema or other structures that have or
1861 will be defined for specific implementations, but it is meant as guidance for the
1862 information such structures need to capture to generate sufficient description. It is
1863 expected that tools will be developed to assist the user in populating description and
1864 auto-filling many of these fields, and in that context, this model provides guidance to the
1865 tool developers.

1866 In Figure 25 each class has an associated value specifier or is made up of components
1867 that will eventually resolve to a value specifier. For example, Description has several
1868 components, one of which is Categorization, which would have an associated a value
1869 specifier.

1870 A value specifier consists of

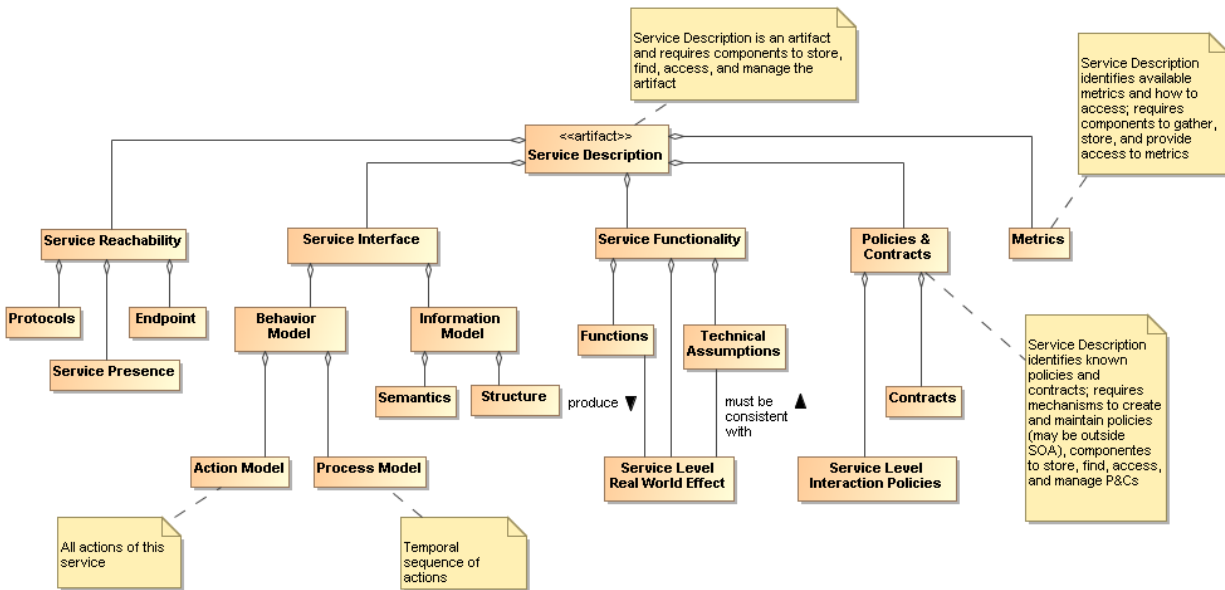
- 1871 • a collection of value sets with associated property-value pairs, pointers to such value
1872 sets, or pointers to descriptions that eventually resolve to value sets that describe
1873 the component; and
- 1874 • attributes that qualify the value specifier and the value sets it contains.

1875 The qualifying attributes for the value specifier include

- 1876 • an optional identifier that would allow the value set to be defined, accessed, and
1877 reused elsewhere;

- 1878 • provenance information that identifies the party (individual, **role**, or organization) that
1879 has responsibility for assigning the value sets to any description component;
- 1880 • an optional source of the value set, if appropriate and meaningful, e.g. if a particular
1881 data source is mandated.
- 1882 If the value specifier is contained within a higher-level component, (such as Service
1883 Description containing Service Functionality), the component may inherit values from
1884 the attributes from its container.
- 1885 Note, provenance as a qualifying attribute of a value specifier is different from
1886 provenance as part of an instance of Description. Provenance for a service identifies
1887 those who own and are responsible for the service, as described in Section **Error!**
1888 **Reference source not found..** Provenance for a value specifier identifies who is
1889 responsible for choosing and assigning values to the value sets that comprise the value
1890 specifier. It is assumed that granularity at the value specifier level is sufficient and
1891 provenance is not required for each value set.
- 1892 The value set also has attributes that define its structure and semantics.
- 1893 • The semantics of the value set property should be associated with a semantic
1894 context conveying the meaning of the property within the execution context, where
1895 the semantic context could vary from a free text definition to a formal ontology.
- 1896 • For numeric values, the structure would provide the numeric format of the value and
1897 the “semantics” would be conveyed by a dimensional unit with an identifier to an
1898 authoritative source defining the dimensional unit and preferred mechanisms for its
1899 conversion to other dimensional units of like type.
- 1900 • For nonnumeric values, the structure would provide the data structure for the value
1901 representation and the semantics would be an associated semantic model.
- 1902 • For pointers, architectural guidelines would define the preferred addressing scheme.
- 1903 The value specifier may indicate a default semantic model for its component value sets
1904 and the individual value sets may provide an override.
- 1905 The property-value pair construct is introduced for the value set to emphasize the need
1906 to identify unambiguously both what is being specified and what is a consistent
1907 associated value. The further qualifying of Structure and Semantics in the Set
1908 Attributes allows for flexibility in defining the form of the associated values.

1909 **4.1.1.3 Model Elements Specific to Service Description**



1910
1911 *Figure 26 Service Description*

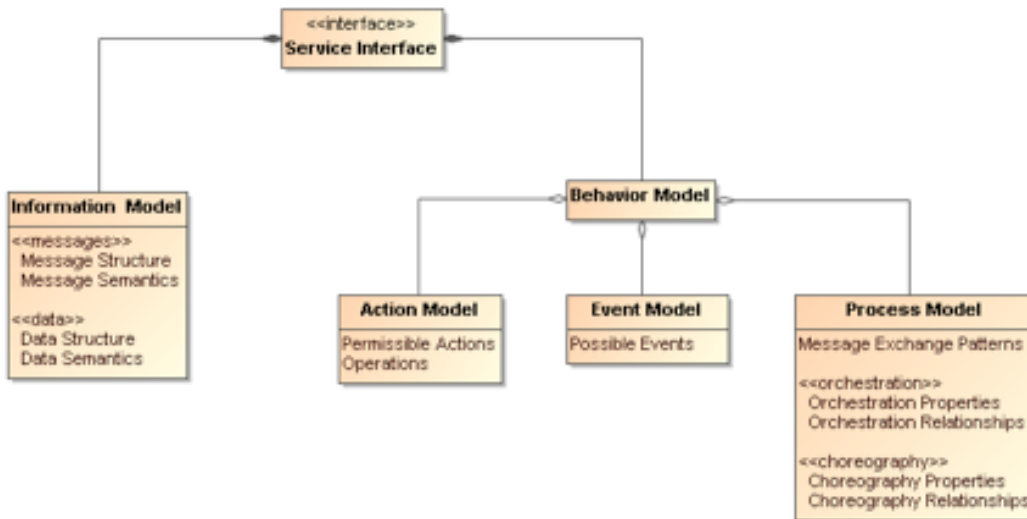
1912 The major elements for the Service Description subclass follow directly from the areas
1913 discussed in the Reference Model. Here, we discuss the detail shown in *Figure 26* and
1914 the purpose served by each element of service description.

1915 Note, the intent in the subsections that follow is to describe how a particular element,
1916 such as the service interface, is reflected in the service description, not to elaborate on
1917 the details of that element. Other sections of the Reference Model and this Reference
1918 Architecture describe the “physics” of each element whereas the service description
1919 subsections will only touch on the meta aspects.

1920 **4.1.1.3.1 Service Interface**

1921 As noted in the Reference Model, the service interface is the means for interacting with
1922 a service. For this Reference Architecture and as shown in Section 4.3 the service
1923 interface will support an exchange of messages, where

- 1924
- 1925 • the message conforms to a referenceable message exchange pattern (MEP),
 - 1926 • the message payload conforms to the structure and semantics of the indicated information model,
 - 1927 • the messages are used to denote **events** or actions against the service, where
 - 1928 the actions are specified in the **action** model and any required sequencing of
 - 1929 actions is specified in the process model.



1930
1931 *Figure 27 Service Interface*

1932 Note we distinguish the structure and semantics of the message from that of the
1933 underlying protocol that conveys the message. The message structure may include
1934 nested structures that are independently defined, such as an enclosing envelope
1935 structure and an enclosed data structure.

1936 These aspects of messages are discussed in more detail in Section 4.3

1937 **4.1.1.3.2 Service Reachability**

1938 Service reachability, as modeled in Section 4.2.2.3 enables service **participants** to
1939 locate and interact with one another. To support service reachability, the service
1940 description should indicate the endpoints to which a **service consumer** can direct
1941 messages to invoke actions and the protocol to be used for message exchange using
1942 that endpoint.

1943 As applied in general to an **action**, the endpoint is the conceptual location where one
1944 applies an **action**; with respect to service description, it is the actual address where a
1945 message is sent.

1946 In addition, the service description should provide information on collected metrics for
1947 service presence; see Section 4.1.1.3.4 for the discussion of metrics as part of service
1948 description.

1949 **4.1.1.3.3 Service Functionality**

1950 While the service interface and service reachability are concerned with the mechanics
1951 of using a service, service functionality and performance metrics (discussed in Section
1952 4.1.1.3.4) describe what can be expected when interacting with a service. Service
1953 Functionality, shown in *Figure 26* as part of the overall Service Description model and
1954 extended in *Figure 28*, is an unambiguous expression of service function(s) and the **real**
1955 **world effects** of invoking the function. The Functions likely represent business activities
1956 in some domain that produce the desired **real world effects**.

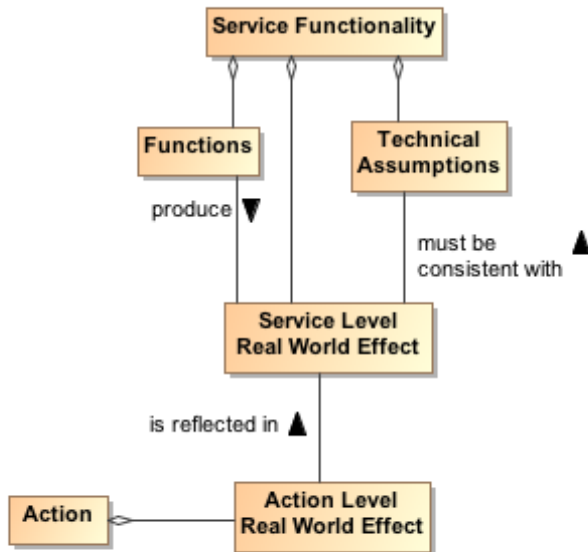


Figure 28 Service Functionality

1957
1958

The Service Functionality may also be constrained by Technical Assumptions that underlie the effects that can result. Technical assumptions are defined as domain specific restrictions and may express underlying physical limitations, such as flow speeds must be below sonic velocity or disk access that cannot be faster than the maximum for its host drive. Technical assumptions are likely related to the underlying capability accessed by the service. In any case, the **real world effects** must be consistent with the Technical Assumptions.

In *Figure 26* and *Figure 28*, we specifically refer to Service Level and Action Level **real world effects**.

1968 Service Level Real World Effect

1969 A service level **real world effect** is a specific change in shared state or
1970 information returned as a result of interacting with a service.

1971 Action Level Real World Effect

1972 An action level **real world effect** is a specific change in shared state or
1973 information returned as a result of performing a specific **action** against a service.

1974 Service description describes the service as a whole while the component aspects
1975 should contribute to that whole. Thus, while individual Actions may contribute to the
1976 **real world effects** to be realized from interaction with the service, there would be a
1977 serious disconnect for Actions to contribute **real world effects** that could not
1978 consistently be reflected in the Service Level Real World Effects and thus the Service
1979 Functionality. The relationship to Action Level Real World Effects and the implications
1980 on defining the scope of a service are discussed in Section 4.1.2.1.

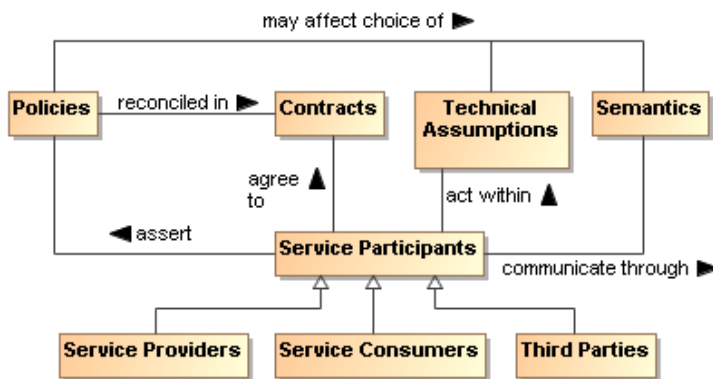
1981 Elements of Service Functionality may be expressed as natural language text, reference
1982 to an existing taxonomy of functions, or reference to a more formal knowledge capture
1983 providing richer description and context.

1984 **4.1.1.3.4 Policies and Contracts, Metrics, and Compliance Records**

1985 Policies prescribe the conditions and constraints for interacting with a service and
1986 impact the willingness to continue visibility with the other **participants**. Whereas
1987 technical assumptions are statements of “physical” fact, policies are subjective
1988 assertions made by the service provider (sometimes as passed on from higher
1989 authorities).

1990 The service description provides a central location for identifying what policies have
1991 been asserted by the service provider. The specific representation of the **policy**, e.g. in
1992 some formal **policy** language, is likely done outside of the service description and the
1993 service description would reference the normative definition of the **policy**.

1994 Policies may also be asserted by other service **participants**, as illustrated by the model
1995 shown in Figure 29. Policies that are generally applicable to any interaction with the
1996 service are likely to be asserted by the service provider and included in the Policies and
1997 Contracts section of the service description. Conversely, policies that are asserted by
1998 specific consumers or consumer communities would likely be identified as part of a
1999 description’s Annotations from 3rd parties (see Section 4.1.1.1.4) because these would
2000 be specific to those parties and not a general aspect of the service being described.



2001
2002 *Figure 29 Model for Policies and Contracts as related to Service Participants*

2003 In *Figure 26* and *Figure 30*, we specifically refer to Service Level Interaction Policies. In a
2004 similar manner to that discussed for Service Level vs. Action Level Real World Effects in
2005 Section 4.1.1.3.3, individual Actions may have associated policies stating conditions for
2006 performing the **action**, but these must be reflected in and be consistent with the policies
2007 made visible at the service level and thus the description of the service as a whole. The
2008 relationship to Action Level Policies and the implications on defining the scope of a
2009 service are discussed in Section 4.1.2.1.




Figure 30 Action-Level and Service-Level Policies

2010
2011

2012 As noted in Figure 29, the policies asserted may affect the allowable Technical
2013 Assumptions that can be embodied in services or their underlying capabilities and may
2014 affect the semantics that can be used. For example of the former, there may be a
2015 **policy** that specifies the surge capacity to be accommodated by a server, and a service
2016 that designs for a smaller capacity would not be appropriate to use. For the latter, a
2017 **policy** may require that only services using a community-sponsored vocabulary can be
2018 used.

2019 Contracts are agreements among the service **participants**. The contract may reconcile
2020 inconsistent policies asserted by the **participants** or may specify details of the
2021 interaction. Service level agreements (SLAs) are one commonly used category of
2022 contracts.

2023 References to contracts under which the service can be used may also be included in
2024 the service description. As with policies, the specific representation of the contract, e.g.
2025 in some formal contract language, is likely done outside of the service description and
2026 the service description would reference the normative definition of the contract. Policies
2027 and contracts are discussed further in Section 4.4.

2028 The definition and later enforcement of policies and contracts are predicated on the
2029 existence of metrics; the relationships among the relevant concepts are shown in the
2030 model in Figure 31. Performance Metrics identify quantities that characterize the speed
2031 and quality of realizing the **real world effects** produced via the SOA service; in
2032 addition, policies and contracts may depend on nonperformance metrics, such as
2033 whether a license is in place to use the service. Some of these metrics reflect the
2034 underlying capability, e.g. a SOA service cannot respond in two seconds if the
2035 underlying capability is expected to take five seconds to do its processing; some
2036 metrics reflect the implementation of the SOA service, e.g. what level of caching is
2037 present to minimize data access requests across the network.

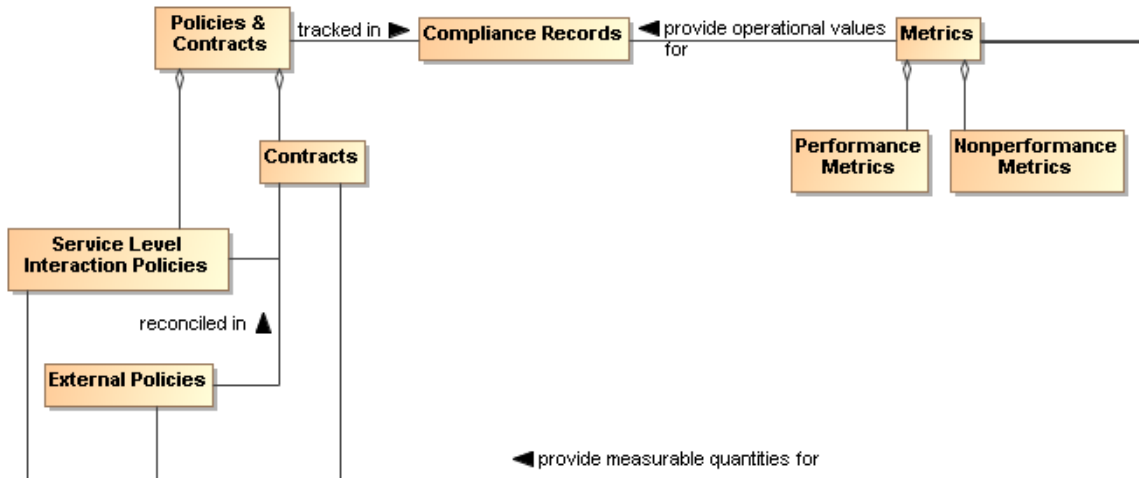


Figure 31 Policies and Contracts, Metrics, and Compliance Records

2038
2039

2040 As with many quantities, the metrics associated with a service are not themselves
2041 defined by this Service Description because it is not known *a priori* which metrics are
2042 being collected or otherwise checked by the services, the SOA infrastructure, or other
2043 resources that participate in the SOA interactions. However, the service description
2044 SHOULD provide a placeholder (possibly through a link to an externally compiled list)
2045 for identifying which metrics are available and how these can be accessed.

2046 The use of metrics to evaluate compliance is discussed in Section **Error! Reference**
2047 **source not found..** The results of compliance evaluation SHOULD be maintained in
2048 compliance records and the means to access the compliance records SHOULD be
2049 included in the Policies and Contracts portion of the service description. For example,
2050 the description may be in the form of static information (e.g. over the first year of
2051 operation, this service had a 91% availability), a link to a dynamically generated metric
2052 (e.g. over the past 30 days, the service has had a 93.3% availability), or access to a
2053 dynamic means to check the service for current availability (e.g. a ping). The
2054 relationship between service presence and the presence of the individual actions that
2055 can be invoked is discussed under Reachability in Section 4.2.2.3.

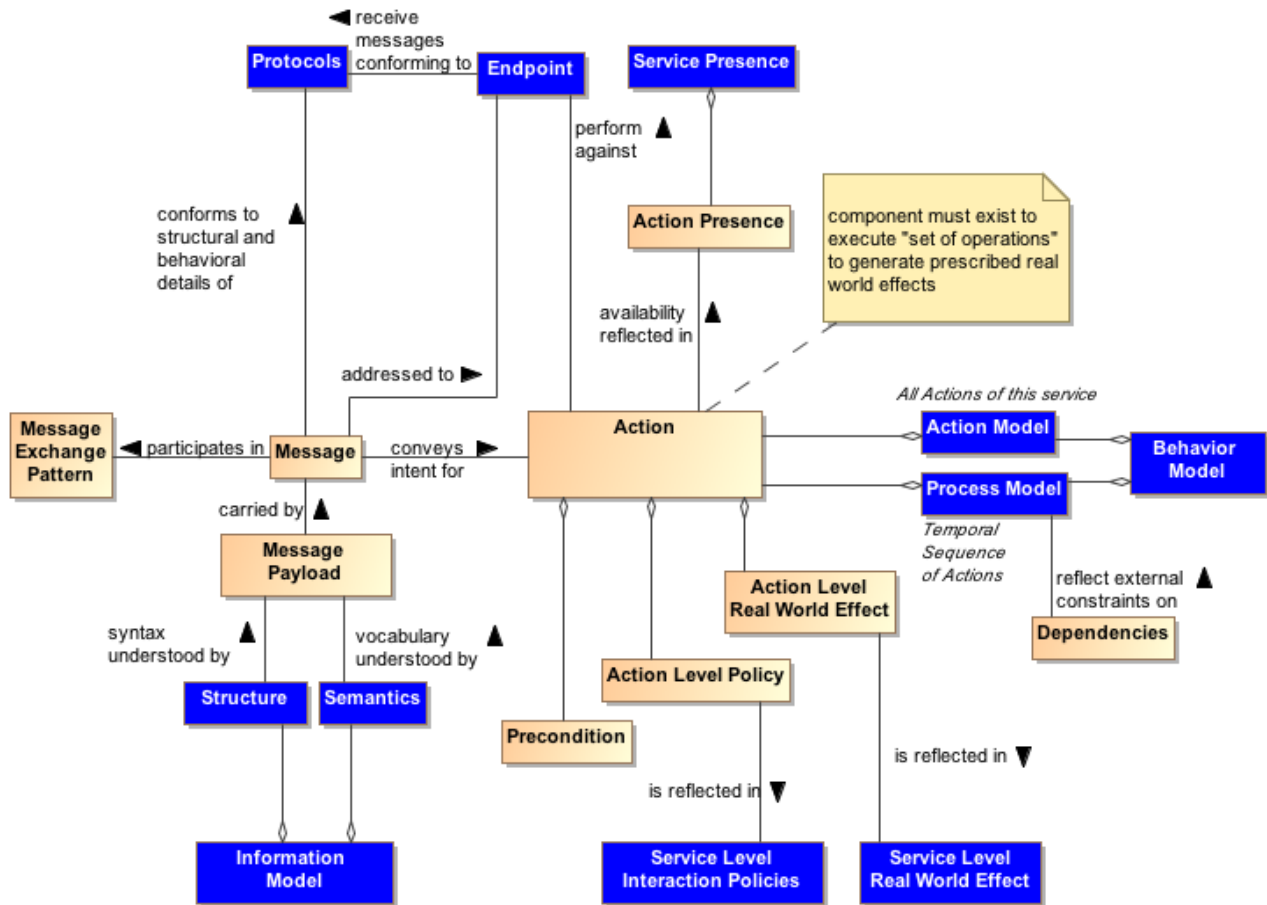
2056 Note, even when policies relate the perspective of a single **participant, policy**
2057 compliance can be measured and policies may be enforceable without contractual
2058 agreement with other **participants**. This should be reflected in the **policy**, contract,
2059 and compliance record information maintained in the service description.

2060 4.1.2 Use Of Service Description

2061 4.1.2.1 Service Description in support of Service Interaction

2062 If we assume we have awareness, i.e. access to relevant descriptions, the service
2063 **participants** must still establish willingness and presence to ensure full visibility (See
2064 Section 4.2) and to interact with the service. Service description provides necessary
2065 information for many aspects of preparing for and carrying through with interaction.
2066 Recall the fundamental definition of service is a mechanism to access an underlying
2067 capability; the service description describes this mechanism and its use. It lays the

2068 groundwork for what can occur, whereas service interaction defines the specifics
 2069 through which occurrences are realized.



2070
 2071 *Figure 32 Relationship Between Action and Service Description Components*

2072 Figure 32 combines the models in the subsections of Section 4.1.1 to concisely relate
 2073 **action** and the relevant components of Service Description. The purpose of Figure 32 is
 2074 to demonstrate that the components of service description go beyond arbitrary
 2075 documentation and form the critical set of information needed to define the what and
 2076 how of **action**. In Figure 32, the leaf nodes from Figure 26 are shown in blue.

2077 **action** is invoked via a Message where the structure and behavioral details of the
 2078 message conform to an identified Protocol and is directed to the address of the
 2079 identified endpoint, and the message payload conforms to the service Information
 2080 Model.

2081 The availability of an **action** is reflected in the Action Presence and each Action
 2082 Presence contributes to the overall Service Presence; this is discussed further in
 2083 Section 4.2.2.3. Each **action** has its own endpoint and also its own protocols associated
 2084 with the endpoint¹⁵ and to what extent, e.g. current or average availability, there is

¹⁵ This is analogous to a WSDL 2.0 interface operation (WSDL 1.1 portType) having one or more defined bindings and the service identifies the endpoints (WSDL 1.1 ports) corresponding to the bindings.

2085 presence for the **action** through that endpoint. The endpoint and service presence are
2086 also part of the service description.

2087 An **action** may have preconditions where a Precondition is something that needs to be
2088 in place before an **action** can occur, e.g. confirmation of a precursor **action**. Whether
2089 preconditions are satisfied is evaluated when someone tries to perform the **action** and
2090 not before. Presence for an **action** means someone can initiate it and is independent of
2091 whether the preconditions are satisfied. However, the successful completion of the
2092 **action** may depend on whether its preconditions were satisfied.

2093 Analogous to the relationship between actions and preconditions, the Process Model
2094 may imply Dependencies for succeeding steps in a process, e.g. that a previous step
2095 has successfully completed, or may be isolated to a given step. An example of the
2096 latter would be a dependency that the host server has scheduled maintenance and
2097 access attempts at these times would fail. Dependencies related to the process model
2098 do not affect the presence of a service although these may affect whether the business
2099 function successfully completes.

2100 The conditions under which an **action** can be invoked may depend on policies
2101 associated with the **action**. The Action Level Policies MUST be reflected in the Service
2102 Level Interaction Policies because such policies may be critical to determining whether
2103 the conditions for use of the service are consistent with the policies asserted by the
2104 **service consumer**. The service level interaction policies are included in the service
2105 description.

2106 Similarly, the result of invoking an **action** is one or more **real world effects**, and the
2107 Action Level Real World Effects MUST be reflected in the Service Level Real World
2108 Effect included in the service description. The unambiguous expression of action level
2109 policies and **real world effects** as service counterparts is necessary to adequately
2110 understand what constitutes the service interaction.

2111 An adequate service description MUST provide a consumer with information needed to
2112 determine if the service policies and the (business) functions and service-level real
2113 world effects are of interest and there is nothing in the technical assumptions that
2114 preclude use of the service.

2115 Note at this level, the business functions are not concerned with the action or process
2116 models. These models are detailed separately.

2117 The service description is not intended to be isolated documentation but rather an
2118 integral part of service use. Changes in service description SHOULD immediately be
2119 made known to consumers and potential consumers.

2120 **4.1.2.1.1 Description and Invoking Actions Against a Service**

2121 At this point, let us assume the descriptions were sufficient to establish willingness; see
2122 Section 4.2.2.2. Figure 32 indicates the service endpoint establishes where to actually
2123 carry out the interaction. This is where we start considering the **action** and process
2124 models.

2125 The action model identifies the multiple actions a user can perform against a service
2126 and the user would perform these in the context of the process model as specified or
2127 referenced under the Service Interface portion of Service Description. For a given
2128 business function, there is a corresponding process model, where any process model

2129 may involve multiple actions. From the above discussion of model elements of
2130 description we may conclude (1) actions have reachability information, including
2131 endpoint and presence, (2) presence of service is some aggregation of presence of its
2132 actions, (3) action preconditions and service dependencies do not affect presence
2133 although these may affect successful completion.

2134 Having established visibility, the interaction can proceed. Given a business function, the
2135 consumer knows what will be accomplished (the service functionality), the conditions
2136 under which interaction will proceed (service policies and contracts), and the process
2137 that must be followed (the process model). The remaining question is how does the
2138 description information for structure and semantics enable interaction.

2139 We have established the importance of the process model in identifying relevant actions
2140 and their sequence. Interaction proceeds through messages and thus it is the syntax
2141 and semantics of the messages with which we are here concerned. A common
2142 approach is to define the structure and semantics that can appear as part of a message;
2143 then assemble the pieces into messages; and, associate messages with actions.
2144 Actions make use of structure and semantics as defined in the information model to
2145 describe its legal messages.

