



---

# Reference Architecture Foundation for Service Oriented Architecture Version 1.0

**Draft 05 30 Apr 2011**

**Technical Committee:**

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

**Chair(s):**

Ken Laskey, MITRE Corporation

**Editor(s):**

Peter Brown, Individual Member, [peter@peterfbrown.com](mailto:peter@peterfbrown.com)  
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 Member, [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 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.

---

<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

**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 .....	9
1.1	Context for Reference Architecture for SOA .....	9
1.1.1	What is a Reference Architecture? .....	9
1.1.2	What is this Reference Architecture? .....	9
1.1.3	Relationship to the OASIS Reference Model for SOA .....	10
1.1.4	Relationship to other Reference Architectures .....	10
1.1.5	Expectations set by this Reference Architecture Foundation .....	10
1.2	Service Oriented Architecture – An Ecosystems Perspective .....	11
1.3	Viewpoints, Views and Models .....	11
1.3.1	ANSI/IEEE 1471-2000::ISO/IEC 42010-2007 .....	11
1.3.2	UML Modeling Notation .....	12
1.4	SOA-RAF Viewpoints .....	13
1.4.1	<i>Participation in a SOA Ecosystem Viewpoint</i> .....	13
1.4.2	<i>Realization of a SOA Ecosystem Viewpoint</i> .....	13
1.4.3	<i>Ownership in a SOA Ecosystem Viewpoint</i> .....	14
1.5	Terminology .....	14
1.5.1	Usage of Terms .....	14
1.6	References .....	14
1.6.1	Normative References .....	14
1.6.2	Non-Normative References .....	14
2	Architectural Goals and Principles .....	16
2.1	Goals and Critical Success Factors of the Reference Architecture Foundation .....	16
2.1.1	Goals .....	17
2.1.2	Critical Success Factors .....	17
2.2	Principles of this Reference Architecture Foundation .....	18
3	<i>Participation in a SOA Ecosystem view</i> .....	20
3.1	Social Structure in a SOA Ecosystem Model .....	21
3.1.1	Participants, Actors and Delegates .....	23
3.1.2	Roles in Social Structures .....	24
3.1.3	Resource and Ownership .....	27
3.1.4	Trust and Risk .....	28
	Assessing Trust and Risk .....	29
3.1.5	Policies and Contracts .....	30
3.1.6	Communication .....	31
3.1.7	Semantics and Semantic Engagement .....	32
3.2	Action in a SOA Ecosystem Model .....	32
3.2.1	Needs, Requirements and Capabilities .....	33
3.2.2	Services Reflecting Business .....	35
3.2.3	Action, Communication and Joint Action .....	35
3.2.4	State, Shared State and Real-World Effect .....	36
3.3	Architectural Implications .....	37
3.3.1	Social structures .....	37
3.3.2	Resource and Ownership .....	37

3.3.3	Policies and Contracts .....	38
3.3.4	Communications as a Means of Mediating Action .....	38
3.3.5	Semantics .....	38
3.3.6	Trust and Risk .....	38
3.3.7	Needs, Requirements and Capabilities .....	39
3.3.8	The Importance of Action .....	39
4	<i>Realization of a SOA Ecosystem view</i> .....	40
4.1	Service Description Model .....	40
4.1.1	The Model for Service Description .....	41
4.1.2	Use Of Service Description .....	49
4.1.3	Relationship to Other Description Models .....	54
4.1.4	Architectural Implications .....	54
4.2	Service Visibility Model .....	56
4.2.1	Visibility to Business .....	56
4.2.2	Visibility .....	57
4.2.3	Architectural Implications .....	62
4.3	Interacting with Services Model .....	62
4.3.1	Interaction Dependencies .....	63
4.3.2	Actions and Events .....	63
4.3.3	Message Exchange .....	64
4.3.4	Composition of Services .....	66
4.3.5	Architectural Implications of Interacting with Services .....	70
4.4	Policies and Contracts Model .....	71
4.4.1	Policy and Contract Representation .....	71
4.4.2	Policy and Contract Enforcement .....	72
4.4.3	Architectural Implications .....	73
5	<i>Ownership in a SOA Ecosystem View</i> .....	75
5.1	Governance Model .....	75
5.1.1	Understanding Governance .....	75
5.1.2	A Generic Model for Governance .....	76
5.1.3	Governance Applied to SOA .....	81
5.1.4	Architectural Implications of SOA Governance .....	85
5.2	Security Model .....	86
5.2.1	Secure Interaction Concepts .....	86
5.2.2	Where SOA Security is Different .....	89
5.2.3	Security Threats .....	89
5.2.4	Security Responses .....	90
5.2.5	Architectural Implications of SOA Security .....	91
5.3	Management Model .....	92
5.4	SOA Testing Model .....	101
5.4.1	Traditional Software Testing as Basis for SOA Testing .....	101
5.4.2	Testing and the SOA Ecosystem .....	102
5.4.3	Elements of SOA Testing .....	103
5.4.4	Testing SOA Services .....	107
5.4.5	Architectural Implications for SOA Testing .....	110

6	Conformance .....	112
A.	Acknowledgements .....	113
B.	Index of Defined Terms .....	114
C.	The Unified Modeling Language, UML.....	117
D.	Critical Factors Analysis .....	118
D.1	Goals .....	118
	Critical Success Factors.....	118
	Requirements .....	118
	CFA Diagrams .....	118
E.	Relationship to other SOA Open Standards .....	119

---

## Table of Figures

Figure 1 Critical Factors Analysis of the Reference Architecture .....	16
Figure 2 Model elements described in the Participation in a SOA Ecosystem view.....	21
Figure 3 Social Structure.....	22
Figure 4 Actors, Participants and Delegates.....	23
Figure 5 Role in Social Structures .....	24
Figure 6 Participant Roles in a Service .....	26
Figure 7 Resources.....	27
Figure 8 Willingness and Trust.....	29
Figure 9 Policies and Contracts .....	30
Figure 10 Relationship between Need, Requirement and Capability .....	34
Figure 11 Model Elements Described in the Realization of a SOA Ecosystem view .....	40
<i>Figure 12 General Description .....</i>	<i>42</i>
Figure 13 Representation of a Description .....	43
<i>Figure 14 Service Description .....</i>	<i>45</i>
Figure 15 Service Interface .....	46
<i>Figure 16 Service Functionality.....</i>	<i>47</i>
Figure 17 Model for Policies and Contracts as related to Service Participants .....	48
Figure 18 Action-Level and Service-Level Policies.....	48
Figure 19 Policies and Contracts, Metrics, and Compliance Records.....	49
Figure 20 Relationship Between Action and Service Description Components .....	50
Figure 21 Execution Context.....	53
Figure 22 Interaction Description .....	53
Figure 23 Visibility to Business .....	57
Figure 24 Mediated Service Awareness .....	59
Figure 25 Awareness in a SOA Ecosystem .....	60
Figure 26 Business, Description and Willingness .....	61
Figure 27 Service Reachability .....	61
Figure 27 Interaction dependencies.....	63
Figure 28 A "message" conveys either an action or an event.....	63
Figure 29 Fundamental SOA message exchange patterns (MEPs).....	65
Figure 30 Simple model of service composition. ....	67
Figure 31 Abstract example of orchestration of service-oriented business process. ....	68
Figure 32 Abstract example of choreography of service-oriented business collaboration. ....	69
Figure 34 Policies and Contracts .....	71
Figure 35 Model Elements Described in the Ownership in a SOA Ecosystem View.....	75
Figure 36 Motivating governance model.....	77
Figure 37 Setting up governance model .....	78
Figure 38 Carrying out governance model.....	79
Figure 39 Ensuring governance compliance model.....	80
Figure 40 Relationship among types of governance .....	82
Figure 41 Confidentiality and Integrity .....	87

Figure 42 Authentication ..... 87  
Figure 43 Authorization ..... 88  
Figure 1- SOA Reference Architecture Positioning (from “Navigating the SOA Open Standards  
Landscape Around Architecture, © OASIS, OMG, The Open Group)..... 120



---

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

### 1.1.2 What is this Reference Architecture?

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

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

---

<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.

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

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

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

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

### 60 **1.1.3 Relationship to the OASIS Reference Model for SOA**

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

65 The SOA-RAF goes further in that it shows how SOA-based systems can be realized – albeit in an  
66 abstract way. As noted above, SOA-based systems are better thought of as dynamic systems rather than  
67 stand-alone software products. Consequently, how they are used and managed is at least as important  
68 architecturally as how they are constructed.

### 69 **1.1.4 Relationship to other Reference Architectures**

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

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

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

### 82 **1.1.5 Expectations set by this Reference Architecture Foundation**

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

## 88 1.2 Service Oriented Architecture – An Ecosystems Perspective

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

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

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

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

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

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

## 112 1.3 Viewpoints, Views and Models

### 113 1.3.1 ANSI/IEEE 1471-2000::ISO/IEC 42010-2007

114 The SOA-RAF uses and follows the IEEE “Recommended Practice for Architectural Description of  
115 Software-Intensive Systems” [ANSI/IEEE 1471] and [ISO/IEC 42010]. An architectural description  
116 conforming to this standard must include the following six (6) elements:

- 117 1. Architectural description identification, version, and overview information
- 118 2. Identification of the system stakeholders and their concerns judged to be relevant to the  
119 architecture
- 120 3. Specifications of each viewpoint that has been selected to organize the representation of the  
121 architecture and the rationale for those selections
- 122 4. One or more architectural views
- 123 5. A record of all known inconsistencies among the architectural description’s required constituents
- 124 6. A rationale for selection of the architecture (in particular, showing how the architecture supports  
125 the identified stakeholders’ concerns).

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

#### 127 **Architecture**

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

---

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

130 **Architectural Description**

131 A collection of products that document the architecture.

132 **System**

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

134 **System Stakeholder**

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

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

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

144 **View**

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

146 **Viewpoint**

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

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

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

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

161 **Model**

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

163 All architectural models used in a particular view are developed using the methods established by the  
164 architectural viewpoint associated with that view. An architectural model may participate in more than one  
165 view but a view must conform to a single viewpoint.

166 **1.3.2 UML Modeling Notation**

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

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

174 **1.4 SOA-RAF Viewpoints**

175 The RAF uses three views that conform to three viewpoints: *Participation in a SOA Ecosystem*,  
 176 *Realization of a SOA Ecosystem*, and *Ownership in a SOA Ecosystem*. There is a one-to-one  
 177 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, securing, and testing 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, security, and testing 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

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

179 **1.4.1 Participation in a SOA Ecosystem Viewpoint**

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

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

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

190 **1.4.2 Realization of a SOA Ecosystem Viewpoint**

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

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

### 197 1.4.3 Ownership in a SOA Ecosystem Viewpoint

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

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

## 206 1.5 Terminology

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

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

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

### 214 1.5.1 Usage of Terms

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

## 220 1.6 References

### 221 1.6.1 Normative References

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

225 **[ISO/IEC 42010]** International Organization for Standardization and International Electrotechnical  
226 Commission, *System and software engineering — Recommended practice for*  
227 *architectural description of software-intensive systems*, July 15, 2007.

228 **[RFC2119]** S. Bradner, *Key words for use in RFCs to Indicate Requirement Levels*,  
229 <http://www.ietf.org/rfc/rfc2119.txt>, IETF RFC 2119, March 1997.

230 **[SOA-RM]** OASIS Standard, "Reference Model for Service Oriented Architecture 1.0, 12  
231 October 2006. <http://docs.oasis-open.org/soa-rm/v1.0/soa-rm.pdf>

232 **[UML 2]** *Unified Modeling Language: Superstructure*, Ver. 2.1.1, OMG Adopted  
233 Specification, OMG document formal/2007-02-05, Object Management Group,  
234 Needham, MA, February 5, 2007.

235 **[WA]** Architecture of the World Wide Web, W3C, 2004. <http://www.w3.org/TR/webarch>.

236 **[WSA]** David Booth, et al., "Web Services Architecture", W3C Working Group Note,  
237 World Wide Web Consortium (W3C) (Massachusetts Institute of Technology,  
238 European Research Consortium for Informatics and Mathematics, Keio  
239 University), February, 2004. [http://www.w3.org/TR/2004/NOTE-ws-arch-  
240 20040211/](http://www.w3.org/TR/2004/NOTE-ws-arch-20040211/)

### 241 1.6.2 Non-Normative References

242 **[BLOOMBERG/SCHMELZER]** Jason Bloomberg and Ronald Schmelzer, *Service Orient or Be*  
243 *Doomed!*, John Wiley & Sons: Hoboken, NJ, 2006.

244 [COX] D. E. Cox and H. Kreger, "Management of the service-oriented architecture life  
245 cycle," "IBM Systems Journal" ""44"", No. 4, 709-726, 2005

246 [ERA] A. Fattah, "Enterprise Reference Architecture," paper presented at 22<sup>nd</sup>  
247 Enterprise Architecture Practitioners Conference, London, UK, April 2009.

248 [ITU-T Rec. X.700 | ISO/IEC 10746-3:1996(E)] Information processing systems—Open Systems  
249 Interconnection—Basic Reference Model—Part 4: Management Framework",  
250 International Telecommunication Union, International Organization for  
251 Standardization and International Electrotechnical Commission, Geneva,  
252 Switzerland, 1989.

253 [NEWCOMER/LOMOW] Eric Newcomer and Greg Lomow, *Understanding SOA with Web Services*,  
254 Addison-Wesley: Upper Saddle River, NJ, 2005.

255 [OECD] Organization for Economic Cooperation and Development, Directorate for  
256 Financial, Fiscal and Enterprise Affairs, OECD Principles of Corporate  
257 Governance, SG/CG(99) 5 and 219, April 1999.

258 [TOGAF v9] *The Open Group Architecture Framework (TOGAF) Version 9 Enterprise Edition*,  
259 The Open Group, Doc Number: G091, February 2009.

260 [WEILL] Harvard Business School Press, IT Governance: How Top Performers Manage  
261 IT Decision Rights for Superior Results, Peter Weill and Jeanne W. Ross, 2004

262 [DAMIANOU] Nicodemos C. Damianou , Thesis - A Policy Framework for Management of  
263 Distributed Systems, University of London, Department of Computing, 2002.

264 [LEVESON] Nancy G. Leveson, *Safeware: System Safety and Computers*, Addison-Wesley  
265 Professional, Addison-Wesley Publishing Company, Inc.: Boston, pg. 181, 1995.

266 [STEEL/NAGAPPAN/LAI] Christopher Steel and Ramesh Nagappan and Ray Lai, *core Security  
267 Patterns: Best Practices and Strategies for J2EE, Web Services and Identity  
268 Management*, Prentice Hall: 2005

269 [ISO/IEC 27002] International Organization for Standardization and International Electrotechnical  
270 Commission, *Information technology -- Security techniques – Code of practice  
271 for information security management*, 2007

272 [SOA NAV] Heather Kreger and Jeff Estefan (Eds.), "Navigating the SOA Open Standards  
273 Landscape Around Architecture," Joint Paper, The Open Group, OASIS, and  
274 OMG, July 2009.  
275 <http://www.opengroup.org/projects/soa/uploads/40/20044/W096.pdf>

276 **2 Architectural Goals and Principles**

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

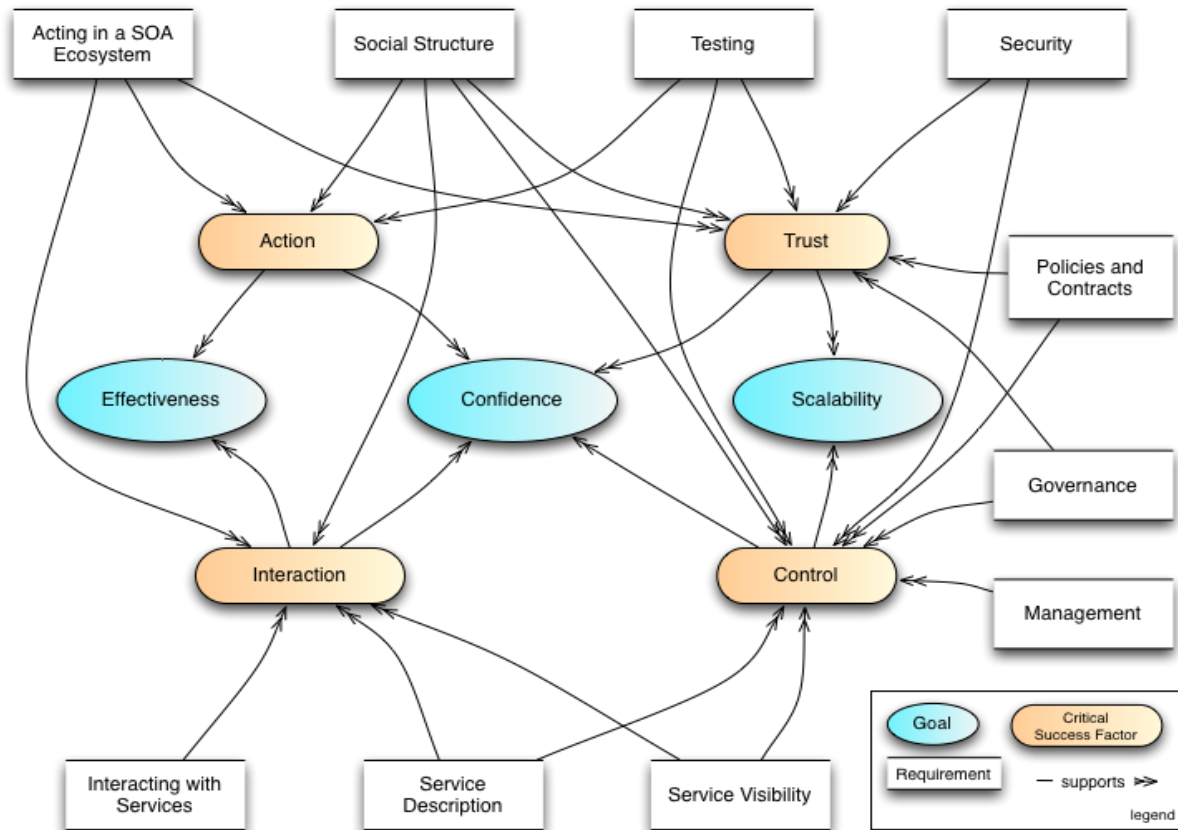
279 **2.1 Goals and Critical Success Factors of the Reference Architecture Foundation**

280 There are three principal goals:

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

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

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



293  
294 *Figure 1 Critical Factors Analysis of the Reference Architecture*



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

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

## 302 **2.1.1 Goals**

### 303 **2.1.1.1 Effectiveness**

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

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

### 310 **2.1.1.2 Confidence**

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

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

### 317 **2.1.1.3 Scalability**

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

## 322 **2.1.2 Critical Success Factors**

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

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

### 329 **2.1.2.1 Action**

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

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

### 336 **2.1.2.2 Trust**

337 The viability of a SOA ecosystem depends on participants being able to effectively measure the  
338 trustworthiness of the system and of participants. Trust is a private assessment of a participant's belief in  
339 the integrity and reliability of the SOA ecosystem (see Section 3.1.4).

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

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

### 348 **2.1.2.3 Interaction**

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

### 352 **2.1.2.4 Control**

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

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

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

## 364 **2.2 Principles of this Reference Architecture Foundation**

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

### 366 **Technology Neutrality**

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

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

375 **Implications:** The Reference Architecture Foundation must be technology neutral, meaning that we  
376 assume that technology will continue to evolve, and that over the lifetime of this  
377 architecture that multiple, potentially competing technologies will co-exist. Another  
378 immediate implication of technology independence is that greater effort on the part of  
379 architects and other decision makers to construct systems based on this architecture is  
380 needed.

### 381 **Parsimony**

382 **Statement:** Parsimony refers to economy of design, avoiding complexity where possible and  
383 minimizing the number of components and relationships needed.

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

387 **Implications:** Parsimoniously designed systems tend to have fewer but better targeted features.

388 **Distinction of Concerns**

389 Statement: Distinction of Concerns refers to the ability to cleanly identify and separate out the  
390 concerns of specific stakeholders in such a way that it is possible to create architectural  
391 models that reflect those stakeholders' viewpoint. In this way, an individual stakeholder or  
392 a set of stakeholders that share common concerns only see those models that directly  
393 address their respective areas of interest.

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

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

408 **Applicability**

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

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

414 Implications: The Reference Architecture Foundation needs to be relevant to the problem of matching  
415 needs and capabilities under disparate domains of ownership; to the concepts of "Intranet  
416 SOA" (SOA within the enterprise) as well as "Internet SOA" (SOA outside the enterprise);  
417 to the concept of "Extranet SOA" (SOA within the extended enterprise, i.e., SOA with  
418 suppliers and trading partners); and finally, to "net-centric SOA" or "Internet-ready SOA."

---

### 419 3 Participation in a SOA Ecosystem view

420  
421  
422  
423  
424  
425  
426  
427  
428  
429  
430

#### 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

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

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

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

450 The subsections below expand on the completely abstract reference model by identifying more fully and  
451 with more specificity what challenges need to be addressed in order to successfully accomplish SOA.  
452 Although this section does not provide a specific recipe, it does identify the important things that need to  
453 be thought about and resolved within an ecosystem context.

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

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

---

<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)

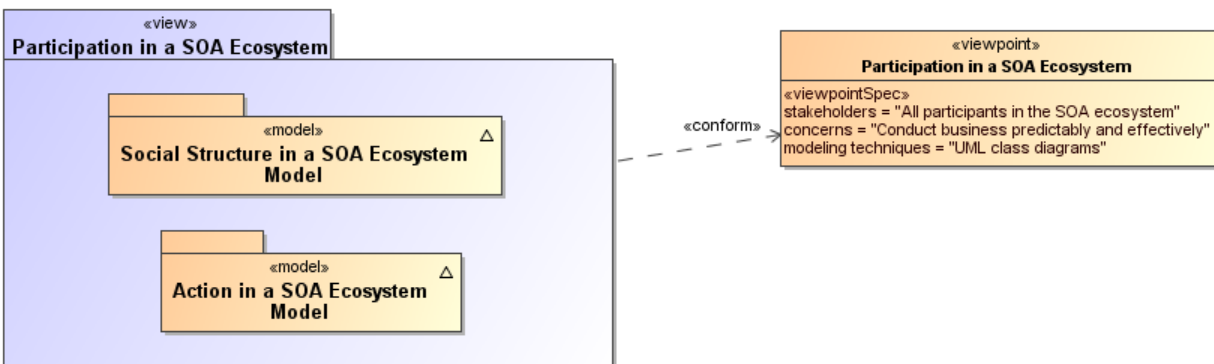
465 This section describes how a SOA-based system behaves when participants may be in different  
466 organizations, with different rules and expectations, and assumes that the primary motivation for  
467 participants to interact with each other is to achieve **objectives** –to get things done.

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

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

480 The main models in this view are:

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

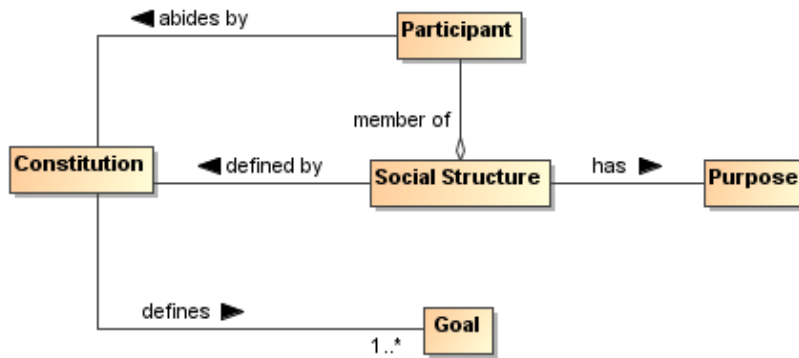


487  
488 Figure 2 Model elements described in the Participation in a SOA Ecosystem view

### 489 3.1 Social Structure in a SOA Ecosystem Model

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

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



500  
501 *Figure 3 Social Structure*

502 **Social Structure**

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

505 A social structure represents a collection of participants, but a collection that is brought together for a  
506 purpose. There may be a large number of different kinds of relationships between participants in a social  
507 structure. The organizing principle for these relationships is the social structure’s purpose.

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

511 A social structure has a purpose – the overarching reason for which it exists. All social structures are  
512 established with implied or explicitly defined purpose. The purpose is usually reflected in specific goals  
513 laid down in the social structure’s constitution or other ‘charter’.

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

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

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

523 **Enterprise**

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

527 **Peer Group**

528 A peer group is a social structure withno discernable leadership structure, that may or may not  
529 have internally established goals, but is identifiable as the locus of interaction between participants  
530 with individual goals and who are considered peers of one another.

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

---

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

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

540 Social structures involved in a particular interaction are not always explicitly identified. For example, when  
541 a customer buys a book over the Internet, the social structure that determines the validity of the  
542 transaction is often the legal framework of the region associated with the book vendor. Such legal  
543 jurisdiction qualification is typically buried in the fine print of the service description.

#### 544 **Constitution**


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

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

### 553 **3.1.1 Participants, Actors and Delegates**

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

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



562  
563 *Figure 4 Actors, Participants and Delegates*

#### 564 **Stakeholder**

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

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

#### 572 **Actor**

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

574 **Participant**

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

577 **Delegate**

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

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

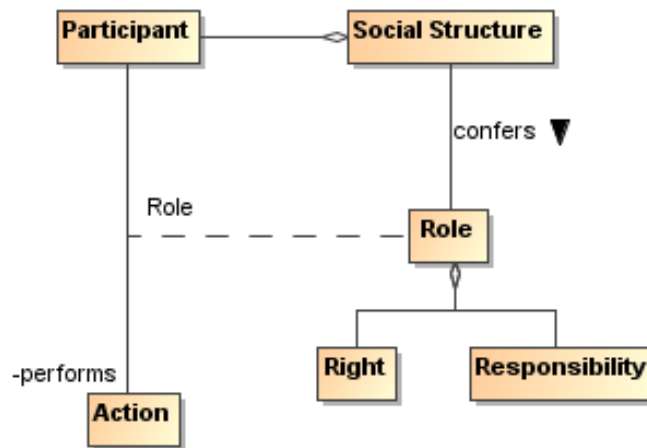
580 Many stakeholders and actors operate in a SOA ecosystem, including software agents that permit people  
581 to offer, and interact with, services; delegates that represent the interests of other participants; or security  
582 agents charged with managing the security of the ecosystem. Note that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

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

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

592 **3.1.2 Roles in Social Structures**

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



597  
598 *Figure 5 Role in Social Structures*

599 **Role**

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

602 A role is not immutable and is often time-bound. A participant can have one or more roles concurrently  
603 and may change them over time and in different contexts, even over the course of a particular interaction.  
604 One participant with appropriate authority in the social structure may formally *designate a role* for another

---

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



605 participant, with associated rights and responsibilities, and that authority may even qualify a period during  
606 which the designated role may be valid.

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

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

### 618 **Authority**

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

### 620 **Right**

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

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

### 625 **Responsibility**

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

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

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

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

### 644 **3.1.2.1 Service Roles**

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

650

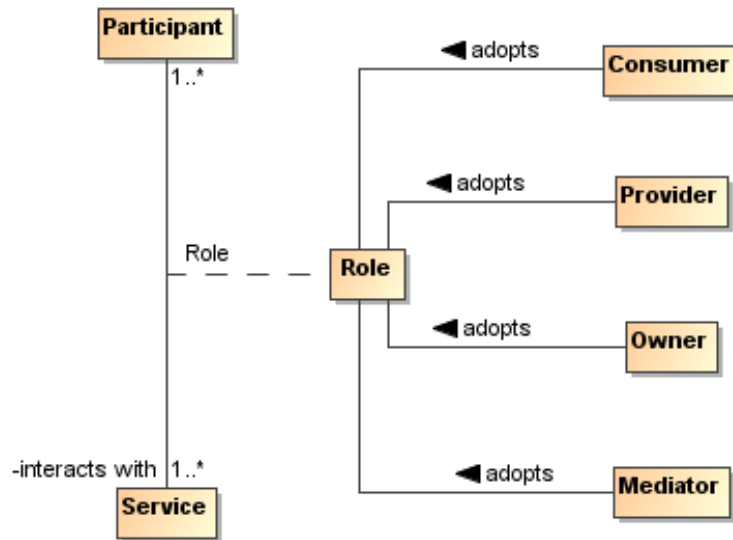


Figure 6 Participant Roles in a Service

651

652

653 **Provider**

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

655 **Consumer**

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

658 **Mediator**

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

661 **Owner**

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


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

667 The roles of service provider and service consumer are often seen as symmetrical, which is also not  
668 entirely correct. A consumer tends to express a 'Need' in non-formal terms: "I want to buy that book". The  
669 type of 'Need' that a service is intended to fulfill has to be formalized and encapsulated by designers and  
670 developers as a 'Requirement'. This Requirement should then be reflected in the target service, as a  
671 'Capability' that, when accessed via a service, delivers a 'Real World Effect' to an arbitrary user: "The  
672 chosen book is ordered for the user" It thus satisfies the need that has been defined for an archetypal  
673 user. Specific and particular users may not experience a need exactly as captured by the service: "I don't  
674 want to pay that much for the book", "I wanted an eBook version", etc. There can therefore be a process  
675 of implicit and explicit negotiation between the user and the service, aimed at finding a 'best fit' between  
676 the user's specific need and the capabilities of the service that are available and consistent with the  
677 service provider's offering. This process may continue up until the point that the user is able to accept  
678 what is on offer as being the best fit and finally 'invokes' the service. 'Execution context' has thus been  
679 established. This is explored in more detail later on. Service mediation by a participant can take many  
680 forms and may invoke and use other services in order to fulfill such mediation. For example, it might use a  
681 service registry in order to identify possible service partners; or, in our book-buying example, it might  
682 provide a price comparison service, suggest alternative suppliers, different language editions or delivery  
683 options.

### 684 3.1.3 Resource and Ownership

#### 685 3.1.3.1 Resource

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



688  
689  
690 *Figure 7 Resources*

#### 691 Resource

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

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

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

#### 699 Identifier

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

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

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

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

#### 714 3.1.3.2 Ownership

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

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

719 **Ownership**

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

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

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

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

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

741 **Ownership Boundary**

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

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

751 **3.1.4 Trust and Risk**

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

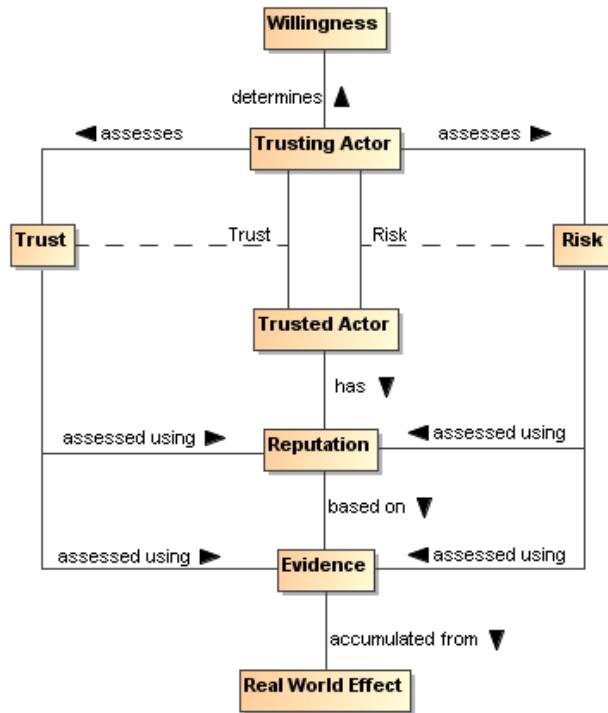


Figure 8 Willingness and Trust

753

754

## 755 Willingness

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

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

## 761 Trust

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

## 764 Risk

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

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

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

## 776 Assessing Trust and Risk

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

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

783 Assessment of trust is rarely binary. An actor is not completely trusted or untrusted. There is typically  
784 some degree of uncertainty in the accuracy or completeness of the evidence or the assessment.

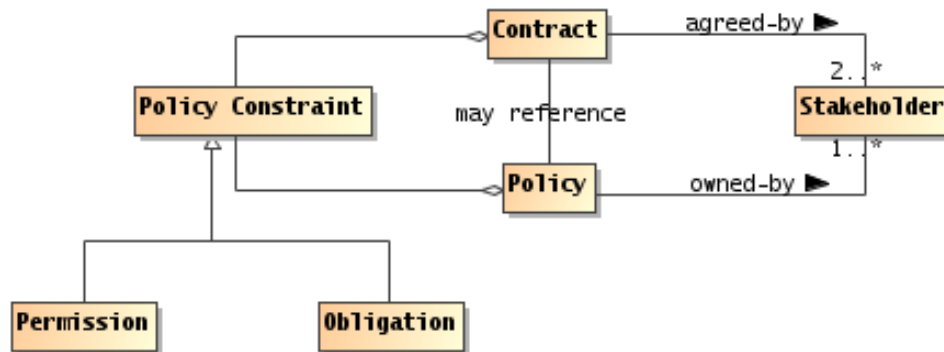
785 Similarly, there may be uncertainty in the amount and potential consequences of risk.

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

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

### 798 3.1.5 Policies and Contracts

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



804  
805 *Figure 9 Policies and Contracts*

#### 806 **Policy**

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

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

#### 814 **Contract**

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

818 A service provider's policy may become a service provider/consumer contract when a service consumer  
819 agrees to the provider's policy. That agreement may be formal, or may be informal. If a consumer's

820 policy and a providers policy are mutually exclusive, then some form of negotiation or mediation to  
821 resolve the mutual exclusion before the service consumer/provider interaction can occur.

822 Both policies and contracts imply a desire to see constraints respected and enforced. Policies are owned  
823 by individual (or aggregate) stakeholders, and contracts are owned by the parties to the contract; these  
824 stakeholders are responsible for ensuring that any constraints in the policy or contract are enforced –  
825 although, of course, the actual enforcement may be delegated to a different mechanism. A contract does  
826 not necessarily oblige the contracting parties to act (for example to use a service) but it does constraint  
827 how they act if and when action covered by the contract occurs (for example, when a service is invoked  
828 and used).

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

### 830 **Permission**

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

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

### 837 **Obligation**

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

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

843 An obligation can also be a requirement to to *maintain* a given state. This may range from a requirement  
844 to maintain a minimum balance on an account to a requirement that a service provider 'remember' that a  
845 particular service consumer is logged in.

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

## 849 **3.1.6 Communication**

### 850 **Communication**

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

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

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

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

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

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

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

### 872 **3.1.7 Semantics and Semantic Engagement**

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

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

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

#### 888 **Semantic Engagement**

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

891 Different actors have differing capabilities and requirements for understanding assertions. This is true for  
892 both human and non-human actors. For example, a purchase order process does not require that a  
893 message forwarding agent 'understand' the purchase order, but a processing agent does need to  
894 'understand' the purchase order in order to know what to do with the order once received.

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

### 900 **3.2 Action in a SOA Ecosystem Model**

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

907 An action model is put forth a-priori by the service provider, and is effectively a promise by the service  
908 provider that the actions identified in the action model and invoked consistent with the process model will  
909 result in the described real world effect. Action model is basically a description of the actions that the  
910 service is willing to do on behalf of another. They should be associated with a real-world effect. The  
911 potential service consumer is interested in accessing or acquiring the real-world effect, and the action  
912 model identifies the actions that the service consumer will have to be a party to in order to access or  
913 generate the real-world effect.

---

<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.



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

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

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

929 All of these things are important to getting to the core of participants’ interest in a SOA ecosystem: the  
930 ability to leverage resources or capabilities to achieve a desired outcome, and in particular where those  
931 resources or capabilities do not belong to them or are beyond their direct control. i.e., that are outside of  
932 their ownership boundary.

933 In order to use such resources, participants must be able to identify their own needs in the form of  
934 requirements, identify and compose into a business solution those resources or capabilities that will meet  
935 their needs, and engage in joint action – the coordinated set of actions that participants pursue in order to  
936 achieve measurable results in furtherance of their goals.

937 In order to act in a way that is appropriate and consistent both to their own goals, objectives and policies,  
938 and those of others, participants must also communicate with each other.

939 A key aspect of joint action revolves around the trust that both parties must exhibit in order to participate  
940 in the joint action. The willingness to act and a mutual understanding of both the information exchanged  
941 and the expected results is the particular focus of Sections **Error! Reference source not found.**6 and  
942 3.1.7.

### 943 **3.2.1 Needs, Requirements and Capabilities**

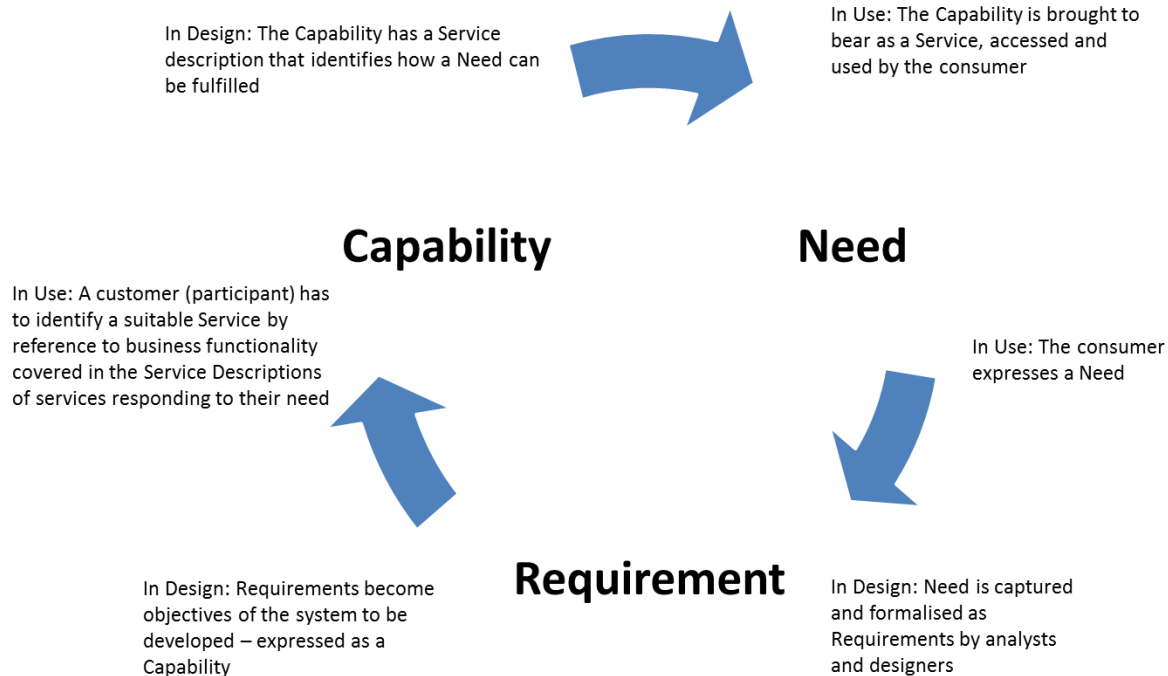
944 Participants in a SOA ecosystem often need other participants to *do* something, leveraging a capability  
945 that they do not themselves possess. For example, a customer requiring a book may call upon a service  
946 provider to deliver the book. Likewise, the service provider needs the customer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

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

### 955 **Business functionality**

956 Business functionality is a defined set of business-aligned tasks that provide recognizable  
 957 business value to 'consumer' stakeholders and possibly others in the SOA ecosystem.  
 958

Figure 10 Relationship between Need, Requirement and Capability



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

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

971 **Need**

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

975 **Requirement**

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

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

981 **Capability**

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

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

### 986 3.2.2 Services Reflecting Business

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

#### 991 **Business solution**

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

#### 994 **Composability**

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

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

1000 To achieve composability, capabilities must be identified that serve as building blocks for business  
1001 solutions. In a SOA ecosystem, these building blocks are captured as services representing well-defined  
1002 business functions, operating under well-defined policies and other constraints, and generating well-  
1003 defined real world effects. These service building blocks should be relatively stable so as not to force  
1004 repeated changes in the compositions that utilize them, but should also embody SOA attributes that  
1005 readily support creating compositions that can be varied to reflect changing circumstances.

