



---

# Reference Architecture Foundation for Service Oriented Architecture Version 1.0

**Draft 05 27 Apr 2011**

**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](mailto:jeffrey.a.estefan@jpl.nasa.gov)

Ken Laskey, MITRE Corporation, [klaskey@mitre.org](mailto:klaskey@mitre.org)

Francis G. McCabe, Individual, [fmccabe@gmail.com](mailto:fmccabe@gmail.com)

Danny Thornton, Northrop Grumman, [danny.thornton@ngc.com](mailto: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 (SOA-RAF). 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 specific SOA concrete architectures can be built.

The focus of the SOA-RAF is on an approach to integrating business with the information technology needed to support it. These issues are always present but are all the more important when business integration involves crossing ownership boundaries.

The SOA-RAF follows the recommended practice of describing architecture in terms of models, views, and viewpoints, as prescribed in the ANSI<sup>1</sup>/IEEE<sup>2</sup> 1471-2000, (now ISO<sup>3</sup>/IEC<sup>4</sup> 42010-2007) Standard. The SOA-RAF is of value to

---

<sup>1</sup> American National Standards Institute

<sup>2</sup> Institute of Electrical and Electronics Engineers

<sup>3</sup> International Organization for Standardization

<sup>4</sup> International Electrotechnical Commission

Enterprise Architects, Business and IT Architects as well as CIOs and other senior executives involved in strategic business and IT planning.

The SOA-RAF has three main views: the *Participation in a SOA Ecosystem* view which focuses on the way that participants are part of a Service Oriented Architecture ecosystem; the *Realization of a SOA Ecosystem* view which addresses the requirements for constructing a SOA-based system in a SOA ecosystem; and the *Ownership in a SOA Ecosystem* view which focuses on what is meant to own a SOA-based system.

**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](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](http://www.oasis-open.org/committees/tc_home.php?wg_abbrev=soa-rm).

---

## 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.

---

# 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 .....	13
1.2	Service Oriented Architecture – An Ecosystems Perspective .....	13
1.3	Viewpoints, Views and Models .....	14
1.3.1	ANSI/IEEE 1471-2000::ISO/IEC 42010-2007 .....	14
1.3.2	UML Modeling Notation .....	16
1.4	SOA-RAF Viewpoints .....	16
1.4.1	<i>Participation in a SOA Ecosystem</i> viewpoint .....	17
1.4.2	<i>Realization of a SOA Ecosystem</i> viewpoint .....	18
1.4.3	<i>Ownership in a SOA Ecosystem</i> viewpoint .....	18
1.5	Terminology .....	18
1.5.1	Usage of Terms .....	18
1.6	References .....	18
1.6.1	Normative References .....	18
1.6.2	Non-Normative References .....	19
2	Architectural Goals and Principles .....	21
2.1	Goals and Critical Success Factors of the Reference Architecture Foundation ...	21
2.1.1	Goals .....	22
2.1.2	Critical Success Factors .....	23
2.2	Principles of this Reference Architecture Foundation .....	24
3	<i>Participation in a SOA Ecosystem</i> view .....	27
3.1	Social Structure in a SOA Ecosystem Model .....	29
3.1.1	Participants, Actors and Delegates .....	31
3.1.2	Roles in Social Structures .....	32
3.1.3	Resource and Ownership .....	36
3.1.4	Policies and Contracts .....	40
3.1.5	Communication .....	41
3.1.6	Semantics and Semantic Engagement .....	42
3.1.7	Trust and Risk .....	<b>Error! Bookmark not defined.</b>
3.2	Action in a SOA Ecosystem Model .....	43
3.2.1	Needs, Requirements and Capabilities .....	44

3.2.2	Services Reflecting Business .....	46
3.2.3	Action, Communication and Joint Action .....	47
3.2.4	State, Shared State and Real-World Effect .....	48
3.3	Architectural Implications .....	49
3.3.1	Social structures .....	49
3.3.2	The Importance of Action .....	49
3.3.3	Communications as a Means of Mediating Action .....	50
3.3.4	Semantics .....	50
3.3.5	Trust and Risk .....	51
3.3.6	Policies and Contracts .....	51
4	<i>Realization of a SOA Ecosystem view</i> .....	52
4.1	Service Description Model .....	52
4.1.1	The Model for Service Description .....	54
4.1.2	Use Of Service Description .....	64
4.1.3	Relationship to Other Description Models .....	70
4.1.4	Architectural Implications .....	71
4.2	Service Visibility Model .....	73
4.2.1	Visibility to Business .....	73
4.2.2	Visibility .....	75
4.2.3	Architectural Implications .....	80
4.3	Interacting with Services Model .....	81
4.3.1	Interaction Dependencies .....	81
4.3.2	Actions and Events .....	82
4.3.3	Message Exchange .....	83
4.3.4	Composition of Services .....	86
4.3.5	Architectural Implications of Interacting with Services .....	90
4.4	Policies and Contracts Model .....	92
4.4.1	Policy and Contract Representation .....	92
4.4.2	Policy and Contract Enforcement .....	94
4.4.3	Architectural Implications .....	95
5	<i>Ownership in a SOA Ecosystem View</i> .....	97
5.1	Governance Model .....	97
5.1.1	Understanding Governance .....	98
5.1.2	A Generic Model for Governance .....	99
5.1.3	Governance Applied to SOA .....	104
5.1.4	Architectural Implications of SOA Governance .....	110
5.2	Security Model .....	111
5.2.1	Secure Interaction Concepts .....	112
5.2.2	Where SOA Security is Different .....	115

5.2.3 Security Threats .....	115
5.2.4 Security Responses .....	116
5.2.5 Architectural Implications of SOA Security .....	119
5.3 Management Model .....	119
5.3.1 Management and Governance.....	129
5.3.2 Management Contracts and Policies.....	129
5.3.3 Management Infrastructure .....	129
5.3.4 Service Life-cycle .....	129
5.4 SOA Testing Model.....	129
5.4.1 Traditional Software Testing as Basis for SOA Testing.....	129
5.4.2 Testing and the SOA Ecosystem .....	131
5.4.3 Elements of SOA Testing.....	132
5.4.4 Testing SOA Services .....	138
5.4.5 Architectural Implications for SOA Testing .....	142
6 Conformance.....	144
A. Acknowledgements .....	145
B. Index of Defined Terms.....	146
C. The Unified Modeling Language, UML.....	149
D. Critical Factors Analysis .....	150
D.1 Goals.....	150
Critical Success Factors.....	150
Requirements.....	150
CFA Diagrams.....	150

---

---

## Table of Figures

Figure 1- SOA Open Standards Work.....	13
Figure 2 Critical Factors Analysis of the Reference Architecture .....	22
Figure 3 Model elements described in the Participation in a SOA ecosystem view .....	29
Figure 4 Social Structure.....	29
Figure 5 Actors, Participants and Delegates .....	31
Figure 7 Service Roles .....	35
Figure 8 Policies and Contracts .....	40
Figure 9 Resources.....	<b>Error! Bookmark not defined.</b>
Figure 10 Resource Ownership.....	<b>Error! Bookmark not defined.</b>
Figure 11 Actions, Real World Effect and Events.....	<b>Error! Bookmark not defined.</b>
<i>Figure 12 Trusting Actor and Willingness.....</i>	<b>Error! Bookmark not defined.</b>
Figure 13 Communication as Joint Action.....	<b>Error! Bookmark not defined.</b>
Figure 14 Model Elements Described in the Realization of a SOA ecosystem view .....	52
<i>Figure 15 General Description .....</i>	55
Figure 16 Representation of a Description .....	57
<i>Figure 17 Service Description .....</i>	59
Figure 18 Service Interface .....	60
<i>Figure 19 Service Functionality.....</i>	61
Figure 20 Model for Policies and Contracts as related to Service Participants .....	62
Figure 21 Action-Level and Service-Level Policies .....	63
Figure 22 Policies and Contracts, Metrics, and Compliance Records.....	64
Figure 23 Relationship Between Action and Service Description Components.....	65
Figure 24 Execution Context .....	69
Figure 25 Interaction Description .....	70
Figure 26 Visibility to Business.....	74
Figure 27 Mediated Service Awareness.....	76
Figure 28 Awareness In a SOA Ecosystem .....	78
Figure 29 Business, Description and Willingness.....	79
Figure 30 Service Reachability.....	80
Figure 18 Interaction dependencies.....	82
Figure 19 A "message" conveys either an action or an event.....	82
Figure 20 Fundamental SOA message exchange patterns (MEPs).....	84
Figure 21 Simple model of service composition.....	86
Figure 22 Abstract example of orchestration of service-oriented business process.....	88



Figure 23 Abstract example of choreography of service-oriented business collaboration. .....	90
Figure 37 Policies and Contracts .....	93
Figure 38 Model Elements Described in the Ownership in a SOA ecosystem View .....	97
Figure 39 Motivating governance model .....	100
Figure 40 Setting up governance model.....	101
Figure 41 Carrying out governance model .....	102
Figure 42 Ensuring governance compliance model .....	103
Figure 43 Relationship among types of governance .....	106
Figure 44 Confidentiality and Integrity .....	113
Figure 45 Authentication .....	113
Figure 46 Authorization .....	114
Figure 47 Managing resources in a SOA .....	129
Figure 48 Example UML class diagram—Resources.....	149



---

# 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.<sup>5</sup>

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/Developers - 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 of interest independent of the technologies, protocols, and products that are used to implement a specific solution for 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, while staying independent of any particular solution but instead applies to a class of solutions.

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

---

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

37 generally will not completely specify all the technologies, components and their  
38 relationships in sufficient detail to enable direct implementation.

### 39 1.1.2 What is this Reference Architecture?

40 There is a continuum of architectures, from the most abstract to the most detailed. This  
41 Reference Architecture is an abstract realization of SOA, focusing on the elements and  
42 their relationships needed to enable SOA-based systems to be used, realized and  
43 owned while avoiding reliance on specific concrete technologies. It is therefore at the  
44 more abstract end of the continuum, described in [TOGAF v9] as a “foundation  
45 architecture”. It is nonetheless a *reference* architecture as it remains solution-  
46 independent. It is defined therefore as a *Reference Architecture Foundation*, because it  
47 takes a first principles approach to architectural modeling of SOA-based systems.

48 While requirements are addressed more fully in Section 2, the SOA-RAF makes key  
49 assumptions that SOA-based systems involve:

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

57 Even in apparently homogenous structures, such as within a single organization,  
58 different groups and departments nonetheless often have ownership boundaries  
59 between them. This reflects organizational reality as well as the real motivations and  
60 desires of the people running those organizations.

61 Such an environment as described above is an *ecosystem* and, specifically in the  
62 context of SOA-based systems, is a **SOA ecosystem**. This concept of an ecosystem  
63 perspective of SOA is elaborated further in Section 1.2.

64 This SOA-RAF shows how Service Oriented Architecture fits into the life of users and  
65 stakeholders, how SOA-based systems may be realized effectively, and what is  
66 involved in owning and managing them. This serves two purposes: to ensure that SOA-  
67 based systems take account of the specific constraints of a SOA ecosystem, and to  
68 allow the audience to focus on the high-level issues without becoming over-burdened  
69 with details of a particular implementation technology.

### 70 1.1.3 Relationship to the OASIS Reference Model for SOA

71 The OASIS Reference Model for Service Oriented Architecture identifies the key  
72 characteristics of SOA and defines many of the important concepts needed to  
73 understand what SOA is and what makes it important. The Reference Architecture  
74 Foundation takes the Reference Model as its starting point, in particular the vocabulary  
75 and definition of important terms and concepts.

76 The SOA-RAF goes further in that it shows how SOA-based systems can be realized –  
77 albeit in an abstract way. As noted above, SOA-based systems are better thought of as  
78 dynamic systems rather than stand-alone software products. Consequently, how they

79 are used and managed is at least as important architecturally as how they are  
80 constructed.

#### 81 **1.1.4 Relationship to other Reference Architectures**

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

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

91 In addition, this Reference Architecture Foundation, as the title illustrates, is intended to  
92 provide foundational models on which to build other reference architectures and  
93 eventual concrete architectures. The relationship to other industry reference  
94 architectures for SOA and related SOA open standards is described in Appendix E.

#### 95 **1.1.5 Expectations set by this Reference Architecture Foundation**

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

### 102 **1.2 Service Oriented Architecture – An Ecosystems Perspective**

103 Many systems cannot be completely understood by a simple decomposition into parts  
104 and subsystems – in particular when many autonomous parts of the system are  
105 governing interactions. We need also to understand the context within which the system  
106 functions and the participants involved in making it function. This is the **ecosystem**. For  
107 example, a biological ecosystem is a self-sustaining and dynamic association of plants,  
108 animals, and the physical environment in which they live. Understanding an ecosystem  
109 often requires a holistic perspective that considers the relationships between the  
110 elements of the system and their environment at least as important as the individual  
111 parts of the system.

112 This Reference Architecture Foundation views the SOA architectural paradigm from an  
113 ecosystems perspective: whereas a system will be a capability developed to fulfill a  
114 defined set of needs, a SOA ecosystem is a space in which people, processes and  
115 machines act together to deliver those capabilities as services.

116 Viewed as whole, a SOA ecosystem is a network of discrete processes and machines  
117 that, together with a community of people, creates, uses, and governs specific services  
118 as well as external suppliers of resources required by those services.

119 In a SOA ecosystem there may not be any single person or organization that is really "in  
120 control" or "in charge" of the whole although there are identifiable stakeholders who  
121 have influence within the community and control over aspects of the overall system.

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

- 123 • a SOA is a paradigm for *exchange of value* between independently acting  
124 *participants*;
- 125 • participants (and stakeholders in general) have legitimate claims to *ownership* of  
126 resources that are made available via the SOA; and
- 127 • the behavior and performance of the participants are subject to *rules of engagement*  
128 which are captured in a series of policies and contracts.

## 129 **1.3 Viewpoints, Views and Models**

### 130 **1.3.1 ANSI/IEEE 1471-2000::ISO/IEC 42010-2007**

131 The SOA-RAF uses and follows the IEEE "Recommended Practice for Architectural  
132 Description of Software-Intensive Systems" [**ANSI/IEEE 1471**] and [**ISO/IEC 42010**].  
133 An architectural description conforming to this standard must include the following six  
134 (6) elements:

- 135 1. Architectural description identification, version, and overview information
- 136 2. Identification of the system stakeholders and their concerns judged to be relevant  
137 to the architecture
- 138 3. Specifications of each viewpoint that has been selected to organize the  
139 representation of the architecture and the rationale for those selections
- 140 4. One or more architectural views
- 141 5. A record of all known inconsistencies among the architectural description's  
142 required constituents
- 143 6. A rationale for selection of the architecture (in particular, showing how the  
144 architecture supports the identified stakeholders' concerns).

145 The standard defines the following terms<sup>6</sup>:

#### 146 **Architecture**

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

#### 150 **Architectural Description**

151 A collection of products that document the architecture.

---

<sup>6</sup> See <http://www.iso-architecture.org/ieee-1471/conceptual-framework.html> for a diagram of the standard's Conceptual Framework

152 **System**

153 A collection of components organized to accomplish a specific function or set of  
154 functions.

155 **System Stakeholder**

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

158 A stakeholder's concern should not be confused with either a need or a formal  
159 requirement. A concern, as understood here, is an area or topic of interest. Within that  
160 concern, system stakeholders may have many different requirements. In other words,  
161 something that is of interest or importance is not the same as something that is  
162 obligatory or of necessity **[TOGAF v9]**.

163 When describing architectures, it is important to identify stakeholder concerns and  
164 associate them with viewpoints to insure that those concerns are addressed in some  
165 manner by the models that comprise the views on the architecture. The standard  
166 defines views and viewpoints as follows:

167 **View**

168 A representation of the whole system from the perspective of a related set of  
169 concerns.

170 **Viewpoint**

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

174 In other words, a view is what the stakeholders see whereas the viewpoint defines the  
175 perspective from which the view is taken and the methods for, and constraints upon,  
176 modeling that view.

177 It is important to note that viewpoints are independent of a particular system (or  
178 solutions). In this way, the architect can select a set of candidate viewpoints first, or  
179 create new viewpoints, and then use those viewpoints to construct specific views that  
180 will be used to organize the architectural description. A view, on the other hand, is  
181 specific to a particular system. Therefore, the practice of creating an architectural  
182 description involves first selecting the viewpoints and then using those viewpoints to  
183 construct specific views for a particular system or subsystem. Note that the standard  
184 requires that each view corresponds to exactly one viewpoint. This helps maintain  
185 consistency among architectural views which is a normative requirement of the  
186 standard.

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

188 **Model**

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

191 All architectural models used in a particular view are developed using the methods  
192 established by the architectural viewpoint associated with that view. An architectural

193 model may participate in more than one view but a view must conform to a single  
194 viewpoint.

### 195 **1.3.2 UML Modeling Notation**

196 An open standard modeling language is used to help visualize structural and behavioral  
197 architectural concepts. Although many architecture description languages exist, we  
198 have adopted the Unified Modeling Language™ 2 (UML® 2) **[UML 2]** as the main  
199 viewpoint modeling language. Normative UML is used unless otherwise stated but it  
200 should be noted that it can only partially describe the concepts in each model – it is  
201 important to read the text in order to gain a more complete understanding of the  
202 concepts being described in each section..

203 Appendix B introduces the UML notation that is used in this document.

### 204 **1.4 SOA-RAF Viewpoints**

205 The RAF uses three views that conform to three viewpoints: *Participation in a SOA*  
206 *Ecosystem*, *Realization of a SOA Ecosystem*, and *Ownership in a SOA Ecosystem*.  
207 There is a one-to-one correspondence between viewpoints and views (see Table 1).



Viewpoint Element	Viewpoint		
	<i>Participation in a SOA Ecosystem</i>	<i>Realization of a SOA Ecosystem</i>	<i>Ownership in a SOA Ecosystem</i>
Main concepts covered	Captures what is meant for people to participate in a SOA ecosystem.	Captures what is meant to realize a SOA-based system in a SOA ecosystem.	Captures what is meant to own a SOA-based system in a SOA ecosystem
Stakeholders addressed	All participants in the SOA ecosystem	Those involved in the design, development and deployment of SOA-based systems	Those involved in governing, managing and securing SOA-based systems
Concerns addressed	Understanding ecosystem constraints and contexts in which business can be conducted predictably and effectively.	Effective construction of SOA-based systems.	Processes to ensure governance, management and security of SOA-based systems.
Modeling Techniques used	UML class diagrams	UML class, sequence,, component, activity, communication, and composite structure diagrams	UML class and communication diagrams

208 *Table 1 Viewpoint specifications for the OASIS Reference Architecture Foundation for SOA*

### 209 **1.4.1 Participation in a SOA Ecosystem viewpoint**

210 This viewpoint captures what a SOA ecosystem is, as an environment for people to  
 211 conduct their business. We do not limit the applicability of such an ecosystem to  
 212 commercial and enterprise systems. We use the term business to include any  
 213 transactional activity between multiple users.

214 All stakeholders in the ecosystem have concerns addressed by this viewpoint. The  
 215 primary concern for people is to ensure that they can conduct their business effectively  
 216 and safely in accordance with the SOA paradigm. The primary concern of decision  
 217 makers is the relationships between people and organizations using systems for which  
 218 they, as decision makers, are responsible but which they may not entirely own, and for  
 219 which they may not own all of the components of the system.

220 Given SOA's value in allowing people to access, manage and provide services across  
 221 ownership boundaries, we must explicitly identify those boundaries and the implications  
 222 of crossing them.

## 223 **1.4.2 Realization of a SOA Ecosystem viewpoint**

224 This viewpoint focuses on the infrastructure elements that are needed to support the  
225 construction of SOA-based systems. From this viewpoint, we are concerned with the  
226 application of well-understood technologies available to system architects to realize the  
227 SOA vision of managing systems and services that cross ownership boundaries.

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

## 230 **1.4.3 Ownership in a SOA Ecosystem viewpoint**

231 This viewpoint addresses the concerns involved in owning and managing a SOA as  
232 opposed to using one or building one. Many of these concerns are not easily  
233 addressed by automation; instead, they often involve people-oriented processes such  
234 as governance bodies.

235 Owning a SOA-based system implies being able to manage an evolving system. It  
236 involves playing an active role in a wider ecosystem. This viewpoint is concerned with  
237 how systems are managed effectively, how decisions are made and promulgated to the  
238 required end points; how to ensure that people may use the system effectively; and how  
239 the system can be protected against, and recover from consequences of, malicious  
240 intent.

## 241 **1.5 Terminology**

242 The keywords “MUST”, “MUST NOT”, “REQUIRED”, “SHALL”, “SHALL NOT”,  
243 “SHOULD”, “SHOULD NOT”, “RECOMMENDED”, “MAY”, and “OPTIONAL” in this  
244 document are to be interpreted as described in **[RFC2119]**.

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

246 The terms “SOA-RAF”, “this Reference Architecture” and “Reference Architecture  
247 Foundation” refer to this document, while “the Reference Model” refers to the OASIS  
248 Reference Model for Service Oriented Architecture”. **[SOA-RM]**.

### 249 **1.5.1 Usage of Terms**

250 Certain terms used in this document to denote concepts with formal definitions and are  
251 used with specific meanings. Where reference is made to a formally defined concept  
252 and the prescribed meaning is intended, we use a **bold font**. The first time these terms  
253 are used, they are also hyperlinked to their definition in the Glossary that appears as  
254 Appendix B to the document. Where a more colloquial or informal meaning is intended,  
255 these words are used without special emphasis.

## 256 **1.6 References**

### 257 **1.6.1 Normative References**

258 **[ANSI/IEEE 1471]** *IEEE Recommended Practice for Architectural Description of*  
259 *Software-Intensive Systems*, American National Standards  
260 Institute/Institute for Electrical and Electronics Engineers,  
261 September 21, 2000.

- 262 **[ISO/IEC 42010]** International Organization for Standardization and  
263 International Electrotechnical Commission, *System and software*  
264 *engineering — Recommended practice for architectural description*  
265 *of software-intensive systems*, July 15, 2007.
- 266 **[RFC2119]** S. Bradner, *Key words for use in RFCs to Indicate Requirement*  
267 *Levels*, <http://www.ietf.org/rfc/rfc2119.txt>, IETF RFC 2119, March  
268 1997.
- 269 **[SOA-RM]** OASIS Standard, "Reference Model for Service Oriented  
270 Architecture 1.0, 12 October 2006. [http://docs.oasis-open.org/soa-](http://docs.oasis-open.org/soa-rm/v1.0/soa-rm.pdf)  
271 [rm/v1.0/soa-rm.pdf](http://docs.oasis-open.org/soa-rm/v1.0/soa-rm.pdf)
- 272 **[UML 2]** *Unified Modeling Language: Superstructure*, Ver. 2.1.1, OMG  
273 Adopted Specification, OMG document formal/2007-02-05, Object  
274 Management Group, Needham, MA, February 5, 2007.
- 275 **[WA]** Architecture of the World Wide Web, W3C, 2004.  
276 <http://www.w3.org/TR/webarch>.
- 277 **[WSA]** David Booth, et al., "Web Services Architecture", W3C Working  
278 Group Note, World Wide Web Consortium (W3C) (Massachusetts  
279 Institute of Technology, European Research Consortium for  
280 Informatics and Mathematics, Keio University), February, 2004.  
281 <http://www.w3.org/TR/2004/NOTE-ws-arch-20040211/>

## 282 1.6.2 Non-Normative References

- 283 **[BLOOMBERG/SCHMELZER]** Jason Bloomberg and Ronald Schmelzer, *Service*  
284 *Orient or Be Doomed!*, John Wiley & Sons: Hoboken, NJ, 2006.
- 285 **[COX]** D. E. Cox and H. Kreger, "Management of the service-oriented  
286 architecture life cycle," "IBM Systems Journal" "44", No. 4, 709-  
287 726, 2005
- 288 **[ERA]** A. Fattah, "Enterprise Reference Architecture," paper presented at  
289 22<sup>nd</sup> Enterprise Architecture Practitioners Conference, London, UK,  
290 April 2009.
- 291 **[ITU-T Rec. X.700 | ISO/IEC 10746-3:1996(E)]** Information processing systems—  
292 Open Systems Interconnection—Basic Reference Model—Part 4:  
293 Management Framework", International Telecommunication Union,  
294 International Organization for Standardization and International  
295 Electrotechnical Commission, Geneva, Switzerland, 1989.
- 296 **[NEWCOMER/LOMOW]** Eric Newcomer and Greg Lomow, *Understanding SOA with*  
297 *Web Services*, Addison-Wesley: Upper Saddle River, NJ, 2005.
- 298 **[OECD]** Organization for Economic Cooperation and Development,  
299 Directorate for Financial, Fiscal and Enterprise Affairs, OECD  
300 Principles of Corporate Governance, SG/CG(99) 5 and 219, April  
301 1999.
- 302 **[TOGAF v9]** *The Open Group Architecture Framework (TOGAF) Version 9*  
303 *Enterprise Edition*, The Open Group, Doc Number: G091, February  
304 2009.

- 305 **[WEILL]** Harvard Business School Press, *IT Governance: How Top*  
306 *Performers Manage IT Decision Rights for Superior Results*, Peter  
307 Weill and Jeanne W. Ross, 2004
- 308 **[DAMIANOU]** Nicodemos C. Damianou , Thesis - A Policy Framework for  
309 Management of Distributed Systems, University of London,  
310 Department of Computing, 2002.
- 311 **[LEVESON]** Nancy G. Leveson, *Safeware: System Safety and Computers*,  
312 Addison-Wesley Professional, Addison-Wesley Publishing  
313 Company, Inc.: Boston, pg. 181, 1995.
- 314 **[STEEL/NAGAPPAN/LAI]** Christopher Steel and Ramesh Nagappan and Ray Lai,  
315 *core Security Patterns:Best Practices and Strategies for J2EE, Web*  
316 *Services and Identity Management*, Prentice Hall: 2005
- 317 **[ISO/IEC 27002]** International Organization for Standardization and  
318 International Electrotechnical Commission, *Information technology*  
319 *-- Security techniques – Code of practice for information security*  
320 *management*, 2007
- 321 **[SOA NAV]** Heather Kreger and Jeff Estefan (Eds.), “Navigating the SOA Open  
322 Standards Landscape Around Architecture,” Joint Paper, The Open  
323 Group, OASIS, and OMG, July 2009.  
324 <http://www.opengroup.org/projects/soa/uploads/40/20044/W096.pdf>

---

## 325 2 Architectural Goals and Principles

326 This section identifies the goals of this Reference Architecture Foundation and the  
327 architectural principles that underpin it.

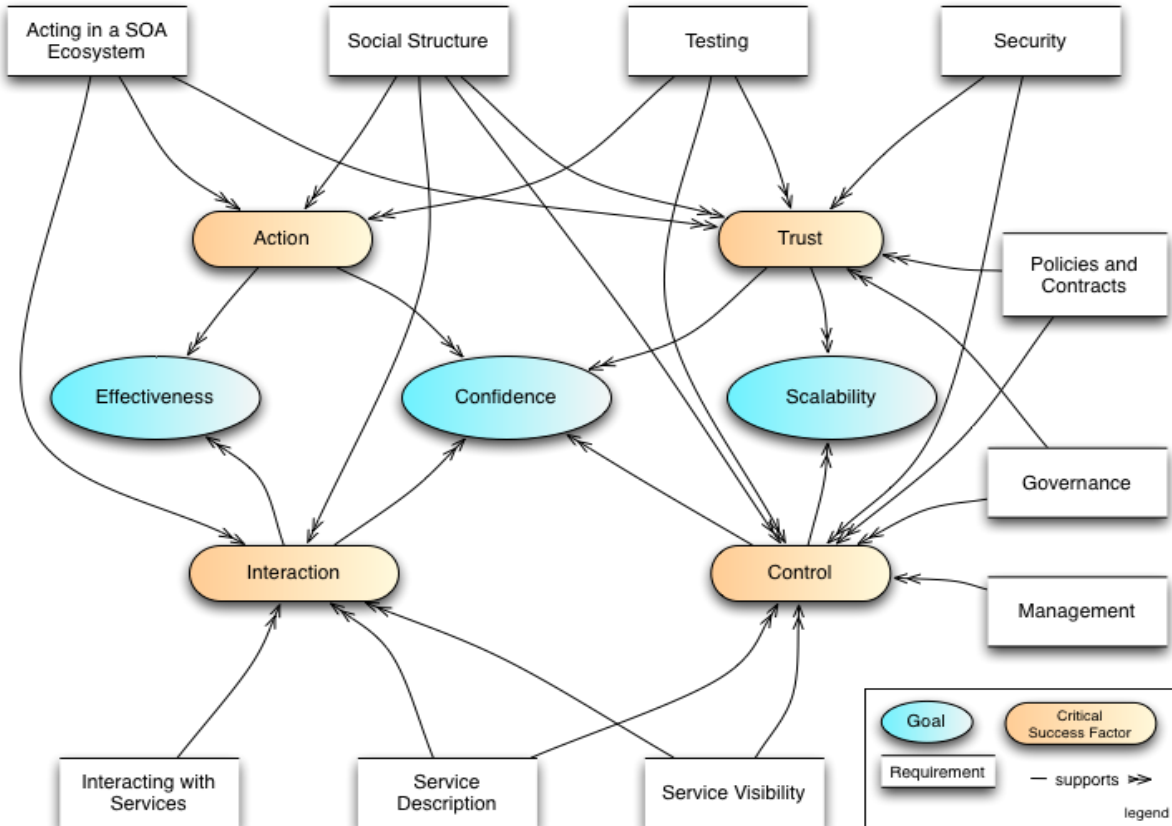
### 328 2.1 Goals and Critical Success Factors of the Reference Architecture 329 Foundation

330 There are three principal goals:

- 331 1. to show how SOA-based systems can effectively bring participants with needs  
332 ('consumers') to interact with participants offering appropriate capabilities as  
333 services ('producers');
- 334 2. for participants to have a clearly understood level of confidence as they interact  
335 using SOA-based systems; and
- 336 3. for SOA-based systems to be scaled for small or large systems as needed.

337 There are four factors critical to the achievement of these goals:

- 338 1. **Action:** an account of participants' action within the ecosystem;
- 339 2. **Trust:** an account of how participants' internal perceptions of the reliability of  
340 others guide their behavior (i.e., the trust that participants may or may not have in  
341 others)
- 342 3. **Interaction:** an account of how participants can interact with each other; and
- 343 4. **Control:** an account of how the management and governance of the entire SOA  
344 ecosystem can be arranged.



345  
346 *Figure 1 Critical Factors Analysis of the Reference Architecture*

347 Figure 1 represents a Critical Factors Analysis (CFA) diagram demonstrating the  
348 relationship between the primary goals of this reference architecture, critical factors that  
349 determine the success of the architecture and individual elements that need to be  
350 modeled.

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

## 356 2.1.1 Goals

### 357 2.1.1.1 Effectiveness

358 A primary purpose of the SOA-RAF is to show how SOA-based systems ensure that  
359 participants can use the facilities of the system to meet their needs. This does not imply  
360 that every need has a SOA solution, but for those needs that can benefit, we look at  
361 what is needed to use the SOA paradigm effectively.

362 The key factors that govern effectiveness from a participant's perspective are actions  
363 undertaken— especially across ownership boundaries — with other participants in the  
364 ecosystem and lead to measurable results.

### 365 **2.1.1.2 Confidence**

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

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

### 374 **2.1.1.3 Scalability**

375 The third goal of this reference architecture is scalability. In architectural terms, we  
376 determine scalability in terms of the smooth growth of complex systems as the number  
377 and complexity of services and interactions between participants increases. Another  
378 measure of scalability is the ease with which interactions can cross ownership  
379 boundaries.

## 380 **2.1.2 Critical Success Factors**

381 A critical success factor (CSF) is a property of the intended system, or a sub-goal that  
382 directly supports a goal and there is strong belief that without it the goal is unattainable.  
383 CSFs are not necessarily measurable in themselves. As illustrated in Figure 1, CSFs  
384 can be associated with more than one goal.

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

### 389 **2.1.2.1 Action**

390 Participants' principal mode of participation in a SOA ecosystem is action; typically  
391 action in the interest of achieving some desired real world effect. Understanding how  
392 action is related to SOA is thus critical to the paradigm.

393 Action is, of course, pervasive in the ecosystem; and many models in the SOA-RAF  
394 address aspects of action. However, action is the central theme of the models labeled  
395 "Action in a Social Context" and "Action in a SOA Ecosystem".

### 396 **2.1.2.2 Trust**

397 The viability of a SOA ecosystem depends on participants being able to effectively  
398 measure the trustworthiness of the system and of participants. Trust is a private  
399 assessment of a participant's belief in the integrity and reliability of the SOA ecosystem  
400 (see Section **Error! Reference source not found.**).

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

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

### 410 **2.1.2.3 Interaction**

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

### 415 **2.1.2.4 Control**

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

421 The governance of SOA-based systems requires decision makers to be able to set  
422 policies about participants, services, and their relationships. It requires an ability to  
423 ensure that policies are effectively described and enforced. It also requires an effective  
424 means of measuring the historical and current performances of services and  
425 participants.

426 The scope of management of SOA-based systems is constrained by the existence of  
427 multiple ownership domains.

## 428 **2.2 Principles of this Reference Architecture Foundation**

429 The following principles serve as core tenets that guided the evolution of this reference  
430 architecture.

### 431 **Technology Neutrality**

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

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

441 **Implications:** The Reference Architecture Foundation must be technology neutral,  
442 meaning that we assume that technology will continue to evolve, and that  
443 over the lifetime of this architecture that multiple, potentially competing  
444 technologies will co-exist. Another immediate implication of technology  
445 independence is that greater effort on the part of architects and other



446 decision makers to construct systems based on this architecture is  
447 needed.

#### 448 **Parsimony**

449 Statement: Parsimony refers to economy of design, avoiding complexity where  
450 possible and minimizing the number of components and relationships  
451 needed.

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

456 Implications: Parsimoniously designed systems tend to have fewer but better targeted  
457 features.

#### 458 **Distinction of Concerns**

459 Statement: Distinction of Concerns refers to the ability to cleanly identify and separate  
460 out the concerns of specific stakeholders in such a way that it is possible  
461 to create architectural models that reflect those stakeholders’ viewpoint. In  
462 this way, an individual stakeholder or a set of stakeholders that share  
463 common concerns only see those models that directly address their  
464 respective areas of interest.

465 Rationale: As SOA-based systems become more mainstream and increasingly  
466 complex, it will be important for the architecture to be able to scale. Trying  
467 to maintain a single, monolithic architecture description that incorporates  
468 all models to address all possible system stakeholders and their  
469 associated concerns will not only rapidly become unmanageable with  
470 rising system complexity, but it will become unusable as well.

471 Implications: This is a core tenet that drives this reference architecture to adopt the  
472 notion of architectural viewpoints and corresponding views. A *viewpoint*  
473 provides the formalization of the groupings of models representing one set  
474 of concerns relative to an architecture, while a *view* is the actual  
475 representation of a particular system. The ability to leverage an industry  
476 standard that formalizes this notion of architectural viewpoints and views  
477 helps us better ground these concepts for not only the developers of this  
478 reference architecture but also for its readers. The IEEE Recommended  
479 Practice for Architectural Description of Software-Intensive Systems  
480 **[ANSI/IEEE 1471-2000::ISO/IEC 42010-2007]** is the standard that serves  
481 as the basis for the structure and organization of this document.

#### 482 **Applicability**

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

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

488 Implications: The Reference Architecture Foundation needs to be relevant to the  
489 problem of matching needs and capabilities under disparate domains of  
490 ownership; to the concepts of “Intranet SOA” (SOA within the enterprise)  
491 as well as “Internet SOA” (SOA outside the enterprise); to the concept of  
492 “Extranet SOA” (SOA within the extended enterprise, i.e., SOA with  
493 suppliers and trading partners); and finally, to “net-centric SOA” or  
494 “Internet-ready SOA.”

---

### 495 3 Participation in a SOA Ecosystem view

496  
497  
498  
499  
500  
501  
502  
503  
504  
505  
506

#### 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

507 The OASIS SOA Reference Model defines Service Oriented Architecture as “a  
508 paradigm for organizing and utilizing distributed capabilities that may be under the  
509 control of different ownership domains” and services as “the mechanism by which  
510 needs and capabilities are brought together”. The central focus of SOA is “the task or  
511 business function – getting something done.”

512 Together, these ideas describe an environment in which business functions (realised in  
513 the form of services) address business needs. Service implementations utilize  
514 capabilities to produce specific (real world) effects that fulfill those business needs. Both  
515 those using the services, and the capabilities themselves, may be distributed across  
516 ownership domains, with different policies and conditions of use in force. The role of a  
517 service in the SOA context is to enable effective business solutions in a distributed  
518 environment. SOA is thus a paradigm that guides the identification, design,  
519 implementation (i.e. organization), and utilization of such services.

520 The *Participation in a SOA Ecosystem* view in the SOA-RAF focuses on the constraints  
521 and context in which people<sup>7</sup> conduct business using a SOA-based system. By  
522 business we mean any shared activity entered into whose **objective** is to satisfy  
523 particular **needs** of each person. The OASIS SOA RM defines SOA as “a paradigm for  
524 organizing and utilizing distributed capabilities that may be under the control of different  
525 ownership domains.” To put it another way, to effectively employ the SOA paradigm,  
526 the architecture must take into account the fact and implications of different ownership  
527 domains, and how best to organize and utilize capabilities that are distributed across  
528 those different ownership domains. These are the main architectural issues that the  
529 Participating in a SOA Ecosystem view tries to address.

530 The subsections below expand on the completely abstract reference model by  
531 identifying more fully and with more specificity what challenges need to be addressed in  
532 order to successfully accomplish SOA. Although this section does not provide a specific

---

<sup>7</sup> ‘People’ and ‘person’ must be understood as both human actors and ‘legal persons’, such as companies, who have rights and responsibilities similar to ‘natural persons’ (humans)

533 recipe, it does identify the important things that need to be thought about and resolved  
534 within an ecosystem context.

535 The people actively participating in a SOA-based system, together with others who may  
536 potentially benefit from the services delivered by the system, together constitute the  
537 **stakeholders**. The stakeholders, the system and the environment (or context) within  
538 which they all operate, taken together forms the **SOA ecosystem**. That ecosystem may  
539 reflect the SOA-based activities within a particular enterprise or of a wider network of  
540 one or more enterprises and individuals. Although a SOA-based system is essentially an  
541 IT concern, it is nonetheless a system engineered deliberately to be able to function in a  
542 SOA ecosystem. In this context, a service is the mechanism that brings a SOA-based  
543 system capability together with stakeholder needs in the wider ecosystem. This is  
544 explored in more detail in Section 3.2.2 below.

545 Furthermore, this *Participation in a SOA Ecosystem* view helps us understand the  
546 importance of execution context – the set of technical and business elements that allow  
547 interaction to occur in, and thus business to be conducted using, a SOA-based system.

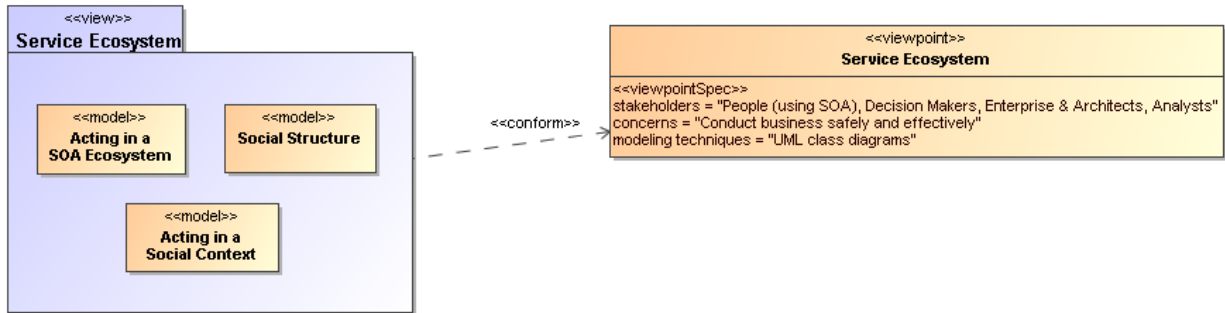
548 This section describes how a SOA-based system behaves when participants may be in  
549 different organizations, with different rules and expectations, and assumes that the  
550 primary motivation for participants to interact with each other is to achieve **objectives** –  
551 to get things done.

552 The dominant mode of communication within a SOA ecosystem is electronic, supported  
553 by IT resources and artifacts. The stakeholders are nonetheless people: since there is  
554 inherent indirection involved when people and systems interact using electronic means,  
555 we lay the foundations for how *communication* can be used to represent and enable  
556 action. However, it is important to understand that these communications are usually a  
557 means to an end and not the primary interest of the participants of the ecosystem.

558 Several interdependent concerns are important in our view of a SOA-ecosystem. The  
559 ecosystem includes stakeholders who are participants in the development, deployment  
560 and governance and use of a system and its services; or who may not participate but  
561 are nonetheless affected by the system. **Actors** – whether stakeholder **participants**  
562 or delegates who act only on behalf of participants (without themselves having any  
563 stake in the ecosystem) – are engaged in **actions** which have an impact on the real  
564 world and whose meaning and intent are determined by implied or agreed-to semantics.

565 The main models in this view are:

- 566 • the **Social Structure in a SOA Ecosystem Model** introduces the key elements  
567 that underlie the relationships between participants and that must be considered  
568 as pre-conditions in order to effectively bring needs and capabilities together  
569 across ownership boundaries;
- 570 • the **Action in a SOA Ecosystem Model** introduces the key concepts involved in  
571 service actions, and shows how joint action and real-world effect are what is  
572 being aimed for in a SOA ecosystem..

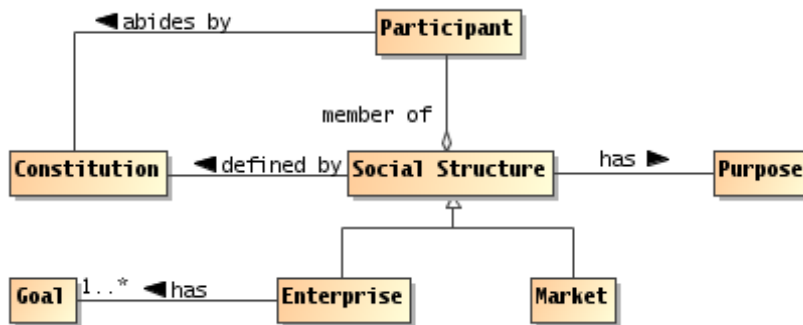


573  
574 Figure 2 Model elements described in the Participation in a SOA Ecosystem view

575 **3.1 Social Structure in a SOA Ecosystem Model**

576 The actions undertaken by participants in a SOA ecosystem are performed in a *social*  
577 *context* that defines the relationships between the participants. That context is the social  
578 structure. In order to achieve success in SOA, the overall social structure in which the  
579 SOA effort is to be undertaken must be taken into consideration. Ownership boundaries  
580 and their implications can only be understood and addressed within the context of the  
581 larger social structure within which they exist and the nature of the relationships  
582 between the different participants in that structure.

583 The primary function of the Social Structure Model is to explain the relationships  
584 between an individual participant and the social context of that participant. The model  
585 also helps in defining and understanding the implications of crossing ownership  
586 boundaries. It is, for example, the foundation for understanding security, governance  
587 and management in the SOA ecosystem.



588  
589 Figure 3 Social Structure

590 **Social Structure**

591 A social structure<sup>8</sup> is a nexus of relationships amongst participants brought  
592 together for a specific purpose. (Social structures are sometimes referred to as  
593 social institutions.)

594 A social structure represents a collection of participants, but a collection that is brought  
595 together for a purpose. There may be a large number of different kinds of relationships

<sup>8</sup> Social structures are sometimes referred to as social institutions.

596 between participants in a social structure. The organizing principle for these  
597 relationships is the social structure's purpose.

598 A social structure may have any number of participants, and a given participant can be  
599 a member of multiple social structures. Thus, there may be interaction among social  
600 structures, sometimes resulting in disagreements when the premises of the social  
601 structures do not align.

602 A social structure has a purpose – the overarching reason for which it exists. All social  
603 structures are established with implied or explicitly defined purpose. The purpose is  
604 usually reflected in specific goals laid down in the social structure's constitution or other  
605 'charter'.

606 A social structure can take different forms. For example, an enterprise is a common kind  
607 of social structure that embodies a form of hierarchic organization; an online chat room  
608 represents a social structure of peers that is very loose. A market represents a social  
609 structure of buyers and sellers. The legal frameworks of entire countries and regions  
610 also count as social structures.

611 The RAF is concerned primarily with social structures that reflect relationships amongst  
612 **participants** in SOA ecosystems, notably:

- 613 • the enterprise social structure which is composed internally of many participants but  
614 that has sufficient cohesiveness to be considered as a potential stakeholder in its  
615 own right; and
- 616 • the peer group which governs relationship between participants within an  
617 ecosystem..

### 618 **Enterprise**

619 An enterprise is a social structure with an identifiable leadership structure, and  
620 that has internally established goals that reflect a defined purpose. It can act as a  
621 participant within other social structures, including other enterprises and is  
622 represented by members of its leadership structure.

### 623 **Peer Group**

624 A peer group is a social structure with no discernable leadership structure, that  
625 may or may not have internally established goals, but is identifiable as the locus of  
626 interaction between participants with individual goals and who are considered  
627 peers of one another.

628 Many interactions between participants take place within social structures. Depending  
629 on the scale and internal structure of an enterprise social structure, these interactions  
630 may or may not cross ownership boundaries (an enterprise can itself be composed of  
631 sub-enterprises). However, interactions between participants within a peer social  
632 structure inherently cross ownership boundaries.

633 The nature and extent of the interactions that take place will reflect, often implicitly,  
634 degrees of trust between participants and the very specific circumstances of each  
635 participant at the time, and over the course, of the interactions. It is in the nature of an  
636 SOA ecosystem that these relationships are rendered more explicit and are formalized  
637 and form a central part of what the SOA-RM refers to as "Execution Context".