2146 The process model identifies actions to be performed against a service and the
2147 sequence for performing the actions. For a given **action**, the Reachability portion of
2148 description indicates the protocol bindings that are available, the endpoint
2149 corresponding to a binding, and whether there is presence at that endpoint. The
2150 interaction with actions is through messages that conform to the structure and
2151 semantics defined in the information model and the message sequence conforming to
2152 the **action**'s identified MEP. The result is some portion of the **real world effect** that will
2153 need to be assessed and/or processed (e.g. if an error exists, that part that covers the
2154 error processing would be invoked).

2155 **4.1.2.1.2 The Question of Multiple Business Functions**

2156 Action level effects and policies **MUST** be reflected at the service level for service
2157 description to support visibility.

2158 It is assumed that a SOA service represents an identifiable business function to which
2159 policies can be applied and from which desired business effects can be obtained. While
2160 contemporary discussions of SOA services and supporting standards do not constrain
2161 what actions or combinations of actions can or should be defined for a service, this
2162 Reference Architecture considers the implications of service description in defining the
2163 range of actions appropriate for an individual SOA service.

2164 Consider the situation if a given SOA service is the container for multiple independent
2165 (but loosely related) business functions. These are not multiple effects from a single
2166 function but multiple functions with potentially different sets of effects for each function.
2167 A service can have multiple actions a user may perform against it, and this does not
2168 change with multiple business functions. As an individual business function corresponds
2169 to a process model, so multiple business functions imply multiple process models. The
2170 same **action** may be used in multiple process models but the aggregated service
2171 presence would be specific to each business function because the components being
2172 aggregated will likely be different between process models. In summary, for a service
2173 with multiple business functions, each function has (1) its own process model and

2174 dependencies, (2) its own aggregated presence, and (3) possibly its own list of policies
2175 and **real world effect** s.

2176 A common variation on this theme is for a single service to have multiple endpoints for
2177 different levels of quality of service (QoS). Different QoS imply separate statements of
2178 policy, separate endpoints, possibly separate dependencies, and so on. One could say
2179 the QoS variation does not require this because there can be a single QoS policy that
2180 encompasses the variations. and all other aspects of the service would be the same
2181 except for the endpoint used for each QoS. However, the different aspects of policy at
2182 the service level would need to be mapped to endpoints, and this introduces an
2183 undesirable level of coupling across the elements of description. In addition, it is
2184 obvious that description at the service level can become very complicated if the number
2185 of combinations is allowed to grow.

2186 One could imagine a service description that is basically a container for **action**
2187 descriptions, where each action description is self contained; however, this would lead
2188 to duplication of description components across actions. If common description
2189 components are factored, this either is limited to components common across all
2190 actions or requires complicated tagging to capture the components that often but do not
2191 universally apply.

2192 If a provider cannot describe a service as a whole but must describe every **action**, this
2193 leads to the situation where it may be extremely difficult to construct a clear and concise
2194 service description that can effectively support discovery and use without tedious logic
2195 to process the description and assemble the available permutations. In effect, if
2196 adequate description of an **action** begins to look like description of a service, it may be
2197 best to have it as a separate service.

2198 Recall, more than one service can access the same underlying capability, and this is
2199 appropriate if a different **real world effect** is to be exposed. Along these lines, one can
2200 argue that different QoS are different services because getting a response in one
2201 minute rather than one hour is more than a QoS difference; it is a fundamental
2202 difference in the business function being provided.

2203 As a best practice, a criteria for whether a service is appropriately scoped may be the
2204 ease or difficulty in creating an unambiguous service description. A consequence of
2205 having tightly-scoped services is there will be a greater reliance on combining services,
2206 i.e. more fundamental business functions, to create more advanced business functions.
2207 This is consistent with the principles of service oriented architecture and is the basic
2208 position of the Reference Architecture, although not an absolute requirement.
2209 Combining services increases the reliance on understanding and implementing the
2210 concepts of orchestration, choreography, and other approaches yet to be developed;
2211 these are discussed in more detail in section 4.4 Interacting with Services.

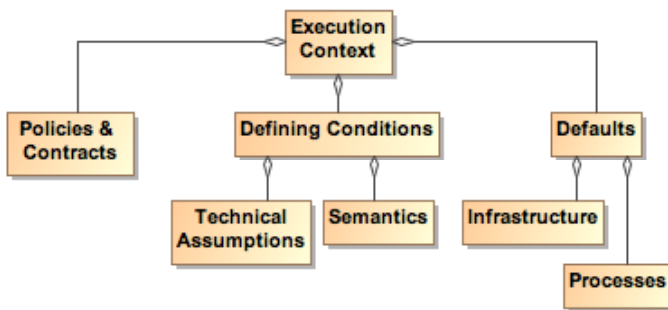
2212 **4.1.2.1.3 Service Description, Execution Context, and Service Interaction**

2213 The service description **MUST** provide sufficient information to support service visibility,
2214 including the willingness of service **participants** to interact. However, the corresponding
2215 descriptions for providers and consumers may both contain policies, technical
2216 assumptions, constraints on semantics, and other technical and procedural conditions
2217 that must be aligned to define the terms of willingness. The agreements which
2218 encapsulate the necessary alignment form the basis upon which interactions may

2219 proceed – in the Reference Model, this collection of agreements and the necessary
2220 environmental support establish the execution context.

2221 To illustrate the concept of the execution context, consider a Web-based system for
2222 timecard entry. For an employee onsite at an employer facility, the execution context
2223 requires a computer connected to the local network and the employee must enter their
2224 network ID and password. Relevant policies include that the employee must maintain
2225 the most recent anti-virus software and virus definitions for any computer connected to
2226 the network.

2227 For the same employee connecting from offsite, the execution context specifies the
2228 need for a computer with installed VPN software and a security token to negotiate the
2229 VPN connection. The execution context also includes proxy settings as needed to
2230 connect to the offsite network. The employee must still comply with the requirements for
2231 onsite computers and access, but the offsite execution context includes additional items
2232 before the employee can access the same underlying capability and realize the same
2233 **real world effect** s, i.e. the timecard entries.

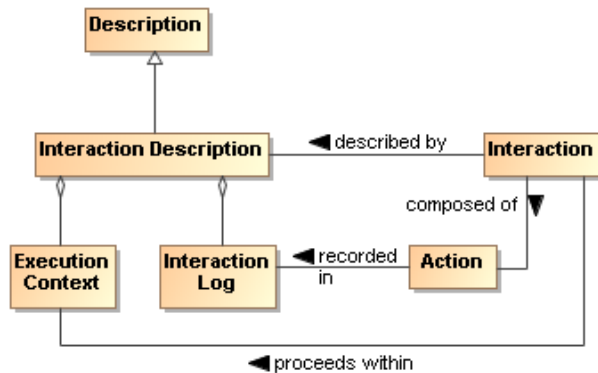


2234
2235 *Figure 33 Execution Context*

2236 Figure 33 shows a few broad categories found in execution context. These are not
2237 meant to be comprehensive. Other items may need to be included to collect a sufficient
2238 description of the interaction conditions. Any other items not explicitly noted in the
2239 model but needed to set the environment SHOULD be included in the execution
2240 context.

2241 While the execution context captures the conditions under which interaction can occur,
2242 it does not capture the specific service invocations that do occur in a specific interaction.
2243 A service interaction as modeled in Figure 32 introduces the concept of an Interaction
2244 Description which is composed of both the Execution Context and an Interaction Log.
2245 The execution context specifies the set of conditions under which the interaction occurs
2246 and the interaction log captures the sequence of service interactions that occur within
2247 the execution context. This sequence should follow the Process Model but can include
2248 details beyond those specified there. For example, the Process Model may specify an
2249 **action** that results in identifying a data source, and the identified source is used in a
2250 subsequent **action**. The Interaction Log would record the specific data source used.

2251 The execution context can be thought of as the container in which the interaction occurs
2252 and the interaction log captures what happens inside the container. This combination is
2253 needed to support auditability and repeatability of the interactions.



2254
2255 *Figure 34 Interaction Description*

2256 SOA allows flexibility to accomplish repeatability or reusability. One benefit of this is that
2257 a service can be updated without disrupting the user experience of the service. So,
2258 Google can improve their ranking algorithm without notifying the user about the details
2259 of the update.

2260 However, it may also be vital for the consumer to be able to recreate past results or to
2261 generate consistent results in the future, and information such as what conditions, which
2262 services, and which versions of those services are used is indispensable in retracing
2263 one's path. The interaction log is a critical part of the resulting **real world effects**
2264 because it defines how the effects were generated and possibly the meaning of
2265 observed effects. This increases in importance as dynamic composability becomes
2266 more feasible. In essence, a result has limited value if one does not know how it was
2267 generated.

2268 The interaction log SHOULD be a detailed trace for a specific interaction, and its reuse
2269 is limited to duplicating that interaction. An execution context can act as a template for
2270 identical or similar interactions. Any given execution context MAY define the conditions
2271 of future interactions.

2272 Such uses of execution context imply (1) a standardized format for capturing execution
2273 context and (2) a subclass of general description could be defined to support visibility of
2274 saved execution contexts. The specifics of the relevant formats and descriptions are
2275 beyond the scope of this Reference Architecture.

2276 A service description is unlikely to track interaction descriptions or the constituent
2277 execution contexts or interaction logs that include mention of the service. However, as
2278 appropriate, linking to specific instances of either of these could be done through
2279 associated annotations.

2280 **4.1.3 Relationship to Other Description Models**

2281 While the representation shown in Figure 25 is derived from considerations related to
2282 service description, it is acknowledged that other metadata standards are relevant and
2283 should, as possible, be incorporated into this work. Two standards of particular
2284 relevance are the Dublin Core Metadata Initiative (DCMI) and ISO 11179, especially
2285 Part 5.

2286 When the service description (or even the general description class) is considered as
2287 the DCMI "resource", Figure 25 aligns nicely with the DCMI resource model. While

2288 some differences exist, these are mostly in areas where DCMI goes into detail that is
2289 considered beyond the scope of the current Reference Architecture. For example,
2290 DCMI defines classes of “shared semantics” whereas this Reference Architecture
2291 considers that an identification of relevant semantic models is sufficient. Likewise, the
2292 DCMI “description model” goes into the details of possible syntax encodings whereas
2293 for the Reference Architecture it is sufficient to identify the relevant formats.

2294 With respect to ISO 11179 Part 5, the metadata fields defined in that reference may be
2295 used without prejudice as the properties in Figure 25. Additionally, other defined
2296 metadata sets may be used by the service provider if the other sets are considered
2297 more appropriate, i.e. it is fundamental to this Reference Architecture to identify the
2298 need and the means to make vocabulary declarations explicit but it is beyond the scope
2299 to specify which vocabularies are to be used. In addition, the identification of domain of
2300 the properties and range of the values has not been included in the current Reference
2301 Architecture discussion, but the text of ISO 11179 Part 5 can be used consistently with
2302 the model prescribed in this document.

2303 Description as defined in the context of this Reference Architecture considers a wide
2304 range of applicability and support of the principles of service oriented architecture.
2305 Other metadata models can be used in concert with the model presented here because
2306 most of these focus on a finer level of detail that is outside the present scope, and so
2307 provide a level of implementation guidance that can be applied as appropriate.

2308 **4.1.4 Architectural Implications**

2309 The description of service description indicates numerous architectural implications on
2310 the SOA ecosystem:

- 2311 • Description will change over time and its contents will reflect changing needs and
2312 context. This requires the existence of:
 - 2313 ○ mechanisms to support the storage, referencing, and access to normative
2314 definitions of one or more versioning schemes that may be applied to
2315 identify different aggregations of descriptive information, where the
2316 different schemes may be versions of a versioning scheme itself;
 - 2317 ○ configuration management mechanisms to capture the contents of the
2318 each aggregation and apply a unique identifier in a manner consistent with
2319 an identified versioning scheme;
 - 2320 ○ one or more mechanisms to support the storage, referencing, and access
2321 to conversion relationships between versioning schemes, and the
2322 mechanisms to carry out such conversions.
- 2323 • Description makes use of defined semantics, where the semantics may be used
2324 for categorization or providing other property and value information for
2325 description classes. This requires the existence of:
 - 2326 ○ semantic models that provide normative descriptions of the utilized terms,
2327 where the models may range from a simple dictionary of terms to an
2328 ontology showing complex relationships and capable of supporting
2329 enhanced reasoning;
 - 2330 ○ mechanisms to support the storage, referencing, and access to these
2331 semantic models;

- 2332
- 2333
- 2334
- 2335
- 2336
- 2337
- 2338
- 2339
- 2340
- 2341
- 2342
- 2343
- 2344
- 2345
- 2346
- 2347
- 2348
- 2349
- 2350
- 2351
- 2352
- 2353
- 2354
- 2355
- 2356
- 2357
- 2358
- 2359
- 2360
- 2361
- 2362
- 2363
- 2364
- 2365
- 2366
- 2367
- 2368
- 2369
- 2370
- 2371
- 2372
- 2373
- 2374
- 2375
- 2376
- 2377
- configuration management mechanisms to capture the normative description of each semantic model and to apply a unique identifier in a manner consistent with an identified versioning scheme;
 - one or more mechanisms to support the storage, referencing, and access to conversion relationships between semantic models, and the mechanisms to carry out such conversions.
 - Descriptions include reference to policies defining conditions of use and optionally contracts representing agreement on policies and other conditions. This requires the existence of (as also enumerated under governance):
 - descriptions to enable the **policy** modules to be visible, where the description includes a unique identifier for the policy and a sufficient, and preferably a machine processible, representation of the meaning of terms used to describe the **policy**, its functions, and its effects;
 - one or more discovery mechanisms that enable searching for policies that best meet the search criteria specified by the service **participant**; where the discovery mechanism will have access to the individual **policy** descriptions, possibly through some repository mechanism;
 - accessible storage of policies and **policy** descriptions, so service **participants** can access, examine, and use the policies as defined.
 - Descriptions include references to metrics which describe the operational characteristics of the subjects being described. This requires the existence of (as partially enumerated under governance):
 - the infrastructure monitoring and reporting information on SOA resources;
 - possible interface requirements to make accessible metrics information generated or most easily accessed by the service itself;
 - mechanisms to catalog and enable discovery of which metrics are available for a described resources and information on how these metrics can be accessed;
 - mechanisms to catalog and enable discovery of compliance records associated with policies and contracts that are based on these metrics.
 - Descriptions of the interactions are important for enabling auditability and repeatability, thereby establishing a context for results and support for understanding observed change in performance or results. This requires the existence of:
 - one or more mechanisms to capture, describe, store, discover, and retrieve interaction logs, execution contexts, and the combined interaction descriptions;
 - one or more mechanisms for attaching to any results the means to identify and retrieve the interaction description under which the results were generated.
 - Descriptions may capture very focused information subsets or can be an aggregate of numerous component descriptions. Service description is an example of a likely aggregate for which manual maintenance of all aspects would not be feasible. This requires the existence of:
 - tools to facilitate identifying description elements that are to be aggregated to assemble the composite description;

- 2378 ○ tools to facilitate identifying the sources of information to associate with
2379 the description elements;
- 2380 ○ tools to collect the identified description elements and their associated
2381 sources into a standard, referenceable format that can support general
2382 access and understanding;
- 2383 ○ tools to automatically update the composite description as the component
2384 sources change, and to consistently apply versioning schemes to identify
2385 the new description contents and the type and significance of change that
2386 occurred.
- 2387 • Descriptions provide up-to-date information on what a **resource** is, the
2388 conditions for interacting with the **resource**, and the results of such interactions.
2389 As such, the description is the source of vital information in establishing
2390 willingness to interact with a **resource**, reachability to make interaction possible,
2391 and compliance with relevant conditions of use. This requires the existence of:
 - 2392 ○ one or more discovery mechanisms that enable searching for described
2393 **resources** that best meet the criteria specified by a service **participant**,
2394 where the discovery mechanism will have access to individual
2395 descriptions, possibly through some repository mechanism;
 - 2396 ○ tools to appropriately track users of the descriptions and notify them when
2397 a new version of the description is available.

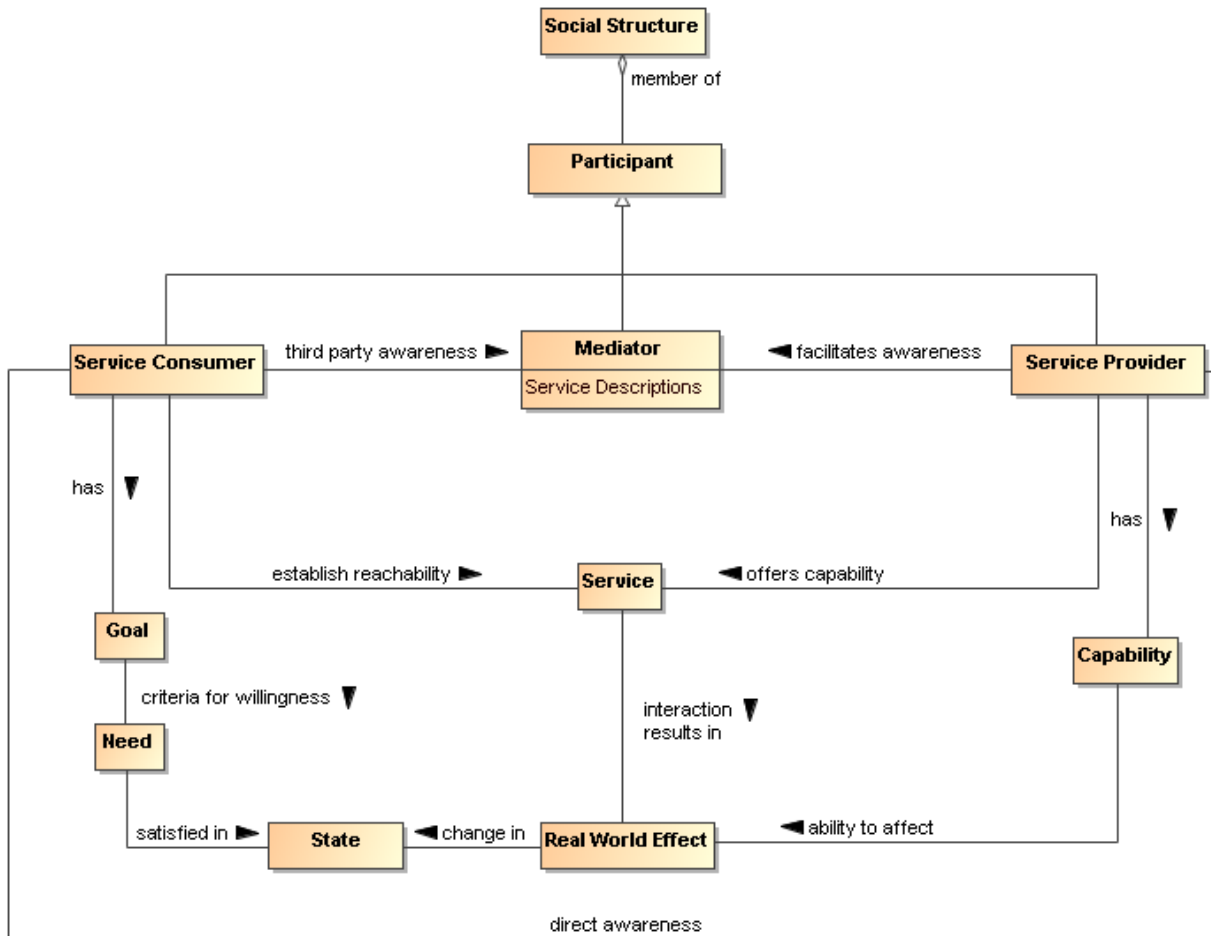
2398 4.2 Service Visibility Model

2399 One of the key requirements for **participants** interacting with each other in the context
2400 of a SOA is achieving visibility: before services can interoperate, the **participants** have
2401 to be visible to each other using whatever means are appropriate. The Reference Model
2402 analyzes visibility in terms of awareness, willingness, and reachability. In this section,
2403 we explore how visibility may be achieved.

2404 4.2.1 Visibility to Business

2405 The relationship of visibility to the SOA ecosystem encompasses both human **social**
2406 **structures** and automated IT mechanisms. Figure 35 depicts a business setting that is
2407 a basis for visibility as related to the **social structure** Model in the Service Ecosystem
2408 View (see Section **Error! Reference source not found.**). **Service consumers** and
2409 service providers may have direct awareness or mediated awareness where mediated
2410 awareness is achieved through some third party. A consumer's willingness to use a
2411 service is reflected by the consumer's presumption of satisfying goals and needs based
2412 on the description of the service. Service providers offer capabilities that have **real**
2413 **world effects** that result in a change in state of the consumer. Reachability of the
2414 service by the consumer leads to interactions that change the state of the consumer.
2415 The consumer can measure the change of state to determine if the claims made by
2416 description and the **real world effects** of consuming the service meet the consumer's
2417 needs.

2418



2419
2420 *Figure 35 Visibility to Business*

2421 Visibility and interoperability in a SOA ecosystem requires more than location and
2422 interface information. A meta-model for this broader view of visibility is depicted in
2423 Section 4.1. In addition to providing improved awareness of service capabilities through
2424 description of information such as reachability, behavior models, information models,
2425 functionality, and metrics, the service description may contain policies valuable for
2426 determination of willingness to interact.

2427 A mediator of service descriptions may provide event notifications to both consumers
2428 and providers about information relating to service descriptions. One example of this
2429 capability is a publish/subscribe model where the mediator allows consumers to
2430 subscribe to service description version changes made by the provider. Likewise, the
2431 mediator may provide notifications to the provider of consumers that have subscribed to
2432 service description updates.

2433 Another important business capability in a SOA environment is the ability to narrow
2434 visibility to trusted members within a **social structure**. Mediators for awareness may
2435 provide **policy** based access to service descriptions allowing for the dynamic formation
2436 of awareness between trusted members.

2437 4.2.2 Visibility

2438 Attaining visibility is described in terms of steps that lead to visibility. While there can be
2439 many contexts for visibility within a single **social structure**, the same general steps can
2440 be applied to each of the contexts to accomplish visibility.

2441 Attaining SOA visibility requires

- 2442 • service description creation and maintenance,
- 2443 • processes and mechanisms for achieving awareness of and accessing descriptions,
- 2444 • processes and mechanisms for establishing willingness of **participants**,
- 2445 • processes and mechanisms to determine reachability.

2446 Visibility may occur in stages, i.e. a **participant** can become aware enough to look or
2447 ask for further description, and with this description, the **participant** can decide on
2448 willingness, possibly requiring additional description. For example, if a potential
2449 consumer has a need for a tree cutting (business) service, the consumer can use a web
2450 search engine to find web sites of providers. The web search engine (a mediator) gives
2451 the consumer links to relevant web pages and the consumer can access those
2452 descriptions. For those prospective providers that satisfy the consumer's criteria, the
2453 consumer's willingness to interact increases. The consumer likely contacts several tree
2454 services to get detailed cost information (or arrange for an estimate) and may ask for
2455 references (further description). Likely, the consumer will establish full visibility and
2456 proceed with the interaction with a tree service who mutually establishes visibility.

2457 4.2.2.1 Awareness

2458 A service **participant** is aware of another **participant** if it has access to a description of
2459 that **participant** with sufficient completeness to establish the other requirements of
2460 visibility.

2461 Awareness is inherently a function of a **participant**; awareness can be established
2462 without any **action** on the part of the target **participant** other than the target providing
2463 appropriate descriptions. Awareness is often discussed in terms of consumer
2464 awareness of providers but the concepts are equally valid for provider awareness of
2465 consumers.

2466 Awareness can be decomposed into the creation of descriptions, making them
2467 available, and discovering the descriptions. Discovery can be initiated or it can be by
2468 notification. Initiated discovery for business may require formalization of the required
2469 capabilities and **resources** to achieve business goals.

2470 Achieving awareness in a SOA can range from word of mouth to formal service
2471 descriptions in a standards-based registry-repository. Some other examples of
2472 achieving awareness in a SOA are the use of a web page containing description
2473 information, email notifications of descriptions, and document based descriptions.

2474 A mediator as discussed for awareness is a third party **participant** that provides
2475 awareness to one or more consumers of one or more services. Direct awareness is
2476 awareness between a consumer and provider without the use of a third party.

2477 Direct awareness may be the result of having previously established an execution
2478 context, or direct awareness may include determining the presence of services and then

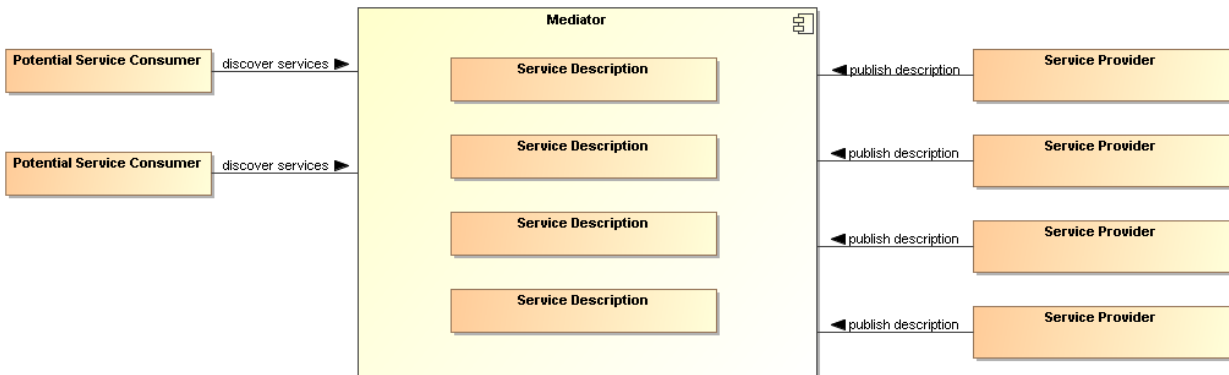
2479 querying the service directly for description. As an example, a priori visibility of some
2480 sensor device may provide the means for interaction or a query for standardized sensor
2481 device metadata may be broadcast to multiple locations. If acknowledged, the service
2482 interface for the device may directly provide description to a consumer so the consumer
2483 can determine willingness to interact.

2484 The same medium for awareness may be direct in one context and may be mediated in
2485 another context. For example, a service provider may maintain a web site with links to
2486 the provider's descriptions of services giving the consumers direct awareness to the
2487 provider's services. Alternatively, a community may maintain a mediated web site with
2488 links to various provider descriptions of services for any number of consumers. More
2489 than one mediator may be involved, as different mediators may specialize in different
2490 mediation functions.

2491 Descriptions may be formal or informal. Section 4.1, provides a comprehensive model
2492 for service description that can be applied to formal registry/repositories used to
2493 mediate visibility. Using consistent description taxonomies and standards based
2494 mediated awareness helps provide more effective awareness.

2495 4.2.2.1.1 Mediated Awareness

2496 Mediated awareness promotes loose coupling by keeping the consumers and services
2497 from explicitly referring to each other and the descriptions. Mediation lets interaction
2498 vary independently. Rather than all potential **service consumers** being informed on a
2499 continual basis about all services, there is a known or agreed upon facility or location
2500 that houses the service description.



2501
2502 *Figure 36 Mediated Service Awareness*

2503 In Figure 36, the potential **service consumers** perform queries or are notified in order
2504 to locate those services that satisfy their needs. As an example, the telephone book is a
2505 mediated registry where individuals perform manual searches to locate services (i.e. the
2506 yellow pages). The telephone book is also a mediated registry for solicitors to find and
2507 notify potential customers (i.e. the white pages).

2508 In mediated service awareness for large and dynamic numbers of **service consumers**
2509 and service providers, the benefits typically far outweigh the management issues
2510 associated with it. Some of the benefits of mediated service awareness are

- 2511 • Potential **service consumers** have a known location for searching thereby
2512 eliminating needless and random searches

2513 • Typically a consortium of interested parties (or a sufficiently large corporation) signs
2514 up to host the mediation facility

2515 • Standardized tools and methods can be developed and promulgated to promote
2516 interoperability and ease of use.