1006 The SOA paradigm emphasizes both composition of services and opacity of how a given service is  
1007 implemented. With respect to opacity, the SOA-RM states that the service could carry out its described  
1008 functionality through one or more automated and/or manual processes that in turn could invoke other  
1009 available services.

1010 Any composition can itself be made available as a service and the details of the business functionality,  
1011 conditions of use, and effects are among the information documented in its service description.

1012 For services to be useful as composable building blocks in the SOA ecosystem, the services should,  
1013 whenever possible, deliver capability that is applicable to multiple needs. Simply providing a Web Service  
1014 interface for an existing IT artifact does not, in general, create opportunities for sharing business  
1015 functions. Furthermore, the use of tools to auto-generate service software interfaces will not guarantee  
1016 services that can effectively be used within compositions if the underlying code represents programming  
1017 constructs rather than business functions. In such cases, services that tightly reflect the software details  
1018 will be as brittle to change as the underlying code and will not exhibit the undefined but intuitive  
1019 characteristic of loose coupling.

### 1020 3.2.3 Action, Communication and Joint Action

1021 In general terms, entities act in order to achieve their goals. However, the form of action that is of most  
1022 interest within a SOA ecosystem is that involving interaction across ownership boundaries (between more  
1023 than one actor) – **joint action**.

#### 1024 3.2.3.1 Action and Actors

##### 1025 **Action**

1026 An action is the application of intent to cause an effect.

1027 The aspect of action that distinguishes it from mere force or accident is that someone *intends* that the  
1028 action achieves a desired objective or effect. This definition of action is very general. In the case of SOA,  
1029 we are mostly concerned with actions that take place within a system and have specific effects on the

1030 SOA ecosystem – what we call **Real World Effects**. The actual real world effect of an action, however,  
1031 may go beyond the intended effect.

1032 Objectives refer to real world effects that participants believe are achievable by a specific action or set of  
1033 actions that deliver appropriate changes in shared state. In contrast, a goal is not expressed in terms of  
1034 specific action but rather in terms of desired end state.

1035 For example, someone may wish to have enough light to read a book. In order to satisfy that goal, the  
1036 reader walks over to flip a light switch. The *objective* is to change the state of the light bulb, by turning on  
1037 the lamp, whereas the *goal* is to be able to read. The *real world effect* is more light being available to  
1038 enable the person to read.

1039 While an effect is any measurable change resulting from an action, a SOA ecosystem is concerned more  
1040 specifically with real world effects.

#### 1041 **Real World Effect**

1042 A real world effect is a measurable change to the shared state of pertinent entities, relevant to  
1043 and experienced by specific stakeholders of an ecosystem.

1044 This implies measurable change in the overall state of the SOA ecosystem. [In practice, however, it is  
1045 specific state changes of certain entities that are relevant to particular participants that constitute the real  
1046 world effect as experienced by those participants.](#)

#### 1047 **3.2.3.2 Communication and Joint Actions**

1048 In this Reference Architecture Foundation, we are concerned with two levels of activity: as communication  
1049 and as participants engaged in joint actions to use and offer services.

1050 In order for multiple actors to participate in a joint action, they must each act according to their role within  
1051 the joint action. This is achieved through communication and messaging.

1052 Communication – the formulation, transmission, receipt and interpretation of messages – is the  
1053 foundation of all joint actions within the SOA ecosystem, given the inherent separation – often across  
1054 ownership boundaries – of actors in the system.

1055 Communication between actors requires that they play the roles of ‘sender’ or ‘receiver’ of messages as  
1056 appropriate to a particular action – although it is not necessarily required that they both be active  
1057 simultaneously.

1058 An actor sends a message in order to communicate with other actors. The communication itself is often  
1059 not intended as part of the desired real world effect but rather includes messages that seek to establish,  
1060 manage, monitor, report on, and guide the joint action throughout its execution.

1061 Like communication, joint action usually involves different actors. However, joint action – resulting from  
1062 the deliberate actions undertaken by different actors – *intentionally* impacts shared state within the  
1063 system leading to real world effects.

#### 1064 **Joint Action**

1065 Joint action is the coordinated set of actions involving the efforts of two or more actors to achieve  
1066 an effect.

1067 Note that the effect of a joint action is *not* always equivalent to one or more effects of the individual  
1068 actions of the participating actors, i.e., it may be more than the sum of the parts.

1069 Different viewpoints lead to either communication or joint action as being considered most important. For  
1070 example, from the viewpoint of ecosystem security, the integrity of the communications may be dominant;  
1071 from the viewpoint of ecosystem governance, the integrity of the joint action may be dominant.

#### 1072 **3.2.4 State, Shared State and Real-World Effect**

##### 1073 **State**

1074 State is the condition of an entity at a particular time.

1075 State is characterized by a set of facts that is true of the entity. In principle, the total state of an entity (or  
1076 the world as a whole) is unbounded. In practice, we are concerned only with a subset of the State of an  
1077 entity that is measurable and useful in a given context.

1078 For example, the total state of a lightbulb includes the temperature of the filament of the bulb. It also  
1079 includes a great deal of other state – the composition of the glass, the dirt that is on the bulb’s surface  
1080 and so on. However, an actor may be primarily interested in whether the bulb is ‘on’ or ‘off’ and not on the  
1081 amount of dirt accumulated. That actor’s characterization of the state of the bulb reduces to the fact: ‘bulb  
1082 is now on’.

1083 In a SOA ecosystem, there is a distinction between the set of facts about an entity that only that entity can  
1084 access – the so-called Private State – and the set of facts that may be accessible to other actors in the  
1085 SOA-based system – the public or Shared State.

#### 1086 **Private State**

1087 The private state is that part of of an entity’s state that is knowable by, and accessible to, only  
1088 that entity.

#### 1089 **Shared State**

1090 Shared state is that part of an entity’s state that is knowable by, and may be accessible to, other  
1091 actors.

1092 Note that shared state does not imply that the state *is* accessible to *all* actors. It simply refers to that  
1093 subset of state that *may* be accessed by *other* actors. Generally this will be the case when actors need to  
1094 participate in joint actions.

1095 It is the aggregation of the shared states of pertinent entities that constitutes the desired effect of a joint  
1096 action. Thus the change to this shared state is what is experienced in the wider ecosystem as a real world  
1097 effect

### 1098 **3.3 Architectural Implications**

#### 1099 **3.3.1 Social structures**

1100 A SOA ecosystem’s participants are organized into various forms of social structure. Not all social  
1101 structures are hierarchical: a SOA ecosystem should be able to incorporate peer-to-peer forms of  
1102 organization as well as hierarchic structures. In addition, it should be possible to identify and manage any  
1103 constitutional agreements that define the social structures present in a SOA ecosystem.

- 1104 • Different social structures have different rules of engagement
  - 1105 ○ Techniques for expressing constitutions are important
- 1106 • social structures have roles and members
  - 1107 ○ Techniques for identifying, managing members of social structures
  - 1108 ○ Techniques for describing roles and role adoption
- 1109 • social structures may be complex
  - 1110 ○ Child social structures’ constitutions depend on their parent constitutions
- 1111 • Social structures overlap and interact
  - 1112 ○ A given actor may be member of multiple social structures
  - 1113 ○ Social structures may be associated with different jurisdictions
  - 1114 ○ Social structures may involved in disputes with one another
    - 1115 ■ Requiring conflict resolution
  - 1116 ○ Social structures inform and limit the “kinds” of governance that can be effectively  
1117 deployed

#### 1118 **3.3.2 Resource and Ownership**

1119 Communication about and between, visibility into, and leveraging of resources requires the unambiguous  
1120 identification of those resources. Ensuring unambiguous identities implies

- 1121 • Mechanism for assigning and guaranteeing uniqueness of globally unique identifiers

- 1122 • Identifying the extent of the enterprise over which the identifier needs to be understandable and  
1123 unique  
1124 • Mechanism and framework for ensuring the long-livedness of identifiers (i.e., they cannot just  
1125 change arbitrarily)

### 1126 3.3.3 Policies and Contracts

- 1127 • Policies are constraints  
1128 ○ It is necessary to be able to express required policies  
1129 ○ It is necessary to be able to enforce the constraints  
1130 ○ It is necessary to manage potentially large numbers of policies  
1131 • Policies have owners  
1132 ○ The right to establish policies is an aspect of the social structure.  
1133 • Policies may not be consistent with one another  
1134 ○ Policy conflict resolution techniques  
1135 • Agreements are constraints agreed to  
1136 ○ Contracts often need to be enforced by mechanisms of the social structure

### 1137 3.3.4 Communications as a Means of Mediating Action

1138 Using message exchange for mediating action implies

- 1139 • Ensuring correct identification of the structure of messages:  
1140 ○ Identifying the syntax of the message;  
1141 ○ Identifying the vocabularies used in the communication  
1142 ○ Identifying the higher-level structure such as the illocutionary form of the communication  
1143 • A principal objective of communication is to mediate action  
1144 ○ Messages convey actions and events  
1145 ○ Receiving a message is an action, but is not the same action as the action conveyed by  
1146 the message  
1147 ○ Actions are associated with objectives of the actors involved  
1148 ■ Explicit representation of objectives may facilitate automated processing of  
1149 messages  
1150 ○ An actor agreeing to adopt an objective becomes responsible for that objective

### 1151 3.3.5 Semantics

1152 Semantics is pervasive in a SOA ecosystem. There are many forms of utterance that are relevant to the  
1153 ecosystem: apart from communicated content there are policy statements, goals, purposes, descriptions,  
1154 and agreements which are all forms of utterance.

1155 The operation of the SOA ecosystem is significantly enhanced if

- 1156 • A careful distinction is made between public semantics and private semantics. In particular, it  
1157 MUST be possible for actors to process content such as communications, descriptions and  
1158 policies solely on the basis of the public semantics of those utterances.  
1159 • A well founded semantics ensures that any assertions that are essential to the operator of the  
1160 ecosystem (such as policy statements, and descriptions) have carefully chosen written  
1161 expressions and associated decision procedures.  
1162 • The role of vocabularies as a focal point for multiple actors to be able to understand each other is  
1163 critical. While no two actors can fully share their interpretation of elements of vocabularies,  
1164 ensuring that they do understand the public meaning of vocabularies' elements is essential.

### 1165 3.3.6 Trust and Risk

1166 In traditional systems, the balance between trust and risk is achieved by severely restricting interactions  
1167 and by controlling the participants of a system.

1168 It is important that actors are able to explicitly reason about both trust and risk in order to effectively  
1169 participate in a SOA ecosystem. The more open and public the SOA ecosystem is, the more important it

1170 is for actors to be able to reason about their participation.

### 1171 **3.3.7 Needs, Requirements and Capabilities**

1172 In the process of capturing needs as requirements, and the subsequent requirements decomposition and  
1173 allocation processes need to be informed by capabilities that already exist.

- 1174 • Architecture needs to
- 1175 ○ Take into account existing capabilities available as services

### 1176 **3.3.8 The Importance of Action**

1177 Participants participate in a SOA ecosystem in order to get their needs met. This involves action; both  
1178 individual actions and joint actions.

1179 Any architectural realization of a SOA ecosystem should address:

- 1180 • How actions are modeled:
- 1181 ○ Identifying the performer or agent of the action;
- 1182 ○ the target of the action; and the
- 1183 ○ verb of the action.

1184 Any explicit models of joint action should take into account

- 1185 • The choreography that defines the joint action.
- 1186 • The potential for multiple joint actions to be layered on top of each other

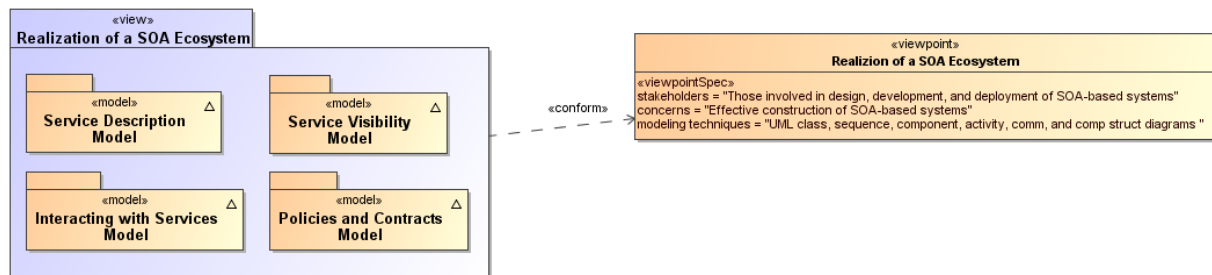
## 1187 4 Realization of a SOA Ecosystem view

1188  
1189  
1190

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

1191 The *Realization of a SOA Ecosystem* view focuses on the infrastructure elements that are needed in  
1192 order to support the discovery and interaction with services. The key questions asked are "What are  
1193 services, what support is needed and how are they realized?"

1194 The models in this view include the Service Description Model, the Service Visibility Model, the Interacting  
1195 with Services Model, and the Policies and Contracts Model.



1196  
1197

Figure 11 Model Elements Described in the Realization of a SOA Ecosystem view

1198 The Service Description Model informs the participants of what services exist and the conditions under  
1199 which these can be used. Some of those conditions follow from policies and agreements on policy that  
1200 flow from the Policies and Contracts Model. The information in the service description as augmented by  
1201 details of policy provides the basis for visibility as defined in the SOA Reference Model and captured in  
1202 the Service Visibility Model. Finally, the process by which services as described are used under the  
1203 defined conditions and agreements is described in the Interacting with Services Model.

### 1204 4.1 Service Description Model

1205 A service description is an artifact, usually document-based, that defines or references the information  
1206 needed to use, deploy, manage and otherwise control a service. This includes not only the information  
1207 and behavior models associated with a service to define the service interface but also includes  
1208 information needed to decide whether the service is appropriate for the current needs of the service  
1209 consumer. Thus, the service description will also include information such as service reachability, service  
1210 functionality, and the policies and contracts associated with a service.

1211 A service description artifact may be a single document or it may be an interlinked set of documents. For  
1212 the purposes of this model, differences in representation are to be ignored, but the implications of a "web  
1213 of documents" is discussed later in this section.

1214 There are several points to note regarding the following discussion of service description:

- 1215 • The Reference Model states that one of the hallmarks of SOA is the large amount of associated  
1216 description. The model presented below focuses on the description of services but it is equally  
1217 important to consider the descriptions of the consumer, other participants, and needed resources  
1218 other than services.
- 1219 • Descriptions are inherently incomplete but may be determined as *sufficient* when it is possible for the  
1220 participants to access and use the described services based only on the descriptions provided. This  
1221 means that, at one end of the spectrum, a description along the lines of "*That service on that*  
1222 *machine*" may be sufficient for the intended audience. On the other extreme, a service description  
1223 with a machine-process-able description of the semantics of its operations and real world effects may  
1224 be required for services accessed via automated service discovery and planning systems.
- 1225 • Descriptions come with context, i.e. a given description comprises information needed to adequately  
1226 support the context. For example, a list of items can define a version of a service, but for many



1227 contexts an indicated version number is sufficient without the detailed list. The current model focuses  
1228 on the description needed by a service consumer to understand what the service does, under what  
1229 conditions the service will do it, how well does the service do it, and what steps are needed by the  
1230 consumer to initiate and complete a service interaction. Such information also enables the service  
1231 provider to clearly specify what is being provided and the intended conditions of use.

1232 • Descriptions change over time as, for example, the ingredients and nutrition information for food  
1233 labeling continues to evolve. A requirement for transparency of transactions may require additional  
1234 description for those associated contexts.

1235 • Description always proceeds from a basis of what is considered "common knowledge". This may be  
1236 social conventions that are commonly expected or possibly codified in law. It is impossible to describe  
1237 everything and it can be expected that a mechanism as far reaching as SOA will also connect entities  
1238 where there is inconsistent "common" knowledge.

1239 • Descriptions will become the collection point of information related to a service or any other resource,  
1240 but it is not necessarily the originating point or the motivation for generating this information. In  
1241 particular, given a SOA service as the access to an underlying capability, the service may point to  
1242 some of the capability's previously generated description, e.g. a service providing access to a data  
1243 store may reference update records that indicate the freshness of the data.

1244 • Descriptions of the provider and consumer are the essential building blocks for establishing the  
1245 execution context of an interaction.

1246 These points emphasize that there is no one "right" description for all contexts and for all time. Several  
1247 descriptions for the same subject may exist at the same time, and this emphasizes the importance of the  
1248 description referencing source material maintained by that material's owner rather than having multiple  
1249 copies that become out of synch and inconsistent.

1250 It may also prove useful for a description assembled for one context to cross-reference description  
1251 assembled for another context as a way of referencing ancillary information without overburdening any  
1252 single description. Rather than a single artifact, description can be thought of as a web of documents that  
1253 enhance the total available description.

1254 This Reference Architecture Foundation uses the term service description for consistency with the  
1255 concept defined in the Reference Model. Some SOA literature treats the idea of a "service contract" as  
1256 equivalent to service description. In the SOA-RAF, the term service description is preferred. Replacing  
1257 service description with service contract implies just one side of the interaction is governing and misses  
1258 the point that a single set of policies identified by a service description may lead to numerous contracts,  
1259 i.e. service level agreements, leveraging the same description.

## 1260 **4.1.1 The Model for Service Description**

1261 *Figure 12* shows Service Description as a subclass of the general Description class, where Description is  
1262 a subclass of the resource class as defined in Section 3.1.5.1. In addition, each resource is assumed to  
1263 have a description. The following section discusses the relationships among elements of general  
1264 description and the subsequent sections focus on service description itself. Other descriptions, such as  
1265 those of participants, are important to SOA but are not individually elaborated in this document.

### 1266 **4.1.1.1 Elements Common to General Description**

1267 The general Description class is composed of a number of elements that are expected to be common  
1268 among all specialized descriptions supporting a service oriented architecture. A registry often contains a  
1269 subset of the description instance, where the chosen subset is identified as that which facilitates mediated  
1270 discovery. Additional information contained in a more complete description may be needed to initiate and  
1271 continue interaction.

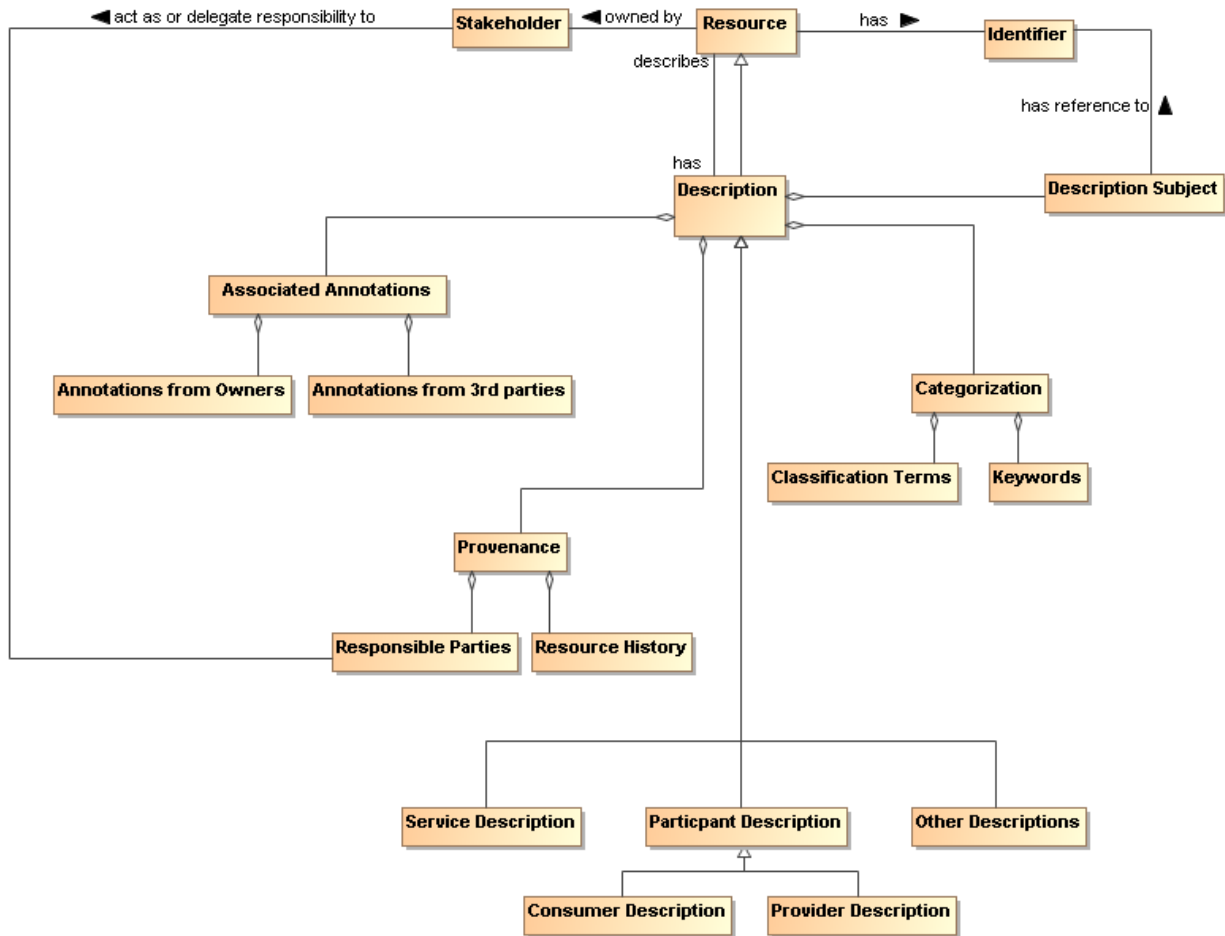


Figure 12 General Description

1272  
1273

#### 1274 4.1.1.1.1 Description Subject

1275 The subject of a description is a resource. The value assigned to the Description Subject class may be of  
1276 any form that provides understanding of what constitutes the resource, but it is often in human-readable  
1277 text. The Description Subject MUST also reference the Identifier of the resource it describes so it can  
1278 unambiguously identify the subject of each description instance.

1279 As a resource, Description also has an identifier with a unique value for each description instance. The  
1280 description instance provides vital information needed to both establish visibility of the resource and to  
1281 support its use in the execution context for the associated interaction. The identifier of the description  
1282 instance allows the description itself to be referenced for discussion, access, or reuse of its content.

#### 1283 4.1.1.1.2 Provenance

1284 While the resource Identifier provides the means to know which subject and subject description are being  
1285 considered, Provenance as related to the Description class provides information that reflects on the  
1286 quality or usability of the subject. Provenance specifically identifies the entity (human, defined role,  
1287 organization, ...) that assumes responsibility for the resource being described and tracks historic  
1288 information that establishes a context for understanding what the resource provides and how it has  
1289 changed over time. Responsibilities may be directly assumed by the stakeholder who owns a resource or  
1290 the Owner may designate Responsible Parties for the various aspects of maintaining the resource and  
1291 provisioning it for use by others. There may be more than one entity identified under Responsible Parties;  
1292 for example, one entity may be responsible for code maintenance while another is responsible for  
1293 provisioning of the executable code. The historical aspects may also have multiple entries, such as when

1294 and how data was collected and when and how it was subsequently processed, and as with other  
1295 elements of description, may provide links to other assets maintained by the resource owner.

#### 1296 4.1.1.1.3 Keywords and Classification Terms

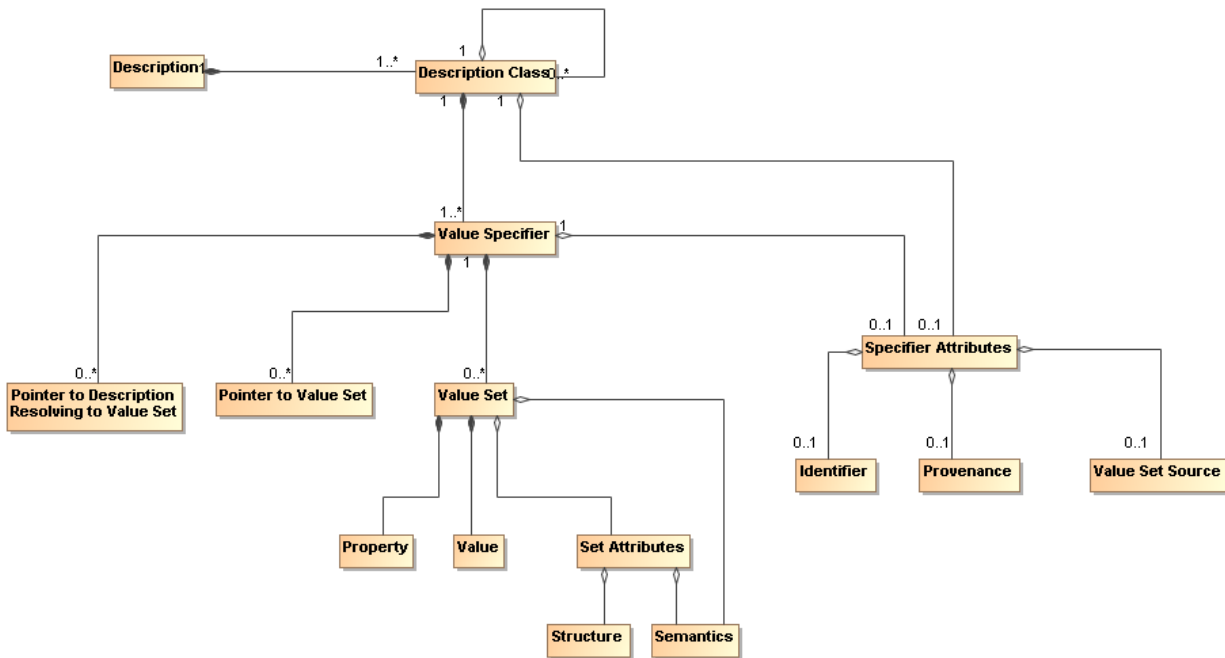
1297 A traditional element of description has been to associate the resource being described with predefined  
1298 keywords or classification taxonomies that derive from referenceable formal definitions and vocabularies.  
1299 This Reference Architecture Foundation does not prescribe which vocabularies or taxonomies may be  
1300 referenced, nor does it limit the number of keywords or classifications that may be associated with the  
1301 resource. It does, however, state that a normative definition SHOULD be referenced, whether that be a  
1302 representation in a formal ontology language, a pointer to an online dictionary, or any other accessible  
1303 source. See Section 4.1.1.2 for further discussion on associating semantics with assigned values.

#### 1304 4.1.1.1.4 Associated Annotations

1305 The general description instance may also reference associated documentation that is in addition to that  
1306 considered necessary in this model. For example, the owner of a service may have documentation on  
1307 best practices for using the service. Alternately, a third party may certify a service based on their own  
1308 criteria and certification process; this may be vital information to other prospective consumers if they were  
1309 willing to accept the certification in lieu of having to perform another certification themselves. Note, while  
1310 the examples of Associated Documentation presented here are related to services, the concept applies  
1311 equally to description of other entities.

#### 1312 4.1.1.2 Assigning Values to Description Instances

1313



1314  
1315 *Figure 13 Representation of a Description*

1316 Figure 12 shows the template for a general description but individual description instances depend on the  
1317 ability to associate meaningful values with the identified elements. Figure 13 shows a model for a  
1318 collection of information that provides for value assignment and traceability for both the value meaning  
1319 and the source of a value. The model is not meant to replace existing or future schema or other  
1320 structures that have or will be defined for specific implementations, but it is meant as guidance for the  
1321 information such structures need to capture to generate sufficient description. It is expected that tools will  
1322 be developed to assist the user in populating description and auto-filling many of these fields, and in that  
1323 context, this model provides guidance to the tool developers.

1324 In Figure 13 each class has an associated value specifier or is made up of components that will  
1325 eventually resolve to a value specifier. For example, Description has several components, one of which is  
1326 Categorization, which would have an associated a value specifier.

1327 A value specifier consists of

- 1328 • a collection of value sets with associated property-value pairs, pointers to such value sets, or pointers  
1329 to descriptions that eventually resolve to value sets that describe the component; and
- 1330 • attributes that qualify the value specifier and the value sets it contains.

1331 The qualifying attributes for the value specifier include

- 1332 • an optional identifier that would allow the value set to be defined, accessed, and reused elsewhere;
- 1333 • provenance information that identifies the party (individual, role, or organization) that has  
1334 responsibility for assigning the value sets to any description component;
- 1335 • an optional source of the value set, if appropriate and meaningful, e.g. if a particular data source is  
1336 mandated.

1337 If the value specifier is contained within a higher-level component, (such as Service Description  
1338 containing Service Functionality), the component may inherit values from the attributes from its container.

1339 Note, provenance as a qualifying attribute of a value specifier is different from provenance as part of an  
1340 instance of Description. Provenance for a service identifies those who own and are responsible for the  
1341 service, as described in Section 3. Provenance for a value specifier identifies who is responsible for  
1342 choosing and assigning values to the value sets that comprise the value specifier. It is assumed that  
1343 granularity at the value specifier level is sufficient and provenance is not required for each value set.

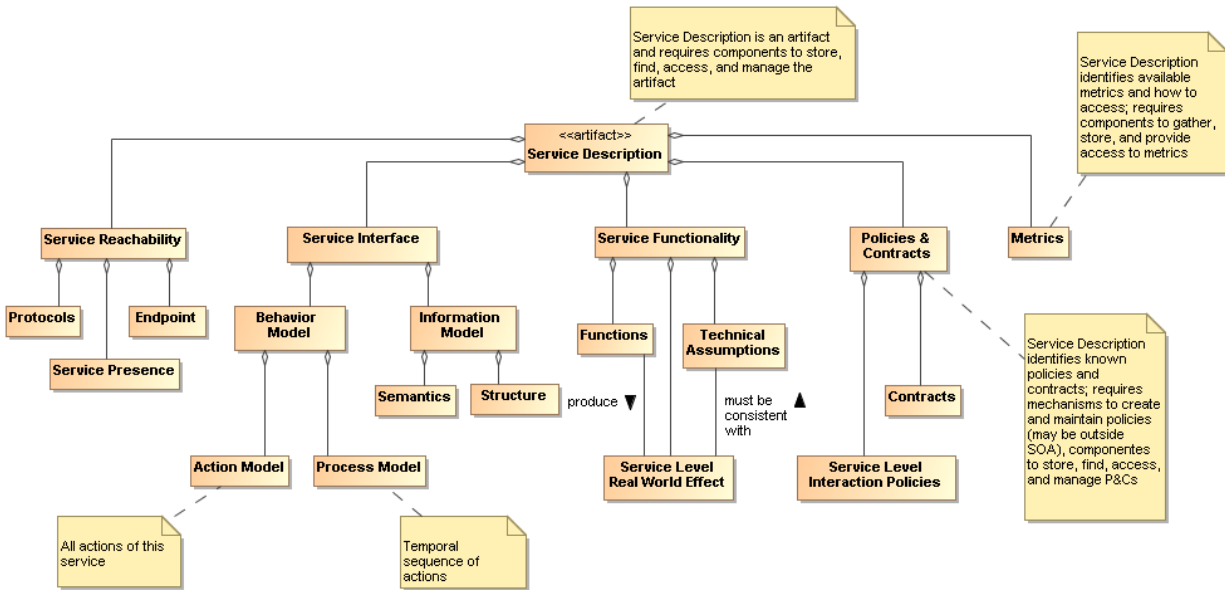
1344 The value set also has attributes that define its structure and semantics.

- 1345 • The semantics of the value set property should be associated with a semantic context conveying the  
1346 meaning of the property within the execution context, where the semantic context could vary from a  
1347 free text definition to a formal ontology.
- 1348 • For numeric values, the structure would provide the numeric format of the value and the “semantics”  
1349 would be conveyed by a dimensional unit with an identifier to an authoritative source defining the  
1350 dimensional unit and preferred mechanisms for its conversion to other dimensional units of like type.
- 1351 • For nonnumeric values, the structure would provide the data structure for the value representation  
1352 and the semantics would be an associated semantic model.
- 1353 • For pointers, architectural guidelines would define the preferred addressing scheme.

1354 The value specifier may indicate a default semantic model for its component value sets and the individual  
1355 value sets may provide an override.

1356 The property-value pair construct is introduced for the value set to emphasize the need to identify  
1357 unambiguously both what is being specified and what is a consistent associated value. The further  
1358 qualifying of Structure and Semantics in the Set Attributes allows for flexibility in defining the form of the  
1359 associated values.

1360 **4.1.1.3 Model Elements Specific to Service Description**



1361  
1362 *Figure 14 Service Description*

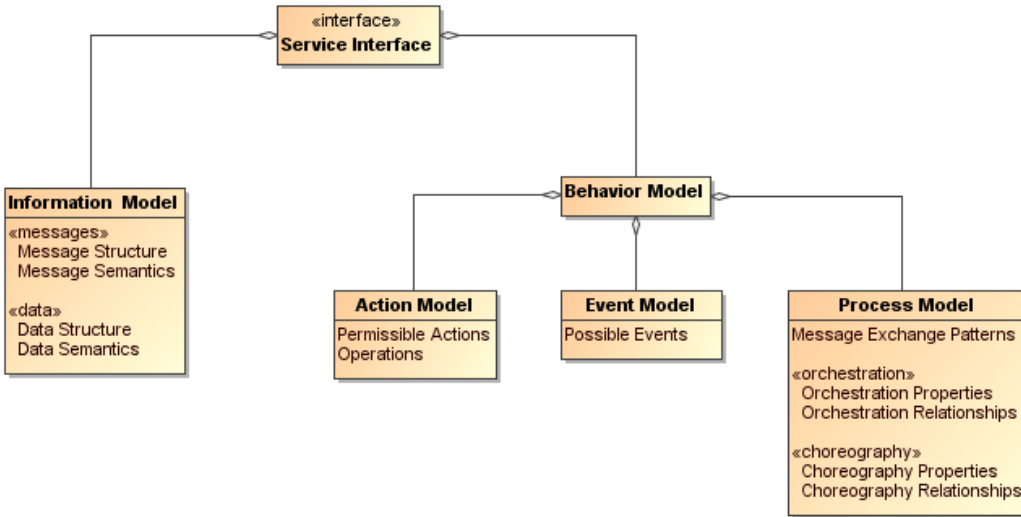
1363 The major elements for the Service Description subclass follow directly from the areas discussed in the  
1364 Reference Model. Here, we discuss the detail shown in *Figure 14* and the purpose served by each element  
1365 of service description.

1366 Note, the intent in the subsections that follow is to describe how a particular element, such as the service  
1367 interface, is reflected in the service description, not to elaborate on the details of that element.

1368 **4.1.1.3.1 Service Interface**

1369 As noted in the Reference Model, the service interface is the means for interacting with a service. For the  
1370 SOA-RAF and as shown in Section 4.3 the service interface will support an exchange of messages,  
1371 where

- 1372 • the message conforms to a referenceable message exchange pattern (MEP),
- 1373 • the message payload conforms to the structure and semantics of the indicated information model,
- 1374 • the messages are used to denote events or actions against the service, where the actions are  
1375 specified in the action model and any required sequencing of actions is specified in the process  
1376 model.



1377  
1378 *Figure 15 Service Interface*

1379 Note we distinguish the structure and semantics of the message from that of the underlying protocol that  
1380 conveys the message. The message structure may include nested structures that are independently  
1381 defined, such as an enclosing envelope structure and an enclosed data structure.  
1382 These aspects of messages are discussed in more detail in Section 4.3

1383 **4.1.1.3.2 Service Reachability**

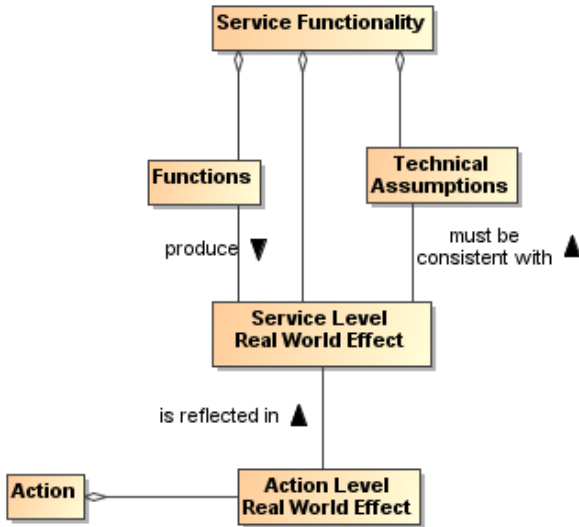
1384 Service reachability, as modeled in Section 4.2.2.3 enables service participants to locate and interact with  
1385 one another. To support service reachability, the service description should indicate the endpoints to  
1386 which a service consumer can direct messages to invoke actions and the protocol to be used for  
1387 message exchange using that endpoint.

1388 As applied in general to an action, the endpoint is the conceptual location where one applies an action;  
1389 with respect to service description, it is the actual address where a message is sent.

1390 In addition, the service description should provide information on collected metrics for service presence;  
1391 see Section 4.1.1.3.4 for the discussion of metrics as part of service description.

1392 **4.1.1.3.3 Service Functionality**

1393 While the service interface and service reachability are concerned with the mechanics of using a service,  
1394 service functionality and performance metrics (discussed in Section 4.1.1.3.4) describe what can be  
1395 expected when interacting with a service. Service Functionality, shown in *Figure 14* as part of the overall  
1396 Service Description model and extended in *Figure 16*, is an unambiguous expression of service function(s)  
1397 and the real world effects of invoking the function. The Functions represent business activities in some  
1398 domain that produce the desired real world effects.



1399  
1400  
1401

Figure 16 Service Functionality

1402 The Service Functionality may also be constrained by Technical Assumptions that underlie the effects  
1403 that can result. Technical assumptions are defined as domain specific restrictions and may express  
1404 underlying physical limitations, such as flow speeds must be below sonic velocity or disk access that  
1405 cannot be faster than the maximum for its host drive. Technical assumptions are related to the underlying  
1406 capability accessed by the service. In any case, the real world effects must be consistent with the  
1407 Technical Assumptions.

1408 In Figure 14 and Figure 16, we specifically refer to Service Level and Action Level real world effects.

1409 **Service Level Real World Effect**

1410 A service level real world effect is a specific change in shared state or information returned as a  
1411 result of interacting with a service.

1412 **Action Level Real World Effect**

1413 An action level real world effect is a specific change in shared state or information returned as a  
1414 result of performing a specific action against a service.

1415 Service description describes the service as a whole while the component aspects should contribute to  
1416 that whole. Thus, while individual Actions may contribute to the real world effects to be realized from  
1417 interaction with the service, there would be a serious disconnect for Actions to contribute real world  
1418 effects that could not consistently be reflected in the Service Level Real World Effects and thus the  
1419 Service Functionality. The relationship to Action Level Real World Effects and the implications on  
1420 defining the scope of a service are discussed in Section 4.1.2.1.