638 Social structures involved in a particular interaction are not always explicitly identified.  
639 For example, when a customer buys a book over the Internet, the social structure that

640 determines the validity of the transaction is often the legal framework of the region  
641 associated with the book vendor. Such legal jurisdiction qualification is typically buried  
642 in the fine print of the service description.

### 643 **Constitution**


644 A constitution is a set of rules, written or unwritten, that spell out the purpose,  
645 goals, scope, and functioning of a social structure.

646 Every social structure functions according to rules by which participants interact with  
647 each other within the structure. In some cases, this is based on an explicit agreement,  
648 in other cases participants behave as though they agree to the constitution without a  
649 formal agreement. In still other cases, participants abide by the rules with some degree  
650 of reluctance – this is an issue raised later on when we discuss governance in SOA-  
651 based systems. In all cases, the constitution may change over time, in those cases of  
652 implicit agreement the change can occur quickly.

### 653 **3.1.1 Participants, Actors and Delegates**

654 Social structures have stakeholders, some of whom may be enterprises. They interact  
655 within the broad ecosystem. Actors operate within a system. The concept of Participant  
656 is particularly important as it reflects the hybrid role of both a Stakeholder (in the  
657 ecosystem), primarily concerned with expressing needs and seeing those needs  
658 fulfilled; and an Actor (in the System), directly involved with system-level activity. This  
659 hybrid role of Participant thus provides a bridge between the ecosystem and the  
660 system.

661 An actor can be either a **participant** (and thus also a stakeholder) – with a stake in the  
662 ecosystem; or a **delegate** (a human actor with no stake in the ecosystem or an  
663 automated agent), acting on behalf of a participant.



664  
665 *Figure 4 Actors, Participants and Delegates*

### 666 **Stakeholder**

667 A stakeholder in the SOA ecosystem is a person with an interest – a ‘stake’ – in  
668 the ecosystem.

669 Note: Not all stakeholders necessarily participate in the SOA ecosystem; indeed, the  
670 interest of non-participant stakeholders may be in realizing the benefits of a well-  
671 functioning ecosystem and not suffering unwanted consequences. They can not all or  
672 always be identified in advance but due account is often taken of such stakeholder  
673 types, including potential customers, beneficiaries, affected third parties, as well as  
674 potential “negative stakeholders” who might deliberately seek a negative impact on the  
675 ecosystem (such as hackers or criminals).

676 **Actor**

677 An actor is a human or non-human agent capable of action within a SOA-based  
678 system.

679 **Participant**

680 A participant is a person<sup>9</sup> who is both a stakeholder in the SOA ecosystem and  
681 an actor in the SOA-based system.

682 **Delegate**

683 A delegate is an actor that is acting on behalf of a participant.

684 A delegate can be a person or an automated or semi-automated agent.

685 Many stakeholders and actors operate in a SOA ecosystem, including software agents  
686 that permit people to offer, and interact with, services; delegates that represent the  
687 interests of other participants; or security agents charged with managing the security of  
688 the ecosystem. Note that automated agents are always delegates, in that they act on  
689 behalf of a stakeholder.

690 In the different models of the RAF, actor is used when it is not important whether the  
691 entity is a delegate or a participant. If the actor is acting on behalf of a stakeholder, then  
692 we use delegate. This underlines the importance of delegation in SOA-based systems,  
693 whether the delegation is of work procedures carried out by human agents who have no  
694 stake in the ecosystem but act on behalf of a participant who does; or whether the  
695 delegation is performed by technology (automation). If the actor is also a stakeholder in  
696 the ecosystem, then we use participant.

697 In order for a delegate to act on behalf of another person, they must be able to act and  
698 have the authority to do so.

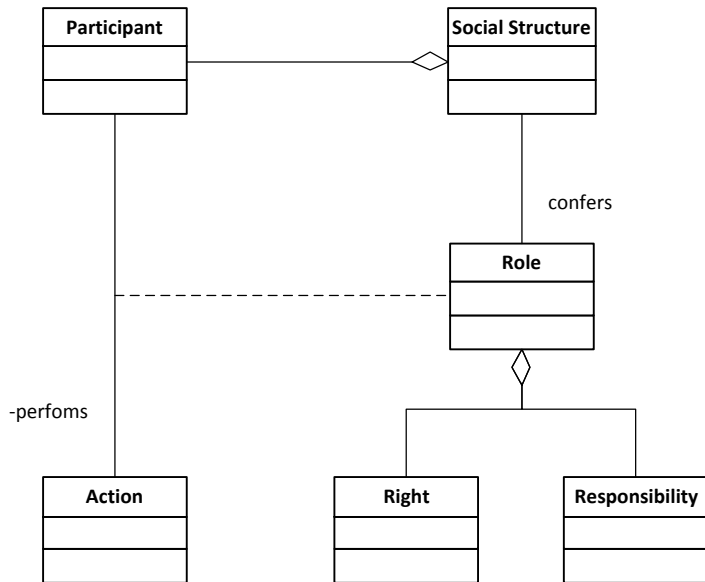
699 **3.1.2 Roles in Social Structures**

700 Social structures are abstractions: a social structure cannot directly perform actions –  
701 only people or automated processes following the instructions of people can actually do  
702 things. However, an actor may act on behalf of a social structure and certainly acts  
703 within a social structure depending on the roles that the actor assumes and the nature  
704 of the relationships between the concerned parties or stakeholders.

---

<sup>9</sup> Again, this can be a 'natural' or 'legal' person





705  
706 *Figure 5 Role in Social Structures*

707 **Role**

708 A role is a type of relationship between a participant and the actions that  
709 participant may performs (or is allowed to perform) within a social structure.

710 A role is not immutable and is often time-bound. A participant can have one or more  
711 roles concurrently and may change them over time and in different contexts, even over  
712 the course of a particular interaction. One participant with appropriate authority in the  
713 social structure may formally *designate a role* for another participant, with associated  
714 rights and responsibilities, and that authority may even qualify a period during which the  
715 designated role may be valid.

716 Conversely, someone who exhibits qualification and skill may *assume a role* without any  
717 formal designation. For example, an office administrator who has demonstrated facility  
718 with personal computers may be known as (and thus assumed to role of) the ‘goto’  
719 person for people who need help with their computers.

720 Although many roles are clearly identified, with appropriate names and definitions of  
721 responsibilities, it is also entirely possible to separately bestow rights, bestow or  
722 assume responsibilities and so on, often in a temporary fashion. For example, when a  
723 company president delegates certain responsibilities on another person, this does not  
724 imply that the other person has become company president. Likewise, a company  
725 president may bestow on someone else her role during a period of time that she is on  
726 vacation or otherwise unreachable, with the understanding that she will re-assume the  
727 role when she returns from vacation.

728 **Authority**

729 Authority is the right or responsibility to act on behalf of an organization or  
730 another person.

731 **Right**

732 A right is a predetermined permission conferred upon an actor that allows them  
733 to perform some action or assume a role in relation to the social structure.

734 Rights can be constrained. For example, sellers might have a general right to refuse  
735 service to potential customers but this right could be constrained so as to be exercised  
736 only when certain criteria are met.

### 737 **Responsibility**

738 A responsibility is a predetermined obligation on a participant to perform some  
739 action or to adopt a stance or role in relation to other actors.

740 Responsibility implies human agency, which is why only participants, as opposed to all  
741 actors (who can be non-human agents) are concerned. even if the consequences of  
742 such responsibility can impact other (human and non-human) actors.

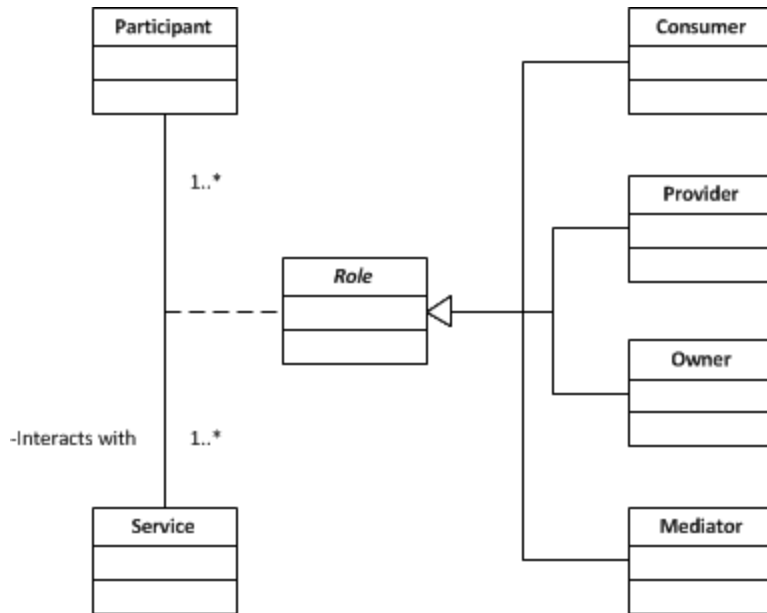
743 Rights, authorities, responsibilities and roles form the foundation for the security model  
744 as well as contributing to the governance model in the 'Ownership in a SOA Ecosystem'  
745 View of the RAF. Rights and responsibilities are similar in structure to permissions and  
746 obligations; except that rights and responsibilities are associated with participants as  
747 opposed to permissions and obligations which are associated with actions.

748 People will assume and perform roles according to their actual or perceived rights and  
749 responsibilities, with or without explicit authority. In the context of a SOA ecosystem,  
750 human abilities and skills are relevant as they equip individuals with knowledge,  
751 information and tools that may be necessary to have meaningful and productive  
752 interactions with a view to achieving a desired outcome. For example, a person who  
753 needs a particular book, and has both the right and responsibility of purchasing the  
754 book from a given bookseller, will not have that need met from the online delegate of  
755 that bookstore if he does not know how to use a web browser. Equally, just because  
756 someone does have the requisite knowledge or skills does not entitle them *per se* to  
757 interact with a specific system.

#### 758 **3.1.2.1 Service Roles**

759 As in roles generically, a participant can play one or more of those roles inherent to the  
760 SOA paradigm in the SOA ecosystem, including as a service consumer, a service  
761 provider, a mediator, and so on, depending on the context. A participant may be playing  
762 a role of a service provider in one relationship while simultaneously playing the role of a  
763 consumer in another. Roles inherent to the SOA paradigm include Consumer, Provider,  
764 and Mediator.

765



766  
767 *Figure 6 Participant Roles in a Service*

768 **Provider**

769 A provider is a role assumed by a participant who is offering a service.

770 **Consumer**

771 A consumer is a role assumed by a participant who is interacting with a service in  
772 order to fulfill a need.

773 **Mediator**

774 A mediator is a role assumed by a participant to facilitate interaction and  
775 connectivity in the offering and use of services.

776 **Owner**

777 An owner is a role assumed by a participant who is claiming and exercising  
778 ownership over a service.

779 It is a common understanding that service interactions are typically initiated by service  
780 consumers, although this is not necessarily true in all situations. Additionally, as with  
781 service providers, several stakeholders may be involved in a service interaction  
782 supporting a given consumer.

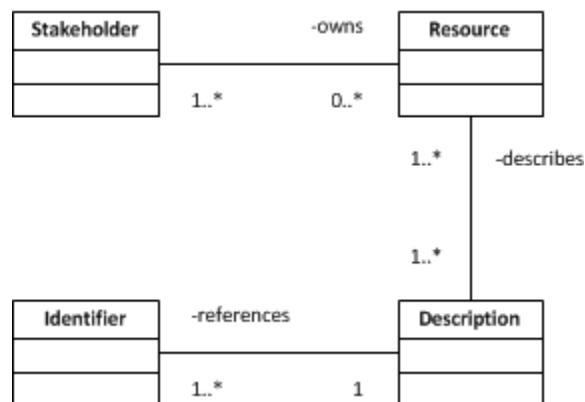
783 The roles of service provider and service consumer are often seen as symmetrical,  
784 which is also not entirely correct. A consumer tends to express a 'Need' in non-formal  
785 terms: "I want to buy that book". The type of 'Need' that a service is intended to fulfill  
786 has to be formalized and encapsulated by designers and developers as a  
787 'Requirement'. This Requirement should then be reflected in the target service, as a  
788 'Capability' that, when accessed via a service, delivers a 'Real World Effect' to an  
789 arbitrary user: "The chosen book is ordered for the user" It thus satisfies the need that  
790 has been defined for an archetypal user. Specific and particular users may not  
791 experience a need exactly as captured by the service: "I don't want to pay that much for  
792 the book", "I wanted an eBook version", etc. There can therefore be a process of implicit  
793 and explicit negotiation between the user and the service, aimed at finding a 'best fit'

794 between the user's specific need and the capabilities of the service that are available  
 795 and consistent with the service provider's offering. This process may continue up until  
 796 the point that the user is able to accept what is on offer as being the best fit and finally  
 797 'invokes' the service. 'Execution context' has thus been established. This is explored in  
 798 more detail later on. Service mediation by a participant can take many forms and may  
 799 invoke and use other services in order to fulfill such mediation. For example, it might  
 800 use a service registry in order to identify possible service partners; or, in our book-  
 801 buying example, it might provide a price comparison service, suggest alternative  
 802 suppliers, different language editions or delivery options.

### 803 3.1.3 Resource and Ownership

#### 804 3.1.3.1 Resource

805 A resource is generally understood as an asset: it has value to someone. Key to this  
 806 concept in a SOA ecosystem is that a resource needs to be identifiable.



807  
 808 *Figure 7 Resources*

### 809 Resource

810 A resource is any identifiable entity that has value to a stakeholder.

811 A resource may be identifiable by different methods but within a SOA ecosystem a  
 812 resource must have at least one well-formed identifier that may be unambiguously  
 813 resolved to the intended resource.

814 Codified (but not *implied*) contracts, policies, obligations, and permissions are all  
 815 examples of resources as are capabilities, services, service descriptions, and SOA-  
 816 based systems. An *implied* policy, contract, obligation or permission would not be a  
 817 resource, even though it may have value to a stakeholder, because it is not an  
 818 identifiable entity.

### 819 Identifier

820 An identifier is any sequence of characters that may be unambiguously resolved  
 821 to identifying a particular resource.

822 **Identifiers** typically require a context in order to establish the connection with the  
 823 resource. In a SOA ecosystem, it is good practice to use globally unique identifiers; for  
 824 example globally unique IRIs.

825 A given resource may have multiple identifiers, with different value for different contexts.

826 The ability to identify a resource is important in interactions to determine such things as  
827 rights and authorizations, to understand what functions are being performed and what  
828 the results mean, and to ensure repeatability or characterize differences with future  
829 interactions. The specific subset of individual characteristics that are necessary and  
830 sufficient in order to unambiguously identify a resource depends on the ecosystem  
831 and/or specific interactions within a system. However, in order to enable visibility and  
832 interaction in a SOA ecosystem, those resources that are important to a given SOA  
833 system must be *unambiguously* identifiable at any moment and in any interaction, many  
834 of which may not be predictable given the operation of systems across ownership  
835 boundaries. The way to achieve this is by using identifiers.

### 836 3.1.3.2 Ownership

837 Ownership is defined as a relationship between a stakeholder and a resource, where  
838 some stakeholder (in a role as **owner**) has certain claims with respect to the resource.

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

#### 841 Ownership

842 Ownership is a particular set of claims, expressed as rights and responsibilities,  
843 that a stakeholder has in relation to a resource; It may include the right to transfer  
844 that ownership, or some subset of rights and responsibilities, to another entity.

845 To own a resource implies taking responsibility for creating, maintaining and, if it is to be  
846 available to others, provisioning the resource. More than one stakeholder may own  
847 different rights or responsibilities associated with a given service, such as one  
848 stakeholder having the responsibility to deploy a capability as a service, another owning  
849 the rights to the profits that result from charging consumers for using the service, and  
850 yet another owning the right to use the service.

851 A stakeholder who owns a resource may delegate some or all of these rights and  
852 responsibilities to others, but typically retains the responsibility to see that the delegated  
853 rights and responsibilities are exercised as intended. There may also be joint  
854 ownership of a resource, where the rights and responsibilities are shared.

855 A crucial property that distinguishes ownership from a more limited **right to use** is the  
856 right to transfer rights and responsibilities totally and irrevocably to another stakeholder.  
857 When a stakeholder uses a resource but does not own the resource, that stakeholder  
858 may not transfer the right to use the resource to a third stakeholder. The owner of the  
859 resource maintains the rights and responsibilities of being able to authorize other  
860 stakeholders to use the owned resource.

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

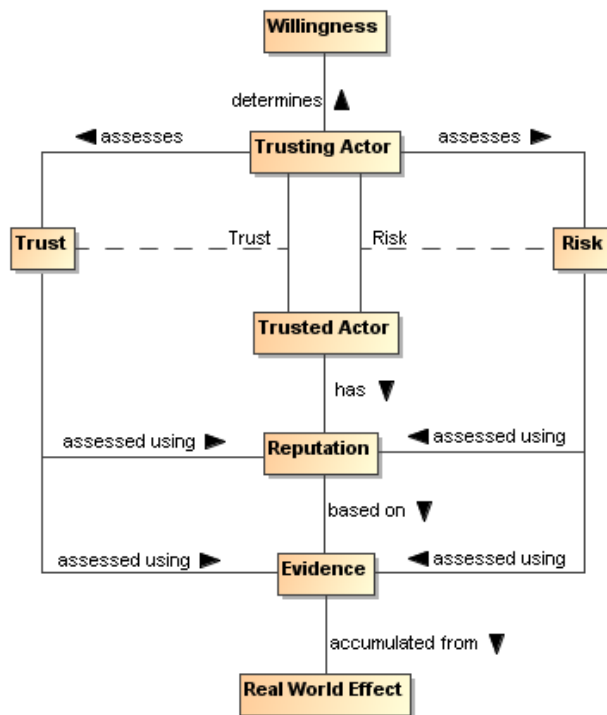
#### 865 Ownership Boundary

866 An ownership boundary is the extent of ownership asserted by a stakeholder  
867 over a set of resources and for which rights and responsibilities are claimed and  
868 (usually) recognized by other stakeholders.

869 In a SOA ecosystem, providers and consumers of services may be, or may be acting on  
 870 behalf of, different owners, and thus the interaction between the provider and the  
 871 consumer of a given service will necessarily cross an ownership boundary. It is  
 872 important to identify these ownership boundaries in a SOA ecosystem, as successfully  
 873 crossing them requires the elements identified in the following sections be addressed.  
 874 Addressing the elements identified in the following sections is referred to in the OASIS  
 875 SOA RM as establishing the execution context.

### 876 3.1.4 Trust and Risk

877 For an interaction to occur each actor must be able and **willing** to participate.



878  
 879 *Figure 8 Willingness and Trust*

### 880 Willingness

881 Willingness is the internal commitment of a human actor to carry out its part of an  
 882 interaction.

883 Willingness to interact is not the same as a willingness to perform requested actions,  
 884 however. For example, a service provider that rejects all attempts to perform a particular  
 885 action may still be fully willing and engaged in interacting with the consumer. Important  
 886 considerations in establishing willingness are both **trust** and **risk**.

### 887 Trust

888 Trust is a private assessment or internal perception of one participant that  
 889 another participant will perform actions in accordance with an assertion regarding  
 890 a desired real world effect.

891 **Risk**

892 Risk is a private assessment or internal perception of the likelihood that certain  
893 undesirable real world effects will result from actions taken, or that the RWE  
894 might not meet certain criteria (e.g., performance), and the consequences or  
895 implications of such.

896 Trust is involved in all interactions – it is necessary for *all* the actors (consumers,  
897 providers, mediators) involved in a given interaction to trust each other at least to the  
898 extent required for continuance of the interaction. The degree and nature of that trust is  
899 likely to be different for each actor, most especially when those actors are in different  
900 ownership boundaries.

901 An actor perceiving risk may take actions to mitigate that risk. At one extreme this will  
902 result in a refusal to interact. Alternately, it may involve adding protection – for example  
903 by using encrypted communication and/or anonymization – to reduce the perception of  
904 risk. Often, standard procedures are put in place to increase trust and to mitigate risk.

905 **Assessing Trust and Risk**

906 The assessments of trust and risk are based on evidence available to the *trusting*  
907 participant. In general, participants will seek evidence directly from the *trusted* actor  
908 (e.g., via documentation provided via the service description) as well as evidence of the  
909 reputation of the trusted actor (e.g., third-party annotations such as consumer  
910 feedback).

911 Trust is based on the confidence that the trusting participant has accurately and  
912 sufficiently gathered and assessed evidence to the degree appropriate for the situation  
913 being assessed.

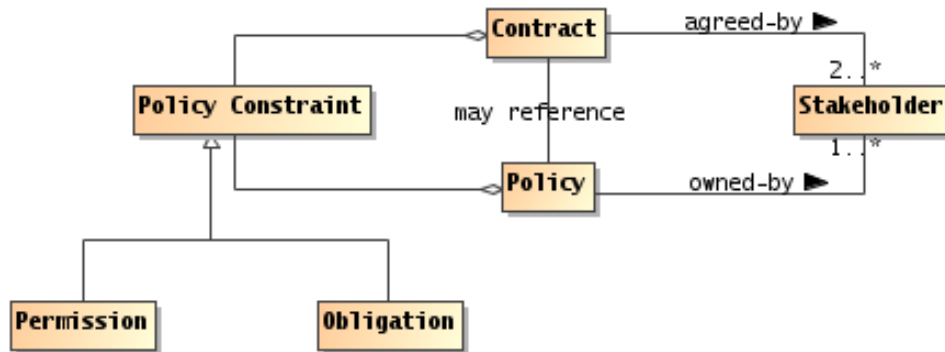
914 Assessment of trust is rarely binary. An actor is not completely trusted or untrusted.  
915 There is typically some degree of uncertainty in the accuracy or completeness of the  
916 evidence or the assessment. Similarly, there may be uncertainty in the amount and  
917 potential consequences of risk.

918 The relevance of trust to interaction depends on the assessment of risk. If there is little  
919 or no perceived risk, or the risk can be covered by another party who accepts  
920 responsibility for it, then the degree of trust may be less or not relevant in assessing  
921 possible actions. For example, most people consider there to be an acceptable level of  
922 risk to privacy when using search engines, and submit queries without any sense of  
923 trust being considered.

924 As perceived risk increases, the issue of trust becomes more of a consideration. For  
925 interactions with a high degree of risk, the trusting participant will typically require  
926 stronger or additional evidence when evaluating the balance between risk and trust. An  
927 example of high-risk is where a consumer's business is dependent on the provider's  
928 service meeting certain availability and security requirements. If the service fails to  
929 meet those requirements, the service consumer will go out of business. In this  
930 example, the consumer will look for evidence that the likelihood of the service not  
931 meeting the performance and security requirements is extremely low.

932 **3.1.5 Policies and Contracts**

933 As noted in the Reference Model, a **policy** represents some commitment and/or  
934 constraint promulgated and enforced by a stakeholder and that stakeholder alone. A  
935 **contract**, on the other hand, represents an agreement by two or more participants.  
936 Enforcement of contracts may or may not be the responsibility of the parties to the  
937 agreement but is usually performed by a stakeholder in the ecosystem (public authority,  
938 legal system, etc.).



939  
940 *Figure 9 Policies and Contracts*

941 **Policy**

942 A policy is an assertion made by a stakeholder which the stakeholder commits to  
943 uphold and, if possible and necessary, enforce through stated constraints.

944 Policies can often be said to be about something – they have an object. For example,  
945 there may be policies about the use of a service. Policies have an **owner** – the  
946 stakeholder who asserts and takes responsibility for the policy. Note that the policy  
947 owner may or may not be the owner of the object of the policy. Thirdly, policies  
948 represent constraints – some measurable limitation on the state or behavior of the  
949 object of the policy, or of the behavior of the stakeholders of the policy.

950 **Contract**

951 A contract represents an agreement made by two or more participants (the  
952 contracting parties) on a set of promises (or contractual terms) together with a  
953 set of constraints that govern their behavior and/or state in fulfilling those  
954 promises.

955 A service provider’s policy may become a service provider/consumer contract when a  
956 service consumer agrees to the provider’s policy. That agreement may be formal, or  
957 may be informal. If a consumer’s policy and a providers policy are mutually exclusive,  
958 then some form of negotiation or mediation to resolve the mutual exclusion before the  
959 service consumer/provider interaction can occur.

960 Both policies and contracts imply a desire to see constraints respected and enforced.  
961 Policies are owned by individual (or aggregate) stakeholders, and contracts are owned  
962 by the parties to the contract; these stakeholders are responsible for ensuring that any  
963 constraints in the policy or contract are enforced – although, of course, the actual  
964 enforcement may be delegated to a different mechanism. A contract does not  
965 necessarily oblige the contracting parties to act (for example to use a service) but it



966 does constraint how they act if and when action covered by the contract occurs (for  
967 example, when a service is invoked and used).

968 Two important types of constraint that are relevant to a SOA ecosystem are permission  
969 and Obligation.

### 970 **Permission**

971 A permission is a constraint that identifies **actions** that an actor is (or is not)  
972 allowed to perform and/or the **states** the actor is (or is not) permitted to be in.

973 Note that permissions are distinct from ability and from authority. Authority refers to the  
974 legitimate nature of an action as performed by an actor on behalf of a social structure  
975 and ability refers to whether an actor has the capacity to perform the action, whereas  
976 permission does not always involve acting on behalf of anyone, nor does it imply or  
977 require the capacity to perform the action.

### 978 **Obligation**

979 An obligation is a constraint that prescribes the actions that an actor must (or  
980 must not) perform and/or the states the actor must (or must not) be in.

981 An example of obligations is the case where the service consumer and provider have  
982 entered into an agreement to provide and consume a service such that the consumer is  
983 obligated to pay for the service and the provider is obligated to provide the service –  
984 based on the terms of the contract.

985 An obligation can also be a requirement to to *maintain* a given state. This may range  
986 from a requirement to maintain a minimum balance on an account to a requirement that  
987 a service provider ‘remember’ that a particular service consumer is logged in.

988 Both permissions and obligations can be identified ahead of time, but only Permissions  
989 can be validated a priori: before the intended action or before entering the constrained  
990 state. Obligations can only be validated a posteriori through some form of auditing or  
991 verification process.

## 992 **3.1.6 Communication**

### 993 **Communication**

994 A communication is a process of reaching mutual understanding, in which  
995 participants not only exchange information as messages but also create and  
996 share meaning..

997 A communication involves one or more actors playing the role of **sender** and at least  
998 one other actor playing the role of **recipient**; all actors must perform their part in order  
999 for the communication to occur.

1000 A given communication may involve any number of **recipients**. In some situations, the  
1001 sender may not be aware of the recipient. However, without both a sender and a  
1002 recipient there is no communication. A given communication does not necessarily  
1003 involve interaction between the actors; it can be a simple one-way transmission  
1004 requiring no further action by the recipient. However, interaction does, necessarily,  
1005 involve communication.

1006 A communication involves a message, which an actor receiving must be able to  
1007 correctly interpret. The extent of that correct interpretation depends on the role of the  
1008 actor and the purpose of the communication.

1009 A communication is not effective unless the recipient can correctly interpret the  
1010 message. However, interpretation can itself be characterized in terms of semantic  
1011 engagement: the proper understanding of a message in a given context.

1012 We can characterize the necessary modes of interpretation in terms of a shared  
1013 understanding of a common vocabulary and of the purpose of the communication. More  
1014 formally, we can say that a communication has a combination of message and purpose.

1015 Interactions between service consumers and providers do not need to resemble human  
1016 speech. Machine-machine communication is typically highly stylized in form, it may  
1017 have particular forms and it may involve particular terms not found in everyday human  
1018 communication.

### 1019 **3.1.7 Semantics and Semantic Engagement**

1020 A SOA ecosystem is a space in which actors need to share understanding<sup>10</sup> as well as  
1021 sharing actions. Indeed, such shared understanding is a pre-requisite to a joint action  
1022 being carried out as intended. It is vital to a trusted and effective ecosystem. Semantics  
1023 are therefore pervasive throughout SOA ecosystems and important in communicative  
1024 actions described above, as well as a driver for policies and other aspects of the  
1025 ecosystem.

1026 In order to arrive at shared understanding, an actor must effectively process and  
1027 understand assertions in a manner appropriate to the particular context. An assertion, in  
1028 general, is a measurable and explicit statement made by an actor. In a SOA ecosystem,  
1029 in particular, assertions are concerned with the 'what' and the 'why' of the state of the  
1030 ecosystem and its actors.

1031 Understanding and interpreting those assertions allows other actors to know what may  
1032 be expected of them in any particular joint action. An actor can potentially 'understand'  
1033 an assertion in a number of ways, but it is specifically the process of arriving at a *shared*  
1034 understanding that is important in the ecosystem. This process is semantic engagement  
1035 by the actor with the SOA ecosystem. It can be instantaneous or progressively  
1036 achieved. It is important that there is a level of engagement appropriate to the particular  
1037 context.

#### 1038 **Semantic Engagement**

1039 Semantic engagement is the process by which an actor engages with a set of  
1040 assertions based on that actor's interpretation and understanding of those  
1041 assertions.

1042 Different actors have differing capabilities and requirements for understanding  
1043 assertions. This is true for both human and non-human actors. For example, a purchase  
1044 order process does not require that a message forwarding agent 'understand' the

---

<sup>10</sup> We use a mechanical, Turing test-based approach to understanding here: if an actor behaves as though it understands an utterance then we assume that it does understand it.

1045 purchase order, but a processing agent does need to ‘understand’ the purchase order in  
1046 order to know what to with the order once received.

1047 The impact of any assertion can only be fully understood in terms of specific social  
1048 contexts; contexts that necessarily include the actors that are involved. For example, a  
1049 policy statement that governs the actions relating to a particular resource may have a  
1050 different impact or purpose for the participant that owns the resource than for the actor  
1051 that is trying to access it: the former understands the purpose of the policy as a  
1052 statement of enforcement; and the latter understands it as a statement of constraint.

### 1053 **3.2 Action in a SOA Ecosystem Model**

1054 Participants cannot always achieve desired results leveraging resources in their own  
1055 ownership domain; thus generating a need for which they look for and leverage services  
1056 provided by other participants, using resources beyond their ownership and control;  
1057 They identify service providers with which they think they can interact to achieve their  
1058 objective; They thus engage in joint action with those other actors (service providers) in  
1059 order to bring about the desired outcome; the SOA ecosystem provides the environment  
1060 to make this happen.

1061 An action model is put forth a-priori by the service provider, and is effectively a promise  
1062 by the service provider that the actions identified in the action model and invoked  
1063 consistent with the process model will result in the described real world effect. Action  
1064 model is basically a description of the actions that the service is willing to do on behalf  
1065 of another. They should be associated with a real-world effect. The potential service  
1066 consumer is interested in accessing or acquiring the real-world effect, and the action  
1067 model identifies the actions that the service consumer will have to be a party to in order  
1068 to access or generate the real-world effect.

1069 When the consumer “invokes” a service, a joint action is started as identified in the  
1070 action model, consistent with the temporal sequence as defined by the process model,  
1071 and where the consumer and the provider are the two parties of the joint action.  
1072 Additionally, the consumer can be assured that the identified real-world effects will be  
1073 accomplished through evidence provided via the service description.

1074 Since the service provider does not know about all potential service consumers, the  
1075 service provider may also describe what additional constraints are necessary in order  
1076 for the service consumer to invoke particular actions, and thus participate in the joint  
1077 action. These additional constraints, along with others that might not be listed, are  
1078 preconditions for the joint action to occur and/or continue (as per the process model),  
1079 and are referred to in the SOA RM as execution context. Execution context goes all the  
1080 way from human beings involved in aligning policies, semantics, network connectivity  
1081 and communication protocols, to the automated negotiation of security protocols and  
1082 end-points as the individual actions proceed through the process model.

1083 Also, it is important to note that both actions and RWE are ‘fractal’ in nature, in the  
1084 sense that they can often be broken down into more and more granularity depending on  
1085 how they are examined and what level of detail is important.

1086 All of these things are important to getting to the core of participants’ interest in a SOA  
1087 ecosystem: the ability to leverage resources or capabilities to achieve a desired

1088 outcome, and in particular where those resources or capabilities do not belong to them  
1089 or are beyond their direct control. i.e., that are outside of their ownership boundary.  
1090 In order to use such resources, participants must be able to identify their own needs in  
1091 the form of requirements, identify and compose into a business solution those resources  
1092 or capabilities that will meet their needs, and engage in joint action – the coordinated  
1093 set of actions that participants pursue in order to achieve measurable results in  
1094 furtherance of their goals.  
1095 In order to act in a way that is appropriate and consistent both to their own goals,  
1096 objectives and policies, and those of others, participants must also communicate with  
1097 each other.  
1098 A key aspect of joint action revolves around the trust that both parties must exhibit in  
1099 order to participate in the joint action. The willingness to act and a mutual understanding  
1100 of both the information exchanged and the expected results is the particular focus of  
1101 Sections **Error! Reference source not found.**6 and 3.1.7.

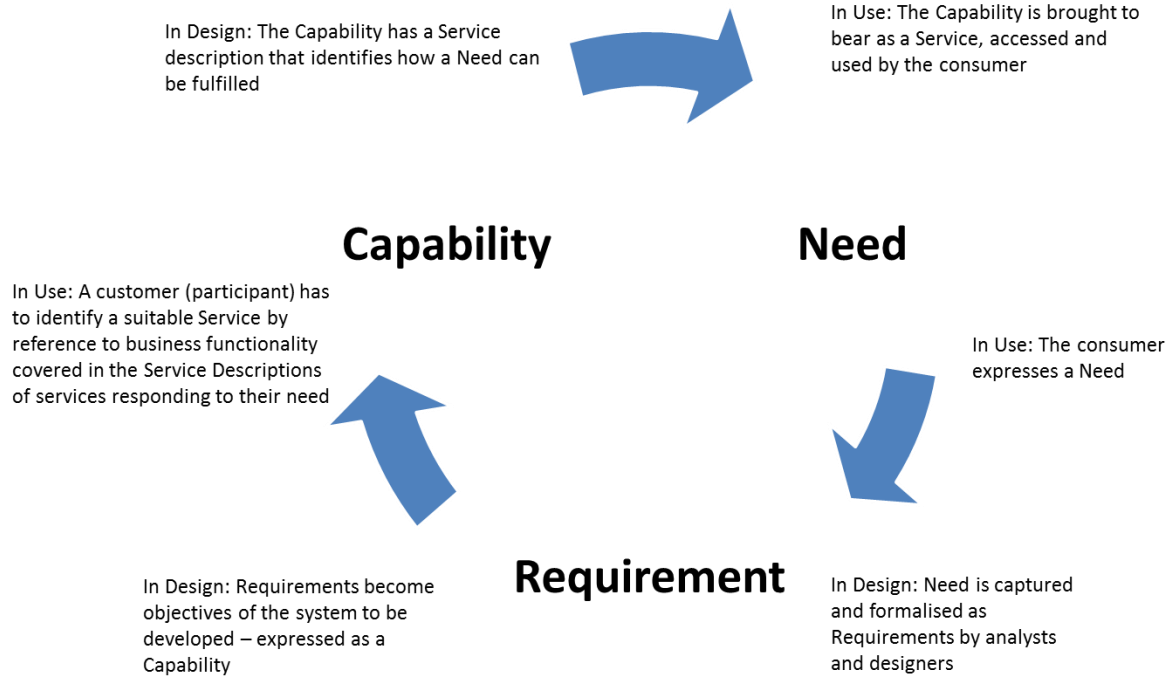
### 1102 **3.2.1 Needs, Requirements and Capabilities**

1103 Participants in a SOA ecosystem often need other participants to *do* something,  
1104 leveraging a capability that they do not themselves possess. For example, a customer  
1105 requiring a book may call upon a service provider to deliver the book. Likewise, the  
1106 service provider needs the customer to pay for it.

1107 There is a reason that participants are engaged in this activity: different participants  
1108 have different **needs** and have or apply different **capabilities** for satisfying them. These  
1109 are core to the concept of a service. The SOA-RM defines a service as “the mechanism  
1110 by which needs and capabilities are brought together”. This idea of services being a  
1111 mechanism “between” needs and capabilities was introduced in order to emphasize  
1112 capability as the notional or existing business functionality that would address a well-  
1113 defined need. Service is therefore the *implementation* of such business functionality  
1114 *such that it is accessible* through a well-defined interface. A capability that is isolated, or  
1115 by itself (i.e., not accessible to potential consumers) is emphatically not a service.

### 1116 **Business functionality**

1117 Business functionality is a defined set of business-aligned tasks that provide  
 1118 recognizable business value to ‘consumer’ stakeholders and possibly others in  
 1119 the SOA ecosystem.



1120 *Figure 10 Need, Requirement and Capability*

1121 The idea of a service in a SOA ecosystem combines business functionality with  
 1122 implementation, including the artifacts needed and made available as IT resources.  
 1123 From the perspective of software developers, a SOA service enables the use of  
 1124 capabilities in an IT context. For the consumer, the service (combining business  
 1125 functionality and implementation) generates intended real world effects. The consumer  
 1126 is not concerned with the underlying artifacts which make that delivery possible.

1127 *Figure 11 - Relationship between Need, Requirement and Capability*

1128 In a SOA context, the consumer (as a stakeholder) expresses a need (“I want to buy a  
 1129 book”) and looks to an appropriate service to fulfill that need and assesses issues such  
 1130 as the trustworthiness, intent and willingness of a particular provider. This ecosystem  
 1131 communication continues up to the point when the consumer is ready to act. The  
 1132 consumer (as an actor now) will then interact with a provider by invoking a service (for  
 1133 example, ordering the book using an online bookseller) and engaging in relevant actions  
 1134 (validating the purchase, submitting billing and delivery details) within the system with a  
 1135 view to achieving the desired Real World Effect (having the book delivered).

1136 **Need**

1137 A need is a general statement expressed by a stakeholder of the lack of  
 1138 something deemed necessary. It may be formalized as one or more  
 1139 **requirements** that must be fulfilled in order to achieve a stated goal.

## 1140 **Requirement**

1141 A requirement is a formal statement of a desired result (a real world effect) that, if  
1142 achieved, will satisfy a need.

1143 This requirement can then be used to create a capability that in turn can be brought to  
1144 bear to satisfy that need. Both the requirement and the capability to fulfill it are  
1145 expressed in terms of desired real world effect.

## 1146 **Capability**

1147 A capability is an ability to achieve a real world effect.

1148 The Reference Model makes a distinction between a capability (as a potential to  
1149 generate a real world effect) and the ability of bringing that capability to bear (via a  
1150 realized service) as the realization of the real world effect.

## 1151 **3.2.2 Services Reflecting Business**

1152 The SOA paradigm often emphasizes the prescribed interface through which service  
1153 interaction is accomplished. While this enables predictable integration in the sense of  
1154 traditional software development, the prescribed interface alone does not guarantee that  
1155 services will be composable into business solutions.

## 1156 **Business solution**

1157 A **business solution** is a set of defined interactions that combine implemented  
1158 or notional business functionality in order to address a set of business needs.

## 1159 **Composability**

1160 **Composability** is the ability to combine individual services, each providing  
1161 defined business functionality, so as to provide more complex business solutions.

1162 Composability is important because many of the benefits of a SOA approach assume  
1163 multiple uses for services, and multiple use requires that the service deliver a business  
1164 function that is reusable in multiple business solutions.

1165 To achieve composability, capabilities must be identified that serve as building blocks  
1166 for business solutions. In a SOA ecosystem, these building blocks are captured as  
1167 services representing well-defined business functions, operating under well-defined  
1168 policies and other constraints, and generating well-defined real world effects. These  
1169 service building blocks should be relatively stable so as not to force repeated changes  
1170 in the compositions that utilize them, but should also embody SOA attributes that readily  
1171 support creating compositions that can be varied to reflect changing circumstances.

1172 The SOA paradigm emphasizes both composition of services and opacity of how a  
1173 given service is implemented. With respect to opacity, the SOA-RM states that the  
1174 service could carry out its described functionality through one or more automated and/or  
1175 manual processes that in turn could invoke other available services.

1176 Any composition can itself be made available as a service and the details of the  
1177 business functionality, conditions of use, and effects are among the information  
1178 documented in its service description.

1179 For services to be useful as composable building blocks in the SOA ecosystem, the  
1180 services should, whenever possible, deliver capability that is applicable to multiple  
1181 needs. Simply providing a Web Service interface for an existing IT artifact does not, in

1182 general, create opportunities for sharing business functions. Furthermore, the use of  
1183 tools to auto-generate service software interfaces will not guarantee services than can  
1184 effectively be used within compositions if the underlying code represents programming  
1185 constructs rather than business functions. In such cases, services that tightly reflect the  
1186 software details will be as brittle to change as the underlying code and will not exhibit  
1187 the undefined but intuitive characteristic of loose coupling.

### 1188 **3.2.3 Action, Communication and Joint Action**

1189 In general terms, entities act in order to achieve their goals. However, the form of action  
1190 that is of most interest within a SOA ecosystem is that involving interaction across  
1191 ownership boundaries (between more than one actor) – **joint action**.

#### 1192 **3.2.3.1 Action and Actors**

##### 1193 **Action**

1194 An action is the application of intent to cause an effect.

1195 The aspect of action that distinguishes it from mere force or accident is that someone  
1196 *intends* that the action achieves a desired objective or effect. This definition of action is  
1197 very general. In the case of SOA, we are mostly concerned with actions that take place  
1198 within a system and have specific effects on the SOA ecosystem – what we call **Real**  
1199 **World Effects**. The actual real world effect of an action, however, may go beyond the  
1200 intended effect.

1201 Objectives refer to real world effects that participants believe are achievable by a  
1202 specific action or set of actions that deliver appropriate changes in shared state. In  
1203 contrast, a goal is not expressed in terms of specific action but rather in terms of desired  
1204 end state.

1205 For example, someone may wish to have enough light to read a book. In order to satisfy  
1206 that goal, the reader walks over to flip a light switch. The *objective* is to change the state  
1207 of the light bulb, by turning on the lamp, whereas the *goal* is to be able to read. The *real*  
1208 *world effect* is more light being available to enable the person to read.

1209 While an effect is any measurable change resulting from an action, a SOA ecosystem is  
1210 concerned more specifically with real world effects.

##### 1211 **Real World Effect**

1212 A real world effect is a measurable change to the shared state of pertinent  
1213 entities, relevant to and experienced by specific stakeholders of an ecosystem.

1214 This implies measurable change in the overall state of the SOA ecosystem. [In practice,](#)  
1215 [however, it is specific state changes of certain entities that are relevant to particular](#)  
1216 [participants that constitute the real world effect as experienced by those participants.](#)

#### 1217 **3.2.3.2 Communication and Joint Actions**

1218 In this Reference Architecture Foundation, we are concerned with two levels of activity:  
1219 as communication and as participants engaged in joint actions to use and offer services.

1220 In order for multiple actors to participate in a joint action, they must each act according  
1221 to their role within the joint action. This is achieved through communication and  
1222 messaging.

1223 Communication – the formulation, transmission, receipt and interpretation of messages  
1224 – is the foundation of all joint actions within the SOA ecosystem, given the inherent  
1225 separation – often across ownership boundaries – of actors in the system.

1226 Communication between actors requires that they play the roles of ‘sender’ or ‘receiver’  
1227 of messages as appropriate to a particular action – although it is not necessarily  
1228 required that they both be active simultaneously.

1229 An actor sends a message in order to communicate with other actors. The  
1230 communication itself is often not intended as part of the desired real world effect but  
1231 rather includes messages that seek to establish, manage, monitor, report on, and guide  
1232 the joint action throughout its execution.

1233 Like communication, joint action usually involves different actors. However, joint action  
1234 – resulting from the deliberate actions undertaken by different actors – *intentionally*  
1235 impacts shared state within the system leading to real world effects.

### 1236 **Joint Action**

1237 Joint action is the coordinated set of actions involving the efforts of two or more  
1238 actors to achieve an effect.

1239 Note that the effect of a joint action is *not* always equivalent to one or more effects of  
1240 the individual actions of the participating actors, i.e., it may be more than the sum of the  
1241 parts.

1242 Different viewpoints lead to either communication or joint action as being considered  
1243 most important. For example, from the viewpoint of ecosystem security, the integrity of  
1244 the communications may be dominant; from the viewpoint of ecosystem governance,  
1245 the integrity of the joint action may be dominant.

## 1246 **3.2.4 State, Shared State and Real-World Effect**

### 1247 **State**

1248 State is the condition of an entity at a particular time.

1249 State is characterized by a set of facts that is true of the entity. In principle, the total  
1250 state of an entity (or the world as a whole) is unbounded. In practice, we are concerned  
1251 only with a subset of the State of an entity that is measurable and useful in a given  
1252 context.

1253 For example, the total state of a lightbulb includes the temperature of the filament of the  
1254 bulb. It also includes a great deal of other state – the composition of the glass, the dirt  
1255 that is on the bulb’s surface and so on. However, an actor may be primarily interested in  
1256 whether the bulb is ‘on’ or ‘off’ and not on the amount of dirt accumulated. That actor’s  
1257 characterization of the state of the bulb reduces to the fact: ‘bulb is now on’.

1258 In a SOA ecosystem, there is a distinction between the set of facts about an entity that  
1259 only that entity can access – the so-called Private State – and the set of facts that may  
1260 be accessible to other actors in the SOA-based system – the public or Shared State.



1261 **Private State**

1262 The private state is that part of of an entity's state that is knowable by, and  
1263 accessible to, only that entity.

1264 **Shared State**

1265 Shared state is that part of an entity's state that is knowable by, and may be  
1266 accessible to, other actors.

1267 Note that shared state does not imply that the state *is* accessible to *all* actors. It simply  
1268 refers to that subset of state that *may* be accessed by *other* actors. Generally this will  
1269 be the case when actors need to participate in joint actions.

1270 It is the aggregation of the shared states of pertinent entities that constitutes the desired  
1271 effect of a joint action. Thus the change to this shared state is what is experienced in  
1272 the wider ecosystem as a real world effect

1273 **3.3 Architectural Implications**

1274 **3.3.1 Social structures**

1275 A SOA ecosystem's participants are organized into various forms of social structure.  
1276 Not all social structures are hierarchical: a SOA ecosystem should be able to  
1277 incorporate peer-to-peer forms of organization as well as hierarchic structures. In  
1278 addition, it should be possible to identify and manage any constitutional agreements  
1279 that define the social structures present in a SOA ecosystem.