2517 However, mediated awareness can have some risks associated with it:

2518 • A single point of failure. If the central mediation service fails then a potentially large
2519 number of service providers and consumers will be adversely affected.

2520 • A single point of control. If the central mediation service is owned by, or controlled
2521 by, someone other than the **service consumers** and/or providers then the latter
2522 may be put at a competitive disadvantage based on policies of the discovery
2523 provider.

2524 A common mechanism for mediated awareness is a registry-repository. The registry
2525 stores links or pointers to service description artifacts. The repository in this example is
2526 the storage location for the service description artifacts. Service descriptions can be
2527 pushed (publish/subscribe for example) or pulled from the register-repository mediator.

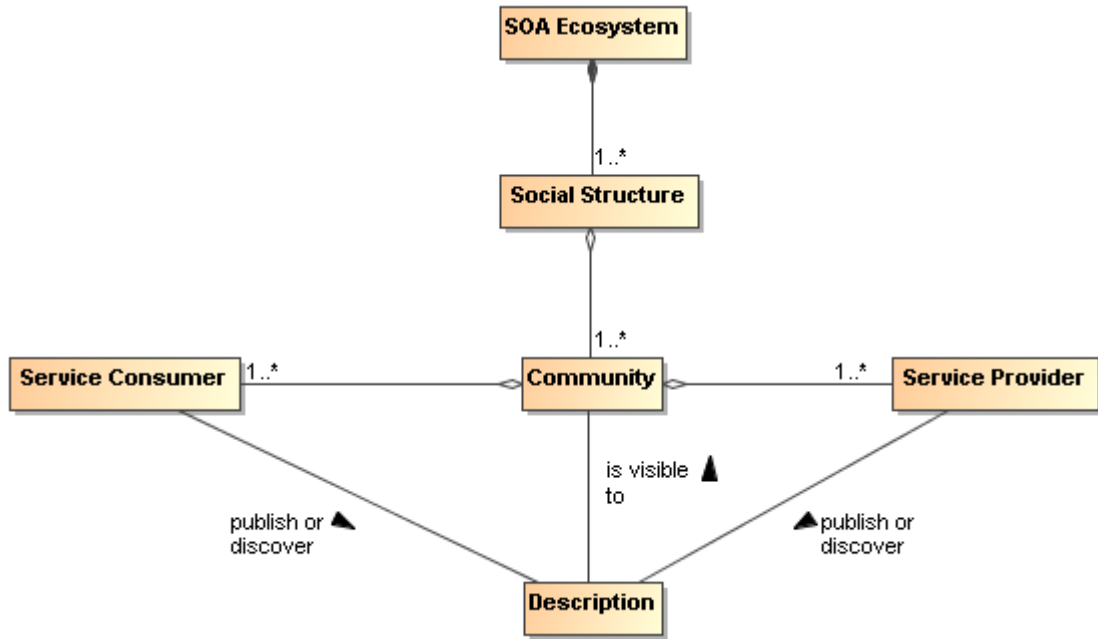
2528 The registry is like a card catalog at the library and a repository is like the shelves for
2529 the books. Standardized metadata describing repository content can be stored as
2530 registry objects in a registry and any type of content can be stored as repository items in
2531 a repository. The registry may be constructed such that description items stored within
2532 the mediation facility repository will have intrinsic links in the registry while description
2533 items stored outside the mediation facility will have extrinsic links in the registry.

2534 When independent but like SOA IT mechanisms interoperate with one another, the IT
2535 mechanisms may be referred to as federated.

2536 **4.2.2.1.2 Awareness in Complex Social Structures**

2537 Awareness applies to one or more communities within one or more **social structures**
2538 where a community consists of at least one description provider and one description
2539 consumer. These communities may be part of the same **social structure** or be part of
2540 different ones.

2541 In Figure 37, awareness can be within a single community, multiple communities, or all
2542 communities in the **social structure**. The **social structure** can encourage or restrict
2543 awareness through its policies, and these policies can affect **participant** willingness.
2544 The information about policies should be incorporated in the relevant descriptions. The
2545 **social structure** also governs the conditions for establishing contracts, the results of
2546 which will be reflected in the execution context if interaction is to proceed.



2547
2548 *Figure 37 Awareness In a SOA Ecosystem*

2549 IT **policy**/contract mechanisms can be used by visibility mechanisms to provide
2550 awareness between communities. The IT mechanisms for awareness may incorporate
2551 trust mechanisms to assure awareness between trusted communities. For example,
2552 government organizations will often want to limit awareness of an organization’s
2553 services to specific communities of interest.

2554 Another common business model for awareness is maximizing awareness to
2555 communities within the **social structure**, the traditional market place business model. A
2556 centralized mediator often arises as a provider for this global visibility, a gatekeeper of
2557 visibility so to speak. For example, Google is a centralized mediator for accessing
2558 information on the web. As another example, television networks have centralized
2559 entities providing a level of awareness to communities that otherwise could not be
2560 achieved without going through the television network.

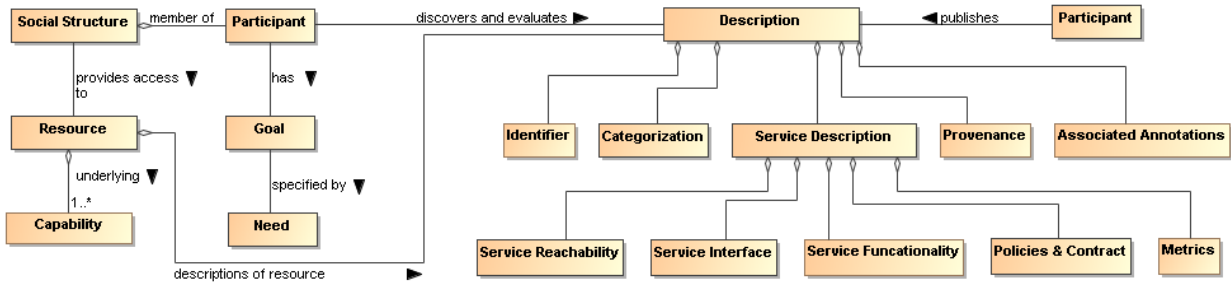
2561 However, mediators have motivations, and they may be selective in which information
2562 they choose to make available to potential consumers. For example, in a secure
2563 environment, the mediator may enforce security policies and make information
2564 selectively available depending on the security clearance of the consumers.

2565 **4.2.2.2 Willingness**

2566 Having achieved awareness, **participants** use descriptions to help determine their
2567 willingness to interact with another **participant**. Both awareness and willingness are
2568 determined prior to consumer/provider interaction.

2569

2570



2571

2572 *Figure 38 Business, Description and Willingness*

2573 Figure 38 relates elements of the Service Ecosystem View, and elements from the
2574 Service Description Model to willingness. By having a willingness to interact within a
2575 particular **social structure**, the social structure provides the **participant** access to
2576 capabilities based on conditions the **social structure** finds appropriate for its context.
2577 The **participant** can use these capabilities to satisfy goals and objectives as specified
2578 by the **participant's** needs.

2579 In Figure 38, information used to determine willingness is defined by Description.
2580 Information referenced by Description may come from many sources. For example, a
2581 mediator for descriptions may provide 3rd party annotations for reputation. Another
2582 source for reputation may be a **participant's** own history of interactions with another
2583 **participant**.

2584 A **participant** will inspect functionality for potential satisfaction of needs. Identity is
2585 associated with any **participant**, however, identity may or may not be verified. If
2586 available, **participant** reputation may be a deciding factor for willingness to interact.
2587 Policies and contracts referenced by the description may be particularly important to
2588 determine the agreements and **commitments** required for business interactions.
2589 Provenance may be used for verification of authenticity of a **resource**.

2590 Mechanisms that aid in determining willingness will likely make use of the artifacts
2591 referenced by descriptions of services. Mechanisms for establishing willingness could
2592 be as simple as rendering service description information for human consumption to
2593 automated evaluation of functionality, policies, and contracts by a rules engine. The
2594 rules engine for determining willingness could operate as a **policy decision procedure**
2595 as defined in Section 4.4.

2596 **4.2.2.3 Reachability**

2597 Reachability involves knowing the endpoint, protocol, and presence of a service. At a
2598 minimum, reachability requires information about the location of the service and the
2599 protocol describing the means of communication.

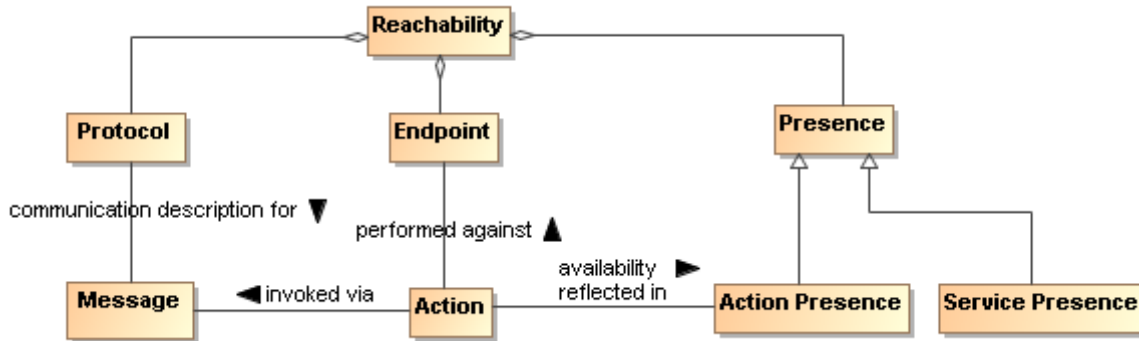


Figure 39 Service Reachability

Endpoint

An endpoint is a reference-able entity, processor or **resource** against which an **action** can be performed.

Protocol

A protocol is a structured means by which service interaction is regulated.

Presence

Presence is the measurement of reachability of a service at a particular point in time.

A protocol defines a structured method of communication with a service. Presence is determined by interaction through a communication protocol. Presence may not be known in many cases until the act of interaction begins. To overcome this problem, IT mechanisms may make use of presence protocols to provide the current up/down status of a service.

Service reachability enables service **participants** to locate and interact with one another. Each **action** may have its own endpoint and also its own protocols associated with the endpoint and whether there is presence for the **action** through that endpoint. Presence of a service is an aggregation of the presence of the service's actions, and the service level may aggregate to some degraded or restricted presence if some **action** presence is not confirmed. For example, if error processing actions are not available, the service can still provide required functionality if no error processing is needed. This implies reachability relates to each **action** as well as applying to the service/business as a whole.

4.2.3 Architectural Implications

Visibility in a SOA ecosystem has the following architectural implications on mechanisms providing support for awareness, willingness, and reachability:

- Mechanisms providing support for awareness will likely have the following minimum capabilities:
 - creation of Description, preferably conforming to a standard Description format and structure;

- 2633 ○ publishing of Description directly to a consumer or through a third party
- 2634 mediator;
- 2635 ○ discovery of Description, preferably conforming to a standard for Description
- 2636 discovery;
- 2637 ○ notification of Description updates or notification of the addition of new and
- 2638 relevant Descriptions;
- 2639 ○ classification of Description elements according to standardized classification
- 2640 schemes.
- 2641 • In a SOA ecosystem with complex **social structures**, awareness may be provided
- 2642 for specific communities of interest. The architectural mechanisms for providing
- 2643 awareness to communities of interest will require support for:
 - 2644 ○ policies that allow dynamic formation of communities of interest;
 - 2645 ○ trust that awareness can be provided for and only for specific communities of
 - 2646 interest, the bases of which is typically built on keying and encryption
 - 2647 technology.
- 2648 • The architectural mechanisms for determining willingness to interact will require
- 2649 support for:
 - 2650 ○ verification of identity and credentials of the provider and/or consumer;
 - 2651 ○ access to and understanding of description;
 - 2652 ○ inspection of functionality and capabilities;
 - 2653 ○ inspection of policies and/or contracts.
- 2654 • The architectural mechanisms for establishing reachability will require support for:
 - 2655 ○ the location or address of an endpoint;
 - 2656 ○ verification and use of a service interface by means of a communication
 - 2657 protocol;
 - 2658 ○ determination of presence with an endpoint which may only be determined at
 - 2659 the point of interaction but may be further aided by the use of a presence
 - 2660 protocol for which the endpoints actively participate.

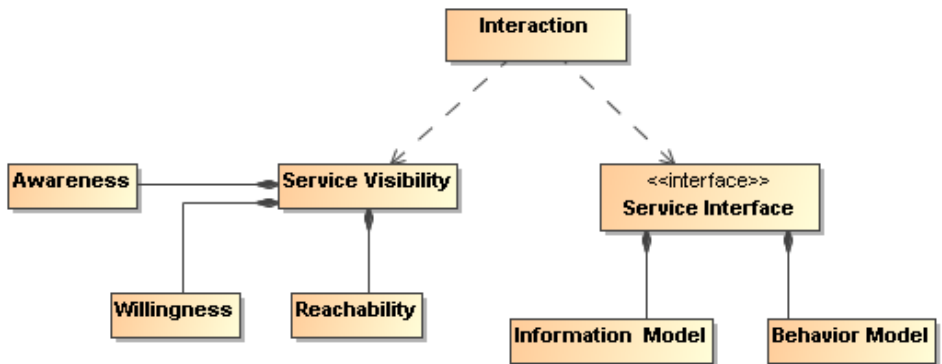
2661 4.3 Interacting with Services Model

2662 Interaction is the activity involved in using a service to access capability in order to
2663 achieve a particular desired **real world effect**, where real world effect is the actual
2664 *result* of using a service. An interaction can be characterized by a sequence of actions.
2665 Consequently, interacting with a service, i.e. performing actions against the service—
2666 usually mediated by a series of message exchanges—involves actions performed by
2667 the service. Different modes of interaction are possible such as modifying the shared
2668 state of a **resource**. Note that a **participant** (or **delegate** acting on behalf of the
2669 **participant**) can be the sender of a message, the receiver of a message, or both.

2670 4.3.1 Interaction Dependencies

2671 Recall from the Reference Model that service visibility is the capacity for those with
2672 needs and those with capabilities to be able to interact with each other, and that the
2673 service interface is the means by which the underlying capabilities of a service are
2674 accessed. Ideally, the details of the underlying service implementation are abstracted
2675 away by the service interface. [Service] interaction therefore has a direct dependency
2676 on the visibility of the service as well as its implementation-neutral interface (see Figure
2677 40). Service visibility is composed of awareness, willingness, and reachability and

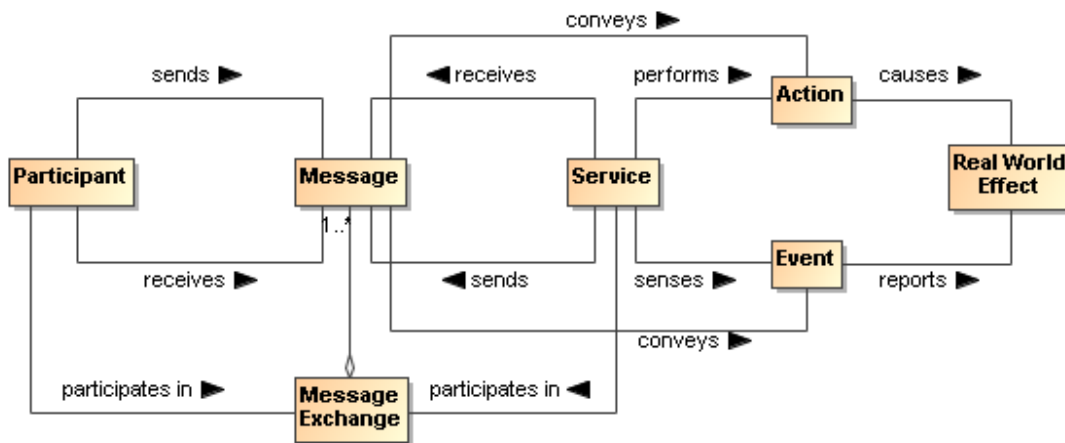
2678 service interface is composed of the information and behavior models. Service visibility
 2679 is modeled in Section 4.2 while service interface is modeled in Section 4.1.



2680
 2681 *Figure 40 Interaction dependencies.*

2682 **4.3.2 Actions and Events**

2683 For purposes of this Reference Architecture, the authors have committed to the use of
 2684 message exchange between service **participants** to denote actions performed against
 2685 and by the service, and to denote **events** that report on **real world effects** that are
 2686 caused by the service actions. A visual model of the relationship between these
 2687 concepts is shown in Figure 41.



2688
 2689 *Figure 41 A "message" conveys either an action or an event.*

2690 A **message** conveys either an **action** or an **event**. In other words, both actions and
 2691 **events**, realized by the SOA services, are denoted by the messages. The Reference
 2692 Model states that the action model characterizes the "permissible set of actions that
 2693 may be invoked against a service." We extend that notion here to include **events** as
 2694 part of the event model and that messages denote either actions or notification of
 2695 **events**.

2696 In Section **Error! Reference source not found.**, we saw that **participants** interact with
 2697 each other in order to perform actions. An **action** is not itself the same thing as the
 2698 result of performing the **action**. When an **action** is performed against a service, the **real**
 2699 **world effect** that results is reported in the form of notification of **events**.

2700 4.3.3 Message Exchange

2701 *Message exchange* is the means by which service **participants** (or their **delegates**)
2702 interact with each other. There are two primary modes of interaction: joint actions that
2703 cause **real world effects**, and notification of **events** that report real world effects.¹⁶

2704 A message exchange is used to affect an **action** when the messages contain the
2705 appropriately formatted content that should be interpreted as joint action and the
2706 **delegates** involved interpret the message appropriately.

2707 A message exchange is also used to communicate **event** notifications. An **event** is an
2708 occurrence that is of interest to some **participant**; in our case when some **real world**
2709 **effect** has occurred. Just as action messages will have formatting requirements, so will
2710 event notification messages. In this way, the Information Model of a service must
2711 specify the syntax (structure), and semantics (meaning) of the action messages and
2712 event notification messages as part of a service interface. It must also specify the
2713 syntax and semantics of any data that is carried as part of a payload of the action or
2714 event notification message. The Information Model is described in greater detail in the
2715 Service Description Model (see Section 4.1).

2716 In addition to the Information Model that describes the syntax and semantics of the
2717 messages and data payloads, exception conditions and error handling in the event of
2718 faults (e.g., network outages, improper message formats, etc.) must be specified or
2719 referenced as part of the Service Description.

2720 When a message is interpreted as an **action**, the correct interpretation typically requires
2721 the receiver to perform a set of operations. These *operations* represent the sequence
2722 of actions (often private) a service must perform in order to validly participate in a given
2723 joint action.

2724 Similarly, the correct consequence of realizing a **real world effect** may be to initiate the
2725 reporting of that real world effect via an event notification.

2726 **Message Exchange**

2727 The means by which joint actions and event notifications are coordinated by
2728 service **participants** (or **delegates**).

2729 **Operations**

2730 The sequence of actions a service must perform in order to validly participate in a
2731 given joint action.

2732 4.3.3.1 Message Exchange Patterns (MEPs)

2733 As stated earlier, this Reference Architecture commits to the use of message exchange
2734 to denote actions against the services, and to denote notification of **events** that report
2735 on **real world effects** that arise from those actions.

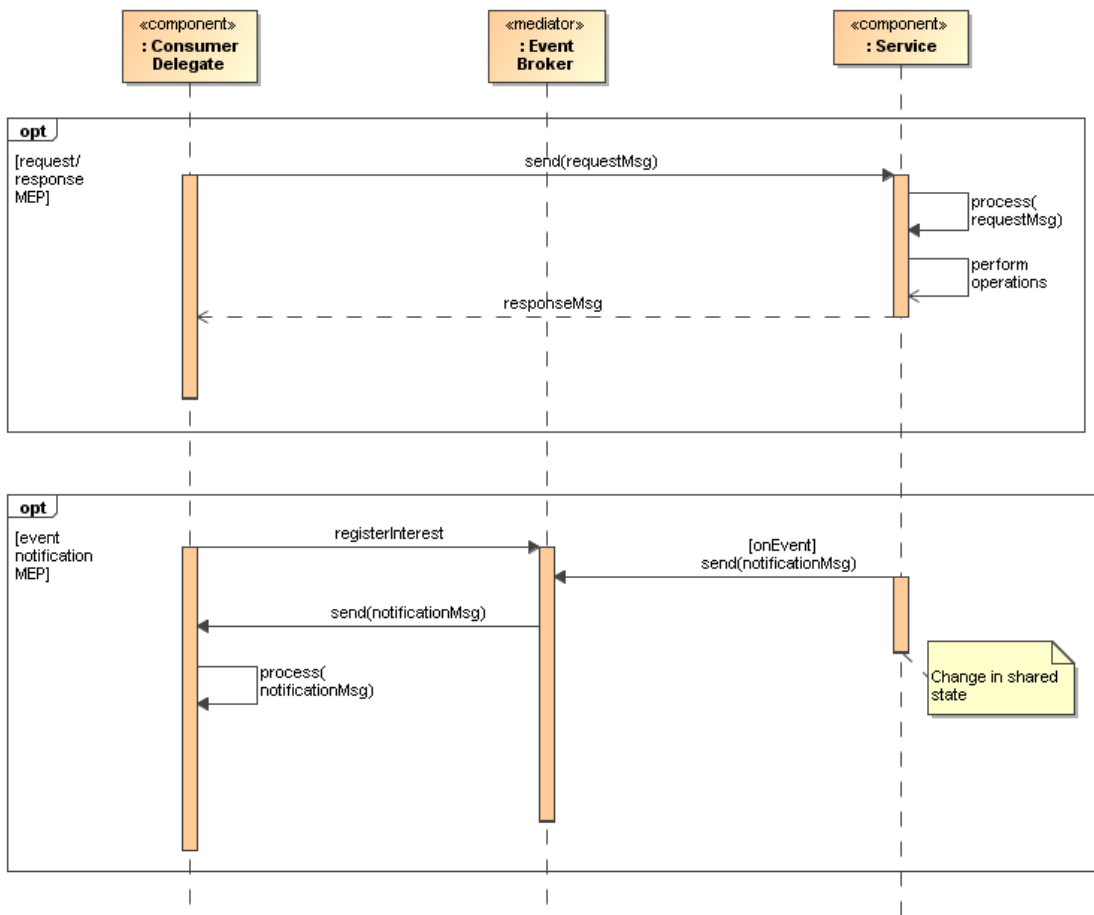
2736 Based on these assumptions, the basic temporal aspect of service interaction can be
2737 characterized by two fundamental message exchange patterns (MEPs):

¹⁶ The notion of “joint” in joint action implies that you have to have a speaker *and* a listener in order to interact.

- 2738 • Request/response to represent how actions cause a **real world effect**
- 2739 • Event notification to represent how **events** report a **real world effect**

2740 This is by no means a complete list of all possible MEPs used for inter- or intra-
 2741 enterprise messaging but it does represent those that are most commonly used in
 2742 exchange of information and reporting changes in state both within organizations and
 2743 across organizational boundaries, a hallmark of a SOA.

2744 Recall from the Reference Model that the Process Model characterizes “the temporal
 2745 relationships between and temporal properties of actions and **events** associated with
 2746 interacting with the service.” Thus, MEPs are a key element of the Process Model. The
 2747 meta-level aspects of the Process Model (just as with the Action Model) are provided as
 2748 part of the Service Description Model (see Section 4.1).



2749
 2750 *Figure 42 Fundamental SOA message exchange patterns (MEPs)*

2751 In the UML sequence diagram shown in Figure 42 it is assumed that the service
 2752 **participants** (consumer and provider) have delegated message handling to hardware
 2753 or software delegates acting on their behalf. In the case of the **service consumer**, this
 2754 is represented by the *Consumer Delegate* component. In the case of the service
 2755 provider, the **delegate** is represented by the *Service* component. The message
 2756 interchange model illustrated represents a logical view of the MEPs and not a physical
 2757 view. In other words, specific hosts, network protocols, and underlying messaging
 2758 system are not shown as these tend to be implementation specific. Although such

2759 implementation-specific elements are considered outside the scope of this Reference
2760 Architecture, they are important considerations in modeling the SOA execution context.
2761 Recall from the Reference Model that the *execution context* of a service interaction is
2762 “the set of infrastructure elements, process entities, policy assertions and agreements
2763 that are identified as part of an instantiated service interaction, and thus forms a path
2764 between those with needs and those with capabilities.”

2765 4.3.3.2 Request/Response MEP

2766 In a request/response MEP, the Consumer Delegate component sends a request
2767 message to the Service component. The Service component then processes the
2768 request message. Based on the content of the message, the Service component
2769 performs the service operations. Following the completion of these operations, a
2770 response message is returned to the Consumer Delegate component. The response
2771 could be that a step in a process is complete, the initiation of a follow-on operation, or
2772 the return of requested information.¹⁷

2773 Although the sequence diagram shows a *synchronous* interaction (because the sender
2774 of the request message, i.e., Consumer Delegate, is blocked from continued processing
2775 until a response is returned from the Service) other variations of request/response are
2776 valid, including *asynchronous* (non-blocking) interaction through use of queues,
2777 channels, or other messaging techniques.

2778 What is important to convey here is that the request/response MEP represents **action**,
2779 which causes a **real world effect**, irrespective of the underlying messaging techniques
2780 and messaging infrastructure used to implement the request/response MEP.

2781 4.3.3.3 Event Notification MEP

2782 An **event** is made visible to interested consumers by means of an event notification
2783 message exchange that reports a **real world effect**; specifically, a change in shared
2784 state between service **participants**. The basic event notification MEP takes the form of
2785 a one-way message sent by a notifier component (in this case, the Service component)
2786 and received by components with an interest in the **event** (here, the Consumer
2787 Delegate component).


2788 Often the sending component may not be fully aware of all the components that will
2789 receive the notification; particularly in so-called publish/subscribe (“pub/sub”) situations.
2790 In event notification message exchanges, it is rare to have a tightly-coupled link
2791 between the sending and the receiving component(s) for a number of practical reasons.
2792 One of the most common is the potential for network outages or communication

¹⁷ There are cases when a response is not always desired and this would be an example of a “one-way” MEP. Similarly, while not shown here, there are cases when some type of “callback” MEP is required in which the consumer agent is actually exposed as a service itself and is able to process incoming messages from another service.

2793 interrupts that can result in loss of notification of **events**. Therefore, a third-party
2794 mediator component is often used to decouple the sending and receiving components .
2795 Although this is typically an implementation issue, because this type of third-party
2796 decoupling is so common in event-driven systems, we felt that for this Reference
2797 Architecture, it was warranted for use in modeling this type of message exchange. This
2798 third-party intermediary is shown in Figure 42 as an Event Broker mediator. As with the
2799 request/response MEP, no distinction is made between synchronous versus
2800 asynchronous communication, although asynchronous message exchange is illustrated
2801 in the UML sequence diagram depicted in Figure 42 .

2802 4.3.4 Composition of Services

2803 Composition of services is the act of aggregating or “composing” a single service from
2804 one or more other services. A simple model of service composition is illustrated in
2805 Figure 43.



2806
2807 *Figure 43 Simple model of service composition.*

2808 Here, Service A is a service that has an exposed interface IServiceA, which is available
2809 to the Consumer Delegate and relies on two other services in its implementation. The
2810 Consumer Delegate does not know that Services B and C are used by Service A, or
2811 whether they are used in serial or parallel, or if their operations succeed or fail. The
2812 Consumer Delegate only cares about the success or failure of Service A. The exposed
2813 interfaces of Services B and C (IService B and IServiceC) are not necessarily hidden
2814 from the Consumer Delegate; only the fact that these services are used as part of the
2815 composition of Service A. In this example, there is no practical reason the Consumer
2816 Delegate could not interact with Service B or Service C in some other interaction
2817 scenario.

2818 It is possible for a service composition to be opaque from one perspective and
2819 transparent from another. For example, a service may appear to be a single service
2820 from the Consumer’s Delegate’s perspective, but is transparently composed of one or
2821 more services from a service management perspective. A Service Management Service
2822 needs to be able to have visibility into the composition in order to properly manage the
2823 dependencies between the services used in constructing the composite service—
2824 including managing the service’s lifecycle. The subject of services as management
2825 entities is described and modeled in the Owing Service Oriented Architectures View of
2826 this Reference Architecture and will not be further elaborated in this section. The point
2827 to be made here is that there can be different levels of opaqueness or transparency
2828 when it comes to visibility of service composition.

