



Reference Architecture Foundation for Service Oriented Architecture Version 1.0

Committee Specification Draft 03

06 July 2011

Specification URIs:

This version:

<http://docs.oasis-open.org/soa-rm/soa-ra/v1.0/csd03/soa-ra-v1.0-csd03.pdf> (Authoritative)
<http://docs.oasis-open.org/soa-rm/soa-ra/v1.0/csd03/soa-ra-v1.0-csd03.html>
<http://docs.oasis-open.org/soa-rm/soa-ra/v1.0/csd03/soa-ra-v1.0-csd03.doc>

Previous version:

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

Latest version:

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

Technical Committee:

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

Chair:

Ken Laskey (klaskey@mitre.org), MITRE Corporation

Editors:

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

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/IEEE 1471-2000 (now ISO/IEC 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.

Status:

This document was last revised or approved by the OASIS Service Oriented Architecture Reference Model TC on the above date. The level of approval is also listed above. Check the "Latest 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/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>).

Citation format:

When referencing this specification the following citation format should be used:

[SOA-RAF]

Reference Architecture Foundation for Service Oriented Architecture Version 1.0. 06 July 2011. OASIS Committee Specification Draft 03. <http://docs.oasis-open.org/soa-rm/soa-ra/v1.0/csd03/soa-ra-v1.0-csd03.html>.

Notices

Copyright © OASIS Open 2011. 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.

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?	10
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	11
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.....	13
1.4	SOA-RAF Viewpoints	13
1.4.1	Participation in a SOA Ecosystem Viewpoint.....	14
1.4.2	Realization of a SOA Ecosystem Viewpoint.....	14
1.4.3	Ownership in a SOA Ecosystem Viewpoint	14
1.5	Terminology	14
1.5.1	Usage of Terms.....	14
1.6	References.....	15
1.6.1	Normative References.....	15
1.6.2	Non-Normative References.....	15
2	Architectural Goals and Principles	17
2.1	Goals and Critical Success Factors of the Reference Architecture Foundation	17
2.1.1	Goals	18
2.1.2	Critical Success Factors.....	18
2.2	Principles of this Reference Architecture Foundation	19
3	Participation in a SOA Ecosystem View.....	21
3.1	Social Structure in a SOA Ecosystem Model	23
3.1.1	Participants, Actors and Delegates	24
3.1.2	Roles in Social Structures	26
3.1.3	Resource and Ownership.....	29
3.1.4	Trust and Risk	30
3.1.5	Policies and Contracts.....	32
3.1.6	Communication.....	33
3.1.7	Semantics and Semantic Engagement.....	33
3.2	Action in a SOA Ecosystem Model.....	34
3.2.1	Needs, Requirements and Capabilities.....	35
3.2.2	Services Reflecting Business	36
3.2.3	Action, Communication and Joint Action.....	36
3.2.4	State, Shared State and Real-World Effect.....	37
3.3	Architectural Implications.....	38
3.3.1	Social structures.....	38
3.3.2	Resource and Ownership.....	38
3.3.3	Policies and Contracts.....	39

3.3.4	Communications as a Means of Mediating Action	39
3.3.5	Semantics	39
3.3.6	Trust and Risk	39
3.3.7	Needs, Requirements and Capabilities	40
3.3.8	The Importance of Action	40
4	Realization of a SOA Ecosystem view	41
4.1	Service Description Model	41
4.1.1	The Model for Service Description	42
4.1.2	Use of Service Description	50
4.1.3	Relationship to Other Description Models	55
4.1.4	Architectural Implications	55
4.2	Service Visibility Model	57
4.2.1	Visibility to Business	57
4.2.2	Visibility	58
4.2.3	Architectural Implications	63
4.3	Interacting with Services Model	63
4.3.1	Interaction Dependencies	64
4.3.2	Actions and Events	64
4.3.3	Message Exchange	65
4.3.4	Composition of Services	68
4.3.5	Architectural Implications of Interacting with Services	72
4.4	Policies and Contracts Model	73
4.4.1	Policy and Contract Representation	73
4.4.2	Policy and Contract Enforcement	74
4.4.3	Architectural Implications	75
5	Ownership in a SOA Ecosystem View	76
5.1	Governance Model	76
5.1.1	Understanding Governance	76
5.1.2	A Generic Model for Governance	78
5.1.3	Governance Applied to SOA	82
5.1.4	Architectural Implications of SOA Governance	86
5.2	Security Model	86
5.2.1	Secure Interaction Concepts	87
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	92
5.3	Management Model	93
5.3.1	Management	93
5.3.2	Management Means and Relationships	96
5.3.3	Management and Governance	97
5.3.4	Management and Contracts	97
5.3.5	Management for Monitoring and Reporting	101
5.3.6	Management for Infrastructure	101
5.3.7	Architectural Implication of the Management Model	102

5.4 SOA Testing Model.....	102
5.4.1 Traditional Software Testing as Basis for SOA Testing	102
5.4.2 Testing and the SOA Ecosystem	103
5.4.3 Elements of SOA Testing	104
5.4.4 Testing SOA Services	108
5.4.5 Architectural Implications for SOA Testing.....	111
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
D.2 Critical Success Factors.....	118
D.3 Requirements	118
D.4 CFA Diagrams.....	118
E. Relationship to other SOA Open Standards	119

Table of Figures

Figure 1 - Critical Factors Analysis of the Reference Architecture	17
Figure 2 - Model elements described in the Participation in a SOA Ecosystem view	22
Figure 3 - Social Structure	23
Figure 4 - Actors, Participants and Delegates	25
Figure 5 - Role in Social Structures	26
Figure 6 - Participant Roles in a Service.....	28
Figure 7 - Resources.....	29
Figure 8 - Willingness and Trust	31
Figure 9 - Policies and Contracts	32
Figure 10 - Model Elements Described in the Realization of a SOA Ecosystem view	41
Figure 11 - General Description	43
Figure 12 - Representation of a Description	44
Figure 13 - Service Description.....	46
Figure 14 - Service Interface.....	47
Figure 15 - Service Functionality	48
Figure 16 - Model for Policies and Contracts as related to Service Participants	49
Figure 17 - Policies and Contracts, Metrics, and Compliance Records	50
Figure 18 - Relationship between Action and Components of Service Description Modelx.....	51
Figure 19 - Execution Context.....	54
Figure 20 - Interaction Description	54
Figure 21 - Visibility to Business	58
Figure 22 - Mediated Service Awareness	60
Figure 23 - Awareness in a SOA Ecosystem.....	61
Figure 24 - Business, Description and Willingness.....	62
Figure 25 - Service Reachability	62
Figure 26 - Interaction dependencies	64
Figure 27 - A "message" denotes either an action or an event	65
Figure 28 - Fundamental SOA message exchange patterns (MEPs).....	66
Figure 29 - Simple model of service composition	68
Figure 30 - Abstract example of orchestration of service-oriented business process	70
Figure 31 - Abstract example of choreography of service-oriented business collaboration	71
Figure 32 - Policies and Contracts	73
Figure 33 - Model Elements Described in the Ownership in a SOA Ecosystem View	76
Figure 34 - Motivating Governance.....	78
Figure 35 - Setting Up Governance	79
Figure 36 - Carrying Out Governance.....	80
Figure 37 - Ensuring Governance Compliance.....	81
Figure 38 - Relationship Among Types of Governance.....	83
Figure 39 - Authentication	88
Figure 40 - Authorization.....	88

Figure 41 - Manageability capabilities in SOA ecosystem.....	94
Figure 42 - Management Means and Relationships in SOA ecosystem	96
Figure 43 - Management of the service interaction	99
Figure 44 - Example UML class diagram—Resources.....	117
Figure 45 - 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 area 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.¹

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

33 1.1.2 What is this Reference Architecture?