- 1280 • Different social structures have different rules of engagement
  - 1281 ○ Techniques for expressing constitutions are important
- 1282 • social structures have roles and members
  - 1283 ○ Techniques for identifying, managing members of social structures
  - 1284 ○ Techniques for describing roles and role adoption
- 1285 • social structures may be complex
  - 1286 ○ Child social structures' constitutions depend on their parent constitutions
- 1287 • Social structures overlap and interact
  - 1288 ○ A given actor may be member of multiple social structures
  - 1289 ○ Social structures may be associated with different jurisdictions
  - 1290 ○ Social structures may involved in disputes with one another
    - 1291 ▪ Requiring conflict resolution
  - 1292 ○ Social structures inform and limit the "kinds" of governance that can be
  - 1293 effectively deployed

1294 **3.3.2 Resource and Ownership**

1295 Communication about and between, visibility into, and leveraging of resources requires  
1296 the unambiguous identification of those resources. Ensuring unambiguous identities  
1297 implies

- 1298 • Mechanism for assigning and guaranteeing uniqueness of globally unique  
1299 identifiers
- 1300 • Identifying the extent of the enterprise over which the identifier needs to be  
1301 understandable and unique

- 1302       • Mechanism and framework for ensuring the long-livedness of identifiers (i.e., they  
1303       cannot just change arbitrarily)

### 1304   **3.3.3 Policies and Contracts**

- 1305       • Policies are constraints
- 1306           ○ It is necessary to be able to express required policies
  - 1307           ○ It is necessary to be able to enforce the constraints
  - 1308           ○ It is necessary to manage potentially large numbers of policies
- 1309       • Policies have owners
- 1310           ○ The right to establish policies is an aspect of the social structure.
- 1311       • Policies may not be consistent with one another
- 1312           ○ Policy conflict resolution techniques
- 1313       • Agreements are constraints agreed to
- 1314           ○ Contracts often need to be enforced by mechanisms of the social structure

### 1315   **3.3.4 Communications as a Means of Mediating Action**

1316   Using message exchange for mediating action implies

- 1317       • Ensuring correct identification of the structure of messages:
- 1318           ○ Identifying the syntax of the message;
  - 1319           ○ Identifying the vocabularies used in the communication
  - 1320           ○ Identifying the higher-level structure such as the illocutionary form of the  
1321           communication
- 1322       • A principal objective of communication is to mediate action
- 1323           ○ Messages convey actions and events
  - 1324           ○ Receiving a message is an action, but is not the same action as the action  
1325           conveyed by the message
  - 1326           ○ Actions are associated with objectives of the actors involved
  - 1327           ▪ Explicit representation of objectives may facilitate automated  
1328           processing of messages
  - 1329           ○ An actor agreeing to adopt an objective becomes responsible for that  
1330           objective

### 1331   **3.3.5 Semantics**

1332   Semantics is pervasive in a SOA ecosystem. There are many forms of utterance that  
1333   are relevant to the ecosystem: apart from communicated content there are policy  
1334   statements, goals, purposes, descriptions, and agreements which are all forms of  
1335   utterance.

1336   The operation of the SOA ecosystem is significantly enhanced if

- 1337       • A careful distinction is made between public semantics and private semantics. In  
1338       particular, it **MUST** be possible for actors to process content such as  
1339       communications, descriptions and policies solely on the basis of the public  
1340       semantics of those utterances.
- 1341       • A well founded semantics ensures that any assertions that are essential to the  
1342       operator of the ecosystem (such as policy statements, and descriptions) have  
1343       carefully chosen written expressions and associated decision procedures.

- 1344
- 1345
- 1346
- 1347
- The role of vocabularies as a focal point for multiple actors to be able to understand each other is critical. While no two actors can fully share their interpretation of elements of vocabularies, ensuring that they do understand the public meaning of vocabularies' elements is essential.

### 1348 **3.3.6 Trust and Risk**

1349 In traditional systems, the balance between trust and risk is achieved by severely  
1350 restricting interactions and by controlling the participants of a system.

1351 It is important that actors are able to explicitly reason about both trust and risk in order  
1352 to effectively participate in a SOA ecosystem. The more open and public the SOA  
1353 ecosystem is, the more important it is for actors to be able to reason about their  
1354 participation.

### 1355 **3.3.7 Needs, Requirements and Capabilities**

1356 In the process of capturing needs as requirements, and the subsequent requirements  
1357 decomposition and allocation processes need to be informed by capabilities that already  
1358 exist.

- 1359
- 1360
- Architecture needs to
    - Take into account existing capabilities available as services

### 1361 **3.3.8 The Importance of Action**

1362 Participants participate in a SOA ecosystem in order to get their needs met. This  
1363 involves action; both individual actions and joint actions.

1364 Any architectural realization of a SOA ecosystem should address:

- 1365
- 1366
- 1367
- 1368
- How actions are modeled:
    - Identifying the performer or agent of the action;
    - the target of the action; and the
    - verb of the action.

1369 Any explicit models of joint action should take into account

- 1370
- 1371
- The choreography that defines the joint action.
  - The potential for multiple joint actions to be layered on top of each other

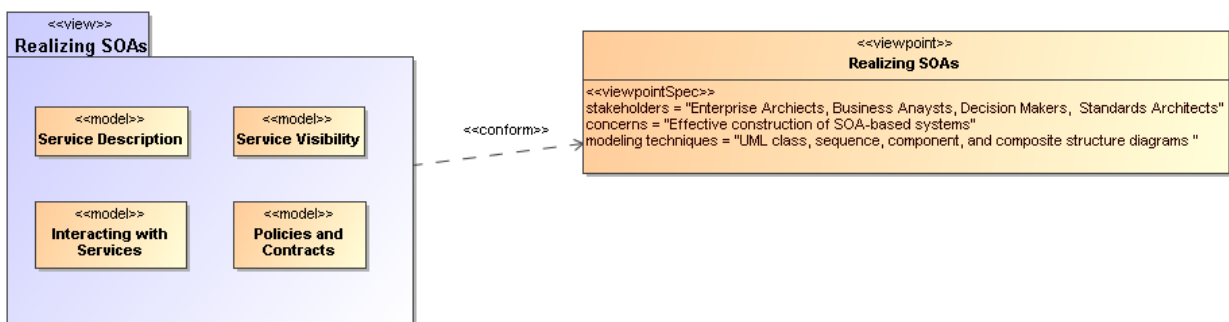
## 1372 4 Realization of a SOA Ecosystem view

1373  
1374  
1375

*Make everything as simple as possible but no simpler.*  
Albert Einstein

1376 The *Realization of a SOA Ecosystem* view focuses on the infrastructure elements that  
1377 are needed in order to support the discovery and interaction with services. The key  
1378 questions asked are "What are services, what support is needed and how are they  
1379 realized?"

1380 The models in this view include the Service Description Model, the Service Visibility  
1381 Model, the Interacting with Services Model, and the Policies and Contracts Model.



1382  
1383

Figure 12 Model Elements Described in the Realization of a SOA Ecosystem view

1384 The Service Description Model informs the participants of what services exist and the  
1385 conditions under which these can be used. Some of those conditions follow from  
1386 policies and agreements on policy that flow from the Policies and Contracts Model. The  
1387 information in the service description as augmented by details of policy provides the  
1388 basis for visibility as defined in the SOA Reference Model and captured in the Service  
1389 Visibility Model. Finally, the process by which services as described are used under the  
1390 defined conditions and agreements is described in the Interacting with Services Model.

### 1391 4.1 Service Description Model

1392 A service description is an artifact, usually document-based, that defines or references  
1393 the information needed to use, deploy, manage and otherwise control a service. This  
1394 includes not only the information and behavior models associated with a service to  
1395 define the service interface but also includes information needed to decide whether the  
1396 service is appropriate for the current needs of the service consumer. Thus, the service  
1397 description will also include information such as service reachability, service  
1398 functionality, and the policies and contracts associated with a service.

1399 A service description artifact may be a single document or it may be an interlinked set of  
1400 documents. For the purposes of this model, differences in representation are to be  
1401 ignored, but the implications of a "web of documents" is discussed later in this section.

1402 There are several points to note regarding the following discussion of service  
1403 description:

- 1404 • The Reference Model states that one of the hallmarks of SOA is the large amount of  
1405 associated description. The model presented below focuses on the description of  
1406 services but it is equally important to consider the descriptions of the consumer,  
1407 other participants, and needed resources other than services.
- 1408 • Descriptions are inherently incomplete but may be determined as *sufficient* when it is  
1409 possible for the participants to access and use the described services based only on  
1410 the descriptions provided. This means that, at one end of the spectrum, a description  
1411 along the lines of “*That service on that machine*” may be sufficient for the intended  
1412 audience. On the other extreme, a service description with a machine-process-able  
1413 description of the semantics of its operations and real world effects may be required  
1414 for services accessed via automated service discovery and planning systems.
- 1415 • Descriptions come with context, i.e. a given description comprises information  
1416 needed to adequately support the context. For example, a list of items can define a  
1417 version of a service, but for many contexts an indicated version number is sufficient  
1418 without the detailed list. The current model focuses on the description needed by a  
1419 service consumer to understand what the service does, under what conditions he  
1420 service will do it, how well does the service do it, and what steps are needed by the  
1421 consumer to initiate and complete a service interaction. Such information also  
1422 enables the service provider to clearly specify what is being provided and the  
1423 intended conditions of use.
- 1424 • Descriptions change over time as, for example, the ingredients and nutrition  
1425 information for food labeling continues to evolve. A requirement for transparency of  
1426 transactions may require additional description for those associated contexts.
- 1427 • Description always proceeds from a basis of what is considered “common  
1428 knowledge”. This may be social conventions that are commonly expected or possibly  
1429 codified in law. It is impossible to describe everything and it can be expected that a  
1430 mechanism as far reaching as SOA will also connect entities where there is  
1431 inconsistent “common” knowledge.
- 1432 • Descriptions will become the collection point of information related to a service or  
1433 any other resource, but it is not necessarily the originating point or the motivation for  
1434 generating this information. In particular, given a SOA service as the access to an  
1435 underlying capability, the service may point to some of the capability’s previously  
1436 generated description, e.g. a service providing access to a data store may reference  
1437 update records that indicate the freshness of the data.
- 1438 • Descriptions of the provider and consumer are the essential building blocks for  
1439 establishing the execution context of an interaction.

1440 These points emphasize that there is no one “right” description for all contexts and for  
1441 all time. Several descriptions for the same subject may exist at the same time, and this  
1442 emphasizes the importance of the description referencing source material maintained  
1443 by that material’s owner rather than having multiple copies that become out of synch  
1444 and inconsistent.

1445 It may also prove useful for a description assembled for one context to cross-reference  
1446 description assembled for another context as a way of referencing ancillary information  
1447 without overburdening any single description. Rather than a single artifact, description  
1448 can be thought of as a web of documents that enhance the total available description.

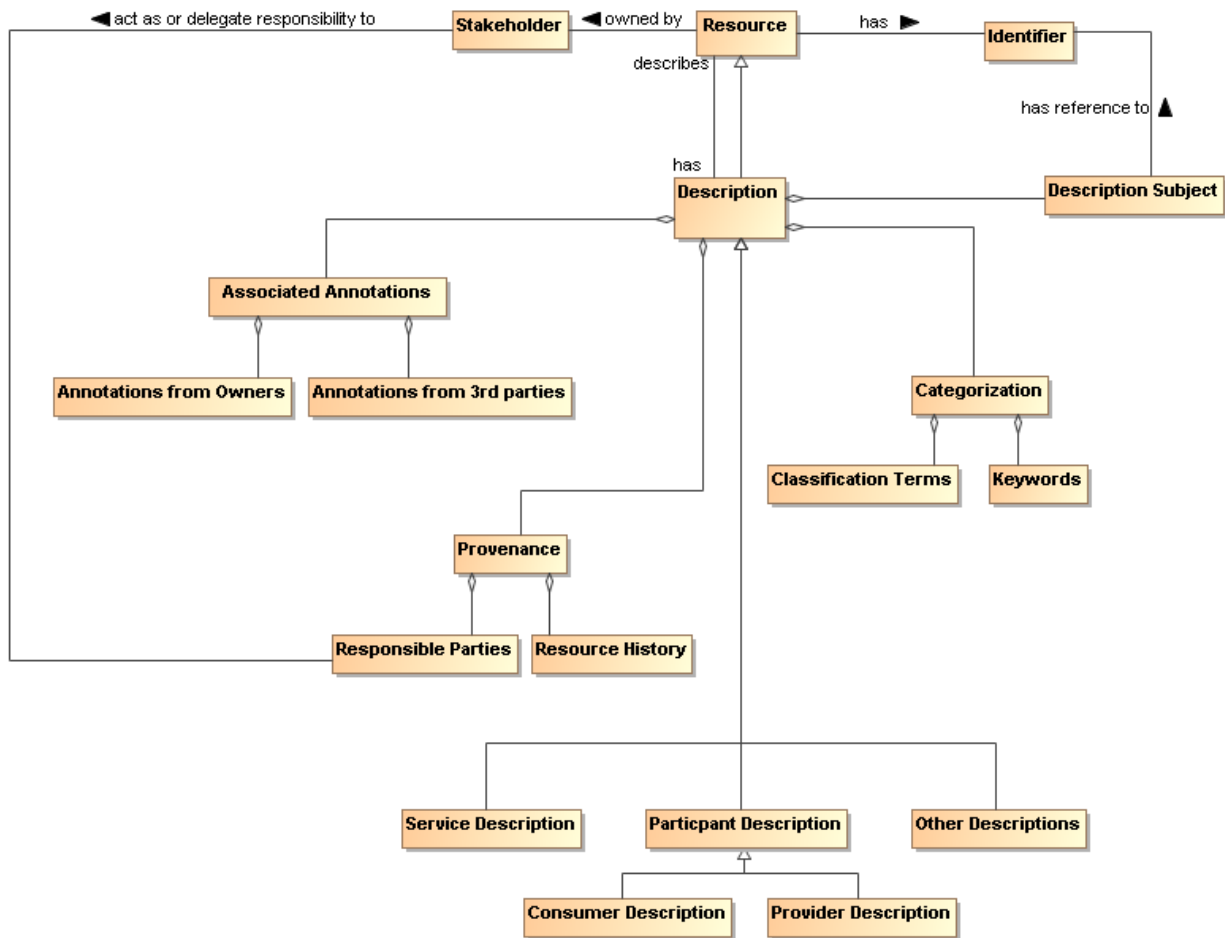
1449 This Reference Architecture Foundation uses the term service description for  
1450 consistency with the concept defined in the Reference Model. Some SOA literature  
1451 treats the idea of a “service contract” as equivalent to service description. In the SOA-  
1452 RAF, the term service description is preferred. Replacing service description with  
1453 service contract implies just one side of the interaction is governing and misses the  
1454 point that a single set of policies identified by a service description may lead to  
1455 numerous contracts, i.e. service level agreements, leveraging the same description.

#### 1456 **4.1.1 The Model for Service Description**

1457 *Figure 13* shows Service Description as a subclass of the general Description class,  
1458 where Description is a subclass of the resource class as defined in Section 3.1.5.1. In  
1459 addition, each resource is assumed to have a description. The following section  
1460 discusses the relationships among elements of general description and the subsequent  
1461 sections focus on service description itself. Other descriptions, such as those of  
1462 participants, are important to SOA but are not individually elaborated in this document.

##### 1463 **4.1.1.1 Elements Common to General Description**

1464 The general Description class is composed of a number of elements that are expected  
1465 to be common among all specialized descriptions supporting a service oriented  
1466 architecture. A registry often contains a subset of the description instance, where the  
1467 chosen subset is identified as that which facilitates mediated discovery. Additional  
1468 information contained in a more complete description may be needed to initiate and  
1469 continue interaction.



1470  
1471 *Figure 13 General Description*

1472 **4.1.1.1.1 Description Subject**

1473 The subject of a description is a resource. The value assigned to the Description  
1474 Subject class may be of any form that provides understanding of what constitutes the  
1475 resource, but it is often in human-readable text. The Description Subject MUST also  
1476 reference the Identifier of the resource it describes so it can unambiguously identify the  
1477 subject of each description instance.

1478 As a resource, Description also has an identifier with a unique value for each  
1479 description instance. The description instance provides vital information needed to both  
1480 establish visibility of the resource and to support its use in the execution context for the  
1481 associated interaction. The identifier of the description instance allows the description  
1482 itself to be referenced for discussion, access, or reuse of its content.

1483 **4.1.1.1.2 Provenance**

1484 While the resource Identifier provides the means to know which subject and subject  
1485 description are being considered, Provenance as related to the Description class  
1486 provides information that reflects on the quality or usability of the subject. Provenance  
1487 specifically identifies the entity (human, defined role, organization, ...) that assumes  
1488 responsibility for the resource being described and tracks historic information that

1489 establishes a context for understanding what the resource provides and how it has  
1490 changed over time. Responsibilities may be directly assumed by the stakeholder who  
1491 owns a resource or the Owner may designate Responsible Parties for the various  
1492 aspects of maintaining the resource and provisioning it for use by others. There may be  
1493 more than one entity identified under Responsible Parties; for example, one entity may  
1494 be responsible for code maintenance while another is responsible for provisioning of the  
1495 executable code. The historical aspects may also have multiple entries, such as when  
1496 and how data was collected and when and how it was subsequently processed, and as  
1497 with other elements of description, may provide links to other assets maintained by the  
1498 resource owner.

#### 1499 **4.1.1.1.3 Keywords and Classification Terms**

1500 A traditional element of description has been to associate the resource being described  
1501 with predefined keywords or classification taxonomies that derive from referenceable  
1502 formal definitions and vocabularies. This Reference Architecture Foundation does not  
1503 prescribe which vocabularies or taxonomies may be referenced, nor does it limit the  
1504 number of keywords or classifications that may be associated with the resource. It  
1505 does, however, state that a normative definition SHOULD be referenced, whether that  
1506 be a representation in a formal ontology language, a pointer to an online dictionary, or  
1507 any other accessible source. See Section 4.1.1.2 for further discussion on associating  
1508 semantics with assigned values.

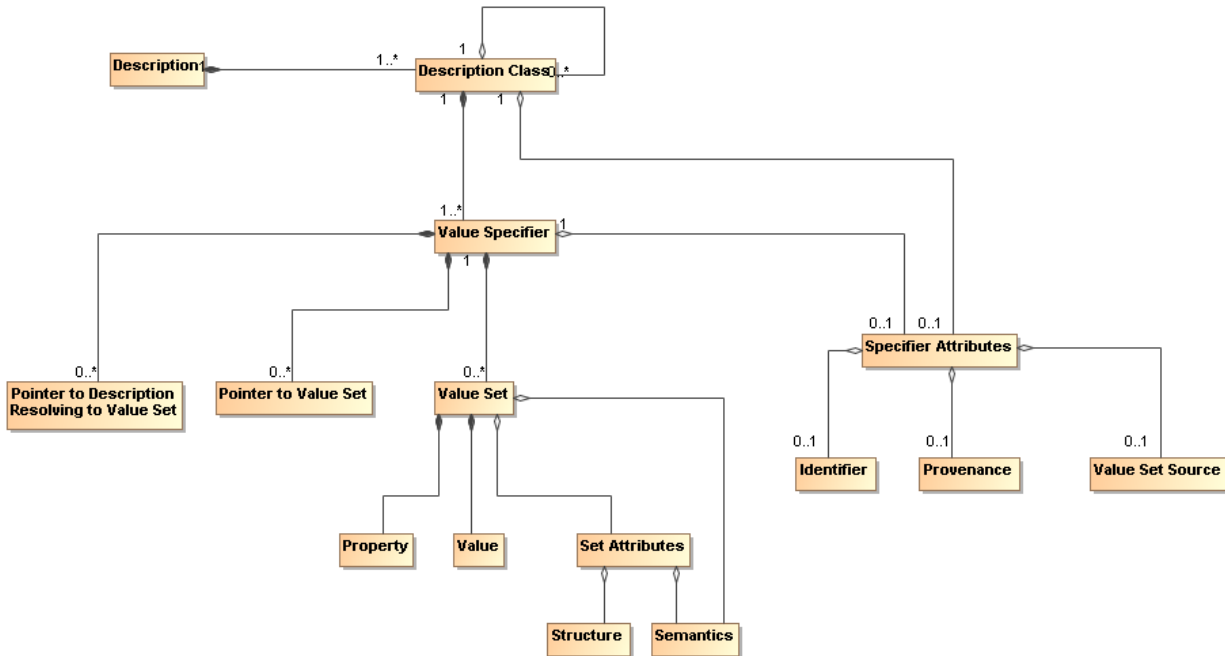
#### 1509 **4.1.1.1.4 Associated Annotations**

1510 The general description instance may also reference associated documentation that is  
1511 in addition to that considered necessary in this model. For example, the owner of a  
1512 service may have documentation on best practices for using the service. Alternately, a  
1513 third party may certify a service based on their own criteria and certification process;  
1514 this may be vital information to other prospective consumers if they were willing to  
1515 accept the certification in lieu of having to perform another certification themselves.  
1516 Note, while the examples of Associated Documentation presented here are related to  
1517 services, the concept applies equally to description of other entities.



1518 **4.1.1.2 Assigning Values to Description Instances**

1519



1520

1521 *Figure 14 Representation of a Description*

1522 Figure 13 shows the template for a general description but individual description  
 1523 instances depend on the ability to associate meaningful values with the identified  
 1524 elements. Figure 14 shows a model for a collection of information that provides for value  
 1525 assignment and traceability for both the value meaning and the source of a value. The  
 1526 model is not meant to replace existing or future schema or other structures that have or  
 1527 will be defined for specific implementations, but it is meant as guidance for the  
 1528 information such structures need to capture to generate sufficient description. It is  
 1529 expected that tools will be developed to assist the user in populating description and  
 1530 auto-filling many of these fields, and in that context, this model provides guidance to the  
 1531 tool developers.

1532 In Figure 14 each class has an associated value specifier or is made up of components  
 1533 that will eventually resolve to a value specifier. For example, **Description** has several  
 1534 components, one of which is **Categorization**, which would have an associated a value  
 1535 specifier.

1536 A value specifier consists of

- 1537 • a collection of value sets with associated property-value pairs, pointers to such value  
 1538 sets, or pointers to descriptions that eventually resolve to value sets that describe  
 1539 the component; and
- 1540 • attributes that qualify the value specifier and the value sets it contains.

1541 The qualifying attributes for the value specifier include

- 1542 • an optional identifier that would allow the value set to be defined, accessed, and  
 1543 reused elsewhere;

1544 • provenance information that identifies the party (individual, role, or organization) that  
1545 has responsibility for assigning the value sets to any description component;

1546 • an optional source of the value set, if appropriate and meaningful, e.g. if a particular  
1547 data source is mandated.

1548 If the value specifier is contained within a higher-level component, (such as Service  
1549 Description containing Service Functionality), the component may inherit values from  
1550 the attributes from its container.

1551 Note, provenance as a qualifying attribute of a value specifier is different from  
1552 provenance as part of an instance of Description. Provenance for a service identifies  
1553 those who own and are responsible for the service, as described in Section 3.  
1554 Provenance for a value specifier identifies who is responsible for choosing and  
1555 assigning values to the value sets that comprise the value specifier. It is assumed that  
1556 granularity at the value specifier level is sufficient and provenance is not required for  
1557 each value set.

1558 The value set also has attributes that define its structure and semantics.

1559 • The semantics of the value set property should be associated with a semantic  
1560 context conveying the meaning of the property within the execution context, where  
1561 the semantic context could vary from a free text definition to a formal ontology.

1562 • For numeric values, the structure would provide the numeric format of the value and  
1563 the “semantics” would be conveyed by a dimensional unit with an identifier to an  
1564 authoritative source defining the dimensional unit and preferred mechanisms for its  
1565 conversion to other dimensional units of like type.

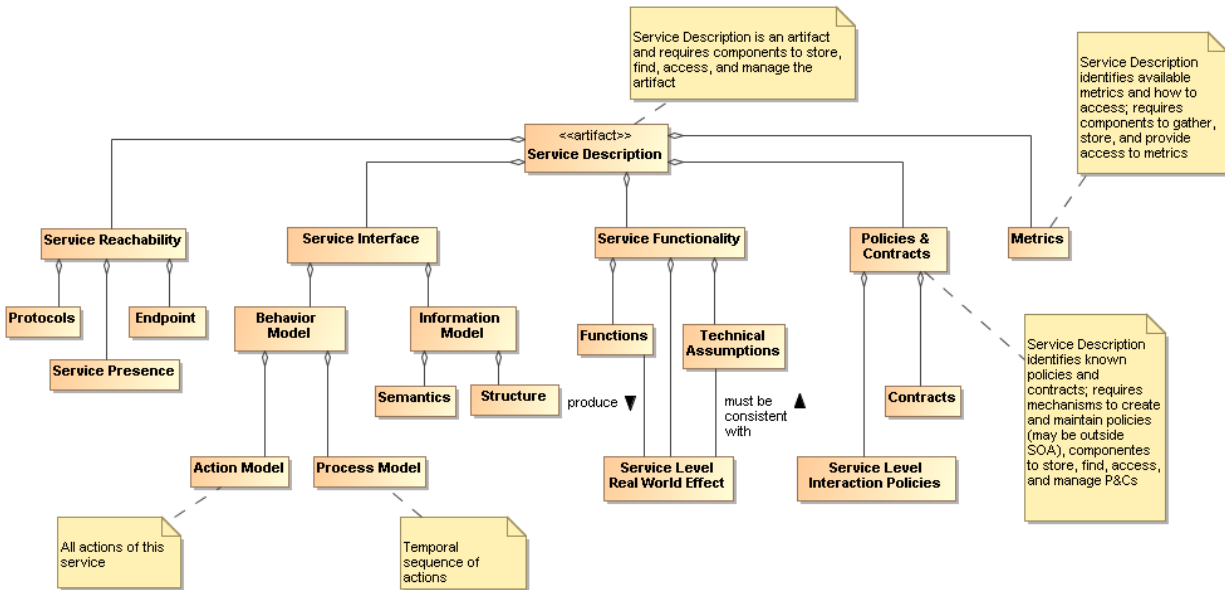
1566 • For nonnumeric values, the structure would provide the data structure for the value  
1567 representation and the semantics would be an associated semantic model.

1568 • For pointers, architectural guidelines would define the preferred addressing scheme.

1569 The value specifier may indicate a default semantic model for its component value sets  
1570 and the individual value sets may provide an override.

1571 The property-value pair construct is introduced for the value set to emphasize the need  
1572 to identify unambiguously both what is being specified and what is a consistent  
1573 associated value. The further qualifying of Structure and Semantics in the Set  
1574 Attributes allows for flexibility in defining the form of the associated values.

1575 **4.1.1.3 Model Elements Specific to Service Description**



1576  
1577 *Figure 15 Service Description*

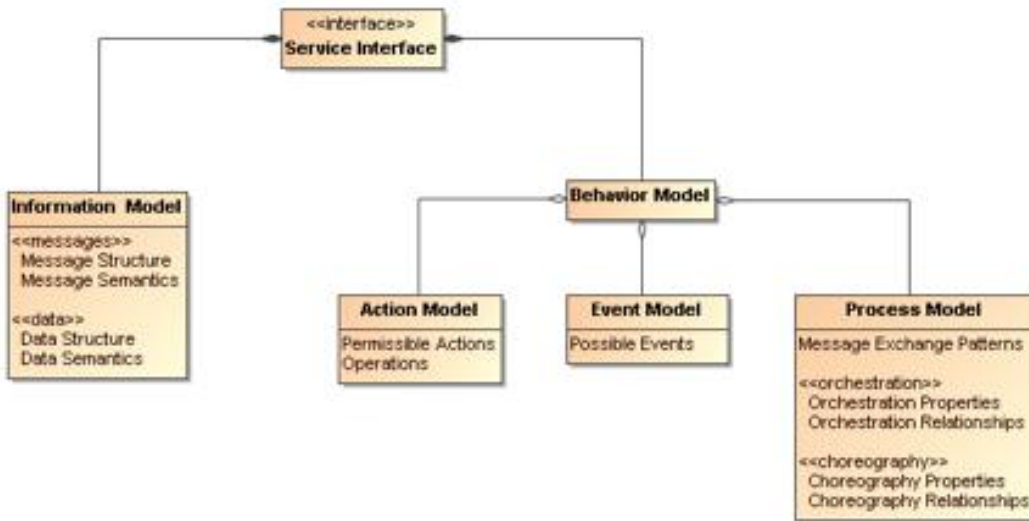
1578 The major elements for the Service Description subclass follow directly from the areas  
1579 discussed in the Reference Model. Here, we discuss the detail shown in *Figure 15* and  
1580 the purpose served by each element of service description.

1581 Note, the intent in the subsections that follow is to describe how a particular element,  
1582 such as the service interface, is reflected in the service description, not to elaborate on  
1583 the details of that element.

1584 **4.1.1.3.1 Service Interface**

1585 As noted in the Reference Model, the service interface is the means for interacting with  
1586 a service. For the SOA-RAF and as shown in Section 4.3 the service interface will  
1587 support an exchange of messages, where

- 1588 • the message conforms to a referenceable message exchange pattern (MEP),
- 1589 • the message payload conforms to the structure and semantics of the indicated  
1590 information model,
- 1591 • the messages are used to denote events or actions against the service, where  
1592 the actions are specified in the action model and any required sequencing of  
1593 actions is specified in the process model.



1594  
1595 *Figure 16 Service Interface*

1596 Note we distinguish the structure and semantics of the message from that of the  
1597 underlying protocol that conveys the message. The message structure may include  
1598 nested structures that are independently defined, such as an enclosing envelope  
1599 structure and an enclosed data structure.

1600 These aspects of messages are discussed in more detail in Section 4.3

#### 1601 **4.1.1.3.2 Service Reachability**

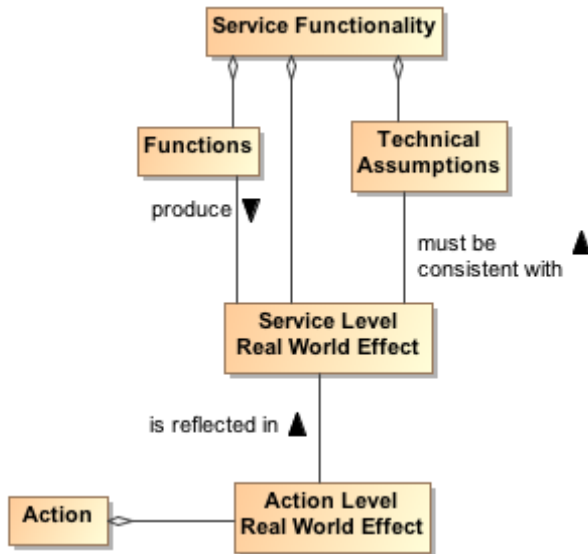
1602 Service reachability, as modeled in Section 4.2.2.3 enables service participants to  
1603 locate and interact with one another. To support service reachability, the service  
1604 description should indicate the endpoints to which a service consumer can direct  
1605 messages to invoke actions and the protocol to be used for message exchange using  
1606 that endpoint.

1607 As applied in general to an action, the endpoint is the conceptual location where one  
1608 applies an action; with respect to service description, it is the actual address where a  
1609 message is sent.

1610 In addition, the service description should provide information on collected metrics for  
1611 service presence; see Section 4.1.1.3.4 for the discussion of metrics as part of service  
1612 description.

#### 1613 **4.1.1.3.3 Service Functionality**

1614 While the service interface and service reachability are concerned with the mechanics  
1615 of using a service, service functionality and performance metrics (discussed in Section  
1616 4.1.1.3.4) describe what can be expected when interacting with a service. Service  
1617 Functionality, shown in *Figure 15* as part of the overall Service Description model and  
1618 extended in *Figure 17*, is an unambiguous expression of service function(s) and the real  
1619 world effects of invoking the function. The Functions represent business activities in  
1620 some domain that produce the desired real world effects.



1621  
1622 *Figure 17 Service Functionality*

1623 The Service Functionality may also be constrained by Technical Assumptions that  
1624 underlie the effects that can result. Technical assumptions are defined as domain  
1625 specific restrictions and may express underlying physical limitations, such as flow  
1626 speeds must be below sonic velocity or disk access that cannot be faster than the  
1627 maximum for its host drive. Technical assumptions are related to the underlying  
1628 capability accessed by the service. In any case, the real world effects must be  
1629 consistent with the Technical Assumptions.

1630 In *Figure 15* and *Figure 17*, we specifically refer to Service Level and Action Level real world  
1631 effects.

1632 **Service Level Real World Effect**

1633 A service level real world effect is a specific change in shared state or  
1634 information returned as a result of interacting with a service.

1635 **Action Level Real World Effect**

1636 An action level real world effect is a specific change in shared state or  
1637 information returned as a result of performing a specific action against a service.

1638 Service description describes the service as a whole while the component aspects  
1639 should contribute to that whole. Thus, while individual Actions may contribute to the  
1640 real world effects to be realized from interaction with the service, there would be a  
1641 serious disconnect for Actions to contribute real world effects that could not consistently  
1642 be reflected in the Service Level Real World Effects and thus the Service Functionality.  
1643 The relationship to Action Level Real World Effects and the implications on defining the  
1644 scope of a service are discussed in Section 4.1.2.1.

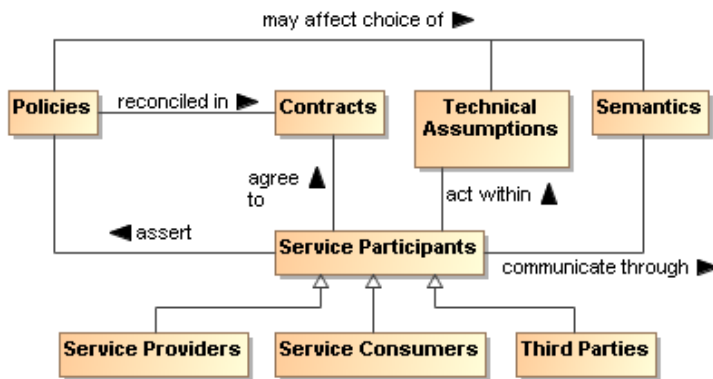
1645 Elements of Service Functionality may be expressed as natural language text, reference  
1646 to an existing taxonomy of functions, or reference to a more formal knowledge capture  
1647 providing richer description and context.

1648 **4.1.1.3.4 Policies and Contracts, Metrics, and Compliance Records**

1649 Policies prescribe the conditions and constraints for interacting with a service and  
1650 impact the willingness to continue visibility with the other participants. Whereas  
1651 technical assumptions are statements of “physical” fact, policies are subjective  
1652 assertions made by the service provider (sometimes as passed on from higher  
1653 authorities).

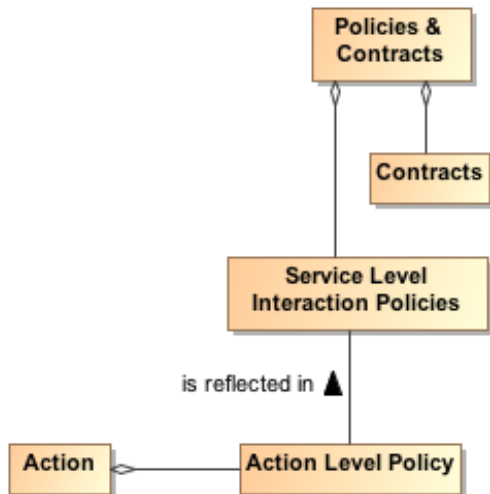
1654 The service description provides a central location for identifying what policies have  
1655 been asserted by the service provider. The specific representation of the policy, e.g. in  
1656 some formal policy language, is likely done outside of the service description and the  
1657 service description would reference the normative definition of the policy.

1658 Policies may also be asserted by other service participants, as illustrated by the model  
1659 shown in Figure 18. Policies that are generally applicable to any interaction with the  
1660 service are asserted by the service provider and included in the Policies and Contracts  
1661 section of the service description. Conversely, policies that are asserted by specific  
1662 consumers or consumer communities would be identified as part of a description’s  
1663 Annotations from 3<sup>rd</sup> parties (see Section 4.1.1.1.4) because these would be specific to  
1664 those parties and not a general aspect of the service being described.



1665  
1666 *Figure 18 Model for Policies and Contracts as related to Service Participants*

1667 In *Figure 15* and *Figure 19*, we specifically refer to Service Level Interaction Policies. In a  
1668 similar manner to that discussed for Service Level vs. Action Level Real World Effects in  
1669 Section 4.1.1.3.3, individual Actions may have associated policies stating conditions for  
1670 performing the action, but these must be reflected in and be consistent with the policies  
1671 made visible at the service level and thus the description of the service as a whole. The  
1672 relationship to Action Level Policies and the implications on defining the scope of a  
1673 service are discussed in Section 4.1.2.1.



1674

1675

*Figure 19 Action-Level and Service-Level Policies*

1676

1677

1678

1679

1680

1681

As noted in Figure 18, the policies asserted may affect the allowable Technical Assumptions that can be embodied in services or their underlying capabilities and may affect the semantics that can be used. For example of the former, there may be a policy that specifies the surge capacity to be accommodated by a server, and a service that designs for a smaller capacity would not be appropriate to use. For the latter, a policy may require that only services using a community-sponsored vocabulary can be used.

1682

1683

1684

1685

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

1686

1687

1688

1689

1690

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

1691

1692

1693

1694

1695

1696

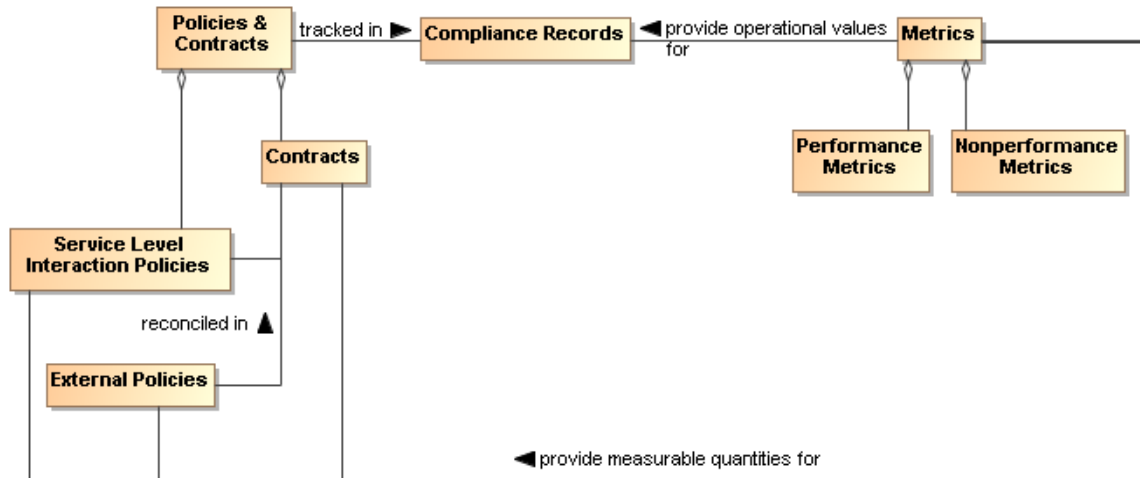
1697

1698

1699

1700

The definition and later enforcement of policies and contracts are predicated on the existence of metrics; the relationships among the relevant concepts are shown in the model in Figure 20. Performance Metrics identify quantities that characterize the speed and quality of realizing the real world effects produced using the SOA service; in addition, policies and contracts may depend on nonperformance metrics, such as whether a license is in place to use the service. Some of these metrics reflect the underlying capability, e.g. a SOA service cannot respond in two seconds if the underlying capability is expected to take five seconds to do its processing; some metrics reflect the implementation of the SOA service, e.g. what level of caching is present to minimize data access requests across the network.



1701  
1702 *Figure 20 Policies and Contracts, Metrics, and Compliance Records*

1703 As with many quantities, the metrics associated with a service are not themselves  
1704 defined by this Service Description because it is not known *a priori* which metrics are  
1705 being collected or otherwise checked by the services, the SOA infrastructure, or other  
1706 resources that participate in the SOA interactions. However, the service description  
1707 SHOULD provide a placeholder (possibly through a link to an externally compiled list)  
1708 for identifying which metrics are available and how these can be accessed.

1709 The use of metrics to evaluate compliance is discussed in Section **Error! Reference**  
1710 **source not found..** The results of compliance evaluation SHOULD be maintained in  
1711 compliance records and the means to access the compliance records SHOULD be  
1712 included in the Policies and Contracts portion of the service description. For example,  
1713 the description may be in the form of static information (e.g. over the first year of  
1714 operation, this service had a 91% availability), a link to a dynamically generated metric  
1715 (e.g. over the past 30 days, the service has had a 93.3% availability), or access to a  
1716 dynamic means to check the service for current availability (e.g. a ping). The  
1717 relationship between service presence and the presence of the individual actions that  
1718 can be invoked is discussed under Reachability in Section 4.2.2.3.

1719 Note, even when policies relate the perspective of a single participant, policy  
1720 compliance can be measured and policies may be enforceable without contractual  
1721 agreement with other participants. This should be reflected in the policy, contract, and  
1722 compliance record information maintained in the service description.

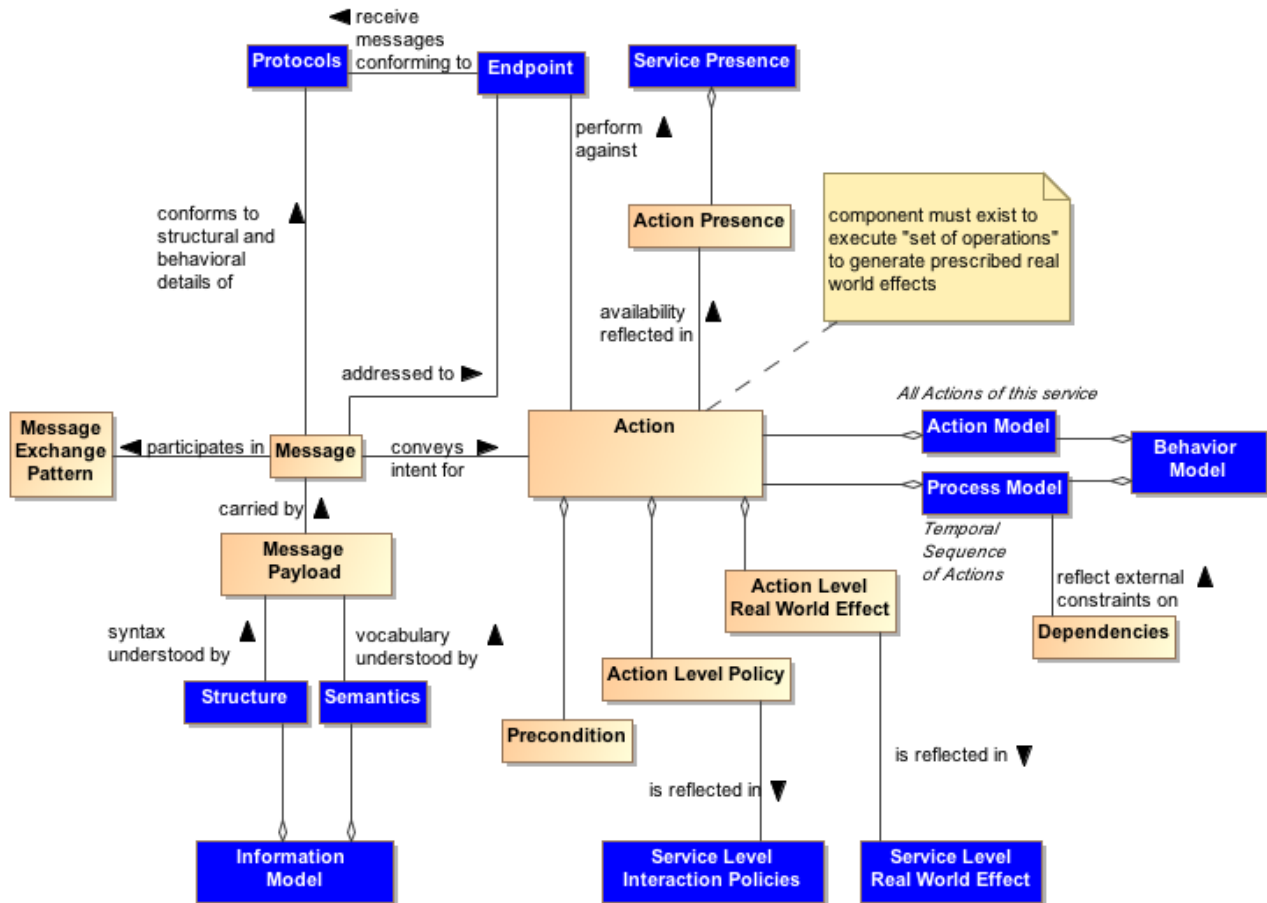
## 1723 4.1.2 Use Of Service Description

### 1724 4.1.2.1 Service Description in support of Service Interaction

1725 If we assume we have awareness, i.e. access to relevant descriptions, the service  
1726 participants must still establish willingness and presence to ensure full visibility (See  
1727 Section 4.2) and to interact with the service. Service description provides necessary  
1728 information for many aspects of preparing for and carrying through with interaction.  
1729 Recall the fundamental definition of service is a mechanism to access an underlying  
1730 capability; the service description describes this mechanism and its use. It lays the



1731 groundwork for what can occur, whereas service interaction defines the specifics  
 1732 through which occurrences are realized.



1733  
 1734 Figure 21 Relationship Between Action and Service Description Components

1735 Figure 21 combines the models in the subsections of Section 4.1.1 to concisely relate  
 1736 action and the relevant components of Service Description. The purpose of Figure 21 is  
 1737 to demonstrate that the components of service description go beyond arbitrary  
 1738 documentation and form the critical set of information needed to define the what and  
 1739 how of action. In Figure 21, the leaf nodes from Figure 15 are shown in blue.