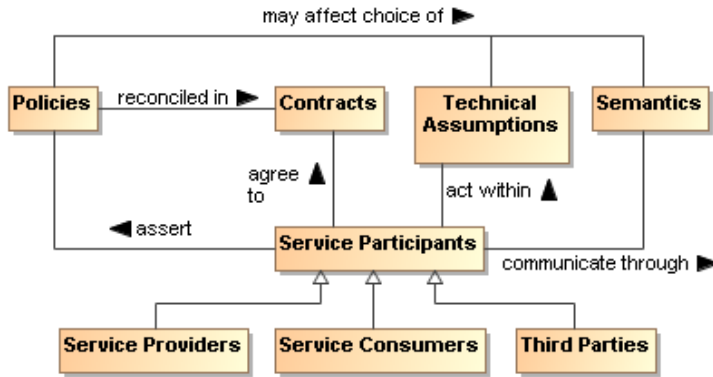
1421 Elements of Service Functionality may be expressed as natural language text, reference to an existing  
1422 taxonomy of functions, or reference to a more formal knowledge capture providing richer description and  
1423 context.

1424 **4.1.1.3.4 Policies and Contracts, Metrics, and Compliance Records**

1425 Policies prescribe the conditions and constraints for interacting with a service and impact the willingness  
1426 to continue visibility with the other participants. Whereas technical assumptions are statements of  
1427 “physical” fact, policies are subjective assertions made by the service provider (sometimes as passed on  
1428 from higher authorities).

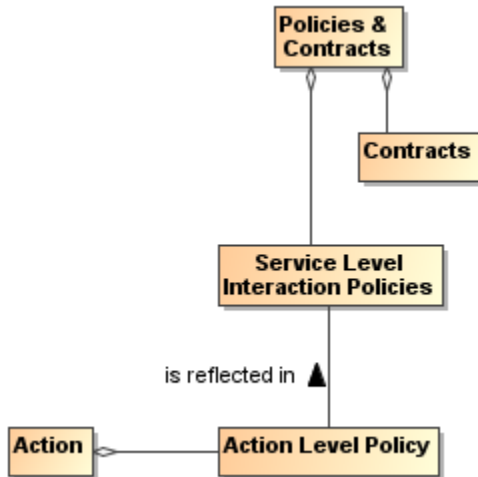
1429 The service description provides a central location for identifying what policies have been asserted by the  
1430 service provider. The specific representation of the policy, e.g. in some formal policy language, is likely  
1431 done outside of the service description and the service description would reference the normative  
1432 definition of the policy.

1433 Policies may also be asserted by other service participants, as illustrated by the model shown in Figure  
 1434 17. Policies that are generally applicable to any interaction with the service are asserted by the service  
 1435 provider and included in the Policies and Contracts section of the service description. Conversely,  
 1436 policies that are asserted by specific consumers or consumer communities would be identified as part of  
 1437 a description's Annotations from 3<sup>rd</sup> parties (see Section 4.1.1.1.4) because these would be specific to  
 1438 those parties and not a general aspect of the service being described.



1439  
 1440 *Figure 17 Model for Policies and Contracts as related to Service Participants*

1441 In *Figure 14* and *Figure 18*, we specifically refer to Service Level Interaction Policies. In a similar manner to  
 1442 that discussed for Service Level vs. Action Level Real World Effects in Section 4.1.1.3.3, individual  
 1443 Actions may have associated policies stating conditions for performing the action, but these must be  
 1444 reflected in and be consistent with the policies made visible at the service level and thus the description of  
 1445 the service as a whole. The relationship to Action Level Policies and the implications on defining the  
 1446 scope of a service are discussed in Section 4.1.2.1.



1447  
 1448  
 1449 *Figure 18 Action-Level and Service-Level Policies*

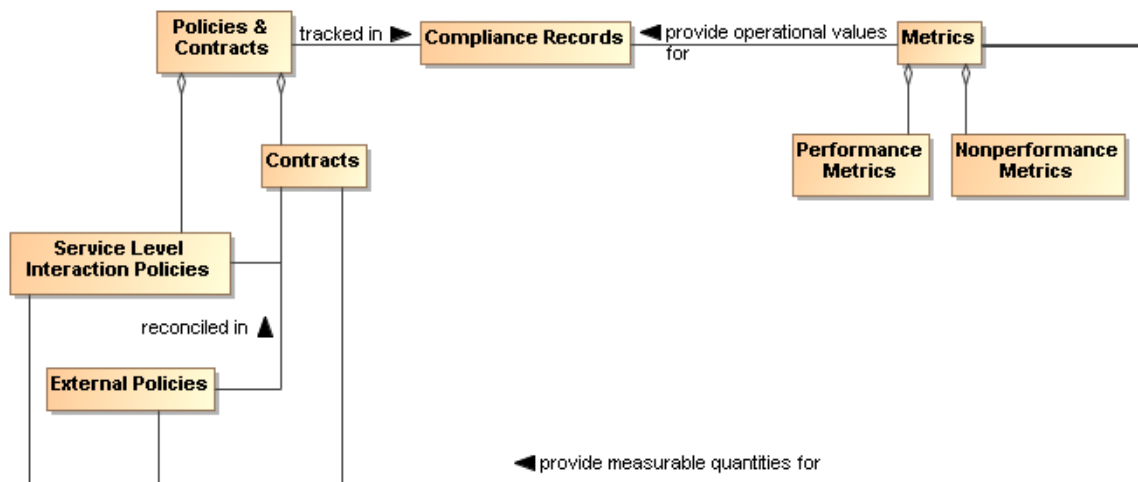
1450 As noted in *Figure 17*, the policies asserted may affect the allowable Technical Assumptions that can be  
 1451 embodied in services or their underlying capabilities and may affect the semantics that can be used. For  
 1452 example of the former, there may be a policy that specifies the surge capacity to be accommodated by a  
 1453 server, and a service that designs for a smaller capacity would not be appropriate to use. For the latter, a  
 1454 policy may require that only services using a community-sponsored vocabulary can be used.

1455 Contracts are agreements among the service participants. The contract may reconcile inconsistent  
 1456 policies asserted by the participants or may specify details of the interaction. Service level agreements  
 1457 (SLAs) are one commonly used category of contracts.



1458 References to contracts under which the service can be used may also be included in the service  
 1459 description. As with policies, the specific representation of the contract, e.g. in some formal contract  
 1460 language, is likely done outside of the service description and the service description would reference the  
 1461 normative definition of the contract. Policies and contracts are discussed further in Section 4.4.

1462 The definition and later enforcement of policies and contracts are predicated on the existence of metrics;  
 1463 the relationships among the relevant concepts are shown in the model in Figure 19. Performance Metrics  
 1464 identify quantities that characterize the speed and quality of realizing the real world effects produced  
 1465 using the SOA service; in addition, policies and contracts may depend on nonperformance metrics, such  
 1466 as whether a license is in place to use the service. Some of these metrics reflect the underlying  
 1467 capability, e.g. a SOA service cannot respond in two seconds if the underlying capability is expected to  
 1468 take five seconds to do its processing; some metrics reflect the implementation of the SOA service, e.g.  
 1469 what level of caching is present to minimize data access requests across the network.



1470  
 1471 *Figure 19 Policies and Contracts, Metrics, and Compliance Records*

1472 As with many quantities, the metrics associated with a service are not themselves defined by this Service  
 1473 Description because it is not known *a priori* which metrics are being collected or otherwise checked by the  
 1474 services, the SOA infrastructure, or other resources that participate in the SOA interactions. However,  
 1475 the service description SHOULD provide a placeholder (possibly through a link to an externally compiled  
 1476 list) for identifying which metrics are available and how these can be accessed.

1477 The use of metrics to evaluate compliance is discussed in Section **Error! Reference source not found.**  
 1478 he results of compliance evaluation SHOULD be maintained in compliance records and the means to  
 1479 access the compliance records SHOULD be included in the Policies and Contracts portion of the service  
 1480 description. For example, the description may be in the form of static information (e.g. over the first year  
 1481 of operation, this service had a 91% availability), a link to a dynamically generated metric (e.g. over the  
 1482 past 30 days, the service has had a 93.3% availability), or access to a dynamic means to check the  
 1483 service for current availability (e.g. a ping). The relationship between service presence and the presence  
 1484 of the individual actions that can be invoked is discussed under Reachability in Section 4.2.2.3.

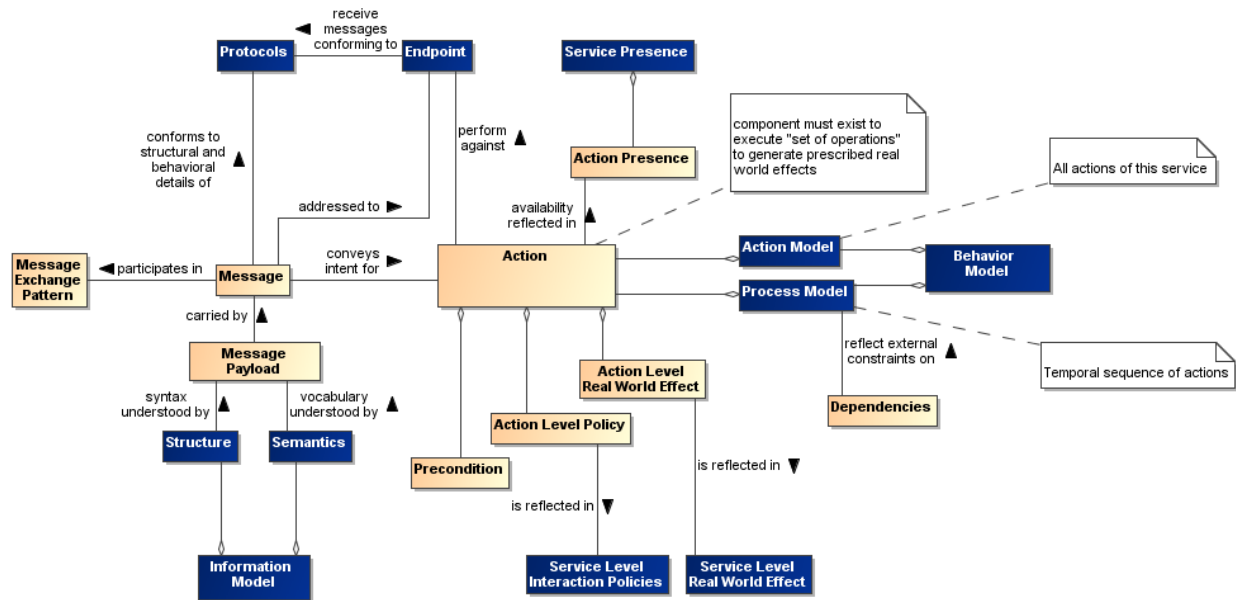
1485 Note, even when policies relate the perspective of a single participant, policy compliance can be  
 1486 measured and policies may be enforceable without contractual agreement with other participants. This  
 1487 should be reflected in the policy, contract, and compliance record information maintained in the service  
 1488 description.

## 1489 4.1.2 Use Of Service Description

### 1490 4.1.2.1 Service Description in support of Service Interaction

1491 If we assume we have awareness, i.e. access to relevant descriptions, the service participants must still  
 1492 establish willingness and presence to ensure full visibility (See Section 4.2) and to interact with the  
 1493 service. Service description provides necessary information for many aspects of preparing for and

1494 carrying through with interaction. Recall the fundamental definition of service is a mechanism to access  
 1495 an underlying capability; the service description describes this mechanism and its use. It lays the  
 1496 groundwork for what can occur, whereas service interaction defines the specifics through which  
 1497 occurrences are realized.



1498  
 1499 *Figure 20 Relationship Between Action and Service Description Components*

1500 Figure 20 combines the models in the subsections of Section 4.1.1 to concisely relate action and the  
 1501 relevant components of Service Description. The purpose of Figure 20 is to demonstrate that the  
 1502 components of service description go beyond arbitrary documentation and form the critical set of  
 1503 information needed to define the what and how of action. In Figure 20, the leaf nodes from *Figure 14* are  
 1504 shown in blue.

1505 action is invoked via a Message where the structure and behavioral details of the message conform to an  
 1506 identified Protocol and is directed to the address of the identified endpoint, and the message payload  
 1507 conforms to the service Information Model.

1508 The availability of an action is reflected in the Action Presence and each Action Presence contributes to  
 1509 the overall Service Presence; this is discussed further in Section 4.2.2.3. Each action has its own  
 1510 endpoint and also its own protocols associated with the endpoint<sup>11</sup> and to what extent, e.g. current or  
 1511 average availability, there is presence for the action through that endpoint. The endpoint and service  
 1512 presence are also part of the service description.

1513 An action may have preconditions where a Precondition is something that needs to be in place before an  
 1514 action can occur, e.g. confirmation of a precursor action. Whether preconditions are satisfied is evaluated  
 1515 when someone tries to perform the action and not before. Presence for an action means someone can  
 1516 initiate it and is independent of whether the preconditions are satisfied. However, the successful  
 1517 completion of the action may depend on whether its preconditions were satisfied.

1518 Analogous to the relationship between actions and preconditions, the Process Model may imply  
 1519 Dependencies for succeeding steps in a process, e.g. that a previous step has successfully completed, or  
 1520 may be isolated to a given step. An example of the latter would be a dependency that the host server has  
 1521 scheduled maintenance and access attempts at these times would fail. Dependencies related to the  
 1522 process model do not affect the presence of a service although these may affect whether the business  
 1523 function successfully completes.

<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.

1524 The conditions under which an action can be invoked may depend on policies associated with the action.  
1525 The Action Level Policies MUST be reflected in the Service Level Interaction Policies because such  
1526 policies may be critical to determining whether the conditions for use of the service are consistent with the  
1527 policies asserted by the service consumer. The service level interaction policies are included in the  
1528 service description.

1529 Similarly, the result of invoking an action is one or more real world effects, and the Action Level Real  
1530 World Effects MUST be reflected in the Service Level Real World Effect included in the service  
1531 description. The unambiguous expression of action level policies and real world effects as service  
1532 counterparts is necessary to adequately understand what constitutes the service interaction.

1533 An adequate service description MUST provide a consumer with information needed to determine if the  
1534 service policies and the (business) functions and service-level real world effects are of interest and there  
1535 is nothing in the technical assumptions that preclude use of the service.

1536 Note at this level, the business functions are not concerned with the action or process models. These  
1537 models are detailed separately.

1538 The service description is not intended to be isolated documentation but rather an integral part of service  
1539 use. Changes in service description SHOULD immediately be made known to consumers and potential  
1540 consumers.

#### 1541 **4.1.2.1.1 Description and Invoking Actions Against a Service**

1542 At this point, let us assume the descriptions were sufficient to establish willingness; see Section 4.2.2.2.  
1543 Figure 20 indicates the service endpoint establishes where to actually carry out the interaction. This is  
1544 where we start considering the action and process models.

1545 The action model identifies the multiple actions a user can perform against a service and the user would  
1546 perform these in the context of the process model as specified or referenced under the Service Interface  
1547 portion of Service Description. For a given business function, there is a corresponding process model,  
1548 where any process model may involve multiple actions. From the above discussion of model elements of  
1549 description we may conclude (1) actions have reachability information, including endpoint and presence,  
1550 (2) presence of service is some aggregation of presence of its actions, (3) action preconditions and  
1551 service dependencies do not affect presence although these may affect successful completion.

1552 Having established visibility, the interaction can proceed. Given a business function, the consumer knows  
1553 what will be accomplished (the service functionality), the conditions under which interaction will proceed  
1554 (service policies and contracts), and the process that must be followed (the process model). The  
1555 remaining question is how does the description information for structure and semantics enable  
1556 interaction.

1557 We have established the importance of the process model in identifying relevant actions and their  
1558 sequence. Interaction proceeds through messages and thus it is the syntax and semantics of the  
1559 messages with which we are here concerned. A common approach is to define the structure and  
1560 semantics that can appear as part of a message; then assemble the pieces into messages; and,  
1561 associate messages with actions. Actions make use of structure and semantics as defined in the  
1562 information model to describe its legal messages.

1563 The process model identifies actions to be performed against a service and the sequence for performing  
1564 the actions. For a given action, the Reachability portion of description indicates the protocol bindings that  
1565 are available, the endpoint corresponding to a binding, and whether there is presence at that endpoint.  
1566 The interaction with actions is through messages that conform to the structure and semantics defined in  
1567 the information model and the message sequence conforming to the action's identified MEP. The result  
1568 is some portion of the real world effect that must be assessed and/or processed (e.g. if an error exists,  
1569 that part that covers the error processing would be invoked).

#### 1570 **4.1.2.1.2 The Question of Multiple Business Functions**

1571 Action level effects and policies MUST be reflected at the service level for service description to support  
1572 visibility.

1573 It is assumed that a SOA service represents an identifiable business function to which policies can be  
1574 applied and from which desired business effects can be obtained. While contemporary discussions of

1575 SOA services and supporting standards do not constrain what actions or combinations of actions can or  
1576 should be defined for a service, the SOA-RAF considers the implications of service description in defining  
1577 the range of actions appropriate for an individual SOA service.

1578 Consider the situation if a given SOA service is the container for multiple independent (but loosely  
1579 related) business functions. These are not multiple effects from a single function but multiple functions  
1580 with potentially different sets of effects for each function. A service can have multiple actions a user may  
1581 perform against it, and this does not change with multiple business functions. As an individual business  
1582 function corresponds to a process model, so multiple business functions imply multiple process models.  
1583 The same action may be used in multiple process models but the aggregated service presence would be  
1584 specific to each business function because the components being aggregated may be different between  
1585 process models. In summary, for a service with multiple business functions, each function has (1) its own  
1586 process model and dependencies, (2) its own aggregated presence, and (3) possibly its own list of  
1587 policies and real world effects.

1588 A common variation on this theme is for a single service to have multiple endpoints for different levels of  
1589 quality of service (QoS). Different QoS imply separate statements of policy, separate endpoints, possibly  
1590 separate dependencies, and so on. One could say the QoS variation does not require this because there  
1591 can be a single QoS policy that encompasses the variations. and all other aspects of the service would be  
1592 the same except for the endpoint used for each QoS. However, the different aspects of policy at the  
1593 service level would need to be mapped to endpoints, and this introduces an undesirable level of coupling  
1594 across the elements of description. In addition, it is obvious that description at the service level can  
1595 become very complicated if the number of combinations is allowed to grow.

1596 One could imagine a service description that is basically a container for action descriptions, where each  
1597 action description is self contained; however, this would lead to duplication of description components  
1598 across actions. If common description components are factored, this either is limited to components  
1599 common across all actions or requires complicated tagging to capture the components that often but do  
1600 not universally apply.

1601 If a provider cannot describe a service as a whole but must describe every action, this leads to the  
1602 situation where it may be extremely difficult to construct a clear and concise service description that can  
1603 effectively support discovery and use without tedious logic to process the description and assemble the  
1604 available permutations. In effect, if adequate description of an action begins to look like description of a  
1605 service, it may be best to have it as a separate service.

1606 Recall, more than one service can access the same underlying capability, and this is appropriate if a  
1607 different real world effect is to be exposed. Along these lines, one can argue that different QoS are  
1608 different services because getting a response in one minute rather than one hour is more than a QoS  
1609 difference; it is a fundamental difference in the business function being provided.

1610 As a best practice, a criteria for whether a service is appropriately scoped may be the ease or difficulty in  
1611 creating an unambiguous service description. A consequence of having tightly-scoped services is there  
1612 will be a greater reliance on combining services, i.e. more fundamental business functions, to create more  
1613 advanced business functions. This is consistent with the principles of service oriented architecture and is  
1614 the basic position of the Reference Architecture, although not an absolute requirement. Combining  
1615 services increases the reliance on understanding and implementing the concepts of orchestration,  
1616 choreography, and other approaches yet to be developed; these are discussed in more detail in section  
1617 4.4 Interacting with Services.

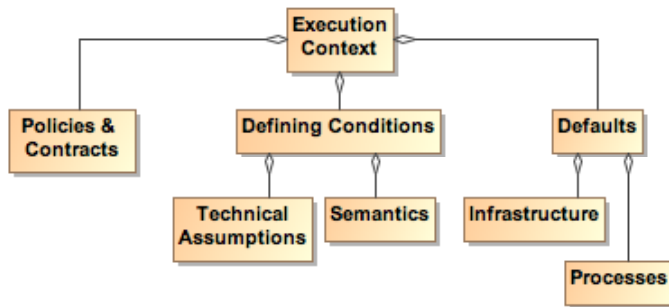
#### 1618 **4.1.2.1.3 Service Description, Execution Context, and Service Interaction**

1619 The service description MUST provide sufficient information to support service visibility, including the  
1620 willingness of service participants to interact. However, the corresponding descriptions for providers and  
1621 consumers may both contain policies, technical assumptions, constraints on semantics, and other  
1622 technical and procedural conditions that must be aligned to define the terms of willingness. The  
1623 agreements which encapsulate the necessary alignment form the basis upon which interactions may  
1624 proceed – in the Reference Model, this collection of agreements and the necessary environmental  
1625 support establish the execution context.

1626 To illustrate the concept of the execution context, consider a Web-based system for timecard entry. For  
1627 an employee onsite at an employer facility, the execution context requires a computer connected to the

1628 local network and the employee must enter their network ID and password. Relevant policies include that  
 1629 the employee must maintain the most recent anti-virus software and virus definitions for any computer  
 1630 connected to the network.

1631 For the same employee connecting from offsite, the execution context specifies the need for a computer  
 1632 with installed VPN software and a security token to negotiate the VPN connection. The execution context  
 1633 also includes proxy settings as needed to connect to the offsite network. The employee must still comply  
 1634 with the requirements for onsite computers and access, but the offsite execution context includes  
 1635 additional items before the employee can access the same underlying capability and realize the same  
 1636 real world effects, i.e. the timecard entries.

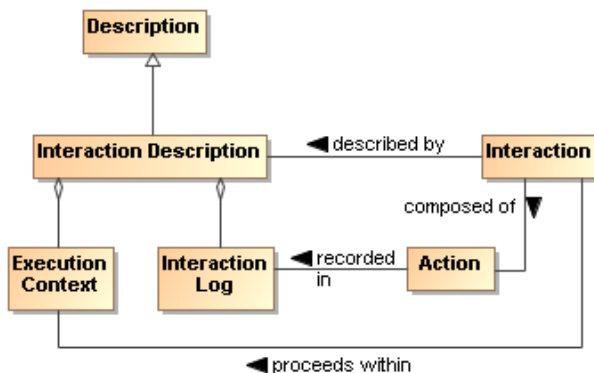


1637  
 1638 *Figure 21 Execution Context*

1639 Figure 21 shows a few broad categories found in execution context. These are not meant to be  
 1640 comprehensive. Other items may need to be included to collect a sufficient description of the interaction  
 1641 conditions. Any other items not explicitly noted in the model but needed to set the environment SHOULD  
 1642 be included in the execution context.

1643 While the execution context captures the conditions under which interaction can occur, it does not capture  
 1644 the specific service invocations that do occur in a specific interaction. A service interaction as modeled in  
 1645 Figure 20 introduces the concept of an Interaction Description which is composed of both the Execution  
 1646 Context and an Interaction Log. The execution context specifies the set of conditions under which the  
 1647 interaction occurs and the interaction log captures the sequence of service interactions that occur within  
 1648 the execution context. This sequence should follow the Process Model but can include details beyond  
 1649 those specified there. For example, the Process Model may specify an action that results in identifying a  
 1650 data source, and the identified source is used in a subsequent action. The Interaction Log would record  
 1651 the specific data source used.

1652 The execution context can be thought of as the container in which the interaction occurs and the  
 1653 interaction log captures what happens inside the container. This combination is needed to support  
 1654 auditability and repeatability of the interactions.



1655  
 1656 *Figure 22 Interaction Description*

1657 SOA allows flexibility to accomplish repeatability or reusability. One benefit of this is that a service can be  
1658 updated without disrupting the user experience of the service. So, Google can improve their ranking  
1659 algorithm without notifying the user about the details of the update.

1660 However, it may also be vital for the consumer to be able to recreate past results or to generate  
1661 consistent results in the future, and information such as what conditions, which services, and which  
1662 versions of those services are used is indispensable in retracing one's path. The interaction log is a  
1663 critical part of the resulting real world effects because it defines how the effects were generated and  
1664 possibly the meaning of observed effects. This increases in importance as dynamic composability  
1665 becomes more feasible. In essence, a result has limited value if one does not know how it was  
1666 generated.

1667 The interaction log SHOULD be a detailed trace for a specific interaction, and its reuse is limited to  
1668 duplicating that interaction. An execution context can act as a template for identical or similar  
1669 interactions. Any given execution context MAY define the conditions of future interactions.

1670 Such uses of execution context imply (1) a standardized format for capturing execution context and (2) a  
1671 subclass of general description could be defined to support visibility of saved execution contexts. The  
1672 specifics of the relevant formats and descriptions are beyond the scope of this document.

1673 A service description is unlikely to track interaction descriptions or the constituent execution contexts or  
1674 interaction logs that include mention of the service. However, as appropriate, linking to specific instances  
1675 of either of these could be done through associated annotations.

#### 1676 4.1.3 Relationship to Other Description Models

1677 While the representation shown in Figure 13 is derived from considerations related to service description,  
1678 it is acknowledged that other metadata standards are relevant and should, as possible, be incorporated  
1679 into this work. Two standards of particular relevance are the Dublin Core Metadata Initiative (DCMI) and  
1680 ISO 11179, especially Part 5.

1681 When the service description (or even the general description class) is considered as the DCMI  
1682 "resource", Figure 13 aligns nicely with the DCMI resource model. While some differences exist, these  
1683 are mostly in areas where DCMI goes into detail that is considered beyond the scope of the current  
1684 Reference Architecture. For example, DCMI defines classes of "shared semantics" whereas this  
1685 Reference Architecture Framework considers that an identification of relevant semantic models is  
1686 sufficient. Likewise, the DCMI "description model" goes into the details of possible syntax encodings  
1687 whereas for the Reference Architecture Framework it is sufficient to identify the relevant formats.

1688 With respect to ISO 11179 Part 5, the metadata fields defined in that reference may be used without  
1689 prejudice as the properties in Figure 13. Additionally, other defined metadata sets may be used by the  
1690 service provider if the other sets are considered more appropriate, i.e. it is fundamental to this reference  
1691 architecture to identify the need and the means to make vocabulary declarations explicit but it is beyond  
1692 the scope to specify which vocabularies are to be used. In addition, the identification of domain of the  
1693 properties and range of the values has not been included in the current Reference Architecture  
1694 discussion, but the text of ISO 11179 Part 5 can be used consistently with the model prescribed in this  
1695 document.

1696 Description as defined here considers a wide range of applicability and support of the principles of service  
1697 oriented architecture. Other metadata models can be used in concert with the model presented here  
1698 because most of these focus on a finer level of detail that is outside the present scope, and so provide a  
1699 level of implementation guidance that can be applied as appropriate.

#### 1700 4.1.4 Architectural Implications

1701 The description of service description indicates numerous architectural implications on the SOA  
1702 ecosystem:

- 1703 • It changes over time and its contents will reflect changing needs and context. This requires the  
1704 existence of:
  - 1705 ○ mechanisms to support the storage, referencing, and access to normative definitions of one  
1706 or more versioning schemes that may be applied to identify different aggregations of

- 1707 descriptive information, where the different schemes may be versions of a versioning scheme  
 1708 itself;
- 1709 ○ configuration management mechanisms to capture the contents of the each aggregation and  
 1710 apply a unique identifier in a manner consistent with an identified versioning scheme;
  - 1711 ○ one or more mechanisms to support the storage, referencing, and access to conversion  
 1712 relationships between versioning schemes, and the mechanisms to carry out such  
 1713 conversions.
- 1714 ● Description makes use of defined semantics, where the semantics may be used for categorization or  
 1715 providing other property and value information for description classes. This requires the existence of:
    - 1716 ○ semantic models that provide normative descriptions of the utilized terms, where the models  
 1717 may range from a simple dictionary of terms to an ontology showing complex relationships  
 1718 and capable of supporting enhanced reasoning;
    - 1719 ○ mechanisms to support the storage, referencing, and access to these semantic models;
    - 1720 ○ configuration management mechanisms to capture the normative description of each  
 1721 semantic model and to apply a unique identifier in a manner consistent with an identified  
 1722 versioning scheme;
    - 1723 ○ one or more mechanisms to support the storage, referencing, and access to conversion  
 1724 relationships between semantic models, and the mechanisms to carry out such conversions.
  - 1725 ● Descriptions include reference to policies defining conditions of use and optionally contracts  
 1726 representing agreement on policies and other conditions. This requires the existence of (as also  
 1727 enumerated under governance):
    - 1728 ○ descriptions to enable the policy modules to be visible, where the description includes a  
 1729 unique identifier for the policy and a sufficient, and preferably a machine processible,  
 1730 representation of the meaning of terms used to describe the policy, its functions, and its  
 1731 effects;
    - 1732 ○ one or more discovery mechanisms that enable searching for policies that best meet the  
 1733 search criteria specified by the service participant; where the discovery mechanism has  
 1734 access to the individual policy descriptions, possibly through some repository mechanism;
    - 1735 ○ accessible storage of policies and policy descriptions, so service participants can access,  
 1736 examine, and use the policies as defined.
  - 1737 ● Descriptions include references to metrics which describe the operational characteristics of the  
 1738 subjects being described. This requires the existence of (as partially enumerated under governance):
    - 1739 ○ the infrastructure monitoring and reporting information on SOA resources;
    - 1740 ○ possible interface requirements to make accessible metrics information generated or most  
 1741 easily accessed by the service itself;
    - 1742 ○ mechanisms to catalog and enable discovery of which metrics are available for a described  
 1743 resources and information on how these metrics can be accessed;
    - 1744 ○ mechanisms to catalog and enable discovery of compliance records associated with policies  
 1745 and contracts that are based on these metrics.
  - 1746 ● Descriptions of the interactions are important for enabling auditability and repeatability, thereby  
 1747 establishing a context for results and support for understanding observed change in performance or  
 1748 results. This requires the existence of:
    - 1749 ○ one or more mechanisms to capture, describe, store, discover, and retrieve interaction logs,  
 1750 execution contexts, and the combined interaction descriptions;
    - 1751 ○ one or more mechanisms for attaching to any results the means to identify and retrieve the  
 1752 interaction description under which the results were generated.
  - 1753 ● Descriptions may capture very focused information subsets or can be an aggregate of numerous  
 1754 component descriptions. Service description is an example of an aggregate for which manual  
 1755 maintenance of the whole would not be feasible. This requires the existence of:
    - 1756 ○ tools to facilitate identifying description elements that are to be aggregated to assemble the  
 1757 composite description;
    - 1758 ○ tools to facilitate identifying the sources of information to associate with the description  
 1759 elements;
    - 1760 ○ tools to collect the identified description elements and their associated sources into a  
 1761 standard, referenceable format that can support general access and understanding;

- 1762  
1763  
1764  
1765  
1766  
1767  
1768  
1769  
1770  
1771  
1772  
1773
- tools to automatically update the composite description as the component sources change, and to consistently apply versioning schemes to identify the new description contents and the type and significance of change that occurred.
  - Descriptions provide up-to-date information on what a resource is, the conditions for interacting with the resource, and the results of such interactions. As such, the description is the source of vital information in establishing willingness to interact with a resource, reachability to make interaction possible, and compliance with relevant conditions of use. This requires the existence of:
    - one or more discovery mechanisms that enable searching for described resources that best meet the criteria specified by a service participant, where the discovery mechanism has access to individual descriptions, possibly through some repository mechanism;
    - tools to appropriately track users of the descriptions and notify them when a new version of the description is available.

## 1774 4.2 Service Visibility Model

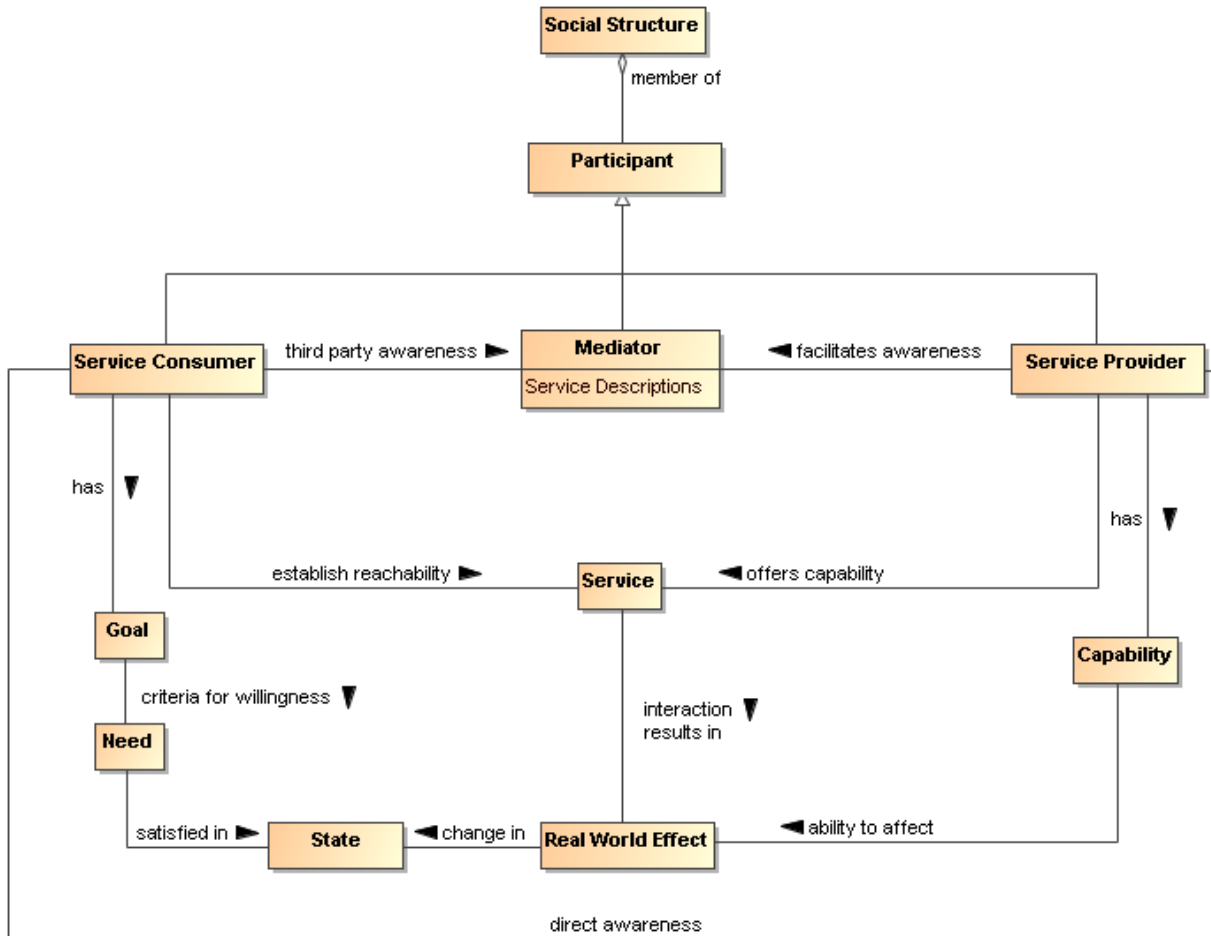
1775 One of the key requirements for participants interacting with each other in the context of a SOA is  
1776 achieving visibility: before services can interoperate, the participants have to be visible to each other  
1777 using whatever means are appropriate. The Reference Model analyzes visibility in terms of awareness,  
1778 willingness, and reachability. In this section, we explore how visibility may be achieved.

### 1779 4.2.1 Visibility to Business

1780 The relationship of visibility to the SOA ecosystem encompasses both human social structures and  
1781 automated IT mechanisms. Figure 23 depicts a business setting that is a basis for visibility as related to  
1782 the social structure Model in the *Participation in a SOA Ecosystem* view (see Section **Error! Reference**  
1783 **source not found.**). Service consumers and service providers may have direct awareness or mediated  
1784 awareness where mediated awareness is achieved through some third party. A consumer's willingness to  
1785 use a service is reflected by the consumer's presumption of satisfying goals and needs based on the  
1786 description of the service. Service providers offer capabilities that have real world effects that result in a  
1787 change in state of the consumer. Reachability of the service by the consumer leads to interactions that  
1788 change the state of the consumer. The consumer can measure the change of state to determine if the  
1789 claims made by description and the real world effects of consuming the service meet the consumer's  
1790 needs.

1791





1792  
1793

Figure 23 Visibility to Business

1794 Visibility and interoperability in a SOA ecosystem requires more than location and interface information.  
1795 A meta-model for this broader view of visibility is depicted in Section 4.1. In addition to providing  
1796 improved awareness of service capabilities through description of information such as reachability,  
1797 behavior models, information models, functionality, and metrics, the service description may contain  
1798 policies valuable for determination of willingness to interact.

1799 A mediator of service descriptions may provide event notifications to both consumers and providers about  
1800 information relating to service descriptions. One example of this capability is a publish/subscribe model  
1801 where the mediator allows consumers to subscribe to service description version changes made by the  
1802 provider. Likewise, the mediator may provide notifications to the provider of consumers that have  
1803 subscribed to service description updates.

1804 Another important business capability in a SOA environment is the ability to narrow visibility to trusted  
1805 members within a social structure. Mediators for awareness may provide policy based access to service  
1806 descriptions allowing for the dynamic formation of awareness between trusted members.

## 1807 4.2.2 Visibility

1808 Attaining visibility is described in terms of steps that lead to visibility. While there can be many contexts  
1809 for visibility within a single social structure, the same general steps can be applied to each of the contexts  
1810 to accomplish visibility.

1811 Attaining SOA visibility requires

- 1812 • service description creation and maintenance,
- 1813 • processes and mechanisms for achieving awareness of and accessing descriptions,

- 1814 • processes and mechanisms for establishing willingness of participants,  
1815 • processes and mechanisms to determine reachability.

1816 Visibility may occur in stages, i.e. a participant can become aware enough to look or ask for further  
1817 description, and with this description, the participant can decide on willingness, possibly requiring  
1818 additional description. For example, if a potential consumer has a need for a tree cutting (business)  
1819 service, the consumer can use a web search engine to find web sites of providers. The web search  
1820 engine (a mediator) gives the consumer links to relevant web pages and the consumer can access those  
1821 descriptions. For those prospective providers that satisfy the consumer's criteria, the consumer's  
1822 willingness to interact increases. The consumer may contact several tree services to get detailed cost  
1823 information (or arrange for an estimate) and may ask for references (further description). The consumer is  
1824 likely to establish full visibility and proceed with interaction with the tree service who mutually establishes  
1825 visibility.

#### 1826 **4.2.2.1 Awareness**

1827 A service participant is aware of another participant if it has access to a description of that participant with  
1828 sufficient completeness to establish the other requirements of visibility.

1829 Awareness is inherently a function of a participant; awareness can be established without any action on  
1830 the part of the target participant other than the target providing appropriate descriptions. Awareness is  
1831 often discussed in terms of consumer awareness of providers but the concepts are equally valid for  
1832 provider awareness of consumers.

1833 Awareness can be decomposed into the creation of descriptions, making them available, and discovering  
1834 the descriptions. Discovery can be initiated or it can be by notification. Initiated discovery for business  
1835 may require formalization of the required capabilities and resources to achieve business goals.

1836 Achieving awareness in a SOA can range from word of mouth to formal service descriptions in a  
1837 standards-based registry-repository. Some other examples of achieving awareness in a SOA are the  
1838 use of a web page containing description information, email notifications of descriptions, and document  
1839 based descriptions.