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

41 While requirements are addressed more fully in Section 0, the SOA-RAF makes key assumptions that
42 SOA-based systems involve:

- 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 several other industry reference architectures for SOA and related SOA open
81 standards is 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 89 Perspective

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

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

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

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

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

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

113 1.3 Viewpoints, Views and Models

114 1.3.1 ANSI/IEEE 1471-2000::ISO/IEC 42010-2007

- 115 1. The SOA-RAF uses and follows the IEEE "Recommended Practice for Architectural Description
116 of Software-Intensive Systems" [ANSI/IEEE 1471] and [ISO/IEC 42010]. An architectural
117 description conforming to this standard must include the following six (6) elements:
- 118 2. Architectural description identification, version, and overview information

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

127 The standard defines the following terms²:

128 **Architecture**

129 The fundamental organization of a system embodied in its components, their relationships to

130 each other, and to the environment, and the principles guiding its design and evolution.

131 **Architectural Description**

132 A collection of products that document the architecture.

133 **System**

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

135 **System Stakeholder**

136 A system stakeholder is an individual, team, or organization (or classes thereof) with interests in,

137 or concerns relative to, a system.

138 A stakeholder's concern should not be confused with either a need or a formal requirement. A concern,

139 as understood here, is an area or topic of interest. Within that concern, system stakeholders may have

140 many different requirements. In other words, something that is of interest or importance is not the same

141 as something that is obligatory or of necessity [TOGAF v9].

142 When describing architectures, it is important to identify stakeholder concerns and associate them with

143 viewpoints to insure that those concerns are addressed in some manner by the models that comprise the

144 views on the architecture. The standard defines views and viewpoints as follows:

145 **View**

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

147 **Viewpoint**

148 A specification of the conventions for constructing and using a view. A pattern or template from

149 which to develop individual views by establishing the purposes and audience for a view and the

150 techniques for its creation and analysis.

151 In other words, a view is what the stakeholders see whereas the viewpoint defines the perspective from

152 which the view is taken and the methods for, and constraints upon, modeling that view.

153 It is important to note that viewpoints are independent of a particular system (or solutions). In this way,

154 the architect can select a set of candidate viewpoints first, or create new viewpoints, and then use those

155 viewpoints to construct specific views that will be used to organize the architectural description. A view,

156 on the other hand, is specific to a particular system. Therefore, the practice of creating an architectural

² See <http://www.iso-architecture.org/ieee-1471/conceptual-framework.html> for a diagram of the standard's Conceptual Framework

157 description involves first selecting the viewpoints and then using those viewpoints to construct specific
 158 views for a particular system or subsystem. Note that the standard requires that each view corresponds to
 159 exactly one viewpoint. This helps maintain consistency among architectural views which is a normative
 160 requirement of the standard.

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

162 **Model**

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

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

167 **1.3.2 UML Modeling Notation**

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

174 Appendix The Unified Modeling Language, UML introduces the UML notation that is used in this
 175 document.

176 **1.4 SOA-RAF Viewpoints**

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

Viewpoint Element	Viewpoint		
	Participation in a SOA Ecosystem	Realization of a SOA Ecosystem	Ownership in a SOA Ecosystem
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

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

181 **1.4.1 Participation in a SOA Ecosystem Viewpoint**

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

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

190 Given SOA's value in allowing people to access, manage and provide services across [ownership](#)
191 [boundaries](#), we must explicitly identify those boundaries and the implications of crossing them.

192 **1.4.2 Realization of a SOA Ecosystem Viewpoint**

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

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

199 **1.4.3 Ownership in a SOA Ecosystem Viewpoint**

200 This viewpoint addresses the concerns involved in owning and managing SOA-based systems within the
201 SOA ecosystem. Many of these concerns are not easily addressed by automation; instead, they often
202 involve people-oriented processes such as governance bodies.

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

208 **1.5 Terminology**

209 The keywords "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD
210 NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described
211 in **[RFC2119]**.

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

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

216 **1.5.1 Usage of Terms**

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

222

1.6 References

223

1.6.1 Normative References

- 224 [ANSI/IEEE 1471] *IEEE Recommended Practice for Architectural Description of Software-Intensive*
225 *Systems*, American National Standards Institute/Institute for Electrical and
226 Electronics Engineers, September 21, 2000.
- 227 [ISO/IEC 42010] International Organization for Standardization and International Electrotechnical
228 Commission, *System and software engineering — Recommended practice for*
229 *architectural description of software-intensive systems*, July 15, 2007.
- 230 [RFC2119] S. Bradner, *Key words for use in RFCs to Indicate Requirement Levels*,
231 <http://www.ietf.org/rfc/rfc2119.txt>, IETF RFC 2119, March 1997.
- 232 [SOA-RM] OASIS Standard, "Reference Model for Service Oriented Architecture 1.0, 12
233 October 2006. <http://docs.oasis-open.org/soa-rm/v1.0/soa-rm.pdf>
- 234 [UML 2] *Unified Modeling Language: Superstructure*, Ver. 2.1.1, OMG Adopted
235 Specification, OMG document formal/2007-02-05, Object Management Group,
236 Needham, MA, February 5, 2007.
- 237 [WA] Architecture of the World Wide Web, W3C, 2004. <http://www.w3.org/TR/webarch>.
- 238 [WSA] David Booth, et al., "Web Services Architecture", W3C Working Group Note,
239 World Wide Web Consortium (W3C) (Massachusetts Institute of Technology,
240 European Research Consortium for Informatics and Mathematics, Keio
241 University), February, 2004. <http://www.w3.org/TR/2004/NOTE-ws-arch-20040211/>
- 242