1740 action is invoked via a Message where the structure and behavioral details of the  
 1741 message conform to an identified Protocol and is directed to the address of the  
 1742 identified endpoint, and the message payload conforms to the service Information  
 1743 Model.

1744 The availability of an action is reflected in the Action Presence and each Action  
 1745 Presence contributes to the overall Service Presence; this is discussed further in  
 1746 Section 4.2.2.3. Each action has its own endpoint and also its own protocols associated  
 1747 with the endpoint<sup>11</sup> and to what extent, e.g. current or average availability, there is

<sup>11</sup> 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.

1748 presence for the action through that endpoint. The endpoint and service presence are  
1749 also part of the service description.

1750 An action may have preconditions where a Precondition is something that needs to be  
1751 in place before an action can occur, e.g. confirmation of a precursor action. Whether  
1752 preconditions are satisfied is evaluated when someone tries to perform the action and  
1753 not before. Presence for an action means someone can initiate it and is independent of  
1754 whether the preconditions are satisfied. However, the successful completion of the  
1755 action may depend on whether its preconditions were satisfied.

1756 Analogous to the relationship between actions and preconditions, the Process Model  
1757 may imply Dependencies for succeeding steps in a process, e.g. that a previous step  
1758 has successfully completed, or may be isolated to a given step. An example of the  
1759 latter would be a dependency that the host server has scheduled maintenance and  
1760 access attempts at these times would fail. Dependencies related to the process model  
1761 do not affect the presence of a service although these may affect whether the business  
1762 function successfully completes.

1763 The conditions under which an action can be invoked may depend on policies  
1764 associated with the action. The Action Level Policies MUST be reflected in the Service  
1765 Level Interaction Policies because such policies may be critical to determining whether  
1766 the conditions for use of the service are consistent with the policies asserted by the  
1767 service consumer. The service level interaction policies are included in the service  
1768 description.

1769 Similarly, the result of invoking an action is one or more real world effects, and the  
1770 Action Level Real World Effects MUST be reflected in the Service Level Real World  
1771 Effect included in the service description. The unambiguous expression of action level  
1772 policies and real world effects as service counterparts is necessary to adequately  
1773 understand what constitutes the service interaction.

1774 An adequate service description MUST provide a consumer with information needed to  
1775 determine if the service policies and the (business) functions and service-level real  
1776 world effects are of interest and there is nothing in the technical assumptions that  
1777 preclude use of the service.

1778 Note at this level, the business functions are not concerned with the action or process  
1779 models. These models are detailed separately.

1780 The service description is not intended to be isolated documentation but rather an  
1781 integral part of service use. Changes in service description SHOULD immediately be  
1782 made known to consumers and potential consumers.

#### 1783 **4.1.2.1.1 Description and Invoking Actions Against a Service**

1784 At this point, let us assume the descriptions were sufficient to establish willingness; see  
1785 Section 4.2.2.2. Figure 21 indicates the service endpoint establishes where to actually  
1786 carry out the interaction. This is where we start considering the action and process  
1787 models.

1788 The action model identifies the multiple actions a user can perform against a service  
1789 and the user would perform these in the context of the process model as specified or  
1790 referenced under the Service Interface portion of Service Description. For a given  
1791 business function, there is a corresponding process model, where any process model

1792 may involve multiple actions. From the above discussion of model elements of  
1793 description we may conclude (1) actions have reachability information, including  
1794 endpoint and presence, (2) presence of service is some aggregation of presence of its  
1795 actions, (3) action preconditions and service dependencies do not affect presence  
1796 although these may affect successful completion.

1797 Having established visibility, the interaction can proceed. Given a business function, the  
1798 consumer knows what will be accomplished (the service functionality), the conditions  
1799 under which interaction will proceed (service policies and contracts), and the process  
1800 that must be followed (the process model). The remaining question is how does the  
1801 description information for structure and semantics enable interaction.

1802 We have established the importance of the process model in identifying relevant actions  
1803 and their sequence. Interaction proceeds through messages and thus it is the syntax  
1804 and semantics of the messages with which we are here concerned. A common  
1805 approach is to define the structure and semantics that can appear as part of a message;  
1806 then assemble the pieces into messages; and, associate messages with actions.  
1807 Actions make use of structure and semantics as defined in the information model to  
1808 describe its legal messages.

1809 The process model identifies actions to be performed against a service and the  
1810 sequence for performing the actions. For a given action, the Reachability portion of  
1811 description indicates the protocol bindings that are available, the endpoint  
1812 corresponding to a binding, and whether there is presence at that endpoint. The  
1813 interaction with actions is through messages that conform to the structure and  
1814 semantics defined in the information model and the message sequence conforming to  
1815 the action's identified MEP. The result is some portion of the real world effect that must  
1816 be assessed and/or processed (e.g. if an error exists, that part that covers the error  
1817 processing would be invoked).

#### 1818 **4.1.2.1.2 The Question of Multiple Business Functions**

1819 Action level effects and policies **MUST** be reflected at the service level for service  
1820 description to support visibility.

1821 It is assumed that a SOA service represents an identifiable business function to which  
1822 policies can be applied and from which desired business effects can be obtained. While  
1823 contemporary discussions of SOA services and supporting standards do not constrain  
1824 what actions or combinations of actions can or should be defined for a service, the  
1825 SOA-RAF considers the implications of service description in defining the range of  
1826 actions appropriate for an individual SOA service.

1827 Consider the situation if a given SOA service is the container for multiple independent  
1828 (but loosely related) business functions. These are not multiple effects from a single  
1829 function but multiple functions with potentially different sets of effects for each function.  
1830 A service can have multiple actions a user may perform against it, and this does not  
1831 change with multiple business functions. As an individual business function corresponds  
1832 to a process model, so multiple business functions imply multiple process models. The  
1833 same action may be used in multiple process models but the aggregated service  
1834 presence would be specific to each business function because the components being  
1835 aggregated may be different between process models. In summary, for a service with  
1836 multiple business functions, each function has (1) its own process model and

1837 dependencies, (2) its own aggregated presence, and (3) possibly its own list of policies  
1838 and real world effects.

1839 A common variation on this theme is for a single service to have multiple endpoints for  
1840 different levels of quality of service (QoS). Different QoS imply separate statements of  
1841 policy, separate endpoints, possibly separate dependencies, and so on. One could say  
1842 the QoS variation does not require this because there can be a single QoS policy that  
1843 encompasses the variations. and all other aspects of the service would be the same  
1844 except for the endpoint used for each QoS. However, the different aspects of policy at  
1845 the service level would need to be mapped to endpoints, and this introduces an  
1846 undesirable level of coupling across the elements of description. In addition, it is  
1847 obvious that description at the service level can become very complicated if the number  
1848 of combinations is allowed to grow.

1849 One could imagine a service description that is basically a container for action  
1850 descriptions, where each action description is self contained; however, this would lead  
1851 to duplication of description components across actions. If common description  
1852 components are factored, this either is limited to components common across all  
1853 actions or requires complicated tagging to capture the components that often but do not  
1854 universally apply.

1855 If a provider cannot describe a service as a whole but must describe every action, this  
1856 leads to the situation where it may be extremely difficult to construct a clear and concise  
1857 service description that can effectively support discovery and use without tedious logic  
1858 to process the description and assemble the available permutations. In effect, if  
1859 adequate description of an action begins to look like description of a service, it may be  
1860 best to have it as a separate service.

1861 Recall, more than one service can access the same underlying capability, and this is  
1862 appropriate if a different real world effect is to be exposed. Along these lines, one can  
1863 argue that different QoS are different services because getting a response in one  
1864 minute rather than one hour is more than a QoS difference; it is a fundamental  
1865 difference in the business function being provided.

1866 As a best practice, a criteria for whether a service is appropriately scoped may be the  
1867 ease or difficulty in creating an unambiguous service description. A consequence of  
1868 having tightly-scoped services is there will be a greater reliance on combining services,  
1869 i.e. more fundamental business functions, to create more advanced business functions.  
1870 This is consistent with the principles of service oriented architecture and is the basic  
1871 position of the Reference Architecture, although not an absolute requirement.  
1872 Combining services increases the reliance on understanding and implementing the  
1873 concepts of orchestration, choreography, and other approaches yet to be developed;  
1874 these are discussed in more detail in section 4.4 Interacting with Services.

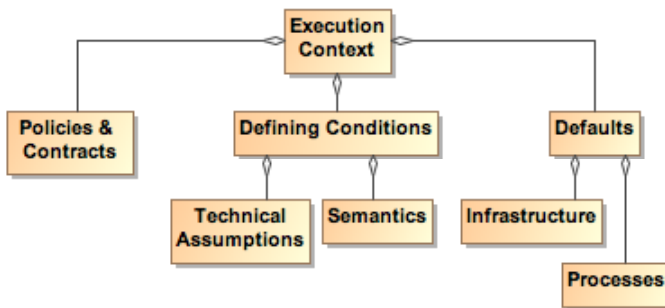
#### 1875 **4.1.2.1.3 Service Description, Execution Context, and Service Interaction**

1876 The service description MUST provide sufficient information to support service visibility,  
1877 including the willingness of service participants to interact. However, the corresponding  
1878 descriptions for providers and consumers may both contain policies, technical  
1879 assumptions, constraints on semantics, and other technical and procedural conditions  
1880 that must be aligned to define the terms of willingness. The agreements which  
1881 encapsulate the necessary alignment form the basis upon which interactions may

1882 proceed – in the Reference Model, this collection of agreements and the necessary  
1883 environmental support establish the execution context.

1884 To illustrate the concept of the execution context, consider a Web-based system for  
1885 timecard entry. For an employee onsite at an employer facility, the execution context  
1886 requires a computer connected to the local network and the employee must enter their  
1887 network ID and password. Relevant policies include that the employee must maintain  
1888 the most recent anti-virus software and virus definitions for any computer connected to  
1889 the network.

1890 For the same employee connecting from offsite, the execution context specifies the  
1891 need for a computer with installed VPN software and a security token to negotiate the  
1892 VPN connection. The execution context also includes proxy settings as needed to  
1893 connect to the offsite network. The employee must still comply with the requirements for  
1894 onsite computers and access, but the offsite execution context includes additional items  
1895 before the employee can access the same underlying capability and realize the same  
1896 real world effects, i.e. the timecard entries.

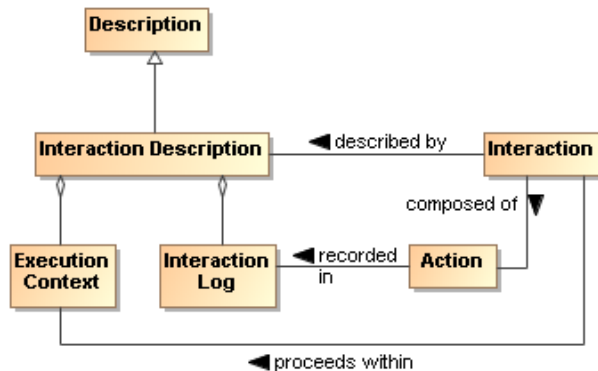


1897  
1898 *Figure 22 Execution Context*

1899 Figure 22 shows a few broad categories found in execution context. These are not  
1900 meant to be comprehensive. Other items may need to be included to collect a sufficient  
1901 description of the interaction conditions. Any other items not explicitly noted in the  
1902 model but needed to set the environment SHOULD be included in the execution  
1903 context.

1904 While the execution context captures the conditions under which interaction can occur,  
1905 it does not capture the specific service invocations that do occur in a specific interaction.  
1906 A service interaction as modeled in Figure 21 introduces the concept of an Interaction  
1907 Description which is composed of both the Execution Context and an Interaction Log.  
1908 The execution context specifies the set of conditions under which the interaction occurs  
1909 and the interaction log captures the sequence of service interactions that occur within  
1910 the execution context. This sequence should follow the Process Model but can include  
1911 details beyond those specified there. For example, the Process Model may specify an  
1912 action that results in identifying a data source, and the identified source is used in a  
1913 subsequent action. The Interaction Log would record the specific data source used.

1914 The execution context can be thought of as the container in which the interaction occurs  
1915 and the interaction log captures what happens inside the container. This combination is  
1916 needed to support auditability and repeatability of the interactions.



1917  
1918 *Figure 23 Interaction Description*

1919 SOA allows flexibility to accomplish repeatability or reusability. One benefit of this is that  
1920 a service can be updated without disrupting the user experience of the service. So,  
1921 Google can improve their ranking algorithm without notifying the user about the details  
1922 of the update.

1923 However, it may also be vital for the consumer to be able to recreate past results or to  
1924 generate consistent results in the future, and information such as what conditions, which  
1925 services, and which versions of those services are used is indispensable in retracing  
1926 one’s path. The interaction log is a critical part of the resulting real world effects  
1927 because it defines how the effects were generated and possibly the meaning of  
1928 observed effects. This increases in importance as dynamic composability becomes  
1929 more feasible. In essence, a result has limited value if one does not know how it was  
1930 generated.

1931 The interaction log SHOULD be a detailed trace for a specific interaction, and its reuse  
1932 is limited to duplicating that interaction. An execution context can act as a template for  
1933 identical or similar interactions. Any given execution context MAY define the conditions  
1934 of future interactions.

1935 Such uses of execution context imply (1) a standardized format for capturing execution  
1936 context and (2) a subclass of general description could be defined to support visibility of  
1937 saved execution contexts. The specifics of the relevant formats and descriptions are  
1938 beyond the scope of this document.

1939 A service description is unlikely to track interaction descriptions or the constituent  
1940 execution contexts or interaction logs that include mention of the service. However, as  
1941 appropriate, linking to specific instances of either of these could be done through  
1942 associated annotations.

1943 **4.1.3 Relationship to Other Description Models**

1944 While the representation shown in Figure 14 is derived from considerations related to  
1945 service description, it is acknowledged that other metadata standards are relevant and  
1946 should, as possible, be incorporated into this work. Two standards of particular  
1947 relevance are the Dublin Core Metadata Initiative (DCMI) and ISO 11179, especially  
1948 Part 5.

1949 When the service description (or even the general description class) is considered as  
1950 the DCMI “resource”, Figure 14 aligns nicely with the DCMI resource model. While

1951 some differences exist, these are mostly in areas where DCMI goes into detail that is  
1952 considered beyond the scope of the current Reference Architecture. For example,  
1953 DCMI defines classes of “shared semantics” whereas this Reference Architecture  
1954 Framework considers that an identification of relevant semantic models is sufficient.  
1955 Likewise, the DCMI “description model” goes into the details of possible syntax  
1956 encodings whereas for the Reference Architecture Framework it is sufficient to identify  
1957 the relevant formats.

1958 With respect to ISO 11179 Part 5, the metadata fields defined in that reference may be  
1959 used without prejudice as the properties in Figure 14. Additionally, other defined  
1960 metadata sets may be used by the service provider if the other sets are considered  
1961 more appropriate, i.e. it is fundamental to this reference architecture to identify the need  
1962 and the means to make vocabulary declarations explicit but it is beyond the scope to  
1963 specify which vocabularies are to be used. In addition, the identification of domain of  
1964 the properties and range of the values has not been included in the current Reference  
1965 Architecture discussion, but the text of ISO 11179 Part 5 can be used consistently with  
1966 the model prescribed in this document.

1967 Description as defined here considers a wide range of applicability and support of the  
1968 principles of service oriented architecture. Other metadata models can be used in  
1969 concert with the model presented here because most of these focus on a finer level of  
1970 detail that is outside the present scope, and so provide a level of implementation  
1971 guidance that can be applied as appropriate.

#### 1972 **4.1.4 Architectural Implications**

1973 The description of service description indicates numerous architectural implications on  
1974 the SOA ecosystem:

- 1975 • It changes over time and its contents will reflect changing needs and context. This  
1976 requires the existence of:
  - 1977 ○ mechanisms to support the storage, referencing, and access to normative  
1978 definitions of one or more versioning schemes that may be applied to identify  
1979 different aggregations of descriptive information, where the different schemes  
1980 may be versions of a versioning scheme itself;
  - 1981 ○ configuration management mechanisms to capture the contents of the each  
1982 aggregation and apply a unique identifier in a manner consistent with an  
1983 identified versioning scheme;
  - 1984 ○ one or more mechanisms to support the storage, referencing, and access to  
1985 conversion relationships between versioning schemes, and the mechanisms  
1986 to carry out such conversions.
- 1987 • Description makes use of defined semantics, where the semantics may be used for  
1988 categorization or providing other property and value information for description  
1989 classes. This requires the existence of:
  - 1990 ○ semantic models that provide normative descriptions of the utilized terms,  
1991 where the models may range from a simple dictionary of terms to an ontology  
1992 showing complex relationships and capable of supporting enhanced  
1993 reasoning;
  - 1994 ○ mechanisms to support the storage, referencing, and access to these  
1995 semantic models;

- 1996 ○ configuration management mechanisms to capture the normative description
- 1997 of each semantic model and to apply a unique identifier in a manner
- 1998 consistent with an identified versioning scheme;
- 1999 ○ one or more mechanisms to support the storage, referencing, and access to
- 2000 conversion relationships between semantic models, and the mechanisms to
- 2001 carry out such conversions.
- 2002 ● Descriptions include reference to policies defining conditions of use and optionally
- 2003 contracts representing agreement on policies and other conditions. This requires the
- 2004 existence of (as also enumerated under governance):
- 2005 ○ descriptions to enable the policy modules to be visible, where the description
- 2006 includes a unique identifier for the policy and a sufficient, and preferably a
- 2007 machine processible, representation of the meaning of terms used to describe
- 2008 the policy, its functions, and its effects;
- 2009 ○ one or more discovery mechanisms that enable searching for policies that
- 2010 best meet the search criteria specified by the service participant; where the
- 2011 discovery mechanism has access to the individual policy descriptions,
- 2012 possibly through some repository mechanism;
- 2013 ○ accessible storage of policies and policy descriptions, so service participants
- 2014 can access, examine, and use the policies as defined.
- 2015 ● Descriptions include references to metrics which describe the operational
- 2016 characteristics of the subjects being described. This requires the existence of (as
- 2017 partially enumerated under governance):
- 2018 ○ the infrastructure monitoring and reporting information on SOA resources;
- 2019 ○ possible interface requirements to make accessible metrics information
- 2020 generated or most easily accessed by the service itself;
- 2021 ○ mechanisms to catalog and enable discovery of which metrics are available
- 2022 for a described resources and information on how these metrics can be
- 2023 accessed;
- 2024 ○ mechanisms to catalog and enable discovery of compliance records
- 2025 associated with policies and contracts that are based on these metrics.
- 2026 ● Descriptions of the interactions are important for enabling auditability and
- 2027 repeatability, thereby establishing a context for results and support for understanding
- 2028 observed change in performance or results. This requires the existence of:
- 2029 ○ one or more mechanisms to capture, describe, store, discover, and retrieve
- 2030 interaction logs, execution contexts, and the combined interaction
- 2031 descriptions;
- 2032 ○ one or more mechanisms for attaching to any results the means to identify
- 2033 and retrieve the interaction description under which the results were
- 2034 generated.
- 2035 ● Descriptions may capture very focused information subsets or can be an aggregate
- 2036 of numerous component descriptions. Service description is an example of an
- 2037 aggregate for which manual maintenance of the whole would not be feasible. This
- 2038 requires the existence of:
- 2039 ○ tools to facilitate identifying description elements that are to be aggregated to
- 2040 assemble the composite description;



- 2041 ○ tools to facilitate identifying the sources of information to associate with the  
2042 description elements;
- 2043 ○ tools to collect the identified description elements and their associated  
2044 sources into a standard, referenceable format that can support general  
2045 access and understanding;
- 2046 ○ tools to automatically update the composite description as the component  
2047 sources change, and to consistently apply versioning schemes to identify the  
2048 new description contents and the type and significance of change that  
2049 occurred.
- 2050 ● Descriptions provide up-to-date information on what a resource is, the conditions for  
2051 interacting with the resource, and the results of such interactions. As such, the  
2052 description is the source of vital information in establishing willingness to interact  
2053 with a resource, reachability to make interaction possible, and compliance with  
2054 relevant conditions of use. This requires the existence of:
  - 2055 ○ one or more discovery mechanisms that enable searching for described  
2056 resources that best meet the criteria specified by a service participant, where  
2057 the discovery mechanism has access to individual descriptions, possibly  
2058 through some repository mechanism;
  - 2059 ○ tools to appropriately track users of the descriptions and notify them when a  
2060 new version of the description is available.

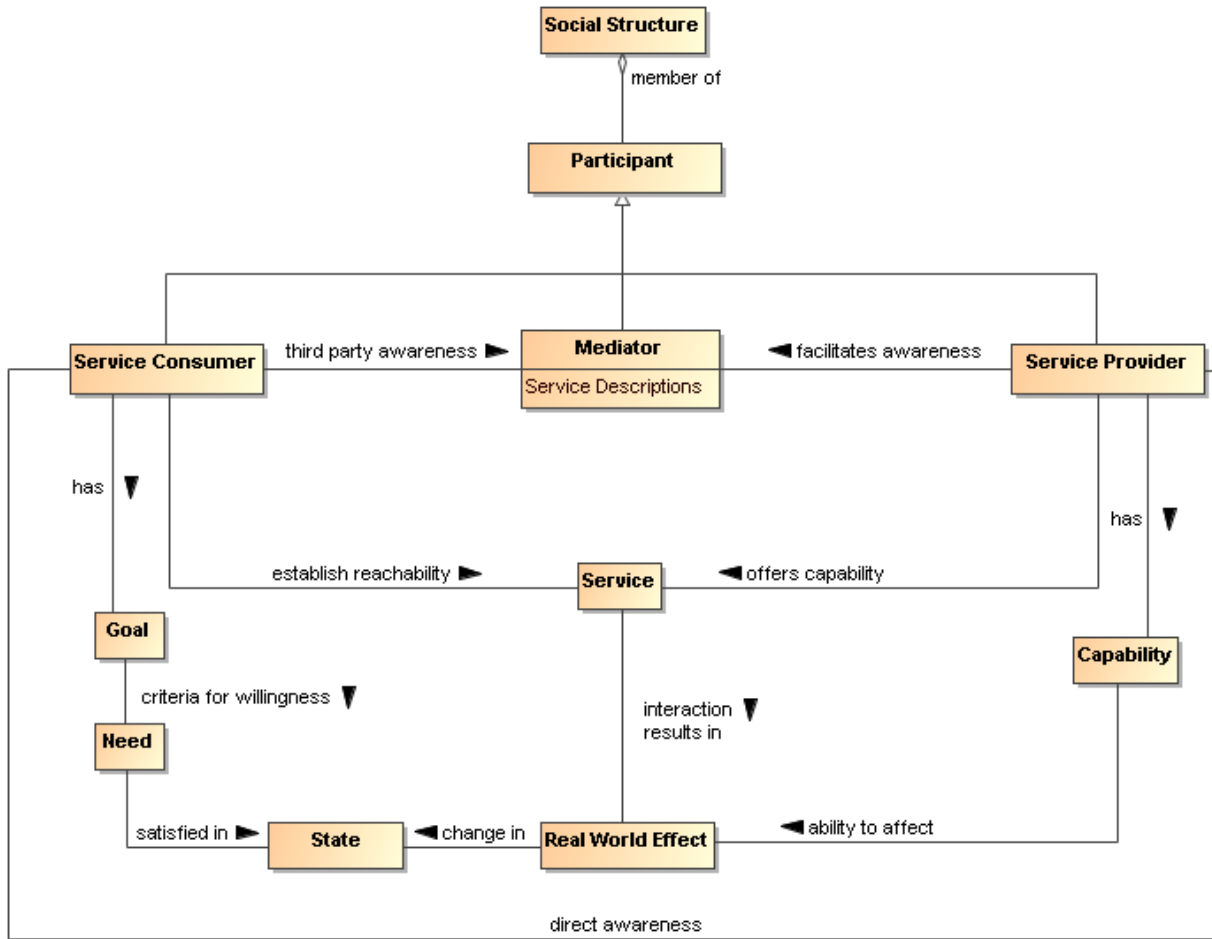
## 2061 4.2 Service Visibility Model

2062 One of the key requirements for participants interacting with each other in the context of  
2063 a SOA is achieving visibility: before services can interoperate, the participants have to  
2064 be visible to each other using whatever means are appropriate. The Reference Model  
2065 analyzes visibility in terms of awareness, willingness, and reachability. In this section,  
2066 we explore how visibility may be achieved.

### 2067 4.2.1 Visibility to Business

2068 The relationship of visibility to the SOA ecosystem encompasses both human social  
2069 structures and automated IT mechanisms. Figure 24 depicts a business setting that is a  
2070 basis for visibility as related to the social structure Model in the *Participation in a SOA*  
2071 *Ecosystem* view (see Section **Error! Reference source not found.**). Service  
2072 consumers and service providers may have direct awareness or mediated awareness  
2073 where mediated awareness is achieved through some third party. A consumer's  
2074 willingness to use a service is reflected by the consumer's presumption of satisfying  
2075 goals and needs based on the description of the service. Service providers offer  
2076 capabilities that have real world effects that result in a change in state of the consumer.  
2077 Reachability of the service by the consumer leads to interactions that change the state  
2078 of the consumer. The consumer can measure the change of state to determine if the  
2079 claims made by description and the real world effects of consuming the service meet  
2080 the consumer's needs.

2081



2082  
2083 *Figure 24 Visibility to Business*

2084 Visibility and interoperability in a SOA ecosystem requires more than location and  
2085 interface information. A meta-model for this broader view of visibility is depicted in  
2086 Section 4.1. In addition to providing improved awareness of service capabilities through  
2087 description of information such as reachability, behavior models, information models,  
2088 functionality, and metrics, the service description may contain policies valuable for  
2089 determination of willingness to interact.

2090 A mediator of service descriptions may provide event notifications to both consumers  
2091 and providers about information relating to service descriptions. One example of this  
2092 capability is a publish/subscribe model where the mediator allows consumers to  
2093 subscribe to service description version changes made by the provider. Likewise, the  
2094 mediator may provide notifications to the provider of consumers that have subscribed to  
2095 service description updates.

2096 Another important business capability in a SOA environment is the ability to narrow  
2097 visibility to trusted members within a social structure. Mediators for awareness may  
2098 provide policy based access to service descriptions allowing for the dynamic formation  
2099 of awareness between trusted members.

## 2100 **4.2.2 Visibility**

2101 Attaining visibility is described in terms of steps that lead to visibility. While there can be  
2102 many contexts for visibility within a single social structure, the same general steps can  
2103 be applied to each of the contexts to accomplish visibility.

2104 Attaining SOA visibility requires

- 2105 • service description creation and maintenance,
- 2106 • processes and mechanisms for achieving awareness of and accessing descriptions,
- 2107 • processes and mechanisms for establishing willingness of participants,
- 2108 • processes and mechanisms to determine reachability.

2109 Visibility may occur in stages, i.e. a participant can become aware enough to look or ask  
2110 for further description, and with this description, the participant can decide on  
2111 willingness, possibly requiring additional description. For example, if a potential  
2112 consumer has a need for a tree cutting (business) service, the consumer can use a web  
2113 search engine to find web sites of providers. The web search engine (a mediator) gives  
2114 the consumer links to relevant web pages and the consumer can access those  
2115 descriptions. For those prospective providers that satisfy the consumer's criteria, the  
2116 consumer's willingness to interact increases. The consumer may contact several tree  
2117 services to get detailed cost information (or arrange for an estimate) and may ask for  
2118 references (further description). The consumer is likely to establish full visibility and  
2119 proceed with interaction with the tree service who mutually establishes visibility.

### 2120 **4.2.2.1 Awareness**

2121 A service participant is aware of another participant if it has access to a description of  
2122 that participant with sufficient completeness to establish the other requirements of  
2123 visibility.

2124 Awareness is inherently a function of a participant; awareness can be established  
2125 without any action on the part of the target participant other than the target providing  
2126 appropriate descriptions. Awareness is often discussed in terms of consumer  
2127 awareness of providers but the concepts are equally valid for provider awareness of  
2128 consumers.

2129 Awareness can be decomposed into the creation of descriptions, making them  
2130 available, and discovering the descriptions. Discovery can be initiated or it can be by  
2131 notification. Initiated discovery for business may require formalization of the required  
2132 capabilities and resources to achieve business goals.

2133 Achieving awareness in a SOA can range from word of mouth to formal service  
2134 descriptions in a standards-based registry-repository. Some other examples of  
2135 achieving awareness in a SOA are the use of a web page containing description  
2136 information, email notifications of descriptions, and document based descriptions.

2137 A mediator as discussed for awareness is a third party participant that provides  
2138 awareness to one or more consumers of one or more services. Direct awareness is  
2139 awareness between a consumer and provider without the use of a third party.

2140 Direct awareness may be the result of having previously established an execution  
2141 context, or direct awareness may include determining the presence of services and then

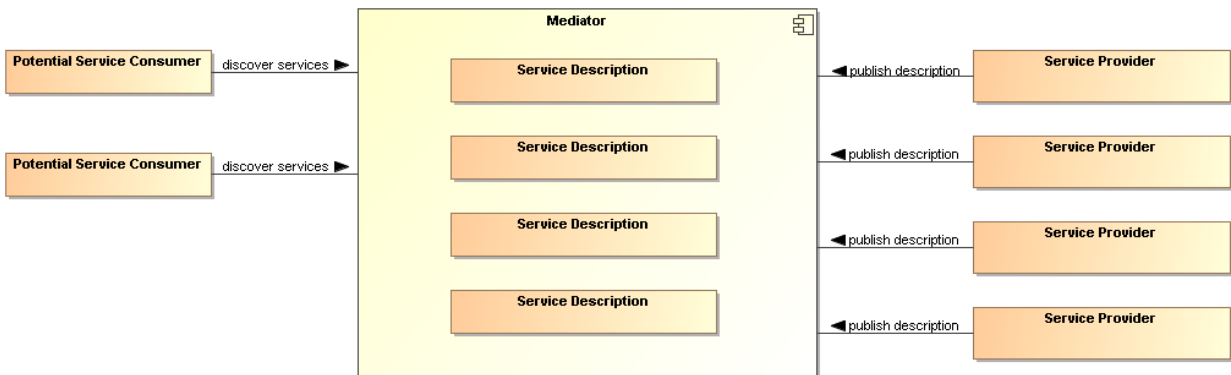
2142 querying the service directly for description. As an example, a priori visibility of some  
2143 sensor device may provide the means for interaction or a query for standardized sensor  
2144 device metadata may be broadcast to multiple locations. If acknowledged, the service  
2145 interface for the device may directly provide description to a consumer so the consumer  
2146 can determine willingness to interact.

2147 The same medium for awareness may be direct in one context and may be mediated in  
2148 another context. For example, a service provider may maintain a web site with links to  
2149 the provider's descriptions of services giving the consumers direct awareness to the  
2150 provider's services. Alternatively, a community may maintain a mediated web site with  
2151 links to various provider descriptions of services for any number of consumers. More  
2152 than one mediator may be involved, as different mediators may specialize in different  
2153 mediation functions.

2154 Descriptions may be formal or informal. Section 4.1, provides a comprehensive model  
2155 for service description that can be applied to formal registry/repositories used to  
2156 mediate visibility. Using consistent description taxonomies and standards based  
2157 mediated awareness helps provide more effective awareness.

#### 2158 4.2.2.1.1 Mediated Awareness

2159 Mediated awareness promotes loose coupling by keeping the consumers and services  
2160 from explicitly referring to each other and the descriptions. Mediation lets interaction  
2161 vary independently. Rather than all potential service consumers being informed on a  
2162 continual basis about all services, there is a known or agreed upon facility or location  
2163 that houses the service description.



2164  
2165 *Figure 25 Mediated Service Awareness*

2166 In Figure 25, the potential service consumers perform queries or are notified in order to  
2167 locate those services that satisfy their needs. As an example, the telephone book is a  
2168 mediated registry where individuals perform manual searches to locate services (i.e. the  
2169 yellow pages). The telephone book is also a mediated registry for solicitors to find and  
2170 notify potential customers (i.e. the white pages).

2171 In mediated service awareness for large and dynamic numbers of service consumers  
2172 and service providers, the benefits typically far outweigh the management issues  
2173 associated with it. Some of the benefits of mediated service awareness are

- 2174 • Potential service consumers have a known location for searching thereby eliminating  
2175 needless and random searches

2176 • Typically a consortium of interested parties (or a sufficiently large corporation) signs  
2177 up to host the mediation facility

2178 • Standardized tools and methods can be developed and promulgated to promote  
2179 interoperability and ease of use.

2180 However, mediated awareness can have some risks associated with it:

2181 • A single point of failure. If the central mediation service fails then a large number of  
2182 service providers and consumers are potentially adversely affected.

2183 • A single point of control. If the central mediation service is owned by, or controlled  
2184 by, someone other than the service consumers and/or providers then the latter may  
2185 be put at a competitive disadvantage based on policies of the discovery provider.

2186 A common mechanism for mediated awareness is a registry-repository. The registry  
2187 stores links or pointers to service description artifacts. The repository in this example is  
2188 the storage location for the service description artifacts. Service descriptions can be  
2189 pushed (publish/subscribe for example) or pulled from the register-repository mediator.

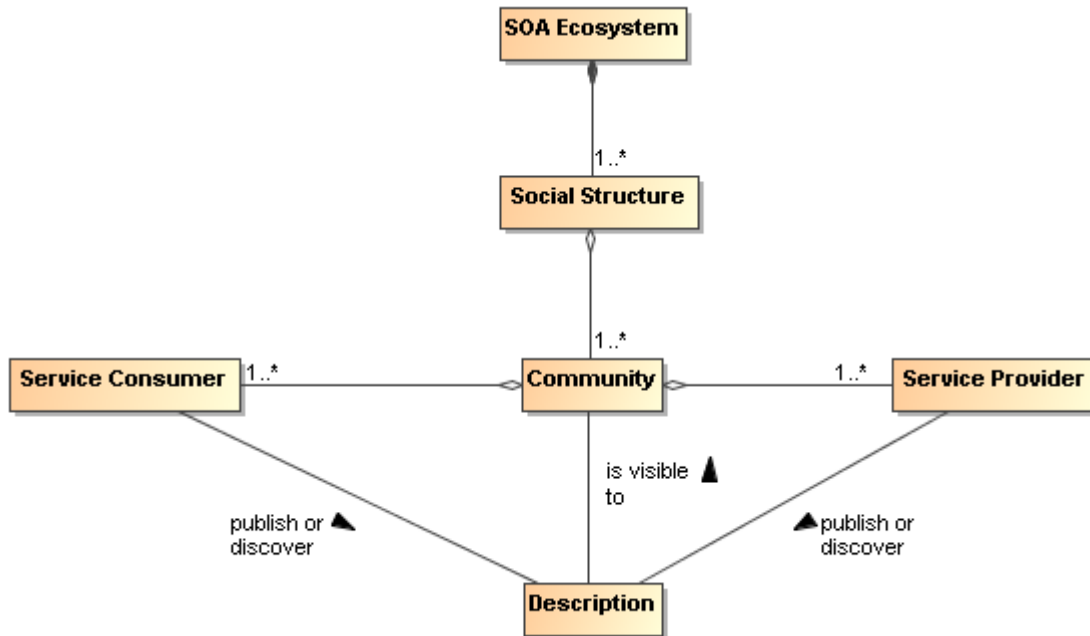
2190 The registry is like a card catalog at the library and a repository is like the shelves for  
2191 the books. Standardized metadata describing repository content can be stored as  
2192 registry objects in a registry and any type of content can be stored as repository items in  
2193 a repository. The registry may be constructed such that description items stored within  
2194 the mediation facility repository has intrinsic links in the registry while description items  
2195 stored outside the mediation facility have extrinsic links in the registry.

2196 When independent but like SOA IT mechanisms interoperate with one another, the IT  
2197 mechanisms may be referred to as federated.

#### 2198 **4.2.2.1.2 Awareness in Complex Social Structures**

2199 Awareness applies to one or more communities within one or more social structures  
2200 where a community consists of at least one description provider and one description  
2201 consumer. These communities may be part of the same social structure or be part of  
2202 different ones.

2203 In Figure 26, awareness can be within a single community, multiple communities, or all  
2204 communities in the social structure. The social structure can encourage or restrict  
2205 awareness through its policies, and these policies can affect participant willingness. The  
2206 information about policies should be incorporated in the relevant descriptions. The  
2207 social structure also governs the conditions for establishing contracts, the results of  
2208 which will be reflected in the execution context if interaction is to proceed.



2209  
2210 *Figure 26 Awareness in a SOA Ecosystem*

2211 IT policy/contract mechanisms can be used by visibility mechanisms to provide  
2212 awareness between communities. The IT mechanisms for awareness may incorporate  
2213 trust mechanisms to assure awareness between trusted communities. For example,  
2214 government organizations may want to limit awareness of an organization’s services to  
2215 specific communities of interest.

2216 Another common business model for awareness is maximizing awareness to  
2217 communities within the social structure, the traditional market place business model. A  
2218 centralized mediator often arises as a provider for this global visibility, a gatekeeper of  
2219 visibility so to speak. For example, Google is a centralized mediator for accessing  
2220 information on the web. As another example, television networks have centralized  
2221 entities providing a level of awareness to communities that otherwise could not be  
2222 achieved without going through the television network.

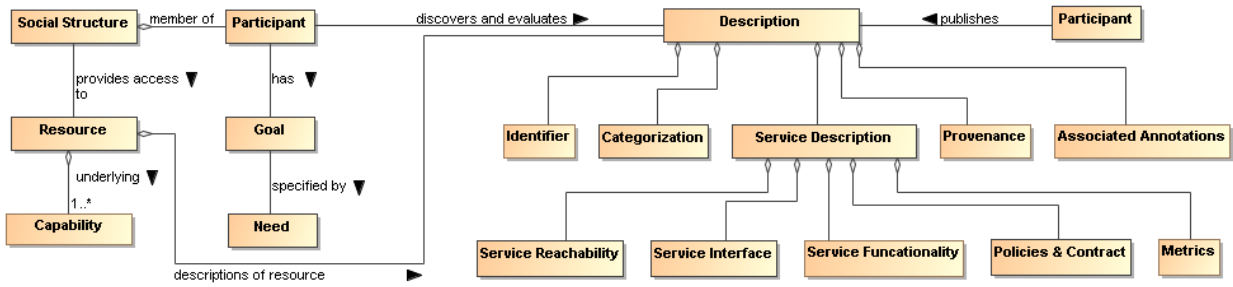
2223 However, mediators have motivations, and they may be selective in which information  
2224 they choose to make available to potential consumers. For example, in a secure  
2225 environment, the mediator may enforce security policies and make information  
2226 selectively available depending on the security clearance of the consumers.

2227 **4.2.2.2 Willingness**

2228 Having achieved awareness, participants use descriptions to help determine their  
2229 willingness to interact with another participant. Both awareness and willingness are  
2230 determined prior to consumer/provider interaction.

2231

2232



2233

2234

Figure 27 Business, Description and Willingness

2235 Figure 27 relates elements of the *Participation in a SOA Ecosystem* view, and elements  
 2236 from the Service Description Model to willingness. By having a willingness to interact  
 2237 within a particular social structure, the social structure provides the participant access to  
 2238 capabilities based on conditions the social structure finds appropriate for its context.  
 2239 The participant can use these capabilities to satisfy goals and objectives as specified by  
 2240 the participant's needs.

2241 In Figure 27, information used to determine willingness is defined by Description.  
 2242 Information referenced by Description may come from many sources. For example, a  
 2243 mediator for descriptions may provide 3rd party annotations for reputation. Another  
 2244 source for reputation may be a participant's own history of interactions with another  
 2245 participant.

2246 A participant inspects functionality for potential satisfaction of needs. Identity is  
 2247 associated with any participant, however, identity may or may not be verified. If  
 2248 available, participant reputation may be a deciding factor for willingness to interact.  
 2249 Policies and contracts referenced by the description may be particularly important to  
 2250 determine the agreements and commitments required for business interactions.  
 2251 Provenance may be used for verification of authenticity of a resource.

2252 Mechanisms that aid in determining willingness make use of the artifacts referenced by  
 2253 descriptions of services. Mechanisms for establishing willingness could be as simple as  
 2254 rendering service description information for human consumption to automated  
 2255 evaluation of functionality, policies, and contracts by a rules engine. The rules engine  
 2256 for determining willingness could operate as a policy decision procedure as defined in  
 2257 Section 4.4.

### 2258 4.2.2.3 Reachability

2259 Reachability involves knowing the endpoint, protocol, and presence of a service. At a  
 2260 minimum, reachability requires information about the location of the service and the  
 2261 protocol describing the means of communication.

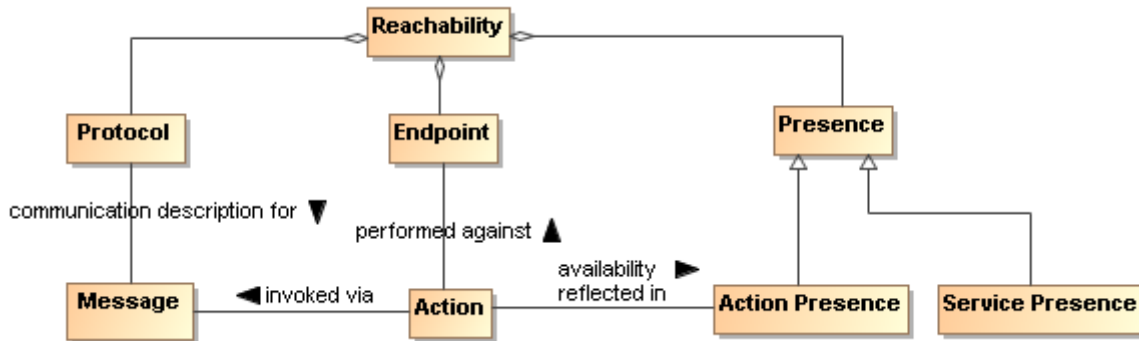


Figure 28 Service Reachability

## 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 have the following minimum capabilities:
  - creation of Description, preferably conforming to a standard Description format and structure;



- 2294 ○ publishing of Description directly to a consumer or through a third party
- 2295 mediator;
- 2296 ○ discovery of Description, preferably conforming to a standard for Description
- 2297 discovery;
- 2298 ○ notification of Description updates or notification of the addition of new and
- 2299 relevant Descriptions;
- 2300 ○ classification of Description elements according to standardized classification
- 2301 schemes.
- 2302 ● In a SOA ecosystem with complex social structures, awareness may be provided for
- 2303 specific communities of interest. The architectural mechanisms for providing
- 2304 awareness to communities of interest require support for:
  - 2305 ○ policies that allow dynamic formation of communities of interest;
  - 2306 ○ trust that awareness can be provided for and only for specific communities of
  - 2307 interest, the bases of which is typically built on keying and encryption
  - 2308 technology.
- 2309 ● The architectural mechanisms for determining willingness to interact require support
- 2310 for:
  - 2311 ○ verification of identity and credentials of the provider and/or consumer;
  - 2312 ○ access to and understanding of description;
  - 2313 ○ inspection of functionality and capabilities;
  - 2314 ○ inspection of policies and/or contracts.
- 2315 ● The architectural mechanisms for establishing reachability require support for:
  - 2316 ○ the location or address of an endpoint;
  - 2317 ○ verification and use of a service interface by means of a communication
  - 2318 protocol;
  - 2319 ○ determination of presence with an endpoint which may only be determined at
  - 2320 the point of interaction but may be further aided by the use of a presence
  - 2321 protocol for which the endpoints actively participate.

## 2322 4.3 Interacting with Services Model

2323 Interaction is the activity involved in using a service to access capability in order to

2324 achieve a particular desired real world effect, where real world effect is the actual *result*

2325 of using a service. An interaction can be characterized by a sequence of actions.

2326 Consequently, interacting with a service, i.e. performing actions against the service—

2327 usually mediated by a series of message exchanges—involves actions performed by

2328 the service. Different modes of interaction are possible such as modifying the shared

2329 state of a resource. Note that a participant (or delegate acting on behalf of the

2330 participant) can be the sender of a message, the receiver of a message, or both.

### 2331 4.3.1 Interaction Dependencies

2332 Recall from the Reference Model that service visibility is the capacity for those with

2333 needs and those with capabilities to be able to interact with each other, and that the

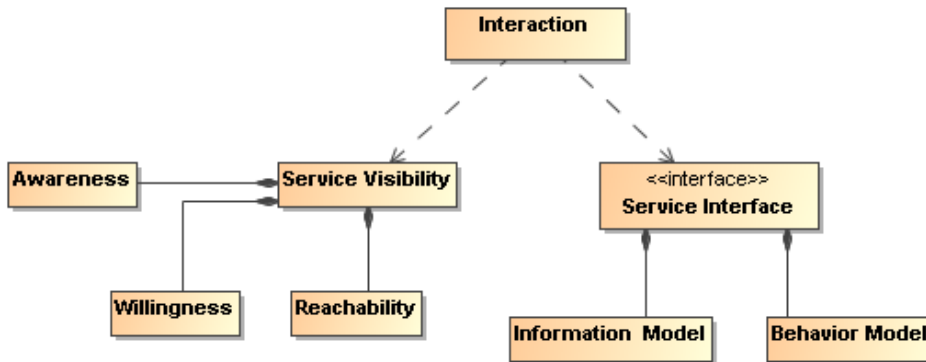
2334 service interface is the means by which the underlying capabilities of a service are

2335 accessed. Ideally, the details of the underlying service implementation are abstracted

2336 away by the service interface. [Service] interaction therefore has a direct dependency

2337 on the visibility of the service as well as its implementation-neutral interface (see Figure

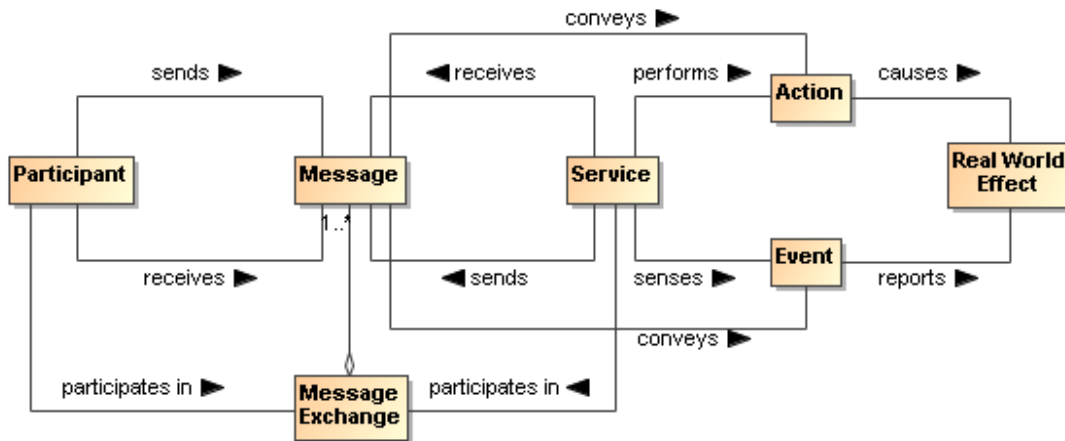
2338 29). Service visibility is composed of awareness, willingness, and reachability and  
 2339 service interface is composed of the information and behavior models. Service visibility  
 2340 is modeled in Section 4.2 while service interface is modeled in Section 4.1.