1840 A mediator as discussed for awareness is a third party participant that provides awareness to one or  
1841 more consumers of one or more services. Direct awareness is awareness between a consumer and  
1842 provider without the use of a third party.

1843 Direct awareness may be the result of having previously established an execution context, or direct  
1844 awareness may include determining the presence of services and then querying the service directly for  
1845 description. As an example, a priori visibility of some sensor device may provide the means for interaction  
1846 or a query for standardized sensor device metadata may be broadcast to multiple locations. If  
1847 acknowledged, the service interface for the device may directly provide description to a consumer so the  
1848 consumer can determine willingness to interact.

1849 The same medium for awareness may be direct in one context and may be mediated in another context.  
1850 For example, a service provider may maintain a web site with links to the provider's descriptions of  
1851 services giving the consumers direct awareness to the provider's services. Alternatively, a community  
1852 may maintain a mediated web site with links to various provider descriptions of services for any number of  
1853 consumers. More than one mediator may be involved, as different mediators may specialize in different  
1854 mediation functions.

1855 Descriptions may be formal or informal. Section 4.1, provides a comprehensive model for service  
1856 description that can be applied to formal registry/repositories used to mediate visibility. Using consistent  
1857 description taxonomies and standards based mediated awareness helps provide more effective  
1858 awareness.

#### 1859 **4.2.2.1.1 Mediated Awareness**

1860 Mediated awareness promotes loose coupling by keeping the consumers and services from explicitly  
1861 referring to each other and the descriptions. Mediation lets interaction vary independently. Rather than all  
1862 potential service consumers being informed on a continual basis about all services, there is a known or  
1863 agreed upon facility or location that houses the service description.

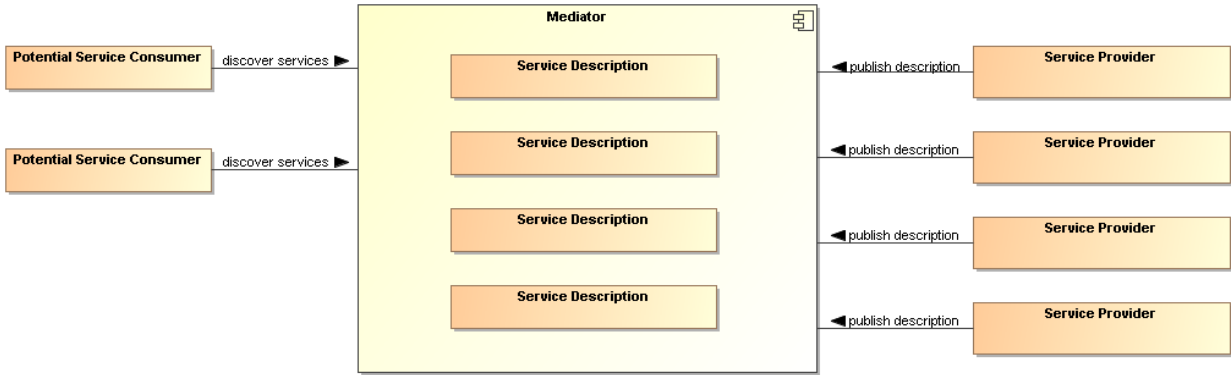


Figure 24 Mediated Service Awareness

1864  
1865  
1866 In Figure 24, the potential service consumers perform queries or are notified in order to locate those  
1867 services that satisfy their needs. As an example, the telephone book is a mediated registry where  
1868 individuals perform manual searches to locate services (i.e. the yellow pages). The telephone book is  
1869 also a mediated registry for solicitors to find and notify potential customers (i.e. the white pages).

1870 In mediated service awareness for large and dynamic numbers of service consumers and service  
1871 providers, the benefits typically far outweigh the management issues associated with it. Some of the  
1872 benefits of mediated service awareness are

- 1873 • Potential service consumers have a known location for searching thereby eliminating needless and  
1874 random searches
- 1875 • Typically a consortium of interested parties (or a sufficiently large corporation) signs up to host the  
1876 mediation facility
- 1877 • Standardized tools and methods can be developed and promulgated to promote interoperability and  
1878 ease of use.

1879 However, mediated awareness can have some risks associated with it:

- 1880 • A single point of failure. If the central mediation service fails then a large number of service providers  
1881 and consumers are potentially adversely affected.
- 1882 • A single point of control. If the central mediation service is owned by, or controlled by, someone other  
1883 than the service consumers and/or providers then the latter may be put at a competitive disadvantage  
1884 based on policies of the discovery provider.

1885 A common mechanism for mediated awareness is a registry-repository. The registry stores links or  
1886 pointers to service description artifacts. The repository in this example is the storage location for the  
1887 service description artifacts. Service descriptions can be pushed (publish/subscribe for example) or pulled  
1888 from the register-repository mediator.

1889 The registry is like a card catalog at the library and a repository is like the shelves for the books.  
1890 Standardized metadata describing repository content can be stored as registry objects in a registry and  
1891 any type of content can be stored as repository items in a repository. The registry may be constructed  
1892 such that description items stored within the mediation facility repository has intrinsic links in the registry  
1893 while description items stored outside the mediation facility have extrinsic links in the registry.

1894 When independent but like SOA IT mechanisms interoperate with one another, the IT mechanisms may  
1895 be referred to as federated.

#### 1896 4.2.2.1.2 Awareness in Complex Social Structures

1897 Awareness applies to one or more communities within one or more social structures where a community  
1898 consists of at least one description provider and one description consumer. These communities may be  
1899 part of the same social structure or be part of different ones.

1900 In Figure 25, awareness can be within a single community, multiple communities, or all communities in  
1901 the social structure. The social structure can encourage or restrict awareness through its policies, and  
1902 these policies can affect participant willingness. The information about policies should be incorporated in

1903 the relevant descriptions. The social structure also governs the conditions for establishing contracts, the  
 1904 results of which will be reflected in the execution context if interaction is to proceed.

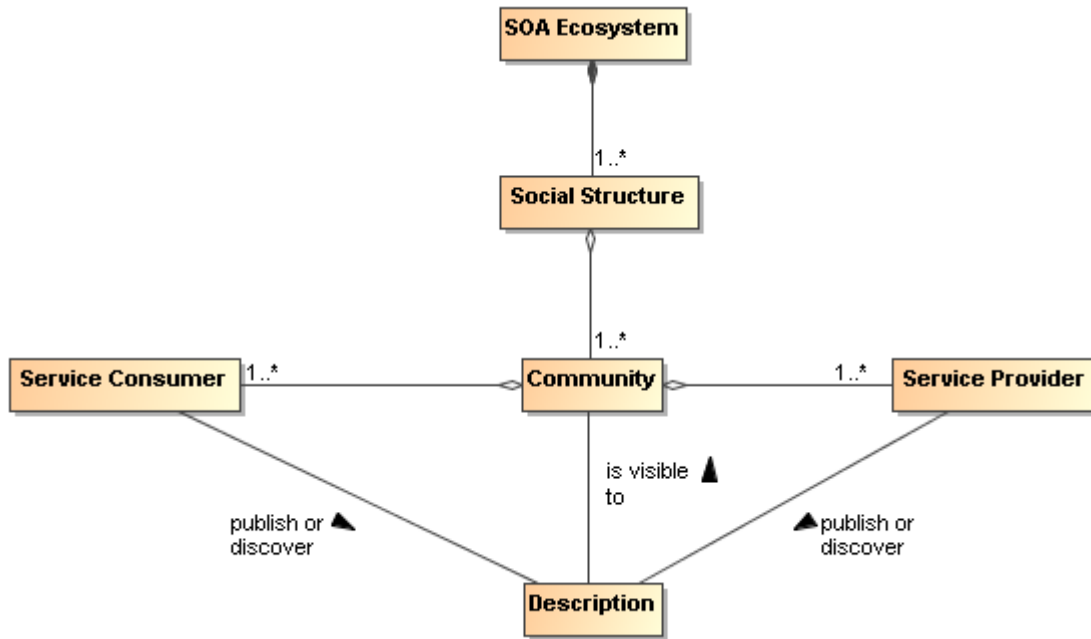


Figure 25 Awareness in a SOA Ecosystem

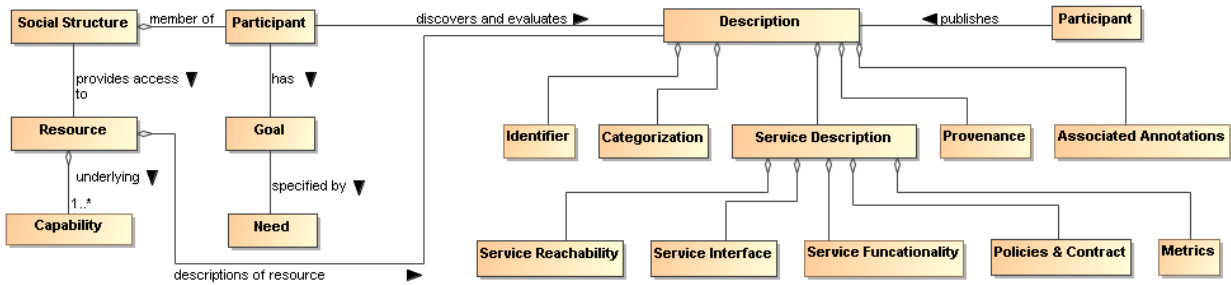
1905  
 1906  
 1907 IT policy/contract mechanisms can be used by visibility mechanisms to provide awareness between  
 1908 communities. The IT mechanisms for awareness may incorporate trust mechanisms to assure  
 1909 awareness between trusted communities. For example, government organizations may want to limit  
 1910 awareness of an organization's services to specific communities of interest.  
 1911 Another common business model for awareness is maximizing awareness to communities within the  
 1912 social structure, the traditional market place business model. A centralized mediator often arises as a  
 1913 provider for this global visibility, a gatekeeper of visibility so to speak. For example, Google is a  
 1914 centralized mediator for accessing information on the web. As another example, television networks have  
 1915 centralized entities providing a level of awareness to communities that otherwise could not be achieved  
 1916 without going through the television network.  
 1917 However, mediators have motivations, and they may be selective in which information they choose to  
 1918 make available to potential consumers. For example, in a secure environment, the mediator may enforce  
 1919 security policies and make information selectively available depending on the security clearance of the  
 1920 consumers.

1921 **4.2.2.2 Willingness**

1922 Having achieved awareness, participants use descriptions to help determine their willingness to interact  
 1923 with another participant. Both awareness and willingness are determined prior to consumer/provider  
 1924 interaction.

1925

1926



1927

1928

Figure 26 Business, Description and Willingness

1929 Figure 26 relates elements of the *Participation in a SOA Ecosystem* view, and elements from the Service  
 1930 Description Model to willingness. By having a willingness to interact within a particular social structure,  
 1931 the social structure provides the participant access to capabilities based on conditions the social structure  
 1932 finds appropriate for its context. The participant can use these capabilities to satisfy goals and objectives  
 1933 as specified by the participant's needs.

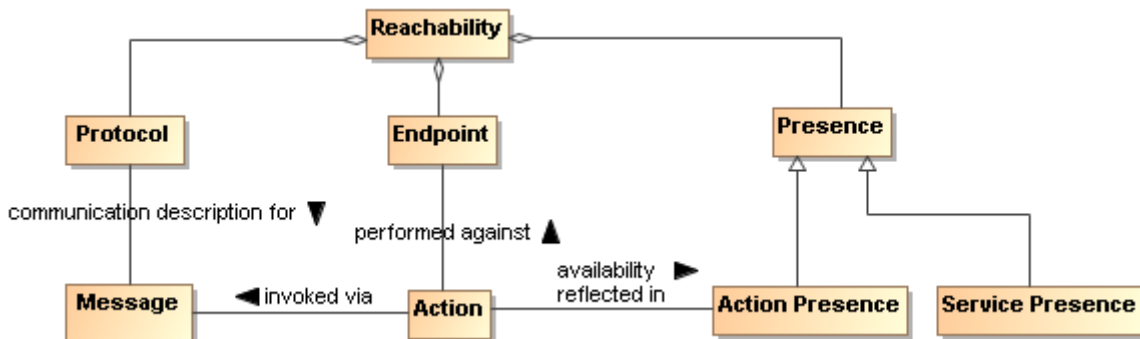
1934 In Figure 26, information used to determine willingness is defined by Description. Information referenced  
 1935 by Description may come from many sources. For example, a mediator for descriptions may provide 3rd  
 1936 party annotations for reputation. Another source for reputation may be a participant's own history of  
 1937 interactions with another participant.

1938 A participant inspects functionality for potential satisfaction of needs. Identity is associated with any  
 1939 participant, however, identity may or may not be verified. If available, participant reputation may be a  
 1940 deciding factor for willingness to interact. Policies and contracts referenced by the description may be  
 1941 particularly important to determine the agreements and commitments required for business interactions.  
 1942 Provenance may be used for verification of authenticity of a resource.

1943 Mechanisms that aid in determining willingness make use of the artifacts referenced by descriptions of  
 1944 services. Mechanisms for establishing willingness could be as simple as rendering service description  
 1945 information for human consumption to automated evaluation of functionality, policies, and contracts by a  
 1946 rules engine. The rules engine for determining willingness could operate as a policy decision procedure  
 1947 as defined in Section 4.4.

### 1948 4.2.2.3 Reachability

1949 Reachability involves knowing the endpoint, protocol, and presence of a service. At a minimum,  
 1950 reachability requires information about the location of the service and the protocol describing the means  
 1951 of communication.



1952

1953

Figure 27 Service Reachability

1954

### 1955 Endpoint

1956 An endpoint is a reference-able entity, processor or resource against which an action can be  
 1957 performed.

1958 **Protocol**

1959 A protocol is a structured means by which service interaction is regulated.

1960 **Presence**

1961 Presence is the measurement of reachability of a service at a particular point in time.

1962 A protocol defines a structured method of communication with a service. Presence is determined by  
1963 interaction through a communication protocol. Presence may not be known in many cases until the act of  
1964 interaction begins. To overcome this problem, IT mechanisms may make use of presence protocols to  
1965 provide the current up/down status of a service.

1966 Service reachability enables service participants to locate and interact with one another. Each action may  
1967 have its own endpoint and also its own protocols associated with the endpoint and whether there is  
1968 presence for the action through that endpoint. Presence of a service is an aggregation of the presence of  
1969 the service's actions, and the service level may aggregate to some degraded or restricted presence if  
1970 some action presence is not confirmed. For example, if error processing actions are not available, the  
1971 service can still provide required functionality if no error processing is needed. This implies reachability  
1972 relates to each action as well as applying to the service/business as a whole.

1973 **4.2.3 Architectural Implications**

1974 Visibility in a SOA ecosystem has the following architectural implications on mechanisms providing  
1975 support for awareness, willingness, and reachability:

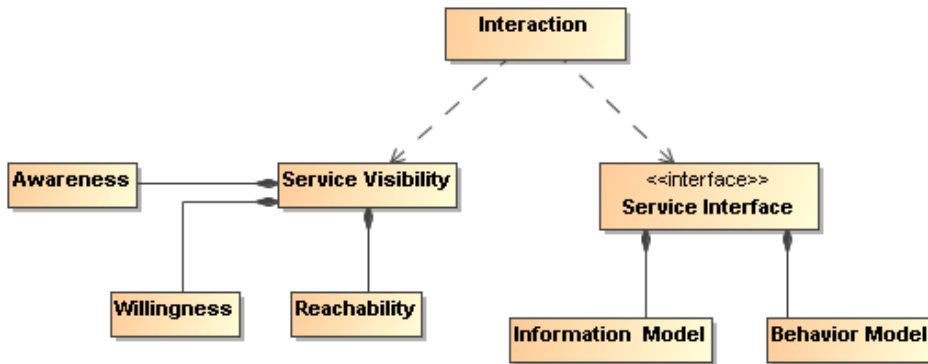
- 1976 • Mechanisms providing support for awareness have the following minimum capabilities:
  - 1977 ○ creation of Description, preferably conforming to a standard Description format and structure;
  - 1978 ○ publishing of Description directly to a consumer or through a third party mediator;
  - 1979 ○ discovery of Description, preferably conforming to a standard for Description discovery;
  - 1980 ○ notification of Description updates or notification of the addition of new and relevant  
1981 Descriptions;
  - 1982 ○ classification of Description elements according to standardized classification schemes.
- 1983 • In a SOA ecosystem with complex social structures, awareness may be provided for specific  
1984 communities of interest. The architectural mechanisms for providing awareness to communities of  
1985 interest require support for:
  - 1986 ○ policies that allow dynamic formation of communities of interest;
  - 1987 ○ trust that awareness can be provided for and only for specific communities of interest, the  
1988 bases of which is typically built on keying and encryption technology.
- 1989 • The architectural mechanisms for determining willingness to interact require support for:
  - 1990 ○ verification of identity and credentials of the provider and/or consumer;
  - 1991 ○ access to and understanding of description;
  - 1992 ○ inspection of functionality and capabilities;
  - 1993 ○ inspection of policies and/or contracts.
- 1994 • The architectural mechanisms for establishing reachability require support for:
  - 1995 ○ the location or address of an endpoint;
  - 1996 ○ verification and use of a service interface by means of a communication protocol;
  - 1997 ○ determination of presence with an endpoint which may only be determined at the point of  
1998 interaction but may be further aided by the use of a presence protocol for which the endpoints  
1999 actively participate.

2000 **4.3 Interacting with Services Model**

2001 Interaction is the activity involved in using a service to access capability in order to achieve a particular  
2002 desired real world effect, where real world effect is the actual *result* of using a service. An interaction can  
2003 be characterized by a sequence of actions. Consequently, interacting with a service, i.e. performing  
2004 actions against the service—usually mediated by a series of message exchanges—involves actions  
2005 performed by the service. Different modes of interaction are possible such as modifying the shared state  
2006 of a resource. Note that a participant (or delegate acting on behalf of the participant) can be the sender  
2007 of a message, the receiver of a message, or both.

2008 **4.3.1 Interaction Dependencies**

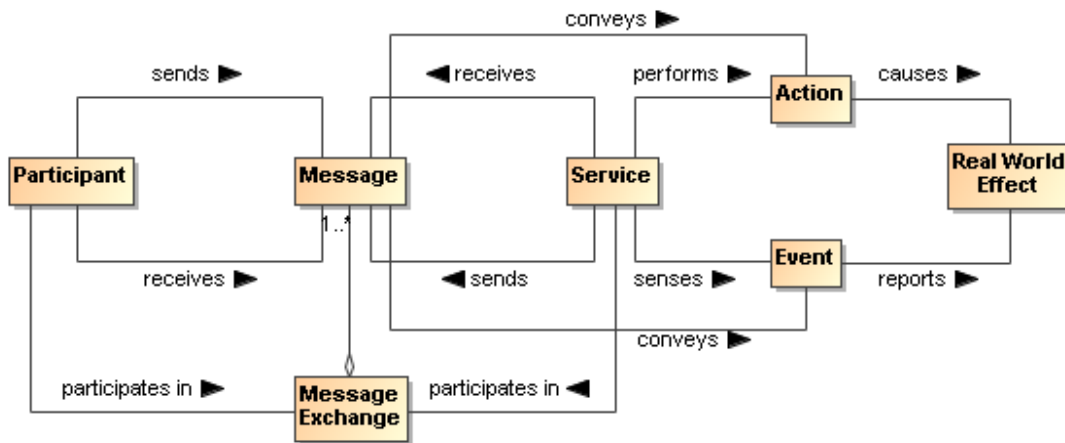
2009 Recall from the Reference Model that service visibility is the capacity for those with needs and those with  
 2010 capabilities to be able to interact with each other, and that the service interface is the means by which the  
 2011 underlying capabilities of a service are accessed. Ideally, the details of the underlying service  
 2012 implementation are abstracted away by the service interface. [Service] interaction therefore has a direct  
 2013 dependency on the visibility of the service as well as its implementation-neutral interface (see Figure 28).  
 2014 Service visibility is composed of awareness, willingness, and reachability and service interface is  
 2015 composed of the information and behavior models. Service visibility is modeled in Section 4.2 while  
 2016 service interface is modeled in Section 4.1.



2017  
 2018 *Figure 28 Interaction dependencies.*

2019 **4.3.2 Actions and Events**

2020 For purposes of the SOA-RAF, the authors have committed to the use of message exchange between  
 2021 service participants to denote actions performed against and by the service, and to denote events that  
 2022 report on real world effects that are caused by the service actions. A visual model of the relationship  
 2023 between these concepts is shown in Figure 29.



2024  
 2025 *Figure 29 A "message" conveys either an action or an event.*

2026 A *message* conveys either an action or an event. In other words, both actions and events, realized by the  
 2027 SOA services, are denoted by the messages. The Reference Model states that the action model  
 2028 characterizes the “permissible set of actions that may be invoked against a service.” We extend that  
 2029 notion here to include events as part of the event model and that messages denote either actions or  
 2030 notification of events.

2031 In Section **Error! Reference source not found.**, we saw that participants interact with each other in  
 2032 order to perform actions. An action is not itself the same thing as the result of performing the action. When  
 2033 an action is performed against a service, the real world effect that results is reported in the form of  
 2034 notification of events.

### 2035 4.3.3 Message Exchange

2036 *Message exchange* is the means by which service participants (or their delegates) interact with each  
2037 other. There are two primary modes of interaction: joint actions that cause real world effects, and  
2038 notification of events that report real world effects.<sup>12</sup>

2039 A message exchange is used to affect an action when the messages contain the appropriately formatted  
2040 content that should be interpreted as joint action and the delegates involved interpret the message  
2041 appropriately.

2042 A message exchange is also used to communicate event notifications. An event is an occurrence that is  
2043 of interest to some participant; in our case when some real world effect has occurred. Just as action  
2044 messages have formatting requirements, so do event notification messages. In this way, the Information  
2045 Model of a service must specify the syntax (structure), and semantics (meaning) of the action messages  
2046 and event notification messages as part of a service interface. It must also specify the syntax and  
2047 semantics of any data that is carried as part of a payload of the action or event notification message. The  
2048 Information Model is described in greater detail in the Service Description Model (see Section 4.1).

2049 In addition to the Information Model that describes the syntax and semantics of the messages and data  
2050 payloads, exception conditions and error handling in the event of faults (e.g., network outages, improper  
2051 message formats, etc.) must be specified or referenced as part of the Service Description.

2052 When a message is interpreted as an action, the correct interpretation typically requires the receiver to  
2053 perform a set of operations. These *operations* represent the sequence of actions (often private) a service  
2054 must perform in order to validly participate in a given joint action.

2055 Similarly, the correct consequence of realizing a real world effect may be to initiate the reporting of that  
2056 real world effect via an event notification.

#### 2057 **Message Exchange**

2058         The means by which joint action and event notifications are coordinated by service participants  
2059         (or delegates).

#### 2060 **Operations**

2061         The sequence of actions a service must perform in order to validly participate in a given joint  
2062         action.

### 2063 4.3.3.1 Message Exchange Patterns (MEPs)

2064 The SOA-RAF commits to the use of message exchange to denote actions against the services, and to  
2065 denote notification of events that report on real world effects that arise from those actions.

2066 Based on these assumptions, the basic temporal aspect of service interaction can be characterized by  
2067 two fundamental message exchange patterns (MEPs):

- 2068 • Request/response to represent how actions cause a real world effect
- 2069 • Event notification to represent how events report a real world effect

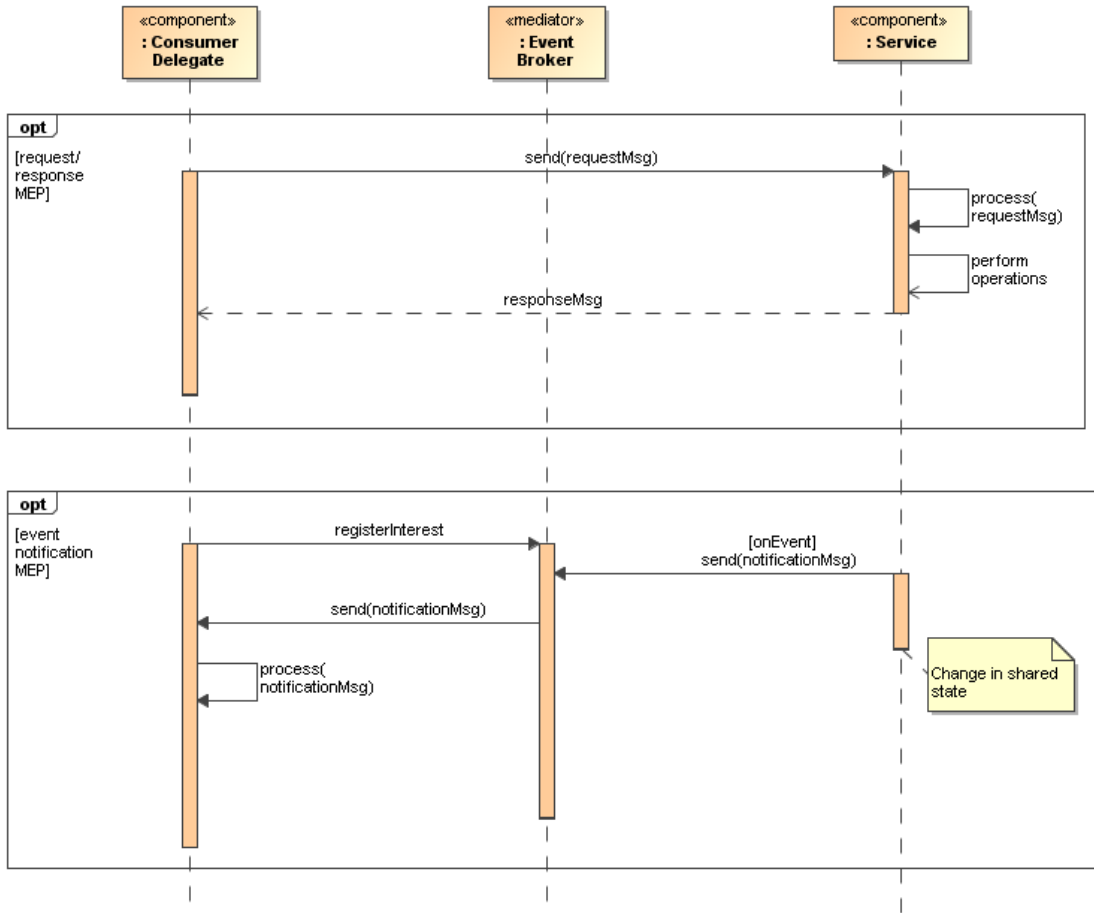
2070 This is by no means a complete list of all possible MEPs used for inter- or intra-enterprise messaging but  
2071 it does represent those that are most commonly used in exchange of information and reporting changes  
2072 in state both within organizations and across organizational boundaries, a hallmark of a SOA.

2073 Recall from the Reference Model that the Process Model characterizes “the temporal relationships  
2074 between and temporal properties of actions and events associated with interacting with the service.”  
2075 Thus, MEPs are a key element of the Process Model. The meta-level aspects of the Process Model (just  
2076 as with the Action Model) are provided as part of the Service Description Model (see Section 4.1).

---

<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.





2077  
2078

Figure 30 Fundamental SOA message exchange patterns (MEPs)

2079 In the UML sequence diagram shown in Figure 30 it is assumed that the service participants (consumer  
 2080 and provider) have delegated message handling to hardware or software delegates acting on their behalf.  
 2081 In the case of the service consumer, this is represented by the *Consumer Delegate* component. In the  
 2082 case of the service provider, the delegate is represented by the *Service* component. The message  
 2083 interchange model illustrated represents a logical view of the MEPs and not a physical view. In other  
 2084 words, specific hosts, network protocols, and underlying messaging system are not shown as these tend  
 2085 to be implementation specific. Although such implementation-specific elements are considered outside  
 2086 the scope of this document, they are important considerations in modeling the SOA execution context.  
 2087 Recall from the Reference Model that the *execution context* of a service interaction is “the set of  
 2088 infrastructure elements, process entities, policy assertions and agreements that are identified as part of  
 2089 an instantiated service interaction, and thus forms a path between those with needs and those with  
 2090 capabilities.”

#### 2091 4.3.3.2 Request/Response MEP

2092 In a request/response MEP, the Consumer Delegate component sends a request message to the Service  
 2093 component. The Service component then processes the request message. Based on the content of the  
 2094 message, the Service component performs the service operations. Following the completion of these  
 2095 operations, a response message is returned to the Consumer Delegate component. The response could

2096 be that a step in a process is complete, the initiation of a follow-on operation, or the return of requested  
2097 information.<sup>13</sup>

2098 Although the sequence diagram shows a *synchronous* interaction (because the sender of the request  
2099 message, i.e., Consumer Delegate, is blocked from continued processing until a response is returned  
2100 from the Service) other variations of request/response are valid, including *asynchronous* (non-blocking)  
2101 interaction through use of queues, channels, or other messaging techniques.

2102 What is important to convey here is that the request/response MEP represents action, which causes a  
2103 real world effect, irrespective of the underlying messaging techniques and messaging infrastructure used  
2104 to implement the request/response MEP.

#### 2105 4.3.3.3 Event Notification MEP


2106 An event is made visible to interested consumers by means of an event notification message exchange  
2107 that reports a real world effect; specifically, a change in shared state between service participants. The  
2108 basic event notification MEP takes the form of a one-way message sent by a notifier component (in this  
2109 case, the Service component) and received by components with an interest in the event (here, the  
2110 Consumer Delegate component).

2111 Often the sending component may not be fully aware of all the components that receive the notification;  
2112 particularly in so-called publish/subscribe (“pub/sub”) situations. In event notification message  
2113 exchanges, it is rare to have a tightly-coupled link between the sending and the receiving component(s)  
2114 for a number of practical reasons. One of the most common is the potential for network outages or  
2115 communication interrupts that can result in loss of notification of events. Therefore, a third-party mediator  
2116 component is often used to decouple the sending and receiving components .

2117 Although this is typically an implementation issue, because this type of third-party decoupling is so  
2118 common in event-driven systems, it is warranted for use in modeling this type of message exchange in  
2119 the SOA-RAF. This third-party intermediary is shown in Figure 30 as an Event Broker mediator. As with  
2120 the request/response MEP, no distinction is made between synchronous versus asynchronous  
2121 communication, although asynchronous message exchange is illustrated in the UML sequence diagram  
2122 depicted in Figure 30 .

#### 2123 4.3.4 Composition of Services

2124 Composition of services is the act of aggregating or “composing” a single service from one or more other  
2125 services. A simple model of service composition is illustrated in Figure 31.



2126

---

<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.

2127 *Figure 31 Simple model of service composition.*

2128 Here, Service A is a service that has an exposed interface IServiceA, which is available to the Consumer  
2129 Delegate and relies on two other services in its implementation. The Consumer Delegate does not know  
2130 that Services B and C are used by Service A, or whether they are used in serial or parallel, or if their  
2131 operations succeed or fail. The Consumer Delegate only cares about the success or failure of Service A.  
2132 The exposed interfaces of Services B and C (IService B and IServiceC) are not necessarily hidden from  
2133 the Consumer Delegate; only the fact that these services are used as part of the composition of Service  
2134 A. In this example, there is no practical reason the Consumer Delegate could not interact with Service B  
2135 or Service C in some other interaction scenario.

2136 It is possible for a service composition to be opaque from one perspective and transparent from another.  
2137 For example, a service may appear to be a single service from the Consumer's Delegate's perspective,  
2138 but is transparently composed of one or more services from a service management perspective. A  
2139 Service Management Service needs to be able to have visibility into the composition in order to properly  
2140 manage the dependencies between the services used in constructing the composite service—including  
2141 managing the service's lifecycle. The subject of services as management entities is described and  
2142 modeled in the *Ownership in a SOA Ecosystem View* of the SOA-RAF and is not further elaborated in this  
2143 section. The point to be made here is that there can be different levels of opacity or transparency  
2144 when it comes to visibility of service composition.

2145 Services can be composed in a variety of ways including direct service-to-service interaction by using  
2146 programming techniques, or they can be aggregated by means of a scripting approach that leverages a  
2147 service composition scripting language. Such scripting approaches are further elaborated in the following  
2148 sub-sections on service-oriented business processes and collaborations.

#### 2149 **4.3.4.1 Service-Oriented Business Processes**

2150 The concepts of business processes and collaborations in the context of transactions and exchanges  
2151 across organizational boundaries are described and modeled as part of the *Participation in a SOA*  
2152 *Ecosystem* view of this reference architecture (see Section **Error! Reference source not found.**). Here,  
2153 e focus on the belief that the principle of composition of services can be applied to business processes  
2154 and collaborations. Of course, business processes and collaborations traditionally represent complex,  
2155 multi-step business functions that may involve multiple participants, including internal users, external  
2156 customers, and trading partners. Therefore, such complexities cannot simply be ignored when  
2157 transforming traditional business processes and collaborations to their service-oriented variants.

#### 2158 **Business Processes**

2159 Business processes are a set of one or more linked activities that are performed to achieve a  
2160 certain business outcome.

2161 Service orientation as applied to business processes (i.e., “service-oriented business processes”) means  
2162 that the aggregation or composition of all of the abstracted activities, flows, and rules that govern a  
2163 business process can themselves be abstracted as a service **[BLOOMBERG/SCHMELZER]**.

2164 When business processes are abstracted in this manner and accessed through SOA services, all of the  
2165 concepts used to describe and model composition of services that were articulated in Section 4.3.4 apply.  
2166 There are some important differences from a composite service that represents an abstraction of a  
2167 business process from a composite service that represents a single-step business interaction. As stated  
2168 earlier, business processes have temporal properties and can range from short-lived processes that  
2169 execute on the order of minutes or hours to long-lived processes that can execute for weeks, months, or  
2170 even years. Further, these processes may involve many participants. These are important  
2171 considerations for the consumer of a service-oriented business process and these temporal properties  
2172 must be articulated as part of the meta-level aspects of the service-oriented business process in its  
2173 Service Description, along with the meta-level aspects of any sub-processes that may be of use or need  
2174 to be visible to the service consumer.

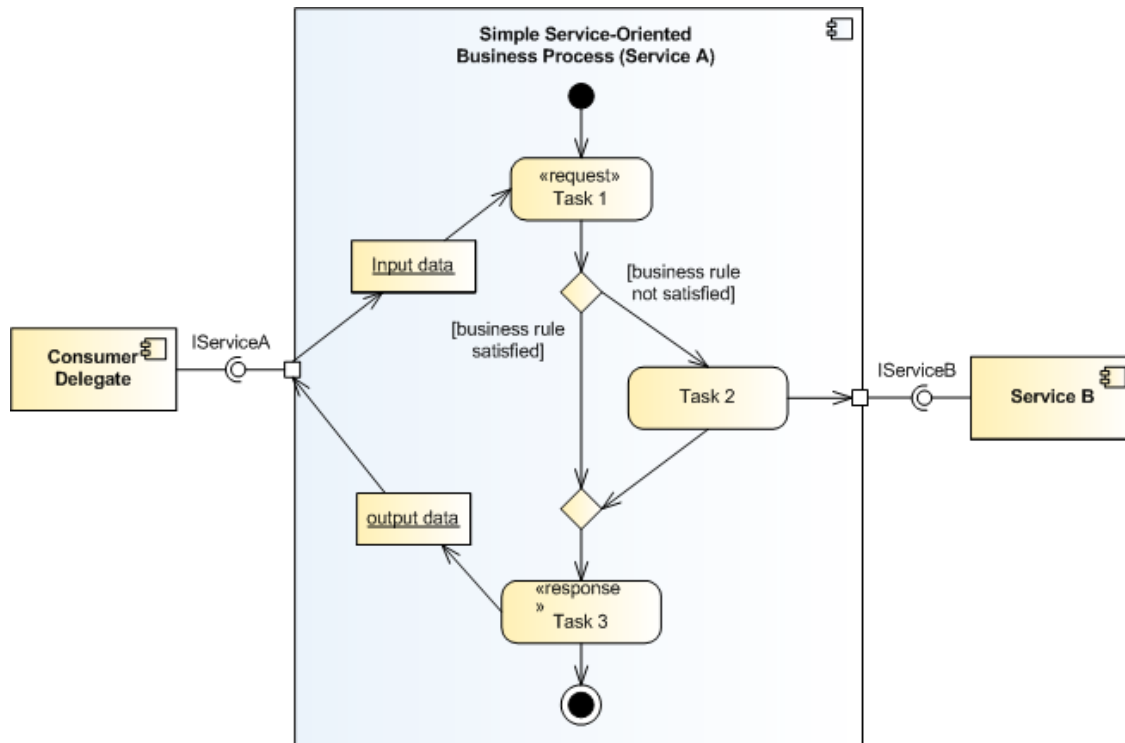
2175 In addition, a workflow activity represents a unit of work that some entity acting in a described role (i.e.,  
2176 role player) is asked to perform. Activities can be broken down into steps with each step representing a  
2177 task for the role player to perform. A technique that is used to compose service-oriented business  
2178 processes that are hierarchical (top-down) and self-contained in nature is known as *orchestration*.

2179 **Orchestration**

2180 A technique used to compose service-oriented business processes that are executed and  
2181 coordinated by an actor acting as “conductor.”

2182 An orchestration is typically implemented using a scripting approach to compose service-oriented  
2183 business processes. This typically involves use of a standards-based orchestration scripting language.  
2184 In terms of automation, an orchestration can be mechanized using a business process orchestration  
2185 engine, which is a hardware or software component (delegate) responsible for acting in the role of central  
2186 conductor/coordinator responsible for executing the flows that comprise the orchestration.

2187 A simple generic example of such an orchestration is illustrated in Figure 32.



2188  
2189 *Figure 32 Abstract example of orchestration of service-oriented business process.*

2190 Here, we use a UML activity diagram to model the simple service-oriented business process as it allows  
2191 us to capture the major elements of business processes such as the set of related tasks to be performed,  
2192 linking between tasks in a logical flow, data that is passed between tasks, and any relevant business  
2193 rules that govern the transitions between tasks. A task is a unit of work that an individual, system, or  
2194 organization performs and can be accomplished in one or more steps or subtasks. While subtasks can  
2195 be readily modeled, they are not illustrated in the orchestration model in Figure 32..

2196 This particular example is based on a request/response MEP and captures how one particular task (Task  
2197 2) actually utilizes an externally-provided service, Service B. The entire service-oriented business  
2198 process is exposed as Service A that is accessible via its externally visible interface, IServiceA.

2199 Although not explicitly shown in the orchestration model above, it is assumed that there exists a software  
2200 or hardware component, i.e., orchestration engine that executes the process flow. Recall that a central  
2201 concept to orchestration is that process flow is coordinated and executed by a single conductor;  
2202 hence the name “orchestration.”

2203 **4.3.4.2 Service-Oriented Business Collaborations**

2204 Business collaborations typically represent the interaction involved in executing business transactions,  
2205 where a business transaction is defined in the *Participation in a SOA Ecosystem* view as “a joint action  
2206 engaged in by two or more participants in which resources are exchanged” (see Section **Error!**  
2207 **Reference source not found.**).

2208 It is important to note that business collaborations represent “peer”-style interactions; in other words,  
 2209 peers in a business collaboration act as equals. This means that unlike the orchestration of business  
 2210 processes, there is no single or central entity that coordinates or “conducts” a business collaboration.  
 2211 These peer styles of interactions typically occur between trading partners that span organizational  
 2212 boundaries.