243

1.6.2 Non-Normative References

- 244 [BLOOMBERG/SCHMELZER] Jason Bloomberg and Ronald Schmelzer, *Service Orient or Be*
245 *Doomed!*, John Wiley & Sons: Hoboken, NJ, 2006.
- 246 [COX] D. E. Cox and H. Kreger, "Management of the service-oriented architecture life
247 cycle," "IBM Systems Journal" "44", No. 4, 709-726, 2005
- 248 [DCMI] Dublin Core reference
- 249 [ERA] A. Fattah, "Enterprise Reference Architecture," paper presented at 22nd
250 Enterprise Architecture Practitioners Conference, London, UK, April 2009.
- 251 [IEEE-829] *IEEE Standard for Software Test Documentation*, Institute for Electrical and
252 Electronics Engineers, 16 September 1998
- 253 [ISO 11179] ISO 11179 reference
- 254 [ITU-T Rec. X.700 | ISO/IEC 10746-3:1996(E)]
255 Information processing systems—Open Systems Interconnection—Basic
256 Reference Model—Part 4: Management Framework", International
257 Telecommunication Union, International Organization for Standardization and
258 International Electrotechnical Commission, Geneva, Switzerland, 1989.
- 259 [NEWCOMER/LOMOW]
260 Eric Newcomer and Greg Lomow, *Understanding SOA with Web Services*,
261 Addison-Wesley: Upper Saddle River, NJ, 2005.
- 262 [OECD] Organization for Economic Cooperation and Development, Directorate for
263 Financial, Fiscal and Enterprise Affairs, OECD Principles of Corporate
264 Governance, SG/CG(99) 5 and 219, April 1999.
- 265 [TOGAF v9] The Open Group Architecture Framework (TOGAF) Version 9 Enterprise Edition,
266 The Open Group, Doc Number: G091, February 2009.
- 267 [WEILL] Harvard Business School Press, IT Governance: How Top Performers Manage
268 IT Decision Rights for Superior Results, Peter Weill and Jeanne W. Ross, 2004
- 269 [DAMIANOU] Nicodemos C. Damianou, Thesis - A Policy Framework for Management of
270 Distributed Systems, University of London, Department of Computing, 2002.

271 **[LEVESON]** Nancy G. Leveson, *Safeware: System Safety and Computers*, Addison-Wesley
272 Professional, Addison-Wesley Publishing Company, Inc.: Boston, pg. 181, 1995.
273 **[STEEL/NAGAPPAN/LAI]**
274 Christopher Steel and Ramesh Nagappan and Ray Lai, *core Security*
275 *Patterns: Best Practices and Strategies for J2EE, Web Services and Identity*
276 *Management*, Prentice Hall: 2005
277 **[ISO/IEC 27002]** International Organization for Standardization and International Electrotechnical
278 Commission, *Information technology -- Security techniques – Code of practice*
279 *for information security management*, 2007
280 **[SOA NAV]** Heather Kreger and Jeff Estefan (Eds.), "Navigating the SOA Open Standards
281 Landscape Around Architecture," Joint Paper, The Open Group, OASIS, and
282 OMG, July 2009. [http://www.oasis-](http://www.oasis-open.org/committees/download.php/32911/wp_soa_harmonize_d1.pdf)
283 [open.org/committees/download.php/32911/wp_soa_harmonize_d1.pdf](http://www.oasis-open.org/committees/download.php/32911/wp_soa_harmonize_d1.pdf)

284 **2 Architectural Goals and Principles**

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

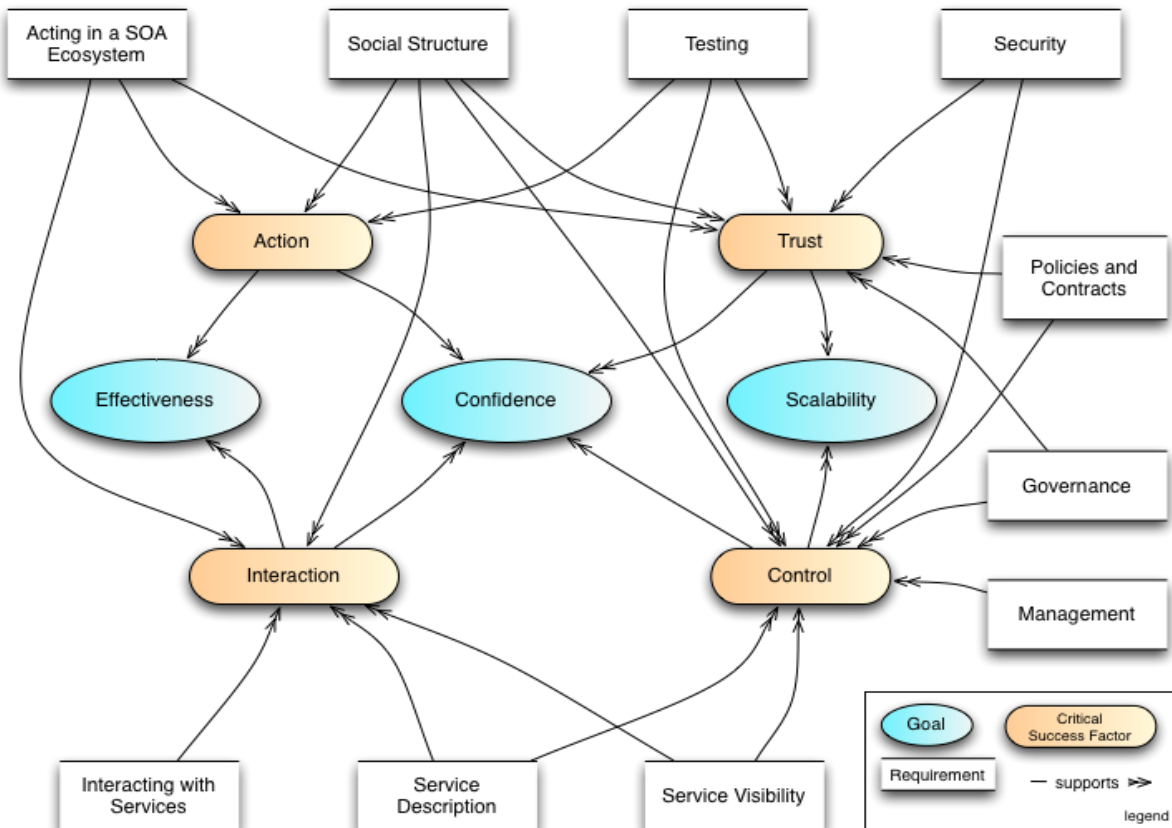
287 **2.1 Goals and Critical Success Factors of the Reference**
288 **Architecture Foundation**

289 There are three principal goals:

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

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

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



302
303 *Figure 1 - Critical Factors Analysis of the Reference Architecture*

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

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

311 **2.1.1 Goals**

312 **2.1.1.1 Effectiveness**

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

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

319 **2.1.1.2 Confidence**

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

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

326 **2.1.1.3 Scalability**

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

331 **2.1.2 Critical Success Factors**

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

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

338 **2.1.2.1 Action**

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

342 **2.1.2.2 Trust**

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

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

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

354 **2.1.2.3 Interaction**

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

358 **2.1.2.4 Control**

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

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

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

370 **2.2 Principles of this Reference Architecture Foundation**

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

372 **Technology Neutrality**

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

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

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

387 **Parsimony**

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

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

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

394 **Distinction of Concerns**

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

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

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

414 **Applicability**

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

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

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

3 Participation in a SOA Ecosystem View

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

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

Together, these ideas describe an environment in which business functions (realized in the form of services) address business needs. Service implementations utilize capabilities to produce specific (real world) effects that fulfill those business needs. Both the people³ using the services, and the capabilities themselves, may be distributed across ownership domains, with different policies and conditions of use in force– this environment is referred to as a **SOA Ecosystem**.