2341  
 2342 *Figure 29 Interaction dependencies.*

2343 **4.3.2 Actions and Events**

2344 For purposes of the SOA-RAF, the authors have committed to the use of message  
 2345 exchange between service participants to denote actions performed against and by the  
 2346 service, and to denote events that report on real world effects that are caused by the  
 2347 service actions. A visual model of the relationship between these concepts is shown in  
 2348 Figure 30.



2349  
 2350 *Figure 30 A "message" conveys either an action or an event.*

2351 A message conveys either an action or an event. In other words, both actions and  
 2352 events, realized by the SOA services, are denoted by the messages. The Reference  
 2353 Model states that the action model characterizes the "permissible set of actions that  
 2354 may be invoked against a service." We extend that notion here to include events as  
 2355 part of the event model and that messages denote either actions or notification of  
 2356 events.

2357 In Section **Error! Reference source not found.**, we saw that participants interact with  
 2358 each other in order to perform actions. An action is not itself the same thing as the  
 2359 result of performing the action. When an action is performed against a service, the real  
 2360 world effect that results is reported in the form of notification of events.

### 2361 4.3.3 Message Exchange

2362 *Message exchange* is the means by which service participants (or their delegates)  
2363 interact with each other. There are two primary modes of interaction: joint actions that  
2364 cause real world effects, and notification of events that report real world effects.<sup>12</sup>

2365 A message exchange is used to affect an action when the messages contain the  
2366 appropriately formatted content that should be interpreted as joint action and the  
2367 delegates involved interpret the message appropriately.

2368 A message exchange is also used to communicate event notifications. An event is an  
2369 occurrence that is of interest to some participant; in our case when some real world  
2370 effect has occurred. Just as action messages have formatting requirements, so do  
2371 event notification messages. In this way, the Information Model of a service must  
2372 specify the syntax (structure), and semantics (meaning) of the action messages and  
2373 event notification messages as part of a service interface. It must also specify the  
2374 syntax and semantics of any data that is carried as part of a payload of the action or  
2375 event notification message. The Information Model is described in greater detail in the  
2376 Service Description Model (see Section 4.1).

2377 In addition to the Information Model that describes the syntax and semantics of the  
2378 messages and data payloads, exception conditions and error handling in the event of  
2379 faults (e.g., network outages, improper message formats, etc.) must be specified or  
2380 referenced as part of the Service Description.

2381 When a message is interpreted as an action, the correct interpretation typically requires  
2382 the receiver to perform a set of operations. These *operations* represent the sequence  
2383 of actions (often private) a service must perform in order to validly participate in a given  
2384 joint action.

2385 Similarly, the correct consequence of realizing a real world effect may be to initiate the  
2386 reporting of that real world effect via an event notification.

#### 2387 **Message Exchange**

2388       The means by which joint action and event notifications are coordinated by  
2389       service participants (or delegates).

#### 2390 **Operations**

2391       The sequence of actions a service must perform in order to validly participate in a  
2392       given joint action.

### 2393 4.3.3.1 Message Exchange Patterns (MEPs)

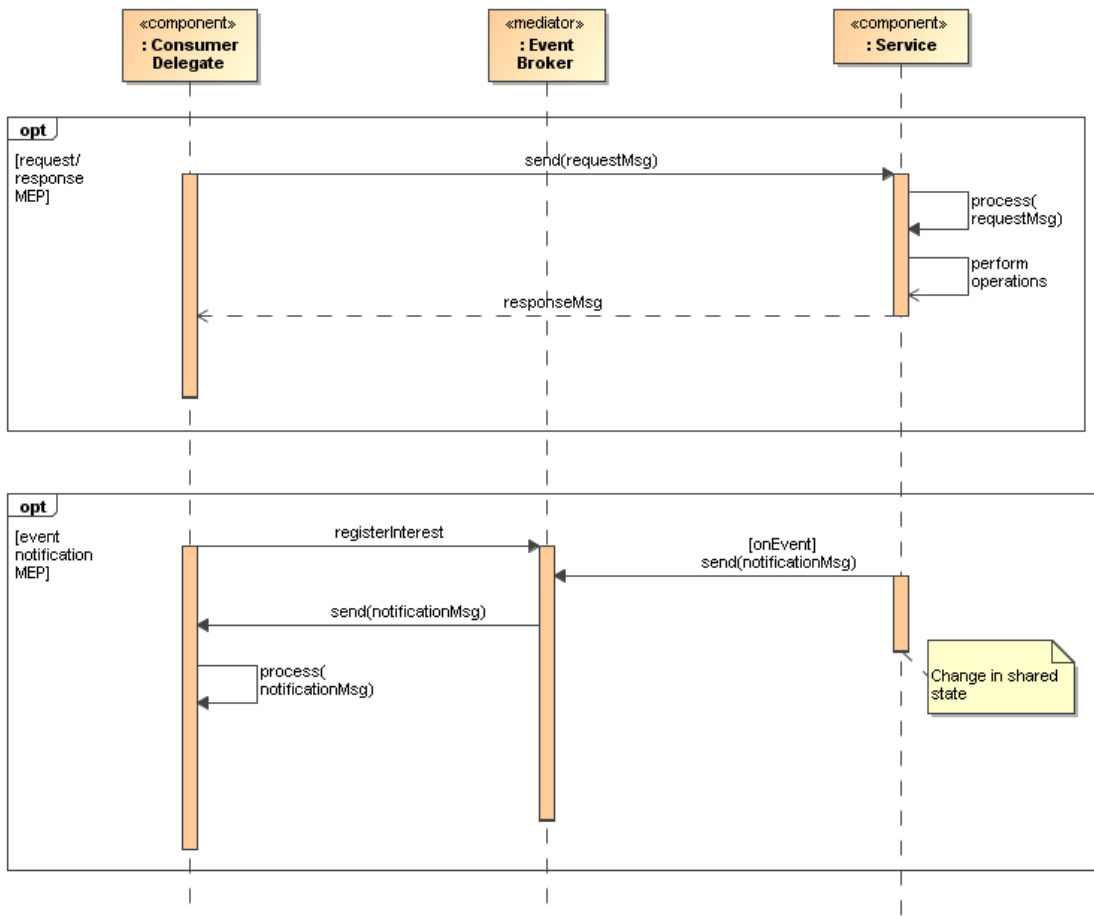
2394 The SOA-RAF commits to the use of message exchange to denote actions against the  
2395 services, and to denote notification of events that report on real world effects that arise  
2396 from those actions.

2397 Based on these assumptions, the basic temporal aspect of service interaction can be  
2398 characterized by two fundamental message exchange patterns (MEPs):

---

<sup>12</sup> The notion of “joint” in joint action implies that you have to have a speaker *and* a listener in order to interact.

- 2399 • Request/response to represent how actions cause a real world effect
  - 2400 • Event notification to represent how events report a real world effect
- 2401 This is by no means a complete list of all possible MEPs used for inter- or intra-  
 2402 enterprise messaging but it does represent those that are most commonly used in  
 2403 exchange of information and reporting changes in state both within organizations and  
 2404 across organizational boundaries, a hallmark of a SOA.
- 2405 Recall from the Reference Model that the Process Model characterizes “the temporal  
 2406 relationships between and temporal properties of actions and events associated with  
 2407 interacting with the service.” Thus, MEPs are a key element of the Process Model. The  
 2408 meta-level aspects of the Process Model (just as with the Action Model) are provided as  
 2409 part of the Service Description Model (see Section 4.1).



2410  
 2411 *Figure 31 Fundamental SOA message exchange patterns (MEPs)*

2412 In the UML sequence diagram shown in Figure 31 it is assumed that the service  
 2413 participants (consumer and provider) have delegated message handling to hardware or  
 2414 software delegates acting on their behalf. In the case of the service consumer, this is  
 2415 represented by the *Consumer Delegate* component. In the case of the service provider,  
 2416 the delegate is represented by the *Service* component. The message interchange  
 2417 model illustrated represents a logical view of the MEPs and not a physical view. In  
 2418 other words, specific hosts, network protocols, and underlying messaging system are  
 2419 not shown as these tend to be implementation specific. Although such implementation-

2420 specific elements are considered outside the scope of this document, they are important  
2421 considerations in modeling the SOA execution context. Recall from the Reference  
2422 Model that the *execution context* of a service interaction is “the set of infrastructure  
2423 elements, process entities, policy assertions and agreements that are identified as part  
2424 of an instantiated service interaction, and thus forms a path between those with needs  
2425 and those with capabilities.”

#### 2426 **4.3.3.2 Request/Response MEP**

2427 In a request/response MEP, the Consumer Delegate component sends a request  
2428 message to the Service component. The Service component then processes the  
2429 request message. Based on the content of the message, the Service component  
2430 performs the service operations. Following the completion of these operations, a  
2431 response message is returned to the Consumer Delegate component. The response  
2432 could be that a step in a process is complete, the initiation of a follow-on operation, or  
2433 the return of requested information.<sup>13</sup>

2434 Although the sequence diagram shows a *synchronous* interaction (because the sender  
2435 of the request message, i.e., Consumer Delegate, is blocked from continued processing  
2436 until a response is returned from the Service) other variations of request/response are  
2437 valid, including *asynchronous* (non-blocking) interaction through use of queues,  
2438 channels, or other messaging techniques.

2439 What is important to convey here is that the request/response MEP represents action,  
2440 which causes a real world effect, irrespective of the underlying messaging techniques  
2441 and messaging infrastructure used to implement the request/response MEP.

#### 2442 **4.3.3.3 Event Notification MEP**

2443 An event is made visible to interested consumers by means of an event notification  
2444 message exchange that reports a real world effect; specifically, a change in shared  
2445 state between service participants. The basic event notification MEP takes the form of a  
2446 one-way message sent by a notifier component (in this case, the Service component)  
2447 and received by components with an interest in the event (here, the Consumer Delegate  
2448 component).

2449 Often the sending component may not be fully aware of all the components that receive  
2450 the notification; particularly in so-called publish/subscribe (“pub/sub”) situations. In  
2451 event notification message exchanges, it is rare to have a tightly-coupled link between  
2452 the sending and the receiving component(s) for a number of practical reasons. One of  
2453 the most common is the potential for network outages or communication interrupts that

---


<sup>13</sup> 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.

2454 can result in loss of notification of events. Therefore, a third-party mediator component  
2455 is often used to decouple the sending and receiving components .

2456 Although this is typically an implementation issue, because this type of third-party  
2457 decoupling is so common in event-driven systems, it is warranted for use in modeling  
2458 this type of message exchange in the SOA-RAF. This third-party intermediary is shown  
2459 in Figure 31 as an Event Broker mediator. As with the request/response MEP, no  
2460 distinction is made between synchronous versus asynchronous communication,  
2461 although asynchronous message exchange is illustrated in the UML sequence diagram  
2462 depicted in Figure 31 .

#### 2463 4.3.4 Composition of Services

2464 Composition of services is the act of aggregating or “composing” a single service from  
2465 one or more other services. A simple model of service composition is illustrated in  
2466 Figure 32.



2467  
2468 *Figure 32 Simple model of service composition.*

2469 Here, Service A is a service that has an exposed interface IServiceA, which is available  
2470 to the Consumer Delegate and relies on two other services in its implementation. The  
2471 Consumer Delegate does not know that Services B and C are used by Service A, or  
2472 whether they are used in serial or parallel, or if their operations succeed or fail. The  
2473 Consumer Delegate only cares about the success or failure of Service A. The exposed  
2474 interfaces of Services B and C (IService B and IServiceC) are not necessarily hidden  
2475 from the Consumer Delegate; only the fact that these services are used as part of the  
2476 composition of Service A. In this example, there is no practical reason the Consumer  
2477 Delegate could not interact with Service B or Service C in some other interaction  
2478 scenario.

2479 It is possible for a service composition to be opaque from one perspective and  
2480 transparent from another. For example, a service may appear to be a single service  
2481 from the Consumer’s Delegate’s perspective, but is transparently composed of one or  
2482 more services from a service management perspective. A Service Management Service  
2483 needs to be able to have visibility into the composition in order to properly manage the  
2484 dependencies between the services used in constructing the composite service—  
2485 including managing the service’s lifecycle. The subject of services as management  
2486 entities is described and modeled in the *Ownership in a SOA Ecosystem View* of the  
2487 SOA-RAF and is not further elaborated in this section. The point to be made here is  
2488 that there can be different levels of opaqueness or transparency when it comes to  
2489 visibility of service composition.

2490 Services can be composed in a variety of ways including direct service-to-service  
2491 interaction by using programming techniques, or they can be aggregated by means of a  
2492 scripting approach that leverages a service composition scripting language. Such  
2493 scripting approaches are further elaborated in the following sub-sections on service-  
2494 oriented business processes and collaborations.

#### 2495 **4.3.4.1 Service-Oriented Business Processes**

2496 The concepts of business processes and collaborations in the context of transactions  
2497 and exchanges across organizational boundaries are described and modeled as part of  
2498 the *Participation in a SOA Ecosystem* view of this reference architecture (see Section  
2499 **Error! Reference source not found.**). Here, we focus on the belief that the principle of  
2500 composition of services can be applied to business processes and collaborations. Of  
2501 course, business processes and collaborations traditionally represent complex, multi-  
2502 step business functions that may involve multiple participants, including internal users,  
2503 external customers, and trading partners. Therefore, such complexities cannot simply  
2504 be ignored when transforming traditional business processes and collaborations to their  
2505 service-oriented variants.

#### 2506 **Business Processes**

2507 Business processes are a set of one or more linked activities that are performed  
2508 to achieve a certain business outcome.

2509 Service orientation as applied to business processes (i.e., “service-oriented business  
2510 processes”) means that the aggregation or composition of all of the abstracted activities,  
2511 flows, and rules that govern a business process can themselves be abstracted as a  
2512 service **[BLOOMBERG/SCHMELZER]**.

2513 When business processes are abstracted in this manner and accessed through SOA  
2514 services, all of the concepts used to describe and model composition of services that  
2515 were articulated in Section 4.3.4 apply. There are some important differences from a  
2516 composite service that represents an abstraction of a business process from a  
2517 composite service that represents a single-step business interaction. As stated earlier,  
2518 business processes have temporal properties and can range from short-lived processes  
2519 that execute on the order of minutes or hours to long-lived processes that can execute  
2520 for weeks, months, or even years. Further, these processes may involve many  
2521 participants. These are important considerations for the consumer of a service-oriented  
2522 business process and these temporal properties must be articulated as part of the meta-  
2523 level aspects of the service-oriented business process in its Service Description, along  
2524 with the meta-level aspects of any sub-processes that may be of use or need to be  
2525 visible to the service consumer.

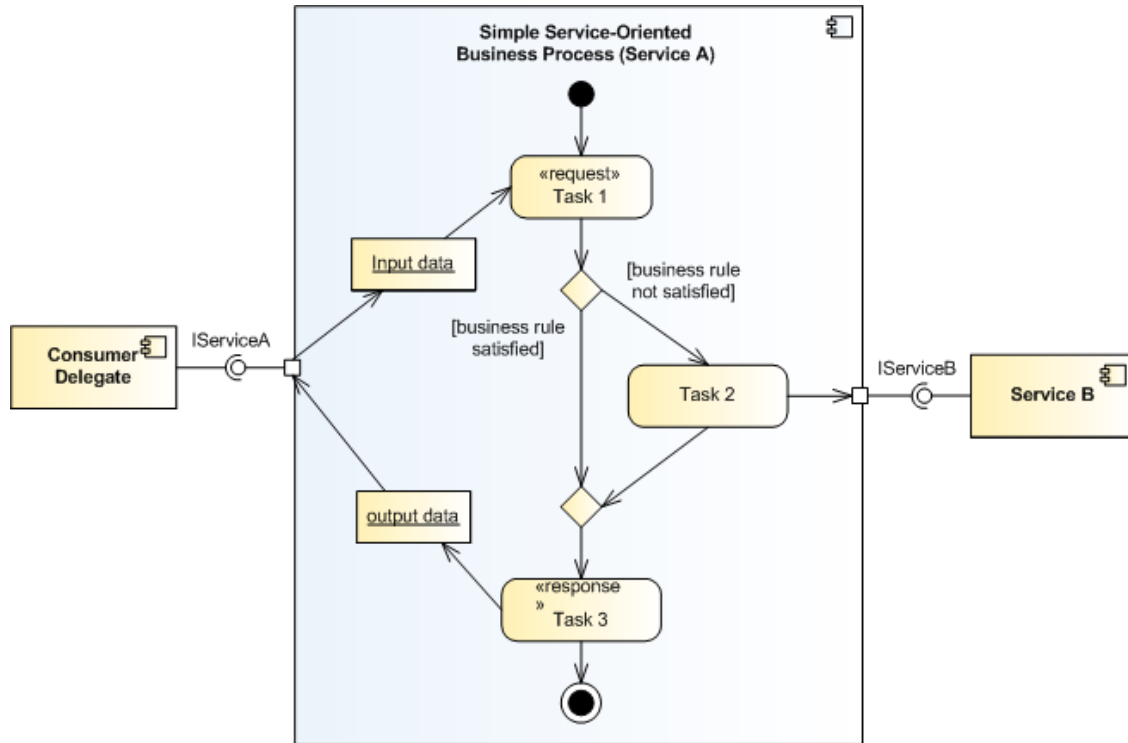
2526 In addition, a workflow activity represents a unit of work that some entity acting in a  
2527 described role (i.e., role player) is asked to perform. Activities can be broken down into  
2528 steps with each step representing a task for the role player to perform. A technique that  
2529 is used to compose service-oriented business processes that are hierarchical (top-  
2530 down) and self-contained in nature is known as *orchestration*.

#### 2531 **Orchestration**

2532 A technique used to compose service-oriented business processes that are  
2533 executed and coordinated by an actor acting as “conductor.”

2534 An orchestration is typically implemented using a scripting approach to compose  
 2535 service-oriented business processes. This typically involves use of a standards-based  
 2536 orchestration scripting language. In terms of automation, an orchestration can be  
 2537 mechanized using a business process orchestration engine, which is a hardware or  
 2538 software component (delegate) responsible for acting in the role of central  
 2539 conductor/coordinator responsible for executing the flows that comprise the  
 2540 orchestration.

2541 A simple generic example of such an orchestration is illustrated in Figure 33.



2542  
 2543 *Figure 33 Abstract example of orchestration of service-oriented business process.*

2544 Here, we use a UML activity diagram to model the simple service-oriented business  
 2545 process as it allows us to capture the major elements of business processes such as  
 2546 the set of related tasks to be performed, linking between tasks in a logical flow, data that  
 2547 is passed between tasks, and any relevant business rules that govern the transitions  
 2548 between tasks. A task is a unit of work that an individual, system, or organization  
 2549 performs and can be accomplished in one or more steps or subtasks. While subtasks  
 2550 can be readily modeled, they are not illustrated in the orchestration model In Figure 33..

2551 This particular example is based on a request/response MEP and captures how one  
 2552 particular task (Task 2) actually utilizes an externally-provided service, Service B. The  
 2553 entire service-oriented business process is exposed as Service A that is accessible via  
 2554 its externally visible interface, IServiceA.

2555 Although not explicitly shown in the orchestration model above, it is assumed that there  
 2556 exists a software or hardware component, i.e., orchestration engine that executes the  
 2557 process flow. Recall that a central concept to orchestration is that process flow is  
 2558 coordinated and executed by a single conductor delegate; hence the name  
 2559 “orchestration.”



#### 2560 4.3.4.2 Service-Oriented Business Collaborations

2561 Business collaborations typically represent the interaction involved in executing  
2562 business transactions, where a business transaction is defined in the *Participation in a*  
2563 *SOA Ecosystem* view as “a joint action engaged in by two or more participants in which  
2564 resources are exchanged” (see Section **Error! Reference source not found.**).

2565 It is important to note that business collaborations represent “peer”-style interactions; in  
2566 other words, peers in a business collaboration act as equals. This means that unlike  
2567 the orchestration of business processes, there is no single or central entity that  
2568 coordinates or “conducts” a business collaboration. These peer styles of interactions  
2569 typically occur between trading partners that span organizational boundaries.

2570 Business collaborations can also be service-enabled. For purposes of this Reference  
2571 Architecture Foundation, we refer to these as “service-oriented business collaborations.”  
2572 Service-oriented business collaborations do not necessarily imply exposing the entire  
2573 peer-style business collaboration as a service itself but rather the collaboration uses  
2574 service-based interchanges.

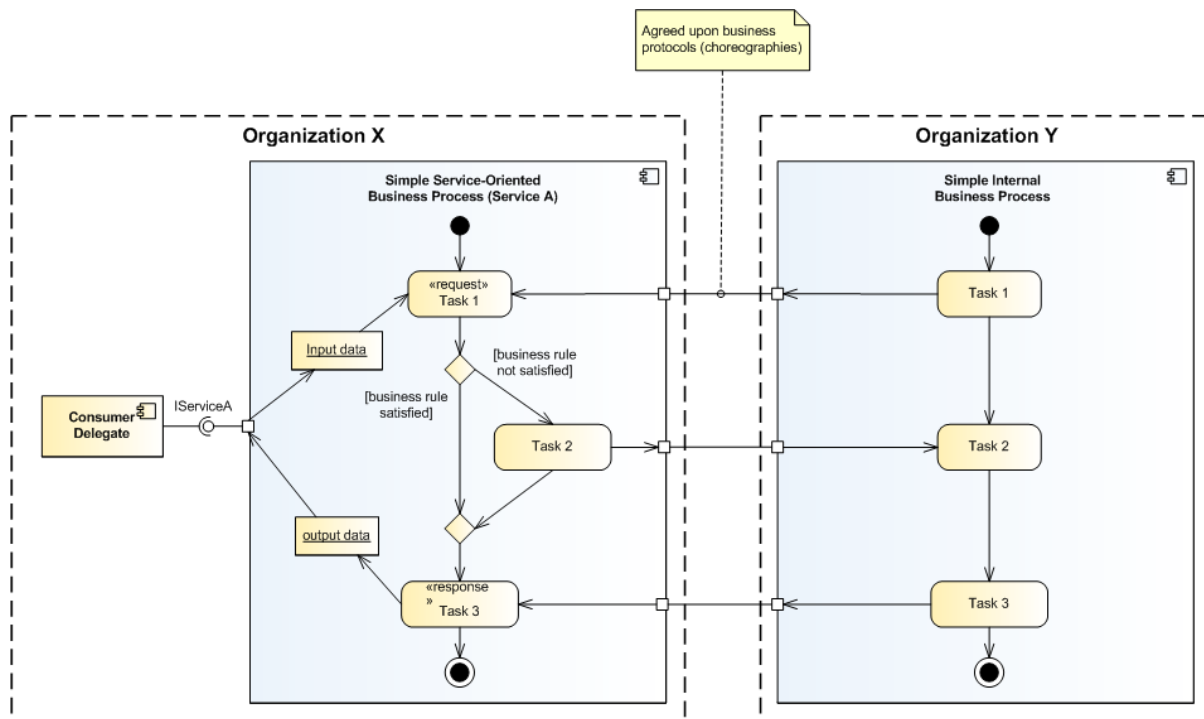
2575 The technique that is used to compose service-oriented business collaborations in  
2576 which multiple parties collaborate in a peer-style as part of some larger business  
2577 transaction by exchanging messages with trading partners and external organizations  
2578 (e.g., suppliers) is known as *choreography* [NEWCOMER/LOMOW].

#### 2579 **Choreography**

2580 A technique used to characterize service-oriented business collaborations based  
2581 on ordered message exchanges between peer entities in order to achieve a  
2582 common business goal.

2583 Choreography differs from orchestration primarily in that each party in a business  
2584 collaboration describes its part in the service interaction. Note that choreography as we  
2585 have defined it here should not be confused with the term *process choreography*, which  
2586 is defined in the *Participation in a SOA Ecosystem* view as “the description of the  
2587 possible interactions that may take place between two or more participants to fulfill an  
2588 objective.” This is an example of domain-specific nomenclature that often leads to  
2589 confusion and why we are making note of it here.

2590 A simple generic example of a choreography is illustrated in Figure 34



2591  
2592

Figure 34 Abstract example of choreography of service-oriented business collaboration.

2593 This example, which is a variant of the orchestration example illustrated earlier in Figure  
2594 33 adds trust boundaries between two organizations; namely, Organization X and  
2595 Organization Y. It is assumed that these two organizations are peer entities that have  
2596 an interest in a business collaboration, for example, Organization X and Organization Y  
2597 could be trading partners. Organization X retains the service-oriented business process  
2598 Service A, which is exposed to internal consumers via its provided service interface,  
2599 IServiceA. Organization Y also has a business process that is involved in the business  
2600 collaboration; however, for this example, it is an internal business process that is not  
2601 exposed to potential consumers either within or outside its organizational boundary.

2602 The scripting language that is used for the choreography needs to define how and when  
2603 to pass control from one trading partner to another, i.e., Organization X and  
2604 Organization Y. Defining the business protocols used in the business collaboration  
2605 involves precisely specifying the visible message exchange behavior of each of the  
2606 parties involved in the protocol, without revealing internal implementation details  
2607 **[NEWCOMER/LOMOW]**.

2608 In a peer-style business collaboration, a choreography scripting language must be  
2609 capable of describing the coordination of those service-oriented processes that cross  
2610 organizational boundaries.

### 2611 4.3.5 Architectural Implications of Interacting with Services

2612 Interacting with Services has the following architectural implications on mechanisms  
2613 that facilitate service interaction:

- 2614 • A well-defined service Information Model that:
  - 2615 ○ describes the syntax and semantics of the messages used to denote actions
  - 2616 ○ and events;

- 2617 ○ describes the syntax and semantics of the data payload(s) contained within
- 2618 messages;
- 2619 ○ documents exception conditions in the event of faults due to network outages,
- 2620 improper message/data formats, etc.;
- 2621 ○ is both human readable and machine processable;
- 2622 ○ is referenceable from the Service Description artifact.
- 2623 ● A well-defined service Behavior Model that:
  - 2624 ○ characterizes the knowledge of the actions invokes against the service and
  - 2625 events that report real world effects as a result of those actions;
  - 2626 ○ characterizes the temporal relationships and temporal properties of actions
  - 2627 and events associated in a service interaction;
  - 2628 ○ describe activities involved in a workflow activity that represents a unit of
  - 2629 work;
  - 2630 ○ describes the role (s) that a role player performs in a service-oriented
  - 2631 business process or service-oriented business collaboration;
  - 2632 ○ is both human readable and machine processable;
  - 2633 ○ is referenceable from the Service Description artifact.
- 2634 ● Service composition mechanisms to support orchestration of service-oriented
- 2635 business processes and choreography of service-oriented business collaborations
- 2636 such as:
  - 2637 ○ Declarative and programmatic compositional languages;
  - 2638 ○ Orchestration and/or choreography engines that support multi-step
  - 2639 processes as part of a short-lived or long-lived business transaction;
  - 2640 ○ Orchestration and/or choreography engines that support compensating
  - 2641 transactions in the presences of exception and fault conditions.
- 2642 ● Infrastructure services that provides mechanisms to support service interaction,
- 2643 including but not limited to:
  - 2644 ○ mediation services such as message and event brokers, providers, and/or
  - 2645 buses that provide message translation/transformation, gateway
  - 2646 capability, message persistence, reliable message delivery, and/or
  - 2647 intelligent routing semantics;
  - 2648 ○ binding services that support translation and transformation of multiple
  - 2649 application-level protocols to standard network transport protocols;
  - 2650 ○ auditing and logging services that provide a data store and mechanism to
  - 2651 record information related to service interaction activity such as message
  - 2652 traffic patterns, security violations, and service contract and policy
  - 2653 violations
  - 2654 ○ security services that abstract techniques such as public key
  - 2655 cryptography, secure networks, virus protection, etc., which provide
  - 2656 protection against common security threats in a SOA ecosystem;
  - 2657 ○ monitoring services such as hardware and software mechanisms that both
  - 2658 monitor the performance of systems that host services and network traffic
  - 2659 during service interaction, and are capable of generating regular
  - 2660 monitoring reports.
- 2661 ● A layered and tiered service component architecture that supports multiple message
- 2662 exchange patterns (MEPs) in order to:

- 2663 ○ promote the industry best practice of separation of concerns that facilitates
- 2664 flexibility in the presence of changing business requirements;
- 2665 ○ promote the industry best practice of separation of roles in a service
- 2666 development lifecycle such that subject matter experts and teams are
- 2667 structured along areas of expertise;
- 2668 ○ support numerous standard interaction patterns, peer-to-peer interaction
- 2669 patterns, enterprise integration patterns, and business-to-business
- 2670 integration patterns.

## 2671 4.4 Policies and Contracts Model

2672 A common phenomenon of many machines and systems is that the scope of potential  
2673 behavior is much broader than is actually needed for a particular circumstance. This is  
2674 especially true of a system as powerful as a SOA ecosystem. As a result, the behavior  
2675 and performance of the system tend to be under-constrained by the implementation;  
2676 instead, the actual behavior is expressed by means of policies of some form. Policies  
2677 define the choices that stakeholders make; these choices are used to guide the actual  
2678 behavior of the system to the desired behavior and performance.

2679 As noted in Section 3.1.5 a policy is a constraint of some form that is promulgated by a  
2680 stakeholder who has the responsibility of ensuring that the constraint is enforced. In  
2681 contrast, contracts are **agreements** between participants. However, like policies, it is a  
2682 necessary part of contracts that they are enforceable.

2683 While responsibility for enforcement may differ, both contracts and policies share a  
2684 common characteristic – there is a **constraint** that must be enforced. In both cases the  
2685 mechanisms needed to enforce policy constraints are likely to be identical; in this model  
2686 we focus on the issues involved in representing policies and contracts and on some of  
2687 the principles behind their enforcement.

### 2688 4.4.1 Policy and Contract Representation

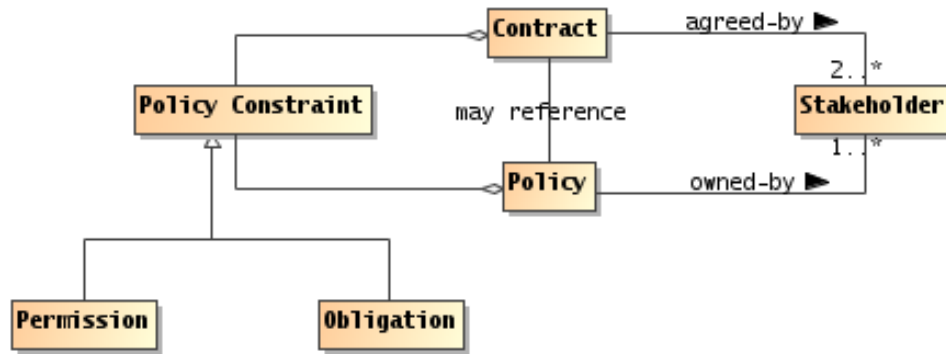
2689 A **policy constraint** is a specific kind of constraint: the ontology of policies and  
2690 contracts includes the core concepts of permission, obligation, owner, subject. In  
2691 addition, it may be necessary to be able combine policy constraints and to be able to  
2692 resolve policy conflicts.

#### 2693 4.4.1.1 Policy Framework

##### 2694 Policy Framework

2695 A policy framework is a language in which policy constraints may be expressed.

2696 A policy framework combines a syntax for expressing policy constraints together with a  
2697 decision procedure for determining if a policy constraint is satisfied.



2698  
2699 *Figure 35 Policies and Contracts*

2700 We can characterize (caricature) a policy framework in terms of a logical framework and  
2701 an ontology of policies. The policy ontology details specific kinds of policy constraints  
2702 that can be expressed; and the logical framework is a 'glue' that allows us to express  
2703 combinations of policies.

2704 **Logical Framework**

2705 A logical framework is a linguistic framework consisting of a syntax – a way of  
2706 writing expressions – and a semantics – a way of interpreting the expressions.

2707 **Policy Ontology**

2708 A policy ontology is a formalization of a set of concepts that are relevant to  
2709 forming policy expressions.

2710 For example, a policy ontology that allows to identify simple constraints – such as the  
2711 existence of a property, or that a value of a property should be compared to a fixed  
2712 value – is often enough to express many basic constraints.

2713 Included in many policy ontologies are the basic signals of permissions and obligations.  
2714 Some policy frameworks are sufficiently constrained that there is not possibility of  
2715 representing an obligation; in which case there is often no need to 'call out' the  
2716 distinction between permissions and obligations.

2717 The logical framework is also a strong determiner of the expressivity of the policy  
2718 framework. The richer the logical framework, the richer the set of policy constraints that  
2719 can be expressed. However, there is a strong inverse correlation between expressivity  
2720 and ease and efficiency of implementation.

2721 In the discussion that follows we assume the following basic policy ontology:

2722 **Policy Owner**

2723 A policy owner is a stakeholder that asserts and enforces the policy.

2724 **Policy Subject**

2725 A policy subject is an actor who is subject to the constraints of a policy or  
2726 contract.

2727 **Policy Constraint**

2728 A policy constraint is a measurable proposition that characterizes the constraint  
2729 that the policy is about.

2730 **Policy Object**

2731 A policy object is an identifiable state, action or resource that is potentially  
2732 constrained by the policy.

2733 **4.4.2 Policy and Contract Enforcement**

2734 The enforcement of policy constraints has to address two core problems: how to  
2735 enforce the atomic policy constraints, and how to enforce combinations of policy  
2736 constraints. In addition, it is necessary to address the resolution of policy conflicts.

2737 **4.4.2.1 Enforcing Simple Policy Constraints**

2738 The two primary kinds of policy constraint – permission and obligation – naturally lead to  
2739 different styles of enforcement. A permission constraint must typically be enforced *prior*  
2740 to the policy subject invoking the **policy object**. On the hand, an obligation constraint  
2741 must typically be enforced post-facto through some form of auditing process and  
2742 remedial action.

2743 For example, if a communications policy required that all communication be encrypted,  
2744 this is enforceable at the point of communication: any attempt to communicate a  
2745 message that is not encrypted can be blocked.

2746 Similarly, an obligation to pay for services rendered is enforced by ensuring that  
2747 payment arrives within a reasonable period of time. Invoices are monitored for prompt  
2748 (or lack of) payment.

2749 The key concepts in enforcing both forms of policy constraint are the policy decision and  
2750 the policy enforcement.

2751 **Policy Decision**

2752 A policy decision is a determination as to whether a given policy constraint is  
2753 satisfied or not.

2754 A policy decision is effectively a measurement of some state – typically a portion of the  
2755 SOA ecosystem's **shared state**. This implies a certain *timeliness* in the measuring: a  
2756 measurement that is too early or is too late does not actually help in determining if the  
2757 policy constraint is satisfied appropriately.

2758 **Policy Enforcement**

2759 A policy enforcement is the use of a mechanism to limit the behavior and/or state  
2760 of policy subjects to comply with a policy decision.

2761 A policy enforcement implies the use of some mechanism to ensure compliance with a  
2762 policy decision. The range of mechanisms is completely dependent on the kinds of  
2763 atomic policy constraints that the policy framework may support. As noted above, the  
2764 two primary styles of constraint – permission and **obligation** –lead to different styles of  
2765 enforcement.

2766 **4.4.2.2 Enforcing Policy Combinations**

2767 Enforcing policy combinations is primarily an elaboration of enforcing simple policy  
2768 constraints. The process of policy decisions is enhanced to allow a measurement to  
2769 involve combinations of policy constraints and the process of policy enforcement may

2770 need to be enhanced to coordinate the enforcement of multiple policy constraints  
2771 simultaneously.

#### 2772 **4.4.2.3 Conflict Resolution**

2773 Whenever it is possible that more than one policy constraint applies in a given situation,  
2774 there is the potential that the policies themselves are not mutually consistent. For  
2775 example, a policy that requires communication to be encrypted and a policy that  
2776 requires an administrator to read every communication conflict with each other – the two  
2777 policies cannot both be satisfied.

2778 In general, with sufficiently rich policy frameworks, it is not possible to always resolve  
2779 policy conflicts automatically. However, a reasonable approach is to augment the policy  
2780 decision process with simple policy conflict resolution rules; with the potential for  
2781 *escalating* a policy conflict to human adjudication.

#### 2782 **Policy Conflict**

2783 A policy conflict exists between two or more policies in a policy decision process  
2784 if it is not possible to satisfy all the policies that apply.

#### 2785 **Policy Conflict Resolution**

2786 A policy conflict resolution rule is a way of determining which policy should  
2787 prevail in a policy conflict.

2788 The inevitable consequence of policy conflicts is that it is not possible to guarantee that  
2789 all policies are satisfied at all times. This, in turn, implies a certain *flexibility* in the  
2790 application of policy constraints: they will not always be honored.

#### 2791 **4.4.3 Architectural Implications**

2792 The key choices that must be made in a system of policies center around the policy  
2793 framework and policy enforcement mechanisms

- 2794 • There SHOULD be a standard policy framework that is adopted across the SOA  
2795 ecosystem:
  - 2796 ○ This framework MUST permit the expression of simple policy constraints
  - 2797 ○ The framework MAY allow (to a varying extent) the combination of policy  
2798 constraints, including
    - 2799 • Both positive and negative constraints
    - 2800 • Conjunctions and disjunctions of constraints
    - 2801 • The quantification of constraints
  - 2802 ○ The framework MUST at least allow the policy subject and the policy object to  
2803 be identified as well as the policy constraint.
  - 2804 ○ The framework MAY allow further structuring of policies into modules,  
2805 inheritance between policies and so on.
- 2806 • There SHOULD be mechanisms that facilitate the application of policies:
  - 2807 ○ There SHOULD be mechanisms that allow policy decisions to be made,  
2808 consistent with the policy frameworks and with the state of the SOA  
2809 ecosystem.
  - 2810 ○ There SHOULD be mechanisms to enforce policy decisions

- 2811
- 2812
- 2813
- 2814
- 2815
- 2816
- 2817
- 2818
- 2819
- 2820
- 2821
- 2822
- There SHOULD be mechanisms to support the measurement of whether certain policy constraints are satisfied or not, or to what degree they are satisfied.
  - Such enforcement mechanisms MAY include support for both permission-style constraints and obligation-style constraints.
  - Enforcement mechanisms MAY support the simultaneous enforcement of multiple policy constraints across multiple points in the SOA ecosystem.
  - There SHOULD be mechanisms to resolve policy conflicts
    - This MAY involve escalating policy conflicts to human adjudication.
  - There SHOULD be mechanisms that support the management and promulgation of policies.



## 5 Ownership in a SOA Ecosystem 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 provides 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.

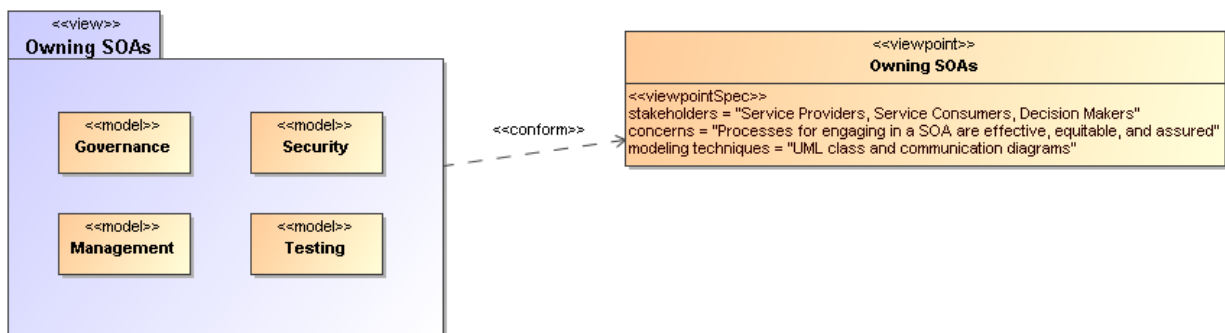


Figure 36 Model Elements Described in the Ownership in a SOA Ecosystem View

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.

## 2852 5.1.1 Understanding Governance

### 2853 5.1.1.1 Terminology

2854 Governance is about making decisions that are aligned with the overall organizational  
2855 strategy and culture of the enterprise. **[Gartner]** It specifies the decision rights and  
2856 accountability framework to encourage desirable behaviors **[Weill/Ross-MIT Sloan**  
2857 **School]** towards realizing the strategy and defines incentives (positive or negative)  
2858 towards that end. It is less about overt control and strict adherence to rules, and more  
2859 about guidance and effective and equitable usage of resources to ensure sustainability  
2860 of an organization's strategic objectives. **[TOGAF v8.1]**

2861 To accomplish this, governance requires organizational structure and processes and  
2862 must identify who has authority to define and carry out its mandates. It must address  
2863 the following questions: 1) what decisions must be made to ensure effective  
2864 management and use?, 2) who should make these decisions?, and 3) how will these  
2865 decisions be made and monitored? , and (4) how will these decisions be  
2866 communicated? The intent is to achieve goals, add value, and reduce risk.

2867 Within a single ownership domain such as an enterprise, generally there is a hierarchy  
2868 of governance structures. Some of the more common enterprise governance structures  
2869 include corporate governance, technology governance, IT governance, and architecture  
2870 governance **[TOGAF v8.1]**. These governance structures can exist at multiple levels  
2871 (global, regional, and local) within the overall enterprise.

2872 It is often asserted that SOA governance is a specialization of IT governance as there is  
2873 a natural hierarchy of these types of governance structures; however, the focus of SOA  
2874 governance is less on decisions to ensure effective management and use of IT as it is  
2875 to ensure effective management and use of SOA-based systems. Certainly, SOA  
2876 governance must still answer the basic questions also associated with IT governance,  
2877 i.e., who should make the decisions, and how these decisions will be made and  
2878 monitored.

### 2879 5.1.1.2 Relationship to Management

2880 There is often confusion centered on the relationship between governance and  
2881 management. As described earlier, governance is concerned with decision making.  
2882 Management, on the other hand, is concerned with execution. Put another way,  
2883 governance describes the world as leadership wants it to be; management executes  
2884 activities that intends to make the leadership's desired world a reality. Where  
2885 governance determines who has the authority and responsibility for making decisions  
2886 and the establishment of guidelines for how those decisions should be made,  
2887 management is the actual process of making, implementing, and measuring the impact  
2888 of those decisions **[Loeb]**. Consequently, governance and management work in  
2889 concert to ensure a well-balanced and functioning organization as well as an ecosystem  
2890 of inter-related organizations. In the sections that follow, we elaborate further on the  
2891 relationship between governance and management in terms of setting and enforcing  
2892 service policies, contracts, and standards as well as addressing issues surrounding  
2893 regulatory compliance.

### 2894 5.1.1.3 Why is SOA Governance Important?

2895 One of the hallmarks of SOA that distinguishes it from other architectural paradigms for  
2896 distributed computing is the ability to provide a uniform means to offer, discover, interact

2897 with and use capabilities (as well the ability to compose new capabilities from existing  
2898 ones) all in an environment that transcends domains of ownership. Consequently,  
2899 ownership, and issues surrounding it, such as obtaining acceptable terms and  
2900 conditions (T&Cs) in a contract, is one of the primary topics for SOA governance.  
2901 Generally, IT governance does not include T&Cs, for example, as a condition of use as  
2902 its primary concern.

2903 Just as other architectural paradigms, technologies, and approaches to IT are subject to  
2904 change and evolution, so too is SOA. Setting policies that allow change management  
2905 and evolution, establishing strategies for change, resolving disputes that arise, and  
2906 ensuring that SOA-based systems continue to fulfill the goals of the business are all  
2907 reasons why governance is important to SOA.

#### 2908 **5.1.1.4 Governance Stakeholders and Concerns**

2909 As noted in Section **Error! Reference source not found.** the participants in a service  
2910 interaction include the service provider, the service consumer, and other interested or  
2911 unintentional third parties. Depending on the circumstances, it may also include the  
2912 owners of the underlying capabilities that the SOA services access. Governance must  
2913 establish the policies and rules under which duties and responsibilities are defined and  
2914 the expectations of participants are grounded. The expectations include transparency  
2915 in aspects where transparency is mandated, trust in the impartial and consistent  
2916 application of governance, and assurance of reliable and robust behavior throughout the  
2917 SOA ecosystem.

#### 2918 **5.1.2 A Generic Model for Governance**

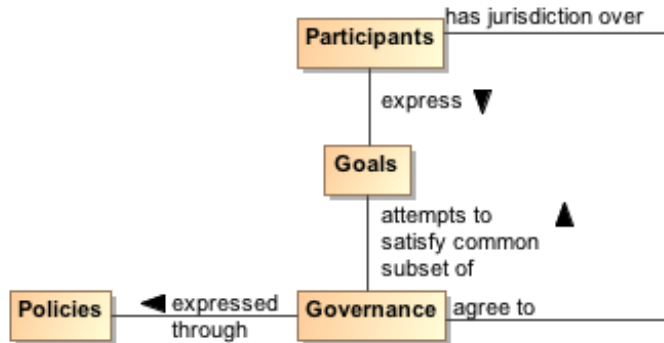
##### 2919 **Governance**

2920 Governance is the prescribing of conditions and constraints consistent with  
2921 satisfying common goals and the structures and processes needed to define and  
2922 respond to actions taken towards realizing those goals.