2829 Services can be composed in a variety of ways including direct service-to-service
2830 interaction by using programming techniques, or they can be aggregated by means of a
2831 scripting approach that leverages a service composition scripting language. Such
2832 scripting approaches are further elaborated in the following sub-sections on service-
2833 oriented business processes and collaborations.

2834 **4.3.4.1 Service-Oriented Business Processes**

2835 The concepts of business processes and collaborations in the context of transactions
2836 and exchanges across organizational boundaries are described and modeled as part of
2837 the Service Ecosystem View of this Reference Architecture (see Section 3). Here, we
2838 focus on the belief that the principle of composition of services can be applied to
2839 business processes and collaborations. Of course, business processes and
2840 collaborations traditionally represent complex, multi-step business functions that may
2841 involve multiple **participants**, including internal users, external customers, and trading
2842 partners. Therefore, such complexities cannot simply be ignored when transforming
2843 traditional business processes and collaborations to their service-oriented variants.

2844 **Business Processes**

2845 Business processes are a set of one or more linked activities that are performed
2846 to achieve a certain business outcome.

2847 Service orientation as applied to business processes (i.e., “service-oriented business
2848 processes”) means that the aggregation or composition of all of the abstracted activities,
2849 flows, and rules that govern a business process can themselves be abstracted as a
2850 service **[BLOOMBERG/SCHMELZER]**.

2851 When business processes are abstracted in this manner and accessed through SOA
2852 services, all of the concepts used to describe and model composition of services that
2853 were articulated in Section 4.3.4 apply. There are some important differences from a
2854 composite service that represents an abstraction of a business process from a
2855 composite service that represents a single-step business interaction. As stated earlier,
2856 business processes have temporal properties and can range from short-lived processes
2857 that execute on the order of minutes or hours to long-lived processes that can execute
2858 for weeks, months, or even years. Further, these processes may involve many
2859 **participants**. These are important considerations for the consumer of a service-
2860 oriented business process and these temporal properties must be articulated as part of
2861 the meta-level aspects of the service-oriented business process in its Service
2862 Description, along with the meta-level aspects of any sub-processes that may be of use
2863 or need to be visible to the **service consumer**.

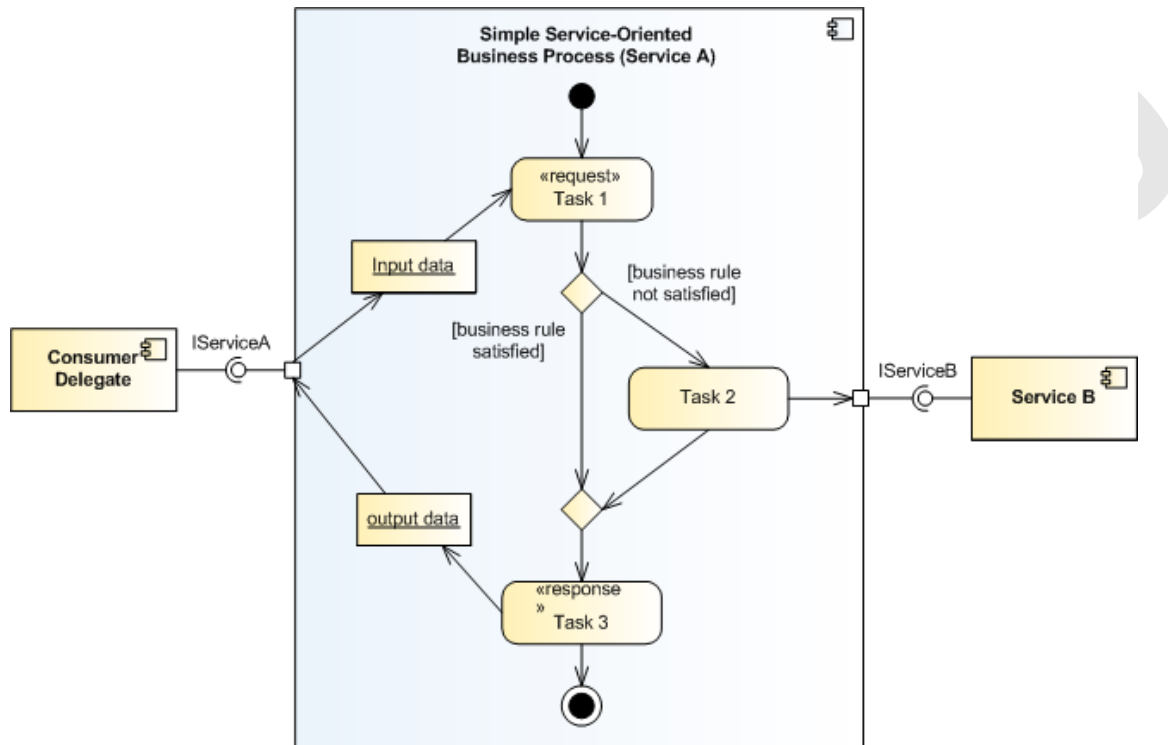
2864 In addition, a workflow activity represents a unit of work that some entity acting in a
2865 described **role** (i.e., **role player**) is asked to perform. Activities can be broken down
2866 into steps with each step representing a task for the **role player** to perform. A
2867 technique that is used to compose service-oriented business processes that are
2868 hierarchical (top-down) and self-contained in nature is known as *orchestration*.

2869 **Orchestration**

2870 A technique used to compose service-oriented business processes that are
2871 executed and coordinated by an **actor** acting as “conductor.”

2872 An orchestration is typically implemented using a scripting approach to compose
 2873 service-oriented business processes. This typically involves use of a standards-based
 2874 orchestration scripting language. In terms of automation, an orchestration can be
 2875 mechanized using a business process orchestration engine, which is a hardware or
 2876 software component (**delegate**) responsible for acting in the role of central
 2877 conductor/coordinator responsible for executing the flows that comprise the
 2878 orchestration.

2879 A simple generic example of such an orchestration is illustrated in Figure 44.



2880
 2881 *Figure 44 Abstract example of orchestration of service-oriented business process.*

2882 Here, we use a UML activity diagram to model the simple service-oriented business
 2883 process as it allows us to capture the major elements of business processes such as
 2884 the set of related tasks to be performed, linking between tasks in a logical flow, data that
 2885 is passed between tasks, and any relevant business rules that govern the transitions
 2886 between tasks. A task is a unit of work that an individual, system, or organization
 2887 performs and can be accomplished in one or more steps or subtasks. While subtasks
 2888 can be readily modeled, they are not illustrated in the orchestration model In Figure 44..

2889 This particular example is based on a request/response MEP and captures how one
 2890 particular task (Task 2) actually utilizes an externally-provided service, Service B. The
 2891 entire service-oriented business process is exposed as Service A that is accessible via
 2892 its externally visible interface, IServiceA.

2893 Although not explicitly shown in the orchestration model above, it is assumed that there
 2894 exists a software or hardware component, i.e., orchestration engine that executes the
 2895 process flow. Recall that a central concept to orchestration is that process flow is
 2896 coordinated and executed by a single conductor delegate; hence the name
 2897 “orchestration.”

2898 4.3.4.2 Service-Oriented Business Collaborations

2899 Business collaborations typically represent the interaction involved in executing
2900 business transactions, where a *business transaction* is defined in the Service
2901 Ecosystem View as “a joint action engaged in by two or more **participants** in which
2902 **resources** are exchanged” (see Section 3.2.3).

2903 It is important to note that business collaborations represent “**peer**”-style interactions; in
2904 other words, **peers** in a business collaboration act as equals. This means that unlike
2905 the orchestration of business processes, there is no single or central entity that
2906 coordinates or “conducts” a business collaboration. These peer styles of interactions
2907 typically occur between trading partners that span organizational boundaries.

2908 Business collaborations can also be service-enabled. For purposes of this Reference
2909 Architecture, we refer to these as “service-oriented business collaborations.” Service-
2910 oriented business collaborations do not necessarily imply exposing the entire peer-style
2911 business collaboration as a service itself but rather the collaboration uses service-based
2912 interchanges.

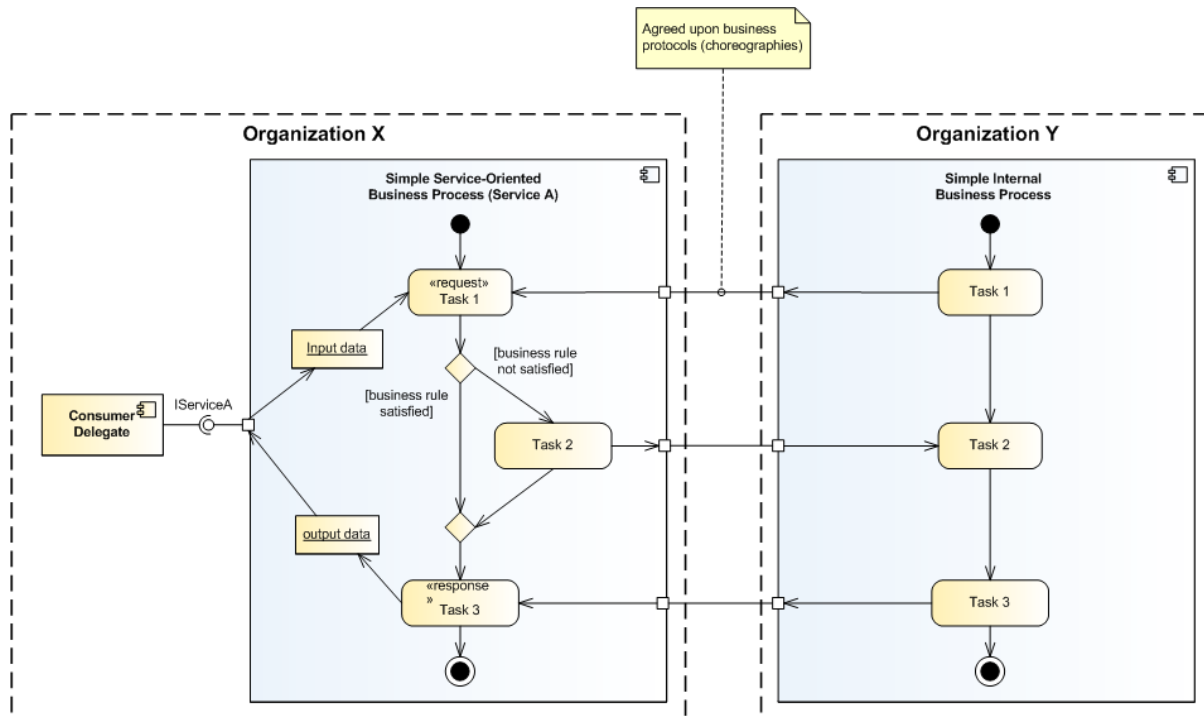
2913 The technique that is used to compose service-oriented business collaborations in
2914 which multiple parties collaborate in a peer-style as part of some larger business
2915 transaction by exchanging messages with trading partners and external organizations
2916 (e.g., suppliers) is known as *choreography* [NEWCOMER/LOMOW].

2917 **Choreography**

2918 A technique used to characterize service-oriented business collaborations based
2919 on ordered message exchanges between **peer** entities in order to achieve a
2920 common business goal.

2921 Choreography differs from orchestration primarily in that each party in a business
2922 collaboration describes its part in the service interaction. Note that choreography as we
2923 have defined it here should not be confused with the term *process choreography*, which
2924 is defined in the Service Ecosystem View as “the description of the possible interactions
2925 that may take place between two or more **participants** to fulfill an objective.” This is an
2926 example of domain-specific nomenclature that often leads to confusion and why we are
2927 making note of it here.

2928 A simple generic example of a choreography is illustrated in Figure 45



2929
2930

Figure 45 Abstract example of choreography of service-oriented business collaboration.

2931 This example, which is a variant of the orchestration example illustrated earlier in Figure
2932 44 adds trust boundaries between two organizations; namely, Organization X and
2933 Organization Y. It is assumed that these two organizations are peer entities that have
2934 an interest in a business collaboration, for example, Organization X and Organization Y
2935 could be trading partners. Organization X retains the service-oriented business process
2936 Service A, which is exposed to internal consumers via its provided service interface,
2937 IServiceA. Organization Y also has a business process that is involved in the business
2938 collaboration; however, for this example, it is an internal business process that is not
2939 exposed to potential consumers either within or outside its organizational boundary.

2940 The scripting language that is used for the choreography needs to define how and when
2941 to pass control from one trading partner to another, i.e., Organization X and
2942 Organization Y. Defining the business protocols used in the business collaboration
2943 involves precisely specifying the visible message exchange behavior of each of the
2944 parties involved in the protocol, without revealing internal implementation details
2945 **[NEWCOMER/LOMOW]**.

2946 In a peer-style business collaboration, a choreography scripting language must be
2947 capable of describing the coordination of those service-oriented processes that cross
2948 organizational boundaries.

2949 4.3.5 Architectural Implications of Interacting with Services

2950 Interacting with Services has the following architectural implications on mechanisms
2951 that facilitate service interaction:

- 2952 • A well-defined service Information Model that:
 - 2953 ○ describes the syntax and semantics of the messages used to denote actions
 - 2954 and **events**;

- 2955 ○ describes the syntax and semantics of the data payload(s) contained within
- 2956 messages;
- 2957 ○ documents exception conditions in the event of faults due to network outages,
- 2958 improper message/data formats, etc.;
- 2959 ○ is both human readable and machine processable;
- 2960 ○ is referenceable from the Service Description artifact.
- 2961 • A well-defined service Behavior Model that:
 - 2962 ○ characterizes the knowledge of the actions invokes against the service and
 - 2963 **events** that report **real world effects** as a result of those actions;
 - 2964 ○ characterizes the temporal relationships and temporal properties of actions
 - 2965 and **events** associated in a service interaction;
 - 2966 ○ describe activities involved in a workflow activity that represents a unit of
 - 2967 work;
 - 2968 ○ describes the **role** (s) that a **role player** performs in a service-oriented
 - 2969 business process or service-oriented business collaboration;
 - 2970 ○ is both human readable and machine processable;
 - 2971 ○ is referenceable from the Service Description artifact.
- 2972 • Service composition mechanisms to support orchestration of service-oriented
- 2973 business processes and choreography of service-oriented business collaborations
- 2974 such as:
 - 2975 ○ Declarative and programmatic compositional languages;
 - 2976 ○ Orchestration and/or choreography engines that support multi-step
 - 2977 processes as part of a short-lived or long-lived business transaction;
 - 2978 ○ Orchestration and/or choreography engines that support compensating
 - 2979 transactions in the presences of exception and fault conditions.
- 2980 • Infrastructure services that provides mechanisms to support service interaction,
- 2981 including but not limited to:
 - 2982 ○ mediation services such as message and event brokers, providers, and/or
 - 2983 buses that provide message translation/transformation, gateway
 - 2984 capability, message persistence, reliable message delivery, and/or
 - 2985 intelligent routing semantics;
 - 2986 ○ binding services that support translation and transformation of multiple
 - 2987 application-level protocols to standard network transport protocols;
 - 2988 ○ auditing and logging services that provide a data store and mechanism to
 - 2989 record information related to service interaction activity such as message
 - 2990 traffic patterns, security violations, and service contract and **policy**
 - 2991 violations
 - 2992 ○ security services that abstract techniques such as public key
 - 2993 cryptography, secure networks, virus protection, etc., which provide
 - 2994 protection against common security threats in a SOA ecosystem;
 - 2995 ○ monitoring services such as hardware and software mechanisms that both
 - 2996 monitor the performance of systems that host services and network traffic
 - 2997 during service interaction, and are capable of generating regular
 - 2998 monitoring reports.
- 2999 • A layered and tiered service component architecture that supports multiple message
- 3000 exchange patterns (MEPs) in order to:

- 3001 ○ promote the industry best practice of separation of concerns that facilitates
- 3002 flexibility in the presence of changing business requirements;
- 3003 ○ promote the industry best practice of separation of **roles** in a service
- 3004 development lifecycle such that subject matter experts and teams are
- 3005 structured along areas of expertise;
- 3006 ○ support numerous standard interaction patterns, peer-to-peer interaction
- 3007 patterns, enterprise integration patterns, and business-to-business
- 3008 integration patterns.

3009 4.4 Policies and Contracts Model

3010 A common phenomenon of many machines and systems is that the scope of potential
3011 behavior is much broader than is actually needed for a particular circumstance. This is
3012 especially true of a system as powerful as a SOA ecosystem. As a result, the behavior
3013 and performance of the system tend to be under-constrained by the implementation;
3014 instead, the actual behavior is expressed by means of **policies** of some form. Policies
3015 define the choices that **stakeholders** make; these choices are used to guide the actual
3016 behavior of the system to the desired behavior and performance.

3017 As noted in Section 3.1.4 a **policy** is a constraint of some form that is promulgated by a
3018 **stakeholder** who has the responsibility of ensuring that the constraint is enforced. In
3019 contrast, **contracts** are **agreements** between **participants**. However, like **policies**, it is
3020 a necessary part of **contracts** that they are enforceable.

3021 While responsibility for enforcement may differ, both **contracts** and **policies** share a
3022 common characteristic – there is a **constraint** that must be enforced. In both cases the
3023 mechanisms needed to enforce **policy constraints** will be largely identical; in this
3024 model we focus on the issues involved in representing **policies** and **contracts** and on
3025 some of the principles behind their enforcement.

3026 4.4.1 Policy and Contract Representation

3027 A **policy constraint** is a specific kind of constraint: the ontology of **policies** and
3028 **contracts** includes the core concepts of **permission**, **obligation**, **owner**, **subject**. In
3029 addition, it may be necessary to be able combine **policy constraints** and to be able to
3030 resolve **policy conflicts**.

3031 4.4.1.1 Policy Framework

3032 Policy Framework

3033 A **policy framework** is a language in which **policy constraints** may be
3034 expressed.

3035 A **policy framework** combines a syntax for expressing **policy constraints** together
3036 with a **decision procedure** for determining if a **policy constraint** is satisfied.

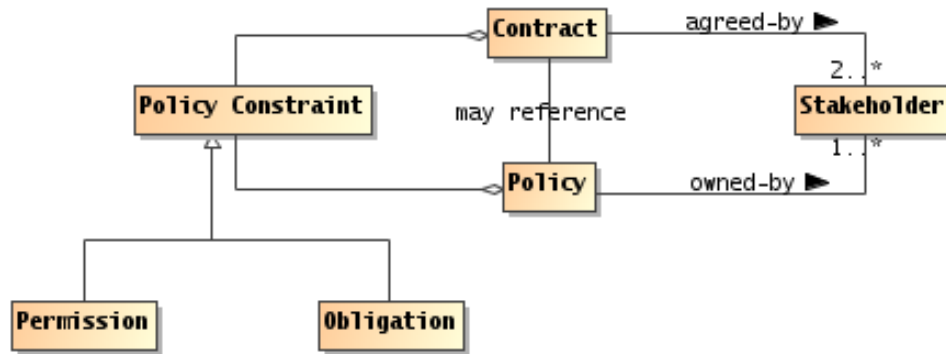


Figure 46 Policies and Contracts

We can characterize (caricature) a **policy framework** in terms of a **logical framework** and an **ontology** of **policies**. The **policy ontology** details specific kinds of **policy constraints** that can be expressed; and the **logical framework** is a ‘glue’ that allows us to express combinations of **policies**.

Logical Framework

A **logical framework** is a linguistic framework consisting of a syntax – a way of writing expressions – and a semantics – a way of interpreting the expressions.

Policy Ontology

A **policy ontology** is a formalization of a set of concepts that are relevant to forming policy expressions.

For example, a **policy ontology** that allows to identify simple constraints – such as the existence of a property, or that a value of a property should be compared to a fixed value – is often enough to express many basic constraints.

Included in many **policy ontologies** are the basic signals of **permissions** and **obligations**. Some **policy frameworks** are sufficiently constrained that there is not possibility of representing an **obligation**; in which case there is often no need to ‘call out’ the distinction between **permissions** and **obligations**.

The **logical framework** is also a strong determiner of the expressivity of the **policy framework**. The richer the **logical framework**, the richer the set of **policy constraints** that can be expressed. However, there is a strong inverse correlation between expressivity and ease and efficiency of implementation.

In the discussion that follows we assume the following basic **policy ontology**:

Policy Owner

A **policy owner** is a **stakeholder** that asserts and enforces the **policy**.

Policy Subject

A **policy subject** is an **actor** who is subject to the constraints of a **policy** or **contract**.

Policy Constraint

A **policy constraint** is a measurable **proposition** that characterizes the constraint that the policy is about.

3069 **Policy Object**

3070 A **policy object** is an identifiable **state**, **action** or **resource** that is potentially
3071 constrained by the **policy**.

3072 **Permission**

3073 A **permission** constraint governs the ability of an **actor** to perform an **action** or
3074 to enter some specified **state**.

3075 Note that **permissions** are distinct from **ability** and from **authority**. **Authority** refers to
3076 the legitimate nature of an **action**, whereas **permission** refers to the **right** to perform
3077 the **action**.

3078 **Obligation**

3079 An **obligation** constraint governs the requirement that a **participant** or other
3080 **actor** should perform an **action** or maintain some specified **state**.

3081 For example, once the **service consumer** and provider have entered into an
3082 agreement to provide and consume a service, both **participants** incur **obligations**: the
3083 consumer is obligated to pay for the service and the provider is obligated to provide the
3084 service.

3085 **Obligations** to maintain state may range from a requirement to maintain a minimum
3086 balance on an account through a requirement that a service provider 'remember' that a
3087 particular **service consumer** is logged in.

3088 **Obligations** and **permissions** have a positive form and a negative form. A positive
3089 **permission** refers to something that a **policy subject** may do, a negative **permission**
3090 refers to something the **policy subject** may not do.

3091 These definitions are replicated from Section 3.1.4.

3092 **4.4.2 Policy and Contract Enforcement**

3093 The enforcement of **policy constraints** has to address two core problems: how to
3094 enforce the atomic policy constraints, and how to enforce combinations of **policy**
3095 **constraints**. In addition, it is necessary to address the resolution of **policy conflicts**.

3096 **4.4.2.1 Enforcing Simple Policy Constraints**

3097 The two primary kinds of **policy constraint** – **permission** and **obligation** – naturally
3098 lead to different styles of enforcement. A **permission** constraint must typically be
3099 enforced *prior* to the **policy subject** invoking the **policy object**. On the hand, an
3100 **obligation** constraint must typically be enforced post-facto through some form of
3101 auditing process and remedial action.

3102 For example, if a communications policy required that all communication be encrypted,
3103 this is enforceable at the point of communication: any attempt to communicate a
3104 message that is not encrypted can be blocked.

3105 Similarly, an **obligation** to pay for services rendered is enforced by ensuring that
3106 payment arrives within a reasonable period of time. Invoices are monitored for prompt
3107 (or lack of) payment.

3108 The key concepts in enforcing both forms of **policy constraint** are the **policy decision**
3109 and the **policy enforcement**.

3110 Policy Decision

3111 A **policy decision** is a determination as to whether a given **policy constraint** is
3112 satisfied or not.

3113 A **policy decision** is effectively a **measurement** of some state – typically a portion of
3114 the SOA ecosystem’s **shared state**. This implies a certain *timeliness* in the measuring:
3115 a measurement that is too early or is too late does not actually help in determining if the
3116 **policy constraint** is satisfied appropriately.

3117 Policy Enforcement

3118 A **policy enforcement** is the use of a mechanism to limit the behavior and/or
3119 state of **policy subjects** to comply with a **policy decision**.

3120 A **policy enforcement** implies the use of some mechanism to ensure compliance with a
3121 **policy decision**. The range of mechanisms is completely dependent on the kinds of
3122 atomic **policy constraints** that the **policy framework** may support. As noted above,
3123 the two primary styles of constraint – **permission** and **obligation** – will lead to different
3124 styles of enforcement.

3125 4.4.2.2 Enforcing Policy Combinations

3126 Enforcing policy combinations is primarily an elaboration of enforcing simple **policy**
3127 **constraints**. The process of **policy decisions** is enhanced to allow a measurement to
3128 involve combinations of **policy constraints** and the process of **policy enforcement**
3129 may need to be enhanced to coordinate the enforcement of multiple **policy constraints**
3130 simultaneously.

3131 4.4.2.3 Conflict Resolution

3132 Whenever it is possible that more than one **policy constraint** applies in a given
3133 situation, there is the potential that the **policies** themselves are not mutually consistent.
3134 For example, a policy that requires communication to be encrypted and a policy that
3135 requires an administrator to read every communication are likely to be in conflict with
3136 each other – the two policies cannot both be satisfied.

3137 In general, with sufficiently rich **policy frameworks**, it is not possible to always resolve
3138 policy conflicts automatically. However, a reasonable approach is to augment the **policy**
3139 **decision** process with simple **policy conflict resolution** rules; with the potential for
3140 *escalating* a **policy conflict** to human adjudication.

3141 Policy Conflict

3142 A **policy conflict** exists between two or more **policies** in a **policy decision**
3143 process if it is not possible to satisfy all the **policies** that apply.

3144 Policy Conflict Resolution

3145 A **policy conflict resolution** rule is a way of determining which **policy** should
3146 prevail in a **policy conflict**.

3147 The inevitable consequence of **policy conflicts** is that it is not possible to guarantee
3148 that all **policies** are satisfied at all times. This, in turn, implies a certain *flexibility* in the
3149 application of **policy constraints**: they will not always be honored.

3150 4.4.3 Architectural Implications

3151 The key choices that must be made in a system of **policies** center around the **policy**
3152 **framework** and **policy enforcement** mechanisms

- 3153 • There SHOULD be a standard **policy framework** that is adopted across the SOA
3154 ecosystem:
 - 3155 ○ This framework MUST permit the expression of simple **policy constraints**
 - 3156 ○ This framework MAY allow (to a varying extent) the combination of **policy**
3157 **constraints**, including
 - 3158 • Both positive and negative constraints
 - 3159 • Conjunctions and disjunctions of constraints
 - 3160 • The quantification of constraints
 - 3161 ○ The framework MUST at least allow the **policy subject** and the **policy object**
3162 to be identified as well as the **policy constraint**.
 - 3163 ○ The framework MAY allow further structuring of **policies** into modules,
3164 inheritance between **policies** and so on.
- 3165 • There SHOULD be mechanisms that facilitate the application of **policies**:
 - 3166 ○ There SHOULD be mechanisms that allow **policy decisions** to be made,
3167 consistent with the **policy frameworks** and with the state of the SOA
3168 ecosystem.
 - 3169 ○ There SHOULD be mechanisms to enforce policy decisions
 - 3170 • There SHOULD be mechanisms to support the **measurement** of
3171 whether certain **policy constraints** are satisfied or not, or to what
3172 degree they are satisfied.
 - 3173 • Such enforcement mechanisms MAY include support for both
3174 **permission-style constraints** and **obligation-style constraints**.
 - 3175 • Enforcement mechanisms MAY support the simultaneous enforcement
3176 of multiple **policy constraints** across multiple points in the SOA
3177 ecosystem.
 - 3178 ○ There SHOULD be mechanisms to resolve **policy conflicts**
 - 3179 • This MAY involve escalating **policy conflicts** to human adjudication.
 - 3180 ○ There SHOULD be mechanisms that support the management and
3181 promulgation of **policies**.

5 Owning Service Oriented Architectures View

*Governments are instituted among Men,
deriving their just power from the consent of the governed*
American Declaration of Independence

The *Owning Service Oriented Architectures View* focuses on the issues, requirements and **responsibilities** involved in owning a SOA-based system.

Owning a SOA-based system raises significantly different challenges to owning other complex systems -- such as Enterprise suites -- because there are strong limits on the control and **authority** of any one party when a system spans multiple ownership domains.

Even when a SOA-based system is deployed internally within an organization, there are multiple internal **stakeholders** involved and there may not be a simple hierarchy of control and management. Thus, an early consideration of how multiple boundaries affect SOA-based systems will provide a firm foundation for dealing with them in whatever form they are found rather than debating whether the boundaries should exist.

This view focuses on the Governance of SOA-based systems, on the security challenges involved in running a SOA-based system and the management challenges.