The role of a service in the SOA context is to enable effective business solutions in this environment. Any technology system created to deliver a service in such an environment is referred to as a **SOA-based system**. SOA is thus a paradigm that guides the identification, design, implementation (i.e., organization), and utilization of such services.

A SOA-based system is concerned with how actors in a system interact to deliver a specific result - the delivery of a capability or real-world effect. The SOA ecosystem is concerned with all potential stakeholders and the roles that they can play; how some stakeholders’ needs are satisfied by other stakeholders’ solutions; how stakeholders assess risk; how they relate to each other through policies and contracts; and how they communicate and establish relationships of trust in the processes leading to the delivery of a specific result.

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

³ ‘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 The subsections below expand on the completely abstract reference model by identifying more fully and
466 with more specificity what challenges need to be addressed in order to successfully apply the SOA
467 paradigm. Although this section does not provide a specific recipe, it does identify the important things
468 that need to be thought about and resolved within an ecosystem context.

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

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

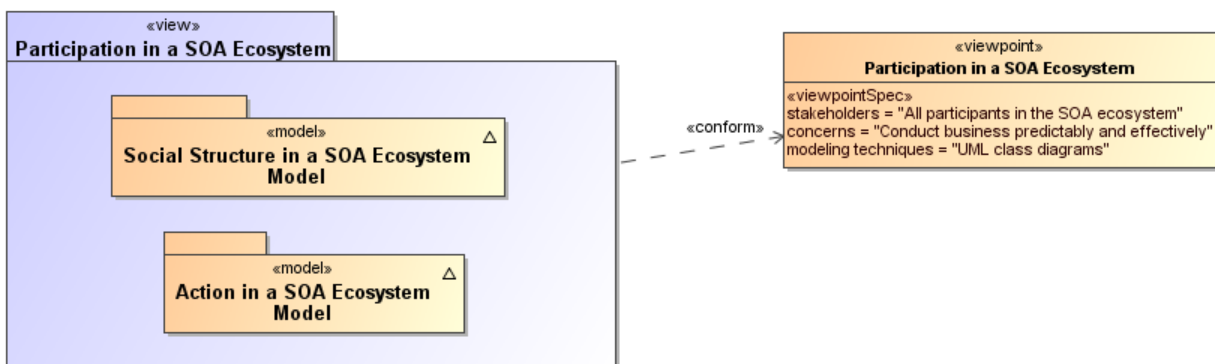
480 This view describes how a SOA-based system behaves when participants may be in different
481 organizations, with different rules and expectations, and assumes that the primary motivation for
482 participants to interact with each other is to achieve **objectives** –to get things done.

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

489 Several interdependent concerns are important in our view of a SOA-ecosystem. The ecosystem includes
490 stakeholders who are participants in the development, deployment and governance and use of a system
491 and its services; or who may not participate but are nonetheless affected by the system. **Actors** –
492 whether stakeholder **participants** or delegates who act only on behalf of participants (without themselves
493 having any stake in the actions that they have been tasked to perform) – are engaged in **actions** which
494 have an impact on the real world and whose meaning and intent are determined by implied or agreed-to
495 semantics.

496 The main models in this view are:

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

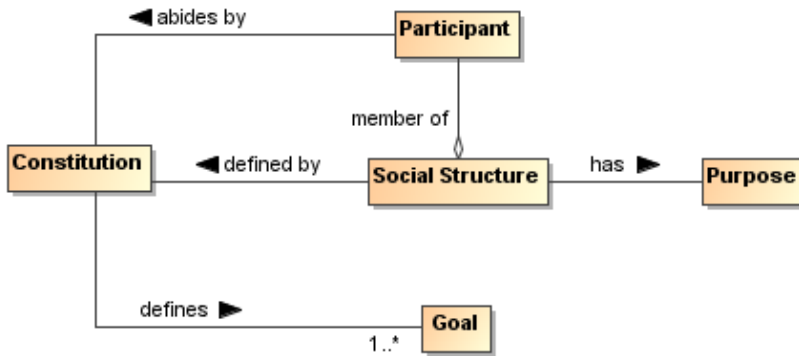


503
504 Figure 2 - Model elements described in the Participation in a SOA Ecosystem view

505 **3.1 Social Structure in a SOA Ecosystem Model**

506 The actions undertaken by **participants** in a SOA ecosystem are performed in a *social context* that defines
507 the relationships between the **participants**. That context is the **social structure**. In order to achieve
508 success in applying the SOA paradigm, the overall social structure in which the SOA effort is to be
509 undertaken must be taken into consideration. Ownership boundaries and their implications can only be
510 understood and addressed within the context of the larger social structure within which they exist and the
511 nature of the relationships between the different participants in that structure.

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



516
517 *Figure 3 - Social Structure*

518 **Social Structure**

519 A **social structure**⁴ is a nexus of relationships amongst **participants** brought together for a specific
520 **purpose**.

521 A social structure represents a collection of **participants** and is established with an implied or explicitly
522 defined purpose. The purpose is usually reflected in specific goals laid down in the social structure's
523 **constitution** or other 'charter'.

524 A social structure may have any number of participants and a large number of different relationships may
525 exist among participants. The organizing principle for these relationships is the social structure's
526 purpose. In addition, a given participant can be a member of multiple social structures. Thus, there may
527 be interaction among social structures, sometimes resulting in disagreements when the premises of the
528 social structures do not align.

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

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

⁴ Social structures are sometimes referred to as social institutions.

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

538 **Enterprise**

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

542 **Peer Group**

543 A **peer group** is a social structure with no discernable leadership structure, that may or may not
544 have internally established goals, but is identifiable as the locus of interaction between participants
545 with individual goals seeking common outcomes and who are considered peers of one another.

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

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

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

559 **Constitution**

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

562 Every social structure functions according to rules by which **participants** interact with each other within the
563 structure. In some cases, this is based on an explicit agreement, in other cases participants behave as
564 though they agree to the **constitution** without a formal agreement. In still other cases, participants abide
565 by the rules with some degree of reluctance, such as governance of SOA-based systems, covered below.
566 In all cases, the constitution may change over time, in those cases of implicit agreement the change can
567 occur quickly.

568 **3.1.1 Participants, Actors and Delegates**


569 **Social structures** have **Stakeholders** – people – some of whom may be enterprises. They interact within
570 the broad SOA ecosystem. **Actors** on the other hand operate strictly within a SOA-based system.

571 There is also the concept of **Participant** which is particularly important as it reflects the hybrid role of a
572 person who is both a Stakeholder in the ecosystem (and thus primarily concerned with expressing needs
573 and seeing those needs fulfilled) *and* an Actor in the System (and thus directly involved with system-level
574 activity).

575 A stakeholder can be either a participant (and thus also an actor with a specific functional role in a SOA-
576 based system); or a non-participant – someone who, without participating, nonetheless has something at
577 stake within the ecosystem.

578 An actor can be either a **participant** (and thus also a stakeholder with a stake in the ecosystem); or a
579 **delegate** – a human actor with no stake in the specific action delegated or some automated agent –
580 acting on behalf of a participant.