2923 The following is a generic model of governance represented by segmented models that  
2924 begin with motivation and proceed through measuring compliance. It is not all-  
2925 encompassing but a focused subset that captures the aspects necessary to describe  
2926 governance for SOA. It does not imply that practical application of governance is a  
2927 single, isolated instance of these models; in reality, there may be hierarchical and  
2928 parallel chains of governance that deal with different aspects or focus on different goals.  
2929 This is discussed further in section 5.1.2.5. The defined models are simultaneously  
2930 applicable to each of the overlapping instances.

2931 A given enterprise may already have portions of these models in place. To a large  
2932 extent, the models shown here are not specific to SOA; discussions on direct  
2933 applicability begin in section 5.1.3.

2934 **5.1.2.1 Motivating Governance**



2935  
2936 *Figure 37 Motivating governance model*

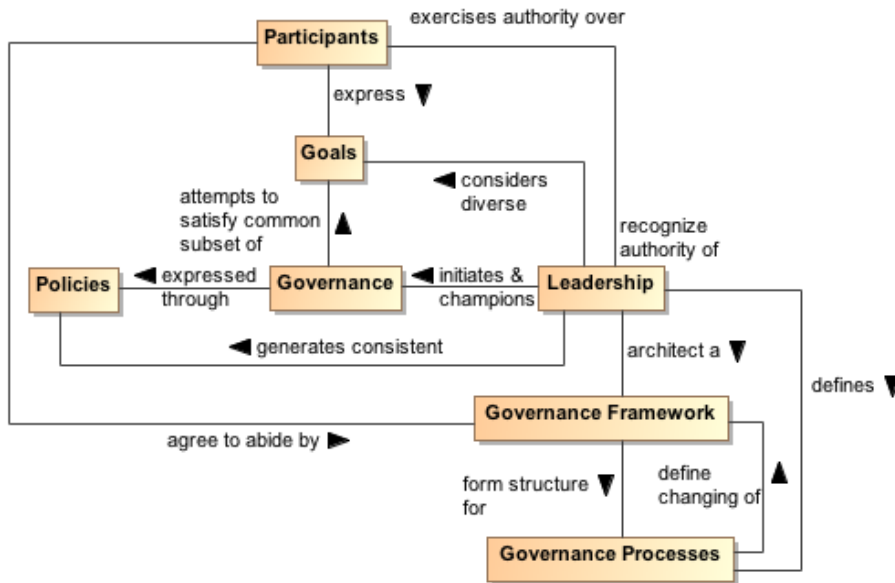
2937 An organizational domain such as an enterprise is made up of participants who may be  
2938 individuals or groups of individuals forming smaller organizational units within the  
2939 enterprise. The overall business strategy should be consistent with the Goals of the  
2940 participants; otherwise, the business strategy would not provide value to the participants  
2941 and governance towards those ends becomes difficult if not impossible. This is not to  
2942 say that an instance of governance simultaneously satisfies all the goals of all the  
2943 participants; rather, the goals of any governance instance must sufficiently satisfy a  
2944 useful subset of each participant's goals so as to provide value and ensure the  
2945 cooperation of all the participants.

2946 A policy is the formal characterization of the conditions and constraints that governance  
2947 deems as necessary to realize the goals which it is attempting to satisfy. Policy may  
2948 identify required conditions or actions or may prescribe limitations or other constraints  
2949 on permitted conditions or actions. For example, a policy may prescribe that  
2950 safeguards must be in place to prevent unauthorized access to sensitive material. It  
2951 may also prohibit use of computers for activities unrelated to the specified work  
2952 assignment. Policy is made operational through the promulgating and implementing of  
2953 Rules and Regulations (as defined in section 5.1.2.3).

2954 As noted in section 4.4.2, policy may be asserted by any participant or on behalf of the  
2955 participant by its organization. Part of the purpose of governance is to arbitrate among  
2956 diverse goals of participants and diverse policies articulated to realize those goals. The  
2957 intent is to form a consistent whole that allows governance to minimize ambiguity about  
2958 its purpose. While resolving all ambiguity would be an ideal, it is unlikely that all  
2959 inconsistencies will be identified and resolved before governance becomes operational.

2960 For governance to have effective jurisdiction over participants, there must be some  
2961 degree of agreement by all participants that they will abide by the governance  
2962 mandates. A minimal degree of agreement often presages participants who “slow-roll” if  
2963 not actively reject complying with Policies that express the specifics of governance.

2964 **5.1.2.2 Setting Up Governance**



2965  
2966 *Figure 38 Setting up governance model*

2967 **Leadership**

2968 Leadership is the entity who has the responsibility and authority to generate  
2969 consistent policies through which the goals of governance can be expressed and  
2970 to define and champion the structures and processes through which governance  
2971 is realized.

2972 **Governance Framework**

2973 The Governance Framework is a set of organizational structures that enable  
2974 governance to be consistently defined, clarified, and as needed, modified to  
2975 respond to changes in its domain of concern.

2976 **Governance Processes**

2977 Governance Processes are the defined set of activities that are performed within  
2978 the Governance Framework to enable the consistent definition, application, and  
2979 as needed, modification of Rules that organize and regulate the activities of  
2980 participants for the fulfillment of expressed policies. (See section 5.1.2.3 for  
2981 elaboration on the relationship of Governance Processes and Rules.)

2982 As noted earlier, governance requires an appropriate organizational structure and  
2983 identification of who has authority to make governance decisions. In Figure 38, the  
2984 entity with governance authority is designated the Leadership. This is someone,  
2985 possibly one or more of the participants, that participants recognize as having authority  
2986 for a given purpose or over a given set of issues or concerns.

2987 The Leadership is responsible for prescribing or delegating a working group to prescribe  
2988 the Governance Framework that forms the structure for Governance Processes which  
2989 define how governance is to be carried out. This does not itself define the specifics of  
2990 how governance is to be applied, but it does provide an unambiguous set of procedures

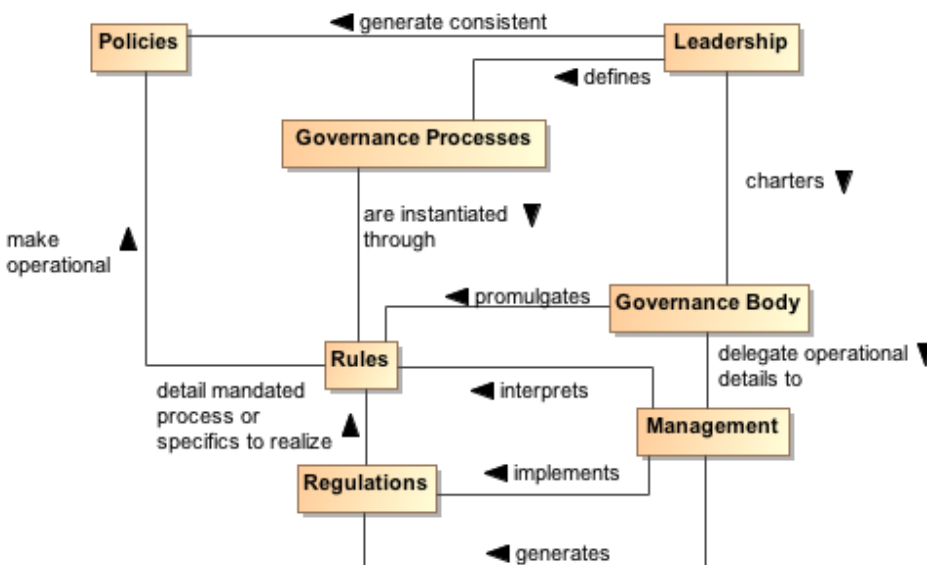
2991 that should ensure consistent actions which participants agree are fair and account for  
2992 sufficient input on the subjects to which governance is applied.

2993 The participants may be part of the working group that codifies the Governance  
2994 Framework and Processes. When complete, the participants must acknowledge and  
2995 agree to abide by the products generated through application of this structure.

2996 The Governance Framework and Processes are often documented in the charter of a  
2997 body created or designated to oversee governance. This is discussed further in the  
2998 next section. Note that the Governance Processes should also include those necessary  
2999 to modify the Governance Framework itself.

3000 An important function of Leadership is not only to initiate but also be the consistent  
3001 champion of governance. Those responsible for carrying out governance mandates  
3002 must have Leadership who makes it clear to participants that expressed Policies are  
3003 seen as a means to realizing established goals and that compliance with governance is  
3004 required.

### 3005 5.1.2.3 Carrying Out Governance



3006  
3007 *Figure 39 Carrying out governance model*

#### 3008 **Rule**

3009 A Rule is a prescribed guide for carrying out activities and processes leading to  
3010 desired results, e.g. the operational realization of policies.

#### 3011 **Regulation**

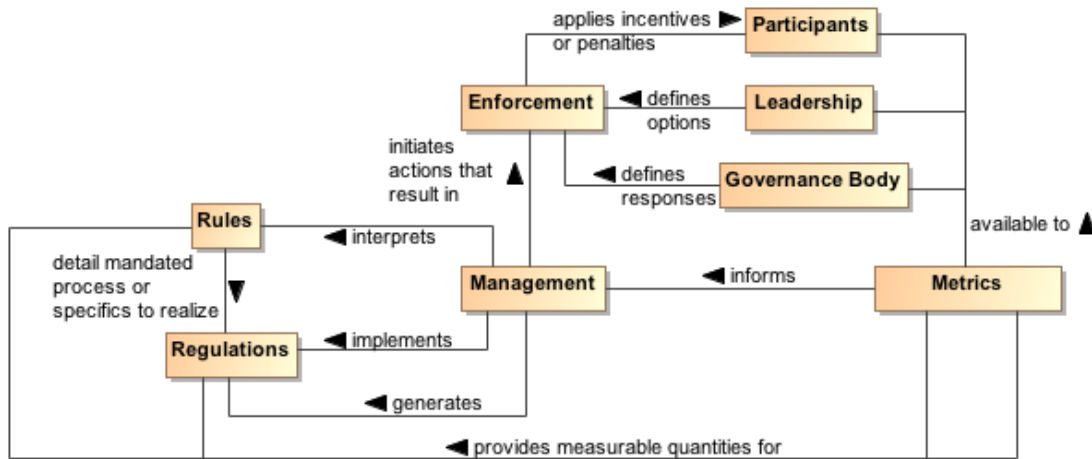
3012 A Regulation is a mandated process or the specific details that derive from the  
3013 interpretation of Rules and lead to measureable quantities against which  
3014 compliance can be measured.

3015 To carry out governance, Leadership charts a Governance Body to promulgate the  
3016 Rules needed to make the Policies operational. The Governance Body acts in line with  
3017 Governance Processes for its rule-making process and other functions. Whereas  
3018 Governance is the setting of Policies and defining the Rules that provide an operational

3019 context for Policies, the operational details of governance may be delegated by the  
 3020 Governance Body to Management. Management generates Regulations that specify  
 3021 details for Rules and other procedures to implement both Rules and Regulations. For  
 3022 example, Leadership could set a Policy that all authorized parties should have access to  
 3023 data, the Governance Body would promulgate a Rule that PKI certificates are required  
 3024 to establish identity of authorized parties, and Management can specify a Regulation of  
 3025 who it deems to be a recognized PKI issuing body. In summary, Policy is a predicate to  
 3026 be satisfied and Rules prescribe the activities by which that satisfying occurs. A number  
 3027 of rules may be required to satisfy a given policy; the carrying out of a rule may  
 3028 contribute to several policies being realized.

3029 Whereas the Governance Framework and Processes are fundamental for having  
 3030 participants acknowledge and commit to compliance with governance, the Rules and  
 3031 Regulations provide operational constraints which may require resource commitments  
 3032 or other levies on the participants. It is important for participants to consider the  
 3033 framework and processes to be fair, unambiguous, and capable of being carried out in a  
 3034 consistent manner and to have an opportunity to formally accept or ratify this situation.  
 3035 Rules and Regulations, however, do not require individual acceptance by any given  
 3036 participant although some level of community comment may be part of the Governance  
 3037 Processes. Having agreed to governance, the participants are bound to comply or be  
 3038 subject to prescribed mechanisms for enforcement.

3039 **5.1.2.4 Ensuring Governance Compliance**



3040  
 3041 *Figure 40 Ensuring governance compliance model*

3042 Setting Rules and Regulations does not ensure effective governance unless compliance  
 3043 can be measured and Rules and Regulations can be enforced. Metrics are those  
 3044 conditions and quantities that can be measured to characterize actions and results.  
 3045 Rules and Regulations MUST be based on collected Metrics or there is no means for  
 3046 Management to assess compliance. The Metrics are available to the participants, the  
 3047 Leadership, and the Governance Body so what is measured and the results of  
 3048 measurement are clear to everyone.

3049 The Leadership in its relationship with participants has certain options that can be used  
 3050 for Enforcement. A common option may be to effect future funding. The Governance  
 3051 Body defines specific enforcement responses, such as what degree of compliance is

3052 necessary for full funding to be restored. It is up to Management to identify compliance  
3053 shortfalls and to initiate the Enforcement process.

3054 Note, enforcement does not strictly need to be negative consequences. Management  
3055 can use Metrics to identify exemplars of compliance and Leadership can provide  
3056 options for rewarding the participants. The Governance Body defines awards or other  
3057 incentives.

#### 3058 **5.1.2.5 Considerations for Multiple Governance Chains**

3059 As noted in section 5.1.2, instances of the governance model often occur as a tiered  
3060 arrangement, with governance at some level delegating specific authority and  
3061 responsibility to accomplish a focused portion of the original level's mandate. For  
3062 example, a corporation may encompass several lines of business and each line of  
3063 business governs its own affairs in a manner that is consistent with and contributes to  
3064 the goals of the parent organization. Within the line of business, an IT group may be  
3065 given the mandate to provide and maintain IT resources, giving rise to IT governance.

3066 In addition to tiered governance, there may be multiple governance chains working in  
3067 parallel. For example, a company making widgets has policies intended to ensure they  
3068 make high quality widgets and make an impressive profit for their shareholders. On the  
3069 other hand, Sarbanes-Oxley is a parallel governance chain in the United States that  
3070 specifies how the management must handle its accounting and information that needs  
3071 to be given to its shareholders. The parallel chains may just be additive or may be in  
3072 conflict and require some harmonization.

3073 Being distributed and representing different ownership domains, a SOA participant falls  
3074 under the jurisdiction of multiple governance domains simultaneously and may  
3075 individually need to resolve consequent conflicts. The governance domains may  
3076 specify precedence for governance conformance or it may fall to the discretion of the  
3077 participant to decide on the course of actions they believe appropriate.

#### 3078 **5.1.3 Governance Applied to SOA**

##### 3079 **5.1.3.1 Where SOA Governance is Different**

3080 SOA governance is often discussed in terms of IT governance, but rather than a parent-  
3081 child relationship, Figure 41 shows the two as siblings of the general governance  
3082 described in section 5.1.2. There are obvious dependencies and a need for coordination  
3083 between the two, but the idea of aligning IT with business already demonstrates that  
3084 resource providers and resource consumers must be working towards common goals if  
3085 they are to be productive and efficient. While SOA governance is shown to be active in  
3086 the area of infrastructure, it is a specialized concern for having a dependable platform to  
3087 support service interaction; a range of traditional IT issues is therefore out of scope of  
3088 this document. A SOA governance plan for an enterprise will not of itself resolve  
3089 shortcomings with the enterprise's IT governance.

3090 Governance in the context of SOA is that organization of services: that promotes their  
3091 visibility; that facilitates interaction among service participants; and that directs that the  
3092 results of service interactions are those real world effects as described within the  
3093 service description and constrained by policies and contracts as assembled in the  
3094 execution context.



3095 SOA governance must specifically account for control across different ownership  
3096 domains, i.e. all the participants may not be under the jurisdiction of a single  
3097 governance authority. However, for governance to be effective, the participants must  
3098 agree to recognize the authority of the Governance Body and must operate within the  
3099 Governance Framework and through the Governance Processes so defined.

3100 SOA governance must account for interactions across ownership boundaries, which  
3101 may also imply across enterprise governance boundaries. For such situations,  
3102 governance emphasizes the need for agreement that some Governance Framework  
3103 and Governance Processes have jurisdiction, and the governance defined must satisfy  
3104 the Goals of the participants for cooperation to continue. A standards development  
3105 organization such as OASIS is an example of voluntary agreement to governance over  
3106 a limited domain to satisfy common goals.

3107 The specifics discussed in the figures in the previous sections are equally applicable to  
3108 governance across ownership boundaries as it is within a single boundary. There is a  
3109 charter agreed to when participants become members of the organization, and this  
3110 charter sets up the structures and processes that will be followed. Leadership may be  
3111 shared by the leadership of the overall organization and the leadership of individual  
3112 groups themselves chartered per the Governance Processes. There are  
3113 Rules/Regulations specific to individual efforts for which participants agree to local  
3114 goals, and Enforcement can be loss of voting rights or under extreme circumstances,  
3115 expulsion from the group.

3116 Thus, the major difference for SOA governance is an appreciation for the cooperative  
3117 nature of the enterprise and its reliance on furthering common goals if productive  
3118 participation is to continue.

### 3119 **5.1.3.2 What Must be Governed**

3120 An expected benefit of employing SOA principles is the ability to quickly bring resources  
3121 to bear to deal with unexpected and evolving situations. This requires a great deal of  
3122 confidence in the underlying capabilities that can be accessed and in the services that  
3123 enable the access. It also requires considerable flexibility in the ways these resources  
3124 can be employed. Thus, SOA governance requires establishing confidence and trust  
3125 while instituting a solid framework that enables flexibility, indicating a combination of  
3126 strict control over a limited set of foundational aspects but minimum constraints beyond  
3127 those bounds.

3128

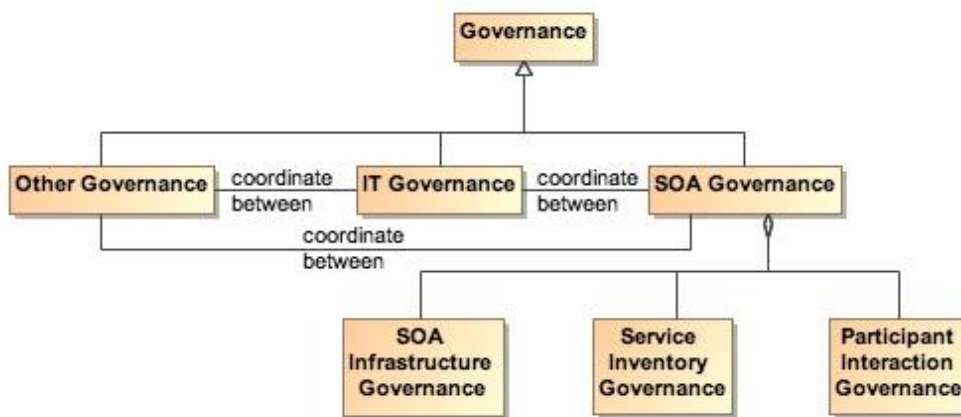


Figure 41 Relationship among types of governance

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

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

### 5.1.3.2.1 Governance of SOA Infrastructure

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

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

### 5.1.3.2.2 Governance of the Service Inventory

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

3160 characteristics from experiences with distributed computing but also evolve with SOA  
3161 experience. A major requirement for ensuring well-behaved services is collecting  
3162 sufficient metrics to know how the service affects the SOA infrastructure and whether it  
3163 complies with established infrastructure policies.

3164 Another common concern of service approval is whether there is a possibility of  
3165 duplication of function by multiple services. Some governance models talk to a tightly  
3166 controlled environment where a primary concern is to avoid any service duplication.  
3167 Other governance models talk to a market of services where the consumers have wide  
3168 choices. For the latter, it is anticipated that the better services will emerge from market  
3169 consensus and the availability of alternatives will drive innovation.

3170 Some combination of control and openness will emerge, possibly with a different  
3171 appropriate balance for different categories of use. For SOA governance, the issue is  
3172 less which services are approved but rather ensuring that sufficient description is  
3173 available to support informed decisions for appropriate use. Thus, SOA governance  
3174 SHOULD concentrate on identifying the required attributes to adequately describe a  
3175 service, the required target values of the attributes, and the standards for defining the  
3176 meaning of the attributes and their target values. Governance may also specify the  
3177 processes by which the attribute values are measured and the corresponding  
3178 certification that some realized attribute set may imply.

3179 For example, unlimited access for using a service may require a degree of life cycle  
3180 maturity that has demonstrated sufficient testing over a certain size community.  
3181 Alternately, the policy may specify that a service in an earlier phase of its life cycle may  
3182 be made available to a smaller, more technically sophisticated group in order to collect  
3183 the metrics that would eventually allow the service to advance its life cycle status.

3184 This aspect of governance is tightly connected to description because, given a well-  
3185 behaved set of services, it is the responsibility of the consumer (or policies promulgated  
3186 by the consumer's organization) to decide whether a service is sufficient for that  
3187 consumer's intended use. The goal is to avoid global governance specifying criteria that  
3188 are too restrictive or too lax for the local needs of which global governance has little  
3189 insight.

3190 Such an approach to specifying governance allows independent domains to describe  
3191 services in local terms while still having the services available for informed use across  
3192 domains. In addition, changes to the attribute sets within a domain can be similarly  
3193 described, thus supporting the use of newly described resources with the existing ones  
3194 without having to update the description of all the legacy content.

### 3195 **5.1.3.2.3 Governance of Participant Interaction**

3196 Finally, given a reliable services infrastructure and a predictable set of services, the  
3197 third aspect of governance is prescribing what is required during a service interaction.

3198 Governance would specify adherence to service interface and service reachability  
3199 parameters and would require that the result of an interaction MUST correspond to the  
3200 real world effects as contained in the service description. Governance would ensure  
3201 preconditions for service use are satisfied, in particular those related to security aspects  
3202 such as user authentication, authorization, and non-repudiation. If conflicts arise,

3203 governance would specify resolution processes to ensure appropriate agreements,  
3204 policies, and conditions are met.

3205 It would also rely on sufficient monitoring by the SOA infrastructure to ensure services  
3206 remain well-behaved during interactions, e.g. do not use excessive resources or exhibit  
3207 other prohibited behavior. Governance would also require that policy agreements as  
3208 documented in the execution context for the interaction are observed and that the  
3209 results and any after effects are consistent with the agreed policies. Governance will  
3210 focus on more contractual and legal aspects rather than the precursor descriptive  
3211 aspects. SOA governance may prescribe the processes by which SOA-specific policies  
3212 are allowed to change, but there are probably more business-specific policies that will  
3213 be governed by processes outside SOA governance.

### 3214 **5.1.3.3 Overarching Governance Concerns**

3215 There are numerous governance related concerns whose effects span the three areas  
3216 just discussed. One is the area of standards, how these are mandated, and how the  
3217 mandates may change. The Web Services standards stack is an example of relevant  
3218 standards where a significant number are still under development. In addition, while  
3219 there are notional scenarios that guide what standards are being developed, the fact  
3220 that many of these standards do not yet exist precludes operational testing of their  
3221 adequacy or effectiveness as a necessary and sufficient set.

3222 That said, standards are critical to creating a SOA ecosystem where SOA services can  
3223 be introduced, used singularly, and combined with other services to deliver complex  
3224 business functionality. As with other aspects of SOA governance, the Governance  
3225 Body should identify the minimum set felt to be needed and rigorously enforce that that  
3226 set be used where appropriate. The Governance Body must take care to expand and  
3227 evolve the mandated standards in a predictable manner and with sufficient technical  
3228 guidance that new services are able to coexist as much as possible with the old, and  
3229 changes to standards do not cause major disruptions.

3230 Another area that may see increasing activity as SOA expands is additional regulation  
3231 by governments and associated legal institutions. New laws are may deal with  
3232 transactions which are service based, possibly including taxes on the transactions.  
3233 Disclosures laws may mandate certain elements of description so both the consumer  
3234 and provider act in a predictable environment and are protected from ambiguity in intent  
3235 or action. Such laws are spawn rules and regulations that will influence the metrics  
3236 collected for evaluation of compliance.

### 3237 **5.1.3.4 Considerations for SOA Governance**

3238 The Reference Architecture definition of a loosely coupled system is one in which the  
3239 constraints on the interactions between components is minimal: sufficient to permit  
3240 interoperation without additional constraints that may be an artifact of implementation  
3241 technology. While governance experience for standalone systems provides useful  
3242 guides, we must be careful not to apply constraints that would preclude the flexibility,  
3243 agility, and adaptability we expect to realize from a SOA ecosystem.

3244 One of the strengths of SOA is it can make effective use of diversity rather than  
3245 requiring monolithic solutions. Heterogeneous organizations can interact without  
3246 requiring each conforms to uniform tools, representation, and processes. However, with

3247 this diversity comes the need to adequately define those elements necessary for  
3248 consistent interaction among systems and participants, such as which communication  
3249 protocol, what level of security, which vocabulary for payload content of messages. The  
3250 solution is not always to lock down these choices but to standardize alternatives and  
3251 standardize the representations through which an unambiguous identification of the  
3252 alternative chosen can be conveyed. For example, the URI standard specifies the URI  
3253 string, including what protocol is being used, what is the target of the message, and how  
3254 may parameters be attached. It does not limit the available protocols, the semantics of  
3255 the target address, or the parameters that can be transferred. Thus, as with our  
3256 definition of loose coupling, it provides absolute constraints but minimizes which  
3257 constraints it imposes.

3258 There is not a one-size-fits-all governance but a need to understand the types of things  
3259 governance is called upon to do in the context of the goals of SOA. Some communities  
3260 may initially desire and require very stringent governance policies and procedures while  
3261 other see need for very little. Over time, best practices will evolve, resulting in some  
3262 consensus on a sensible minimum and, except in extreme cases where it is  
3263 demonstrated to be necessary, a loosening of strict governance toward the best  
3264 practice mean.

3265 A question of how much governance may center on how much time governance  
3266 activities require versus how quickly is the system being governed expected to respond  
3267 to changing conditions. For large single systems that take years to develop, the  
3268 governance process could move slowly without having a serious negative impact. For  
3269 example, if something takes two years to develop and the steps involved in governance  
3270 take two months to navigate, then the governance can go along in parallel and may not  
3271 have a significant impact on system response to changes. Situations where it takes as  
3272 long to navigate governance requirements as it does to develop a response are  
3273 examples where governance may need to be reevaluated as to whether it facilitates or  
3274 inhibits the desired results. Thus, the speed at which services are expected to appear  
3275 and evolve needs to be considered when deciding the processes for control. The  
3276 added weight of governance should be appropriate for overall goals of the application  
3277 domain and the service environment.

3278 Governance, as with other aspects of any SOA implementation, should start small and  
3279 be conceptualized in a way that keeps it flexible, scalable, and realistic. A set of useful  
3280 guidelines would include:

- 3281 • Do not hardwire things that will inevitably change. For example, develop a  
3282 system that uses the representation of policies rather than code the policies into  
3283 the implementations.
- 3284 • Avoid setting up processes that demo well for three services without considering  
3285 how they may work for 300. Similarly, consider whether the display of status and  
3286 activity for a small number of services will also be effective for an operator in a  
3287 crisis situation looking at dozens of services, each with numerous, sometimes  
3288 overlapping and sometimes differing activities.
- 3289 • Maintain consistency and realism. A service solution responding to a natural  
3290 disaster cannot be expected to complete a 6-week review cycle but be effective  
3291 in a matter of hours.

#### 3292 5.1.4 Architectural Implications of SOA Governance

3293 The description of SOA governance indicates numerous architectural requirements on  
3294 the SOA ecosystem:

- 3295 • Governance is expressed through policies and assumes multiple use of focused  
3296 policy modules that can be employed across many common circumstances. This  
3297 requires the existence of:
  - 3298 ○ descriptions to enable the policy modules to be visible, where the  
3299 description includes a unique identifier for the policy and a sufficient, and  
3300 preferably a machine process-able, representation of the meaning of  
3301 terms used to describe the policy, its functions, and its effects;
  - 3302 ○ one or more discovery mechanisms that enable searching for policies that  
3303 best meet the search criteria specified by the service participant; where  
3304 the discovery mechanism will have access to the individual policy  
3305 descriptions, possibly through some repository mechanism;
  - 3306 ○ accessible storage of policies and policy descriptions, so service  
3307 participants can access, examine, and use the policies as defined.
- 3308 • Governance requires that the participants understand the intent of governance,  
3309 the structures created to define and implement governance, and the processes to  
3310 be followed to make governance operational. This requires the existence of:
  - 3311 ○ an information collection site, such as a Web page or portal, where  
3312 governance information is stored and from which the information is always  
3313 available for access;
  - 3314 ○ a mechanism to inform participants of significant governance events, such  
3315 as changes in policies, rules, or regulations;
  - 3316 ○ accessible storage of the specifics of Governance Processes;
  - 3317 ○ SOA services to access automated implementations of the Governance  
3318 Processes
- 3319 • Governance policies are made operational through rules and regulations. This  
3320 requires the existence of:
  - 3321 ○ descriptions to enable the rules and regulations to be visible, where the  
3322 description includes a unique identifier and a sufficient, and preferably a  
3323 machine process-able, representation of the meaning of terms used to  
3324 describe the rules and regulations;
  - 3325 ○ one or more discovery mechanisms that enable searching for rules and  
3326 regulations that may apply to situations corresponding to the search  
3327 criteria specified by the service participant; where the discovery  
3328 mechanism will have access to the individual descriptions of rules and  
3329 regulations, possibly through some repository mechanism;
  - 3330 ○ accessible storage of rules and regulations and their respective  
3331 descriptions, so service participants can understand and prepare for  
3332 compliance, as defined.
  - 3333 ○ SOA services to access automated implementations of the Governance  
3334 Processes.

- 3335
- 3336
- 3337
- 3338
- 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.

## 3350 5.2 Security Model

3351 Security is one aspect of confidence – the confidence in the integrity, reliability, and  
3352 confidentiality of the system. In particular, security focuses on those aspects of  
3353 assurance that involve the accidental or malign intent of other people to damage or  
3354 compromise trust in the system and on the availability of SOA-based systems to  
3355 perform desired capability.

### 3356 Security

3357 Security concerns the set of mechanisms for ensuring and enhancing trust and  
3358 confidence in the SOA ecosystem.

3359 Providing for security for Service Oriented Architecture is somewhat different than for  
3360 other contexts; although many of the same principles apply equally to SOA and to other  
3361 systems. The fact that SOA embraces crossing ownership boundaries makes the issues  
3362 involved with moving data more visible.

3363 As well as securing the movement of data within and across ownership boundaries,  
3364 security often revolves around resources: the need to guard certain resources against  
3365 inappropriate access – whether reading, writing or otherwise manipulating those  
3366 resources.

3367 Any comprehensive security solution must take into account the people that are using,  
3368 maintaining and managing the SOA. Furthermore, the relationships between them must  
3369 also be incorporated: any security assertions that may be associated with particular  
3370 interactions originate in the people that are behind the interaction.

3371 We analyze security in terms of the social structures that define the legitimate  
3372 permissions, obligations and roles of people in relation to the system, and mechanisms  
3373 that must be put into place to realize a secure system. The former are typically captured  
3374 in a series of security policy statements; the latter in terms of security *guards* that  
3375 ensure that policies are enforced.

3376 How and when to apply these derived security policy mechanisms is directly associated  
3377 with the assessment of the *threat model* and a *security response model*. The threat  
3378 model identifies the kinds of threats that directly impact the message and/or application  
3379 of constraints, and the response model is the proposed mitigation to those threats.  
3380 Properly implemented, the result can be an acceptable level of risk to the safety and  
3381 integrity of the system.

## 3382 **5.2.1 Secure Interaction Concepts**

3383 We can characterize secure interactions in terms of key security concepts [ISO/IEC  
3384 **27002**]: confidentiality, integrity, authentication, authorization, non-repudiation, and  
3385 availability. The concepts for secure interactions are well defined in other standards  
3386 and publications. The security concepts here are not defined but rather related to the  
3387 SOA ecosystem perspective of the SOA-RAF.

### 3388 **5.2.1.1 Confidentiality**

3389 Confidentiality concerns the protection of privacy of participants in their interactions.  
3390 Confidentiality refers to the assurance that unauthorized entities are not able to read  
3391 messages or parts of messages that are transmitted.

3392 Note that confidentiality has degrees: in a completely confidential exchange, third  
3393 parties would not even be aware that a confidential exchange has occurred. In a  
3394 partially confidential exchange, the identities of the participants may be known but the  
3395 content of the exchange obscured.

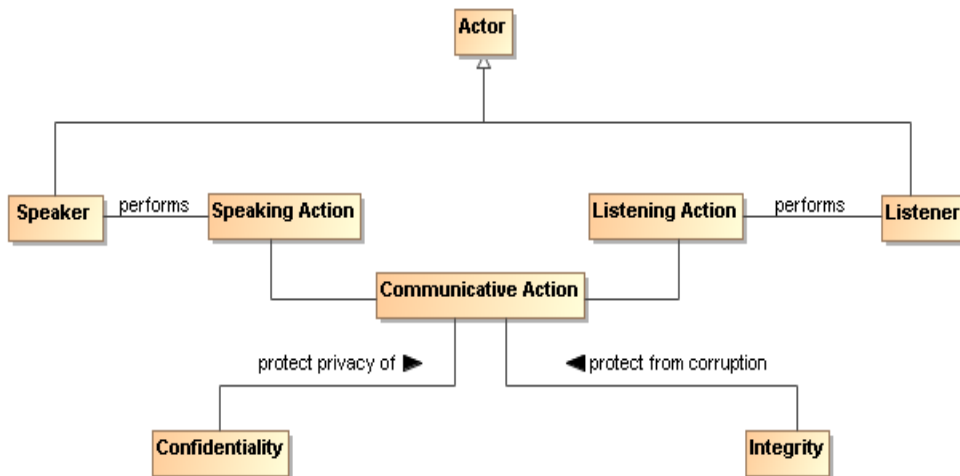
### 3396 **5.2.1.2 Integrity**

3397 Integrity concerns the protection of information that is exchanged – either from  
3398 unauthorized writing or inadvertent corruption. Integrity refers to the assurance that  
3399 information that has been exchanged has not been altered.

3400 Integrity is different from confidentiality in that messages that are sent from one  
3401 participant to another may be obscured to a third party, but the third party may still be  
3402 able to introduce his own content into the exchange without the knowledge of the  
3403 participants.

3404 Figure 42 applies confidentiality and integrity to communicative action.






3405  
3406 *Figure 42 Confidentiality and Integrity*

3407 A communicative action is a joint action involved in the exchange of messages. Section  
3408 5.2.4 describes common computing techniques for providing confidentiality and integrity  
3409 during message exchanges.

### 3410 5.2.1.3 Authentication

3411 Authentication concerns the identity of the participants in an exchange. Authentication  
3412 refers to the means by which one participant can be assured of the identity of other  
3413 participants.

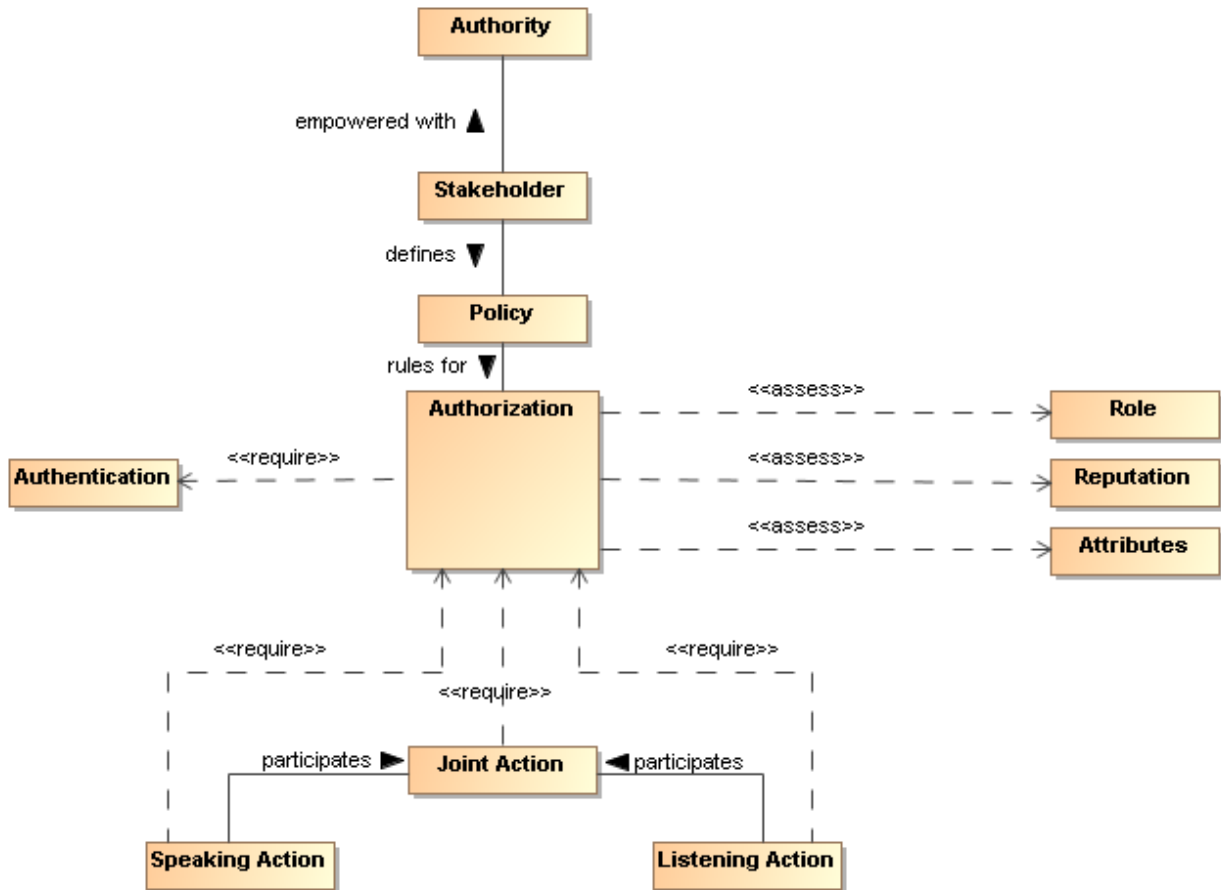
3414 Figure 43 applies authentication to the identity of participants.



3416  
3417 *Figure 43 Authentication*

### 3418 5.2.1.4 Authorization

3419 Authorization concerns the legitimacy of the interaction. Authorization refers to the  
3420 means by which a stakeholder may be assured that the information and actions that are  
3421 exchanged are either explicitly or implicitly approved.



3422  
3423 *Figure 44 Authorization*

3424 The roles and attributes which provide a participant's credentials are expanded to  
3425 include reputation. Reputation often helps determine willingness to interact, for  
3426 example, reviews of a service provider will influence the decision to interact with the  
3427 service provider. The roles, reputation, and attributes are represented as assertions  
3428 measured by authorization decision points.

3429 The role of policy for security is to permit stakeholders to express their choices. In  
3430 Figure 44, a policy is a written constraint and the role, reputation, and attribute  
3431 assertions are evaluated according to the constraints in the authorization policy. A  
3432 combination of security mechanisms and their control via explicit policies can form the  
3433 basis of an authorization solution.

### 3434 **5.2.1.5 Non-repudiation**

3435 Non-repudiation concerns the accountability of participants. To foster trust in the  
3436 performance of a system used to conduct shared activities it is important that the  
3437 participants are not able to later deny their actions: to repudiate them. Non-repudiation  
3438 refers to the means by which a participant may not, at a later time, successfully deny  
3439 having participated in the interaction or having performed the actions as reported by  
3440 other participants.

### 3441 **5.2.1.6 Availability**

3442 Availability concerns the ability of systems to use and offer the services for which they  
3443 were designed. One of the threats against availability is the so-called denial of service  
3444 attack in which attackers attempt to prevent legitimate access to the system.

3445 We differentiate here between general availability – which includes aspects such as  
3446 systems reliability – and availability as a security concept where we need to respond to  
3447 active threats to the system.

### 3448 **5.2.2 Where SOA Security is Different**

3449 The core security concepts are fundamental to all social interactions. The evolution of  
3450 sharing information using a SOA requires the flexibility to dynamically secure computing  
3451 interactions in a computing ecosystem where the owning social groups, roles, and  
3452 authority are constantly changing as described in section 5.1.3.1.

3453 SOA policy-based security can be more adaptive for a computing ecosystem than  
3454 previous computing technologies allow for, and typically involves a greater degree of  
3455 distributed mechanisms.

3456 Standards for security, as is the case with all aspects of SOA, play a large role in  
3457 flexible security on a global scale. SOA security may also involve greater auditing and  
3458 reporting to adhere to regulatory compliance established by governance structures.

### 3459 **5.2.3 Security Threats**

3460 There are a number of ways in which an attacker may attempt to compromise the  
3461 security of a system. The two primary sources of attack are third parties attempting to  
3462 subvert interactions between legitimate participants and an entity that is participating but  
3463 attempting to subvert its partner(s). The latter is particularly important in a SOA where  
3464 there may be multiple ownership boundaries and trust boundaries.

3465 The threat model lists some common threats that relate to the core security concepts  
3466 listed in Section 5.2.1. Each technology choice in the realization of a SOA can  
3467 potentially have many threats to consider.

#### 3468 **Message alteration**

3469 If an attacker is able to modify the content (or even the order) of messages that  
3470 are exchanged without the legitimate participants being aware of it then the  
3471 attacker has successfully compromised the security of the system. In effect, the  
3472 participants may unwittingly serve the needs of the attacker rather than their own.

3473 An attacker may not need to completely replace a message with his own to  
3474 achieve his objective: replacing the identity of the beneficiary of a transaction  
3475 may be enough.

#### 3476 **Message interception**

3477 If an attacker is able to intercept and understand messages exchanged between  
3478 participants, then the attacker may be able to gain advantage. This is probably  
3479 the most commonly understood security threat.

3480 **Man in the middle**

3481 In a man-in-the-middle attack, the legitimate participants believe that they are  
3482 interacting with each other; but are in fact interacting with the attacker. The  
3483 attacker attempts to convince each participant that he is their correspondent;  
3484 whereas in fact he is not.

3485 In a successful man-in-the-middle attack, legitimate participants do not have  
3486 an inaccurate understanding of the state of the other participants. The attacker can  
3487 use this to subvert the intentions of the participants.

3488 **Spoofing**

3489 In a spoofing attack, the attacker convinces a participant that he is really  
3490 someone else – someone that the participant would normally trust.

3491 **Denial of service attack**

3492 In a denial of service (DoS) attack, the attacker attempts to prevent legitimate  
3493 users from making use of the service. A DoS attack is easy to mount and can  
3494 cause considerable harm: by preventing legitimate interactions, or by slowing  
3495 them down enough, the attacker may be able to simultaneously prevent  
3496 legitimate access to a service and to attack the service by another means.

3497 A variation of the DoS attack is the Distributed Denial of Service attack. In a  
3498 DDoS attack the attacker uses multiple agents to the attack the target. In some  
3499 circumstances this can be extremely difficult to counteract effectively.

3500 One of the features of a DoS attack is that it does not require valid interactions to  
3501 be effective: responding to invalid messages also takes resources and that may  
3502 be sufficient to cripple the target.

3503 **Replay attack**

3504 In a replay attack, the attacker captures the message traffic during a legitimate  
3505 interaction and then replays part of it to the target. The target is persuaded that a  
3506 similar transaction to the previous one is being repeated and it responds as  
3507 though it were a legitimate interaction.

3508 A replay attack may not require that the attacker understand any of the individual  
3509 communications; the attacker may have different objectives (for example  
3510 attempting to predict how the target would react to a particular request).

3511 **False repudiation**

3512 In false repudiation, a user completes a normal transaction and then later  
3513 attempts to deny that the transaction occurred. For example, a customer may  
3514 use a service to buy a book using a credit card; then, when the book is delivered,  
3515 refuse to pay the credit card bill claiming that *someone else* must have ordered  
3516 the book.

3517 **5.2.4 Security Responses**

3518 Security goals are never absolute: it is not possible to guarantee 100% confidentiality,  
3519 non-repudiation, etc. However, a well designed and implemented security response  
3520 model can ensure acceptable levels of security risk. For example, using a well-designed

3521 cipher to encrypt messages may make the cost of breaking communications so great  
3522 and so lengthy that the information obtained is valueless.

3523 Performing threat assessments, devising mitigation strategies, and determining  
3524 acceptable levels of risk are the foundation for an effective process to mitigating threats  
3525 in a cost-effective way.<sup>14</sup> The choice in hardware and software to realize a SOA will be  
3526 the basis for threat assessments and mitigation strategies. The stakeholders of a  
3527 specific SOA implementation should determine acceptable levels of risk based on threat  
3528 assessments and the cost of mitigating those threats.

#### 3529 **5.2.4.1 Privacy Enforcement**

3530 The most efficient mechanism to assure confidentiality is the encryption of information.  
3531 Encryption is particularly important when messages must cross trust boundaries;  
3532 especially over the Internet. Note that encryption need not be limited to the content of  
3533 messages: it is possible to obscure even the existence of messages themselves  
3534 through encryption and 'white noise' generation in the communications channel.

3535 The specifics of encryption are beyond the scope of this architecture. However, we are  
3536 concerned about how the connection between privacy-related policies and their  
3537 enforcement is made.

3538 A policy enforcement point for enforcing privacy may take the form of an automatic  
3539 function to encrypt messages as they leave a trust boundary; or perhaps simply  
3540 ensuring that such messages are suitably encrypted.

3541 Any policies relating to the level of encryption being used would then apply to these  
3542 centralized messaging functions.

#### 3543 **5.2.4.2 Integrity Protection**

3544 To protect against message tampering or inadvertent message alteration, and to allow  
3545 the receiver of a message to authenticate the sender, messages may be accompanied  
3546 by a digital signature. Digital signatures provide a means to detect if signed data has  
3547 been altered. This protection can also extend to authentication and non-repudiation of a  
3548 sender.