2213 Business collaborations can also be service-enabled. For purposes of this Reference Architecture  
 2214 Foundation, we refer to these as “service-oriented business collaborations.” Service-oriented business  
 2215 collaborations do not necessarily imply exposing the entire peer-style business collaboration as a service  
 2216 itself but rather the collaboration uses service-based interchanges.

2217 The technique that is used to compose service-oriented business collaborations in which multiple parties  
 2218 collaborate in a peer-style as part of some larger business transaction by exchanging messages with  
 2219 trading partners and external organizations (e.g., suppliers) is known as *choreography*

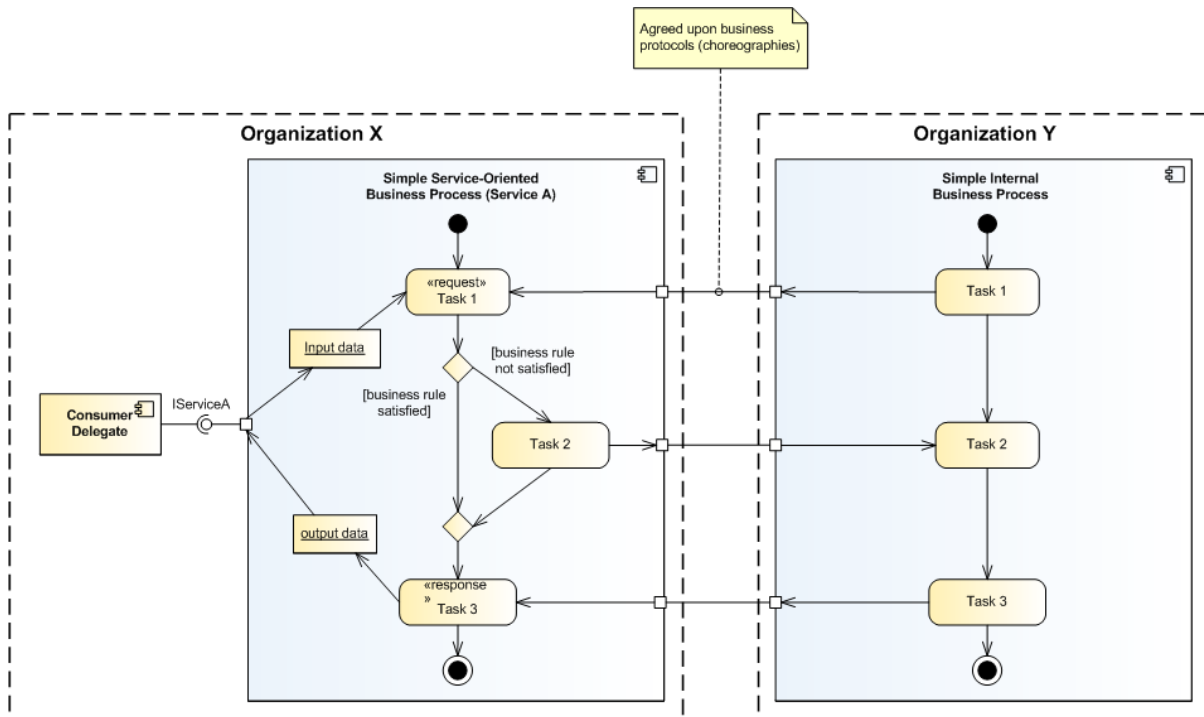
2220 **[NEWCOMER/LOMOW].**

2221 **Choreography**

2222 A technique used to characterize service-oriented business collaborations based on ordered  
 2223 message exchanges between peer entities in order to achieve a common business goal.

2224 Choreography differs from orchestration primarily in that each party in a business collaboration describes  
 2225 its part in the service interaction. Note that choreography as we have defined it here should not be  
 2226 confused with the term *process choreography*, which is defined in the *Participation in a SOA Ecosystem*  
 2227 view as “the description of the possible interactions that may take place between two or more participants  
 2228 to fulfill an objective.” This is an example of domain-specific nomenclature that often leads to confusion  
 2229 and why we are making note of it here.

2230 A simple generic example of a choreography is illustrated in Figure 33



2231 Figure 33 Abstract example of choreography of service-oriented business collaboration.

2232 This example, which is a variant of the orchestration example illustrated earlier in Figure 32 adds trust  
 2233 boundaries between two organizations; namely, Organization X and Organization Y. It is assumed that  
 2234 these two organizations are peer entities that have an interest in a business collaboration, for example,  
 2235 Organization X and Organization Y could be trading partners. Organization X retains the service-oriented  
 2236 business process Service A, which is exposed to internal consumers via its provided service interface,  
 2237 IServiceA. Organization Y also has a business process that is involved in the business collaboration;

2239 however, for this example, it is an internal business process that is not exposed to potential consumers  
2240 either within or outside its organizational boundary.

2241 The scripting language that is used for the choreography needs to define how and when to pass control  
2242 from one trading partner to another, i.e., Organization X and Organization Y. Defining the business  
2243 protocols used in the business collaboration involves precisely specifying the visible message exchange  
2244 behavior of each of the parties involved in the protocol, without revealing internal implementation details  
2245 **[NEWCOMER/LOMOW]**.

2246 In a peer-style business collaboration, a choreography scripting language must be capable of describing  
2247 the coordination of those service-oriented processes that cross organizational boundaries.

## 2248 **4.3.5 Architectural Implications of Interacting with Services**

2249 Interacting with Services has the following architectural implications on mechanisms that facilitate service  
2250 interaction:

- 2251 • A well-defined service Information Model that:
    - 2252 ○ describes the syntax and semantics of the messages used to denote actions and events;
2253 ○ describes the syntax and semantics of the data payload(s) contained within messages;
2254 ○ documents exception conditions in the event of faults due to network outages, improper  
2255 message/data formats, etc.;2256 ○ is both human readable and machine processable;2257 ○ is referenceable from the Service Description artifact.
- 2258 • A well-defined service Behavior Model that:
  - 2259 ○ characterizes the knowledge of the actions invokes against the service and events that report  
2260 real world effects as a result of those actions;
2261 ○ characterizes the temporal relationships and temporal properties of actions and events  
2262 associated in a service interaction;
- 2263 ○ describe activities involved in a workflow activity that represents a unit of work;2264 ○ describes the role (s) that a role player performs in a service-oriented business process or
- 
- 2265 service-oriented business collaboration;2266 ○ is both human readable and machine processable;2267 ○ is referenceable from the Service Description artifact.
- 2268 • Service composition mechanisms to support orchestration of service-oriented business processes and  
2269 choreography of service-oriented business collaborations such as:
  - 2270 ○ Declarative and programmatic compositional languages;
2271 ○ Orchestration and/or choreography engines that support multi-step processes as part of a  
2272 short-lived or long-lived business transaction;
- 2273 ○ Orchestration and/or choreography engines that support compensating transactions in
- 
- 2274 the presences of exception and fault conditions.
- 2275 • Infrastructure services that provides mechanisms to support service interaction, including but not  
2276 limited to:
  - 2277 ○ mediation services such as message and event brokers, providers, and/or buses that  
2278 provide message translation/transformation, gateway capability, message persistence,  
2279 reliable message delivery, and/or intelligent routing semantics;
2280 ○ binding services that support translation and transformation of multiple application-level  
2281 protocols to standard network transport protocols;
- 2282 ○ auditing and logging services that provide a data store and mechanism to record
- 
- 2283 information related to service interaction activity such as message traffic patterns,
- 
- 2284 security violations, and service contract and policy violations2285 ○ security services that abstract techniques such as public key cryptography, secure
- 
- 2286 networks, virus protection, etc., which provide protection against common security threats
- 
- 2287 in a SOA ecosystem;2288 ○ monitoring services such as hardware and software mechanisms that both monitor the
- 
- 2289 performance of systems that host services and network traffic during service interaction,
- 
- 2290 and are capable of generating regular monitoring reports.
- 2291 • A layered and tiered service component architecture that supports multiple message exchange  
2292 patterns (MEPs) in order to:

- 2293 ○ promote the industry best practice of separation of concerns that facilitates flexibility in
- 2294 the presence of changing business requirements;
- 2295 ○ promote the industry best practice of separation of roles in a service development
- 2296 lifecycle such that subject matter experts and teams are structured along areas of
- 2297 expertise;
- 2298 ○ support numerous standard interaction patterns, peer-to-peer interaction patterns,
- 2299 enterprise integration patterns, and business-to-business integration patterns.

## 2300 4.4 Policies and Contracts Model

2301 A common phenomenon of many machines and systems is that the scope of potential behavior is much  
 2302 broader than is actually needed for a particular circumstance. This is especially true of a system as  
 2303 powerful as a SOA ecosystem. As a result, the behavior and performance of the system tend to be  
 2304 under-constrained by the implementation; instead, the actual behavior is expressed by means of policies  
 2305 of some form. Policies define the choices that stakeholders make; these choices are used to guide the  
 2306 actual behavior of the system to the desired behavior and performance.

2307 As noted in Section 3.1.5 a policy is a constraint of some form that is promulgated by a stakeholder who  
 2308 has the responsibility of ensuring that the constraint is enforced. In contrast, contracts are **agreements**  
 2309 between participants. However, like policies, it is a necessary part of contracts that they are enforceable.

2310 While responsibility for enforcement may differ, both contracts and policies share a common characteristic  
 2311 – there is a **constraint** that must be enforced. In both cases the mechanisms needed to enforce policy  
 2312 constraints are likely to be identical; in this model we focus on the issues involved in representing policies  
 2313 and contracts and on some of the principles behind their enforcement.

### 2314 4.4.1 Policy and Contract Representation

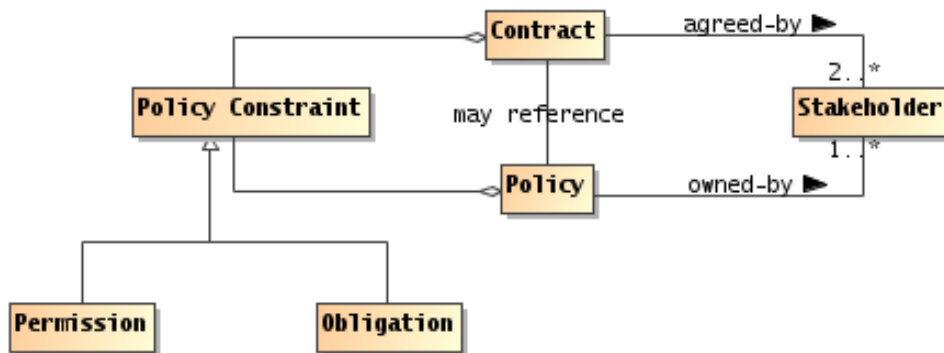
2315 A **policy constraint** is a specific kind of constraint: the ontology of policies and contracts includes the  
 2316 core concepts of permission, obligation, owner, subject. In addition, it may be necessary to be able  
 2317 combine policy constraints and to be able to resolve policy conflicts.

#### 2318 4.4.1.1 Policy Framework

##### 2319 Policy Framework

2320 A policy framework is a language in which policy constraints may be expressed.

2321 A policy framework combines a syntax for expressing policy constraints together with a decision  
 2322 procedure for determining if a policy constraint is satisfied.



2323  
 2324 *Figure 34 Policies and Contracts*

2325 We can characterize (caricature) a policy framework in terms of a logical framework and an ontology of  
 2326 policies. The policy ontology details specific kinds of policy constraints that can be expressed; and the  
 2327 logical framework is a 'glue' that allows us to express combinations of policies.

- 2328 **Logical Framework**
- 2329 A logical framework is a linguistic framework consisting of a syntax – a way of writing expressions  
2330 – and a semantics – a way of interpreting the expressions.
- 2331 **Policy Ontology**
- 2332 A policy ontology is a formalization of a set of concepts that are relevant to forming policy  
2333 expressions.
- 2334 For example, a policy ontology that allows to identify simple constraints – such as the existence of a  
2335 property, or that a value of a property should be compared to a fixed value – is often enough to express  
2336 many basic constraints.
- 2337 Included in many policy ontologies are the basic signals of permissions and obligations. Some policy  
2338 frameworks are sufficiently constrained that there is not possibility of representing an obligation; in which  
2339 case there is often no need to ‘call out’ the distinction between permissions and obligations.
- 2340 The logical framework is also a strong determiner of the expressivity of the policy framework. The richer  
2341 the logical framework, the richer the set of policy constraints that can be expressed. However, there is a  
2342 strong inverse correlation between expressivity and ease and efficiency of implementation.
- 2343 In the discussion that follows we assume the following basic policy ontology:
- 2344 **Policy Owner**
- 2345 A policy owner is a stakeholder that asserts and enforces the policy.
- 2346 **Policy Subject**
- 2347 A policy subject is an actor who is subject to the constraints of a policy or contract.
- 2348 **Policy Constraint**
- 2349 A policy constraint is a measurable proposition that characterizes the constraint that the policy is  
2350 about.
- 2351 **Policy Object**
- 2352 A policy object is an identifiable state, action or resource that is potentially constrained by the  
2353 policy.
- 2354 **4.4.2 Policy and Contract Enforcement**
- 2355 The enforcement of policy constraints has to address two core problems: how to enforce the atomic policy  
2356 constraints, and how to enforce combinations of policy constraints. In addition, it is necessary to address  
2357 the resolution of policy conflicts.
- 2358 **4.4.2.1 Enforcing Simple Policy Constraints**
- 2359 The two primary kinds of policy constraint – permission and obligation – naturally lead to different styles  
2360 of enforcement. A permission constraint must typically be enforced *prior* to the policy subject invoking the  
2361 **policy object**. On the hand, an obligation constraint must typically be enforced post-facto through some  
2362 form of auditing process and remedial action.
- 2363 For example, if a communications policy required that all communication be encrypted, this is enforceable  
2364 at the point of communication: any attempt to communicate a message that is not encrypted can be  
2365 blocked.
- 2366 Similarly, an obligation to pay for services rendered is enforced by ensuring that payment arrives within a  
2367 reasonable period of time. Invoices are monitored for prompt (or lack of) payment.
- 2368 The key concepts in enforcing both forms of policy constraint are the policy decision and the policy  
2369 enforcement.
- 2370 **Policy Decision**
- 2371 A policy decision is a determination as to whether a given policy constraint is satisfied or not.



2372 A policy decision is effectively a measurement of some state – typically a portion of the SOA ecosystem’s  
2373 **shared state**. This implies a certain *timeliness* in the measuring: a measurement that is too early or is too  
2374 late does not actually help in determining if the policy constraint is satisfied appropriately.

#### 2375 **Policy Enforcement**

2376 A policy enforcement is the use of a mechanism to limit the behavior and/or state of policy  
2377 subjects to comply with a policy decision.

2378 A policy enforcement implies the use of some mechanism to ensure compliance with a policy decision.  
2379 The range of mechanisms is completely dependent on the kinds of atomic policy constraints that the  
2380 policy framework may support. As noted above, the two primary styles of constraint – permission and  
2381 **obligation** –lead to different styles of enforcement.

#### 2382 **4.4.2.2 Enforcing Policy Combinations**

2383 Enforcing policy combinations is primarily an elaboration of enforcing simple policy constraints. The  
2384 process of policy decisions is enhanced to allow a measurement to involve combinations of policy  
2385 constraints and the process of policy enforcement may need to be enhanced to coordinate the  
2386 enforcement of multiple policy constraints simultaneously.

#### 2387 **4.4.2.3 Conflict Resolution**

2388 Whenever it is possible that more than one policy constraint applies in a given situation, there is the  
2389 potential that the policies themselves are not mutually consistent. For example, a policy that requires  
2390 communication to be encrypted and a policy that requires an administrator to read every communication  
2391 conflict with each other – the two policies cannot both be satisfied.

2392 In general, with sufficiently rich policy frameworks, it is not possible to always resolve policy conflicts  
2393 automatically. However, a reasonable approach is to augment the policy decision process with simple  
2394 policy conflict resolution rules; with the potential for *escalating* a policy conflict to human adjudication.

#### 2395 **Policy Conflict**

2396 A policy conflict exists between two or more policies in a policy decision process if it is not  
2397 possible to satisfy all the policies that apply.

#### 2398 **Policy Conflict Resolution**

2399 A policy conflict resolution rule is a way of determining which policy should prevail in a policy  
2400 conflict.

2401 The inevitable consequence of policy conflicts is that it is not possible to guarantee that all policies are  
2402 satisfied at all times. This, in turn, implies a certain *flexibility* in the application of policy constraints: they  
2403 will not always be honored.

#### 2404 **4.4.3 Architectural Implications**

2405 The key choices that must be made in a system of policies center around the policy framework and policy  
2406 enforcement mechanisms

- 2407 • There SHOULD be a standard policy framework that is adopted across the SOA ecosystem:
  - 2408 ○ This framework MUST permit the expression of simple policy constraints
  - 2409 ○ The framework MAY allow (to a varying extent) the combination of policy constraints,  
2410 including
    - 2411 • Both positive and negative constraints
    - 2412 • Conjunctions and disjunctions of constraints
    - 2413 • The quantification of constraints
  - 2414 ○ The framework MUST at least allow the policy subject and the policy object to be identified as  
2415 well as the policy constraint.
  - 2416 ○ The framework MAY allow further structuring of policies into modules, inheritance between  
2417 policies and so on.
- 2418 • There SHOULD be mechanisms that facilitate the application of policies:

- 2419 ○ There SHOULD be mechanisms that allow policy decisions to be made, consistent with the
- 2420 policy frameworks and with the state of the SOA ecosystem.
- 2421 ○ There SHOULD be mechanisms to enforce policy decisions
- 2422 ● There SHOULD be mechanisms to support the measurement of whether certain
- 2423 policy constraints are satisfied or not, or to what degree they are satisfied.
- 2424 ● Such enforcement mechanisms MAY include support for both permission-style
- 2425 constraints and obligation-style constraints.
- 2426 ● Enforcement mechanisms MAY support the simultaneous enforcement of multiple
- 2427 policy constraints across multiple points in the SOA ecosystem.
- 2428 ○ There SHOULD be mechanisms to resolve policy conflicts
- 2429 ● This MAY involve escalating policy conflicts to human adjudication.
- 2430 ○ There SHOULD be mechanisms that support the management and promulgation of policies.

## 2431 5 Ownership in a SOA Ecosystem View

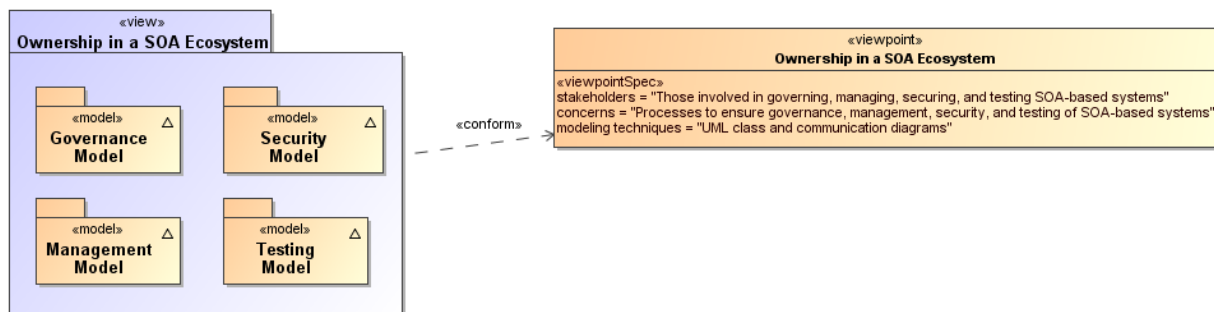
2432 *Governments are instituted among Men,*  
2433 *deriving their just power from the consent of the governed*  
2434 *American Declaration of Independence*  
2435

2436 The *Owning Service Oriented Architectures View* focuses on the issues, requirements and responsibilities  
2437 involved in owning a SOA-based system.

2438 Owning a SOA-based system raises significantly different challenges to owning other complex systems --  
2439 such as Enterprise suites -- because there are strong limits on the control and authority of any one party  
2440 when a system spans multiple ownership domains.

2441 Even when a SOA-based system is deployed internally within an organization, there are multiple internal  
2442 stakeholders involved and there may not be a simple hierarchy of control and management. Thus, an  
2443 early consideration of how multiple boundaries affect SOA-based systems provides a firm foundation for  
2444 dealing with them in whatever form they are found rather than debating whether the boundaries should  
2445 exist.

2446 This view focuses on the Governance of SOA-based systems, on the security challenges involved in  
2447 running a SOA-based system and the management challenges.



2448  
2449 *Figure 35 Model Elements Described in the Ownership in a SOA Ecosystem View*

2450 The following subsections present models of these functions.

### 2451 5.1 Governance Model

2452 The Reference Model defines Service Oriented Architecture as an architectural paradigm for organizing  
2453 and utilizing distributed capabilities that may be under the control of different ownership domains [**SOA-**  
2454 **RM**]. Consequently, it is important that organizations that plan to engage in service interactions adopt  
2455 governance policies and procedures sufficient to ensure that there is standardization across both internal  
2456 and external organizational boundaries to promote the effective creation and use of SOA-based services.

#### 2457 5.1.1 Understanding Governance

##### 2458 5.1.1.1 Terminology

2459 Governance is about making decisions that are aligned with the overall organizational strategy and  
2460 culture of the enterprise. [**Gartner**] It specifies the decision rights and accountability framework to  
2461 encourage desirable behaviors [**Weill/Ross-MIT Sloan School**] towards realizing the strategy and  
2462 defines incentives (positive or negative) towards that end. It is less about overt control and strict  
2463 adherence to rules, and more about guidance and effective and equitable usage of resources to ensure  
2464 sustainability of an organization's strategic objectives. [**TOGAF v8.1**]

2465 To accomplish this, governance requires organizational structure and processes and must identify who  
2466 has authority to define and carry out its mandates. It must address the following questions: 1) what  
2467 decisions must be made to ensure effective management and use?, 2) who should make these

2468 decisions?, and 3) how will these decisions be made and monitored? , and (4) how will these decisions  
2469 be communicated? The intent is to achieve goals, add value, and reduce risk.

2470 Within a single ownership domain such as an enterprise, generally there is a hierarchy of governance  
2471 structures. Some of the more common enterprise governance structures include corporate governance,  
2472 technology governance, IT governance, and architecture governance **[TOGAF v8.1]**. These governance  
2473 structures can exist at multiple levels (global, regional, and local) within the overall enterprise.

2474 It is often asserted that SOA governance is a specialization of IT governance as there is a natural  
2475 hierarchy of these types of governance structures; however, the focus of SOA governance is less on  
2476 decisions to ensure effective management and use of IT as it is to ensure effective management and use  
2477 of SOA-based systems. Certainly, SOA governance must still answer the basic questions also  
2478 associated with IT governance, i.e., who should make the decisions, and how these decisions will be  
2479 made and monitored.

### 2480 **5.1.1.2 Relationship to Management**

2481 There is often confusion centered on the relationship between governance and management. As  
2482 described earlier, governance is concerned with decision making. Management, on the other hand, is  
2483 concerned with execution. Put another way, governance describes the world as leadership wants it to be;  
2484 management executes activities that intends to make the leadership's desired world a reality. Where  
2485 governance determines who has the authority and responsibility for making decisions and the  
2486 establishment of guidelines for how those decisions should be made, management is the actual process  
2487 of making, implementing, and measuring the impact of those decisions **[Loeb]**. Consequently,  
2488 governance and management work in concert to ensure a well-balanced and functioning organization as  
2489 well as an ecosystem of inter-related organizations. In the sections that follow, we elaborate further on  
2490 the relationship between governance and management in terms of setting and enforcing service policies,  
2491 contracts, and standards as well as addressing issues surrounding regulatory compliance.

### 2492 **5.1.1.3 Why is SOA Governance Important?**

2493 One of the hallmarks of SOA that distinguishes it from other architectural paradigms for distributed  
2494 computing is the ability to provide a uniform means to offer, discover, interact with and use capabilities  
2495 (as well the ability to compose new capabilities from existing ones) all in an environment that transcends  
2496 domains of ownership. Consequently, ownership, and issues surrounding it, such as obtaining  
2497 acceptable terms and conditions (T&Cs) in a contract, is one of the primary topics for SOA governance.  
2498 Generally, IT governance does not include T&Cs, for example, as a condition of use as its primary  
2499 concern.

2500 Just as other architectural paradigms, technologies, and approaches to IT are subject to change and  
2501 evolution, so too is SOA. Setting policies that allow change management and evolution, establishing  
2502 strategies for change, resolving disputes that arise, and ensuring that SOA-based systems continue to  
2503 fulfill the goals of the business are all reasons why governance is important to SOA.

### 2504 **5.1.1.4 Governance Stakeholders and Concerns**

2505 As noted in Section **Error! Reference source not found.** the participants in a service interaction include  
2506 he service provider, the service consumer, and other interested or unintentional third parties. Depending  
2507 on the circumstances, it may also include the owners of the underlying capabilities that the SOA services  
2508 access. Governance must establish the policies and rules under which duties and responsibilities are  
2509 defined and the expectations of participants are grounded. The expectations include transparency in  
2510 aspects where transparency is mandated, trust in the impartial and consistent application of governance,  
2511 and assurance of reliable and robust behavior throughout the SOA ecosystem.

## 2512 **5.1.2 A Generic Model for Governance**


### 2513 **Governance**

2514 Governance is the prescribing of conditions and constraints consistent with satisfying common  
2515 goals and the structures and processes needed to define and respond to actions taken towards  
2516 realizing those goals.

2517 The following is a generic model of governance represented by segmented models that begin with  
 2518 motivation and proceed through measuring compliance. It is not all-encompassing but a focused subset  
 2519 that captures the aspects necessary to describe governance for SOA. It does not imply that practical  
 2520 application of governance is a single, isolated instance of these models; in reality, there may be  
 2521 hierarchical and parallel chains of governance that deal with different aspects or focus on different goals.  
 2522 This is discussed further in section 5.1.2.5. The defined models are simultaneously applicable to each of  
 2523 the overlapping instances.

2524 A given enterprise may already have portions of these models in place. To a large extent, the models  
 2525 shown here are not specific to SOA; discussions on direct applicability begin in section 5.1.3.

2526 **5.1.2.1 Motivating Governance**



2527  
 2528 *Figure 36 Motivating governance model*

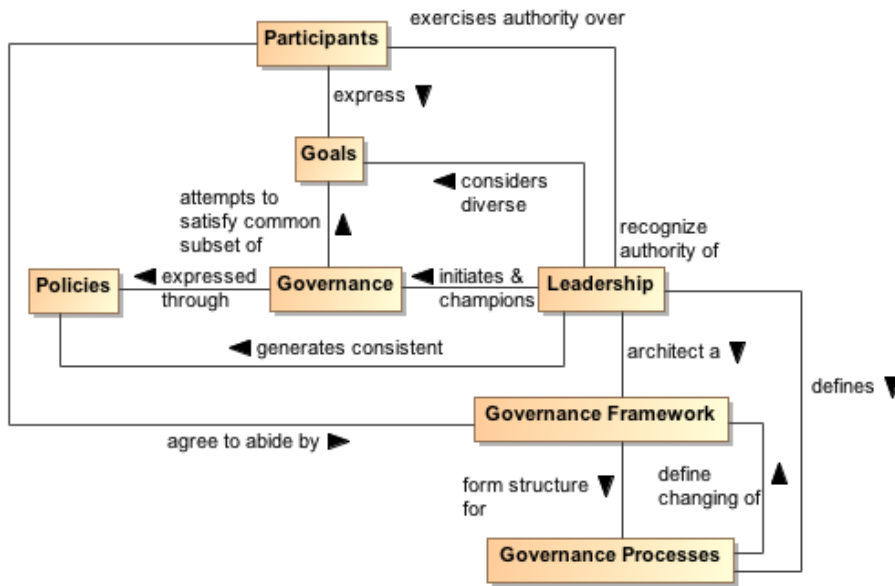
2529 An organizational domain such as an enterprise is made up of participants who may be individuals or  
 2530 groups of individuals forming smaller organizational units within the enterprise. The overall business  
 2531 strategy should be consistent with the Goals of the participants; otherwise, the business strategy would  
 2532 not provide value to the participants and governance towards those ends becomes difficult if not  
 2533 impossible. This is not to say that an instance of governance simultaneously satisfies all the goals of all  
 2534 the participants; rather, the goals of any governance instance must sufficiently satisfy a useful subset of  
 2535 each participant's goals so as to provide value and ensure the cooperation of all the participants.

2536 A policy is the formal characterization of the conditions and constraints that governance deems as  
 2537 necessary to realize the goals which it is attempting to satisfy. Policy may identify required conditions or  
 2538 actions or may prescribe limitations or other constraints on permitted conditions or actions. For example,  
 2539 a policy may prescribe that safeguards must be in place to prevent unauthorized access to sensitive  
 2540 material. It may also prohibit use of computers for activities unrelated to the specified work assignment.  
 2541 Policy is made operational through the promulgating and implementing of Rules and Regulations (as  
 2542 defined in section 5.1.2.3).

2543 As noted in section 4.4.2, policy may be asserted by any participant or on behalf of the participant by its  
 2544 organization. Part of the purpose of governance is to arbitrate among diverse goals of participants and  
 2545 diverse policies articulated to realize those goals. The intent is to form a consistent whole that allows  
 2546 governance to minimize ambiguity about its purpose. While resolving all ambiguity would be an ideal, it is  
 2547 unlikely that all inconsistencies will be identified and resolved before governance becomes operational.

2548 For governance to have effective jurisdiction over participants, there must be some degree of agreement  
 2549 by all participants that they will abide by the governance mandates. A minimal degree of agreement often  
 2550 presages participants who "slow-roll" if not actively reject complying with Policies that express the  
 2551 specifics of governance.

2552 **5.1.2.2 Setting Up Governance**



2553  
2554 *Figure 37 Setting up governance model*

2555 **Leadership**

2556 Leadership is the entity who has the responsibility and authority to generate consistent policies  
2557 through which the goals of governance can be expressed and to define and champion the  
2558 structures and processes through which governance is realized.

2559 **Governance Framework**

2560 The Governance Framework is a set of organizational structures that enable governance to be  
2561 consistently defined, clarified, and as needed, modified to respond to changes in its domain of  
2562 concern.

2563 **Governance Processes**

2564 Governance Processes are the defined set of activities that are performed within the Governance  
2565 Framework to enable the consistent definition, application, and as needed, modification of Rules  
2566 that organize and regulate the activities of participants for the fulfillment of expressed policies.  
2567 (See section 5.1.2.3 for elaboration on the relationship of Governance Processes and Rules.)

2568 As noted earlier, governance requires an appropriate organizational structure and identification of who  
2569 has authority to make governance decisions. In Figure 37, the entity with governance authority is  
2570 designated the Leadership. This is someone, possibly one or more of the participants, that participants  
2571 recognize as having authority for a given purpose or over a given set of issues or concerns.

2572 The Leadership is responsible for prescribing or delegating a working group to prescribe the Governance  
2573 Framework that forms the structure for Governance Processes which define how governance is to be  
2574 carried out. This does not itself define the specifics of how governance is to be applied, but it does  
2575 provide an unambiguous set of procedures that should ensure consistent actions which participants agree  
2576 are fair and account for sufficient input on the subjects to which governance is applied.

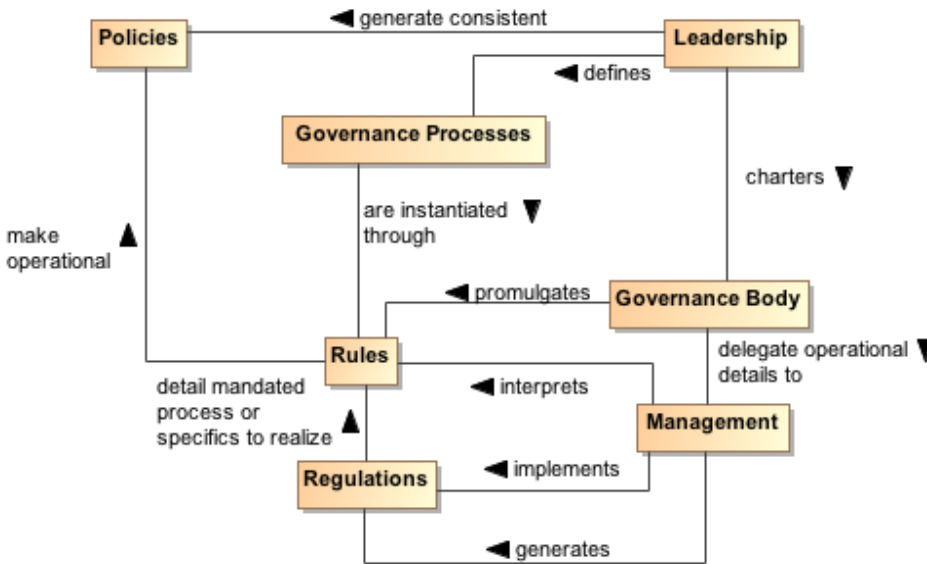
2577 The participants may be part of the working group that codifies the Governance Framework and  
2578 Processes. When complete, the participants must acknowledge and agree to abide by the products  
2579 generated through application of this structure.

2580 The Governance Framework and Processes are often documented in the charter of a body created or  
2581 designated to oversee governance. This is discussed further in the next section. Note that the  
2582 Governance Processes should also include those necessary to modify the Governance Framework itself.

2583 An important function of Leadership is not only to initiate but also be the consistent champion of  
2584 governance. Those responsible for carrying out governance mandates must have Leadership who

2585 makes it clear to participants that expressed Policies are seen as a means to realizing established goals  
2586 and that compliance with governance is required.

### 2587 5.1.2.3 Carrying Out Governance



2588  
2589 *Figure 38 Carrying out governance model*

#### 2590 **Rule**

2591 A Rule is a prescribed guide for carrying out activities and processes leading to desired results,  
2592 e.g. the operational realization of policies.

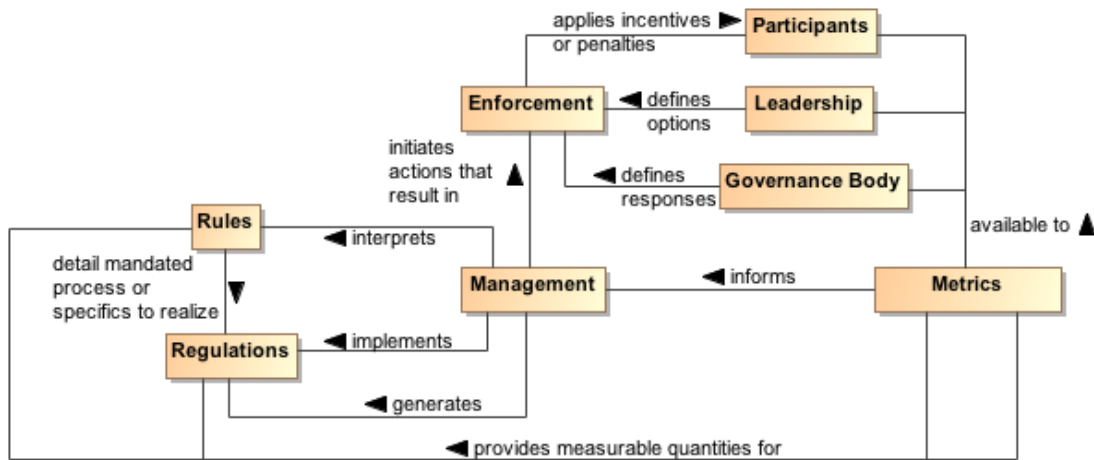
#### 2593 **Regulation**

2594 A Regulation is a mandated process or the specific details that derive from the interpretation of  
2595 Rules and lead to measurable quantities against which compliance can be measured.

2596 To carry out governance, Leadership charts a Governance Body to promulgate the Rules needed to  
2597 make the Policies operational. The Governance Body acts in line with Governance Processes for its rule-  
2598 making process and other functions. Whereas Governance is the setting of Policies and defining the  
2599 Rules that provide an operational context for Policies, the operational details of governance may be  
2600 delegated by the Governance Body to Management. Management generates Regulations that specify  
2601 details for Rules and other procedures to implement both Rules and Regulations. For example,  
2602 Leadership could set a Policy that all authorized parties should have access to data, the Governance  
2603 Body would promulgate a Rule that PKI certificates are required to establish identity of authorized parties,  
2604 and Management can specify a Regulation of who it deems to be a recognized PKI issuing body. In  
2605 summary, Policy is a predicate to be satisfied and Rules prescribe the activities by which that satisfying  
2606 occurs. A number of rules may be required to satisfy a given policy; the carrying out of a rule may  
2607 contribute to several policies being realized.

2608 Whereas the Governance Framework and Processes are fundamental for having participants  
2609 acknowledge and commit to compliance with governance, the Rules and Regulations provide operational  
2610 constraints which may require resource commitments or other levies on the participants. It is important  
2611 for participants to consider the framework and processes to be fair, unambiguous, and capable of being  
2612 carried out in a consistent manner and to have an opportunity to formally accept or ratify this situation.  
2613 Rules and Regulations, however, do not require individual acceptance by any given participant although  
2614 some level of community comment may be part of the Governance Processes. Having agreed to  
2615 governance, the participants are bound to comply or be subject to prescribed mechanisms for  
2616 enforcement.

2617 **5.1.2.4 Ensuring Governance Compliance**



2618  
2619 *Figure 39 Ensuring governance compliance model*

2620 Setting Rules and Regulations does not ensure effective governance unless compliance can be  
2621 measured and Rules and Regulations can be enforced. Metrics are those conditions and quantities that  
2622 can be measured to characterize actions and results. Rules and Regulations MUST be based on  
2623 collected Metrics or there is no means for Management to assess compliance. The Metrics are available  
2624 to the participants, the Leadership, and the Governance Body so what is measured and the results of  
2625 measurement are clear to everyone.

2626 The Leadership in its relationship with participants has certain options that can be used for Enforcement.  
2627 A common option may be to effect future funding. The Governance Body defines specific enforcement  
2628 responses, such as what degree of compliance is necessary for full funding to be restored. It is up to  
2629 Management to identify compliance shortfalls and to initiate the Enforcement process.

2630 Note, enforcement does not strictly need to be negative consequences. Management can use Metrics to  
2631 identify exemplars of compliance and Leadership can provide options for rewarding the participants. The  
2632 Governance Body defines awards or other incentives.

2633 **5.1.2.5 Considerations for Multiple Governance Chains**

2634 As noted in section 5.1.2, instances of the governance model often occur as a tiered arrangement, with  
2635 governance at some level delegating specific authority and responsibility to accomplish a focused portion  
2636 of the original level's mandate. For example, a corporation may encompass several lines of business and  
2637 each line of business governs its own affairs in a manner that is consistent with and contributes to the  
2638 goals of the parent organization. Within the line of business, an IT group may be given the mandate to  
2639 provide and maintain IT resources, giving rise to IT governance.

2640 In addition to tiered governance, there may be multiple governance chains working in parallel. For  
2641 example, a company making widgets has policies intended to ensure they make high quality widgets and  
2642 make an impressive profit for their shareholders. On the other hand, Sarbanes-Oxley is a parallel  
2643 governance chain in the United States that specifies how the management must handle its accounting  
2644 and information that needs to be given to its shareholders. The parallel chains may just be additive or  
2645 may be in conflict and require some harmonization.

2646 Being distributed and representing different ownership domains, a SOA participant falls under the  
2647 jurisdiction of multiple governance domains simultaneously and may individually need to resolve  
2648 consequent conflicts. The governance domains may specify precedence for governance conformance or  
2649 it may fall to the discretion of the participant to decide on the course of actions they believe appropriate.