581 The hybrid role of Participant provides a bridge between the wider (real-world) ecosystem – the world of
582 the stakeholder – and the more specific (usually technology-focused) system – the world of the actor.



583

584 *Figure 4 - Actors, Participants and Delegates*

585 **Stakeholder**

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

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

593 **Actor**

594 An actor is a human or non-human agent capable of [action within a SOA-based system](#).

595 **Participant**

596 A participant is a person⁵ who is both a [stakeholder](#) in the SOA ecosystem and an [actor](#) in the
 597 SOA-based system.

598 **Delegate**

599 A delegate is an [actor](#) that is acting on behalf of a [participant](#).

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

601 Many stakeholders and actors operate in a SOA ecosystem, including software agents that permit people
 602 to offer, and interact with, services; [delegates](#) that represent the interests of other participants; or security
 603 agents charged with managing the security of the ecosystem. Note that automated agents are always
 604 delegates, in that they act on behalf of a stakeholder.

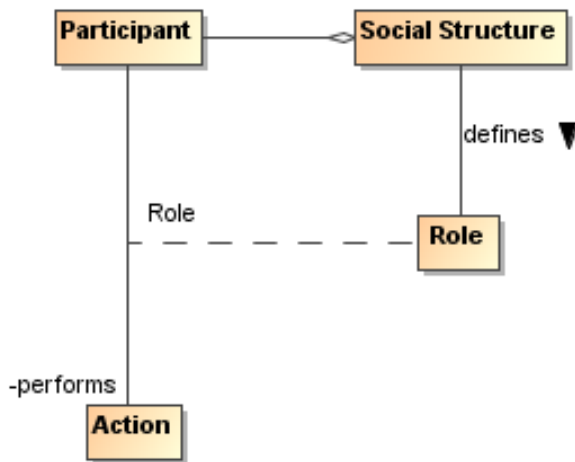
605 In the different models of the RAF, [actor](#) is used when it is not important whether the entity is a [delegate](#)
 606 or a [participant](#). If the [actor](#) is acting on behalf of a stakeholder, then we use [delegate](#). This underlines the
 607 importance of delegation in SOA-based systems, whether the delegation is of work procedures carried
 608 out by human agents who have no stake in the actions with which they are tasked but act on behalf of a
 609 participant who does; or whether the delegation is performed by technology (automation). If the [actor](#) is
 610 also a [stakeholder](#) in the ecosystem, then we use [participant](#).

611 In order for a delegate to act on behalf of another person, they must be able to act and have the [authority](#)
 612 to do so.

⁵ Again, this can be a ‘natural’ or ‘legal’ person

613 **3.1.2 Roles in Social Structures**

614 **Social structures** are abstractions: a **social structure** cannot directly perform **actions** – only people or
615 automated processes following the instructions of people can actually do things. However, an **actor** may
616 act on behalf of a **social structure** and certainly acts within a **social structure** depending on the **roles** that
617 the **actor** assumes and the nature of the relationships between the concerned parties or stakeholders.
618



619
620 *Figure 5 - Role in Social Structures*

621 **Role**

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

624 A role is not immutable and is often time-bound. A **participant** can have one or more roles concurrently
625 and may change them over time and in different contexts, even over the course of a particular interaction.

626
627 One participant with appropriate authority in the **social structure** may formally *designate a role* for another
628 **participant**, with associated **rights** and **responsibilities**, and that authority may even qualify a period during
629 which the designated role may be valid. In addition, while many **roles** are clearly identified, with
630 appropriate names and definitions of **responsibilities**, it is also possible to separately bestow **rights**,
631 bestow or assume **responsibilities** and so on, often in a temporary fashion. For example, when a
632 company president delegates certain responsibilities on another person, this does not imply that the other
633 person has become company president. Likewise, a company president may bestow on someone else
634 her role during a period of time that she is on vacation or otherwise unreachable, with the understanding
635 that she will re-assume the role when she returns from vacation.

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

640 **Authority**

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

642 **Right**

643 A right is a predetermined **permission** conferred upon an **actor** to perform some **action** or assume
644 a **role** in relation to the **social structure**.

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

647 **Responsibility**

648 A responsibility is a predetermined **obligation** on a participant to perform some **action** or assume
649 a role in relation to other participants.

650 Responsibility implies human agency, which is why only participants, as opposed to all actors (who can
651 be non-human agents) are concerned. This applies even if the consequences of such responsibility can
652 impact other (human and non-human) actors. Having authority often implies having responsibility.

653 **Rights, authorities, responsibilities** and **roles** form the foundation for the security model as well as
654 contributing to the governance model in the 'Ownership in a SOA Ecosystem' View of the RAF.

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

663 Two important types of constraints that are relevant to a SOA ecosystem are Permission and Obligation.

664 **Permission**

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

667 Note that **permissions** are distinct from ability and from authority. Authority refers to the legitimate nature
668 of an **action** as performed by an **actor** on behalf of a **social structure**. Ability refers to whether an actor has
669 the capacity to perform the action. **Permission** does not always involve acting on behalf of anyone, nor
670 does it imply or require the capacity to perform the action.

671 **Obligation**

672 An obligation is a constraint that prescribes the **actions** that an **actor** must (or must not) perform
673 and/or the **states** the **actor** must (or must not) attain or maintain.

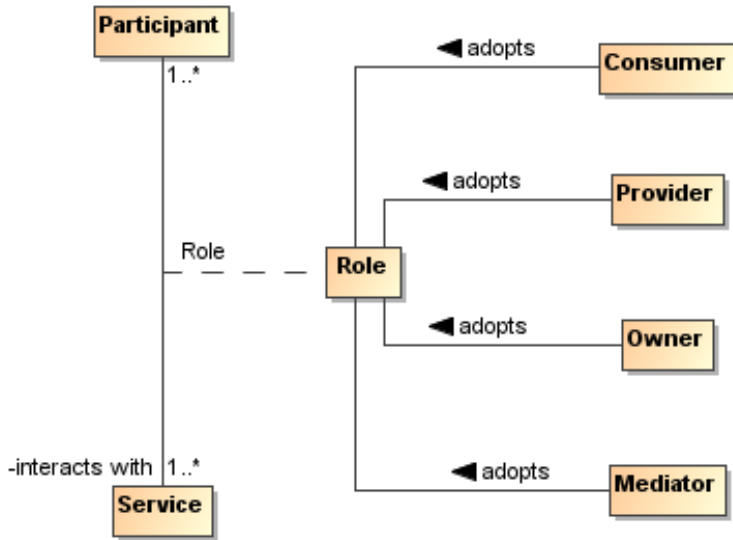
674 An example of obligations is the case where the **service consumer** and provider have entered into an
675 agreement to provide and consume a service such that the consumer is obligated to pay for the service
676 and the provider is obligated to provide the service – based on the terms of the contract.

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

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

683 **3.1.2.1 Service Roles**

684 As in roles generically, a participant can play one or more of those roles inherent to the SOA paradigm in
685 the SOA ecosystem, depending on the context. A participant may be playing a role of a service provider
686 in one relationship while simultaneously playing the role of a consumer in another. Roles inherent to the
687 SOA paradigm include Consumer, Provider, Owner, and Mediator.