3549 A common way a digital signature is generated is with the use of a private key that is  
3550 associated with a public key and a digital certificate. The private key of some entity in  
3551 the system is used to create a digital signature for some set of data. Other entities in the  
3552 system can check the integrity of the signed data set via signature verification  
3553 algorithms. Any changes to the data that was signed will cause signature verification to  
3554 fail, which indicates that integrity of the data set has been compromised.

3555 A party verifying a digital signature must have access to the public key that corresponds  
3556 to the private key used to generate the signature. A digital certificate contains the public

---

<sup>14</sup> 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.

3557 key of the owner, and is itself protected by a digital signature created using the private  
3558 key of the issuing Certificate Authority (CA).

#### 3559 **5.2.4.3 Message Replay Protection**

3560 To protect against replay attacks, messages may contain information that can be used  
3561 to detect replayed messages. The simplest requirement to prevent replay attacks is that  
3562 each message that is ever sent is unique. For example, a message may contain a  
3563 message ID, a timestamp, and the intended destination.

3564 By storing message IDs, and comparing each new message with the store, it becomes  
3565 possible to verify whether a given message has been received before (and therefore  
3566 should be discarded).

3567 The timestamp may be included in the message to help check for message freshness.  
3568 Messages that arrive after their message ID could have been cleared (after receiving  
3569 the same message some time previously) may also have been replayed. A common  
3570 means for representing timestamps is a useful part of an interoperable replay detection  
3571 mechanism.

3572 The destination information is used to determine if the message was misdirected or  
3573 replayed. If the replayed message is sent to a different endpoint than the destination of  
3574 the original message, the replay could go undetected if the message does not contain  
3575 information about the intended destination.

3576 In the case of messages that are replies to prior messages, it is also possible to include  
3577 seed information in the prior messages that is randomly and uniquely generated for  
3578 each message that is sent out. A replay attack can then be detected if the reply does  
3579 not embed the random number that corresponds to the original message.

#### 3580 **5.2.4.4 Auditing and Logging**

3581 False repudiation involves a participant denying that it authorized a previous interaction.  
3582 An effective strategy for responding to such a denial is to maintain careful and complete  
3583 logs of interactions which can be used for auditing purposes. The more detailed and  
3584 comprehensive an audit trail is, the less likely it is that a false repudiation would be  
3585 successful.

3586 The countermeasures assume that the non-repudiation tactic (e.g. digital signatures) is  
3587 not undermined itself. For example, if private key is stolen and used by an adversary,  
3588 even extensive logging cannot assist in rejecting a false repudiation.

3589 Unlike many of the security responses discussed here, it is likely that the scope for  
3590 automation in rejecting a repudiation attempt is limited to careful logging.

#### 3591 **5.2.4.5 Graduated engagement**

3592 The key to managing and responding to DoS attacks is to be careful in the use of  
3593 resources when responding to interaction. Put simply, a system has a choice to respond  
3594 to a communication or to ignore it. In order to avoid vulnerability to DoS attacks a  
3595 service provider should be careful not to commit resources beyond those implied by the  
3596 current state of interactions; this permits a graduation in commitment by the service  
3597 provider that mirrors any commitment on the part of service consumers and attackers  
3598 alike.

## 3599 **5.2.5 Architectural Implications of SOA Security**

3600 Providing SOA security in an ecosystem of governed services has the following  
3601 implications on the policy support and the distributed nature of mechanisms used to  
3602 assure SOA security:

- 3603 • Security expressed through policies have the same architectural implications as  
3604 described in Section 4.4.3 for policies and contracts architectural implications.
- 3605 • Security policies require mechanisms to support security description  
3606 administration, storage, and distribution.
- 3607 • Service descriptions supporting security policies should:
  - 3608 ○ have a meta-structure sufficiently rich to support security policies;
  - 3609 ○ be able to reference one or more security policy artifacts;
  - 3610 ○ have a framework for resolving conflicts between security policies.
- 3611 • The mechanisms that make-up the execution context in secure SOA-based  
3612 systems should:
  - 3613 ○ provide protection of the confidentiality and integrity of message  
3614 exchanges;
  - 3615 ○ be distributed so as to provide centralized or decentralized policy-based  
3616 identification, authentication, and authorization;
  - 3617 ○ ensure service availability to consumers;
  - 3618 ○ be able to scale to support security for a growing ecosystem of services;
  - 3619 ○ be able to support security between different communication technologies;
- 3620 • Common security services include:
  - 3621 ○ services that abstract encryption techniques;
  - 3622 ○ services for auditing and logging interactions and security violations;
  - 3623 ○ services for identification;
  - 3624 ○ services for authentication;
  - 3625 ○ services for authorization;
  - 3626 ○ services for intrusion detection and prevention;
  - 3627 ○ services for availability including support for quality of service  
3628 specifications and metrics.

## 3629 **5.3 Management Model**

### 3630 **Management**

3631 Management is a process of controlling resources in accordance with the policies and principles defined  
3632 by Governance.

3633 There are three separate but linked domains of interest within the management of SOA:

- 3634 1. the management and support of the resources that are involved in any complex structures – of  
3635 which SOA-based solutions are excellent examples;
- 3636 2. the promulgation and enforcement of the policies and service contracts agreed to by the  
3637 stakeholders in SOA ecosystem;

3638 3. the management of the relationships of the participants in SOA-based solutions – both to each  
3639 other and to the services that they use and offer.

3640 There are many artifacts related to management. Historically, systems management capabilities have  
3641 been organized by the “FCAPS” functions (based on ITU-T Rec. M.3400 (02/2000), “TMN Management  
3642 Functions”):

- 3643 • fault management,
- 3644 • configuration management,
- 3645 • account management,
- 3646 • performance and security management.

3647 The primary task of the functional groups is to concentrate on maintaining systems in a trusted, active,  
3648 and accessible state.

3649 In the context of the SOA ecosystem, we see many possible resources that may require management  
3650 such as services, service descriptions, service contracts, policies, roles, relationships, security, people  
3651 and systems that implement services and infrastructure elements. In addition, given the ecosystem  
3652 nature, it is also potentially necessary to manage the business relationships between participants.

3653 Successful operation of a SOA ecosystem requires trust between the stakeholders and the ecosystem  
3654 elements. In contrast, regular systems in technology are not necessarily operated or used in an  
3655 environment requiring trust before the stakeholders make use of the system. Indeed, many of these  
3656 systems exist in hierarchical management structures, within which use may be mandated by legal  
3657 requirement, executive decision, or good business practice in furthering the business’ strategy. Pre-  
3658 condition of trust in the SOA ecosystem roots in both principles of service orientation and distributed  
3659 authoritative ownership of independent services. Even for hierarchical management structures applied  
3660 to a SOA ecosystem, the service use should have contractual basis rather than being mandated.

3661 The trust may be established through agreements/contracts, policies, or implicitly through observation of  
3662 repeated interactions with others. Explicit trust is usually accompanied by formalized documents suitable for  
3663 the management activities. Implicit trust adds fragility to the management of a SOA ecosystem because  
3664 failure to maintain consistent and predictable interactions will undermine the trust between participants  
3665 and within the ecosystem as a whole.

3666 Management in a SOA ecosystem is thus concerned with management taking actions that will establish  
3667 the condition of trust that must be present before engaging in service interactions. These concerns should  
3668 largely be handled within the governance of the ecosystem. The policies, agreements, and practices  
3669 defined through the governance provide the boundaries within which management operates and for which  
3670 management must provide enforcement and feedback. However, governance alone cannot anticipate all  
3671 circumstances and must offer sufficient guidance in areas where anticipation is unclear or for which  
3672 agreement between all stakeholders cannot be reached. Management in these cases must be flexible  
3673 and adaptable to handle unanticipated conditions without unnecessarily breaking trust relationships.

3674 Service management is the process – manual, automated, or a combination – of proactively monitoring  
3675 and controlling the behavior of a service or a set of services. Service management operates under  
3676 constraints attributed to the business and social context. Particularly, special policies may be used for  
3677 governing cross-boundary relationships. Managing solutions that may be used across ownership  
3678 boundaries based on such policies raises issues that are not typically present when managing a service  
3679 within a single ownership domain. For example, care is required in managing a service when the owner of  
3680 the service, the provider of the service, the host of the service and mediators to the service may all  
3681 belong to different stakeholders.

3682 Cross-boundary service management takes place in, at least, the following situations:

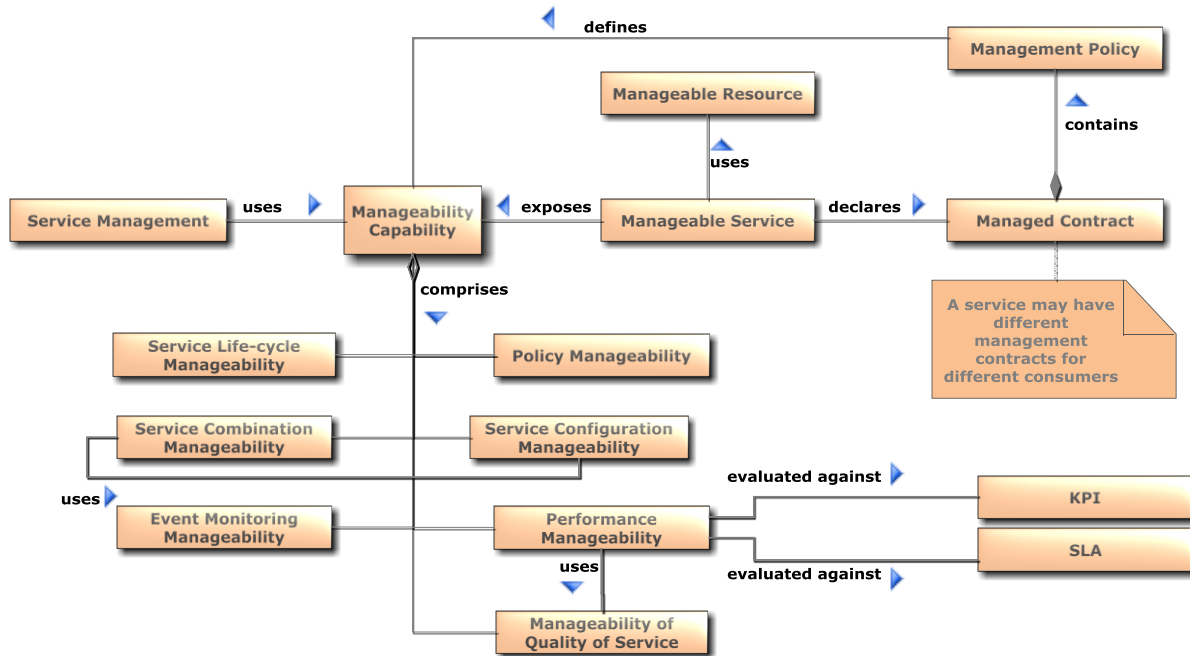
- 3683 • using combinations of services that belong to different ownership realms
- 3684 • using of services that mediate between ownership realms
- 3685 • sharing monitoring and reporting means and results.

3686 These situations are particularly important in ecosystems that are highly decentralized, in which the  
3687 participants interact as peers as well as in the “master-servant” mode.

3688 The management model shown in Figure 46 conveys how the SOA framework applies to managing  
3689 services. Services management operates via service metadata, such as service lifecycles and attributes



3690 associated with service use, that are typically collected in or accessed through the service description.  
 3691 [this Figure to be re-drawn in common style]



3692  
 3693 Figure 46 Manageability capabilities in SOA ecosystem  
 3694

3695 The service metadata of interest is that set of service properties that is manageable. These manageability  
 3696 properties are generally identifiable for any service consumed or supplied within the ecosystem. The  
 3697 necessary existence of these properties within the SOA ecosystem motivates the following definitions:

3698 **Manageability** of a resource is the capability that allows it to be controlled, monitored, and  
 3699 reported on with respect to some property. Note that manageability is not necessarily a part of the  
 3700 managed entities themselves and are generally considered to be external to the managed  
 3701 entities.

3702 Each resource may be managed through a number of aspects of management, and the resources may  
 3703 be grouped to categories based on similarity of managed aspects. For example, the managed aspect  
 3704 relating to configuration manageability is referred to as “Configuration Manageability” for the collection of  
 3705 services. Resources not managed under a particular capability are resources, for which those  
 3706 manageability aspects have no clear meaning or use. As an example, all resources within a SOA  
 3707 ecosystem have a lifecycle that is meaningful within the ecosystem. Thus, all resources are manageable  
 3708 under Lifecycle Manageability. In contrast, not all resources report or handle events. Thus, Event  
 3709 Manageability is only concerned with those resources for which events are meaningful.

3710 **Life-cycle Manageability** of a service typically refers to how the service is created, how it is  
 3711 destroyed and how service versions must be managed. This manageability is the feature of the SOA  
 3712 ecosystem because the service cannot manage its own life cycle.

3713 Another important consideration is that services may have resource requirements that must be  
 3714 established at various points in the services’ life cycles. However actual providers of these resources  
 3715 maybe not known at the time of the service creation and, thus, have to be managed at the service run-  
 3716 time.

3717 **Combination Manageability** of a service addresses management of service  
 3718 characteristics that allow for creating and changing of combinations in which the service  
 3719 participates or that the service combines by itself. Known models of such combinations

3720 are aggregations and compositions. Examples of patterns of combinations are  
3721 choreography and orchestration. Combination Manageability drives implementation of  
3722 the Service Composability Principle of service orientation.

3723  
3724 Service combination manageability resonates with the methodology of process management.  
3725 Combination Manageability may be applied at different phases of the service creation and execution and,  
3726 in some cases, can utilize Configuration Manageability.

3727  
3728 Service combinations contribute the most in delivering business values to the stakeholders and managing  
3729 service combinations is the one of the top-level tasks and features of the SOA ecosystem.

3730  
3731 **Configuration Manageability** of a service allows managing the identity of and the interactions  
3732 among internal elements of the service. Also, Configuration Manageability correlates with the  
3733 management of service versions and configuration of the deployment of new services into the ecosystem.  
3734 Configuration Management differs from the **Combination Manageability** in the scope and scale  
3735 of manageability, and addresses lower level concerns than the architectural  
3736 combination of services.

3737  
3738 **Event Monitoring Manageability** allows managing the categories of events of interest related to  
3739 services and reporting recognized events to the interested stakeholders. Such events may be the ones  
3740 that trigger service invocations as well as execution of particular functionality provided by the service.  
3741 This is one of the key lower-level manageability aspects that the service provider and associated  
3742 stakeholders are primarily interested. Monitored events may be internal or external to the SOA  
3743 ecosystem. For example, a disaster in the oil producing industry, which is outside of the SOA ecosystem  
3744 of the Insurer, can trigger the service's functionality that is responsible for immediate or constant  
3745 monitoring of the oil prices in the oil trading exchanges and, respectively, modify the premium paid by the  
3746 insured oil companies.

3747  
3748 **Performance Manageability** of a service allows controlling the service results, shared and  
3749 sharable real world effects against the business goals and objectives of the service. This manageability  
3750 assumes monitoring of the business performance as well as the management of this monitoring itself.  
3751 Performance Manageability includes business and technical performance manageability means through  
3752 performance criteria set, such as business key performance indicators (KPI) and service-level  
3753 agreements (SLA).

3754  
3755 The performance business- and technical-level characteristics of the service should be known from the  
3756 service contract. The service provider and consumer must be able to monitor and measure these  
3757 characteristics or be informed about the results measured by a third party.

3758  
3759 Performance Manageability is the instrument for providing compliance of the service with its service  
3760 contracts. Performance Manageability utilizes Manageability of Quality of Service.

3761  
3762 **Manageability of Quality of Service** deals with management of service non-functional  
3763 characteristics that may be of significant value to the service consumers and other stakeholders in the  
3764 SOA ecosystem. Classic examples of this include bandwidth offerings associated with a service.  
3765  
3766 Manageability of quality of service assumes that the properties associated with service qualities are  
3767 monitored during the service execution. Results of monitoring may be challenged against SLA and even  
3768 KPI, which results in the continuous validation of how the service contract is preserved by the service  
3769 provider.

3770

3771 **Policy Manageability** allows additions, changes and replacements of the policies associated  
 3772 with a resource in the SOA ecosystem. The ability to manage those policies (such as promulgating  
 3773 policies, retiring policies and ensuring that policy decision points and enforcement points are current)  
 3774 enables the ecosystem to *apply* policies and *evaluate* the results.

3775 Capability to manage, i.e. use particular manageability, requires policies from governance to be translated  
 3776 into the details of rules and regulations and then corresponding measurement and feedback on the  
 3777 specifics.

3778 In the following sub-sections, we describe how the elements of the SOA ecosystem may be managed  
 3779 with integrity.

### 3780 5.3.1. Management Means and Relationships

3781 A minimal set of management for the SOA ecosystem is shown on Figure 47 and elaborated in the  
 3782 following sections.

#### 3783 5.3.1.1. Management Policy

3784 The management of resources within the SOA may be governed by management policies.

3785 In a deployed SOA-based solution, it may well be that different aspects of the management of a  
 3786 given service are managed by different management services. For example, the life-cycle management of  
 3787 services often involves managing service versions. Managing quality of service is often very specific to  
 3788 the service itself; for example, quality of service attributes for a video streaming service are quite different  
 3789 to those for a banking system.

3790 There are additional concepts of management that also apply to IT management:

#### 3791 5.3.1.2. Network Management

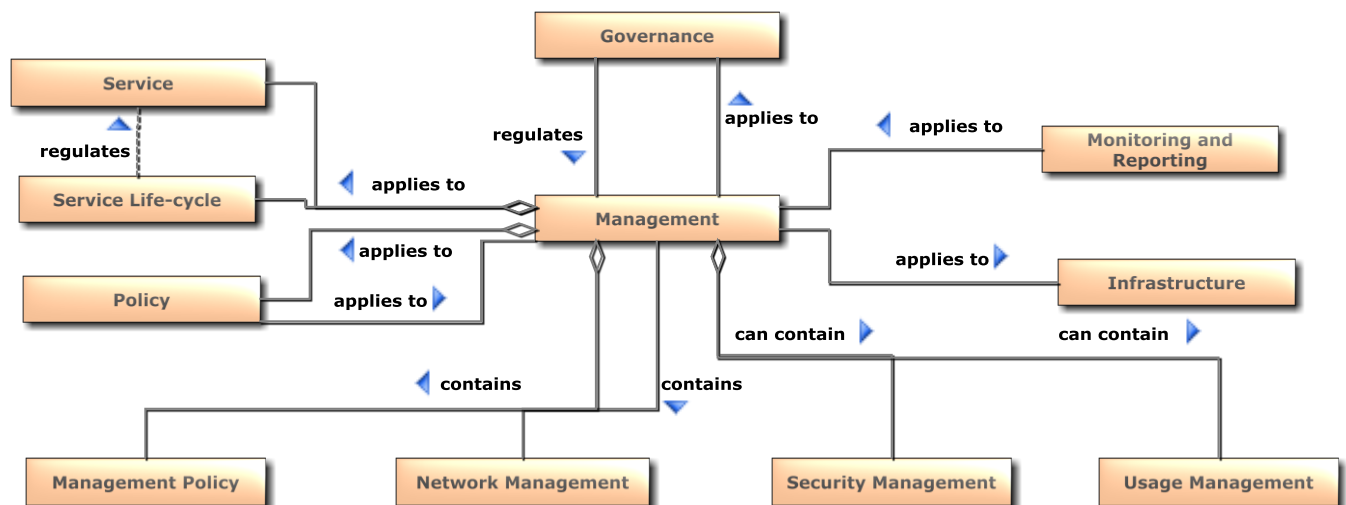
3792 Network management deals with the maintenance and administration of large scale physical networks  
 3793 such as computer networks and telecommunication networks. Specifics of the networks may affect  
 3794 service interactions from performance and operational perspectives.

3795 Network and related system management executes a set of functions required for controlling, planning,  
 3796 deploying, coordinating, and monitoring the distributed services in the SOA ecosystem. However, while  
 3797 recognizing their importance, the specifics of systems management or network management are out of  
 3798 scope for this Reference Architecture Foundation.

3799

3800 *[this Figure to be re-drawn in common style]*

3801



3802

3803

3804 *Figure 47 Management Means and Relationships in SOA ecosystem*

3805

### 3806 **5.3.1.3. Security Management**

3807 Management of the security related to resources includes identification of roles, permissions, access  
3808 rights, and policy attributes defining security boundaries and events that may trigger a security response.

3809 Security management within a SOA ecosystem is essential to maintaining the trust relationships between  
3810 participants residing in different ownership domains. Security management must consider not just the  
3811 internal properties related to interactions between participants but ecosystem properties that preserve the  
3812 integrity of the ecosystem from external threats.

### 3813 **5.3.1.4. Usage Management**

3814 Usage Management applies to management of the use of resources. Usage management includes  
3815 access properties, demand properties, and financial properties. Access properties include how the  
3816 resource is accessed, who is using the resource, and the state of the resource after use. Demand  
3817 properties are concerned with controlling or shaping demand for resources to optimize the overall  
3818 operation of the ecosystem. Financial properties are those associated with assigning costs to the use of  
3819 resources and distributing those cost assignments to the participants in an equitable manner.

## 3820 **5.3.2. Management and Governance**

3821 The primary role of governance in the context of a SOA ecosystem is to foster an atmosphere of  
3822 predictability, trust, and efficiency, and it accomplishes this by allowing the stakeholders to negotiate and  
3823 set the key policies that govern the running of the SOA-based solution. Recall that in an ecosystems  
3824 perspective, the goal of governance is less to have complete fine-grained control but more to enable the  
3825 individual participants to work together.

3826 Policies for a SOA ecosystem will tend to focus on the rules of engagement between participants; for  
3827 example, what kind of interactions are permissible, how will disputes be resolved, and so on. While  
3828 governance may primarily focus on setting policies, management will focus on the realization and  
3829 enforcement of policies. Effective management in the SOA ecosystem requires an ability for governance  
3830 to understand the consequences of its policies, guidelines, and principles, and to adjust those as needed  
3831 when inconsistencies or ambiguity become evident from the operation of the management functions. This  
3832 understanding and adjustment must be facilitated by the results of management and so the mechanisms  
3833 for providing feedback from management into governance must exist.

3834 Governance operates via specialized activities and, thus, should be managed itself. Management to  
3835 operationalize governance utilizes management policies that are included in the Governance Framework  
3836 and Processes, and driven by the enterprise business model, business objectives and strategies. Where  
3837 corporate management policies exist, these are usually guided and directed by the corporate executives.  
3838 In peer relationships, the governing policies are set by either an external entity and accepted by the peers  
3839 or by the peers themselves. This creates the appropriate authoritative level for the policies used for the  
3840 management of the Governance Framework and Processes. Management to operationalize governance  
3841 controls the life-cycle of the governing policies, including procedures and processes, for modifying the  
3842 Governance Framework and Processes.

## 3843 **5.3.3. Management and Contracts**

### 3844 **5.3.3.1 Management for Contracts and Policies**

3845 As we noted above, management can often be viewed as the application of contracts and individual  
3846 policies to ensure the smooth running of the SOA ecosystem. Policies play an important role as the  
3847 guiding constraints for management, as well as artifacts that need to be managed themselves. Service  
3848 contracts also serve as both guiding constraints and artifacts that need to be managed. Policies and  
3849 service contracts specify the service characteristics that have to be monitored, analysed and managed.

### 3850 **5.3.3.2 Contracts**

3851 As described in sections “Participation in a SOA Ecosystem view” and “Realization of a SOA Ecosystem  
3852 view”, there are several types of contractual information in the SOA ecosystem. From the management  
3853 perspective, three basic types of the contractual information relate to:

3854 · relationship between service provider and consumer;

3855 · communication with the service;

3856 · control of the quality of the service execution.

3857 When a consumer prepares to interact with a service, the consumer and the service provider must come  
3858 to agreement on service features and characteristics that will be provided by the service and available to  
3859 the consumer; this agreement is known as a **service contract**.

#### 3860 **Service Contract**

3861 An implicit or an explicit and documented agreement between the service consumer and service  
3862 provider about the use of the service based on

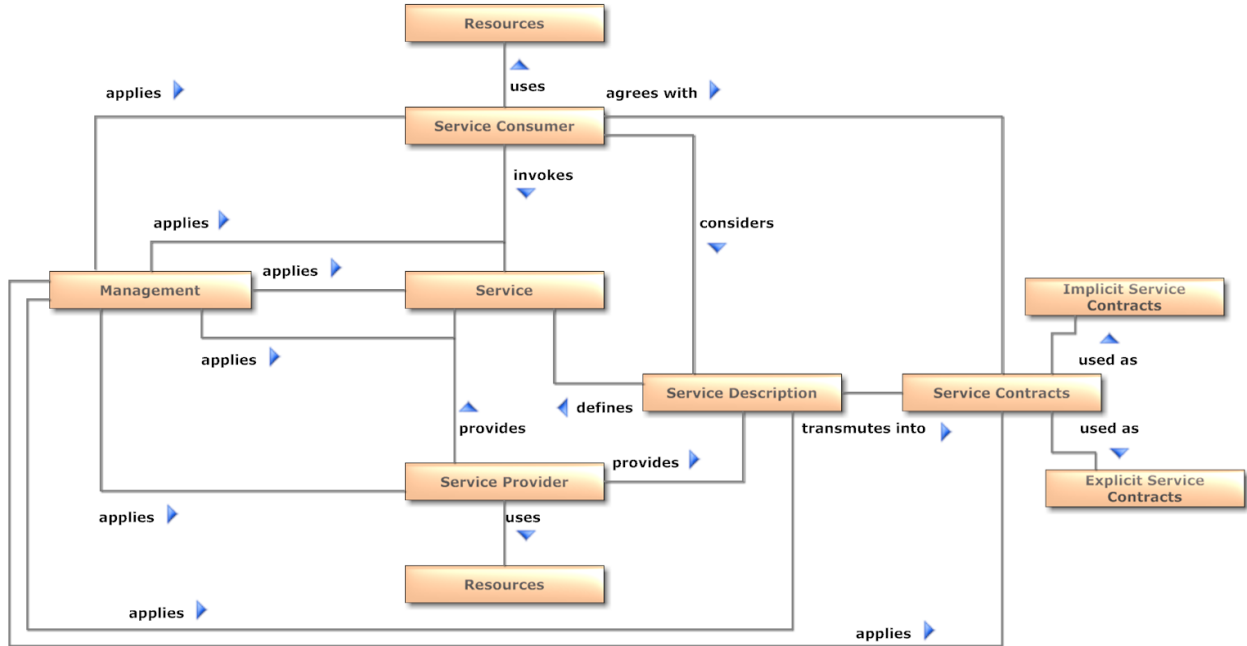
- 3863 • the commitment by a service provider to provide service functionality and results  
3864 consistent with identified real world effects and
- 3865 • the commitment by a service consumer to interact with the service per specific means  
3866 and per specified policies,

3867 where both consumer and provider actions are in the manner described in the service description.

3868 The service description provides the basis for the service contract and, in some situations, may be used  
3869 as an implicit default service contract. In addition, the service description may set mandatory aspects of a  
3870 service contract, e.g. for security services, or may specify acceptable alternatives. As an example of  
3871 alternatives, the service description may identify which versions of a vocabulary will be recognized, and  
3872 the specifics of the contract are satisfied when the consumer uses one of the alternatives. Another  
3873 alternative could have a consumer identifying a policy they require be satisfied, e.g. a standard privacy  
3874 policy on handling personal information, and a provider that is prepared to accept a policy request would  
3875 indicate acceptance as part of the service contract by continuing with the interaction. In each of these  
3876 cases, the actions of the participants are consistent with an implicit service contract without the existence  
3877 of a formal agreement between the participants.

3878 In the case of business services, it is anticipated that the service contract may take an explicit form and  
3879 the agreement between business consumer and business service provider is formalized. Formalization  
3880 requires up-front interactions between service consumer and service provider. In many business  
3881 interactions, especially between business organisations within or across corporate boundaries, a  
3882 consumer needs a contractual assurance from the provider or wants to explicitly indicate choices among  
3883 alternatives, e.g., only use a subset of the business functionality offered by the service and pay a  
3884 prorated cost.

3885



3886

3887

Figure 48 Management of the service interaction

3888

3889

3890

3891

3892

3893

3894

3895

3896

Consequently, an implicit service contract is an agreement (1) on the consumer side with the terms, conditions, features and interaction means specified in the service description "as is" or (2) a selection from alternatives that are made available through mechanisms included in the service description, and neither of these require any a priori interactions between the service consumer and the service provider. An explicit service contract always requires a form of interaction between the service consumer and the service provider prior to the service invocation. This interaction may regard the choice or selection of the subset of the elements of the service description or other alternatives introduced through the formal agreement process that would be applicable to the interaction with the service and affect related joint action.

3897

3898

3899

3900

3901

3902

3903

3904

3905

3906

3907

Any form of explicit contract couples the service consumer and provider. While explicit contract may be necessary or desirable in some cases, such as in supply chain management, commerce often uses a mix of implicit and explicit contracts, and a service provider may offer (via service description) a conditional shift from implicit to explicit contract. For example, Twitter offers an implicit contract on the use of its APIs to any application with the limit on the amount of service invocations; if the application needs to use more invocations, one has to enter into the explicit fee-based contract with the provider. A case where an implicit contract transforms into explicit contract may be illustrated when one buys a new computer and it does not work. The buyer returns the computer for repairing under manufacture warranty as stated by an implicit purchase contract. However, if the repair does not fix the problem and the seller offers a replacement by upgraded model, the buyer may agree to an explicit contract that limits the rights of the buyer to make the explicit agreement public.

3908

3909

3910

3911

3912

Control of the quality of the service execution, often represented as a service level agreement (SLA), is performed by service monitoring systems and includes both technical and operational business controls. SLA is a part of the service contract and, because of individual nature of this type of contracts, may vary from one service contract to another, even for the same consumer. Typically, a particular SLA in the service contract is a concrete instance of the SLA declared in the service description.

3913

3914

3915

3916

3917

3918

3919

Management of the service contracts is based on management policies that may be mentioned in the service description and in the service contracts. Management of the service contracts is mandatory for consumer relationship management. In the case of explicit service contracts, the contracts have to be created, stored, maintained, reviewed/controlled and archived/destroyed as needed. All the activities are management concerns. Explicit service contracts may be stored in specialised repositories that provide appropriate level of security.

3920 Management of the service interfaces is based on several management policies that regulate  
3921     • availability of interfaces specified in the service contracts,  
3922     • accessibility of interfaces,  
3923     • procedures for interface changes,  
3924     • interface versions and well as the versions of all parts of the interfaces, and  
3925     • traceability of the interfaces and their versions back to the service description document.

3926  
3927 Management of the SLA is integral to the management of service monitoring and operational service  
3928 behavior at run-time. A SLA usually enumerates service characteristics and expected performances of  
3929 the service. Since SLA carries connotation of “promise”, monitoring is needed to know if the promise is  
3930 kept. Existence of an SLA itself does not guarantee the consumer will be provided with the service level  
3931 specified in the service contract.

3932 The use of SLA in SOA ecosystem can be wider than just an agreement on technical performances.  
3933 An SLA may contain remedies for situations where the promised service cannot be maintained, or the  
3934 real world effect can't be achieved due to developments subsequent to the agreement. A service  
3935 consumer that acts accordingly to realize the real world effect may be compensated for the breach of the  
3936 SLA if the effect is not realized.

3937 Management of the SLA includes, among others, policies for the SLA changes, updates, and  
3938 replacement. This aspect concerns service Execution Context because the business logic associated with  
3939 a defined interface may differ in different Execution Contexts and affect the overall performance of the  
3940 service.

#### 3941 **5.3.3.3 Policies**

3942 "Although provision of management capabilities enables a service to become manageable, the extent and  
3943 degree of permissible management are defined in management policies that are associated with the  
3944 services. Management policies are used to define the obligations for, and permissions to, managing the  
3945 service" [WSA]. Management policies, in essence, are the realisation of governing rules and regulations.  
3946 As such, some management policies may target services while other policies may target the management  
3947 of the services.

3948 In practice, a policy without any means of enforcing it is vacuous. In the case of management policy, we  
3949 rely on a management infrastructure to realize and enforce management policy.

#### 3950 **5.3.3.4 Service Description and Management**

3951 The service description identifies several management objects such as a set of service interfaces and  
3952 related set of SLAs: service behavioral characteristics and performances specified in the SLA depend on  
3953 the interface type and its Execution Context. In the service description, a service consumer can find  
3954 references to management policies, SLA metrics, and the means of accessing measured values that  
3955 together increase assurance in the service quality. At the same time, service description is an artifact that  
3956 needs to be managed.

3957 In the SOA ecosystem, the service description is the assembled information that describes the service but  
3958 it may be reported or displayed in different presentations. While each separate version of the service has  
3959 one and only one service description, different categories of service consumers may focus their interests  
3960 on different aspects of the service description. Thus, the same service description may be displayed not  
3961 only in different languages but also with different cultural and professional accents in the content.

3962 New service description may be issued to reflect changes and update in the service. If the change in the  
3963 service does not affect its service description, the new service version may have the same service  
3964 description as the previous version except for the updated version identifier. For example, a service  
3965 description may stay the same if bugs were fixed in the service. However, if a change in the service  
3966 influences any aspects of the service quality that can affect the real world effect resulting from  
3967 interactions with the service, the service description must reflect this change even if there are no changes  
3968 to the service interface.

3969 Management of the service description and related explicit service contracts is essential for delivery of the  
3970 service to the consumer satisfaction. This management can also prevent business problems rooted in  
3971 poor communication between the service consumers and the service providers.

3972 Thus, management of the service description contains, among others, management of the service  
3973 description presentations, the life-cycles of the service descriptions, service description distribution  
3974 practices and storage of the service descriptions and related service contracts. Collections of service  
3975 descriptions in the enterprise may manifest a need for specialised registries and/or repositories.  
3976 Depending on the enterprise policies, an allocation of purposes and duties of registries and repositories  
3977 may vary but this topic is beyond the current scope.

#### 3978 **5.3.4. Management for Monitoring and Reporting**

3979 The successful application of management relies on the monitoring and reporting aspects of  
3980 management to enable the control aspect. Monitoring in the context of management consists of  
3981 measuring values of managed aspects and evaluating that measurement in relationship to some  
3982 expectation. Monitoring in a SOA ecosystem is enabled through the use of mechanisms by resources for  
3983 exposing managed aspects. In the SOA framework, this mechanism may be a service for obtaining the  
3984 measurement. Alternatively, the measurement may be monitored by means of event generation  
3985 containing updated values of the managed aspect.

3986 Approaches to monitoring may use a polling strategy in which the measurements are requested from  
3987 resources in periodic intervals, in a pull strategy in which the measurements are requested from  
3988 resources at random times, or in a push strategy in which the measurements are supplied by the  
3989 resource without request. The push strategy can be used in a periodic update approach or in an “update  
3990 on change” approach. Management services must be capable of handling these different approaches to  
3991 monitoring.

3992 Reporting is the complement to monitoring. Where monitoring is responsible for obtaining measurements,  
3993 reporting is responsible for distributing those measurements to interested stakeholders. The separation  
3994 between monitoring and reporting is made to include the possibility that data obtained through monitoring  
3995 might not be used until an event impacting the ecosystem occurs or the measurement requires further  
3996 processing to be useful. In the SOA framework, reporting is provided using services for requesting  
3997 measurement reports. These reports may consist of raw measurement data, formatted collections of  
3998 data, or the results of analysis performed on measurement data from collections of different managed  
3999 aspects. Reporting is also used to support logging and auditing capabilities, where the reporting  
4000 mechanisms create log or audit entries.

#### 4001 **5.3.5 Management for Infrastructure**

4002 All of the properties, policies, interactions, resources, and management are only possible if a SOA  
4003 ecosystem infrastructure provides support for managed capabilities. Each managed capability imposes  
4004 different requirements on the capabilities supplied by the infrastructure in SOA ecosystem and requires  
4005 that those capabilities be usable as services or at the very least be interoperable.

4006 Not providing the full list of infrastructural elements of SOA ecosystem, we list an example of such  
4007 elements here:

- 4008 1. Registries and repositories for services, policies, and related descriptions  
4009 and contracts
- 4010 2. Synchronous and asynchronous communication channels for service  
4011 interactions (e.g., network, e-mail, message routing with ability of mediating  
4012 transport protocols, etc.)
- 4013 3. Recovery capabilities
- 4014 4. Security controls

4015 Also, a SOA ecosystem infrastructure, enabling service management, should support

- 4016 1. Management enforcement and control means
- 4017 2. Monitoring and SLA validation controls
- 4018 3. Testing and Reporting capabilities

4019 Combination of manageability capabilities and infrastructure elements constitutes certain level of SOA  
4020 management maturity. While several maturity models exist, this topic is out of the scope of the document.



## 4021 5.4 SOA Testing Model

4022 *Program testing can be used to show the presence of bugs,*  
4023 *but never to show their absence!*  
4024 Edsger Dijkstra

4025 Testing for SOA combines the typical challenges of software testing and certification  
4026 with the additional needs of accommodating the distributed nature of the resources, the  
4027 greater access of a more unbounded consumer population, and the desired flexibility to  
4028 create new solutions from existing components over which the solution developer has  
4029 little if any control. The purpose of testing is to demonstrate a required level of reliability,  
4030 correctness, and effectiveness that enable prospective consumers to have adequate  
4031 confidence in using a service. Adequacy is defined by the consumer based on the  
4032 consumer's needs and context of use. As the Dijkstra quote points out, absolute  
4033 correctness and completeness cannot be proven by testing; however, for SOA, it is  
4034 critical for the prospective consumer to know what testing has been performed, how it  
4035 has been performed, and what were the results.

### 4036 5.4.1 Traditional Software Testing as Basis for SOA Testing

4037 SOA services are largely software artifacts and can leverage the body of experience  
4038 that has evolved around software testing. IEEE-829 specifies the basic set of software  
4039 test documents while allowing flexibility for tailored use. As such, the document  
4040 structure can also provide guidance to SOA testing.

4041 IEEE-829 covers test specification and test reporting through use of the following  
4042 document types:

- 4043 • *Test plan* documenting the scope (what is to be tested, both which entity and what  
4044 features of the entity), the approach (how it is tested), and the needed resources  
4045 (who does the testing, for how long), with details contained in the:
  - 4046 • *Test design specification*: features to be tested, test conditions (e.g. test cases,  
4047 test procedures needed) and expected results (criteria for passing test); entrance  
4048 and exit criteria
  - 4049 • *Test case specification*: test data used for input and expected output
  - 4050 • *Test procedure specification*: steps required to run the test, including any set-up  
4051 preconditions
- 4052 • *Test item transmittal* to identify the test items being transmitted for testing
- 4053 • *Test log* to record what occurred during test, i.e. which tests run, who ran, what  
4054 order, what happened
- 4055 • *Test incident report* to capture any event that happened during test which requires  
4056 further investigation
- 4057 • *Test summary* as a management report summarizing test run and results,  
4058 conclusions

4059 In summary, IEEE-829 captures (1) what was tested, (2) how it was tested, e.g. the test  
4060 procedure used, and (3) the results of the test.

#### 4061 **5.4.1.1 Types of Testing**

4062 There are numerous aspects of testing that, in total, work to establish that an entity is  
4063 (1) built as required per policies and related specifications prescribed by the entity's  
4064 owner, and (2) delivers the functionality required by its intended users. This is often  
4065 referred to as verification and validation.

4066 Policies, as described in Section 4.4, that are related to testing may prescribe but are  
4067 not limited to the business processes to be followed, the standards with which an  
4068 implementation must comply, and the qualifications of and restrictions on the users. In  
4069 addition to the functional requirements prescribing what an entity does, there may also  
4070 be non-functional performance and/or quality metrics that state how well the entity does  
4071 it. The relation of these policies to SOA testing is discussed further below.

4072 The identification of policies is the purview of governance (section 5.1) and the assuring  
4073 of compliance (including response to noncompliance) with policies is a matter for  
4074 management (section **Error! Reference source not found.**).

#### 4075 **5.4.1.2 Range of Test Conditions**

4076 Test conditions and expected responses are detailed in the test case specification. The  
4077 test conditions should be designed to cover the areas for which the entity's response  
4078 must be documented and may include:

- 4079 • nominal conditions
- 4080 • boundaries and extremes of expected conditions
- 4081 • breaking point where the entity has degraded below a certain level or has  
4082 otherwise ceased effective functioning
- 4083 • random conditions to investigate unidentified dependencies among combinations  
4084 of conditions
- 4085 • errors conditions to test error handling

4086 The specification of how each of these conditions should be tested for SOA resources,  
4087 including the infrastructure elements of the SOA ecosystem, is beyond the scope of this  
4088 document but is an area that evolves along with operational SOA experience.

#### 4089 **5.4.1.3 Configuration Management of Test Artifacts**

4090 The test item transmittal provides an unambiguous identification of the entity being  
4091 tested, thus REQUIRING that the configuration of the entity is appropriately tracked and  
4092 documented. In addition, the test documents (such as those specified by IEEE-829)  
4093 MUST also be under a documented and appropriately audited configuration  
4094 management process, as should other resources used for testing. The description of  
4095 each artifact would follow the general description model as discussed in section 4.1.1.1;  
4096 in particular, it would include a version number for the artifact and reference to the  
4097 documentation describing the versioning scheme from which the version number is  
4098 derived.

4099  
4100 [EDITOR'S NOTE: TO WHAT EXTENT SHOULD CM BE EXPLICITLY INCLUDED IN THE MANAGEMENT  
4101 SECTION?]

## 4102 5.4.2 Testing and the SOA Ecosystem

4103 [EDITOR'S NOTE: THE EMPHASIS THOUGH MUCH OF THE RA IS THE LARGER ECOSYSTEM BUT WE NEED  
4104 WORDS IN SECTION 3 TO ACKNOWLEDGE THE EXISTENCE OF THE ENTERPRISE AND THAT AN  
4105 ENTERPRISE (AS COMMONLY INTERPRETED) IS LIKELY MORE CONSTRAINED AND MORE PRECISELY  
4106 DESCRIBED FOR THE CONTEXT OF THE ENTERPRISE. THE ECOSYSTEM PERSPECTIVE, THOUGH, IS  
4107 STILL APPLICABLE FOR THE FOLLOWING REASONS:

- 4108
- 4109 1. A GIVEN ENTERPRISE MAY COMPRISE NUMEROUS CONSTITUENT ENTERPRISES THAT  
4110 RESEMBLE THE INDEPENDENT ENTITIES DESCRIBED FOR THE ECOSYSTEM. AN ENTERPRISE  
4111 MAY ATTEMPT TO REDUCE VARIATIONS AMONG THE CONSTITUENTS BUT THE *PARTICIPATION IN*  
4112 *A SOA ECOSYSTEM* VIEW ENABLES SOA TO BENEFIT THE ENTERPRISE WITHOUT REQUIRING  
4113 THE ENTERPRISE ISSUES TO BE FULLY RESOLVED.
  - 4114 2. RESOURCES SPECIFICALLY MOTIVATED BY THE CONTEXT OF THE ENTERPRISE CAN BE MORE  
4115 READILY USED IN A DIFFERENT CONTEXT IF ECOSYSTEM CONSIDERATIONS ARE INCLUDED AT  
4116 AN EARLY STAGE. THE CHANGE IN A CONTEXT MAY BE A FUNDAMENTAL CHANGE IN THE  
4117 ENTERPRISE OR THE NEWLY DISCOVERED APPLICABILITY OF ENTERPRISE RESOURCES TO USE  
4118 OUTSIDE THE ENTERPRISE.

4119

4120 IN THIS DOCUMENT, REFERENCE TO THE SOA ECOSYSTEM APPLIES BUT WITH POSSIBLY LESS  
4121 GENERALITY TO AN ENTERPRISE USE OF SOA.]

4122 Testing of SOA artifacts for use in the SOA ecosystem differs from traditional software  
4123 testing for several reasons. First, a highly touted benefit of SOA is to enable  
4124 unanticipated consumers to make use of services for unanticipated purposes.  
4125 Examples of this could include the consumer using a service for a result that was not  
4126 considered the primary one by the provider, or the service may be used in combination  
4127 with other services in a scenario that is different from the one considered when  
4128 designing for the initial target consumer community. It is unlikely that a new consumer  
4129 will push the services back to anything resembling the initial test phase to test the new  
4130 use, and thus additional paradigms for testing are necessary. Some testing may  
4131 depend on the availability of test resources made available as a service outside the  
4132 initial test community, while some testing is likely to be done as part of limited use in the  
4133 operational setting. The potential responsibilities related to such "consumer testing" is  
4134 discussed further below.

4135 Secondly, in addition to consumers who interact with a service to realize the described  
4136 real world effects, the developer community is also intended to be a consumer. In the  
4137 SOA vision of reuse, the developer composes new solutions using existing services,  
4138 where the existing services provides access to some desired real world effects that are  
4139 needed by the new solution. The new solution is a consumer of the existing services,  
4140 enabling repeated interactions with the existing services playing the role of reusable  
4141 components. Note, those components are used at the locations where they individually  
4142 reside and are not typically duplicated for the new solution. The new solution may itself  
4143 be offered as a SOA service, and a consumer of the service composition representing  
4144 the new solution may be totally unaware of the component services being used. (See  
4145 section 4.3.4 for further discussion on service compositions.)

4146 Another difference from traditional testing is that the distributed, unbounded nature of  
4147 the SOA ecosystem makes it unlikely to have an isolated test environment that  
4148 duplicates the operational environment. A traditional testing approach often makes use  
4149 of a test system that is identical to the eventual operational system but isolated for  
4150 testing. After testing is successfully completed, the tested entity would be migrated to

4151 the operational environment, or the test environment may be delivered as part of the  
4152 system to become operational. This is not feasible for the SOA ecosystem as a whole.  
4153 SOA services must be testable in the environment and under the conditions that can be  
4154 encountered in the operational SOA ecosystem. As the ecosystem is in a state of  
4155 constant change, so some level of testing is continuous through the lifetime of the  
4156 service, leveraging utility services used by the ecosystem infrastructure to monitor its  
4157 own health and respond to situations that could lead to degraded performance. This  
4158 implies the test resources must incorporate aspects of the SOA paradigm, and a  
4159 category of services may be created to specifically support and enable effective  
4160 monitoring and continuous testing for resources participating in the SOA ecosystem.  
4161 While SOA within an enterprise may represent a more constrained and predictable  
4162 operational environment, the composability and unanticipated use aspects are highly  
4163 touted within the enterprise. The expanded perspective on testing may not be as  
4164 demanding within an enterprise but fuller consideration of the ecosystem enables the  
4165 enterprise to be more responsive should conditions change.