## 2650 **5.1.3 Governance Applied to SOA**

### 2651 **5.1.3.1 Where SOA Governance is Different**

2652 SOA governance is often discussed in terms of IT governance, but rather than a parent-child relationship,  
2653 Figure 40 shows the two as siblings of the general governance described in section 5.1.2. There are  
2654 obvious dependencies and a need for coordination between the two, but the idea of aligning IT with  
2655 business already demonstrates that resource providers and resource consumers must be working  
2656 towards common goals if they are to be productive and efficient. While SOA governance is shown to be  
2657 active in the area of infrastructure, it is a specialized concern for having a dependable platform to support  
2658 service interaction; a range of traditional IT issues is therefore out of scope of this document. A SOA  
2659 governance plan for an enterprise will not of itself resolve shortcomings with the enterprise's IT  
2660 governance.

2661 Governance in the context of SOA is that organization of services: that promotes their visibility; that  
2662 facilitates interaction among service participants; and that directs that the results of service interactions  
2663 are those real world effects as described within the service description and constrained by policies and  
2664 contracts as assembled in the execution context.

2665 SOA governance must specifically account for control across different ownership domains, i.e. all the  
2666 participants may not be under the jurisdiction of a single governance authority. However, for governance  
2667 to be effective, the participants must agree to recognize the authority of the Governance Body and must  
2668 operate within the Governance Framework and through the Governance Processes so defined.

2669 SOA governance must account for interactions across ownership boundaries, which may also imply  
2670 across enterprise governance boundaries. For such situations, governance emphasizes the need for  
2671 agreement that some Governance Framework and Governance Processes have jurisdiction, and the  
2672 governance defined must satisfy the Goals of the participants for cooperation to continue. A standards  
2673 development organization such as OASIS is an example of voluntary agreement to governance over a  
2674 limited domain to satisfy common goals.

2675 The specifics discussed in the figures in the previous sections are equally applicable to governance  
2676 across ownership boundaries as it is within a single boundary. There is a charter agreed to when  
2677 participants become members of the organization, and this charter sets up the structures and processes  
2678 that will be followed. Leadership may be shared by the leadership of the overall organization and the  
2679 leadership of individual groups themselves chartered per the Governance Processes. There are  
2680 Rules/Regulations specific to individual efforts for which participants agree to local goals, and  
2681 Enforcement can be loss of voting rights or under extreme circumstances, expulsion from the group.

2682 Thus, the major difference for SOA governance is an appreciation for the cooperative nature of the  
2683 enterprise and its reliance on furthering common goals if productive participation is to continue.

### 2684 **5.1.3.2 What Must be Governed**

2685 An expected benefit of employing SOA principles is the ability to quickly bring resources to bear to deal  
2686 with unexpected and evolving situations. This requires a great deal of confidence in the underlying  
2687 capabilities that can be accessed and in the services that enable the access. It also requires  
2688 considerable flexibility in the ways these resources can be employed. Thus, SOA governance requires  
2689 establishing confidence and trust while instituting a solid framework that enables flexibility, indicating a  
2690 combination of strict control over a limited set of foundational aspects but minimum constraints beyond  
2691 those bounds.

2692

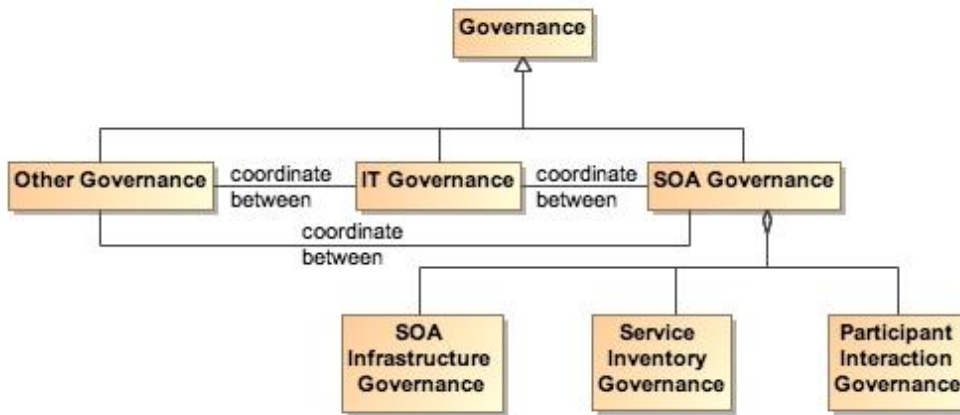


Figure 40 Relationship among types of governance

2693  
2694  
2695  
2696  
2697  
2698  
2699  
2700  
2701

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

### 2702 5.1.3.2.1 Governance of SOA Infrastructure

2703 The SOA infrastructure is likely composed of several families of SOA services that provide access to  
2704 fundamental computing business services. These include, among many others, services such as  
2705 messaging, security, storage, discovery, and mediation. The provisioning of an infrastructure on which  
2706 these services may be accessed and the general realm of those contributing as utility functions of the  
2707 infrastructure are a traditional IT governance concern. In contrast, the focus of SOA governance is how  
2708 the existence and use of the services enables the SOA ecosystem.

2709 By characterizing the environment as containing families of SOA services, the assumption is that there  
2710 may be multiple approaches to providing the business services or variations in the actual business  
2711 services provided. For example, discovery could be based on text search, on metadata search, on  
2712 approximate matches when exact matches are not available, and numerous other variations. The  
2713 underlying implementation of search algorithms are not the purview of SOA governance, but the access  
2714 to the resulting service infrastructure enabling discovery must be stable, reliable, and extremely robust to  
2715 all operating conditions. Such access enables other specialized SOA services to use the infrastructure in  
2716 dependable and predictable ways, and is where governance is important.

### 2717 5.1.3.2.2 Governance of the Service Inventory

2718 Given an infrastructure in which other SOA services can operate, a key governance issue is which SOA  
2719 services to allow in the ecosystem. The major concern SHOULD be a definition of well-behaved services,  
2720 where the required behavior will likely inherit their characteristics from experiences with distributed  
2721 computing but also evolve with SOA experience. A major requirement for ensuring well-behaved services  
2722 is collecting sufficient metrics to know how the service affects the SOA infrastructure and whether it  
2723 complies with established infrastructure policies.

2724 Another common concern of service approval is whether there is a possibility of duplication of function by  
2725 multiple services. Some governance models talk to a tightly controlled environment where a primary  
2726 concern is to avoid any service duplication. Other governance models talk to a market of services where  
2727 the consumers have wide choices. For the latter, it is anticipated that the better services will emerge from  
2728 market consensus and the availability of alternatives will drive innovation.

2729 Some combination of control and openness will emerge, possibly with a different appropriate balance for  
2730 different categories of use. For SOA governance, the issue is less which services are approved but rather  
2731 ensuring that sufficient description is available to support informed decisions for appropriate use. Thus,  
2732 SOA governance SHOULD concentrate on identifying the required attributes to adequately describe a  
2733 service, the required target values of the attributes, and the standards for defining the meaning of the  
2734 attributes and their target values. Governance may also specify the processes by which the attribute  
2735 values are measured and the corresponding certification that some realized attribute set may imply.

2736 For example, unlimited access for using a service may require a degree of life cycle maturity that has  
2737 demonstrated sufficient testing over a certain size community. Alternately, the policy may specify that a  
2738 service in an earlier phase of its life cycle may be made available to a smaller, more technically  
2739 sophisticated group in order to collect the metrics that would eventually allow the service to advance its  
2740 life cycle status.

2741 This aspect of governance is tightly connected to description because, given a well-behaved set of  
2742 services, it is the responsibility of the consumer (or policies promulgated by the consumer's organization)  
2743 to decide whether a service is sufficient for that consumer's intended use. The goal is to avoid global  
2744 governance specifying criteria that are too restrictive or too lax for the local needs of which global  
2745 governance has little insight.

2746 Such an approach to specifying governance allows independent domains to describe services in local  
2747 terms while still having the services available for informed use across domains. In addition, changes to  
2748 the attribute sets within a domain can be similarly described, thus supporting the use of newly described  
2749 resources with the existing ones without having to update the description of all the legacy content.

#### 2750 **5.1.3.2.3 Governance of Participant Interaction**

2751 Finally, given a reliable services infrastructure and a predictable set of services, the third aspect of  
2752 governance is prescribing what is required during a service interaction.

2753 Governance would specify adherence to service interface and service reachability parameters and would  
2754 require that the result of an interaction MUST correspond to the real world effects as contained in the  
2755 service description. Governance would ensure preconditions for service use are satisfied, in particular  
2756 those related to security aspects such as user authentication, authorization, and non-repudiation. If  
2757 conflicts arise, governance would specify resolution processes to ensure appropriate agreements,  
2758 policies, and conditions are met.

2759 It would also rely on sufficient monitoring by the SOA infrastructure to ensure services remain well-  
2760 behaved during interactions, e.g. do not use excessive resources or exhibit other prohibited behavior.  
2761 Governance would also require that policy agreements as documented in the execution context for the  
2762 interaction are observed and that the results and any after effects are consistent with the agreed policies.  
2763 Governance will focus on more contractual and legal aspects rather than the precursor descriptive  
2764 aspects. SOA governance may prescribe the processes by which SOA-specific policies are allowed to  
2765 change, but there are probably more business-specific policies that will be governed by processes  
2766 outside SOA governance.

#### 2767 **5.1.3.3 Overarching Governance Concerns**

2768 There are numerous governance related concerns whose effects span the three areas just discussed.  
2769 One is the area of standards, how these are mandated, and how the mandates may change. The Web  
2770 Services standards stack is an example of relevant standards where a significant number are still under  
2771 development. In addition, while there are notional scenarios that guide what standards are being  
2772 developed, the fact that many of these standards do not yet exist precludes operational testing of their  
2773 adequacy or effectiveness as a necessary and sufficient set.

2774 That said, standards are critical to creating a SOA ecosystem where SOA services can be introduced,  
2775 used singularly, and combined with other services to deliver complex business functionality. As with  
2776 other aspects of SOA governance, the Governance Body should identify the minimum set felt to be  
2777 needed and rigorously enforce that that set be used where appropriate. The Governance Body must take  
2778 care to expand and evolve the mandated standards in a predictable manner and with sufficient technical  
2779 guidance that new services are able to coexist as much as possible with the old, and changes to  
2780 standards do not cause major disruptions.

2781 Another area that may see increasing activity as SOA expands is additional regulation by governments  
2782 and associated legal institutions. New laws are may deal with transactions which are service based,  
2783 possibly including taxes on the transactions. Disclosures laws may mandate certain elements of  
2784 description so both the consumer and provider act in a predictable environment and are protected from  
2785 ambiguity in intent or action. Such laws are spawn rules and regulations that will influence the metrics  
2786 collected for evaluation of compliance.

#### 2787 **5.1.3.4 Considerations for SOA Governance**

2788 The Reference Architecture definition of a loosely coupled system is one in which the constraints on the  
2789 interactions between components is minimal: sufficient to permit interoperation without additional  
2790 constraints that may be an artifact of implementation technology. While governance experience for  
2791 standalone systems provides useful guides, we must be careful not to apply constraints that would  
2792 preclude the flexibility, agility, and adaptability we expect to realize from a SOA ecosystem.

2793 One of the strengths of SOA is it can make effective use of diversity rather than requiring monolithic  
2794 solutions. Heterogeneous organizations can interact without requiring each conforms to uniform tools,  
2795 representation, and processes. However, with this diversity comes the need to adequately define those  
2796 elements necessary for consistent interaction among systems and participants, such as which  
2797 communication protocol, what level of security, which vocabulary for payload content of messages. The  
2798 solution is not always to lock down these choices but to standardize alternatives and standardize the  
2799 representations through which an unambiguous identification of the alternative chosen can be conveyed.  
2800 For example, the URI standard specifies the URI string, including what protocol is being used, what is the  
2801 target of the message, and how may parameters be attached. It does not limit the available protocols, the  
2802 semantics of the target address, or the parameters that can be transferred. Thus, as with our definition of  
2803 loose coupling, it provides absolute constraints but minimizes which constraints it imposes.

2804 There is not a one-size-fits-all governance but a need to understand the types of things governance is  
2805 called upon to do in the context of the goals of SOA. Some communities may initially desire and require  
2806 very stringent governance policies and procedures while other see need for very little. Over time, best  
2807 practices will evolve, resulting in some consensus on a sensible minimum and, except in extreme cases  
2808 where it is demonstrated to be necessary, a loosening of strict governance toward the best practice  
2809 mean.

2810 A question of how much governance may center on how much time governance activities require versus  
2811 how quickly is the system being governed expected to respond to changing conditions. For large single  
2812 systems that take years to develop, the governance process could move slowly without having a serious  
2813 negative impact. For example, if something takes two years to develop and the steps involved in  
2814 governance take two months to navigate, then the governance can go along in parallel and may not have  
2815 a significant impact on system response to changes. Situations where it takes as long to navigate  
2816 governance requirements as it does to develop a response are examples where governance may need to  
2817 be reevaluated as to whether it facilitates or inhibits the desired results. Thus, the speed at which  
2818 services are expected to appear and evolve needs to be considered when deciding the processes for  
2819 control. The added weight of governance should be appropriate for overall goals of the application  
2820 domain and the service environment.

2821 Governance, as with other aspects of any SOA implementation, should start small and be conceptualized  
2822 in a way that keeps it flexible, scalable, and realistic. A set of useful guidelines would include:

- 2823 • Do not hardwire things that will inevitably change. For example, develop a system that uses the  
2824 representation of policies rather than code the policies into the implementations.
- 2825 • Avoid setting up processes that demo well for three services without considering how they may  
2826 work for 300. Similarly, consider whether the display of status and activity for a small number of  
2827 services will also be effective for an operator in a crisis situation looking at dozens of services,  
2828 each with numerous, sometimes overlapping and sometimes differing activities.
- 2829 • Maintain consistency and realism. A service solution responding to a natural disaster cannot be  
2830 expected to complete a 6-week review cycle but be effective in a matter of hours.

#### 2831 5.1.4 Architectural Implications of SOA Governance

2832 The description of SOA governance indicates numerous architectural requirements on the SOA  
2833 ecosystem:

- 2834 • Governance is expressed through policies and assumes multiple use of focused policy modules  
2835 that can be employed across many common circumstances. This requires the existence of:
  - 2836 ○ descriptions to enable the policy modules to be visible, where the description includes a  
2837 unique identifier for the policy and a sufficient, and preferably a machine process-able,  
2838 representation of the meaning of terms used to describe the policy, its functions, and its  
2839 effects;
  - 2840 ○ one or more discovery mechanisms that enable searching for policies that best meet the  
2841 search criteria specified by the service participant; where the discovery mechanism will  
2842 have access to the individual policy descriptions, possibly through some repository  
2843 mechanism;
  - 2844 ○ accessible storage of policies and policy descriptions, so service participants can access,  
2845 examine, and use the policies as defined.
- 2846 • Governance requires that the participants understand the intent of governance, the structures  
2847 created to define and implement governance, and the processes to be followed to make  
2848 governance operational. This requires the existence of:
  - 2849 ○ an information collection site, such as a Web page or portal, where governance  
2850 information is stored and from which the information is always available for access;
  - 2851 ○ a mechanism to inform participants of significant governance events, such as changes in  
2852 policies, rules, or regulations;
  - 2853 ○ accessible storage of the specifics of Governance Processes;
  - 2854 ○ SOA services to access automated implementations of the Governance Processes
- 2855 • Governance policies are made operational through rules and regulations. This requires the  
2856 existence of:
  - 2857 ○ descriptions to enable the rules and regulations to be visible, where the description  
2858 includes a unique identifier and a sufficient, and preferably a machine process-able,  
2859 representation of the meaning of terms used to describe the rules and regulations;
  - 2860 ○ one or more discovery mechanisms that enable searching for rules and regulations that  
2861 may apply to situations corresponding to the search criteria specified by the service  
2862 participant; where the discovery mechanism will have access to the individual  
2863 descriptions of rules and regulations, possibly through some repository mechanism;
  - 2864 ○ accessible storage of rules and regulations and their respective descriptions, so service  
2865 participants can understand and prepare for compliance, as defined.
  - 2866 ○ SOA services to access automated implementations of the Governance Processes.
- 2867 • Governance implies management to define and enforce rules and regulations. Management is  
2868 discussed more specifically in section **Error! Reference source not found.**, but in a parallel to  
2869 overnance, management requires the existence of:
  - 2870 ○ an information collection site, such as a Web page or portal, where management  
2871 information is stored and from which the information is always available for access;
  - 2872 ○ a mechanism to inform participants of significant management events, such as changes  
2873 in rules or regulations;
  - 2874 ○ accessible storage of the specifics of processes followed by management.
- 2875 • Governance relies on metrics to define and measure compliance. This requires the existence of:
  - 2876 ○ the infrastructure monitoring and reporting information on SOA resources;
  - 2877 ○ possible interface requirements to make accessible metrics information generated or  
2878 most easily accessed by the service itself.

## 2879 5.2 Security Model

2880 Security is one aspect of confidence – the confidence in the integrity, reliability, and confidentiality of the  
2881 system. In particular, security focuses on those aspects of assurance that involve the accidental or malign  
2882 intent of other people to damage or compromise trust in the system and on the availability of SOA-based  
2883 systems to perform desired capability.

### 2884 Security

2885 Security concerns the set of mechanisms for ensuring and enhancing trust and confidence in the  
2886 SOA ecosystem.

2887 Providing for security for Service Oriented Architecture is somewhat different than for other contexts;  
2888 although many of the same principles apply equally to SOA and to other systems. The fact that SOA  
2889 embraces crossing ownership boundaries makes the issues involved with moving data more visible.

2890 As well as securing the movement of data within and across ownership boundaries, security often  
2891 revolves around resources: the need to guard certain resources against inappropriate access – whether  
2892 reading, writing or otherwise manipulating those resources.

2893 Any comprehensive security solution must take into account the people that are using, maintaining and  
2894 managing the SOA. Furthermore, the relationships between them must also be incorporated: any security  
2895 assertions that may be associated with particular interactions originate in the people that are behind the  
2896 interaction.

2897 We analyze security in terms of the social structures that define the legitimate permissions, obligations  
2898 and roles of people in relation to the system, and mechanisms that must be put into place to realize a  
2899 secure system. The former are typically captured in a series of security policy statements; the latter in  
2900 terms of security *guards* that ensure that policies are enforced.

2901 How and when to apply these derived security policy mechanisms is directly associated with the  
2902 assessment of the *threat model* and a *security response model*. The threat model identifies the kinds of  
2903 threats that directly impact the message and/or application of constraints, and the response model is the  
2904 proposed mitigation to those threats. Properly implemented, the result can be an acceptable level of risk  
2905 to the safety and integrity of the system.

### 2906 5.2.1 Secure Interaction Concepts

2907 We can characterize secure interactions in terms of key security concepts **[ISO/IEC 27002]**:  
2908 confidentiality, integrity, authentication, authorization, non-repudiation, and availability. The concepts for  
2909 secure interactions are well defined in other standards and publications. The security concepts here are  
2910 not defined but rather related to the SOA ecosystem perspective of the SOA-RAF.

#### 2911 5.2.1.1 Confidentiality

2912 Confidentiality concerns the protection of privacy of participants in their interactions. Confidentiality refers  
2913 to the assurance that unauthorized entities are not able to read messages or parts of messages that are  
2914 transmitted.

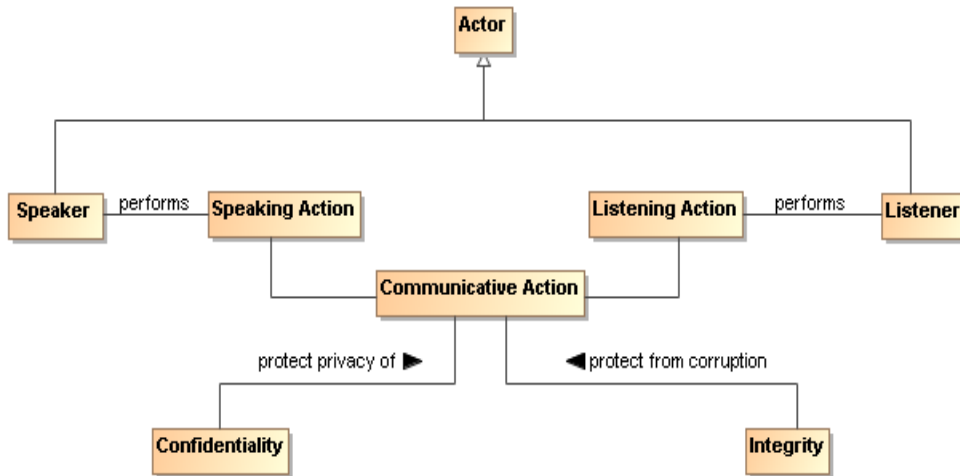
2915 Note that confidentiality has degrees: in a completely confidential exchange, third parties would not even  
2916 be aware that a confidential exchange has occurred. In a partially confidential exchange, the identities of  
2917 the participants may be known but the content of the exchange obscured.

#### 2918 5.2.1.2 Integrity

2919 Integrity concerns the protection of information that is exchanged – either from unauthorized writing or  
2920 inadvertent corruption. Integrity refers to the assurance that information that has been exchanged has not  
2921 been altered.

2922 Integrity is different from confidentiality in that messages that are sent from one participant to another  
2923 may be obscured to a third party, but the third party may still be able to introduce his own content into the  
2924 exchange without the knowledge of the participants.

2925 Figure 41 applies confidentiality and integrity to communicative action.




2926  
2927 *Figure 41 Confidentiality and Integrity*

2928 A communicative action is a joint action involved in the exchange of messages. Section 5.2.4 describes  
2929 common computing techniques for providing confidentiality and integrity during message exchanges.

2930 **5.2.1.3 Authentication**

2931 Authentication concerns the identity of the participants in an exchange. Authentication refers to the  
2932 means by which one participant can be assured of the identity of other participants.

2933 Figure 42 applies authentication to the identity of participants.



2935  
2936 *Figure 42 Authentication*

2937 **5.2.1.4 Authorization**

2938 Authorization concerns the legitimacy of the interaction. Authorization refers to the means by which a  
2939 stakeholder may be assured that the information and actions that are exchanged are either explicitly or  
2940 implicitly approved.




Figure 43 Authorization

2941  
2942

2943 The roles and attributes which provide a participant's credentials are expanded to include reputation.  
2944 Reputation often helps determine willingness to interact, for example, reviews of a service provider will  
2945 influence the decision to interact with the service provider. The roles, reputation, and attributes are  
2946 represented as assertions measured by authorization decision points.

2947 The role of policy for security is to permit stakeholders to express their choices. In Figure 43, a policy is a  
2948 written constraint and the role, reputation, and attribute assertions are evaluated according to the  
2949 constraints in the authorization policy. A combination of security mechanisms and their control via  
2950 explicit policies can form the basis of an authorization solution.

### 2951 5.2.1.5 Non-repudiation

2952 Non-repudiation concerns the accountability of participants. To foster trust in the performance of a system  
2953 used to conduct shared activities it is important that the participants are not able to later deny their  
2954 actions: to repudiate them. Non-repudiation refers to the means by which a participant may not, at a later  
2955 time, successfully deny having participated in the interaction or having performed the actions as reported  
2956 by other participants.

### 2957 5.2.1.6 Availability

2958 Availability concerns the ability of systems to use and offer the services for which they were designed.  
2959 One of the threats against availability is the so-called denial of service attack in which attackers attempt to  
2960 prevent legitimate access to the system.

2961 We differentiate here between general availability – which includes aspects such as systems reliability –  
2962 and availability as a security concept where we need to respond to active threats to the system.



## 2963 **5.2.2 Where SOA Security is Different**

2964 The core security concepts are fundamental to all social interactions. The evolution of sharing  
2965 information using a SOA requires the flexibility to dynamically secure computing interactions in a  
2966 computing ecosystem where the owning social groups, roles, and authority are constantly changing as  
2967 described in section 5.1.3.1.

2968 SOA policy-based security can be more adaptive for a computing ecosystem than previous computing  
2969 technologies allow for, and typically involves a greater degree of distributed mechanisms.

2970 Standards for security, as is the case with all aspects of SOA, play a large role in flexible security on a  
2971 global scale. SOA security may also involve greater auditing and reporting to adhere to regulatory  
2972 compliance established by governance structures.

## 2973 **5.2.3 Security Threats**

2974 There are a number of ways in which an attacker may attempt to compromise the security of a system.  
2975 The two primary sources of attack are third parties attempting to subvert interactions between legitimate  
2976 participants and an entity that is participating but attempting to subvert its partner(s). The latter is  
2977 particularly important in a SOA where there may be multiple ownership boundaries and trust boundaries.

2978 The threat model lists some common threats that relate to the core security concepts listed in Section  
2979 5.2.1. Each technology choice in the realization of a SOA can potentially have many threats to consider.

### 2980 **Message alteration**

2981 If an attacker is able to modify the content (or even the order) of messages that are exchanged  
2982 without the legitimate participants being aware of it then the attacker has successfully  
2983 compromised the security of the system. In effect, the participants may unwittingly serve the  
2984 needs of the attacker rather than their own.

2985 An attacker may not need to completely replace a message with his own to achieve his objective:  
2986 replacing the identity of the beneficiary of a transaction may be enough.

### 2987 **Message interception**

2988 If an attacker is able to intercept and understand messages exchanged between participants,  
2989 then the attacker may be able to gain advantage. This is probably the most commonly understood  
2990 security threat.

### 2991 **Man in the middle**

2992 In a man-in-the-middle attack, the legitimate participants believe that they are interacting with  
2993 each other; but are in fact interacting with the attacker. The attacker attempts to convince each  
2994 participant that he is their correspondent; whereas in fact he is not.

2995 In a successful man-in-the-middle attack, legitimate participants do not have an accurate  
2996 understanding of the state of the other participants. The attacker can use this to subvert the  
2997 intentions of the participants.

### 2998 **Spoofing**

2999 In a spoofing attack, the attacker convinces a participant that he is really someone else –  
3000 someone that the participant would normally trust.

### 3001 **Denial of service attack**

3002 In a denial of service (DoS) attack, the attacker attempts to prevent legitimate users from making  
3003 use of the service. A DoS attack is easy to mount and can cause considerable harm: by  
3004 preventing legitimate interactions, or by slowing them down enough, the attacker may be able to  
3005 simultaneously prevent legitimate access to a service and to attack the service by another  
3006 means.

3007 A variation of the DoS attack is the Distributed Denial of Service attack. In a DDoS attack the  
3008 attacker uses multiple agents to attack the target. In some circumstances this can be  
3009 extremely difficult to counteract effectively.

3010 One of the features of a DoS attack is that it does not require valid interactions to be effective:  
3011 responding to invalid messages also takes resources and that may be sufficient to cripple the  
3012 target.

### 3013 **Replay attack**

3014 In a replay attack, the attacker captures the message traffic during a legitimate interaction and  
3015 then replays part of it to the target. The target is persuaded that a similar transaction to the  
3016 previous one is being repeated and it responds as though it were a legitimate interaction.

3017 A replay attack may not require that the attacker understand any of the individual  
3018 communications; the attacker may have different objectives (for example attempting to predict  
3019 how the target would react to a particular request).

### 3020 **False repudiation**

3021 In false repudiation, a user completes a normal transaction and then later attempts to deny that  
3022 the transaction occurred. For example, a customer may use a service to buy a book using a credit  
3023 card; then, when the book is delivered, refuse to pay the credit card bill claiming that *someone*  
3024 *else* must have ordered the book.

## 3025 **5.2.4 Security Responses**

3026 Security goals are never absolute: it is not possible to guarantee 100% confidentiality, non-repudiation,  
3027 etc. However, a well designed and implemented security response model can ensure acceptable levels of  
3028 security risk. For example, using a well-designed cipher to encrypt messages may make the cost of  
3029 breaking communications so great and so lengthy that the information obtained is valueless.

3030 Performing threat assessments, devising mitigation strategies, and determining acceptable levels of risk  
3031 are the foundation for an effective process to mitigating threats in a cost-effective way.<sup>14</sup> The choice in  
3032 hardware and software to realize a SOA will be the basis for threat assessments and mitigation  
3033 strategies. The stakeholders of a specific SOA implementation should determine acceptable levels of risk  
3034 based on threat assessments and the cost of mitigating those threats.

### 3035 **5.2.4.1 Privacy Enforcement**

3036 The most efficient mechanism to assure confidentiality is the encryption of information. Encryption is  
3037 particularly important when messages must cross trust boundaries; especially over the Internet. Note that  
3038 encryption need not be limited to the content of messages: it is possible to obscure even the existence of  
3039 messages themselves through encryption and 'white noise' generation in the communications channel.

3040 The specifics of encryption are beyond the scope of this architecture. However, we are concerned about  
3041 how the connection between privacy-related policies and their enforcement is made.

3042 A policy enforcement point for enforcing privacy may take the form of an automatic function to encrypt  
3043 messages as they leave a trust boundary; or perhaps simply ensuring that such messages are suitably  
3044 encrypted.

3045 Any policies relating to the level of encryption being used would then apply to these centralized  
3046 messaging functions.

### 3047 **5.2.4.2 Integrity Protection**

3048 To protect against message tampering or inadvertent message alteration, and to allow the receiver of a  
3049 message to authenticate the sender, messages may be accompanied by a digital signature. Digital

---

<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.

3050 signatures provide a means to detect if signed data has been altered. This protection can also extend to  
3051 authentication and non-repudiation of a sender.

3052 A common way a digital signature is generated is with the use of a private key that is associated with a  
3053 public key and a digital certificate. The private key of some entity in the system is used to create a digital  
3054 signature for some set of data. Other entities in the system can check the integrity of the signed data set  
3055 via signature verification algorithms. Any changes to the data that was signed will cause signature  
3056 verification to fail, which indicates that integrity of the data set has been compromised.

3057 A party verifying a digital signature must have access to the public key that corresponds to the private key  
3058 used to generate the signature. A digital certificate contains the public key of the owner, and is itself  
3059 protected by a digital signature created using the private key of the issuing Certificate Authority (CA).

### 3060 **5.2.4.3 Message Replay Protection**

3061 To protect against replay attacks, messages may contain information that can be used to detect replayed  
3062 messages. The simplest requirement to prevent replay attacks is that each message that is ever sent is  
3063 unique. For example, a message may contain a message ID, a timestamp, and the intended destination.

3064 By storing message IDs, and comparing each new message with the store, it becomes possible to verify  
3065 whether a given message has been received before (and therefore should be discarded).

3066 The timestamp may be included in the message to help check for message freshness. Messages that  
3067 arrive after their message ID could have been cleared (after receiving the same message some time  
3068 previously) may also have been replayed. A common means for representing timestamps is a useful part  
3069 of an interoperable replay detection mechanism.

3070 The destination information is used to determine if the message was misdirected or replayed. If the  
3071 replayed message is sent to a different endpoint than the destination of the original message, the replay  
3072 could go undetected if the message does not contain information about the intended destination.

3073 In the case of messages that are replies to prior messages, it is also possible to include seed information  
3074 in the prior messages that is randomly and uniquely generated for each message that is sent out. A  
3075 replay attack can then be detected if the reply does not embed the random number that corresponds to  
3076 the original message.

### 3077 **5.2.4.4 Auditing and Logging**

3078 False repudiation involves a participant denying that it authorized a previous interaction. An effective  
3079 strategy for responding to such a denial is to maintain careful and complete logs of interactions which can  
3080 be used for auditing purposes. The more detailed and comprehensive an audit trail is, the less likely it is  
3081 that a false repudiation would be successful.

3082 The countermeasures assume that the non-repudiation tactic (e.g. digital signatures) is not undermined  
3083 itself. For example, if private key is stolen and used by an adversary, even extensive logging cannot  
3084 assist in rejecting a false repudiation.

3085 Unlike many of the security responses discussed here, it is likely that the scope for automation in  
3086 rejecting a repudiation attempt is limited to careful logging.

### 3087 **5.2.4.5 Graduated engagement**

3088 The key to managing and responding to DoS attacks is to be careful in the use of resources when  
3089 responding to interaction. Put simply, a system has a choice to respond to a communication or to ignore  
3090 it. In order to avoid vulnerability to DoS attacks a service provider should be careful not to commit  
3091 resources beyond those implied by the current state of interactions; this permits a graduation in  
3092 commitment by the service provider that mirrors any commitment on the part of service consumers and  
3093 attackers alike.

## 3094 **5.2.5 Architectural Implications of SOA Security**

3095 Providing SOA security in an ecosystem of governed services has the following implications on the policy  
3096 support and the distributed nature of mechanisms used to assure SOA security:

- 3097 • Security expressed through policies have the same architectural implications as described in  
3098 Section 4.4.3 for policies and contracts architectural implications.
- 3099 • Security policies require mechanisms to support security description administration, storage, and  
3100 distribution.
- 3101 • Service descriptions supporting security policies should:
  - 3102 ○ have a meta-structure sufficiently rich to support security policies;
  - 3103 ○ be able to reference one or more security policy artifacts;
  - 3104 ○ have a framework for resolving conflicts between security policies.
- 3105 • The mechanisms that make-up the execution context in secure SOA-based systems should:
  - 3106 ○ provide protection of the confidentiality and integrity of message exchanges;
  - 3107 ○ be distributed so as to provide centralized or decentralized policy-based identification,  
3108 authentication, and authorization;
  - 3109 ○ ensure service availability to consumers;
  - 3110 ○ be able to scale to support security for a growing ecosystem of services;
  - 3111 ○ be able to support security between different communication technologies;
- 3112 • Common security services include:
  - 3113 ○ services that abstract encryption techniques;
  - 3114 ○ services for auditing and logging interactions and security violations;
  - 3115 ○ services for identification;
  - 3116 ○ services for authentication;
  - 3117 ○ services for authorization;
  - 3118 ○ services for intrusion detection and prevention;
  - 3119 ○ services for availability including support for quality of service specifications and metrics.

## 3120 **5.3 Management Model**

### 3121 **Management**

3122 Management is a process of controlling resources in accordance with the policies and principles defined  
3123 by Governance.

3124 There are three separate but linked domains of interest within the management of SOA:

- 3125 1. the management and support of the resources that are involved in any complex structures – of which  
3126 SOA-based solutions are excellent examples;
- 3127 2. the promulgation and enforcement of the policies and service contracts agreed to by the stakeholders  
3128 in SOA ecosystem;
- 3129 3. the management of the relationships of the participants in SOA-based solutions – both to each other  
3130 and to the services that they use and offer.

3131 There are many artifacts related to management. Historically, systems management capabilities have  
3132 been organized by the “FCAPS” functions (based on ITU-T Rec. M.3400 (02/2000), “TMN Management  
3133 Functions”):

- 3134 • fault management,
- 3135 • configuration management,
- 3136 • account management,
- 3137 • performance and security management.

3138 The primary task of the functional groups is to concentrate on maintaining systems in a trusted, active,  
3139 and accessible state.

3140 In the context of the SOA ecosystem, we see many possible resources that may require management  
3141 such as services, service descriptions, service contracts, policies, roles, relationships, security, people  
3142 and systems that implement services and infrastructure elements. In addition, given the ecosystem  
3143 nature, it is also potentially necessary to manage the business relationships between participants.

3144 Successful operation of a SOA ecosystem requires trust between the stakeholders and the ecosystem  
3145 elements. In contrast, regular systems in technology are not necessarily operated or used in an  
3146 environment requiring trust before the stakeholders make use of the system. Indeed, many of these  
3147 systems exist in hierarchical management structures, within which use may be mandated by legal  
3148 requirement, executive decision, or good business practice in furthering the business' strategy. Pre-  
3149 condition of trust in the SOA ecosystem roots in both principles of service orientation and distributed  
3150 authoritative ownership of independent services. Even for hierarchical management structures applied to  
3151 a SOA ecosystem, the service use should have contractual basis rather than being mandated.

3152 The trust may be established through agreements/contracts, policies, or implicitly through observation of  
3153 repeated interactions with others. Explicit trust is usually accompanied by formalized documents suitable  
3154 for the management activities. Implicit trust adds fragility to the management of a SOA ecosystem  
3155 because failure to maintain consistent and predictable interactions will undermine the trust between  
3156 participants and within the ecosystem as a whole.

3157 Management in a SOA ecosystem is thus concerned with management taking actions that will establish  
3158 the condition of trust that must be present before engaging in service interactions. These concerns should  
3159 largely be handled within the governance of the ecosystem. The policies, agreements, and practices  
3160 defined through the governance provide the boundaries within which management operates and for which  
3161 management must provide enforcement and feedback. However, governance alone cannot anticipate all  
3162 circumstances and must offer sufficient guidance in areas where anticipation is unclear or for which  
3163 agreement between all stakeholders cannot be reached. Management in these cases must be flexible  
3164 and adaptable to handle unanticipated conditions without unnecessarily breaking trust relationships.

3165 Service management is the process – manual, automated, or a combination – of proactively monitoring  
3166 and controlling the behavior of a service or a set of services. Service management operates under  
3167 constraints attributed to the business and social context. Particularly, special policies may be used for  
3168 governing cross-boundary relationships. Managing solutions that may be used across ownership  
3169 boundaries based on such policies raises issues that are not typically present when managing a service  
3170 within a single ownership domain. For example, care is required in managing a service when the owner of  
3171 the service, the provider of the service, the host of the service and mediators to the service may all  
3172 belong to different stakeholders.

3173 Cross-boundary service management takes place in, at least, the following situations:

- 3174 • using combinations of services that belong to different ownership realms
- 3175 • using of services that mediate between ownership realms
- 3176 • sharing monitoring and reporting means and results.

3177 These situations are particularly important in ecosystems that are highly decentralized, in which the  
3178 participants interact as peers as well as in the “master-servant” mode.

3179 The management model shown in Figure 46 conveys how the SOA framework applies to managing  
3180 services. Services management operates via service metadata, such as service lifecycles and attributes  
3181 associated with service use, that are typically collected in or accessed through the service description.

3182 *[this Figure to be re-drawn in common style]*

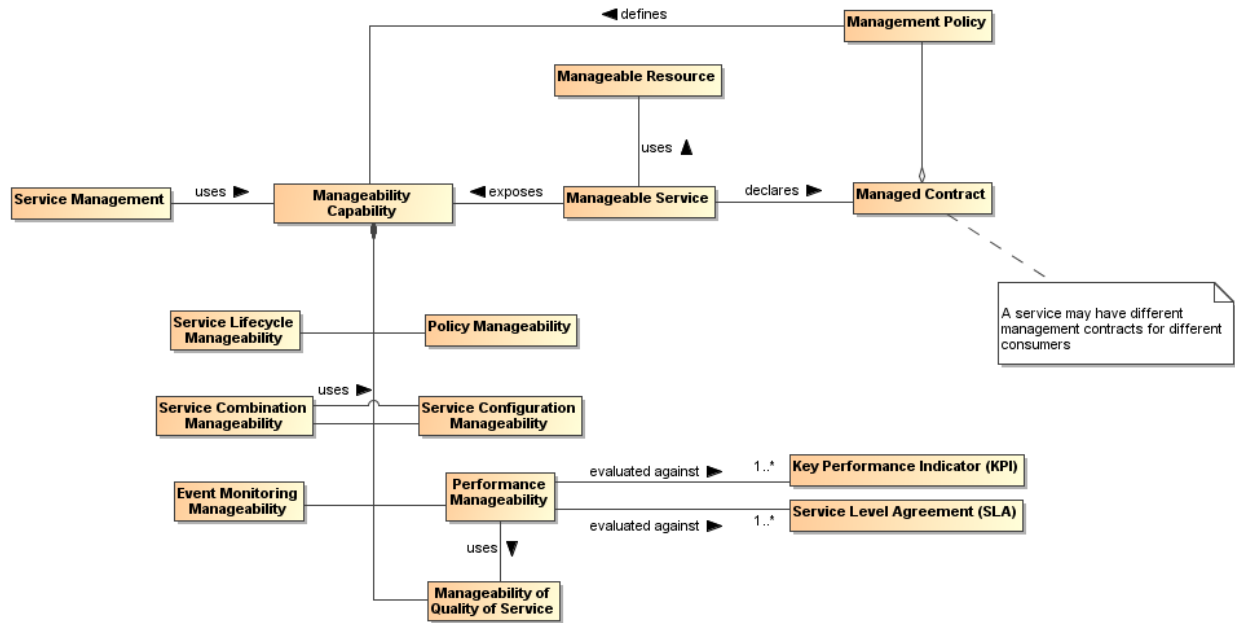


Figure 46 Manageability capabilities in SOA ecosystem

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

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

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

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

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

**Combination Manageability** of a service addresses management of service characteristics that allow for creating and changing of combinations in which the service participates or that the service combines by itself. Known models of such combinations are aggregations and compositions. Examples of patterns of combinations are choreography and orchestration. Combination Manageability drives implementation of the Service Composability Principle of service orientation.