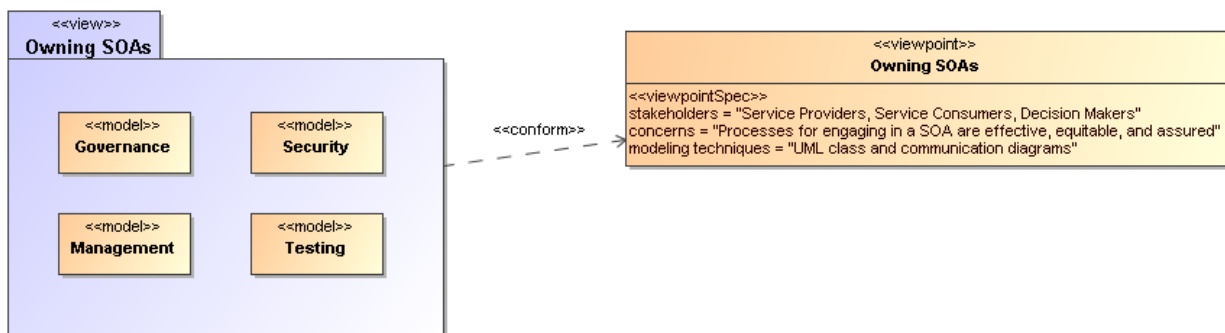


Figure 47 Model Elements Described in the Owning Service Oriented Architectures View

The following subsections present models of these functions.

5.1 Governance Model

The Reference Model defines Service Oriented Architecture as an architectural paradigm for organizing and utilizing distributed capabilities that may be under the control of different ownership domains **[SOA-RM]**. Consequently, it is important that organizations that plan to engage in service interactions adopt governance policies and procedures sufficient to ensure that there is standardization across both internal and external organizational boundaries to promote the effective creation and use of SOA-based services.

3211 5.1.1 Understanding Governance

3212 5.1.1.1 Terminology

3213 Governance is about making decisions that are aligned with the overall organizational
3214 strategy and culture of the enterprise. **[Gartner]** It specifies the decision rights and
3215 accountability framework to encourage desirable behaviors **[Weill/Ross-MIT Sloan**
3216 **School]** towards realizing the strategy and defines incentives (positive or negative)
3217 towards that end. It is less about overt control and strict adherence to rules, and more
3218 about guidance and effective and equitable usage of resources to ensure sustainability
3219 of an organization's strategic objectives. **[TOGAF v8.1]**

3220 To accomplish this, governance requires organizational structure and processes and
3221 must identify who has authority to define and carry out its mandates. It must address
3222 the following questions: 1) what decisions must be made to ensure effective
3223 management and use?, 2) who should make these decisions?, and 3) how will these
3224 decisions be made and monitored? , and (4) how will these decisions be
3225 communicated? The intent is to achieve goals, add value, and reduce risk.

3226 Within a single ownership domain such as an enterprise, generally there is a hierarchy
3227 of governance structures. Some of the more common enterprise governance structures
3228 include corporate governance, technology governance, IT governance, and architecture
3229 governance **[TOGAF v8.1]**. These governance structures can exist at multiple levels
3230 (global, regional, and local) within the overall enterprise.

3231 It is often asserted that SOA governance is a specialization of IT governance as there is
3232 a natural hierarchy of these types of governance structures; however, the focus of SOA
3233 governance is less on decisions to ensure effective management and use of IT as it is
3234 to ensure effective management and use of SOA-based systems. Certainly, SOA
3235 governance must still answer the basic questions also associated with IT governance,
3236 i.e., who should make the decisions, and how these decisions will be made and
3237 monitored.

3238 5.1.1.2 Relationship to Management

3239 There is often confusion centered on the relationship between governance and
3240 management. As described earlier, governance is concerned with decision making.
3241 Management, on the other hand, is concerned with execution. Put another way,
3242 governance describes the world as leadership wants it to be; management executes
3243 activities that intends to make the leadership's desired world a reality. Where
3244 governance determines who has the authority and responsibility for making decisions
3245 and the establishment of guidelines for how those decisions should be made,
3246 management is the actual process of making, implementing, and measuring the impact
3247 of those decisions **[Loeb]**. Consequently, governance and management work in
3248 concert to ensure a well-balanced and functioning organization as well as an ecosystem
3249 of inter-related organizations. In the sections that follow, we elaborate further on the
3250 relationship between governance and management in terms of setting and enforcing
3251 service policies, contracts, and standards as well as addressing issues surrounding
3252 regulatory compliance.

3253 5.1.1.3 Why is SOA Governance Important?

3254 One of the hallmarks of SOA that distinguishes it from other architectural paradigms for
3255 distributed computing is the ability to provide a uniform means to offer, discover, interact

3256 with and use capabilities (as well the ability to compose new capabilities from existing
3257 ones) all in an environment that transcends domains of **ownership**. Consequently,
3258 **ownership**, and issues surrounding it, such as obtaining acceptable terms and
3259 conditions (T&Cs) in a contract, is one of the primary topics for SOA governance.
3260 Generally, IT governance does not include T&Cs, for example, as a condition of use as
3261 its primary concern.

3262 Just as other architectural paradigms, technologies, and approaches to IT are subject to
3263 change and evolution, so too is SOA. Setting policies that allow change management
3264 and evolution, establishing strategies for change, resolving disputes that arise, and
3265 ensuring that SOA-based systems continue to fulfill the goals of the business are all
3266 reasons why governance is important to SOA.

3267 **5.1.1.4 Governance Stakeholders and Concerns**

3268 As noted in Section **Error! Reference source not found.** the **participants** in a service
3269 interaction include the service provider, the **service consumer**, and other interested or
3270 unintentional third parties. Depending on the circumstances, it may also include the
3271 owners of the underlying capabilities that the SOA services access. Governance must
3272 establish the policies and rules under which duties and **responsibilities** are defined
3273 and the expectations of **participants** are grounded. The expectations include
3274 transparency in aspects where transparency is mandated, trust in the impartial and
3275 consistent application of governance, and assurance of reliable and robust behavior
3276 throughout the SOA ecosystem.

3277 **5.1.2 A Generic Model for Governance**

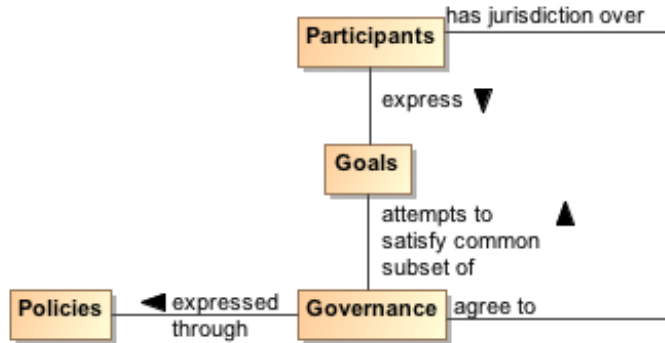
3278 **Governance**

3279 Governance is the prescribing of conditions and constraints consistent with
3280 satisfying common goals and the structures and processes needed to define and
3281 respond to actions taken towards realizing those goals.

3282 The following is a generic model of governance represented by segmented models that
3283 begin with motivation and proceed through measuring compliance. It is not meant to be
3284 an all-encompassing treatise on governance but a focused subset that captures the
3285 aspects necessary to describe governance for SOA. It is also not meant to imply that
3286 practical application of governance is a single, isolated instance of these models; in fact,
3287 there are likely hierarchical chains of governance that apply and possibly parallel chains
3288 that govern different aspects or focus on different goals. This is discussed further in
3289 section 5.1.2.5. The defined models are simultaneously applicable to each of the
3290 overlapping instances.

3291 A given **enterprise** may already have portions of these models in place. To a large
3292 extent, the models shown here are not specific to SOA; discussions on direct
3293 applicability begin in section 5.1.3.

3294 **5.1.2.1 Motivating Governance**



3295
3296 *Figure 48 Motivating governance model*

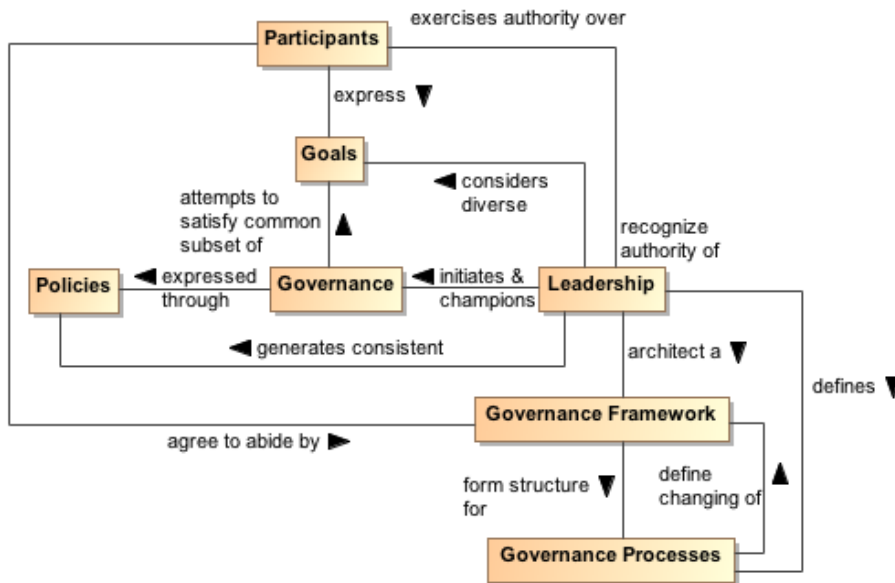
3297 An organizational domain such as an **enterprise** is made up of **participants** who may
 3298 be individuals or groups of individuals forming smaller organizational units within the
 3299 **enterprise**. The overall business strategy should be consistent with the Goals of the
 3300 **participants**; otherwise, the business strategy would not provide value to the
 3301 **participants** and governance towards those ends becomes difficult if not impossible.
 3302 This is not to say that an instance of governance will simultaneously satisfy all the goals
 3303 of all the **participants**; rather, the goals of any governance instance must sufficiently
 3304 satisfy a useful subset of each **participant's** goals so as to provide value and ensure
 3305 the cooperation of all the **participants**.

3306 A policy is the formal characterization of the conditions and constraints that governance
 3307 deems as necessary to realize the goals which it is attempting to satisfy. Policy may
 3308 identify required conditions or actions or may prescribe limitations or other constraints
 3309 on permitted conditions or actions. For example, a policy may prescribe that
 3310 safeguards must be in place to prevent unauthorized access to sensitive material. It
 3311 may also prohibit use of computers for activities unrelated to the specified work
 3312 assignment. Policy is made operational through the promulgating and implementing of
 3313 Rules and Regulations (as defined in section 5.1.2.3).

3314 As noted in section 4.4.2, policy may be asserted by any **participant** or on behalf of the
 3315 **participant** by its organization. Part of the **purpose** of governance is to arbitrate
 3316 among diverse goals of **participants** and diverse policies articulated to realize those
 3317 goals. The intent is to form a consistent whole that allows governance to minimize
 3318 ambiguity about its **purpose**. While resolving all ambiguity would be an ideal, it is
 3319 unlikely that all inconsistencies will be identified and resolved before governance
 3320 becomes operational.

3321 For governance to have effective jurisdiction over **participants**, there must be some
 3322 degree of agreement by each **participant** that it will abide by the governance
 3323 mandates. A minimal degree of agreement often presages **participants** who “slow-roll”
 3324 if not actively reject complying with Policies that express the specifics of governance.

3325 **5.1.2.2 Setting Up Governance**



3326
3327 *Figure 49 Setting up governance model*

3328 **Leadership**

3329 Leadership is the entity who has the responsibility and authority to generate
3330 consistent policies through which the goals of governance can be expressed and
3331 to define and champion the structures and processes through which governance
3332 is realized.

3333 **Governance Framework**

3334 The Governance Framework is a set of organizational structures that enable
3335 governance to be consistently defined, clarified, and as needed, modified to
3336 respond to changes in its domain of concern.

3337 **Governance Processes**

3338 Governance Processes are the defined set of activities that are performed within
3339 the Governance Framework to enable the consistent definition, application, and
3340 as needed, modification of Rules that organize and regulate the activities of
3341 **participants** for the fulfillment of expressed policies. (See section 5.1.2.3 for
3342 elaboration on the relationship of Governance Processes and Rules.)

3343 As noted earlier, governance requires an appropriate organizational structure and
3344 identification of who has governance **authority** to make governance decisions. In Figure 49, the
3345 entity with governance **authority** is designated the Leadership. This is someone,
3346 possibly one or more of the **participants**, that **participants** recognize as having
3347 **authority** for a given **purpose** or over a given set of issues or concerns.

3348 The Leadership is responsible for prescribing or delegating a working group to prescribe
3349 the Governance Framework that forms the structure for Governance Processes which
3350 define how governance is to be carried out. This does not itself define the specifics of
3351 how governance is to be applied, but it does provide an unambiguous set of procedures

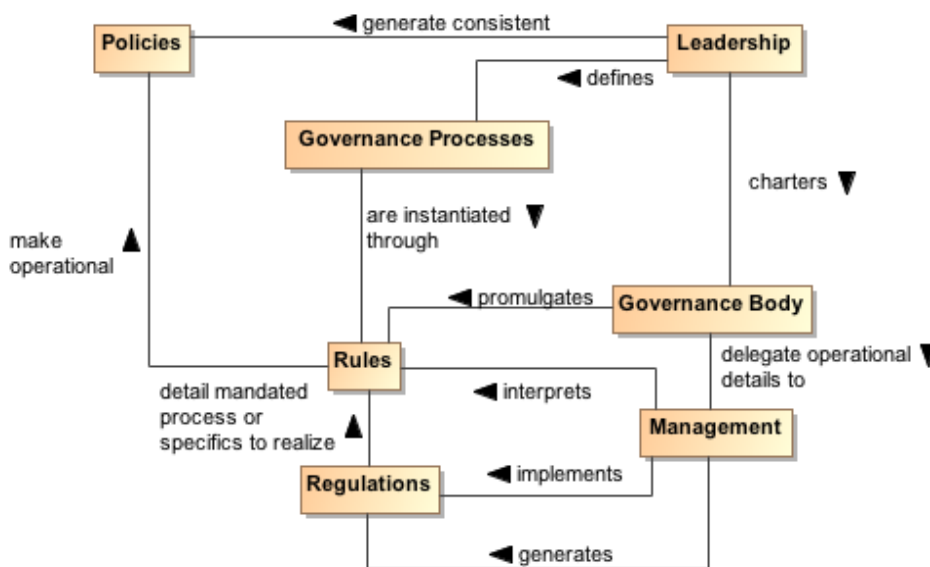
3352 that should ensure consistent actions which **participants** agree are fair and account for
3353 sufficient input on the subjects to which governance will be applied.

3354 The **participants** may be part of the working group that codifies the Governance
3355 Framework and Processes. When complete, the **participants** must acknowledge and
3356 agree to abide by the products generated through application of this structure.

3357 The Governance Framework and Processes are often documented in the charter of a
3358 body created or designated to oversee governance. This is discussed further in the
3359 next section. Note that the Governance Processes should also include those necessary
3360 to modify the Governance Framework itself.

3361 An important function of Leadership is not only to initiate but also be the consistent
3362 champion of governance. Those responsible for carrying out governance mandates
3363 must have Leadership who makes it clear to **participants** that expressed Policies are
3364 seen as a means to realizing established goals and that compliance with governance is
3365 required.

3366 5.1.2.3 Carrying Out Governance



3367
3368 *Figure 50 Carrying out governance model*

3369 **Rule**

3370 A Rule is a prescribed guide for carrying out activities and processes leading to
3371 desired results, e.g. the operational realization of policies.

3372 **Regulation**

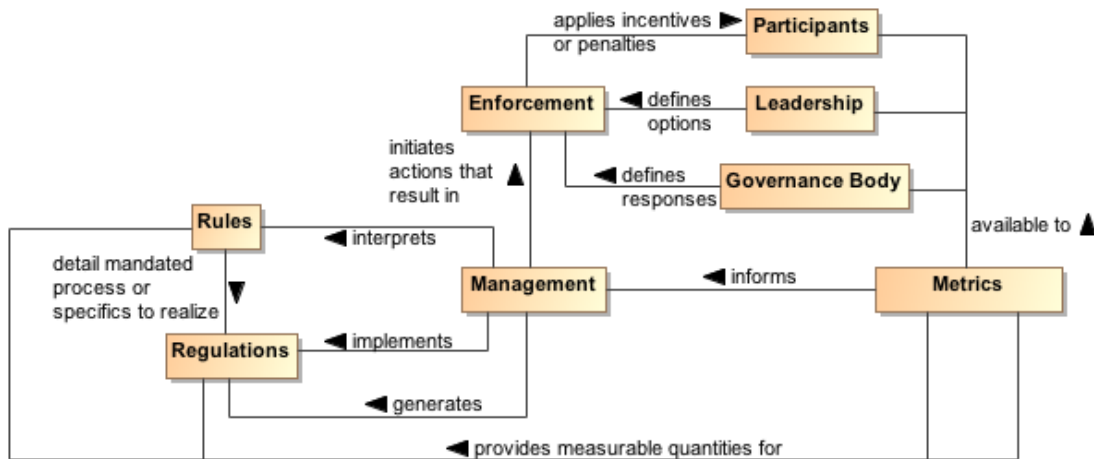
3373 A Regulation is a mandated process or the specific details that derive from the
3374 interpretation of Rules and lead to measurable quantities against which
3375 compliance can be measured.

3376 To carry out governance, Leadership charts a Governance Body to promulgate the
3377 Rules needed to make the Policies operational. The Governance Body acts in line with
3378 Governance Processes for its rule-making process and other functions. Whereas
3379 Governance is the setting of Policies and defining the Rules that provide an operational

3380 context for Policies, the operational details of governance are likely delegated by the
 3381 Governance Body to Management. Management generates Regulations that specify
 3382 details for Rules and other procedures to implement both Rules and Regulations. For
 3383 example, Leadership could set a Policy that all authorized parties should have access to
 3384 data, the Governance Body would promulgate a Rule that PKI certificates are required
 3385 to establish identity of authorized parties, and Management can specify a Regulation of
 3386 who it deems to be a recognized PKI issuing body. In summary, Policy is a predicate to
 3387 be satisfied and Rules prescribe the activities by which that satisfying occurs. A number
 3388 of rules may be required to satisfy a given policy; the carrying out of a rule may
 3389 contribute to several policies being realized.

3390 Whereas the Governance Framework and Processes are fundamental for having
 3391 **participants** acknowledge and commit to compliance with governance, the Rules and
 3392 Regulations provide operational constraints which may require resource **commitments**
 3393 or other levies on the **participants**. It is important for **participants** to consider the
 3394 framework and processes to be fair, unambiguous, and capable of being carried out in a
 3395 consistent manner and to have an opportunity to formally accept or ratify this situation.
 3396 Rules and Regulations, however, do not require individual acceptance by any given
 3397 **participant** although some level of community comment is likely to be part of the
 3398 Governance Processes. Having agreed to governance, the **participants** are bound to
 3399 comply or be subject to prescribed mechanisms for enforcement.

3400 **5.1.2.4 Ensuring Governance Compliance**



3401
 3402 *Figure 51 Ensuring governance compliance model*

3403 Setting Rules and Regulations does not ensure effective governance unless compliance
 3404 can be measured and Rules and Regulations can be enforced. Metrics are those
 3405 conditions and quantities that can be measured to characterize actions and results.
 3406 Rules and Regulations MUST be based on collected Metrics or there will be no way for
 3407 Management to assess compliance. The Metrics are available to the **participants**, the
 3408 Leadership, and the Governance Body so what is measured and the results of
 3409 measurement are clear to everyone.

3410 The Leadership in its relationship with **participants** will have certain options that can be
 3411 used for Enforcement. A common option may be to effect future funding. The
 3412 Governance Body defines specific enforcement responses, such as what degree of

3413 compliance is necessary for full funding to be restored. It is up to Management to
3414 identify compliance shortfalls and to initiate the Enforcement process.

3415 Note, enforcement does not strictly need to be negative consequences. Management
3416 can use Metrics to identify exemplars of compliance and Leadership can provide
3417 options for rewarding the **participants**. It is likely the Governance Body that defines
3418 awards or other incentives.

3419 **5.1.2.5 Considerations for Multiple Governance Chains**

3420 As noted in section 5.1.2, instances of the governance model often occur as a tiered
3421 arrangement, with governance at some level delegating specific **authority** and
3422 responsibility to accomplish a focused portion of the original level's mandate. For
3423 example, a corporation may encompass several lines of business and each line of
3424 business governs its own affairs in a manner that is consistent with and contributes to
3425 the goals of the parent organization. Within the line of business, an IT group may be
3426 given the mandate to provide and maintain IT resources, giving rise to IT governance.

3427 In addition to tiered governance, there are likely to be multiple governance chains
3428 working in parallel. For example, a company making widgets likely has policies intended
3429 to ensure they make high quality widgets and make an impressive profit for their
3430 shareholders. On the other hand, Sarbanes-Oxley is a parallel governance chain in the
3431 United States that specifies how the management must handle its accounting and
3432 information that needs to be given to its shareholders. The parallel chains may just be
3433 additive or may be in conflict and require some harmonization.

3434 Being distributed and representing different ownership domains, a SOA **participant** is
3435 likely under the jurisdiction of multiple governance domains simultaneously and may
3436 individually need to resolve consequent conflicts. The governance domains may
3437 specify precedence for governance conformance or it may fall to the discretion of the
3438 **participant** to decide on the course of actions they believe appropriate.

3439 **5.1.3 Governance Applied to SOA**

3440 **5.1.3.1 Where SOA Governance is Different**

3441 SOA governance is often discussed in terms of IT governance, but rather than a parent-
3442 child relationship, Figure 52 shows the two as siblings of the general governance
3443 described in section 5.1.2. There are obvious dependencies and a need for coordination
3444 between the two, but the idea of aligning IT with business already demonstrates that
3445 resource providers and resource consumers must be working towards common goals if
3446 they are to be productive and efficient. While SOA governance will be shown to be
3447 active in the area of infrastructure, it is a specialized concern for having a dependable
3448 platform to support service interaction; a host of traditional IT issues is considered to be
3449 out of scope. A SOA governance plan for an **enterprise** will not resolve shortcomings
3450 with the enterprise IT governance.

3451 Governance in the context of SOA is that organization of services: that promotes their
3452 visibility; that facilitates interaction among service **participants**; and that directs that the
3453 results of service interactions are those **real world effects** as described within the
3454 service description and constrained by policies and contracts as assembled in the
3455 execution context.

3456 SOA governance must specifically account for control across different ownership
3457 domains, i.e. all the **participants** may not be under the jurisdiction of a single
3458 governance authority. However, for governance to be effective, the **participants** must
3459 agree to recognize the **authority** of the Governance Body and must operate within the
3460 Governance Framework and through the Governance Processes so defined.

3461 SOA governance must account for interactions across **ownership boundaries**, which
3462 likely also implies across enterprise governance boundaries. For such situations,
3463 governance emphasizes the need for agreement that some Governance Framework
3464 and Governance Processes have jurisdiction, and the governance defined must satisfy
3465 the Goals of the **participants** for cooperation to continue. A standards development
3466 organization such as OASIS is an example of voluntary agreement to governance over
3467 a limited domain to satisfy common goals.

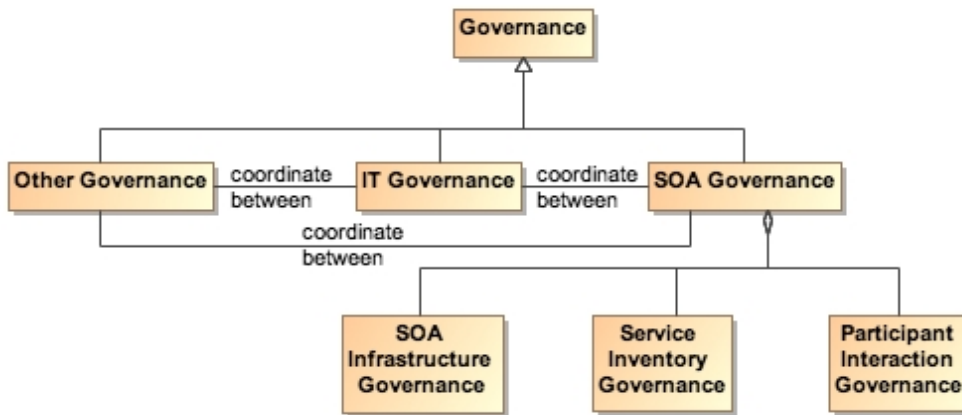
3468 The specifics discussed in the figures in the previous sections are equally applicable to
3469 governance across **ownership boundaries** as it is within a single boundary. There is a
3470 charter agreed to when **participants** become members of the organization, and this
3471 charter sets up the structures and processes that will be followed. Leadership may be
3472 shared by the leadership of the overall organization and the leadership of individual
3473 groups themselves chartered per the Governance Processes. There are
3474 Rules/Regulations specific to individual efforts for which **participants** agree to local
3475 goals, and Enforcement can be loss of voting rights or under extreme circumstances,
3476 expulsion from the group.

3477 Thus, the major difference for SOA governance is an appreciation for the cooperative
3478 nature of the enterprise and its reliance on furthering common goals if productive
3479 participation is to continue.

3480 **5.1.3.2 What Must be Governed**

3481 An expected benefit of employing SOA principles is the ability to quickly bring
3482 **resources** to bear to deal with unexpected and evolving situations. This requires a
3483 great deal of confidence in the underlying capabilities that can be accessed and in the
3484 services that enable the access. It also requires considerable flexibility in the ways
3485 these **resources** can be employed. Thus, SOA governance requires establishing
3486 confidence and trust while instituting a solid framework that enables flexibility, indicating
3487 a combination of strict control over a limited set of foundational aspects but minimum
3488 constraints beyond those bounds.

3489



3490
3491 *Figure 52 Relationship among types of governance*

3492 SOA governance applies to three aspects of service definition and use:

- 3493 • SOA infrastructure – the “plumbing” that provides utility functions that enable and
- 3494 support the use of the service
- 3495 • Service inventory – the requirements on a service to permit it to be accessed
- 3496 within the infrastructure
- 3497 • Participant interaction – the consistent expectations with which all **participants**
- 3498 are expected to comply

3499 **5.1.3.2.1 Governance of SOA Infrastructure**

3500 The SOA infrastructure is likely composed of several families of SOA services that
3501 provide access to fundamental computing business services. These include, among
3502 many others, services such as messaging, security, storage, discovery, and mediation.
3503 The provisioning of an infrastructure on which these services may be accessed and the
3504 general realm of those contributing as utility functions of the infrastructure are a
3505 traditional IT governance concern. In contrast, the focus of SOA governance is how the
3506 existence and use of the services enables the SOA ecosystem.

3507 By characterizing the environment as containing families of SOA services, the
3508 assumption is that there may be multiple approaches to providing the business services
3509 or variations in the actual business services provided. For example, discovery could be
3510 based on text search, on metadata search, on approximate matches when exact
3511 matches are not available, and numerous other variations. The underlying
3512 implementation of search algorithms are not the purview of SOA governance, but the
3513 access to the resulting service infrastructure enabling discovery must be stable, reliable,
3514 and extremely robust to all operating conditions. Such access enables other
3515 specialized SOA services to use the infrastructure in dependable and predictable ways,
3516 and is where governance is important.

3517 **5.1.3.2.2 Governance of the Service Inventory**

3518 Given an infrastructure in which other SOA services can operate, a key governance
3519 issue is which SOA services to allow in the ecosystem. The major concern SHOULD be
3520 a definition of well-behaved services, where the required behavior will likely inherit their

3521 characteristics from experiences with distributed computing but will also evolve with
3522 SOA experience. A major requirement for ensuring well-behaved services is collecting
3523 sufficient metrics to know how the service affects the SOA infrastructure and whether it
3524 complies with established infrastructure policies.

3525 Another common concern of service approval is whether there will be duplication of
3526 function by multiple services. Some governance models talk to a tightly controlled
3527 environment where a primary concern is to avoid any service duplication. Other
3528 governance models talk to a market of services where the consumers have wide
3529 choices. For the latter, it is anticipated that the better services will emerge from market
3530 consensus and the availability of alternatives will drive innovation.