688
689 *Figure 6 - Participant Roles in a Service*

690 **Provider**

691 A provider is a role assumed by a [participant](#) who is offering a service.

692 **Consumer**

693 A consumer is a role assumed by a [participant](#) who is interacting with a service in order to fulfill a
694 need.

695 **Mediator**

696 A mediator is a role assumed by a [participant](#) to facilitate interaction and connectivity in the
697 offering and use of services.

698 **Owner**

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

701 It is a common understanding that service interactions are typically initiated by service consumers,
702 although this is not necessarily true in all situations. Additionally, as with service providers, several
703 [stakeholders](#) may be involved in a service interaction supporting a given consumer.

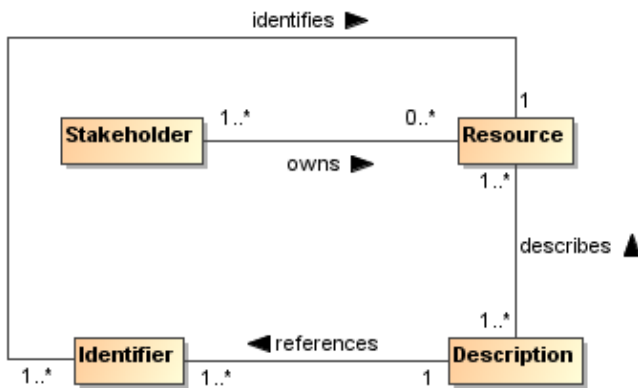
704 The roles of service provider and service consumer are often seen as symmetrical, which is also not
705 entirely correct. A consumer tends to express a 'Need' in non-formal terms: "I want to buy that book". The
706 type of 'Need' that a service is intended to fulfill has to be formalized and encapsulated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

721 3.1.3 Resource and Ownership

722 3.1.3.1 Resource

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

725



726

727 Figure 7 - Resources

728 Resource

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

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

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

736 Identifier

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

739 Identifiers typically require a context in order to establish the connection with the resource. In a SOA
740 ecosystem, it is good practice to use globally unique identifiers; for example globally unique
741 Internationalized Resource Identifiers (IRIs).

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

743 The ability to identify a resource is important in interactions to determine such things as rights and
744 authorizations, to understand what functions are being performed and what the results mean, and to
745 ensure repeatability or characterize differences with future interactions. Many interactions within a SOA
746 ecosystem take place across ownership boundaries and the combination of interactions can be
747 unpredictable. Identifiers provide the means for all resources important to a given SOA system to be
748 unambiguously identifiable at any moment and in any interaction.

749 3.1.3.2 Ownership

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

752 Typically, the ownership relationship is one of control: the owner of a resource can control some aspect
753 of the resource.

754 Ownership

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

758 To own a **resource** implies taking responsibility for creating, maintaining and, if it is to be available to
759 others, provisioning the **resource**. More than one **stakeholder** may own different **rights** or responsibilities
760 associated with a given service, such as one **stakeholder** having the **responsibility** to deploy a capability
761 as a service, another owning the **rights** to the profits that result from charging consumers for using the
762 service, and yet another owning the **right** to use the service. . There may also be joint **ownership** of a
763 **resource**, where the rights and responsibilities are shared.

764 A stakeholder who owns a **resource** may delegate some or all of these **rights** and responsibilities to
765 others, but typically retains the responsibility to see that the delegated **rights** and responsibilities are
766 exercised as intended

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

772 **Ownership** is defined in relation to the **social structure** relative to which the given **rights** and
773 **responsibilities** are exercised. For example, there may be constraints on how **ownership** may be
774 transferred, such as a government may not permit a corporation to transfer assets to a subsidiary in a
775 different jurisdiction.

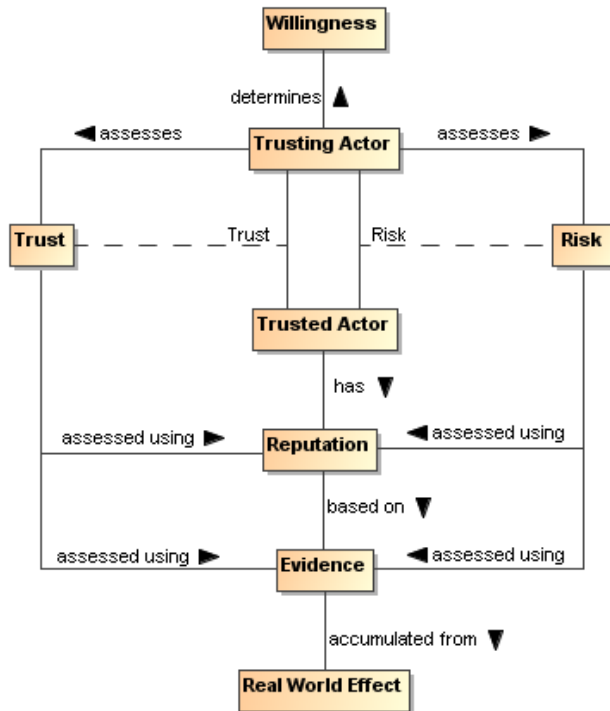
776 **Ownership Boundary**

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

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

786 **3.1.4 Trust and Risk**

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



788

789 *Figure 8 - Willingness and Trust*

790 **Willingness**

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

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

796 **Trust**

797 Trust is a private assessment or internal perception of one actor that another actor will perform
 798 actions in accordance with an assertion regarding a desired **real world effect**.

799 **Risk**

800 Risk is a private assessment or internal perception of the likelihood that certain undesirable **real**
 801 **world effects** will result from actions taken and the consequences or implications of such.

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

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

810 **3.1.4.1 Assessing Trust and Risk**

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

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

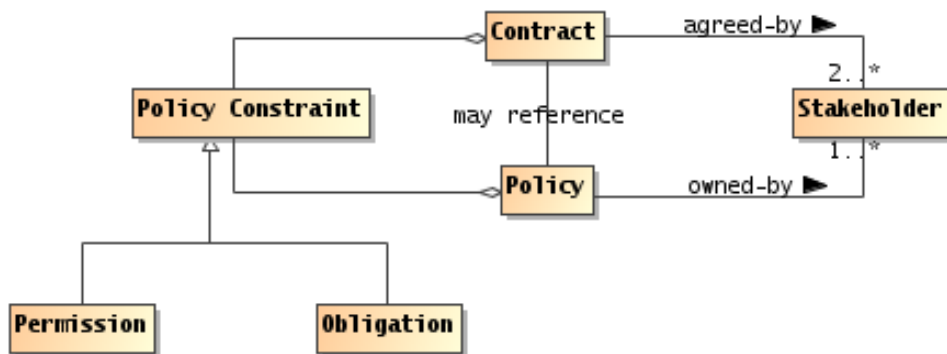
817 Assessment of trust is rarely binary. An **actor** is not completely trusted or untrusted because there is
818 typically some degree of uncertainty in the accuracy or completeness of the evidence or the assessment.
819 Similarly, there may be uncertainty in the amount and potential consequences of risk.