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

3217 Service combinations contribute the most in delivering business values to the stakeholders and managing  
3218 service combinations is the one of the top-level tasks and features of the SOA ecosystem.

3219 **Configuration Manageability** of a service allows managing the identity of and the interactions  
3220 among internal elements of the service. Also, Configuration Manageability correlates with the  
3221 management of service versions and configuration of the deployment of new services into the  
3222 ecosystem. Configuration Management differs from the Combination Manageability in the scope  
3223 and scale of manageability, and addresses lower level concerns than the architectural  
3224 combination of services.

3225 **Event Monitoring Manageability** allows managing the categories of events of interest related to  
3226 services and reporting recognized events to the interested stakeholders. Such events may be the  
3227 ones that trigger service invocations as well as execution of particular functionality provided by  
3228 the service.

3229 This is one of the key lower-level manageability aspects that the service provider and associated  
3230 stakeholders are primarily interested. Monitored events may be internal or external to the SOA  
3231 ecosystem. For example, a disaster in the oil producing industry, which is outside of the SOA ecosystem  
3232 of the Insurer, can trigger the service's functionality that is responsible for immediate or constant  
3233 monitoring of the oil prices in the oil trading exchanges and, respectively, modify the premium paid by the  
3234 insured oil companies.

3235 **Performance Manageability** of a service allows controlling the service results, shared and  
3236 sharable real world effects against the business goals and objectives of the service. This  
3237 manageability assumes monitoring of the business performance as well as the management of  
3238 this monitoring itself. Performance Manageability includes business and technical performance  
3239 manageability means through performance criteria set, such as business key performance  
3240 indicators (KPI) and service-level agreements (SLA).

3241 The performance business- and technical-level characteristics of the service should be known from the  
3242 service contract. The service provider and consumer must be able to monitor and measure these  
3243 characteristics or be informed about the results measured by a third party.

3244 Performance Manageability is the instrument for providing compliance of the service with its service  
3245 contracts. Performance Manageability utilizes Manageability of Quality of Service.

3246 **Manageability of Quality of Service** deals with management of service non-functional  
3247 characteristics that may be of significant value to the service consumers and other stakeholders  
3248 in the SOA ecosystem. Classic examples of this include bandwidth offerings associated with a  
3249 service.

3250 Manageability of quality of service assumes that the properties associated with service qualities are  
3251 monitored during the service execution. Results of monitoring may be challenged against SLA and even  
3252 KPI, which results in the continuous validation of how the service contract is preserved by the service  
3253 provider.

3254 **Policy Manageability** allows additions, changes and replacements of the policies associated  
3255 with a resource in the SOA ecosystem. The ability to manage those policies (such as  
3256 promulgating policies, retiring policies and ensuring that policy decision points and enforcement  
3257 points are current) enables the ecosystem to apply policies and *evaluate* the results.

3258 Capability to manage, i.e. use particular manageability, requires policies from governance to be translated  
3259 into the details of rules and regulations and then corresponding measurement and feedback on the  
3260 specifics.

3261 In the following sub-sections, we describe how the elements of the SOA ecosystem may be managed  
3262 with integrity.

### 3263 **5.3.1. Management Means and Relationships**

3264 A minimal set of management for the SOA ecosystem is shown on Figure 47 and elaborated in the  
3265 following sections.

#### 3266 **5.3.1.1. Management Policy**

3267 The management of resources within the SOA may be governed by management policies. In a deployed  
3268 SOA-based solution, it may well be that different aspects of the management of a given service are

3269 managed by different management services. For example, the life-cycle management of services often  
 3270 involves managing service versions. Managing quality of service is often very specific to the service itself;  
 3271 for example, quality of service attributes for a video streaming service are quite different to those for a  
 3272 banking system.

3273 There are additional concepts of management that also apply to IT management:

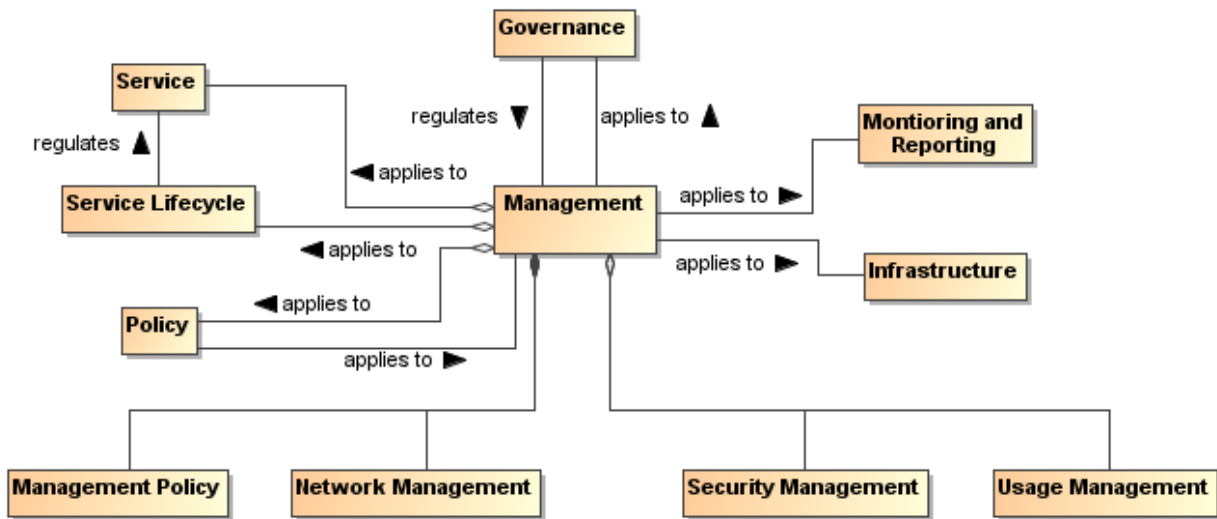
3274 **5.3.1.2. Network Management**

3275 Network management deals with the maintenance and administration of large scale physical networks  
 3276 such as computer networks and telecommunication networks. Specifics of the networks may affect  
 3277 service interactions from performance and operational perspectives.

3278 Network and related system management executes a set of functions required for controlling, planning,  
 3279 deploying, coordinating, and monitoring the distributed services in the SOA ecosystem. However, while  
 3280 recognizing their importance, the specifics of systems management or network management are out of  
 3281 scope for this Reference Architecture Foundation.

3282 *[this Figure to be re-drawn in common style]*

3283



3284  
 3285 *Figure 47 Management Means and Relationships in SOA ecosystem*  
 3286

3287 **5.3.1.3. Security Management**

3288 Management of the security related to resources includes identification of roles, permissions, access  
 3289 rights, and policy attributes defining security boundaries and events that may trigger a security response.

3290 Security management within a SOA ecosystem is essential to maintaining the trust relationships between  
 3291 participants residing in different ownership domains. Security management must consider not just the  
 3292 internal properties related to interactions between participants but ecosystem properties that preserve the  
 3293 integrity of the ecosystem from external threats.

3294 **5.3.1.4. Usage Management**

3295 Usage Management applies to management of the use of resources. Usage management includes  
 3296 access properties, demand properties, and financial properties. Access properties include how the  
 3297 resource is accessed, who is using the resource, and the state of the resource after use. Demand  
 3298 properties are concerned with controlling or shaping demand for resources to optimize the overall  
 3299 operation of the ecosystem. Financial properties are those associated with assigning costs to the use of  
 3300 resources and distributing those cost assignments to the participants in an equitable manner.

3301 **5.3.2. Management and Governance**

3302 The primary role of governance in the context of a SOA ecosystem is to foster an atmosphere of  
 3303 predictability, trust, and efficiency, and it accomplishes this by allowing the stakeholders to negotiate and



3304 set the key policies that govern the running of the SOA-based solution. Recall that in an ecosystems  
3305 perspective, the goal of governance is less to have complete fine-grained control but more to enable the  
3306 individual participants to work together.

3307 Policies for a SOA ecosystem will tend to focus on the rules of engagement between participants; for  
3308 example, what kind of interactions are permissible, how will disputes be resolved, and so on. While  
3309 governance may primarily focus on setting policies, management will focus on the realization and  
3310 enforcement of policies. Effective management in the SOA ecosystem requires an ability for governance  
3311 to understand the consequences of its policies, guidelines, and principles, and to adjust those as needed  
3312 when inconsistencies or ambiguity become evident from the operation of the management functions. This  
3313 understanding and adjustment must be facilitated by the results of management and so the mechanisms  
3314 for providing feedback from management into governance must exist.

3315 Governance operates via specialized activities and, thus, should be managed itself. Management to  
3316 operationalize governance utilizes management policies that are included in the Governance Framework  
3317 and Processes, and driven by the enterprise business model, business objectives and strategies. Where  
3318 corporate management policies exist, these are usually guided and directed by the corporate executives.  
3319 In peer relationships, the governing policies are set by either an external entity and accepted by the peers  
3320 or by the peers themselves. This creates the appropriate authoritative level for the policies used for the  
3321 management of the Governance Framework and Processes. Management to operationalize governance  
3322 controls the life-cycle of the governing policies, including procedures and processes, for modifying the  
3323 Governance Framework and Processes.

### 3324 **5.3.3. Management and Contracts**

#### 3325 **5.3.3.1 Management for Contracts and Policies**

3326 As we noted above, management can often be viewed as the application of contracts and individual  
3327 policies to ensure the smooth running of the SOA ecosystem. Policies play an important role as the  
3328 guiding constraints for management, as well as artifacts that need to be managed themselves. Service  
3329 contracts also serve as both guiding constraints and artifacts that need to be managed. Policies and  
3330 service contracts specify the service characteristics that have to be monitored, analysed and managed.

#### 3331 **5.3.3.2 Contracts**

3332 As described in sections “Participation in a SOA Ecosystem view” and “Realization of a SOA Ecosystem  
3333 view”, there are several types of contractual information in the SOA ecosystem. From the management  
3334 perspective, three basic types of the contractual information relate to:

- 3335 • . relationship between service provider and consumer;
- 3336 • . communication with the service;
- 3337 • . control of the quality of the service execution.

3338 When a consumer prepares to interact with a service, the consumer and the service provider must come  
3339 to agreement on service features and characteristics that will be provided by the service and available to  
3340 the consumer; this agreement is known as a service contract.

#### 3341 **Service Contract**

3342 An implicit or an explicit and documented agreement between the service consumer and service  
3343 provider about the use of the service based on

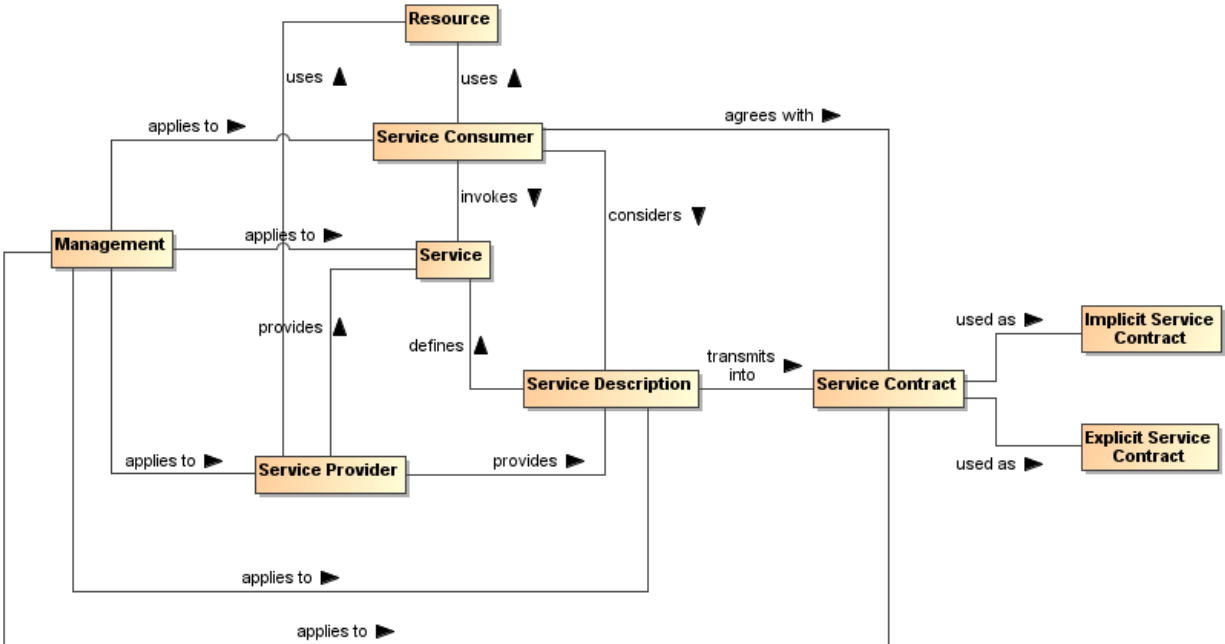
- 3344 • the commitment by a service provider to provide service functionality and results  
3345 consistent with identified real world effects and
- 3346 • the commitment by a service consumer to interact with the service per specific means  
3347 and per specified policies,

3348 where both consumer and provider actions are in the manner described in the service description.

3349 The service description provides the basis for the service contract and, in some situations, may be used  
3350 as an implicit default service contract. In addition, the service description may set mandatory aspects of a  
3351 service contract, e.g. for security services, or may specify acceptable alternatives. As an example of  
3352 alternatives, the service description may identify which versions of a vocabulary will be recognized, and  
3353 the specifics of the contract are satisfied when the consumer uses one of the alternatives. Another  
3354 alternative could have a consumer identifying a policy they require be satisfied, e.g. a standard privacy  
3355 policy on handling personal information, and a provider that is prepared to accept a policy request would

3356 indicate acceptance as part of the service contract by continuing with the interaction. In each of these  
 3357 cases, the actions of the participants are consistent with an implicit service contract without the existence  
 3358 of a formal agreement between the participants.

3359 In the case of business services, it is anticipated that the service contract may take an explicit form and  
 3360 the agreement between business consumer and business service provider is formalized. Formalization  
 3361 requires up-front interactions between service consumer and service provider. In many business  
 3362 interactions, especially between business organisations within or across corporate boundaries, a  
 3363 consumer needs a contractual assurance from the provider or wants to explicitly indicate choices among  
 3364 alternatives, e.g., only use a subset of the business functionality offered by the service and pay a  
 3365 prorated cost.  
 3366



3367  
 3368 *Figure 48 Management of the service interaction*

3369 Consequently, an implicit service contract is an agreement (1) on the consumer side with the terms,  
 3370 conditions, features and interaction means specified in the service description "as is" or (2) a selection  
 3371 from alternatives that are made available through mechanisms included in the service description, and  
 3372 neither of these require any a priori interactions between the service consumer and the service provider.  
 3373 An explicit service contract always requires a form of interaction between the service consumer and the  
 3374 service provider prior to the service invocation. This interaction may regard the choice or selection of the  
 3375 subset of the elements of the service description or other alternatives introduced through the formal  
 3376 agreement process that would be applicable to the interaction with the service and affect related joint  
 3377 action.

3378 Any form of explicit contract couples the service consumer and provider. While explicit contract may be  
 3379 necessary or desirable in some cases, such as in supply chain management, commerce often uses a mix  
 3380 of implicit and explicit contracts, and a service provider may offer (via service description) a conditional  
 3381 shift from implicit to explicit contract. For example, Twitter offers an implicit contract on the use of its APIs  
 3382 to any application with the limit on the amount of service invocations; if the application needs to use more  
 3383 invocations, one has to enter into the explicit fee-based contract with the provider. A case where an  
 3384 implicit contract transforms into explicit contract may be illustrated when one buys a new computer and it  
 3385 does not work. The buyer returns the computer for repairing under manufacture warranty as stated by an  
 3386 implicit purchase contract. However, if the repair does not fix the problem and the seller offers a  
 3387 replacement by upgraded model, the buyer may agree to an explicit contract that limits the rights of the  
 3388 buyer to make the explicit agreement public.

3389 Control of the quality of the service execution, often represented as a service level agreement (SLA), is  
 3390 performed by service monitoring systems and includes both technical and operational business controls.

3391 SLA is a part of the service contract and, because of individual nature of this type of contracts, may vary  
3392 from one service contract to another, even for the same consumer. Typically, a particular SLA in the  
3393 service contract is a concrete instance of the SLA declared in the service description.

3394 Management of the service contracts is based on management policies that may be mentioned in the  
3395 service description and in the service contracts. Management of the service contracts is mandatory for  
3396 consumer relationship management. In the case of explicit service contracts, the contracts have to be  
3397 created, stored, maintained, reviewed/controlled and archived/destroyed as needed. All the activities are  
3398 management concerns. Explicit service contracts may be stored in specialised repositories that provide  
3399 appropriate level of security.

3400

3401 Management of the service interfaces is based on several management policies that regulate

- 3402 • availability of interfaces specified in the service contracts,
- 3403 • accessibility of interfaces,
- 3404 • procedures for interface changes,
- 3405 • interface versions and well as the versions of all parts of the interfaces, and
- 3406 • traceability of the interfaces and their versions back to the service description document.

3407

3408 Management of the SLA is integral to the management of service monitoring and operational service  
3409 behavior at run-time. A SLA usually enumerates service characteristics and expected performances of  
3410 the service. Since SLA carries connotation of “promise”, monitoring is needed to know if the promise is  
3411 kept. Existence of an SLA itself does not guarantee the consumer will be provided with the service level  
3412 specified in the service contract.

3413 The use of SLA in SOA ecosystem can be wider than just an agreement on technical performances.  
3414 An SLA may contain remedies for situations where the promised service cannot be maintained, or the  
3415 real world effect can't be achieved due to developments subsequent to the agreement. A service  
3416 consumer that acts accordingly to realize the real world effect may be compensated for the breach of the  
3417 SLA if the effect is not realized.

3418 Management of the SLA includes, among others, policies for the SLA changes, updates, and  
3419 replacement. This aspect concerns service Execution Context because the business logic associated with  
3420 a defined interface may differ in different Execution Contexts and affect the overall performance of the  
3421 service.

### 3422 5.3.3.3 Policies

3423 "Although provision of management capabilities enables a service to become manageable, the extent and  
3424 degree of permissible management are defined in management policies that are associated with the  
3425 services. Management policies are used to define the obligations for, and permissions to, managing the  
3426 service" [WSA]. Management policies, in essence, are the realisation of governing rules and regulations.  
3427 As such, some management policies may target services while other policies may target the management  
3428 of the services.

3429 In practice, a policy without any means of enforcing it is vacuous. In the case of management policy, we  
3430 rely on a management infrastructure to realize and enforce management policy.

### 3431 5.3.3.4 Service Description and Management

3432 The service description identifies several management objects such as a set of service interfaces and  
3433 related set of SLAs: service behavioral characteristics and performances specified in the SLA depend on  
3434 the interface type and its Execution Context. In the service description, a service consumer can find  
3435 references to management policies, SLA metrics, and the means of accessing measured values that  
3436 together increase assurance in the service quality. At the same time, service description is an artifact that  
3437 needs to be managed.

3438 In the SOA ecosystem, the service description is the assembled information that describes the service but  
3439 it may be reported or displayed in different presentations. While each separate version of the service has  
3440 one and only one service description, different categories of service consumers may focus their interests  
3441 on different aspects of the service description. Thus, the same service description may be displayed not  
3442 only in different languages but also with different cultural and professional accents in the content.

3443 New service description may be issued to reflect changes and update in the service. If the change in the  
3444 service does not affect its service description, the new service version may have the same service  
3445 description as the previous version except for the updated version identifier. For example, a service  
3446 description may stay the same if bugs were fixed in the service. However, if a change in the service  
3447 influences any aspects of the service quality that can affect the real world effect resulting from  
3448 interactions with the service, the service description must reflect this change even if there are no changes  
3449 to the service interface.

3450 Management of the service description and related explicit service contracts is essential for delivery of the  
3451 service to the consumer satisfaction. This management can also prevent business problems rooted in  
3452 poor communication between the service consumers and the service providers.

3453 Thus, management of the service description contains, among others, management of the service  
3454 description presentations, the life-cycles of the service descriptions, service description distribution  
3455 practices and storage of the service descriptions and related service contracts. Collections of service  
3456 descriptions in the enterprise may manifest a need for specialised registries and/or repositories.  
3457 Depending on the enterprise policies, an allocation of purposes and duties of registries and repositories  
3458 may vary but this topic is beyond the current scope.

#### 3459 **5.3.4. Management for Monitoring and Reporting**

3460 The successful application of management relies on the monitoring and reporting aspects of  
3461 management to enable the control aspect. Monitoring in the context of management consists of  
3462 measuring values of managed aspects and evaluating that measurement in relationship to some  
3463 expectation. Monitoring in a SOA ecosystem is enabled through the use of mechanisms by resources for  
3464 exposing managed aspects. In the SOA framework, this mechanism may be a service for obtaining the  
3465 measurement. Alternatively, the measurement may be monitored by means of event generation  
3466 containing updated values of the managed aspect.

3467 Approaches to monitoring may use a polling strategy in which the measurements are requested from  
3468 resources in periodic intervals, in a pull strategy in which the measurements are requested from  
3469 resources at random times, or in a push strategy in which the measurements are supplied by the  
3470 resource without request. The push strategy can be used in a periodic update approach or in an “update  
3471 on change” approach. Management services must be capable of handling these different approaches to  
3472 monitoring.

3473 Reporting is the complement to monitoring. Where monitoring is responsible for obtaining measurements,  
3474 reporting is responsible for distributing those measurements to interested stakeholders. The separation  
3475 between monitoring and reporting is made to include the possibility that data obtained through monitoring  
3476 might not be used until an event impacting the ecosystem occurs or the measurement requires further  
3477 processing to be useful. In the SOA framework, reporting is provided using services for requesting  
3478 measurement reports. These reports may consist of raw measurement data, formatted collections of  
3479 data, or the results of analysis performed on measurement data from collections of different managed  
3480 aspects. Reporting is also used to support logging and auditing capabilities, where the reporting  
3481 mechanisms create log or audit entries.

#### 3482 **5.3.5 Management for Infrastructure**

3483 All of the properties, policies, interactions, resources, and management are only possible if a SOA  
3484 ecosystem infrastructure provides support for managed capabilities. Each managed capability imposes  
3485 different requirements on the capabilities supplied by the infrastructure in SOA ecosystem and requires  
3486 that those capabilities be usable as services or at the very least be interoperable.

3487 Not providing the full list of infrastructural elements of SOA ecosystem, we list an example of such  
3488 elements here:

- 3489 1. Registries and repositories for services, policies, and related descriptions  
3490 and contracts
- 3491 2. Synchronous and asynchronous communication channels for service  
3492 interactions (e.g., network, e-mail, message routing with ability of mediating  
3493 transport protocols, etc.)
- 3494 3. Recovery capabilities
- 3495 4. Security controls

3496 Also, a SOA ecosystem infrastructure, enabling service management, should support

- 3497 1. Management enforcement and control means  
3498 2. Monitoring and SLA validation controls  
3499 3. Testing and Reporting capabilities

3500 Combination of manageability capabilities and infrastructure elements constitutes certain level of SOA  
3501 management maturity. While several maturity models exist, this topic is out of the scope of the document.

## 3502 **5.4 SOA Testing Model**

3503 *Program testing can be used to show the presence of bugs,*  
3504 *but never to show their absence!*  
3505 Edsger Dijkstra

3506 Testing for SOA combines the typical challenges of software testing and certification with the additional  
3507 needs of accommodating the distributed nature of the resources, the greater access of a more  
3508 unbounded consumer population, and the desired flexibility to create new solutions from existing  
3509 components over which the solution developer has little if any control. The purpose of testing is to  
3510 demonstrate a required level of reliability, correctness, and effectiveness that enable prospective  
3511 consumers to have adequate confidence in using a service. Adequacy is defined by the consumer based  
3512 on the consumer's needs and context of use. As the Dijkstra quote points out, absolute correctness and  
3513 completeness cannot be proven by testing; however, for SOA, it is critical for the prospective consumer to  
3514 know what testing has been performed, how it has been performed, and what were the results.

### 3515 **5.4.1 Traditional Software Testing as Basis for SOA Testing**

3516 SOA services are largely software artifacts and can leverage the body of experience that has evolved  
3517 around software testing. IEEE-829 specifies the basic set of software test documents while allowing  
3518 flexibility for tailored use. As such, the document structure can also provide guidance to SOA testing.

3519 IEEE-829 covers test specification and test reporting through use of the following document types:

- 3520 • *Test plan* documenting the scope (what is to be tested, both which entity and what features of the  
3521 entity), the approach (how it is tested), and the needed resources (who does the testing, for how  
3522 long), with details contained in the:
- 3523 • *Test design specification*: features to be tested, test conditions (e.g. test cases, test procedures  
3524 needed) and expected results (criteria for passing test); entrance and exit criteria
  - 3525 • *Test case specification*: test data used for input and expected output
  - 3526 • *Test procedure specification*: steps required to run the test, including any set-up preconditions
- 3527 • *Test item transmittal* to identify the test items being transmitted for testing
- 3528 • *Test log* to record what occurred during test, i.e. which tests run, who ran, what order, what happened
- 3529 • *Test incident report* to capture any event that happened during test which requires further  
3530 investigation
- 3531 • *Test summary* as a management report summarizing test run and results, conclusions

3532 In summary, IEEE-829 captures (1) what was tested, (2) how it was tested, e.g. the test procedure used,  
3533 and (3) the results of the test.

#### 3534 **5.4.1.1 Types of Testing**

3535 There are numerous aspects of testing that, in total, work to establish that an entity is (1) built as required  
3536 per policies and related specifications prescribed by the entity's owner, and (2) delivers the functionality  
3537 required by its intended users. This is often referred to as verification and validation.

3538 Policies, as described in Section 4.4, that are related to testing may prescribe but are not limited to the  
3539 business processes to be followed, the standards with which an implementation must comply, and the  
3540 qualifications of and restrictions on the users. In addition to the functional requirements prescribing what  
3541 an entity does, there may also be non-functional performance and/or quality metrics that state how well  
3542 the entity does it. The relation of these policies to SOA testing is discussed further below.

3543 The identification of policies is the purview of governance (section 5.1) and the assuring of compliance  
3544 (including response to noncompliance) with policies is a matter for management (section **Error!**  
3545 **reference source not found.**).

#### 3546 5.4.1.2 Range of Test Conditions

3547 Test conditions and expected responses are detailed in the test case specification. The test conditions  
3548 should be designed to cover the areas for which the entity's response must be documented and may  
3549 include:

- 3550 • nominal conditions
- 3551 • boundaries and extremes of expected conditions
- 3552 • breaking point where the entity has degraded below a certain level or has otherwise ceased  
3553 effective functioning
- 3554 • random conditions to investigate unidentified dependencies among combinations of conditions
- 3555 • errors conditions to test error handling

3556 The specification of how each of these conditions should be tested for SOA resources, including the  
3557 infrastructure elements of the SOA ecosystem, is beyond the scope of this document but is an area that  
3558 evolves along with operational SOA experience.

#### 3559 5.4.1.3 Configuration Management of Test Artifacts

3560 The test item transmittal provides an unambiguous identification of the entity being tested, thus  
3561 REQUIRING that the configuration of the entity is appropriately tracked and documented. In addition, the  
3562 test documents (such as those specified by IEEE-829) MUST also be under a documented and  
3563 appropriately audited configuration management process, as should other resources used for testing.  
3564 The description of each artifact would follow the general description model as discussed in section  
3565 4.1.1.1; in particular, it would include a version number for the artifact and reference to the documentation  
3566 describing the versioning scheme from which the version number is derived.

3567  
3568 [EDITOR'S NOTE: TO WHAT EXTENT SHOULD CM BE EXPLICITLY INCLUDED IN THE MANAGEMENT SECTION?]

#### 3569 5.4.2 Testing and the SOA Ecosystem

3570 [EDITOR'S NOTE: THE EMPHASIS THOUGH MUCH OF THE RA IS THE LARGER ECOSYSTEM BUT WE NEED WORDS IN SECTION 3 TO  
3571 ACKNOWLEDGE THE EXISTENCE OF THE ENTERPRISE AND THAT AN ENTERPRISE (AS COMMONLY INTERPRETED) IS LIKELY MORE  
3572 CONSTRAINED AND MORE PRECISELY DESCRIBED FOR THE CONTEXT OF THE ENTERPRISE. THE ECOSYSTEM PERSPECTIVE,  
3573 THOUGH, IS STILL APPLICABLE FOR THE FOLLOWING REASONS:

- 3574
- 3575 1. A GIVEN ENTERPRISE MAY COMPRISE NUMEROUS CONSTITUENT ENTERPRISES THAT RESEMBLE THE INDEPENDENT  
3576 ENTITIES DESCRIBED FOR THE ECOSYSTEM. AN ENTERPRISE MAY ATTEMPT TO REDUCE VARIATIONS AMONG THE  
3577 CONSTITUENTS BUT THE *PARTICIPATION IN A SOA ECOSYSTEM* VIEW ENABLES SOA TO BENEFIT THE ENTERPRISE WITHOUT  
3578 REQUIRING THE ENTERPRISE ISSUES TO BE FULLY RESOLVED.
- 3579 2. RESOURCES SPECIFICALLY MOTIVATED BY THE CONTEXT OF THE ENTERPRISE CAN BE MORE READILY USED IN A  
3580 DIFFERENT CONTEXT IF ECOSYSTEM CONSIDERATIONS ARE INCLUDED AT AN EARLY STAGE. THE CHANGE IN A CONTEXT  
3581 MAY BE A FUNDAMENTAL CHANGE IN THE ENTERPRISE OR THE NEWLY DISCOVERED APPLICABILITY OF ENTERPRISE  
3582 RESOURCES TO USE OUTSIDE THE ENTERPRISE.

3583  
3584 IN THIS DOCUMENT, REFERENCE TO THE SOA ECOSYSTEM APPLIES BUT WITH POSSIBLY LESS GENERALITY TO AN ENTERPRISE USE  
3585 OF SOA.]

3586 Testing of SOA artifacts for use in the SOA ecosystem differs from traditional software testing for several  
3587 reasons. First, a highly touted benefit of SOA is to enable unanticipated consumers to make use of  
3588 services for unanticipated purposes. Examples of this could include the consumer using a service for a  
3589 result that was not considered the primary one by the provider, or the service may be used in combination  
3590 with other services in a scenario that is different from the one considered when designing for the initial  
3591 target consumer community. It is unlikely that a new consumer will push the services back to anything  
3592 resembling the initial test phase to test the new use, and thus additional paradigms for testing are  
3593 necessary. Some testing may depend on the availability of test resources made available as a service  
3594 outside the initial test community, while some testing is likely to be done as part of limited use in the

3595 operational setting. The potential responsibilities related to such "consumer testing" is discussed further  
3596 below.

3597 Secondly, in addition to consumers who interact with a service to realize the described real world effects,  
3598 the developer community is also intended to be a consumer. In the SOA vision of reuse, the developer  
3599 composes new solutions using existing services, where the existing services provides access to some  
3600 desired real world effects that are needed by the new solution. The new solution is a consumer of the  
3601 existing services, enabling repeated interactions with the existing services playing the role of reusable  
3602 components. Note, those components are used at the locations where they individually reside and are not  
3603 typically duplicated for the new solution. The new solution may itself be offered as a SOA service, and a  
3604 consumer of the service composition representing the new solution may be totally unaware of the  
3605 component services being used. (See section 4.3.4 for further discussion on service compositions.)

3606 Another difference from traditional testing is that the distributed, unbounded nature of the SOA ecosystem  
3607 makes it unlikely to have an isolated test environment that duplicates the operational environment. A  
3608 traditional testing approach often makes use of a test system that is identical to the eventual operational  
3609 system but isolated for testing. After testing is successfully completed, the tested entity would be  
3610 migrated to the operational environment, or the test environment may be delivered as part of the system  
3611 to become operational. This is not feasible for the SOA ecosystem as a whole.

3612 SOA services must be testable in the environment and under the conditions that can be encountered in  
3613 the operational SOA ecosystem. As the ecosystem is in a state of constant change, so some level of  
3614 testing is continuous through the lifetime of the service, leveraging utility services used by the ecosystem  
3615 infrastructure to monitor its own health and respond to situations that could lead to degraded  
3616 performance. This implies the test resources must incorporate aspects of the SOA paradigm, and a  
3617 category of services may be created to specifically support and enable effective monitoring and  
3618 continuous testing for resources participating in the SOA ecosystem.

3619 While SOA within an enterprise may represent a more constrained and predictable operational  
3620 environment, the composability and unanticipated use aspects are highly touted within the enterprise.  
3621 The expanded perspective on testing may not be as demanding within an enterprise but fuller  
3622 consideration of the ecosystem enables the enterprise to be more responsive should conditions change.

### 3623 **5.4.3 Elements of SOA Testing**

3624 IEEE-829 identifies fundamental aspects of testing, and many of these should carry over to SOA testing:  
3625 in particular, the identification of what is to be tested, how it is to be tested, and by whom the testing is to  
3626 be done. While IEEE-829 identifies a suggested document tree, the availability of these documents in the  
3627 SOA ecosystem is discussed below.

#### 3628 **5.4.3.1 What is to be Tested**

3629 The focus of this discussion is the SOA service. It is recognized that the infrastructure components of  
3630 any SOA environment are likely to also be SOA services and, as such, falls under the same testing  
3631 guidance. Other resources that contribute to a SOA environment may not be SOA services, but are  
3632 expected to satisfy the intent if not the letter of guidance presented here. Specific differences for such  
3633 resources are as yet largely undefined and further elaboration is beyond the scope of the SOA-RAF.

3634 The following discussion often focuses on a singular SOA service but it is implicit that any service may be  
3635 a composite of other services. As such, testing the functionality of a composite service may effectively be  
3636 testing an end-to-end business process that is being provided by the composite service. If new versions  
3637 are available for the component services, appropriate end-to-end testing of the composite may be  
3638 required in order to verify that the composite functionality is still adequately provided. The level of  
3639 required testing of an updated composite depends on policies of those providing the service, policies of  
3640 those using the service, and mission criticality of those depending on the service results.

3641 The SOA service to be tested MUST be unambiguously identified as specified by its applicable  
3642 configuration management scheme. Specifying such a scheme is beyond the scope of the SOA-RAF  
3643 other than to say the scheme should be documented and itself under configuration management.

#### 3644 5.4.3.1.1 Origin of Test Requirements

3645 In the Service Description model (Figure 21), the aspects of a service that need to be described are:

- 3646 • the service functionality and technical assumptions that underlie the functionality;
- 3647 • the policies that describe conditions of use;
- 3648 • the service interface that defines information exchange with the service;
- 3649 • service reachability that identifies how and where message exchange is to occur; and
- 3650 • metrics access for any participant to have information on how a service is performing.

3651 Service testing must provide adequate assurance that each of these aspects is operational as defined.

3652 The information in the service description comes from different sources. The functionality is defined  
3653 through whatever process identifies needs and the community for which these needs are addressed. The  
3654 process may be ad hoc as serves the prospective service owner or strictly governed, but defining the  
3655 functionality is an essential first step in development. It is also an early and ongoing focus of testing to  
3656 ensure the service accurately reflects the described functionality and the described functionality  
3657 accurately addresses the consumer needs.

3658 Policies define the conditions of development and conditions of use for a service and are typically  
3659 specified as part of the governance process. Policies constraining service development, such as coding  
3660 standards and best practices, require appropriate testing and auditing during development to ensure  
3661 compliance. While the governance process identifies development policies, these are likely to originate  
3662 from the technical community responsible for development activities. Policies that define conditions of  
3663 use often define business practices that service owners and providers or those responsible for the SOA  
3664 infrastructure want followed. These policies are initially tested during service development and are  
3665 continuously monitored during the operational lifetime of the service.

3666 The testing of the service interface and service reachability are often related but essentially reflect  
3667 different motivations and needs. The service interface is specified as a joint product of the service  
3668 owners and providers who define service functionality, the prospective consumer community, the service  
3669 developer, and the governance process. The semantics of the information model must align with the  
3670 semantics of those who consume the service in order for there to be meaningful exchange of information.  
3671 The structure of the information is influenced by the consumer semantics and the requirements and  
3672 constraints of the representation as interpreted by the service developer. The service process model that  
3673 defines actions which can be performed against a service and any temporal dependencies derive from  
3674 the defined functionality and may be influenced by the development process. Any of these constraints  
3675 may be identified and expressed as policy through the governance process.

3676 Service reachability conditions are the purview of the service provider who identifies the service endpoint  
3677 and the protocols recognized at the endpoint. These may be constrained by governance decisions on  
3678 how endpoint addresses may be allocated and what protocols should be used.

3679 While the considerations for defining the service interface derive from several sources, testing of the  
3680 service interface is more straightforward and isolated in the testing process. At any point where the  
3681 interface is modified or exposes a new resource, the message exchange should be monitored both to  
3682 ensure the message reaches its intended destination and it is parsed correctly once received. Once an  
3683 interface has been shown to function properly, it is unlikely to fail later unless something fundamental to  
3684 the service changes.

3685 The service interface is also tested when the service endpoint changes. Testing of the endpoint ensures  
3686 message exchange can occur at the time of testing and the initial testing shows the interface is being  
3687 processed properly at the new endpoint. Functioning of a service endpoint at one time does not  
3688 guarantee it is functioning at another time, e.g. the server with the endpoint address may be down,  
3689 making testing of service reachability a continual monitoring function through the life of the service's use  
3690 of the endpoint. Also, while testing of the service endpoint is a necessary and most commonly noted part  
3691 of the test regiment, it is not in itself sufficient to ensure the other aspects of testing discussed in this  
3692 section.

3693 Finally, governance is impossible without the collection of metrics against which service behavior can be  
3694 assessed. Metrics are also a key indicator for consumers to decide if a service is adequate for their  
3695 needs. For instance, the average response time or the recent availability can be determining factors even



3696 if there are no rules or regulations promulgated through the governance process against which these  
3697 metrics are assessed. The available metrics are a combination of those expected by the consumer  
3698 community and those mandated through the governance process. The total set of metrics will evolve  
3699 over time with SOA experience. Testing of the services that gather and provide access to the metrics will  
3700 follow testing as described in this section, but for an individual service, testing will ensure that the metrics  
3701 access indicated in the service description is accurate.