3531 It is likely that some combination of control and openness will emerge, possibly with a
3532 different appropriate balance for different categories of use. For SOA governance, the
3533 issue is less which services are approved but rather ensuring that sufficient description
3534 is available to support informed decisions for appropriate use. Thus, SOA governance
3535 SHOULD concentrate on identifying the required attributes to adequately describe a
3536 service, the required target values of the attributes, and the standards for defining the
3537 meaning of the attributes and their target values. Governance may also specify the
3538 processes by which the attribute values are measured and the corresponding
3539 certification that some realized attribute set may imply.

3540 For example, unlimited access for using a service may require a degree of life cycle
3541 maturity that has demonstrated sufficient testing over a certain size community.
3542 Alternately, the policy may specify that a service in an earlier phase of its life cycle may
3543 be made available to a smaller, more technically sophisticated group in order to collect
3544 the metrics that would eventually allow the service to advance its life cycle status.

3545 This aspect of governance is tightly connected to description because, given a well-
3546 behaved set of services, it is the responsibility of the consumer (or policies promulgated
3547 by the consumer's organization) to decide whether a service is sufficient for that
3548 consumer's intended use. The goal is to avoid global governance specifying criteria that
3549 are too restrictive or too lax for the local needs of which global governance has little
3550 insight.

3551 Such an approach to specifying governance allows independent domains to describe
3552 services in local terms while still having the services available for informed use across
3553 domains. In addition, changes to the attribute sets within a domain can be similarly
3554 described, thus supporting the use of newly described **resources** with the existing ones
3555 without having to update the description of all the legacy content.

3556 **5.1.3.2.3 Governance of Participant Interaction**

3557 Finally, given a reliable services infrastructure and a predictable set of services, the
3558 third aspect of governance is prescribing what is required during a service interaction.

3559 Governance would specify adherence to service interface and service reachability
3560 parameters and would require that the result of an interaction **MUST** correspond to the
3561 **real world effects** as contained in the service description. Governance would ensure
3562 preconditions for service use are satisfied, in particular those related to security aspects
3563 such as user authentication, authorization, and non-repudiation. If conflicts arise,

3564 governance would specify resolution processes to ensure appropriate agreements,
3565 policies, and conditions are met.

3566 It would also rely on sufficient monitoring by the SOA infrastructure to ensure services
3567 remain well-behaved during interactions, e.g. do not use excessive resources or exhibit
3568 other prohibited behavior. Governance would also require that policy agreements as
3569 documented in the execution context for the interaction are observed and that the
3570 results and any after effects are consistent with the agreed policies. It is likely that in
3571 this area the governance will focus on more contractual and legal aspects rather than
3572 the precursor descriptive aspects. SOA governance may prescribe the processes by
3573 which SOA-specific policies are allowed to change, but there are likely more business-
3574 specific policies that will be governed by processes outside SOA governance.

3575 **5.1.3.3 Overarching Governance Concerns**

3576 There are numerous governance related concerns whose effects span the three areas
3577 just discussed. One is the area of standards, how these are mandated, and how the
3578 mandates may change. The Web Services standards stack is an example of relevant
3579 standards where a significant number are still under development. In addition, while
3580 there are notional scenarios that guide what standards are being developed, the fact
3581 that many of these standards do not yet exist precludes operational testing of their
3582 adequacy or effectiveness as a necessary and sufficient set.

3583 That said, standards are critical to creating a SOA ecosystem where SOA services can
3584 be introduced, used singularly, and combined with other services to deliver complex
3585 business functionality. As with other aspects of SOA governance, the Governance
3586 Body should identify the minimum set felt to be needed and rigorously enforce that that
3587 set be used where appropriate. The Governance Body must take care to expand and
3588 evolve the mandated standards in a predictable manner and with sufficient technical
3589 guidance that new services will be able to coexist as much as possible with the old, and
3590 changes to standards do not cause major disruptions.

3591 Another area that may see increasing activity as SOA expands will be additional
3592 regulation by governments and associated legal institutions. New laws are likely that will
3593 deal with transactions which are service based, possibly including taxes on the
3594 transactions. Disclosures laws are likely to mandate certain elements of description so
3595 both the consumer and provider act in a predictable environment and are protected from
3596 ambiguity in intent or **action**. Such laws are likely to spawn rules and regulations that
3597 will influence the metrics collected for evaluation of compliance.

3598 **5.1.3.4 Considerations for SOA Governance**

3599 The Reference Architecture definition of a loosely coupled system is one in which the
3600 constraints on the interactions between components is minimal: sufficient to permit
3601 interoperation without additional constraints that may be an artifact of implementation
3602 technology. While governance experience for standalone systems provides useful
3603 guides, we must be careful not to apply constraints that would preclude the flexibility,
3604 agility, and adaptability we expect to realize from a SOA ecosystem.

3605 One of the strengths of SOA is it can make effective use of diversity rather than
3606 requiring monolithic solutions. Heterogeneous organizations can interact without
3607 requiring each conforms to uniform tools, representation, and processes. However, with

3608 this diversity comes the need to adequately define those elements necessary for
3609 consistent interaction among systems and **participants**, such as which communication
3610 protocol, what level of security, which vocabulary for payload content of messages. The
3611 solution is not always to lock down these choices but to standardize alternatives and
3612 standardize the representations through which an unambiguous identification of the
3613 alternative chosen can be conveyed. For example, the URI standard specifies the URI
3614 string, including what protocol is being used, what is the target of the message, and how
3615 may parameters be attached. It does not limit the available protocols, the semantics of
3616 the target address, or the parameters that can be transferred. Thus, as with our
3617 definition of loose coupling, it provides absolute constraints but minimizes which
3618 constraints it imposes.

3619 There is not a one-size-fits-all governance but a need to understand the types of things
3620 governance will be called on to do in the context of the goals of SOA. It is likely that
3621 some communities will initially desire and require very stringent governance policies and
3622 procedures while other will see need for very little. Over time, best practices will evolve,
3623 likely resulting in some consensus on a sensible minimum and, except in extreme cases
3624 where it is demonstrated to be necessary, a loosening of strict governance toward the
3625 best practice mean.

3626 A question of how much governance may center on how much time governance
3627 activities require versus how quickly is the system being governed expected to respond
3628 to changing conditions. For large single systems that take years to develop, the
3629 governance process could move slowly without having a serious negative impact. For
3630 example, if something takes two years to develop and the steps involved in governance
3631 take two months to navigate, then the governance can go along in parallel and may not
3632 have a significant impact on system response to changes. Situations where it takes as
3633 long to navigate governance requirements as it does to develop a response are
3634 examples where governance may need to be reevaluated as to whether it facilitates or
3635 inhibits the desired results. Thus, the speed at which services are expected to appear
3636 and evolve needs to be considered when deciding the processes for control. The
3637 added weight of governance should be appropriate for overall goals of the application
3638 domain and the service environment.

3639 Governance, as with other aspects of any SOA implementation, should start small and
3640 be conceptualized in a way that keeps it flexible, scalable, and realistic. A set of useful
3641 guidelines would include:

- 3642 • Do not hardwire things that will inevitably change. For example, develop a
3643 system that uses the representation of policies rather than code the policies into
3644 the implementations.
- 3645 • Avoid setting up processes that demo well for three services without considering
3646 how it will work for 300. Similarly, consider whether the display of status and
3647 activity for a small number of services will also be effective for an operator in a
3648 crisis situation looking at dozens of services, each with numerous, sometimes
3649 overlapping and sometimes differing activities.
- 3650 • Maintain consistency and realism. A service solution responding to a natural
3651 disaster cannot be expected to complete a 6-week review cycle but be effective
3652 in a matter of hours.

3653 5.1.4 Architectural Implications of SOA Governance

3654 The description of SOA governance indicates numerous architectural requirements on
3655 the SOA ecosystem:

- 3656 • Governance is expressed through policies and assumes multiple use of focused
3657 policy modules that can be employed across many common circumstances. This
3658 requires the existence of:
 - 3659 ○ descriptions to enable the policy modules to be visible, where the
3660 description includes a unique identifier for the policy and a sufficient, and
3661 preferably a machine process-able, representation of the meaning of
3662 terms used to describe the policy, its functions, and its effects;
 - 3663 ○ one or more discovery mechanisms that enable searching for policies that
3664 best meet the search criteria specified by the service **participant**; where
3665 the discovery mechanism will have access to the individual policy
3666 descriptions, possibly through some repository mechanism;
 - 3667 ○ accessible storage of policies and policy descriptions, so service
3668 **participants** can access, examine, and use the policies as defined.
- 3669 • Governance requires that the **participants** understand the intent of governance,
3670 the structures created to define and implement governance, and the processes to
3671 be followed to make governance operational. This requires the existence of:
 - 3672 ○ an information collection site, such as a Web page or portal, where
3673 governance information is stored and from which the information is always
3674 available for access;
 - 3675 ○ a mechanism to inform **participants** of significant governance **events**,
3676 such as changes in policies, rules, or regulations;
 - 3677 ○ accessible storage of the specifics of Governance Processes;
 - 3678 ○ SOA services to access automated implementations of the Governance
3679 Processes
- 3680 • Governance policies are made operational through rules and regulations. This
3681 requires the existence of:
 - 3682 ○ descriptions to enable the rules and regulations to be visible, where the
3683 description includes a unique identifier and a sufficient, and preferably a
3684 machine process-able, representation of the meaning of terms used to
3685 describe the rules and regulations;
 - 3686 ○ one or more discovery mechanisms that enable searching for rules and
3687 regulations that may apply to situations corresponding to the search
3688 criteria specified by the service **participant**; where the discovery
3689 mechanism will have access to the individual descriptions of rules and
3690 regulations, possibly through some repository mechanism;
 - 3691 ○ accessible storage of rules and regulations and their respective
3692 descriptions, so service **participants** can understand and prepare for
3693 compliance, as defined.
 - 3694 ○ SOA services to access automated implementations of the Governance
3695 Processes.

- 3696
- 3697
- 3698
- 3699
- Governance implies management to define and enforce rules and regulations. Management is discussed more specifically in section **Error! Reference source not found.**, but in a parallel to governance, management requires the existence of:
 - an information collection site, such as a Web page or portal, where management information is stored and from which the information is always available for access;
 - a mechanism to inform **participants** of significant management **events**, such as changes in rules or regulations;
 - accessible storage of the specifics of processes followed by management.
 - Governance relies on metrics to define and measure compliance. This requires the existence of:
 - the infrastructure monitoring and reporting information on SOA resources;
 - possible interface requirements to make accessible metrics information generated or most easily accessed by the service itself.

3711 5.2 Security Model

3712 Security is one aspect of confidence – the confidence in the integrity, reliability, and
3713 confidentiality of the system. In particular, security focuses on those aspects of
3714 assurance that involve the accidental or malign intent of other people to damage or
3715 compromise trust in the system and on the availability of SOA-based systems to
3716 perform desired capability.

3717 Security

3718 Security concerns the set of mechanisms for ensuring and enhancing trust and
3719 confidence in the SOA ecosystem.

3720 Providing for security for Service Oriented Architecture is somewhat different than for
3721 other contexts; although many of the same principles apply equally to SOA and to other
3722 systems. The fact that SOA embraces crossing **ownership boundaries** makes the
3723 issues involved with moving data more visible.

3724 As well as securing the movement of data within and across **ownership boundaries**,
3725 security often revolves around **resources**: the need to guard certain resources against
3726 inappropriate access – whether reading, writing or otherwise manipulating those
3727 resources.

3728 Any comprehensive security solution must take into account the people that are using,
3729 maintaining and managing the SOA. Furthermore, the relationships between them must
3730 also be incorporated: any security assertions that may be associated with particular
3731 interactions originate in the people that are behind the interaction.

3732 We analyze security in terms of the **social structures** that define the legitimate
3733 permissions, **obligations** and **roles** of people in relation to the system, and
3734 mechanisms that must be put into place to realize a secure system. The former are
3735 typically captured in a series of security policy statements; the latter in terms of security
3736 *guards* that ensure that policies are enforced.

3737 How and when to apply these derived security policy mechanisms is directly associated
3738 with the assessment of the *threat model* and a *security response model*. The threat
3739 model identifies the kinds of threats that directly impact the message and/or application
3740 of constraints, and the response model is the proposed mitigation to those threats.
3741 Properly implemented, the result can be an acceptable level of risk to the safety and
3742 integrity of the system.

3743 **5.2.1 Secure Interaction Concepts**

3744 We can characterize secure interactions in terms of key security concepts [ISO/IEC
3745 **27002**]: confidentiality, integrity, authentication, authorization, non-repudiation, and
3746 availability. The concepts for secure interactions are well defined in other standards
3747 and publications. The security concepts here are not defined but rather related to the
3748 SOA ecosystem perspective of this reference architecture foundation.

3749 **5.2.1.1 Confidentiality**

3750 Confidentiality concerns the protection of privacy of **participants** in their interactions.
3751 Confidentiality refers to the assurance that unauthorized entities are not able to read
3752 messages or parts of messages that are transmitted.

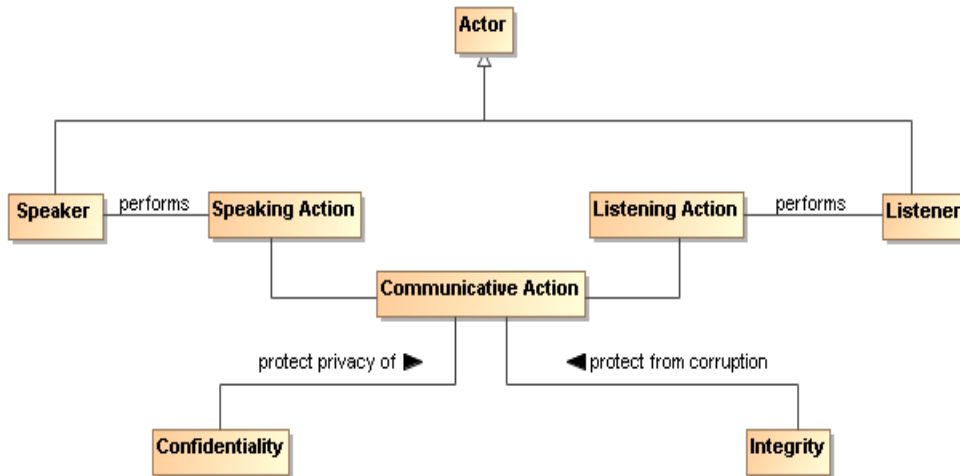
3753 Note that confidentiality has degrees: in a completely confidential exchange, third
3754 parties would not even be aware that a confidential exchange has occurred. In a
3755 partially confidential exchange, the identities of the **participants** may be known but the
3756 content of the exchange obscured.

3757 **5.2.1.2 Integrity**

3758 Integrity concerns the protection of information that is exchanged – either from
3759 unauthorized writing or inadvertent corruption. Integrity refers to the assurance that
3760 information that has been exchanged has not been altered.

3761 Integrity is different from confidentiality in that messages that are sent from one
3762 **participant** to another may be obscured to a third party, but the third party may still be
3763 able to introduce his own content into the exchange without the knowledge of the
3764 **participants**.

3765 Figure 53 applies confidentiality and integrity to **communicative action**.



3766


3767 *Figure 53 Confidentiality and Integrity*

3768 A **communicative action** is a joint action involved in the exchange of messages.
 3769 Section 5.2.4 describes common computing techniques for providing confidentiality and
 3770 integrity during message exchanges.

3771 5.2.1.3 Authentication

3772 Authentication concerns the identity of the **participants** in an exchange. Authentication
 3773 refers to the means by which one **participant** can be assured of the identity of other
 3774 **participants**.

3775 Figure 54 applies authentication to the identity of **participants**.

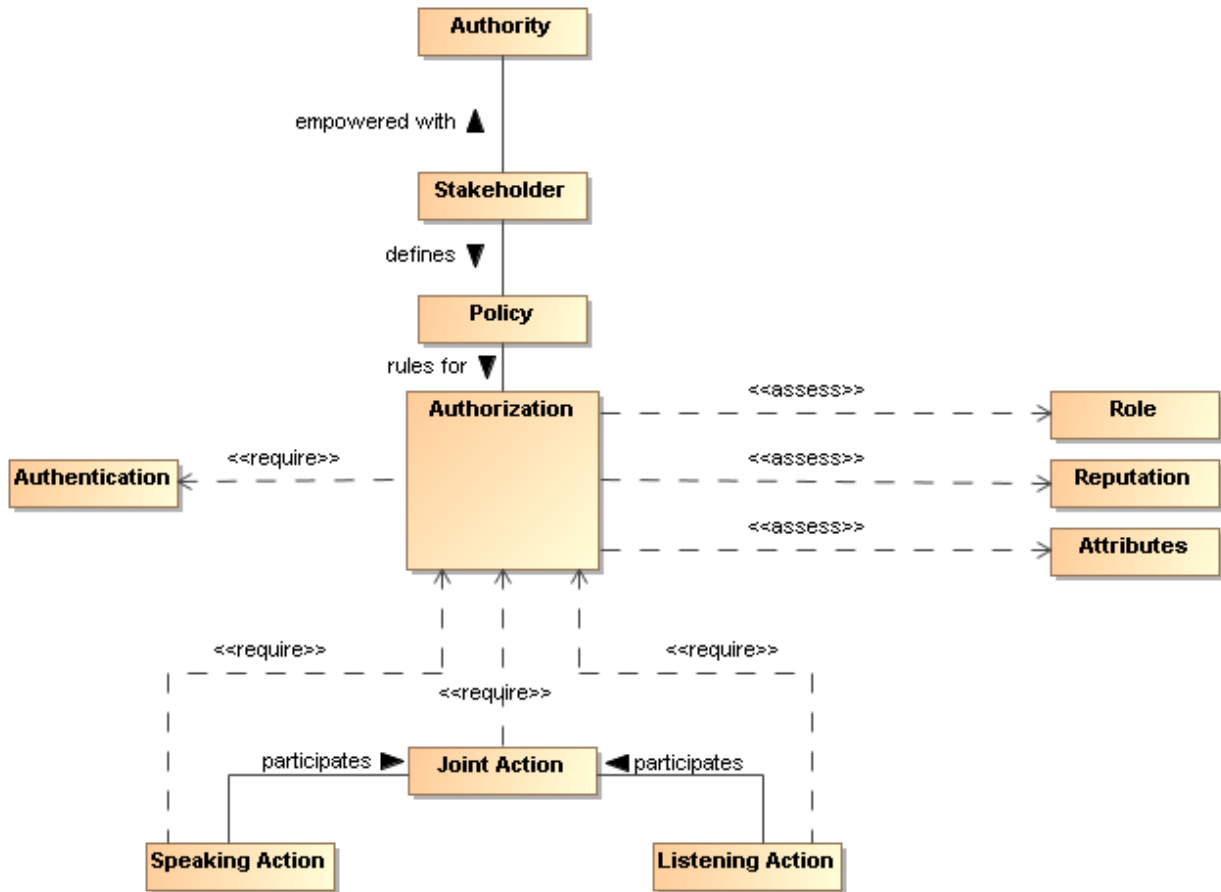


3777

3778 *Figure 54 Authentication*

3779 5.2.1.4 Authorization

3780 Authorization concerns the legitimacy of the interaction. Authorization refers to the
 3781 means by which a **stakeholder** may be assured that the information and actions that
 3782 are exchanged are either explicitly or implicitly approved.



3783
3784 *Figure 55 Authorization*

3785 The **roles** and attributes which provide a **participant's** credentials are expanded to
3786 include reputation. Reputation often helps determine willingness to interact, for
3787 example, reviews of a service provider are likely to influence the decision to interact with
3788 the service provider. The **roles**, reputation, and attributes are represented as
3789 assertions measured by authorization decision points.

3790 The role of policy for security is to permit **stakeholders** to express their choices. In
3791 Figure 55, a policy is a written constraint and the role, reputation, and attribute
3792 assertions are evaluated according to the constraints in the authorization policy. A
3793 combination of security mechanisms and their control via explicit policies can form the
3794 basis of an authorization solution.

3795 **5.2.1.5 Non-repudiation**

3796 Non-repudiation concerns the accountability of **participants**. To foster trust in the
3797 performance of a system used to conduct shared activities it is important that the
3798 **participants** are not able to later deny their actions: to repudiate them. Non-repudiation
3799 refers to the means by which a **participant** may not, at a later time, successfully deny
3800 having participated in the interaction or having performed the actions as reported by
3801 other **participants**.

3802 **5.2.1.6 Availability**

3803 Availability concerns the ability of systems to use and offer the services for which they
3804 were designed. One of the threats against availability is the so-called denial of service
3805 attack in which attackers attempt to prevent legitimate access to the system.

3806 We differentiate here between general availability – which includes aspects such as
3807 systems reliability – and availability as a security concept where we need to respond to
3808 active threats to the system.

3809 **5.2.2 Where SOA Security is Different**

3810 The core security concepts are fundamental to all social interactions. The evolution of
3811 sharing information using a SOA requires the flexibility to dynamically secure computing
3812 interactions in a computing ecosystem where the owning social groups, **roles**, and
3813 **authority** are constantly changing as described in section 5.1.3.1.

3814 SOA policy-based security can be more adaptive for a computing ecosystem than
3815 previous computing technologies allow for, and typically involves a greater degree of
3816 distributed mechanisms.

3817 Standards for security, as is the case with all aspects of SOA, play a large role in
3818 flexible security on a global scale. SOA security may also involve greater auditing and
3819 reporting to adhere to regulatory compliance established by governance structures.

3820 **5.2.3 Security Threats**

3821 There are a number of ways in which an attacker may attempt to compromise the
3822 security of a system. The two primary sources of attack are third parties attempting to
3823 subvert interactions between legitimate **participants** and an entity that is participating
3824 but attempting to subvert its partner(s). The latter is particularly important in a SOA
3825 where there may be multiple **ownership boundaries** and trust boundaries.

3826 The threat model lists some common threats that relate to the core security concepts
3827 listed in Section 5.2.1. Each technology choice in the realization of a SOA can
3828 potentially have many threats to consider.

3829 **Message alteration**

3830 If an attacker is able to modify the content (or even the order) of messages that
3831 are exchanged without the legitimate **participants** being aware of it then the
3832 attacker has successfully compromised the security of the system. In effect, the
3833 **participants** may unwittingly serve the needs of the attacker rather than their
3834 own.

3835 An attacker may not need to completely replace a message with his own to
3836 achieve his objective: replacing the identity of the beneficiary of a transaction
3837 may be enough.

3838 **Message interception**

3839 If an attacker is able to intercept and understand messages exchanged between
3840 **participants**, then the attacker may be able to gain advantage. This is probably
3841 the most commonly understood security threat.

3842 **Man in the middle**

3843 In a man-in-the-middle attack, the legitimate **participants** believe that they are
3844 interacting with each other; but are in fact interacting with the attacker. The
3845 attacker attempts to convince each **participant** that he is their correspondent;
3846 whereas in fact he is not.

3847 In a successful man-in-the-middle attack, legitimate **participants** will often not
3848 have a true understanding of the state of the other **participants**. The attacker
3849 can use this to subvert the intentions of the **participants**.

3850 **Spoofing**

3851 In a spoofing attack, the attacker convinces a **participant** that he is really
3852 someone else – someone that the **participant** would normally trust.

3853 **Denial of service attack**

3854 In a denial of service (DoS) attack, the attacker attempts to prevent legitimate
3855 users from making use of the service. A DoS attack is easy to mount and can
3856 cause considerable harm: by preventing legitimate interactions, or by slowing
3857 them down enough, the attacker may be able to simultaneously prevent
3858 legitimate access to a service and to attack the service by another means.

3859 A variation of the DoS attack is the Distributed Denial of Service attack. In a
3860 DDoS attack the attacker uses multiple agents to the attack the target. In some
3861 circumstances this can be extremely difficult to counteract effectively.

3862 One of the features of a DoS attack is that it does not require valid interactions to
3863 be effective: responding to invalid messages also takes resources and that may
3864 be sufficient to cripple the target.

3865 **Replay attack**

3866 In a replay attack, the attacker captures the message traffic during a legitimate
3867 interaction and then replays part of it to the target. The target is persuaded that a
3868 similar transaction to the previous one is being repeated and it will respond as
3869 though it were a legitimate interaction.

3870 A replay attack may not require that the attacker understand any of the individual
3871 communications; the attacker may have different objectives (for example
3872 attempting to predict how the target would react to a particular request).

3873 **False repudiation**

3874 In false repudiation, a user completes a normal transaction and then later
3875 attempts to deny that the transaction occurred. For example, a customer may
3876 use a service to buy a book using a credit card; then, when the book is delivered,
3877 refuse to pay the credit card bill claiming that *someone else* must have ordered
3878 the book.

3879 **5.2.4 Security Responses**

3880 Security goals are never absolute: it is not possible to guarantee 100% confidentiality,
3881 non-repudiation, etc. However, a well designed and implemented security response
3882 model can ensure acceptable levels of security risk. For example, using a well-designed

3883 cipher to encrypt messages may make the cost of breaking communications so great
3884 and so lengthy that the information obtained is valueless.

3885 Performing threat assessments, devising mitigation strategies, and determining
3886 acceptable levels of risk are the foundation for an effective process to mitigating threats
3887 in a cost-effective way.¹⁸ The choice in hardware and software to realize a SOA will be
3888 the basis for threat assessments and mitigation strategies. The **stakeholders** of a
3889 specific SOA implementation should determine acceptable levels of risk based on threat
3890 assessments and the cost of mitigating those threats.

3891 **5.2.4.1 Privacy Enforcement**

3892 The most efficient mechanism to assure confidentiality is the encryption of information.
3893 Encryption is particularly important when messages must cross trust boundaries;
3894 especially over the Internet. Note that encryption need not be limited to the content of
3895 messages: it is possible to obscure even the existence of messages themselves
3896 through encryption and 'white noise' generation in the communications channel.

3897 The specifics of encryption are beyond the scope of this architecture. However, we are
3898 concerned about how the connection between privacy-related policies and their
3899 enforcement is made.

3900 A policy enforcement point for enforcing privacy may take the form of an automatic
3901 function to encrypt messages as they leave a trust boundary; or perhaps simply
3902 ensuring that such messages are suitably encrypted.

3903 Any policies relating to the level of encryption being used would then apply to these
3904 centralized messaging functions.

3905 **5.2.4.2 Integrity Protection**

3906 To protect against message tampering or inadvertent message alteration, and to allow
3907 the receiver of a message to authenticate the sender, messages may be accompanied
3908 by a digital signature. Digital signatures provide a means to detect if signed data has
3909 been altered. This protection can also extend to authentication and non-repudiation of a
3910 sender.

3911 A common way a digital signature is generated is with the use of a private key that is
3912 associated with a public key and a digital certificate. The private key of some entity in
3913 the system is used to create a digital signature for some set of data. Other entities in the
3914 system can check the integrity of the signed data set via signature verification
3915 algorithms. Any changes to the data that was signed will cause signature verification to
3916 fail, which indicates that integrity of the data set has been compromised.

3917 A party verifying a digital signature must have access to the public key that corresponds
3918 to the private key used to generate the signature. A digital certificate contains the public

¹⁸ In practice, there are perceptions of security from all participants regardless of ownership boundaries. Satisfying security policy often requires asserting sensitive information about the message initiator. The perceptions of this participant about information privacy may be more important than actual security enforcement within the SOA for this stakeholder.

3919 key of the owner, and is itself protected by a digital signature created using the private
3920 key of the issuing Certificate Authority (CA).

3921 **5.2.4.3 Message Replay Protection**

3922 To protect against replay attacks, messages may contain information that can be used
3923 to detect replayed messages. The simplest requirement to prevent replay attacks is that
3924 each message that is ever sent is unique. For example, a message may contain a
3925 message ID, a timestamp, and the intended destination.

3926 By storing message IDs, and comparing each new message with the store, it becomes
3927 possible to verify whether a given message has been received before (and therefore
3928 should be discarded).

3929 The timestamp may be included in the message to help check for message freshness.
3930 Messages that arrive after their message ID could have been cleared (after receiving
3931 the same message some time previously) may also have been replayed. A common
3932 means for representing timestamps is a useful part of an interoperable replay detection
3933 mechanism.