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

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

832 3.1.5 Policies and Contracts

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



838
839 *Figure 9 - Policies and Contracts*

840 Policy

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

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

848 Contract

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

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

854 policy and a provider's policy are mutually exclusive, then some form of negotiation (involving human
855 interactions) or mediation must resolve the mutual exclusion before the service consumer/provider
856 interaction can occur. Note, this also applies if the policy is introduced by the consumer instead of the
857 provider.

858 Both [policies](#) and [contracts](#) imply a desire to see constraints respected and enforced. [Policies](#) are owned
859 by service providers – individual (or aggregate) [stakeholders](#) – and contracts are owned by both service
860 providers and consumers – the parties to the contract; these [stakeholders](#) are responsible for ensuring
861 that any constraints in the [policy](#) or contract are enforced, although the actual enforcement may be
862 delegated to a different mechanism. A contract does not necessarily oblige the contracting parties to act
863 (for example to use a service) but it does constrain how they act if and when the condition covered by the
864 contract occurs (for example, when a service is invoked and used).

865 **3.1.6 Communication**

866 **Communication**

867 A communication is a process of reaching mutual understanding, in which participants not only
868 exchange information as messages but share the meaning of this information.

869 A [communication](#) involves at least one actor in the role of **sender** and at least one other actor in the role
870 of **recipient**. All actors must perform their role in order for the communication to occur.

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

876 A communication involves a message, which an [actor](#) receiving must be able to correctly interpret. The
877 extent of that correct interpretation depends on the [role](#) of the [actor](#) and the purpose of the
878 communication.

879 A communication is not effective unless the recipient can correctly interpret the message (or at least, that
880 part of it which is relevant to the participant). However, interpretation can itself be characterized in terms
881 of [semantic engagement](#): the proper understanding of a message in a given context.

882 We can characterize the necessary modes of interpretation in terms of a shared understanding of a
883 common [vocabulary](#) (or [mediation among vocabularies](#)) and of the purpose of the communication. More
884 formally, we can say that a communication has a combination of message and purpose.

885 Interactions between [service consumers](#) and providers do not need to resemble human speech. Machine-
886 machine communication is typically highly stylized in form, it may have particular forms and it may involve
887 particular terms not found in everyday human communication.

888 **3.1.7 Semantics and Semantic Engagement**

889 A SOA ecosystem is a space in which [actors](#) need to share understanding⁶ as well as sharing actions.
890 Indeed, such shared understanding is a pre-requisite to a joint action being carried out as intended. It is
891 vital to a trusted and effective ecosystem. Semantics are therefore pervasive throughout SOA

⁶ 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.

892 ecosystems and important in communicative actions described above, as well as a driver for [policies](#) and
893 other aspects of the ecosystem.

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

898 Understanding and interpreting those assertions allows other actors to know what may be expected of
899 them in any particular joint action. An actor can potentially ‘understand’ an assertion in a number of ways,
900 but it is specifically the process of arriving at a *shared* understanding that is important in the ecosystem.
901 This process is semantic engagement among the actors in the SOA ecosystem. It can be instantaneous
902 or progressively achieved. It is important that there is a level of engagement appropriate to the particular
903 context.

904 **Semantic Engagement**

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

907 Different [actors](#) have differing capabilities and requirements for understanding [assertions](#). This is true for
908 both human and non-human [actors](#). For example, a purchase order process does not require that a
909 message forwarding agent ‘understand’ the purchase order, but a processing agent does need to
910 ‘understand’ the purchase order in order to know what to do with the order once received.

911 The impact of any [assertion](#) can only be fully understood in terms of specific social contexts that
912 necessarily include the [actors](#) that are involved. For example, a [policy](#) statement that governs the actions
913 relating to a particular resource may have a different impact or purpose for the participant that owns the
914 resource than for the actor that is trying to access it: the former understands the purpose of the [policy](#) as
915 a statement of enforcement - the latter understands it as a statement of constraint.

916 **3.2 Action in a SOA Ecosystem Model**

917 Participants cannot always achieve desired results by leveraging resources in their own ownership
918 domain. This unfulfilled need leads them to seek and leverage services provided by other participants and
919 using resources beyond their ownership and control. The participants identify service providers with which
920 they think they can interact to achieve their objective and engage in joint action with those other actors
921 (service providers) in order to bring about the desired outcome. The SOA ecosystem provides the
922 environment in which this happens.

923 An action model is put forth a-priori by the service provider, and is effectively an undertaking by the
924 service provider that the actions – identified in the action model and invoked consistent with the process
925 model – will result in the described real world effect. The action model describes the actions leading to a
926 real-world effect. A potential service consumer – who is interested in a particular outcome to satisfy their
927 need – must understand those actions as capable of achieving that desired outcome.

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

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

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

943 All of these things are important to getting to the core of participants' concern in a SOA ecosystem: the
944 ability to leverage resources or capabilities to achieve a desired outcome, and in particular where those
945 resources or capabilities do not belong to them or are beyond their direct control. i.e., that are outside of
946 their ownership boundary.

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

951 In order to act in a way that is appropriate and consistent, participants must communicate with each other
952 about their own goals, objectives and policies, and those of others. This is the main concern of Semantic
953 Engagement.

954 A key aspect of joint action revolves around the trust that both parties must exhibit in order to participate
955 in the joint action. The willingness to act and a mutual understanding of both the information exchanged
956 and the expected results is the particular focus of Sections 3.1.4 and 3.1.7.

957 **3.2.1 Needs, Requirements and Capabilities**

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

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

969 Business functionality

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

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

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

984 **Need**

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

988 **Requirement**

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

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

994 **Capability**

995 A capability is an ability to deliver a real world effect.

996 The Reference Model makes a distinction between a capability (as a *potential* to deliver the real world
997 effect) and the ability of bringing that capability to bear (via a realized service) as the realization of the
998 real world effect.

999 **3.2.2 Services Reflecting Business**

1000 The SOA paradigm often emphasizes the interface through which service interaction is accomplished.
1001 While this enables predictable integration in the sense of traditional software development, the prescribed
1002 interface alone does not guarantee that services will be composable into business solutions.

1003 **Business solution**

1004 A **business solution** is a set of defined interactions that combine implemented or notional
1005 business functionality in order to address a set of business needs.

1006 **Composability**

1007 **Composability** is the ability to combine individual services, each providing defined business
1008 functionality, so as to provide more complex business solutions.

1009 To achieve composability, capabilities must be identified that serve as building blocks for business
1010 solutions. In a SOA ecosystem, these building blocks are captured as services representing well-defined
1011 business functions, operating under well-defined policies and other constraints, and generating well-
1012 defined real world effects. These service building blocks should be relatively stable so as not to force
1013 repeated changes in the compositions that utilize them, but should also embody SOA attributes that
1014 readily support creating compositions that can be varied to reflect changing circumstances.

1015 The SOA paradigm emphasizes both composition of services and opacity of how a given service is
1016 implemented. With respect to opacity, the SOA-RM states that the service could carry out its described
1017 functionality through one or more automated and/or manual processes that in turn could invoke other
1018 available services.