### 4166 **5.4.3 Elements of SOA Testing**

4167 IEEE-829 identifies fundamental aspects of testing, and many of these should carry  
4168 over to SOA testing: in particular, the identification of what is to be tested, how it is to be  
4169 tested, and by whom the testing is to be done. While IEEE-829 identifies a suggested  
4170 document tree, the availability of these documents in the SOA ecosystem is discussed  
4171 below.

#### 4172 **5.4.3.1 What is to be Tested**

4173 The focus of this discussion is the SOA service. It is recognized that the infrastructure  
4174 components of any SOA environment are likely to also be SOA services and, as such,  
4175 falls under the same testing guidance. Other resources that contribute to a SOA  
4176 environment may not be SOA services, but are expected to satisfy the intent if not the  
4177 letter of guidance presented here. Specific differences for such resources are as yet  
4178 largely undefined and further elaboration is beyond the scope of the SOA-RAF.

4179 The following discussion often focuses on a singular SOA service but it is implicit that  
4180 any service may be a composite of other services. As such, testing the functionality of a  
4181 composite service may effectively be testing an end-to-end business process that is  
4182 being provided by the composite service. If new versions are available for the  
4183 component services, appropriate end-to-end testing of the composite may be required  
4184 in order to verify that the composite functionality is still adequately provided. The level  
4185 of required testing of an updated composite depends on policies of those providing the  
4186 service, policies of those using the service, and mission criticality of those depending on  
4187 the service results.

4188 The SOA service to be tested MUST be unambiguously identified as specified by its  
4189 applicable configuration management scheme. Specifying such a scheme is beyond  
4190 the scope of the SOA-RAF other than to say the scheme should be documented and  
4191 itself under configuration management.

#### 4192 5.4.3.1.1 Origin of Test Requirements

4193 In the Service Description model (Figure 21), the aspects of a service that need to be  
4194 described are:

- 4195 • the service functionality and technical assumptions that underlie the functionality;
- 4196 • the policies that describe conditions of use;
- 4197 • the service interface that defines information exchange with the service;
- 4198 • service reachability that identifies how and where message exchange is to occur;
- 4199 and
- 4200 • metrics access for any participant to have information on how a service is  
4201 performing.

4202 Service testing must provide adequate assurance that each of these aspects is  
4203 operational as defined.

4204 The information in the service description comes from different sources. The  
4205 functionality is defined through whatever process identifies needs and the community  
4206 for which these needs are addressed. The process may be ad hoc as serves the  
4207 prospective service owner or strictly governed, but defining the functionality is an  
4208 essential first step in development. It is also an early and ongoing focus of testing to  
4209 ensure the service accurately reflects the described functionality and the described  
4210 functionality accurately addresses the consumer needs.

4211 Policies define the conditions of development and conditions of use for a service and  
4212 are typically specified as part of the governance process. Policies constraining service  
4213 development, such as coding standards and best practices, require appropriate testing  
4214 and auditing during development to ensure compliance. While the governance process  
4215 identifies development policies, these are likely to originate from the technical  
4216 community responsible for development activities. Policies that define conditions of use  
4217 often define business practices that service owners and providers or those responsible  
4218 for the SOA infrastructure want followed. These policies are initially tested during  
4219 service development and are continuously monitored during the operational lifetime of  
4220 the service.

4221 The testing of the service interface and service reachability are often related but  
4222 essentially reflect different motivations and needs. The service interface is specified as  
4223 a joint product of the service owners and providers who define service functionality, the  
4224 prospective consumer community, the service developer, and the governance process.  
4225 The semantics of the information model must align with the semantics of those who  
4226 consume the service in order for there to be meaningful exchange of information. The  
4227 structure of the information is influenced by the consumer semantics and the  
4228 requirements and constraints of the representation as interpreted by the service  
4229 developer. The service process model that defines actions which can be performed  
4230 against a service and any temporal dependencies derive from the defined functionality  
4231 and may be influenced by the development process. Any of these constraints may be  
4232 identified and expressed as policy through the governance process.

4233 Service reachability conditions are the purview of the service provider who identifies the  
4234 service endpoint and the protocols recognized at the endpoint. These may be

4235 constrained by governance decisions on how endpoint addresses may be allocated and  
4236 what protocols should be used.

4237 While the considerations for defining the service interface derive from several sources,  
4238 testing of the service interface is more straightforward and isolated in the testing  
4239 process. At any point where the interface is modified or exposes a new resource, the  
4240 message exchange should be monitored both to ensure the message reaches its  
4241 intended destination and it is parsed correctly once received. Once an interface has  
4242 been shown to function properly, it is unlikely to fail later unless something fundamental  
4243 to the service changes.

4244 The service interface is also tested when the service endpoint changes. Testing of the  
4245 endpoint ensures message exchange can occur at the time of testing and the initial  
4246 testing shows the interface is being processed properly at the new endpoint.

4247 Functioning of a service endpoint at one time does not guarantee it is functioning at  
4248 another time, e.g. the server with the endpoint address may be down, making testing of  
4249 service reachability a continual monitoring function through the life of the service's use  
4250 of the endpoint. Also, while testing of the service endpoint is a necessary and most  
4251 commonly noted part of the test regiment, it is not in itself sufficient to ensure the other  
4252 aspects of testing discussed in this section.

4253 Finally, governance is impossible without the collection of metrics against which service  
4254 behavior can be assessed. Metrics are also a key indicator for consumers to decide if a  
4255 service is adequate for their needs. For instance, the average response time or the  
4256 recent availability can be determining factors even if there are no rules or regulations  
4257 promulgated through the governance process against which these metrics are  
4258 assessed. The available metrics are a combination of those expected by the consumer  
4259 community and those mandated through the governance process. The total set of  
4260 metrics will evolve over time with SOA experience. Testing of the services that gather  
4261 and provide access to the metrics will follow testing as described in this section, but for  
4262 an individual service, testing will ensure that the metrics access indicated in the service  
4263 description is accurate.

4264 The individual test requirements highlight aspects of the service that testing must  
4265 consider but testing must establish more than isolated behavior. The emphasis is the  
4266 holistic results of interacting with the service in the SOA environment. Recall that the  
4267 execution context is the set of agreements between a consumer and a provider that  
4268 define the conditions under which service interaction occurs. The agreements are  
4269 expected to be predominantly the acceptance of the standard conditions as enumerated  
4270 by the service provider, but it may include the identification of alternate conditions that  
4271 will govern the interaction.

4272 For example, the provider may prefer a policy where it can sell the contact information  
4273 of its consumers but will honor the request of a consumer to keep such information  
4274 private. The identification of the alternate privacy policy is part of the execution context,  
4275 and it is the application of and compliance with this policy that operational monitoring  
4276 will attempt to measure. The collection of metrics showing this condition is indeed met  
4277 when chosen is considered part of the ongoing testing of the service.

4278 Other variations in the execution context also require monitoring to ensure that different  
4279 combinations of conditions perform together as desired. For example, if a new privacy  
4280 policy takes additional resources to apply, this may affect quality of service and

4281 propagate other effects. These could not be tested during the original testing if the  
4282 alternate policy did not exist at that time.

#### 4283 **5.4.3.1.2 Testing Against Non-Functional Requirements**

4284 Testing against non-functional requirements constitutes testing of business usability of  
4285 the service. In a marketplace of services, non-functional characteristics may be the  
4286 primary differentiator between services that produce essentially the same real world  
4287 effects.

4288 As noted in the previous section, non-functional characteristics are often associated  
4289 with policies or other terms of use and may be collected in service level contracts  
4290 offered by the service providers. Non-functional requirements may also reflect the  
4291 network and hardware infrastructure that support communication with the service, and  
4292 changes may impact quality of service. The service consumer and even the service  
4293 provider may not be aware of all such infrastructure changes but the changes may  
4294 manifest in shared states that impact the usability of the service.

4295 In general, a change in the non-functional requirements results in a change to the  
4296 execution context, but as with any collection of information that constitutes a  
4297 description, the execution context is unable to explicitly capture all non-functional  
4298 requirements that may apply. A change in non-functional requirements, whether  
4299 explicitly part of the execution context or an implicit contributor, may require retesting of  
4300 the service even if its functionality and the implementation of the functionality has not  
4301 changed. Depending on the circumstances, retesting may require a formal recertifying  
4302 of end-to-end behavior or more likely will be part of the continuous monitoring that  
4303 applies throughout the service lifetime.

#### 4304 **5.4.3.1.3 Testing Content and the Interests of Consumers**

4305 As noted in section 5.4.1.1, testing may involve verification of conformance with respect  
4306 to policies and technical specifications and validation with respect to sufficiency of  
4307 functionality to meet some prescribed use. It may also include demonstration of  
4308 performance and quality aspects. For some of these items, such as demonstrating the  
4309 business processes followed in developing the service or the use of standards in  
4310 implementing the service, the testing or relevant auditing is done internal to the service  
4311 development process and follows traditional software testing and quality assurance. If it  
4312 is believed of value to potential consumers, information about such testing could be  
4313 included in the service description. However, it is not required that all test or  
4314 compliance artifacts be available to consumers, as many of the details tested may be  
4315 part of the opacity of the service implementation.

4316 Some aspects of the service being tested will reflect directly on the real world effects  
4317 realized through interaction with the service. In these cases, it is more likely that testing  
4318 results will be directly relevant to potential consumers. For example, if the service was  
4319 designed to correspond to certain elements of a business process or that a certain  
4320 workflow is followed, testing should verify that the real world effects reflect that the  
4321 business process or workflow were satisfactorily captured.

4322 The testing may also need to demonstrate that specified conditions of use are satisfied.  
4323 For example, policies may be asserted that require certain qualifications of or impose  
4324 restrictions on the consumers who may interact with the service. The service testing

4325 must demonstrate that the service independently enforces the policies or it provides the  
4326 required information exchanges with the SOA ecosystem so other resources can ensure  
4327 the specified conditions.

4328 The completeness of the testing, both in terms of the features tested and the range of  
4329 parameters for which response is tested, depends on the context of expected use: the  
4330 more critical the use, the more complete the testing. There are always limits on the  
4331 resources available for testing, if nothing else than the service must be available for use  
4332 in a finite amount of time.

4333 This again emphasizes the need for adequate documentation to be available. If the  
4334 original testing is very thorough, it may be adequate for less demanding uses in the  
4335 future. If the original testing was more constrained, then well-documented test results  
4336 establish the foundation on which further testing can be defined and executed.

#### 4337 **5.4.3.2 How Testing is to be Done**

4338 Testing should follow well-defined methodologies and, if possible, should reuse test  
4339 artifacts that have proven generally useful for past testing. For example, IEEE-829  
4340 notes that test cases are separated from test designs to allow for use in more than one  
4341 design and to allow for reuse in other situations. In the SOA ecosystem, description of  
4342 such artifacts, as with description of a service, enables awareness of the item and  
4343 describes how the artifact may be accessed or used.

4344 As with traditional testing, the specific test procedures and test case inputs are  
4345 important so the tests are unambiguously defined and entities can be retested in the  
4346 future. Automated testing and regression testing may be more important in the SOA  
4347 ecosystem in order to re-verify a service is still acceptable when incorporated in a new  
4348 use. For example, if a new use requires the services to deal with input parameters  
4349 outside the range of initial testing, the tests could be rerun with the new parameters. If  
4350 the testing resources are available to consumers within the SOA ecosystem, the testing  
4351 as designed by test professionals could be consumed through a service accessed by  
4352 consumers, and their results could augment those already in place. This is discussed  
4353 further in the next section.

#### 4354 **5.4.3.3 Who Performs the Testing**

4355 As with any software, the first line of testing is unit testing done by software developers.  
4356 It is likely that initial testing will be done by those developing the software but may also  
4357 be done independently by other developers. For SOA development, unit testing is likely  
4358 confined to a development sandbox isolated from the SOA ecosystem.

4359 SOA testing will differ from traditional software testing in that testing beyond the  
4360 development sandbox must incorporate aspects of the SOA ecosystem, and those  
4361 doing the testing must be familiar with both the characteristics and responses of the  
4362 ecosystem and the tools, especially those available as services, to facilitate and  
4363 standardize testing. Test professionals will know what level of assurance must be  
4364 established as the exposure of the service to the ecosystem and ecosystem to the  
4365 service increases towards operational status. These test professionals may be internal  
4366 resources to an organization or may evolve as a separate discipline provided through  
4367 external contracting.



4368 As noted above, it is unlikely that a complete duplicate of the SOA ecosystem will be  
4369 available for isolated testing, and thus use of ecosystem resources will manifest as a  
4370 transition process rather than a step change from a test environment to an operational  
4371 one. This is especially true for new composite services that incorporate existing  
4372 operational services to achieve the new functionality. The test professionals will need to  
4373 understand the available resources and the ramifications of this transition.

4374 As with current software development, a stage beyond work by test professionals will  
4375 make use of a select group of typical users, commonly referred to as beta testers, to  
4376 report on service response during typical intended use. This establishes fitness by the  
4377 consumers, providing final validation of previously verified processes, requirements, and  
4378 final implementation.

4379 In traditional software development, beta testing is the end of testing for a given version  
4380 of the software. However, although the initial test phase can establish an appropriate  
4381 level of confidence consistent with the designed use for the initial target consumer  
4382 community, the operational service will exist in an evolving ecosystem, and later  
4383 conditions of use may differ from those thought to be sufficient during the initial testing.  
4384 Thus, operational monitoring becomes an extension of testing through the service  
4385 lifetime. This continuous testing will attempt to ensure that a service does not consume  
4386 an inordinate amount of ecosystem resources or display other behavior that degrades  
4387 the ecosystem, but it will not undercover functional errors that may surface over time.

4388 As with any software, it is the responsibility of the consumers to consider the  
4389 reasonableness of solutions in order to spot errors in either the software or the way the  
4390 software is being used. This is especially important for consumers with unanticipated  
4391 uses that may go beyond the original test conditions. It is unlikely the consumers will  
4392 initiate a new round of formal testing unless the new use requires a significantly higher  
4393 level of confidence in the service. Rather the consumer becomes a new extension to  
4394 the testing regiment. Obvious testing would include a sanity check of results during the  
4395 new use. However, if the details of legacy testing are associated with the service  
4396 through the service description and if testing resources are available through automated  
4397 testing services, then the new consumers can rerun and extend previous testing to  
4398 include the extended test conditions. If the test results are acceptable, these can be  
4399 added to the documentation of previous results and become the extended basis for  
4400 future decisions by prospective consumers on the appropriateness of the service. If the  
4401 results are not acceptable or in some way questionable, the responsible party for the  
4402 service or testing professionals can be brought in to decide if remedial action is  
4403 necessary.

#### 4404 **5.4.3.4 How Testing Results are Reported**

4405 For any SOA service, an accurate reporting of the testing a service has undergone and  
4406 the results of the testing is vital to consumers deciding whether a service is appropriate  
4407 for intended use. Appropriateness may be defined by a consumer organization and  
4408 require specific test regiments culminating in a certification; appropriateness could be  
4409 established by accepting testing and certifications that have been conferred by others.

4410 The testing and certification information should be identified in the service description.  
4411 Referring to the general description model of *Figure 13*, tests conducted by or under a  
4412 request from the service owner (see ownership in section **Error! Reference source not**

4413 **found.**) would be captured under Annotations from Owners. Testing done by others,  
4414 such as consumers with unanticipated uses, could be associated through Annotations  
4415 from 3rd Parties. The annotations should clearly indicate what was tested, how the  
4416 testing was done, who did the testing, and the testing results. The clear description of  
4417 each of these artifacts and of standardized testing protocols for various levels of  
4418 sophistication and completeness of testing would enable a common understanding and  
4419 comparison of test coverage. It will also make it more straightforward to conduct and  
4420 report on future testing, facilitating the maintenance of the service description.

4421 Consumer testing and the reporting of results raises additional issues. While stating  
4422 who did the testing is mandatory, there may be formal requirements for authentication of  
4423 the tester to ensure traceability of the testing claims. In some circumstances, persons  
4424 or organizations would not be allowed to state testing claims unless the tester was an  
4425 approved entity. In other cases, ensuring the tester had a valid email may be sufficient.  
4426 In either case, it would be at the discretion of the potential consumer to decide what  
4427 level of authentication was acceptable and which testers are considered authoritative in  
4428 the context of their anticipated use.

4429 Finally, in a world of openly shared information, we would see an ever-expanding set of  
4430 testing information as new uses and new consumers interact with a service. In reality,  
4431 these new uses may represent proprietary processes or classified use that should only  
4432 be available to authorized parties. Testing information, as with other elements of  
4433 description, may require special access controls to ensure appropriate access and use.

#### 4434 **5.4.4 Testing SOA Services**

4435 Testing of SOA services should be consistent with the SOA paradigm. In particular,  
4436 testing resources and artifacts should be visible in support of service interaction  
4437 between providers and consumers, where here the interaction is between the testing  
4438 resource and the tester. In addition, the idea of opacity of the implementation should  
4439 limit the details that need to be available for effective use of the test resources. Testing  
4440 that requires knowledge of the internal structure of the service or its underlying  
4441 capability should be performed as part of unit testing in the development sandbox, and  
4442 should represent a minimum level of confidence before the service begins its transition  
4443 to further testing and eventual operation in the SOA ecosystem.

##### 4444 **5.4.4.1 Progression of SOA Testing**

4445 Software testing is a gradual exercise going from micro inspection to testing macro  
4446 effects. The first step in testing is likely the traditional code reviews. SOA  
4447 considerations would account for the distributed nature of SOA, including issues of  
4448 distributed security and best practices to ensure secure resources. It would also set the  
4449 groundwork for opacity of implementation, hiding programming details and simplifying  
4450 the use of the service.

4451 Code review is likely followed by unit testing in a development sandbox isolated from  
4452 the operational environment. The unit testing is done with full knowledge of the service  
4453 internal structure and knowledge of resources representing underlying capabilities. It  
4454 tests the interface to ensure exchanged messages are as specified in the service  
4455 description and the messages can be parsed and interpreted as intended. Unit testing  
4456 also verifies intended functionality and that the software has dealt correctly with internal

4457 dependencies, such as structure of a file system or access to other dedicated  
4458 resources.

4459 Some aspects of unit testing require external dependencies be satisfied, and this is  
4460 often done using mock objects to substitute for the external resources. In particular, it  
4461 will likely be necessary to include mocks of existing operational services, both those  
4462 provided as part of the SOA infrastructure and services from other providers.

#### 4463 **Service Mock**

4464           A service mock is an entity that mimics some aspect of the performance of an  
4465           operational service without committing to the real world effects that the  
4466           operational service would produce.

4467 Mocks are discussed in detail in sections 5.4.4.3 and 5.4.4.4.

4468 After unit testing has demonstrated an adequate level of confidence in the service, the  
4469 testing must transition from the tightly controlled environment of the development  
4470 sandbox to an environment that more clearly resembles the operational SOA ecosystem  
4471 or, at a minimum, the intended enterprise. While sandbox testing will use simple mocks  
4472 of some aspects of the SOA environment, such as an interface to a security service  
4473 without the security service functionality, the dynamic nature of SOA makes a full  
4474 simulation infeasible to create or maintain. This is especially true when a new  
4475 composite service makes use of operational services provided by others. Thus, at  
4476 some point before testing is complete, the service will need to demonstrate its  
4477 functionality by using resources and dealing with conditions that only exist in the full  
4478 ecosystem or the intended enterprise. Some of these resources may still provide test  
4479 interfaces -- more on this below -- but the interfaces will be accessible using the SOA  
4480 environment and not just implemented for the sandbox.

4481 At this stage, the opacity of the service becomes important as the details of interacting  
4482 with the service now rely on correct use of the service interface and not knowledge of  
4483 the service internals. The workings of the service will only be observable through the  
4484 real world effects realized through service interactions and external indications that  
4485 conditions of use, such as user authentication, are satisfied. Monitoring the behavior of  
4486 the service will depend on service interfaces that expose internal monitoring or provide  
4487 required information to the SOA infrastructure monitoring function. The monitoring  
4488 required to test a new service is likely to have significant overlap with the monitoring the  
4489 SOA infrastructure includes to monitor its own health and to identify and isolate  
4490 behavior outside of acceptable bounds. This is exactly what is needed as part of  
4491 service testing, and it is reasonable to assume that the ecosystem transition includes  
4492 use of operational monitoring rather than solely dedicated monitoring for each service  
4493 being tested.

4494 Use of SOA monitoring resources during the explicit testing phase sets the stage for  
4495 monitoring and a level of continual testing throughout the service lifetime.

#### 4496 **5.4.4.2 Testing Traditional Dependencies vs. Service Interactions**

4497 A SOA service is not required to make use of other operational services beyond what  
4498 may be required for monitoring by the ecosystem infrastructure. The service can  
4499 implement hardcoded dependencies which have been tested in the development  
4500 sandbox through the use of dedicated mocks. While coordination may be required with

4501 real data sources during integration testing, the dependencies can be constrained to  
4502 things that can be tested in a more traditional manner. Policies can also be set to  
4503 restrict access to pre-approved users, and thus the question of unanticipated users and  
4504 unanticipated uses can be eliminated. Operational readiness can be defined in terms of  
4505 what can be proven in isolated testing. While all this may provide more confidence in  
4506 the service for its designed purpose, such a service will not fully participate in the  
4507 benefits or challenges of the ecosystem. This is akin to filling a swimming pool with sea  
4508 water and having someone in the pool say they are swimming in the ocean.

4509 In considering the testing needed for a fully participating service, consider the example  
4510 of a new composite service that combines the real world effects and complies with the  
4511 conditions of use of five existing operational services. The developer of the composite  
4512 service does not own any of the component services and has limited, if any, ability to  
4513 get the distributed owners to do any customization. The developer also is limited by the  
4514 principle of opacity to information comprising the service description, and does not know  
4515 internal details of the component services. The developer of the composite service  
4516 must use the component services as they exist as part of the SOA environment,  
4517 including what is provided to support testing by new users. This introduces  
4518 requirements for what is needed in the way of service mocks.

#### 4519 **5.4.4.3 Use of Service Mocks**

4520 Service mocks enables the tested service to respond to specific features of an  
4521 operational service that is being used as a component. It allows service testing to  
4522 proceed without needing access to or with only limited engagement with the component  
4523 service. Mocks can also mimic difficult to create situations for which it is desired to test  
4524 the new service response. For composite services using multiple component services,  
4525 mocks may be used in combination to function for any number of the components.  
4526 Note, when using service mocks, it is important to remember that it is not the  
4527 component service that is being tested (although anomalous behavior may be  
4528 uncovered during testing) but the use of the component in the new composite.

4529 Individual service mocks can emphasize different features of the component service  
4530 they represent but any given mock does not have to mimic all features. For example, a  
4531 mock of the service interface can echo a sent message and demonstrate the message  
4532 is reaching its intended destination. A mock could go further and parse the sent  
4533 message to demonstrate the message not only reached its destination but was  
4534 understood. As a final step, the mock could report back what actions would have been  
4535 taken by the component service and what real world effects would result. If the  
4536 response mimicked the operational response, functional testing could proceed as if the  
4537 real world effect actually occurred.

4538 There are numerous ways to provide mock functionality. The service mock could be a  
4539 simulation of the operational service and return simulated results in a realistic response  
4540 message or event notification. It is also possible for the operational service to act as its  
4541 own mock and simply not execute the commit stage of its functionality. The service  
4542 mock could use a combination of simulation and service action without commit to  
4543 generate a report of what would have occurred during the defined interaction with the  
4544 operational service.

4545 As the service proceeds through testing, mocks should be systematically replaced by  
4546 the component resources accessed through their operational interfaces. Before beta  
4547 testing begins, end-to-end testing, i.e. proceeding from the beginning of the service  
4548 interaction to the resulting real world results, should be accomplished using component  
4549 resources via their operational interfaces.

#### 4550 **5.4.4.4 Providers of Service Mocks**

4551 In traditional testing, it is often the test professionals who design and develop the  
4552 mocks, but in the distributed world of SOA, this may not be efficient or desirable.

4553 In the development sandbox, it is likely the new service developer or test professionals  
4554 working with the developer will create mocks adequate for unit testing. Given that most  
4555 of this testing is to verify the new service is performing as designed, it is not necessary  
4556 to have high fidelity models of other resources being accessed. In addition, given  
4557 opacity of SOA implementation, the developer of the new service may not have  
4558 sufficient detailed knowledge of a component service to build a detailed mock of the  
4559 component service functionality. Sharing existing mocks at this stage may be possible  
4560 but the mocks would need to be implemented in the sandbox, and for simple models it  
4561 is likely easier to build the mock from scratch.

4562 As testing begins its transition to the wider SOA environment, mocks may be available  
4563 as services. For existing resources, it is possible that an Open Source model could  
4564 evolve where service mocks of available functions can be catalogued and used during  
4565 initial interaction of the tested service and the operational environment. Widely used  
4566 functions may have numerous service mocks, some mimicking detailed conditions  
4567 within the SOA infrastructure. However, the Open Source model is less likely to be  
4568 sufficient for specialty services that are not widely used by a large consumer  
4569 community.

4570 The service developer is probably best qualified for also developing more detailed  
4571 service mocks or for mock modes of operational services. This implies that in addition  
4572 to their operational interfaces, services will routinely provide test interfaces to enable  
4573 service mocks to be used as services. As noted above, a new service developer  
4574 wanting to build a mock of component services is limited to the description provided by  
4575 the component service developer or owner. The description typically will detail real  
4576 world effects and conditions of use but will not provide implementation details, some of  
4577 which may be proprietary. Just as important in the SOA ecosystem, if it becomes  
4578 standard protocol for developers to create service mocks of their own services, a new  
4579 service developer is only responsible for building his own mocks and can expect other  
4580 mocks to be available from other developers. This reduces duplication of effort where  
4581 multiple developers would be trying to build the same mocks from the same insufficient  
4582 information. Finally, a service developer is probably best qualified to know when and  
4583 how a service mock should be updated to reflect modified functionality or message  
4584 exchange.

4585 It is also possible that testing organizations will evolve to provide high-fidelity test  
4586 harnesses for new services. The harnesses would allow new services to plug into a test  
4587 environment and would facilitate accessing mocks of component services. However, it  
4588 will remain a constant challenge for such organizations to capture evolving uses and

4589 characteristics of service interactions in the real SOA environment and maintain the  
4590 fidelity and accuracy of the test systems.

#### 4591 **5.4.4.5 Fundamental Questions for SOA Testing**

4592 In order for the transition to the SOA operational environment to proceed, it is necessary  
4593 to answer two fundamental questions:

- 4594 • Who provides what testing resources for the SOA operational environment, e.g.  
4595 mocks of interfaces, mocks of functionality, monitoring tools?
- 4596 • What testing needs to be accomplished before operational environment  
4597 resources can be accessed for further testing?

4598 The discussion in section 5.4.4.4 notes various levels of sophistication of service mocks  
4599 and different communities are likely to be responsible for different levels. Section  
4600 5.4.4.4 advocates a significant role for service developers, but there needs to be  
4601 community consensus that such mocks are needed and that service developers will  
4602 agree to fulfilling this role. There is also a need for consensus as to what tools should  
4603 be available as services from the SOA infrastructure.

4604 As for use of the service mocks and SOA environment monitoring services, practical  
4605 experience is needed upon which guidelines can be established for when a new service  
4606 has been adequately tested to proceed with a greater level of exposure with the SOA  
4607 environment. Malfunctioning services could cause serious problems if they cannot be  
4608 identified and isolated. On the other hand, without adequate testing under SOA  
4609 operational conditions, it is unlikely that problems can be uncovered and corrected  
4610 before they reach an operational stage.

4611 As noted in section 5.4.4.2, some of these questions can be avoided by restricting  
4612 services to more traditional use scenarios. However, such restriction will limit the  
4613 effectiveness of SOA use and the result will resemble the constraints of traditional  
4614 integration activities we are trying to move beyond.

#### 4615 **5.4.5 Architectural Implications for SOA Testing**

4616 The discussion of SOA Testing indicates numerous architectural implications on the  
4617 SOA ecosystem:

- 4618 • The distributed, boundary-less nature of the SOA ecosystem makes it  
4619 infeasible to create and maintain a single mock of the entire ecosystem to  
4620 support testing activities.
- 4621 • A standard suite of monitoring services needs to be defined, developed,  
4622 and maintained. This should be done in a manner consistent with the evolving  
4623 nature of the ecosystem.
- 4624 • Services should provide interfaces that support access in a test mode.
- 4625 • Testing resources must be described and their descriptions must be  
4626 catalogued in a manner that enables their discovery and access.
- 4627 • Guidelines for testing and ecosystem access need to be established and  
4628 the ecosystem must be able to enforce those guidelines asserted as policies.
- 4629 • Services should be available to support automated testing and regression  
4630 testing.

4631  
4632

- Services should be available to facilitate updating service description by anyone who has performed testing of a service.

---

## 4633 6 Conformance

4634 This Reference Architecture Framework is an abstract architectural description of  
4635 Service Oriented Architecture, which means that it is especially difficult to construct  
4636 tests for conformance to the architecture. In addition, conformance to an architectural  
4637 specification does not, by itself, guarantee any form of interoperability between multiple  
4638 implementations.

4639 However, it *is* possible to decide whether or not a given architecture is conformant to an  
4640 architectural description such as this one. In discussions of conformance we use the  
4641 term **target architecture** to identify the (typically concrete) architecture that may be  
4642 viewable as conforming to the abstract principles outlined in this document.

### 4643 Target Architecture

4644         A target architecture is an architectural description of a system that is intended to  
4645         be viewed as conforming to the SOA-RAF.

4646 While we cannot guarantee interoperability between target architectures (or more  
4647 specifically between applications and systems residing within the ecosystems of those  
4648 target architectures), interoperability between target architectures is promoted by  
4649 conformance to this Reference Architecture Framework as it reduces the semantic  
4650 impedance mismatch between the different ecosystems.

4651 The primary measure of conformance is whether given concepts as described in  
4652 document have corresponding concepts in the target architecture. Such a  
4653 correspondence **MUST** honor the relationships identified within this document for the  
4654 target architecture to be considered conforming.

4655 For example, in Section 3.1.3.1 we identify resource as a key concept. A resource is  
4656 associated with an owner and a number of identifiers. For a target architecture to  
4657 conform to the SOA-RAF, it must be possible to find corresponding concepts of  
4658 resource, identifier and owner within the target architecture: say *entity*, *token* and *user*.  
4659 Furthermore, the relationships between *entity*, *token* and *user* **MUST** mirror the  
4660 relationships between resource, identifier and owner appropriately.

4661 Clearly, such correspondence is simpler if the terminology within the target architecture  
4662 is identical to that in the SOA-RAF. But so long as the 'graph' of concepts and  
4663 relationships is consistent, that is all that is required for the target architecture to  
4664 conform to this Reference Architecture Framework.

4665 [EDITOR'S NOTE: The conformance section is not complete]

4666



---

4667 **A. Acknowledgements**

4668 The following individuals have participated in the work of the technical committee  
4669 responsible for creation of this specification and are gratefully acknowledged:

4670 **Participants:**

4671 Chris Bashioum, MITRE Corporation  
4672 Rex Brooks, Individual Member  
4673 Peter Brown, Individual Member  
4674 Scott Came, Search Group Inc.  
4675 Joseph Chiusano, Booz Allen Hamilton  
4676 Robert Ellinger, Northrop Grumman Corporation  
4677 David Ellis, Sandia National Laboratories  
4678 Jeff A. Estefan, Jet Propulsion Laboratory  
4679 Don Flinn, Individual Member  
4680 Anil John, Johns Hopkins University  
4681 Ken Laskey, MITRE Corporation  
4682 Boris Lublinsky, Nokia Corporation  
4683 Francis G. McCabe, Individual Member  
4684 Christopher McDaniels, USSTRATCOM  
4685 Tom Merkle, Lockheed Martin Corporation  
4686 Jyoti Namjoshi, Patni Computer Systems Ltd.  
4687 Duane Nickull, Adobe Inc.  
4688 James Odell, Associate  
4689 Michael Poulin, Fidelity Investments  
4690 Michael Stiefel, Associate  
4691 Danny Thornton, Northrop Grumman  
4692 Timothy Vibbert, Lockheed Martin Corporation  
4693 Robert Vitello, New York Dept. of Labor

4694 The committee would particularly like to underline the significant contributions made by  
4695 Rex Brooks, Jeff Estefan, Ken Laskey, Boris Lublinsky, Frank McCabe, Michael Poulin  
4696 and Danny Thornton

---

4697 **B. Index of Defined Terms**

4698 The first page number refers to the first use of the term. The second, where necessary,  
4699 refers to the page where the term is formally defined.

4700 Action

4701 Action Level Real World Effect

4702 Actor

4703 Architecture

4704 Architectural Description

4705 Authority

4706 Business Processes

4707 Capability

4708 Choreography

4709 Commitment

4710 Communicative Action

4711 Constitution

4712 Contract

4713 Delegate

4714 Description

4715 Endpoint

4716 Enterprise

4717 Governance

4718 Governance Framework

4719 Governance Processes

4720 Identifier

4721 Identity

4722 Joint Action

4723 Leadership

4724 Life-cycle manageability

4725 Logical Framework

4726 Management

4727 Management Policy

4728 Management Service

4729 Manageability Capability

4730 Message Exchange

4731 **Model**

4732 Obligation  
4733 Objective  
4734 Operations  
4735 Orchestration  
4736 Ownership  
4737 Ownership Boundary  
4738 Participant  
4739 Peer  
4740 Permission  
4741 Policy  
4742 Policy Conflict  
4743 Policy Conflict Resolution  
4744 Policy Constraint  
4745 Policy Decision  
4746 Policy Enforcement  
4747 Policy Framework  
4748 Policy Object  
4749 Policy Ontology  
4750 Policy Owner  
4751 Policy Subject  
4752 Presence  
4753 Private State  
4754 Protocol  
4755 Public Semantics  
4756 Qualification  
4757 Real World Effect  
4758 Regulation  
4759 Resource  
4760 Responsibility  
4761 Right  
4762 Risk  
4763 Role  
4764 Rule  
4765 Security  
4766 Semantic Engagement  
4767 Service Action  
4768 Service Consumer

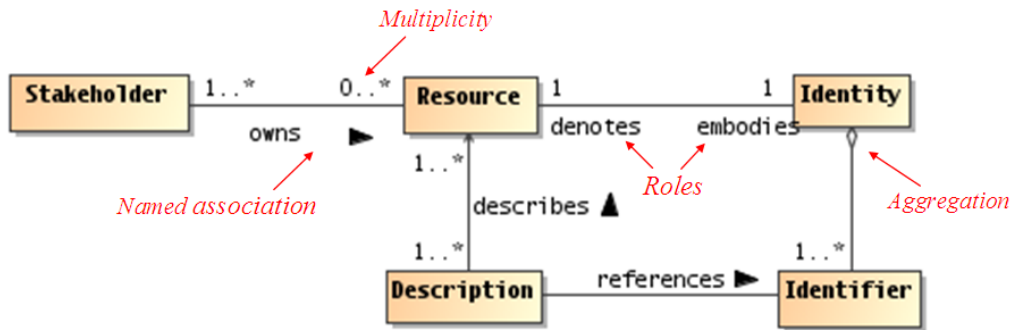
4769 Service Level Real World Effect  
4770 Service Mediator  
4771 Service Provider  
4772 Shared State  
4773 Skill  
4774 Social Structure  
4775 Stakeholder  
4776 State  
4777 System  
4778 System Stakeholder  
4779 Trust  
4780 View  
4781 Viewpoint

4782

## C. The Unified Modeling Language, UML

4783  
4784  
4785  
4786

**Error! Reference source not found.** illustrates an annotated example of a UML class diagram that is used to represent a visual model depiction of the Resources Model in the *Participation in a SOA Ecosystem* view (Section **Error! Reference source not found.**).



4787  
4788

Figure 45 Example UML class diagram—Resources.

4789  
4790  
4791  
4792  
4793  
4794

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

4795  
4796  
4797  
4798

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

4799  
4800  
4801

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

4802  
4803  
4804  
4805  
4806  
4807

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

4808  
4809  
4810  
4811  
4812  
4813  
4814

This is by no means a complete description of the semantics of all diagram elements that comprise a UML class diagram, but rather is intended to serve as an illustrative example for the reader. It should be noted that the SOA-RAF utilizes additional class diagram elements as well as other UML diagram types such as sequence diagrams and component diagrams. The reader who is unfamiliar with the UML is encouraged to review one or more of the many useful online resources and book publications available describing UML (see, for example, [www.uml.org](http://www.uml.org)).

---

## 4815 D. Critical Factors Analysis

4816 A critical factors analysis (CFA) is an analysis of the key properties of a project. A CFA  
4817 is analyzed in terms of the goals of the project, the critical factors that will lead to its  
4818 success and the measurable requirements of the project implementation that support  
4819 the goals of the project. CFA is particularly suitable for capturing quality attributes of a  
4820 project, often referred to as “non-functional” or “other-than-functional” requirements: for  
4821 example, security, scalability, wide-spread adoption, and so on. As such, CFA  
4822 complements rather than attempts to replace other requirements capture techniques.

### 4823 D.1 Goals

4824 A goal is an overall target that you are trying to reach with the project. Typically, goals  
4825 are hard to measure by themselves. Goals are often directed at the potential consumer  
4826 of the product rather than the technology developer.

### 4827 Critical Success Factors

4828 A critical success factor (CSF) is a property, sub-goal that directly supports a goal and  
4829 there is strong belief that without it the goal is unattainable. CSFs themselves are not  
4830 necessarily measurable in themselves.

### 4831 Requirements

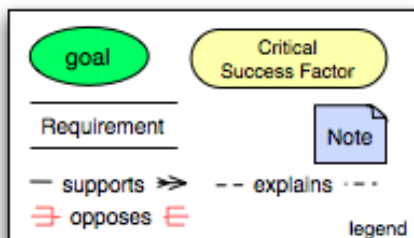
4832 A requirement is a specific measurable property that directly supports a CSF. The key  
4833 here is measurability: it should be possible to unambiguously determine if a requirement  
4834 has been met. While goals are typically directed at consumers of the specification,  
4835 requirements are focused on technical aspects of the specification.

### 4836 CFA Diagrams

4837 It can often be helpful to illustrate graphically the key concepts and relationships  
4838 between them. Such diagrams can act as effective indices into the written descriptions  
4839 of goals etc., but is not intended to replace the text.

4840 The legend:

4841



4842 illustrates the key elements of the graphical notation. Goals are written in round ovals,  
4843 critical success factors are written in round-ended rectangles and requirements are  
4844 written using open-ended rectangles. The arrows show whether a  
4845

4846 CSF/goal/requirement is supported by another element or opposed by it. This highlights  
4847 the potential for conflict in requirements.  
4848

4849

---

## E. Relationship to other SOA Open Standards

4850 The white paper “Navigating the SOA Open Standards Landscape Around Architecture”  
4851 issued jointly by OASIS, OMG, and The Open Group **[SOA-NAV]** was written to help  
4852 the SOA community at large navigate the myriad of overlapping technical products  
4853 produced by these organizations with specific emphasis on the “A” in SOA, i.e.,  
4854 Architecture.

4855 The white paper explains and positions standards for SOA reference models,  
4856 ontologies, reference architectures, maturity models, modeling languages, and  
4857 standards work on SOA governance. It outlines where the works are similar, highlights  
4858 the strengths of each body of work, and touches on how the work can be used together  
4859 in complementary ways. It is also meant as a guide to users for selecting those  
4860 specifications most appropriate for their needs.

4861 While the understanding of SOA and SOA Governance concepts provided by these  
4862 works is similar, the evolving standards are written from different perspectives. Each  
4863 specification supports a similar range of opportunity, but has provided different depths  
4864 of detail for the perspectives on which they focus. Although the definitions and  
4865 expressions may differ, there is agreement on the fundamental concepts of SOA and  
4866 SOA Governance.

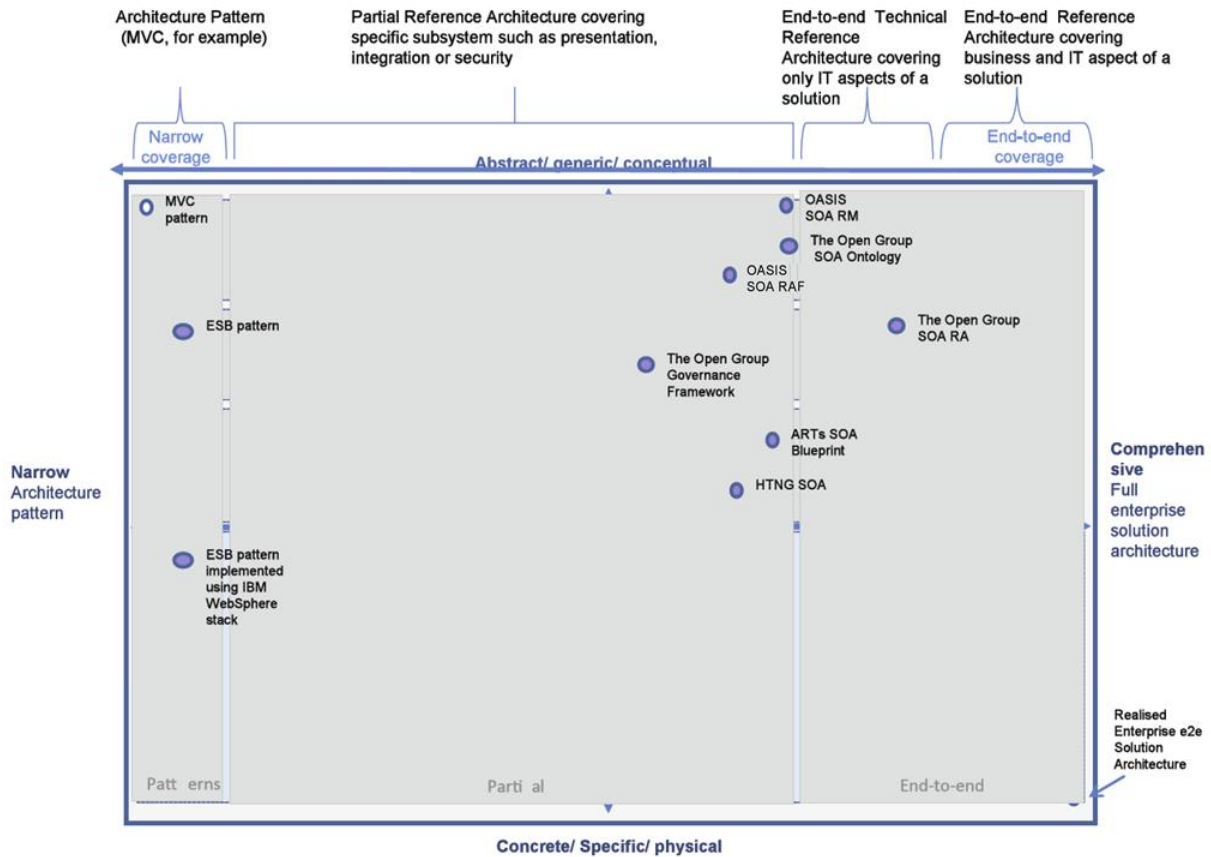
4867 The following is a summary taken from **[SOA-NAV]** of the positioning and guidance on  
4868 the specifications:

- 4869 • The OASIS Reference Model for SOA (SOA RM) is the most abstract of the  
4870 specifications positioned. It is used for understanding core SOA concepts
- 4871 • The Open Group SOA Ontology extends, refines, and formalizes some of the  
4872 core concepts of the SOA RM. It is used for understanding core SOA concepts  
4873 and facilitates a model-driven approach to SOA development.
- 4874 • The OASIS Reference Architecture Foundation for SOA (this document) is an  
4875 abstract, foundational reference architecture addressing a broader ecosystem  
4876 viewpoint for building and interacting within the SOA paradigm. It is used for  
4877 understanding different elements of SOA, the completeness of SOA architectures  
4878 and implementations, and considerations for reaching across ownership  
4879 boundaries where there is no single authoritative entity for SOA and SOA  
4880 governance.
- 4881 • The Open Group SOA Reference Architecture is a layered architecture from  
4882 consumer and provider perspective with cross cutting concerns describing these  
4883 architectural building blocks and principles that support the realizations of SOA. It  
4884 is used for understanding the different elements of SOA, deployment of SOA in  
4885 enterprise, basis for an industry or organizational reference architecture,  
4886 implication of architectural decisions, and positioning of vendor products in a  
4887 SOA context.
- 4888 • The Open Group SOA Governance Framework is a governance domain  
4889 reference model and method. It is for understanding SOA governance in  
4890 organizations. The OASIS Reference Architecture for SOA Foundation contains



- 4891 an abstract discussion of governance principles as applied to SOA across  
 4892 boundaries
- 4893 • The Open Group SOA Integration Maturity Model (OSIMM) is a means to assess  
 4894 an organization’s maturity within a broad SOA spectrum and define a roadmap  
 4895 for incremental adoption. It is used for understanding the level of SOA maturity in  
 4896 an organization
  - 4897 • The Object Management Group SoaML Specification supports services modeling  
 4898 UML extensions. It can be seen as an instantiation of a subset of the Open  
 4899 Group RA used for representing SOA artifacts in UML.

4900 Fortunately, there is a great deal of agreement on the foundational core concepts  
 4901 across the many independent specifications and standards for SOA. This could be best  
 4902 explained by broad and common experience of users of SOA and its maturity in the  
 4903 marketplace. It also provides assurance that investing in SOA-based business and IT  
 4904 transformation initiatives that incorporate and use these specifications and standards  
 4905 helps to mitigate risks that might compromise a successful SOA solution.



4906  
 4907 *Figure 46- SOA Reference Architecture Positioning (from “Navigating the SOA Open Standards Landscape Around*  
 4908 *Architecture, © OASIS, OMG, The Open Group).*