3934 The destination information is used to determine if the message was misdirected or
3935 replayed. If the replayed message is sent to a different endpoint than the destination of
3936 the original message, the replay could go undetected if the message does not contain
3937 information about the intended destination.

3938 In the case of messages that are replies to prior messages, it is also possible to include
3939 seed information in the prior messages that is randomly and uniquely generated for
3940 each message that is sent out. A replay attack can then be detected if the reply does
3941 not embed the random number that corresponds to the original message.

3942 **5.2.4.4 Auditing and Logging**

3943 False repudiation involves a **participant** denying that it authorized a previous
3944 interaction. An effective strategy for responding to such a denial is to maintain careful
3945 and complete logs of interactions which can be used for auditing **purposes**. The more
3946 detailed and comprehensive an audit trail is, the less likely it is that a false repudiation
3947 would be successful.

3948 The countermeasures assume that the non-repudiation tactic (e.g. digital signatures) is
3949 not undermined itself. For example, if private key is stolen and used by an adversary,
3950 even extensive logging cannot assist in rejecting a false repudiation.

3951 Unlike many of the security responses discussed here, it is likely that the scope for
3952 automation in rejecting a repudiation attempt is limited to careful logging.

3953 **5.2.4.5 Graduated engagement**

3954 The key to managing and responding to DoS attacks is to be careful in the use of
3955 **resources** when responding to interaction. Put simply, a system has a choice to
3956 respond to a communication or to ignore it. In order to avoid vulnerability to DoS attacks
3957 a service provider should be careful not to commit **resources** beyond those implied by
3958 the current state of interactions; this permits a graduation in commitment by the service
3959 provider that mirrors any **commitment** on the part of **service consumers** and attackers
3960 alike.

3961 5.2.5 Architectural Implications of SOA Security

3962 Providing SOA security in an ecosystem of governed services has the following
3963 implications on the policy support and the distributed nature of mechanisms used to
3964 assure SOA security:

- 3965 • Security expressed through policies have the same architectural implications as
3966 described in Section 4.4.3 for policies and contracts architectural implications.
- 3967 • Security policies require mechanisms to support security description
3968 administration, storage, and distribution.
- 3969 • Service descriptions supporting security policies should:
 - 3970 ○ have a meta-structure sufficiently rich to support security policies;
 - 3971 ○ be able to reference one or more security policy artifacts;
 - 3972 ○ have a framework for resolving conflicts between security policies.
- 3973 • The mechanisms that make-up the execution context in secure SOA-based
3974 systems should:
 - 3975 ○ provide protection of the confidentiality and integrity of message
3976 exchanges;
 - 3977 ○ be distributed so as to provide centralized or decentralized policy-based
3978 identification, authentication, and authorization;
 - 3979 ○ ensure service availability to consumers;
 - 3980 ○ be able to scale to support security for a growing ecosystem of services;
 - 3981 ○ be able to support security between different communication technologies;
- 3982 • Common security services include:
 - 3983 ○ services that abstract encryption techniques;
 - 3984 ○ services for auditing and logging interactions and security violations;
 - 3985 ○ services for identification;
 - 3986 ○ services for authentication;
 - 3987 ○ services for authorization;
 - 3988 ○ services for intrusion detection and prevention;
 - 3989 ○ services for availability including support for quality of service
3990 specifications and metrics.

3991 5.3 Management Model

3992 Management

3993 Management is the control of the use, configuration, and availability of **resources**
3994 in accordance with the policies of the **stakeholders** involved.

3995 There are three separate but linked domains of interest within the management of SOA-
3996 based systems. The first and most obvious is the management and support of the
3997 **resources** that are involved in any complex system – of which SOA-based systems are
3998 excellent examples. The second is the promulgation and enforcement of the policies
3999 and contracts agreed to by the **stakeholders** in SOA-based systems. The third domain

4000 is the management of the relationships of the **participants** in SOA-based systems –
4001 both to each other and to the services that they use and offer.

4002 There are many artifacts in a large system that may need management. As soon as
4003 there is the possibility of more than one instance of a thing, the issue of managing those
4004 things becomes relevant. Historically, systems management capabilities have been
4005 organized by the following functional groups known as “FCAPS” functions (based on
4006 ITU-T Rec. M.3400 (02/2000), "TMN Management Functions"): Fault management,
4007 configuration management, account management, performance and security
4008 management.


4009 In the context of SOA we see many possible resources that may require management:
4010 services, service descriptions, service capabilities, policies, contracts, **roles**,
4011 relationships, security, and infrastructure elements. In addition, given the ecosystem
4012 nature of SOA, it is also potentially necessary to manage the business relationships
4013 between **participants** in the SOA.

4014 Managing systems that may be used across **ownership boundaries** raises issues that
4015 are not normally present when managing a system within a single ownership domain.
4016 For example, care is required managing a service when the owner of the service, the
4017 provider of the service, the host of the service and access mediators to the service may
4018 all belong to different **stakeholders**. In addition, it may be important to allow **service**
4019 **consumers** to communicate their requirements to the service provider so that they are
4020 satisfied in a timely manner.

4021 A given service may be provided and consumed in more than one version. Version
4022 control of services is important both for service providers and **service consumers** (who
4023 may need to ensure certainty in the version of the service they are interacting with).

4024 In fact, managing a service has quite a few similarities to using a service: suggesting
4025 that we can use the service oriented model to manage SOA-based systems as well as
4026 provide them. A management service would be distinguished from a non-management
4027 service more by the nature of the capabilities involved (i.e., capabilities that relate to
4028 managing services) than by any intrinsic difference.

4029 In this model, we show how the SOA framework may apply to managing services as
4030 well as using and offering them. There are, of course, some special considerations that
4031 apply to service management which we bring out: namely that we will be managing the
4032 life-cycle of services, managing any service level attributes, managing dependencies
4033 between services and so on.



4034
4035 *Figure 56 Managing resources in a SOA*

4036 The core concept in management is that of a manageability capability:

4037 **Manageability Capability**

4038 The manageability capability of a **resource** is the capability that allows it to be
4039 managed with respect to some property. Note that manageability capabilities are
4040 not necessarily part of the managed entities themselves.

4041 Manageability capabilities are the core **resources** that management systems use
4042 to manage: each **resource** that may be managed in some way has a number of
4043 aspects that may be managed. For example, a service's life-cycle may be
4044 manageable, as may its Quality of Service parameter; a policy may also be
4045 managed for life-cycle but Quality of Service would not normally apply.

4046 **Life-cycle manageability**

4047 A manageability capability associated with a **resource** that permits the life cycle
4048 of the **resource** to be managed. As noted above, the life-cycle manageability
4049 capability of a **resource** is unlikely to reside within the resource itself (you cannot
4050 tell a system that is not running to start itself).

4051 The life-cycle management of a **resource** typically refers to how the **resource** is
4052 created, how it is destroyed and what dependencies there might exist that must
4053 be simultaneously managed.

4054 **Configuration manageability**

4055 A capability that permits the configuration of **resources** to be managed. Service
4056 configuration, in particular, may be complex in cases where there are
4057 dependencies between services and other **resources**.

4058 **Event monitoring manageability**

4059 Managing the reporting of events and faults is one of the key lower-level
4060 manageability capabilities.

4061 **Accounting manageability**

4062 A capability associated with **resources** that allows for the use of those
4063 **resources** to be measured and accounted for. This implies that not only can the
4064 use of **resources** be properly measured, but also that those *using* those
4065 **resources** also be properly identified.

4066 Accounting for the use of **resources** by **participants** in the SOA supports the
4067 proper budgeting and allocation of funding by **participants**.

4068 **Quality of service manageability**

4069 A manageability capability associated with a **resource** that permits any quality of
4070 service associated with the **resource** to be managed. Classic examples of this
4071 include bandwidth requirements and offerings associated with a service.

4072 **Business performance manageability**

4073 A manageability capability that is associated with services that permits the
4074 service's business performance to be monitored and managed. In particular, if
4075 there are business-level service level agreements that apply to a service, being
4076 able to monitor and manage those SLAs is an important role for management
4077 systems.

4078 Building support for arbitrary business monitoring is likely to be challenging.
4079 However, given a *measure* for determining a service's compliance to business
4080 service level agreements, management systems can monitor that performance in
4081 a way that is entirely similar to other management tasks.

4082 **Policy manageability**

4083 Where the policies associated with a **resource** may be complex and dynamic, so
4084 those policies themselves may require management. The ability to manage those
4085 policies (such as promulgating policies, retiring policies and ensuring that policy
4086 decision points and enforcement points are current) is a management function.

4087 In the particular case of policies, there is a special relationship between
4088 management and policies. Just like other artifacts, policies require management
4089 in a SOA. However, much of management is about *applying* policies also: where
4090 governance is often about what the policies regarding artifacts and services
4091 should be, a key management role is to ensure that those policies are
4092 consistently applied.

4093 **Management service**

4094 A management service is a service that manages other services and **resources**.

4095 **Management Policy**

4096 A management policy is a policy whose topic is a management topic. Just as with
4097 other aspects of a SOA, the management of **resources** within the SOA may be
4098 governed by management policies, contracts (such as SLAs).

4099 In a deployed system, it may well be that different aspects of the management of a
4100 given service are managed by different management services. For example, the life-
4101 cycle management of services often involves managing dependencies between
4102 services and resource requirements. Managing quality of service is often very specific to
4103 the service itself; for example, quality of service attributes for a video streaming service
4104 are quite different to those for a banking system.

4105 There are additional concepts of management that often also apply to IT management:

4106 **Systems management**

4107 Systems management refers to enterprise-wide maintenance and administration
4108 of distributed computer systems.

4109 **Network management**

4110 Network management refers to the maintenance and administration of large-
4111 scale networks such as computer networks and telecommunication networks.
4112 Systems and network management execute a set of functions required for
4113 controlling, planning, deploying, coordinating, and monitoring the distributed
4114 computer systems and the resources of a network.

4115 However, for the purposes of this Reference Architecture, while recognizing their
4116 importance, we do not focus on systems management or network management.

4117 - the specific identifier is not prescribed by this Reference Architecture but the structure
4118 and semantics of the identifier must be indicated for the identifier value to be properly
4119 used. For example, part of identity may include version identification.

4120 For this, the configuration management plan or similar document from which the version
4121 number is derived must be identified.

4122

4123 **5.3.1 Management and Governance**

4124 The primary role of governance in the context of SOA is to allow the **stakeholders** in
4125 the SOA to be able to negotiate and set the key policies that govern the running of the
4126 system. Recall that in an ecosystems perspective, the goal is less to have complete
4127 fine-grained control but more to enable the individual **participants** to work together.
4128 Policies that are set at the governance of a SOA will tend to focus on the rules of
4129 engagement between **participants** – what kind of interacts are permissible, how to
4130 resolve disputes, and so on.

4131 While governance may be primarily focused on setting policies, management is more
4132 focused on realization and enforcement of policies.

4133 **5.3.2 Management Contracts and Policies**

4134 As we noted above, management can often be viewed as the application of contracts
4135 and policies to ensure the smooth running of the SOA. Policies play an important part
4136 in managing systems both as artifacts that need to be managed and as the guiding
4137 constraints to determine how the SOA should be managed.

4138 **5.3.2.1 Policies**

4139 "Although provision of management capabilities enables a service to become
4140 manageable, the extent and degree of permissible management are defined in
4141 management policies that are associated with the services. Management policies are
4142 used to define the **obligations** for, and permissions to, managing the service." **[WSA]**

4143 On the other hand, a policy without any means of enforcing it is vacuous. In the case of
4144 management policy, we rely on a management infrastructure to realize and enforce
4145 management policy.

4146 **5.3.3 Management Infrastructure**

4147 In order for a service or other **resource** to be manageable there must be a
4148 corresponding manageability capability that can effect that management. The
4149 particulars of this capability will vary somewhat depending on the nature of the
4150 capability. For example, a service life-cycle manageability capability requires the ability
4151 to start a service, to stop the service, and potentially to pause the service. Conversely,
4152 in order to manage document-like artifacts, such as service descriptions, the capability
4153 of storing the artifacts, controlling access to those artifacts, allowing updates of the
4154 artifacts to be deployed are all important capabilities for managing them.

4155

4156 Elements of a basic service management infrastructure should include the following
4157 characteristics:

4158

- 4159 • Integrate with existing security services

- 4160 • Monitoring
- 4161 • Heartbeat and Ping
- 4162 • Alerting
- 4163 • Pause/Restore/Restart Service Access
- 4164 • Logging, Auditing, Non-Repudiation
- 4165 • Runtime Version Management
- 4166 • Complement other infrastructure services (discovery, messaging, mediation)
- 4167
- 4168 * Message Routing and Redirection
- 4169 * Failover
- 4170 * Load-balancing
- 4171
- 4172 * QoS, Management of Service Level Objects and Agreements
- 4173 * Availability
- 4174 * Response Time
- 4175 * Throughput
- 4176
- 4177 • Fault and Exception Management
- 4178

4179 **5.3.4 Service Life-cycle**

4180 Managing a service's life cycle involves managing the establishment of the service,
 4181 managing its steady-state performance, and managing its termination. The most
 4182 obvious feature of this is that a service cannot manage its own life cycle (imagine asking
 4183 a non-functioning service to start). Another important consideration is that services may
 4184 have resource requirements that must be established at various points in the services'
 4185 life cycles. These dependencies may take the form of other services being established;
 4186 possibly even services that are not exposed by the service's own interface.

4189 **5.4 SOA Testing Model**

*Program testing can be used to show the presence of bugs,
 but never to show their absence!*

Edsger Dijkstra

4193 Testing for SOA combines the typical challenges of software testing and certification
 4194 with the additional needs of accommodating the distributed nature of the **resources**, the
 4195 greater access of a more unbounded consumer population, and the desired flexibility to
 4196 create new solutions from existing components over which the solution developer has
 4197 little if any control. The **purpose** of testing is to demonstrate a required level of
 4198 reliability, correctness, and effectiveness that enable prospective consumers to have

4199 adequate confidence in using a service. Adequacy is defined by the consumer based
4200 on the consumer's needs and context of use. As the Dijkstra quote points out, absolute
4201 correctness and completeness cannot be proven by testing; however, for SOA, it is
4202 critical for the prospective consumer to know what testing has been performed, how it
4203 has been performed, and what were the results.

4204 **5.4.1 Traditional Software Testing as Basis for SOA Testing**

4205 SOA services are largely software artifacts and can leverage the body of experience
4206 that has evolved around software testing. IEEE-829 specifies the basic set of software
4207 test documents while allowing flexibility for tailored use. As such, the document
4208 structure can also provide guidance to SOA testing.

4209 IEEE-829 covers test specification and test reporting through use of the following
4210 document types:

- 4211 • *Test plan* documenting the scope (what will be tested, both which entity and what
4212 features of the entity), the approach (how it will be tested), and the needed
4213 **resource** (who will do the testing, for how long), with details contained in the:
 - 4214 • *Test design specification*: features to be tested, test conditions (e.g. test cases,
4215 test procedures needed) and expected results (criteria for passing test); entrance
4216 and exit criteria
 - 4217 • *Test case specification*: test data used for input and expected output
 - 4218 • *Test procedure specification*: steps required to run the test, including any set-up
4219 preconditions
- 4220 • *Test item transmittal* to identify the test items being transmitted for testing
- 4221 • *Test log* to record what occurred during test, i.e. which tests run, who ran, what
4222 order, what happened
- 4223 • *Test incident report* to capture any **event** that happened during test which requires
4224 further investigation
- 4225 • *Test summary* as a management report summarizing test run and results,
4226 conclusions

4227 In summary, IEEE-829 captures (1) what was tested, (2) how it was tested, e.g. the test
4228 procedure used, and (3) the results of the test.

4229 **5.4.1.1 Types of Testing**

4230 There are numerous aspects of testing that, in total, work to establish that an entity is
4231 (1) built as required per policies and related specifications prescribed by the entity's
4232 owner, and (2) delivers the functionality required by its intended users. This is often
4233 referred to as verification and validation.

4234 Policies, as described in Section 4.4, that are related to testing may prescribe but are
4235 not limited to the business processes to be followed, the standards with which an
4236 implementation must comply, and the **qualifications** of and restrictions on the users. In
4237 addition to the functional requirements prescribing what an entity does, there may also
4238 be non-functional performance and/or quality metrics that state how well the entity does
4239 it. The relation of these policies to SOA testing is discussed further below.

4240 The identification of policies is the purview of governance (section 5.1) and the assuring
4241 of compliance (including response to noncompliance) with policies is a matter for
4242 management (section **Error! Reference source not found.**).

4243 **5.4.1.2 Range of Test Conditions**

4244 Test conditions and expected responses are detailed in the test case specification. The
4245 test conditions should be designed to cover the areas for which the entity's response
4246 must be documented and may include:

- 4247 • nominal conditions
- 4248 • boundaries and extremes of expected conditions
- 4249 • breaking point where the entity has degraded below a certain level or has
4250 otherwise ceased effective functioning
- 4251 • random conditions to investigate unidentified dependencies among combinations
4252 of conditions
- 4253 • errors conditions to test error handling

4254 The specification of how each of these conditions should be tested for SOA resources,
4255 including the infrastructure elements of the SOA ecosystem, is beyond the scope of this
4256 Reference Architecture but is an area that will evolve along with operational SOA
4257 experience.

4258 **5.4.1.3 Configuration Management of Test Artifacts**

4259 The test item transmittal provides an unambiguous identification of the entity being
4260 tested, thus REQUIRING that the configuration of the entity is appropriately tracked and
4261 documented. In addition, the test documents (such as those specified by IEEE-829)
4262 MUST also be under a documented and appropriately audited configuration
4263 management process, as should other **resources** used for testing. The description of
4264 each artifact would follow the general description model as discussed in section 4.1.1.1;
4265 in particular, it would include a version number for the artifact and reference to the
4266 documentation describing the versioning scheme from which the version number is
4267 derived.

4268

4269 [EDITOR'S NOTE: TO WHAT EXTENT SHOULD CM BE EXPLICITLY INCLUDED IN THE MANAGEMENT
4270 SECTION?]

4271 **5.4.2 Testing and the SOA Ecosystem**

4272 [EDITOR'S NOTE: THE EMPHASIS THOUGH MUCH OF THE RA IS THE LARGER ECOSYSTEM BUT WE NEED
4273 WORDS IN SECTION 3 TO ACKNOWLEDGE THE EXISTENCE OF THE ENTERPRISE AND THAT AN
4274 ENTERPRISE (AS COMMONLY INTERPRETED) IS LIKELY MORE CONSTRAINED AND MORE PRECISELY
4275 DESCRIBED FOR THE CONTEXT OF THE ENTERPRISE. THE ECOSYSTEM PERSPECTIVE, THOUGH, IS
4276 STILL APPLICABLE FOR THE FOLLOWING REASONS:

4277

- 4278 1. A GIVEN ENTERPRISE MAY COMPRISE NUMEROUS CONSTITUENT ENTERPRISES THAT
4279 RESEMBLE THE INDEPENDENT ENTITIES DESCRIBED FOR THE ECOSYSTEM. AN ENTERPRISE
4280 MAY ATTEMPT TO REDUCE VARIATIONS AMONG THE CONSTITUENTS BUT THE ECOSYSTEM VIEW
4281 ENABLES SOA TO BENEFIT THE ENTERPRISE WITHOUT REQUIRING THE ENTERPRISE ISSUES TO
4282 BE FULLY RESOLVED.

4283 2. RESOURCES SPECIFICALLY MOTIVATED BY THE CONTEXT OF THE ENTERPRISE CAN BE MORE
4284 READILY USED IN A DIFFERENT CONTEXT IF ECOSYSTEM CONSIDERATIONS ARE INCLUDED AT
4285 AN EARLY STAGE. THE CHANGE IN A CONTEXT MAY BE A FUNDAMENTAL CHANGE IN THE
4286 ENTERPRISE OR THE NEWLY DISCOVERED APPLICABILITY OF ENTERPRISE RESOURCES TO USE
4287 OUTSIDE THE ENTERPRISE.

4288
4289 IN THIS REFERENCE ARCHITECTURE, REFERENCE TO THE SOA ECOSYSTEM APPLIES BUT WITH
4290 POSSIBLY LESS GENERALITY TO AN ENTERPRISE USE OF SOA.]

4291 Testing of SOA artifacts for use in the SOA ecosystem differs from traditional software
4292 testing for several reasons. First, a highly touted benefit of SOA is to enable
4293 unanticipated consumers to make use of services for unanticipated purposes.
4294 Examples of this could include the consumer using a service for a result that was not
4295 considered the primary one by the provider, or the service may be used in combination
4296 with other services in a scenario that is different from the one considered when
4297 designing for the initial target consumer community. It is unlikely that a new consumer
4298 will push the services back to anything resembling the initial test phase to test the new
4299 use, and thus additional paradigms for testing are necessary. Some testing may
4300 depend on the availability of test resources made available as a service outside the
4301 initial test community, while some testing is likely to be done as part of limited use in the
4302 operational setting. The potential **responsibilities** related to such "consumer testing" is
4303 discussed further below.

4304 Secondly, in addition to consumers who interact with a service to realize the described
4305 **real world effects**, the developer community is also intended to be a consumer. In the
4306 SOA vision of reuse, the developer will compose new solutions using existing services,
4307 where the existing services provides access to some desired **real world effects** that are
4308 needed by the new solution. The new solution is a consumer of the existing services,
4309 enabling repeated interactions with the existing services playing the role of reusable
4310 components. Note, those components are used at the locations where they individually
4311 reside and are not typically duplicated for the new solution. The new solution may itself
4312 be offered as a SOA service, and a consumer of the service composition representing
4313 the new solution may be totally unaware of the component services being used. (See
4314 section 4.3.4 for further discussion on service compositions.)

4315 Another difference from traditional testing is that the distributed, unbounded nature of
4316 the SOA ecosystem makes it unlikely to have an isolated test environment that
4317 duplicates the operational environment. A traditional testing approach often makes use
4318 of a test system that is identical to the eventual operational system but isolated for
4319 testing. After testing is successfully completed, the tested entity would be migrated to
4320 the operational environment, or the test environment may be delivered as part of the
4321 system to become operational. This is not feasible for the SOA ecosystem as a whole.

4322 SOA services must be testable in the environment and under the conditions that can be
4323 encountered in the operational SOA ecosystem. As the ecosystem is in a state of
4324 constant change, so some level of testing is continuous through the lifetime of the
4325 service, leveraging utility services used by the ecosystem infrastructure to monitor its
4326 own health and respond to situations that could lead to degraded performance. This
4327 implies the test resources must incorporate aspects of the SOA paradigm, and a
4328 category of services may be created to specifically support and enable effective
4329 monitoring and continuous testing for **resources** participating in the SOA ecosystem.

4330 While SOA within an enterprise may represent a more constrained and predictable
4331 operational environment, the composability and unanticipated use aspects are highly
4332 touted within the enterprise. The expanded perspective on testing may not be as
4333 demanding within an enterprise but fuller consideration of the ecosystem enables the
4334 enterprise to be more responsive should conditions change.

4335 **5.4.3 Elements of SOA Testing**

4336 IEEE-829 identifies fundamental aspects of testing, and many of these should carry
4337 over to SOA testing: in particular, the identification of what is to be tested, how it is to be
4338 tested, and by whom the testing is to be done. While IEEE-829 identifies a suggested
4339 document tree, the availability of these documents in the SOA ecosystem is an
4340 additional matter of concern that will be discussed below.

4341 **5.4.3.1 What is to be Tested**

4342 The focus of this discussion is the SOA service. It is recognized that the infrastructure
4343 components of any SOA environment are likely to also be SOA services and, as such,
4344 will fall under the same testing guidance. Other resources that contribute to a SOA
4345 environment may not be SOA services, but will be expected to satisfy the intent if not
4346 the letter of guidance presented here. Specific differences for such **resources** are as
4347 yet largely undefined and further elaboration is beyond the scope of this Reference
4348 Architecture.

4349 The following discussion often focuses on a singular SOA service but it is implicit that
4350 any service may be a composite of other services. As such, testing the functionality of a
4351 composite service may effectively be testing an end-to-end business process that is
4352 being provided by the composite service. If new versions are available for the
4353 component services, appropriate end-to-end testing of the composite may be required
4354 in order to verify that the composite functionality is still adequately provided. The level
4355 of required testing of an updated composite will depend on policies of those providing
4356 the service, policies of those using the service, and mission criticality of those
4357 depending on the service results.

4358 The SOA service to be tested **MUST** be unambiguously identified as specified by its
4359 applicable configuration management scheme. Specifying such a scheme is beyond
4360 the scope of this Reference Architecture other than to say the scheme should be
4361 documented and itself under configuration management.

4362 **5.4.3.1.1 Origin of Test Requirements**

4363 In the Service Description model (Figure 21), the aspects of a service that need to be
4364 described are:

- 4365 • the service functionality and technical assumptions that underlie the functionality;
- 4366 • the policies that describe conditions of use;
- 4367 • the service interface that defines information exchange with the service;
- 4368 • service reachability that identifies how and where message exchange is to occur;
- 4369 and

4370 • metrics access for any **participant** to have information on how a service is
4371 performing.

4372 Service testing must provide adequate assurance that each of these aspects is
4373 operational as defined.

4374 The information in the service description comes from different sources. The
4375 functionality is defined through whatever process identifies needs and the community
4376 for which these needs will be addressed. The process may be ad hoc as serves the
4377 prospective service owner or strictly governed, but defining the functionality is an
4378 essential first step in development. It is also an early and ongoing focus of testing to
4379 ensure the service accurately reflects the described functionality and the described
4380 functionality accurately addresses the consumer needs.

4381 Policies define the conditions of development and conditions of use for a service and
4382 are typically specified as part of the governance process. Policies constraining service
4383 development, such as coding standards and best practices, require appropriate testing
4384 and auditing during development to ensure compliance. While the governance process
4385 will identify development policies, these are likely to originate from the technical
4386 community responsible for development activities. Policies that define conditions of use
4387 often define business practices that service owners and providers or those responsible
4388 for the SOA infrastructure want followed. These policies are initially tested during
4389 service development and are continuously monitored during the operational lifetime of
4390 the service.

4391 The testing of the service interface and service reachability are often related but
4392 essentially reflect different motivations and needs. The service interface is specified as
4393 a joint product of the service owners and providers who define service functionality, the
4394 prospective consumer community, the service developer, and the governance process.
4395 The semantics of the information model must align with the semantics of those who
4396 consume the service in order for there to be meaningful exchange of information. The
4397 structure of the information is influenced by the consumer semantics and the
4398 requirements and constraints of the representation as interpreted by the service
4399 developer. The service process model that defines actions which can be performed
4400 against a service and any temporal dependencies derive from the defined functionality
4401 and may be influenced by the development process. Any of these constraints may be
4402 identified and expressed as policy through the governance process.

4403 Service reachability conditions are the purview of the service provider who identifies the
4404 service endpoint and the protocols recognized at the endpoint. These may be
4405 constrained by governance decisions on how endpoint addresses may be allocated and
4406 what protocols should be used.

4407 While the considerations for defining the service interface derive from several sources,
4408 testing of the service interface is more straightforward and isolated in the testing
4409 process. At any point where the interface is modified or exposes a new **resource**, the
4410 message exchange should be monitored both to ensure the message reaches its
4411 intended destination and it is parsed correctly once received. Once an interface has
4412 been shown to function properly, it is unlikely it will fail later unless something
4413 fundamental to the service changes.

4414 The service interface is also tested when the service endpoint changes. Testing of the
4415 endpoint ensures message exchange can occur at the time of testing and the initial
4416 testing shows the interface is being processed properly at the new endpoint.
4417 Functioning of a service endpoint at one time does not guarantee it is functioning at
4418 another time, e.g. the server with the endpoint address may be down, making testing of
4419 service reachability a continual monitoring function through the life of the service's use
4420 of the endpoint. Also, while testing of the service endpoint is a necessary and most
4421 commonly noted part of the test regiment, it is not in itself sufficient to ensure the other
4422 aspects of testing discussed in this section.