3702 The individual test requirements highlight aspects of the service that testing must consider but testing  
3703 must establish more than isolated behavior. The emphasis is the holistic results of interacting with the  
3704 service in the SOA environment. Recall that the execution context is the set of agreements between a  
3705 consumer and a provider that define the conditions under which service interaction occurs. The  
3706 agreements are expected to be predominantly the acceptance of the standard conditions as enumerated  
3707 by the service provider, but it may include the identification of alternate conditions that will govern the  
3708 interaction.

3709 For example, the provider may prefer a policy where it can sell the contact information of its consumers  
3710 but will honor the request of a consumer to keep such information private. The identification of the  
3711 alternate privacy policy is part of the execution context, and it is the application of and compliance with  
3712 this policy that operational monitoring will attempt to measure. The collection of metrics showing this  
3713 condition is indeed met when chosen is considered part of the ongoing testing of the service.

3714 Other variations in the execution context also require monitoring to ensure that different combinations of  
3715 conditions perform together as desired. For example, if a new privacy policy takes additional resources to  
3716 apply, this may affect quality of service and propagate other effects. These could not be tested during the  
3717 original testing if the alternate policy did not exist at that time.

#### 3718 **5.4.3.1.2 Testing Against Non-Functional Requirements**

3719 Testing against non-functional requirements constitutes testing of business usability of the service. In a  
3720 marketplace of services, non-functional characteristics may be the primary differentiator between services  
3721 that produce essentially the same real world effects.

3722 As noted in the previous section, non-functional characteristics are often associated with policies or other  
3723 terms of use and may be collected in service level contracts offered by the service providers. Non-  
3724 functional requirements may also reflect the network and hardware infrastructure that support  
3725 communication with the service, and changes may impact quality of service. The service consumer and  
3726 even the service provider may not be aware of all such infrastructure changes but the changes may  
3727 manifest in shared states that impact the usability of the service.

3728 In general, a change in the non-functional requirements results in a change to the execution context, but  
3729 as with any collection of information that constitutes a description, the execution context is unable to  
3730 explicitly capture all non-functional requirements that may apply. A change in non-functional  
3731 requirements, whether explicitly part of the execution context or an implicit contributor, may require  
3732 retesting of the service even if its functionality and the implementation of the functionality has not  
3733 changed. Depending on the circumstances, retesting may require a formal recertifying of end-to-end  
3734 behavior or more likely will be part of the continuous monitoring that applies throughout the service  
3735 lifetime.

#### 3736 **5.4.3.1.3 Testing Content and the Interests of Consumers**

3737 As noted in section 5.4.1.1, testing may involve verification of conformance with respect to policies and  
3738 technical specifications and validation with respect to sufficiency of functionality to meet some prescribed  
3739 use. It may also include demonstration of performance and quality aspects. For some of these items,  
3740 such as demonstrating the business processes followed in developing the service or the use of standards  
3741 in implementing the service, the testing or relevant auditing is done internal to the service development  
3742 process and follows traditional software testing and quality assurance. If it is believed of value to  
3743 potential consumers, information about such testing could be included in the service description.  
3744 However, it is not required that all test or compliance artifacts be available to consumers, as many of the  
3745 details tested may be part of the opacity of the service implementation.

3746 Some aspects of the service being tested will reflect directly on the real world effects realized through  
3747 interaction with the service. In these cases, it is more likely that testing results will be directly relevant to

3748 potential consumers. For example, if the service was designed to correspond to certain elements of a  
3749 business process or that a certain workflow is followed, testing should verify that the real world effects  
3750 reflect that the business process or workflow were satisfactorily captured.

3751 The testing may also need to demonstrate that specified conditions of use are satisfied. For example,  
3752 policies may be asserted that require certain qualifications of or impose restrictions on the consumers  
3753 who may interact with the service. The service testing must demonstrate that the service independently  
3754 enforces the policies or it provides the required information exchanges with the SOA ecosystem so other  
3755 resources can ensure the specified conditions.

3756 The completeness of the testing, both in terms of the features tested and the range of parameters for  
3757 which response is tested, depends on the context of expected use: the more critical the use, the more  
3758 complete the testing. There are always limits on the resources available for testing, if nothing else than  
3759 the service must be available for use in a finite amount of time.

3760 This again emphasizes the need for adequate documentation to be available. If the original testing is  
3761 very thorough, it may be adequate for less demanding uses in the future. If the original testing was more  
3762 constrained, then well-documented test results establish the foundation on which further testing can be  
3763 defined and executed.

#### 3764 **5.4.3.2 How Testing is to be Done**

3765 Testing should follow well-defined methodologies and, if possible, should reuse test artifacts that have  
3766 proven generally useful for past testing. For example, IEEE-829 notes that test cases are separated from  
3767 test designs to allow for use in more than one design and to allow for reuse in other situations. In the  
3768 SOA ecosystem, description of such artifacts, as with description of a service, enables awareness of the  
3769 item and describes how the artifact may be accessed or used.

3770 As with traditional testing, the specific test procedures and test case inputs are important so the tests are  
3771 unambiguously defined and entities can be retested in the future. Automated testing and regression  
3772 testing may be more important in the SOA ecosystem in order to re-verify a service is still acceptable  
3773 when incorporated in a new use. For example, if a new use requires the services to deal with input  
3774 parameters outside the range of initial testing, the tests could be rerun with the new parameters. If the  
3775 testing resources are available to consumers within the SOA ecosystem, the testing as designed by test  
3776 professionals could be consumed through a service accessed by consumers, and their results could  
3777 augment those already in place. This is discussed further in the next section.

#### 3778 **5.4.3.3 Who Performs the Testing**

3779 As with any software, the first line of testing is unit testing done by software developers. It is likely that  
3780 initial testing will be done by those developing the software but may also be done independently by other  
3781 developers. For SOA development, unit testing is likely confined to a development sandbox isolated from  
3782 the SOA ecosystem.

3783 SOA testing will differ from traditional software testing in that testing beyond the development sandbox  
3784 must incorporate aspects of the SOA ecosystem, and those doing the testing must be familiar with both  
3785 the characteristics and responses of the ecosystem and the tools, especially those available as services,  
3786 to facilitate and standardize testing. Test professionals will know what level of assurance must be  
3787 established as the exposure of the service to the ecosystem and ecosystem to the service increases  
3788 towards operational status. These test professionals may be internal resources to an organization or may  
3789 evolve as a separate discipline provided through external contracting.

3790 As noted above, it is unlikely that a complete duplicate of the SOA ecosystem will be available for isolated  
3791 testing, and thus use of ecosystem resources will manifest as a transition process rather than a step  
3792 change from a test environment to an operational one. This is especially true for new composite services  
3793 that incorporate existing operational services to achieve the new functionality. The test professionals will  
3794 need to understand the available resources and the ramifications of this transition.

3795 As with current software development, a stage beyond work by test professionals will make use of a  
3796 select group of typical users, commonly referred to as beta testers, to report on service response during  
3797 typical intended use. This establishes fitness by the consumers, providing final validation of previously  
3798 verified processes, requirements, and final implementation.

3799 In traditional software development, beta testing is the end of testing for a given version of the software.  
3800 However, although the initial test phase can establish an appropriate level of confidence consistent with  
3801 the designed use for the initial target consumer community, the operational service will exist in an  
3802 evolving ecosystem, and later conditions of use may differ from those thought to be sufficient during the  
3803 initial testing. Thus, operational monitoring becomes an extension of testing through the service lifetime.  
3804 This continuous testing will attempt to ensure that a service does not consume an inordinate amount of  
3805 ecosystem resources or display other behavior that degrades the ecosystem, but it will not undercover  
3806 functional errors that may surface over time.

3807 As with any software, it is the responsibility of the consumers to consider the reasonableness of solutions  
3808 in order to spot errors in either the software or the way the software is being used. This is especially  
3809 important for consumers with unanticipated uses that may go beyond the original test conditions. It is  
3810 unlikely the consumers will initiate a new round of formal testing unless the new use requires a  
3811 significantly higher level of confidence in the service. Rather the consumer becomes a new extension to  
3812 the testing regiment. Obvious testing would include a sanity check of results during the new use.  
3813 However, if the details of legacy testing are associated with the service through the service description  
3814 and if testing resources are available through automated testing services, then the new consumers can  
3815 rerun and extend previous testing to include the extended test conditions. If the test results are  
3816 acceptable, these can be added to the documentation of previous results and become the extended basis  
3817 for future decisions by prospective consumers on the appropriateness of the service. If the results are not  
3818 acceptable or in some way questionable, the responsible party for the service or testing professionals can  
3819 be brought in to decide if remedial action is necessary.

#### 3820 5.4.3.4 How Testing Results are Reported

3821 For any SOA service, an accurate reporting of the testing a service has undergone and the results of the  
3822 testing is vital to consumers deciding whether a service is appropriate for intended use. Appropriateness  
3823 may be defined by a consumer organization and require specific test regiments culminating in a  
3824 certification; appropriateness could be established by accepting testing and certifications that have been  
3825 conferred by others.

3826 The testing and certification information should be identified in the service description. Referring to the  
3827 general description model of *Figure 12*, tests conducted by or under a request from the service owner (see  
3828 ownership in section **Error! Reference source not found.**) would be captured under Annotations from  
3829 wners. Testing done by others, such as consumers with unanticipated uses, could be associated through  
3830 Annotations from 3rd Parties. The annotations should clearly indicate what was tested, how the testing  
3831 was done, who did the testing, and the testing results. The clear description of each of these artifacts and  
3832 of standardized testing protocols for various levels of sophistication and completeness of testing would  
3833 enable a common understanding and comparison of test coverage. It will also make it more  
3834 straightforward to conduct and report on future testing, facilitating the maintenance of the service  
3835 description.

3836 Consumer testing and the reporting of results raises additional issues. While stating who did the testing  
3837 is mandatory, there may be formal requirements for authentication of the tester to ensure traceability of  
3838 the testing claims. In some circumstances, persons or organizations would not be allowed to state testing  
3839 claims unless the tester was an approved entity. In other cases, ensuring the tester had a valid email  
3840 may be sufficient. In either case, it would be at the discretion of the potential consumer to decide what  
3841 level of authentication was acceptable and which testers are considered authoritative in the context of  
3842 their anticipated use.

3843 Finally, in a world of openly shared information, we would see an ever-expanding set of testing  
3844 information as new uses and new consumers interact with a service. In reality, these new uses may  
3845 represent proprietary processes or classified use that should only be available to authorized parties.  
3846 Testing information, as with other elements of description, may require special access controls to ensure  
3847 appropriate access and use.

#### 3848 5.4.4 Testing SOA Services

3849 Testing of SOA services should be consistent with the SOA paradigm. In particular, testing resources  
3850 and artifacts should be visible in support of service interaction between providers and consumers, where

3851 here the interaction is between the testing resource and the tester. In addition, the idea of opacity of the  
3852 implementation should limit the details that need to be available for effective use of the test resources.  
3853 Testing that requires knowledge of the internal structure of the service or its underlying capability should  
3854 be performed as part of unit testing in the development sandbox, and should represent a minimum level  
3855 of confidence before the service begins its transition to further testing and eventual operation in the SOA  
3856 ecosystem.

#### 3857 **5.4.4.1 Progression of SOA Testing**

3858 Software testing is a gradual exercise going from micro inspection to testing macro effects. The first step  
3859 in testing is likely the traditional code reviews. SOA considerations would account for the distributed  
3860 nature of SOA, including issues of distributed security and best practices to ensure secure resources. It  
3861 would also set the groundwork for opacity of implementation, hiding programming details and simplifying  
3862 the use of the service.

3863 Code review is likely followed by unit testing in a development sandbox isolated from the operational  
3864 environment. The unit testing is done with full knowledge of the service internal structure and knowledge  
3865 of resources representing underlying capabilities. It tests the interface to ensure exchanged messages  
3866 are as specified in the service description and the messages can be parsed and interpreted as intended.  
3867 Unit testing also verifies intended functionality and that the software has dealt correctly with internal  
3868 dependencies, such as structure of a file system or access to other dedicated resources.

3869 Some aspects of unit testing require external dependencies be satisfied, and this is often done using  
3870 mock objects to substitute for the external resources. In particular, it will likely be necessary to include  
3871 mocks of existing operational services, both those provided as part of the SOA infrastructure and services  
3872 from other providers.

#### 3873 **Service Mock**

3874 A service mock is an entity that mimics some aspect of the performance of an operational service  
3875 without committing to the real world effects that the operational service would produce.

3876 Mocks are discussed in detail in sections 5.4.4.3 and 5.4.4.4.

3877 After unit testing has demonstrated an adequate level of confidence in the service, the testing must  
3878 transition from the tightly controlled environment of the development sandbox to an environment that  
3879 more clearly resembles the operational SOA ecosystem or, at a minimum, the intended enterprise. While  
3880 sandbox testing will use simple mocks of some aspects of the SOA environment, such as an interface to  
3881 a security service without the security service functionality, the dynamic nature of SOA makes a full  
3882 simulation infeasible to create or maintain. This is especially true when a new composite service makes  
3883 use of operational services provided by others. Thus, at some point before testing is complete, the  
3884 service will need to demonstrate its functionality by using resources and dealing with conditions that only  
3885 exist in the full ecosystem or the intended enterprise. Some of these resources may still provide test  
3886 interfaces -- more on this below -- but the interfaces will be accessible using the SOA environment and  
3887 not just implemented for the sandbox.

3888 At this stage, the opacity of the service becomes important as the details of interacting with the service  
3889 now rely on correct use of the service interface and not knowledge of the service internals. The workings  
3890 of the service will only be observable through the real world effects realized through service interactions  
3891 and external indications that conditions of use, such as user authentication, are satisfied. Monitoring the  
3892 behavior of the service will depend on service interfaces that expose internal monitoring or provide  
3893 required information to the SOA infrastructure monitoring function. The monitoring required to test a new  
3894 service is likely to have significant overlap with the monitoring the SOA infrastructure includes to monitor  
3895 its own health and to identify and isolate behavior outside of acceptable bounds. This is exactly what is  
3896 needed as part of service testing, and it is reasonable to assume that the ecosystem transition includes  
3897 use of operational monitoring rather than solely dedicated monitoring for each service being tested.

3898 Use of SOA monitoring resources during the explicit testing phase sets the stage for monitoring and a  
3899 level of continual testing throughout the service lifetime.

#### 3900 **5.4.4.2 Testing Traditional Dependencies vs. Service Interactions**

3901 A SOA service is not required to make use of other operational services beyond what may be required for  
3902 monitoring by the ecosystem infrastructure. The service can implement hardcoded dependencies which  
3903 have been tested in the development sandbox through the use of dedicated mocks. While coordination  
3904 may be required with real data sources during integration testing, the dependencies can be constrained to  
3905 things that can be tested in a more traditional manner. Policies can also be set to restrict access to pre-  
3906 approved users, and thus the question of unanticipated users and unanticipated uses can be eliminated.  
3907 Operational readiness can be defined in terms of what can be proven in isolated testing. While all this  
3908 may provide more confidence in the service for its designed purpose, such a service will not fully  
3909 participate in the benefits or challenges of the ecosystem. This is akin to filling a swimming pool with sea  
3910 water and having someone in the pool say they are swimming in the ocean.

3911 In considering the testing needed for a fully participating service, consider the example of a new  
3912 composite service that combines the real world effects and complies with the conditions of use of five  
3913 existing operational services. The developer of the composite service does not own any of the  
3914 component services and has limited, if any, ability to get the distributed owners to do any customization.  
3915 The developer also is limited by the principle of opacity to information comprising the service description,  
3916 and does not know internal details of the component services. The developer of the composite service  
3917 must use the component services as they exist as part of the SOA environment, including what is  
3918 provided to support testing by new users. This introduces requirements for what is needed in the way of  
3919 service mocks.

#### 3920 **5.4.4.3 Use of Service Mocks**

3921 Service mocks enables the tested service to respond to specific features of an operational service that is  
3922 being used as a component. It allows service testing to proceed without needing access to or with only  
3923 limited engagement with the component service. Mocks can also mimic difficult to create situations for  
3924 which it is desired to test the new service response. For composite services using multiple component  
3925 services, mocks may be used in combination to function for any number of the components. Note, when  
3926 using service mocks, it is important to remember that it is not the component service that is being tested  
3927 (although anomalous behavior may be uncovered during testing) but the use of the component in the new  
3928 composite.

3929 Individual service mocks can emphasize different features of the component service they represent but  
3930 any given mock does not have to mimic all features. For example, a mock of the service interface can  
3931 echo a sent message and demonstrate the message is reaching its intended destination. A mock could  
3932 go further and parse the sent message to demonstrate the message not only reached its destination but  
3933 was understood. As a final step, the mock could report back what actions would have been taken by the  
3934 component service and what real world effects would result. If the response mimicked the operational  
3935 response, functional testing could proceed as if the real world effect actually occurred.

3936 There are numerous ways to provide mock functionality. The service mock could be a simulation of the  
3937 operational service and return simulated results in a realistic response message or event notification. It is  
3938 also possible for the operational service to act as its own mock and simply not execute the commit stage  
3939 of its functionality. The service mock could use a combination of simulation and service action without  
3940 commit to generate a report of what would have occurred during the defined interaction with the  
3941 operational service.

3942 As the service proceeds through testing, mocks should be systematically replaced by the component  
3943 resources accessed through their operational interfaces. Before beta testing begins, end-to-end testing,  
3944 i.e. proceeding from the beginning of the service interaction to the resulting real world results, should be  
3945 accomplished using component resources via their operational interfaces.

#### 3946 **5.4.4.4 Providers of Service Mocks**

3947 In traditional testing, it is often the test professionals who design and develop the mocks, but in the  
3948 distributed world of SOA, this may not be efficient or desirable.

3949 In the development sandbox, it is likely the new service developer or test professionals working with the  
3950 developer will create mocks adequate for unit testing. Given that most of this testing is to verify the new  
3951 service is performing as designed, it is not necessary to have high fidelity models of other resources

3952 being accessed. In addition, given opacity of SOA implementation, the developer of the new service may  
3953 not have sufficient detailed knowledge of a component service to build a detailed mock of the component  
3954 service functionality. Sharing existing mocks at this stage may be possible but the mocks would need to  
3955 be implemented in the sandbox, and for simple models it is likely easier to build the mock from scratch.

3956 As testing begins its transition to the wider SOA environment, mocks may be available as services. For  
3957 existing resources, it is possible that an Open Source model could evolve where service mocks of  
3958 available functions can be catalogued and used during initial interaction of the tested service and the  
3959 operational environment. Widely used functions may have numerous service mocks, some mimicking  
3960 detailed conditions within the SOA infrastructure. However, the Open Source model is less likely to be  
3961 sufficient for specialty services that are not widely used by a large consumer community.

3962 The service developer is probably best qualified for also developing more detailed service mocks or for  
3963 mock modes of operational services. This implies that in addition to their operational interfaces, services  
3964 will routinely provide test interfaces to enable service mocks to be used as services. As noted above, a  
3965 new service developer wanting to build a mock of component services is limited to the description  
3966 provided by the component service developer or owner. The description typically will detail real world  
3967 effects and conditions of use but will not provide implementation details, some of which may be  
3968 proprietary. Just as important in the SOA ecosystem, if it becomes standard protocol for developers to  
3969 create service mocks of their own services, a new service developer is only responsible for building his  
3970 own mocks and can expect other mocks to be available from other developers. This reduces duplication  
3971 of effort where multiple developers would be trying to build the same mocks from the same insufficient  
3972 information. Finally, a service developer is probably best qualified to know when and how a service mock  
3973 should be updated to reflect modified functionality or message exchange.

3974 It is also possible that testing organizations will evolve to provide high-fidelity test harnesses for new  
3975 services. The harnesses would allow new services to plug into a test environment and would facilitate  
3976 accessing mocks of component services. However, it will remain a constant challenge for such  
3977 organizations to capture evolving uses and characteristics of service interactions in the real SOA  
3978 environment and maintain the fidelity and accuracy of the test systems.

#### 3979 **5.4.4.5 Fundamental Questions for SOA Testing**

3980 In order for the transition to the SOA operational environment to proceed, it is necessary to answer two  
3981 fundamental questions:

- 3982 • Who provides what testing resources for the SOA operational environment, e.g. mocks of  
3983 interfaces, mocks of functionality, monitoring tools?
- 3984 • What testing needs to be accomplished before operational environment resources can be  
3985 accessed for further testing?

3986 The discussion in section 5.4.4.4 notes various levels of sophistication of service mocks and different  
3987 communities are likely to be responsible for different levels. Section 5.4.4.4 advocates a significant role  
3988 for service developers, but there needs to be community consensus that such mocks are needed and that  
3989 service developers will agree to fulfilling this role. There is also a need for consensus as to what tools  
3990 should be available as services from the SOA infrastructure.

3991 As for use of the service mocks and SOA environment monitoring services, practical experience is  
3992 needed upon which guidelines can be established for when a new service has been adequately tested to  
3993 proceed with a greater level of exposure with the SOA environment. Malfunctioning services could cause  
3994 serious problems if they cannot be identified and isolated. On the other hand, without adequate testing  
3995 under SOA operational conditions, it is unlikely that problems can be uncovered and corrected before  
3996 they reach an operational stage.

3997 As noted in section 5.4.4.2, some of these questions can be avoided by restricting services to more  
3998 traditional use scenarios. However, such restriction will limit the effectiveness of SOA use and the result  
3999 will resemble the constraints of traditional integration activities we are trying to move beyond.

#### 4000 **5.4.5 Architectural Implications for SOA Testing**

4001 The discussion of SOA Testing indicates numerous architectural implications on the SOA ecosystem:

- 4002 • The distributed, boundary-less nature of the SOA ecosystem makes it infeasible to create  
4003 and maintain a single mock of the entire ecosystem to support testing activities.
- 4004 • A standard suite of monitoring services needs to be defined, developed, and maintained.  
4005 This should be done in a manner consistent with the evolving nature of the ecosystem.
- 4006 • Services should provide interfaces that support access in a test mode.
- 4007 • Testing resources must be described and their descriptions must be catalogued in a  
4008 manner that enables their discovery and access.
- 4009 • Guidelines for testing and ecosystem access need to be established and the ecosystem  
4010 must be able to enforce those guidelines asserted as policies.
- 4011 • Services should be available to support automated testing and regression testing.
- 4012 • Services should be available to facilitate updating service description by anyone who has  
4013 performed testing of a service.

---

## 4014 6 Conformance

4015 This Reference Architecture Framework is an abstract architectural description of Service Oriented  
4016 Architecture, which means that it is especially difficult to construct tests for conformance to the  
4017 architecture. In addition, conformance to an architectural specification does not, by itself, guarantee any  
4018 form of interoperability between multiple implementations.

4019 However, it *is* possible to decide whether or not a given architecture is conformant to an architectural  
4020 description such as this one. In discussions of conformance we use the term **target architecture** to  
4021 identify the (typically concrete) architecture that may be viewable as conforming to the abstract principles  
4022 outlined in this document.

### 4023 Target Architecture

4024 A target architecture is an architectural description of a system that is intended to be viewed as  
4025 conforming to the SOA-RAF.

4026 While we cannot guarantee interoperability between target architectures (or more specifically between  
4027 applications and systems residing within the ecosystems of those target architectures), interoperability  
4028 between target architectures is promoted by conformance to this Reference Architecture Framework as it  
4029 reduces the semantic impedance mismatch between the different ecosystems.

4030 The primary measure of conformance is whether given concepts as described in document have  
4031 corresponding concepts in the target architecture. Such a correspondence **MUST** honor the relationships  
4032 identified within this document for the target architecture to be considered conforming.

4033 For example, in Section 3.1.3.1 we identify resource as a key concept. A resource is associated with an  
4034 owner and a number of identifiers. For a target architecture to conform to the SOA-RAF, it must be  
4035 possible to find corresponding concepts of resource, identifier and owner within the target architecture:  
4036 say *entity*, *token* and *user*. Furthermore, the relationships between *entity*, *token* and *user* **MUST** mirror  
4037 the relationships between resource, identifier and owner appropriately.

4038 Clearly, such correspondence is simpler if the terminology within the target architecture is identical to that  
4039 in the SOA-RAF. But so long as the 'graph' of concepts and relationships is consistent, that is all that is  
4040 required for the target architecture to conform to this Reference Architecture Framework.

4041 [EDITOR'S NOTE: The conformance section is not complete]

4042



---

4043 **A. Acknowledgements**

4044 The following individuals have participated in the work of the technical committee responsible for creation  
4045 of this specification and are gratefully acknowledged:

4046 **Participants:**

- 4047 Chris Bashioum, MITRE Corporation
- 4048 Rex Brooks, Individual Member
- 4049 Peter Brown, Individual Member
- 4050 Scott Came, Search Group Inc.
- 4051 Joseph Chiusano, Booz Allen Hamilton
- 4052 Robert Ellinger, Northrop Grumman Corporation
- 4053 David Ellis, Sandia National Laboratories
- 4054 Jeff A. Estefan, Jet Propulsion Laboratory
- 4055 Don Flinn, Individual Member
- 4056 Anil John, Johns Hopkins University
- 4057 Ken Laskey, MITRE Corporation
- 4058 Boris Lublinsky, Nokia Corporation
- 4059 Francis G. McCabe, Individual Member
- 4060 Christopher McDaniels, USSTRATCOM
- 4061 Tom Merkle, Lockheed Martin Corporation
- 4062 Jyoti Namjoshi, Patni Computer Systems Ltd.
- 4063 Duane Nickull, Adobe Inc.
- 4064 James Odell, Associate
- 4065 Michael Poulin, Fidelity Investments
- 4066 Michael Stiefel, Associate
- 4067 Danny Thornton, Northrop Grumman
- 4068 Timothy Vibbert, Lockheed Martin Corporation
- 4069 Robert Vitello, New York Dept. of Labor

4070 The committee would particularly like to underline the significant contributions made by Rex Brooks, Jeff  
4071 Estefan, Ken Laskey, Boris Lublinsky, Frank McCabe, Michael Poulin and Danny Thornton

---

## 4072 B. Index of Defined Terms

- 4073 The first page number refers to the first use of the term. The second, where necessary, refers to the page  
4074 where the term is formally defined.
- 4075 Action
- 4076 Action Level Real World Effect
- 4077 Actor
- 4078 Architecture
- 4079 Architectural Description
- 4080 Authority
- 4081 Business Processes
- 4082 Capability
- 4083 Choreography
- 4084 Commitment
- 4085 Communicative Action
- 4086 Constitution
- 4087 Contract
- 4088 Delegate
- 4089 Description
- 4090 Endpoint
- 4091 Enterprise
- 4092 Governance
- 4093 Governance Framework
- 4094 Governance Processes
- 4095 Identifier
- 4096 Identity
- 4097 Joint Action
- 4098 Leadership
- 4099 Life-cycle manageability
- 4100 Logical Framework
- 4101 Management
- 4102 Management Policy
- 4103 Management Service
- 4104 Manageability Capability
- 4105 Message Exchange
- 4106 **Model**
- 4107 Obligation
- 4108 Objective
- 4109 Operations
- 4110 Orchestration
- 4111 Ownership
- 4112 Ownership Boundary

4113 Participant  
4114 Peer  
4115 Permission  
4116 Policy  
4117 Policy Conflict  
4118 Policy Conflict Resolution  
4119 Policy Constraint  
4120 Policy Decision  
4121 Policy Enforcement  
4122 Policy Framework  
4123 Policy Object  
4124 Policy Ontology  
4125 Policy Owner  
4126 Policy Subject  
4127 Presence  
4128 Private State  
4129 Protocol  
4130 Public Semantics  
4131 Qualification  
4132 Real World Effect  
4133 Regulation  
4134 Resource  
4135 Responsibility  
4136 Right  
4137 Risk  
4138 Role  
4139 Rule  
4140 Security  
4141 Semantic Engagement  
4142 Service Action  
4143 Service Consumer  
4144 Service Level Real World Effect  
4145 Service Mediator  
4146 Service Provider  
4147 Shared State  
4148 Skill  
4149 Social Structure  
4150 Stakeholder  
4151 State  
4152 System  
4153 System Stakeholder  
4154 Trust

4155 View  
4156 Viewpoint

---

## 4157 C. The Unified Modeling Language, UML

4158 **Error! Reference source not found.** illustrates an annotated example of a UML class diagram that is  
4159 sed to represent a visual model depiction of the Resources Model in the *Participation in a SOA*  
4160 *Ecosystem* view (Section **Error! Reference source not found.**).

4161  
4162 *Figure 44 Example UML class diagram—Resources.*

4163 Lines connecting boxes (classifiers) represent associations between things. An association has two roles  
4164 (one in each direction). A role can have cardinality, for example, one or more (“1..\*”) stakeholders own  
4165 zero or more (“0..\*”) resources. The role from classifier A to B is labeled closest to B, and vice versa, for  
4166 example, the role between resource to Identity can be read as resource embodies Identity, and Identity  
4167 denotes a resource.

4168 Mostly, we use named associations, which are denoted with a verb or verb phrase associated with an  
4169 arrowhead. A named association reads from classifier A to B, for example, one or more stakeholders  
4170 owns zero or more resources. Named associations are a very effective way to model relationships  
4171 between concepts.

4172 An open diamond (at the end of an association line) denotes an aggregation, which is a part-of  
4173 relationship, for example, Identifiers are part of Identity (or conversely, Identity is made up of Identifiers).

4174 A stronger form of aggregation is known as composition, which involves using a filled-in diamond at the  
4175 end of an association line (not shown in above diagram). For example, if the association between Identity  
4176 and Identifier were a composition rather than an aggregation as shown, deleting Identity would also  
4177 delete any owned Identifiers. There is also an element of exclusive ownership in a composition  
4178 relationship between classifiers, but this usually refers to specific instances of the owned classes  
4179 (objects).

4180 This is by no means a complete description of the semantics of all diagram elements that comprise a  
4181 UML class diagram, but rather is intended to serve as an illustrative example for the reader. It should be  
4182 noted that the SOA-RAF utilizes additional class diagram elements as well as other UML diagram types  
4183 such as sequence diagrams and component diagrams. The reader who is unfamiliar with the UML is  
4184 encouraged to review one or more of the many useful online resources and book publications available  
4185 describing UML (see, for example, [www.uml.org](http://www.uml.org)).

---

## 4186 D. Critical Factors Analysis

4187 A critical factors analysis (CFA) is an analysis of the key properties of a project. A CFA is analyzed in  
4188 terms of the goals of the project, the critical factors that will lead to its success and the measurable  
4189 requirements of the project implementation that support the goals of the project. CFA is particularly  
4190 suitable for capturing quality attributes of a project, often referred to as “non-functional” or “other-than-  
4191 functional” requirements: for example, security, scalability, wide-spread adoption, and so on. As such,  
4192 CFA complements rather than attempts to replace other requirements capture techniques.

### 4193 D.1 Goals

4194 A goal is an overall target that you are trying to reach with the project. Typically, goals are hard to  
4195 measure by themselves. Goals are often directed at the potential consumer of the product rather than the  
4196 technology developer.

### 4197 Critical Success Factors

4198 A critical success factor (CSF) is a property, sub-goal that directly supports a goal and there is strong  
4199 belief that without it the goal is unattainable. CSFs themselves are not necessarily measurable in  
4200 themselves.

### 4201 Requirements

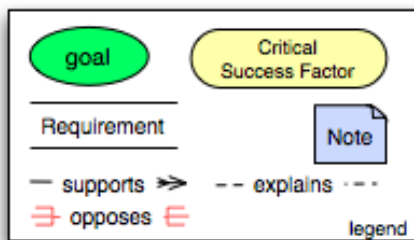
4202 A requirement is a specific measurable property that directly supports a CSF. The key here is  
4203 measurability: it should be possible to unambiguously determine if a requirement has been met. While  
4204 goals are typically directed at consumers of the specification, requirements are focused on technical  
4205 aspects of the specification.

### 4206 CFA Diagrams

4207 It can often be helpful to illustrate graphically the key concepts and relationships between them. Such  
4208 diagrams can act as effective indices into the written descriptions of goals etc., but is not intended to  
4209 replace the text.

4210 The legend:

4211



4212 illustrates the key elements of the graphical notation. Goals are written in round ovals, critical success  
4213 factors are written in round-ended rectangles and requirements are written using open-ended rectangles.  
4214 The arrows show whether a CSF/goal/requirement is supported by another element or opposed by it. This  
4215 highlights the potential for conflict in requirements.  
4216

4217

---

## 4218 E. Relationship to other SOA Open Standards

4219 The white paper “Navigating the SOA Open Standards Landscape Around Architecture” issued jointly by  
4220 OASIS, OMG, and The Open Group **[SOA-NAV]** was written to help the SOA community at large  
4221 navigate the myriad of overlapping technical products produced by these organizations with specific  
4222 emphasis on the “A” in SOA, i.e., Architecture.

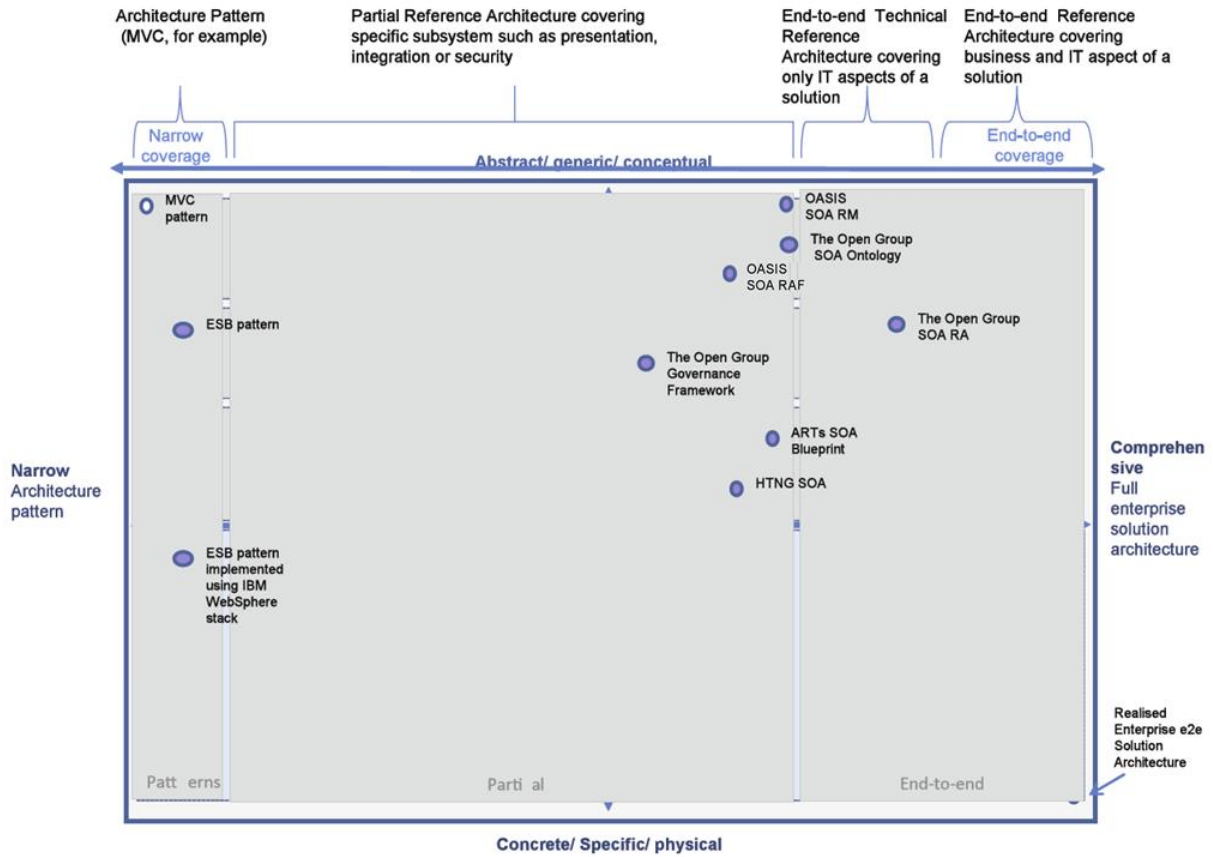
4223 The white paper explains and positions standards for SOA reference models, ontologies, reference  
4224 architectures, maturity models, modeling languages, and standards work on SOA governance. It outlines  
4225 where the works are similar, highlights the strengths of each body of work, and touches on how the work  
4226 can be used together in complementary ways. It is also meant as a guide to users for selecting those  
4227 specifications most appropriate for their needs.

4228 While the understanding of SOA and SOA Governance concepts provided by these works is similar, the  
4229 evolving standards are written from different perspectives. Each specification supports a similar range of  
4230 opportunity, but has provided different depths of detail for the perspectives on which they focus. Although  
4231 the definitions and expressions may differ, there is agreement on the fundamental concepts of SOA and  
4232 SOA Governance.

4233 The following is a summary taken from **[SOA-NAV]** of the positioning and guidance on the specifications:

- 4234 • The OASIS Reference Model for SOA (SOA RM) is the most abstract of the specifications  
4235 positioned. It is used for understanding core SOA concepts
- 4236 • The Open Group SOA Ontology extends, refines, and formalizes some of the core concepts of  
4237 the SOA RM. It is used for understanding core SOA concepts and facilitates a model-driven  
4238 approach to SOA development.
- 4239 • The OASIS Reference Architecture Foundation for SOA (this document) is an abstract,  
4240 foundational reference architecture addressing a broader ecosystem viewpoint for building and  
4241 interacting within the SOA paradigm. It is used for understanding different elements of SOA, the  
4242 completeness of SOA architectures and implementations, and considerations for reaching across  
4243 ownership boundaries where there is no single authoritative entity for SOA and SOA governance.
- 4244 • The Open Group SOA Reference Architecture is a layered architecture from consumer and  
4245 provider perspective with cross cutting concerns describing these architectural building blocks  
4246 and principles that support the realizations of SOA. It is used for understanding the different  
4247 elements of SOA, deployment of SOA in enterprise, basis for an industry or organizational  
4248 reference architecture, implication of architectural decisions, and positioning of vendor products in  
4249 a SOA context.
- 4250 • The Open Group SOA Governance Framework is a governance domain reference model and  
4251 method. It is for understanding SOA governance in organizations. The OASIS Reference  
4252 Architecture for SOA Foundation contains an abstract discussion of governance principles as  
4253 applied to SOA across boundaries
- 4254 • The Open Group SOA Integration Maturity Model (OSIMM) is a means to assess an  
4255 organization’s maturity within a broad SOA spectrum and define a roadmap for incremental  
4256 adoption. It is used for understanding the level of SOA maturity in an organization
- 4257 • The Object Management Group SoaML Specification supports services modeling UML  
4258 extensions. It can be seen as an instantiation of a subset of the Open Group RA used for  
4259 representing SOA artifacts in UML.

4260 Fortunately, there is a great deal of agreement on the foundational core concepts across the many  
4261 independent specifications and standards for SOA. This could be best explained by broad and common  
4262 experience of users of SOA and its maturity in the marketplace. It also provides assurance that investing  
4263 in SOA-based business and IT transformation initiatives that incorporate and use these specifications and  
4264 standards helps to mitigate risks that might compromise a successful SOA solution.



4265  
4266  
4267

Figure 45- SOA Reference Architecture Positioning (from "Navigating the SOA Open Standards Landscape Around Architecture, © OASIS, OMG, The Open Group).