4423 Finally, governance is impossible without the collection of metrics against which service
4424 behavior can be assessed. Metrics are also a key indicator for consumers to decide if a
4425 service is adequate for their needs. For instance, the average response time or the
4426 recent availability can be determining factors even if there are no rules or regulations
4427 promulgated through the governance process against which these metrics are
4428 assessed. The available metrics are a combination of those expected by the consumer
4429 community and those mandated through the governance process. The total set of
4430 metrics will evolve over time with SOA experience. Testing of the services that gather
4431 and provide access to the metrics will follow testing as described in this section, but for
4432 an individual service, testing will ensure that the metrics access indicated in the service
4433 description is accurate.

4434 The individual test requirements highlight aspects of the service that testing must
4435 consider but testing must establish more than isolated behavior. The emphasis is the
4436 holistic results of interacting with the service in the SOA environment. Recall that the
4437 execution context is the set of agreements between a consumer and a provider that
4438 define the conditions under which service interaction occurs. The agreements are
4439 expected to be predominantly the acceptance of the standard conditions as enumerated
4440 by the service provider, but it may include the identification of alternate conditions that
4441 will govern the interaction.

4442 For example, the provider may prefer a policy where it can sell the contact information
4443 of its consumers but will honor the request of a consumer to keep such information
4444 private. The identification of the alternate privacy policy is part of the execution context,
4445 and it is the application of and compliance with this policy that operational monitoring
4446 will attempt to measure. The collection of metrics showing this condition is indeed met
4447 when chosen is considered part of the ongoing testing of the service.

4448 Other variations in the execution context also require monitoring to ensure that different
4449 combinations of conditions perform together as desired. For example, if a new privacy
4450 policy takes additional **resources** to apply, this may affect quality of service and
4451 propagate other effects. These could not be tested during the original testing if the
4452 alternate policy did not exist at that time.

4453 **5.4.3.1.2 Testing Against Non-Functional Requirements**

4454 Testing against non-functional requirements constitutes testing of business usability of
4455 the service. In a marketplace of services, non-functional characteristics may be the
4456 primary differentiator between services that produce essentially the same **real world**
4457 **effects**.

4458 As noted in the previous section, non-functional characteristics are often associated
4459 with policies or other terms of use and may be collected in service level contracts
4460 offered by the service providers. Non-functional requirements may also reflect the
4461 network and hardware infrastructure that support communication with the service, and
4462 changes may impact quality of service. The **service consumer** and even the service
4463 provider may not be aware of all such infrastructure changes but the changes may
4464 manifest in shared states that impact the usability of the service.

4465 In general, a change in the non-functional requirements results in a change to the
4466 execution context, but as with any collection of information that constitutes a
4467 description, the execution context is unable to explicitly capture all non-functional
4468 requirements that may apply. A change in non-functional requirements, whether
4469 explicitly part of the execution context or an implicit contributor, may require retesting of
4470 the service even if its functionality and the implementation of the functionality has not
4471 changed. Depending on the circumstances, retesting may require a formal recertifying
4472 of end-to-end behavior or more likely will be part of the continuous monitoring that
4473 applies throughout the service lifetime.

4474 **5.4.3.1.3 Testing Content and the Interests of Consumers**

4475 As noted in section 5.4.1.1, testing may involve verification of conformance with respect
4476 to policies and technical specifications and validation with respect to sufficiency of
4477 functionality to meet some prescribed use. It may also include demonstration of
4478 performance and quality aspects. For some of these items, such as demonstrating the
4479 business processes followed in developing the service or the use of standards in
4480 implementing the service, the testing or relevant auditing is done internal to the service
4481 development process and follows traditional software testing and quality assurance. If it
4482 is believed of value to potential consumers, information about such testing could be
4483 included in the service description. However, it is not required that all test or
4484 compliance artifacts be available to consumers, as many of the details tested may be
4485 part of the opacity of the service implementation.

4486 Some aspects of the service being tested will reflect directly on the **real world effects**
4487 realized through interaction with the service. In these cases, it is more likely that testing
4488 results will be directly relevant to potential consumers. For example, if the service was
4489 designed to correspond to certain elements of a business process or that a certain
4490 workflow is followed, testing should verify that the **real world effects** reflect that the
4491 business process or workflow were satisfactorily captured.

4492 The testing may also need to demonstrate that specified conditions of use are satisfied.
4493 For example, policies may be asserted that require certain **qualifications** of or impose
4494 restrictions on the consumers who may interact with the service. The service testing
4495 must demonstrate that the service independently enforces the policies or it provides the
4496 required information exchanges with the SOA ecosystem so other **resources** can
4497 ensure the specified conditions.

4498 The completeness of the testing, both in terms of the features tested and the range of
4499 parameters for which response is tested, depends on the context of expected use: the
4500 more critical the use, the more complete the testing. There are always limits on the
4501 **resources** available for testing, if nothing else than the service must be available for
4502 use in a finite amount of time.

4503 This again emphasizes the need for adequate documentation to be available. If the
4504 original testing is very thorough, it may be adequate for less demanding uses in the
4505 future. If the original testing was more constrained, then well-documented test results
4506 establish the foundation on which further testing can be defined and executed.

4507 **5.4.3.2 How Testing is to be Done**

4508 Testing should follow well-defined methodologies and, if possible, should reuse test
4509 artifacts that have proven generally useful for past testing. For example, IEEE-829
4510 notes that test cases are separated from test designs to allow for use in more than one
4511 design and to allow for reuse in other situations. In the SOA ecosystem, description of
4512 such artifacts, as with description of a service, enables awareness of the item and
4513 describes how the artifact may be accessed or used.

4514 As with traditional testing, the specific test procedures and test case inputs are
4515 important so the tests are unambiguously defined and entities can be retested in the
4516 future. Automated testing and regression testing may be more important in the SOA
4517 ecosystem in order to re-verify a service is still acceptable when incorporated in a new
4518 use. For example, if a new use requires the services to deal with input parameters
4519 outside the range of initial testing, the tests could be rerun with the new parameters. If
4520 the testing resources are available to consumers within the SOA ecosystem, the testing
4521 as designed by test professionals could be consumed through a service accessed by
4522 consumers, and their results could augment those already in place. This is discussed
4523 further in the next section.

4524 **5.4.3.3 Who Performs the Testing**

4525 As with any software, the first line of testing is unit testing done by software developers.
4526 It is likely that initial testing will be done by those developing the software but may also
4527 be done independently by other developers. For SOA development, unit testing is likely
4528 confined to a development sandbox isolated from the SOA ecosystem.

4529 SOA testing will differ from traditional software testing in that testing beyond the
4530 development sandbox must incorporate aspects of the SOA ecosystem, and those
4531 doing the testing must be familiar with both the characteristics and responses of the
4532 ecosystem and the tools, especially those available as services, to facilitate and
4533 standardize testing. Test professionals will know what level of assurance must be
4534 established as the exposure of the service to the ecosystem and ecosystem to the
4535 service increases towards operational status. These test professionals may be internal
4536 resources to an organization or may evolve as a separate discipline provided through
4537 external contracting.

4538 As noted above, it is unlikely that a complete duplicate of the SOA ecosystem will be
4539 available for isolated testing, and thus use of ecosystem **resources** will manifest as a
4540 transition process rather than a step change from a test environment to an operational
4541 one. This is especially true for new composite services that incorporate existing
4542 operational services to achieve the new functionality. The test professionals will need to
4543 understand the available resources and the ramifications of this transition.

4544 As with current software development, a stage beyond work by test professionals will
4545 make use of a select group of typical users, commonly referred to as beta testers, to
4546 report on service response during typical intended use. This establishes fitness by the

4547 consumers, providing final validation of previously verified processes, requirements, and
4548 final implementation.

4549 In traditional software development, beta testing is the end of testing for a given version
4550 of the software. However, although the initial test phase can establish an appropriate
4551 level of confidence consistent with the designed use for the initial target consumer
4552 community, the operational service will exist in an evolving ecosystem, and later
4553 conditions of use may differ from those thought to be sufficient during the initial testing.
4554 Thus, operational monitoring becomes an extension of testing through the service
4555 lifetime. This continuous testing will attempt to ensure that a service does not consume
4556 an inordinate amount of ecosystem resources or display other behavior that degrades
4557 the ecosystem, but it will not undercover functional errors that may surface over time.

4558 As with any software, it is the responsibility of the consumers to consider the
4559 reasonableness of solutions in order to spot errors in either the software or the way the
4560 software is being used. This is especially important for consumers with unanticipated
4561 uses that may go beyond the original test conditions. It is unlikely the consumers will
4562 initiate a new round of formal testing unless the new use requires a significantly higher
4563 level of confidence in the service. Rather the consumer becomes a new extension to
4564 the testing regiment. Obvious testing would include a sanity check of results during the
4565 new use. However, if the details of legacy testing are associated with the service
4566 through the service description and if testing resources are available through automated
4567 testing services, then the new consumers can rerun and extend previous testing to
4568 include the extended test conditions. If the test results are acceptable, these can be
4569 added to the documentation of previous results and become the extended basis for
4570 future decisions by prospective consumers on the appropriateness of the service. If the
4571 results are not acceptable or in some way questionable, the responsible party for the
4572 service or testing professionals can be brought in to decide if remedial **action** is
4573 necessary.

4574 **5.4.3.4 How Testing Results are Reported**

4575 For any SOA service, an accurate reporting of the testing a service has undergone and
4576 the results of the testing is vital to consumers deciding whether a service is appropriate
4577 for intended use. Appropriateness may be defined by a consumer organization and
4578 require specific test regiments culminating in a certification; appropriateness could be
4579 established by accepting testing and certifications that have been conferred by others.

4580 The testing and certification information should be identified in the service description.
4581 Referring to the general description model of *Figure 24*, tests conducted by or under a
4582 request from the service owner (see **ownership** in section **Error! Reference source**
4583 **not found.**) would be captured under Annotations from Owners. Testing done by
4584 others, such as consumers with unanticipated uses, could be associated through
4585 Annotations from 3rd Parties. The annotations should clearly indicate what was tested,
4586 how the testing was done, who did the testing, and the testing results. The clear
4587 description of each of these artifacts and of standardized testing protocols for various
4588 levels of sophistication and completeness of testing would enable a common
4589 understanding and comparison of test coverage. It will also make it more
4590 straightforward to conduct and report on future testing, facilitating the maintenance of
4591 the service description.

4592 Consumer testing and the reporting of results raises additional issues. While stating
4593 who did the testing is mandatory, there may be formal requirements for authentication of
4594 the tester to ensure traceability of the testing claims. In some circumstances, persons
4595 or organizations would not be allowed to state testing claims unless the tester was an
4596 approved entity. In other cases, ensuring the tester had a valid email may be sufficient.
4597 In either case, it would be at the discretion of the potential consumer to decide what
4598 level of authentication was acceptable and which testers are considered authoritative in
4599 the context of their anticipated use.

4600 Finally, in a world of openly shared information, we would see an ever-expanding set of
4601 testing information as new uses and new consumers interact with a service. In reality,
4602 these new uses may represent proprietary processes or classified use that should only
4603 be available to authorized parties. Testing information, as with other elements of
4604 description, may require special access controls to ensure appropriate access and use.

4605 **5.4.4 Testing SOA Services**

4606 Testing of SOA services should be consistent with the SOA paradigm. In particular,
4607 testing resources and artifacts should be visible in support of service interaction
4608 between providers and consumers, where here the interaction is between the testing
4609 resource and the tester. In addition, the idea of opacity of the implementation should
4610 limit the details that need to be available for effective use of the test resources. Testing
4611 that requires knowledge of the internal structure of the service or its underlying
4612 capability should be performed as part of unit testing in the development sandbox, and
4613 should represent a minimum level of confidence before the service begins its transition
4614 to further testing and eventual operation in the SOA ecosystem.

4615 **5.4.4.1 Progression of SOA Testing**

4616 Software testing is a gradual exercise going from micro inspection to testing macro
4617 effects. The first step in testing is likely the traditional code reviews. SOA
4618 considerations would account for the distributed nature of SOA, including issues of
4619 distributed security and best practices to ensure secure resources. It would also set the
4620 groundwork for opacity of implementation, hiding programming details and simplifying
4621 the use of the service.

4622 Code review is likely followed by unit testing in a development sandbox isolated from
4623 the operational environment. The unit testing is done with full knowledge of the service
4624 internal structure and knowledge of resources representing underlying capabilities. It
4625 tests the interface to ensure exchanged messages are as specified in the service
4626 description and the messages can be parsed and interpreted as intended. Unit testing
4627 also verifies intended functionality and that the software has dealt correctly with internal
4628 dependencies, such as structure of a file system or access to other dedicated
4629 resources.

4630 Some aspects of unit testing require external dependencies be satisfied, and this is
4631 often done using mock objects to substitute for the external resources. In particular, it
4632 will likely be necessary to include mocks of existing operational services, both those
4633 provided as part of the SOA infrastructure and services from other providers.

4634 **Service Mock**

4635 A service mock is an entity that mimics some aspect of the performance of an
4636 operational service without committing to the **real world effects** that the
4637 operational service would produce.

4638 Mocks are discussed in detail in sections 5.4.4.3 and 5.4.4.4.

4639 After unit testing has demonstrated an adequate level of confidence in the service, the
4640 testing must transition from the tightly controlled environment of the development
4641 sandbox to an environment that more clearly resembles the operational SOA ecosystem
4642 or, at a minimum, the intended enterprise. While sandbox testing will use simple mocks
4643 of some aspects of the SOA environment, such as an interface to a security service
4644 without the security service functionality, the dynamic nature of SOA makes a full
4645 simulation infeasible to create or maintain. This is especially true when a new
4646 composite service makes use of operational services provided by others. Thus, at
4647 some point before testing is complete, the service will need to demonstrate its
4648 functionality by using resources and dealing with conditions that only exist in the full
4649 ecosystem or the intended enterprise. Some of these resources may still provide test
4650 interfaces -- more on this below -- but the interfaces will be accessible via the SOA
4651 environment and not just implemented for the sandbox.

4652 At this stage, the opacity of the service becomes important as the details of interacting
4653 with the service now rely on correct use of the service interface and not knowledge of
4654 the service internals. The workings of the service will only be observable through the
4655 **real world effects** realized through service interactions and external indications that
4656 conditions of use, such as user authentication, are satisfied. Monitoring the behavior of
4657 the service will depend on service interfaces that expose internal monitoring or provide
4658 required information to the SOA infrastructure monitoring function. The monitoring
4659 required to test a new service is likely to have significant overlap with the monitoring the
4660 SOA infrastructure includes to monitor its own health and to identify and isolate
4661 behavior outside of acceptable bounds. This is exactly what is needed as part of
4662 service testing, and it is reasonable to assume that the ecosystem transition includes
4663 use of operational monitoring rather than solely dedicated monitoring for each service
4664 being tested.

4665 Use of SOA monitoring resources during the explicit testing phase sets the stage for
4666 monitoring and a level of continual testing throughout the service lifetime.

4667 **5.4.4.2 Testing Traditional Dependencies vs. Service Interactions**

4668 A SOA service is not required to make use of other operational services beyond what
4669 may be required for monitoring by the ecosystem infrastructure. The service can
4670 implement hardcoded dependencies which have been tested in the development
4671 sandbox through the use of dedicated mocks. While coordination may be required with
4672 real data sources during integration testing, the dependencies can be constrained to
4673 things that can be tested in a more traditional manner. Policies can also be set to
4674 restrict access to pre-approved users, and thus the question of unanticipated users and
4675 unanticipated uses can be eliminated. Operational readiness can be defined in terms of
4676 what can be proven in isolated testing. While all this may provide more confidence in
4677 the service for its designed **purpose**, such a service will not fully participate in the

4678 benefits or challenges of the ecosystem. This is akin to filling a swimming pool with sea
4679 water and having someone in the pool say they are swimming in the ocean.

4680 In considering the testing needed for a fully participating service, consider the example
4681 of a new composite service that combines the **real world effects** and complies with the
4682 conditions of use of five existing operational services. The developer of the composite
4683 service does not own any of the component services and has limited, if any, ability to
4684 get the distributed owners to do any customization. The developer also is limited by the
4685 principle of opacity to information comprising the service description, and does not know
4686 internal details of the component services. The developer of the composite service
4687 must use the component services as they exist as part of the SOA environment,
4688 including what is provided to support testing by new users. This introduces
4689 requirements for what is needed in the way of service mocks.

4690 **5.4.4.3 Use of Service Mocks**

4691 Service mocks enables the tested service to respond to specific features of an
4692 operational service that is being used as a component. It allows service testing to
4693 proceed without needing access to or with only limited engagement with the component
4694 service. Mocks can also mimic difficult to create situations for which it is desired to test
4695 the new service response. For composite services using multiple component services,
4696 mocks may be used in combination to function for any number of the components.

4697 Note, when using service mocks, it is important to remember that it is not the
4698 component service that is being tested (although anomalous behavior may be
4699 uncovered during testing) but the use of the component in the new composite.

4700 Individual service mocks can emphasize different features of the component service
4701 they represent but any given mock does not have to mimic all features. For example, a
4702 mock of the service interface can echo a sent message and demonstrate the message
4703 is reaching its intended destination. A mock could go further and parse the sent
4704 message to demonstrate the message not only reached its destination but was
4705 understood. As a final step, the mock could report back what actions would have been
4706 taken by the component service and what **real world effects** would result. If the
4707 response mimicked the operational response, functional testing could proceed as if the
4708 **real world effect** actually occurred.

4709 There are numerous ways to provide mock functionality. The service mock could be a
4710 simulation of the operational service and return simulated results in a realistic response
4711 message or event notification. It is also possible for the operational service to act as its
4712 own mock and simply not execute the commit stage of its functionality. The service
4713 mock could use a combination of simulation and service action without commit to
4714 generate a report of what would have occurred during the defined interaction with the
4715 operational service.

4716 As the service proceeds through testing, mocks should be systematically replaced by
4717 the component resources accessed through their operational interfaces. Before beta
4718 testing begins, end-to-end testing, i.e. proceeding from the beginning of the service
4719 interaction to the resulting real world results, should be accomplished using component
4720 resources via their operational interfaces.

4721 **5.4.4.4 Providers of Service Mocks**

4722 In traditional testing, it is often the test professionals who design and develop the
4723 mocks, but in the distributed world of SOA, this may not be efficient or desirable.

4724 In the development sandbox, it is likely the new service developer or test professionals
4725 working with the developer will create mocks adequate for unit testing. Given that most
4726 of this testing is to verify the new service is performing as designed, it is not necessary
4727 to have high fidelity models of other resources being accessed. In addition, given
4728 opacity of SOA implementation, the developer of the new service may not have
4729 sufficient detailed knowledge of a component service to build a detailed mock of the
4730 component service functionality. Sharing existing mocks at this stage may be possible
4731 but the mocks would need to be implemented in the sandbox, and for simple models it
4732 is likely easier to build the mock from scratch.

4733 As testing begins its transition to the wider SOA environment, mocks may be available
4734 as services. For existing resources, it is possible that an Open Source model could
4735 evolve where service mocks of available functions can be catalogued and used during
4736 initial interaction of the tested service and the operational environment. Widely used
4737 functions may have numerous service mocks, some mimicking detailed conditions
4738 within the SOA infrastructure. However, the Open Source model is less likely to be
4739 sufficient for specialty services that are not widely used by a large consumer
4740 community.

4741 The service developer is probably best qualified for also developing more detailed
4742 service mocks or for mock modes of operational services. This implies that in addition
4743 to their operational interfaces, services will routinely provide test interfaces to enable
4744 service mocks to be used as services. As noted above, a new service developer
4745 wanting to build a mock of component services is limited to the description provided by
4746 the component service developer or owner. The description typically will detail **real**
4747 **world effects** and conditions of use but will not provide implementation details, some of
4748 which may be proprietary. Just as important in the SOA ecosystem, if it becomes
4749 standard protocol for developers to create service mocks of their own services, a new
4750 service developer is only responsible for building his own mocks and can expect other
4751 mocks to be available from other developers. This reduces duplication of effort where
4752 multiple developers would be trying to build the same mocks from the same insufficient
4753 information. Finally, a service developer is probably best qualified to know when and
4754 how a service mock should be updated to reflect modified functionality or message
4755 exchange.

4756 It is also possible that testing organizations will evolve to provide high-fidelity test
4757 harnesses for new services. The harnesses would allow new services to plug into a test
4758 environment and would facilitate accessing mocks of component services. However, it
4759 will remain a constant challenge for such organizations to capture evolving uses and
4760 characteristics of service interactions in the real SOA environment and maintain the
4761 fidelity and accuracy of the test systems.

4762 **5.4.4.5 Fundamental Questions for SOA Testing**

4763 In order for the transition to the SOA operational environment to proceed, it is necessary
4764 to answer two fundamental questions:

- 4765 • Who provides what testing resources for the SOA operational environment, e.g.
4766 mocks of interfaces, mocks of functionality, monitoring tools?
4767 • What testing needs to be accomplished before operational environment
4768 resources can be accessed for further testing?

4769 The discussion in section 5.4.4.4 notes various levels of sophistication of service mocks
4770 and different communities are likely to be responsible for different levels. Section
4771 5.4.4.4 advocates a significant role for service developers, but there needs to be
4772 community consensus that such mocks are needed and that service developers will
4773 agree to fulfilling this **role**. There is also a need for consensus as to what tools should
4774 be available as services from the SOA infrastructure.

4775 As for use of the service mocks and SOA environment monitoring services, practical
4776 experience is needed upon which guidelines can be established for when a new service
4777 has been adequately tested to proceed with a greater level of exposure with the SOA
4778 environment. Malfunctioning services could cause serious problems if they cannot be
4779 identified and isolated. On the other hand, without adequate testing under SOA
4780 operational conditions, it is unlikely that problems can be uncovered and corrected
4781 before they reach an operational stage.

4782 As noted in section 5.4.4.2, some of these questions can be avoided by restricting
4783 services to more traditional use scenarios. However, such restriction will limit the
4784 effectiveness of SOA use and the result will resemble the constraints of traditional
4785 integration activities we are trying to move beyond.

4786 **5.4.5 Architectural Implications for SOA Testing**

4787 The discussion of SOA Testing indicates numerous architectural implications on the
4788 SOA ecosystem:

- 4789 • The distributed, boundary-less nature of the SOA ecosystem makes it
4790 infeasible to create and maintain a single mock of the entire ecosystem to
4791 support testing activities.
- 4792 • A standard suite of monitoring services needs to be defined, developed,
4793 and maintained. This should be done in a manner consistent with the evolving
4794 nature of the ecosystem.
- 4795 • Services should provide interfaces that support access in a test mode.
- 4796 • Testing resources must be described and their descriptions must be
4797 catalogued in a manner that enables their discovery and access.
- 4798 • Guidelines for testing and ecosystem access need to be established and
4799 the ecosystem must be able to enforce those guidelines asserted as policies.
- 4800 • Services should be available to support automated testing and regression
4801 testing.
- 4802 • Services should be available to facilitate updating service description by
4803 anyone who has performed testing of a service.

4804 6 Conformance

4805 This Reference Architecture is an abstract architectural description of Service Oriented
4806 Architecture, which means that it is especially difficult to construct tests for conformance
4807 to the architecture. In addition, conformance to an architectural specification does not,
4808 by itself, guarantee any form of interoperability between multiple implementations.

4809 However, it *is* possible to decide whether or not a given architecture is conformant to an
4810 architectural description such as this one. In discussions of conformance we use the
4811 term **target architecture** to identify the (typically concrete) architecture that may be
4812 viewable as conforming to the abstract principles outlined in this Reference
4813 Architecture.

4814 Target Architecture

4815 A **target architecture** is an architectural description of a system that is intended
4816 to be viewed as conforming to this Reference Architecture.

4817 While we cannot guarantee interoperability between **target architectures** (or more
4818 specifically between applications and systems residing within the ecosystems of those
4819 **target architectures**), interoperability between **target architectures** is promoted by
4820 conformance to this Reference Architecture as it reduces the semantic impedance
4821 mismatch between the different ecosystems.

4822 The primary measure of conformance is whether given concepts as described in this
4823 Reference Architecture have correspondence with the **target architecture**. Such a
4824 correspondence **MUST** honor the relationships identified within this document for the
4825 **target architecture** to be considered conforming.

4826 For example, in Section 3.1.5.1 we identify **resource** as a key concept. A **resource** is
4827 associated with an **owner** and a number of **identifiers**. For a **target architecture** to
4828 conform to this Reference Architecture, it must be possible to find corresponding
4829 concepts of **resource**, **identifier** and **owner** within the **target architecture**: say *entity*,
4830 *token* and *user*. Furthermore, the relationships between *entity*, *token* and *user* **MUST**
4831 mirror the relationships between **resource**, **identifier** and **owner** appropriately.

4832 Clearly, such correspondence is simpler if the terminology within the **target**
4833 **architecture** is identical to that in this Reference Architecture. But so long as the ‘graph’
4834 of concepts and relationships is consistent, that is all that is required for the **target**
4835 **architecture** to conform to this Reference Architecture.

4836 [EDITOR’S NOTE: The conformance section is not complete]

4837

4838 **A. Acknowledgements**

4839 The following individuals have participated in the creation of this specification and are
4840 gratefully acknowledged:

4841 **Participants:**

4842 Rex Brooks, Individual Member
4843 Peter Brown, Pensive.eu
4844 Scott Came, Search Group Inc.
4845 Joseph Chiusano, Booz Allen Hamilton
4846 Robert Ellinger, Northrop Grumman Corporation
4847 David Ellis, Sandia National Laboratories
4848 Jeff A. Estefan, Jet Propulsion Laboratory
4849 Don Flinn, Individual Member
4850 Anil John, Johns Hopkins University
4851 Ken Laskey, MITRE Corporation
4852 Boris Lublinsky, Nokia Corporation
4853 Francis G. McCabe, Individual Member
4854 Christopher McDaniels, USSTRATCOM
4855 Tom Merkle, Lockheed Martin Corporation
4856 Jyoti Namjoshi, Patni Computer Systems Ltd.
4857 Duane Nickull, Adobe Inc.
4858 James Odell, Associate
4859 Michael Poulin, Fidelity Investments
4860 Michael Stiefel, Associate
4861 Danny Thornton, Northrop Grumman
4862 Timothy Vibbert, Lockheed Martin Corporation
4863 Robert Vitello, New York Dept. of Labor

4864 B. Critical Factors Analysis

4865 A critical factors analysis (CFA) is an analysis of the key properties of a project. A CFA
4866 is analyzed in terms of the goals of the project, the critical factors that will lead to its
4867 success and the measurable requirements of the project implementation that support
4868 the goals of the project. CFA is particularly suitable for capturing quality attributes of a
4869 project, often referred to as “non-functional” or “other-than-functional” requirements: for
4870 example, security, scalability, wide-spread adoption, and so on. As such, CFA
4871 complements rather than attempts to replace other requirements capture techniques.

4872 B.1 Goals

4873 A goal is an overall target that you are trying to reach with the project. Typically, goals
4874 are hard to measure by themselves. Goals are often directed at the potential consumer
4875 of the product rather than the technology developer.

4876 Critical Success Factors

4877 A critical success factor (CSF) is a property, sub-goal that directly supports a goal and
4878 there is strong belief that without it the goal is unattainable. CSFs themselves are not
4879 necessarily measurable in themselves.

4880 Requirements

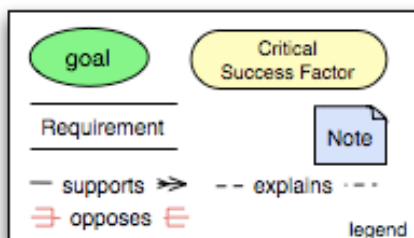
4881 A requirement is a specific measurable property that directly supports a CSF. The key
4882 here is measurability: it should be possible to unambiguously determine if a requirement
4883 has been met. While goals are typically directed at consumers of the specification,
4884 requirements are focused on technical aspects of the specification.

4885 CFA Diagrams

4886 It can often be helpful to illustrate graphically the key concepts and relationships
4887 between them. Such diagrams can act as effective indices into the written descriptions
4888 of goals etc., but is not intended to replace the text.

4889 The legend:

4890



4891 illustrates the key elements of the graphical notation. Goals are written in round ovals,
4892 critical success factors are written in round-ended rectangles and requirements are
4893 written using open-ended rectangles. The arrows show whether a
4894

4895 CSF/goal/requirement is supported by another element or opposed by it. This highlights
4896 the potential for conflict in requirements.

10/26/2010