1019 Any composition can itself be made available as a service and the details of the business functionality,
1020 conditions of use, and effects are among the information documented in its service description.

1021 Composability is important because many of the benefits of a SOA approach assume multiple uses for
1022 services, and multiple use requires that the service deliver a business function that is reusable in multiple
1023 business solutions. Simply providing a Web Service interface for an existing IT artifact does not, in
1024 general, create opportunities for sharing business functions. Furthermore, the use of tools to auto-
1025 generate service software interfaces will not guarantee services that can effectively be used within
1026 compositions if the underlying code represents programming constructs rather than business functions. In
1027 such cases, services that directly expose the software details will be as brittle to change as the underlying
1028 code and will not exhibit the characteristic of loose coupling.

1029 **3.2.3 Action, Communication and Joint Action**

1030 In general terms, entities act in order to achieve their goals. However, the form of [action](#) that is of most
1031 interest within a SOA ecosystem is that involving interaction across ownership boundaries (between more
1032 than one [actor](#)) – **joint action**.

1033 **3.2.3.1 Action and Actors**

1034 **Action**

1035 An [action](#) is the application of intent to cause an effect.

1036 The aspect of action that distinguishes it from mere force or accident is that someone *intends* that the
1037 action achieves a desired objective or effect. This definition of action is very general. In the case of SOA,
1038 we are mostly concerned with actions that take place within a system and have specific effects on the
1039 SOA ecosystem – what we call **Real World Effects**. The actual real world effect of an action, however,
1040 may go beyond the intended effect.

1041 Objectives refer to **real world effects** that participants believe are achievable by a specific **action** or set of
1042 **actions** that deliver appropriate changes in shared state. In contrast, a **goal** is not expressed in terms of
1043 specific **action** but rather in terms of desired end state.

1044 For example, someone may wish to have enough light to read a book. In order to satisfy that goal, the
1045 reader walks over to flip a light switch. The *objective* is to change the state of the light bulb, by turning on
1046 the lamp, whereas the *goal* is to be able to read. The *real world effect* is more light being available to
1047 enable the person to read.

1048 While an effect is any measurable change resulting from an action, a SOA ecosystem is concerned more
1049 specifically with real world effects.

1050 **Real World Effect**

1051 A real world effect is a measurable change to the **shared state of pertinent entities, relevant to**
1052 **and experienced by specific stakeholders of an ecosystem**.

1053 This implies measurable change in the overall state of the SOA ecosystem. In practice, however, it is
1054 specific state changes of certain entities that are relevant to particular participants that constitute the real
1055 world effect as experienced by those participants.

1056 **3.2.3.2 Communication and Joint Actions**

1057 In this Reference Architecture Foundation, we are concerned with two levels of activity: as communication
1058 and as participants engaged in joint actions to use and offer services.

1059 In order for multiple **actors** to participate in a **joint action**, they must each act according to their **role** within
1060 the joint action. This is achieved through communication and messaging.

1061 Communication – the formulation, transmission, receipt and interpretation of messages – is the
1062 foundation of all joint actions within the SOA ecosystem, given the inherent separation – often across
1063 ownership boundaries – of actors in the system.

1064 Communication between **actors** requires that they play the roles of ‘sender’ or ‘receiver’ of messages as
1065 appropriate to a particular action – although it is not necessarily required that they both be active
1066 simultaneously.

1067 An **actor** sends a message **in order** to communicate with other actors. The communication itself is often
1068 not intended as part of the desired real world effect but rather includes messages that seek to establish,
1069 manage, monitor, report on, and guide the joint action throughout its execution.

1070 Like **communication**, joint action usually involves different actors. However, joint action – resulting from
1071 the deliberate actions undertaken by different actors – *intentionally* impacts shared state within the
1072 system leading to real world effects.

1073 **Joint Action**

1074 Joint action is the coordinated set of actions involving the efforts of two or more **actors** to achieve
1075 an effect.

1076 Note that the effect of a joint action is *not* always equivalent to one or more effects of the individual
1077 actions of the participating **actors**, i.e., it may be more than the sum of the parts.

1078 Different viewpoints lead to either communication or joint action as being considered most important. For
1079 example, from the viewpoint of ecosystem security, the integrity of the communications may be dominant;
1080 from the viewpoint of ecosystem governance, the integrity of the joint action may be dominant.

1081 **3.2.4 State, Shared State and Real-World Effect**

1082 **State**

1083 State is the condition of an entity at a particular time.

1084 **State** is characterized by a set of **facts** that is true of the entity. In principle, the total **state** of an entity (or
1085 the world as a whole) is unbounded. In practice, we are concerned only with a subset of the State of an
1086 entity that is measurable and useful in a given context.

1087 For example, the total state of a lightbulb includes the temperature of the filament of the bulb. It also
1088 includes a great deal of other state – the composition of the glass, the dirt that is on the bulb’s surface
1089 and so on. However, an **actor** may be primarily interested in whether the bulb is ‘on’ or ‘off’ and not on the
1090 amount of dirt accumulated. That actor’s characterization of the state of the bulb reduces to the fact: ‘bulb
1091 is now on’.

1092 In a SOA ecosystem, there is a distinction between the set of facts about an entity that only that entity can
1093 access – the so-called Private State – and the set of facts that may be accessible to other actors in the
1094 SOA-based system – the public or Shared State.

1095 **Private State**

1096 The private state is that part of of an entity’s **state** that is knowable by, and accessible to, only
1097 that entity.

1098 **Shared State**

1099 Shared state is that part of an entity’s **state** that is knowable by, and may be accessible to, other
1100 **actors**.

1101 Note that shared state does not imply that the state *is* accessible to *all* actors. It simply refers to that
1102 subset of state that *may* be accessed by *other* actors. Generally this will be the case when actors need to
1103 participate in joint actions.

1104 It is the aggregation of the shared states of pertinent entities that constitutes the desired effect of a joint
1105 action. Thus the change to this shared state is what is experienced in the wider ecosystem as a real world
1106 effect

1107 **3.3 Architectural Implications**

1108 **3.3.1 Social structures**

1109 A SOA ecosystem’s **participants** are organized into various forms of **social structure**. Not all **social**
1110 **structures** are hierarchical: a SOA ecosystem should be able to incorporate peer-to-peer forms of
1111 organization as well as hierarchic structures. In addition, it should be possible to identify and manage any
1112 constitutional agreements that define the **social structures** present in a SOA ecosystem.

- 1113 • Different **social structures** have different rules of engagement but predictable behavior is one of
1114 the underpinnings of trust. This therefore requires mechanisms to:
 - 1115 ○ express constitutions and other organizing principles of participants;
 - 1116 ○ inherit rules of engagement from parent to child social structures.
- 1117 • **Social structures** have **roles** and members and this impacts who may be authorized to act and in
1118 what circumstances. This requires mechanisms to:
 - 1119 ○ identify and manage members of **social structures**
 - 1120 ○ Identify and manage attributes of the members
 - 1121 ○ describe **roles** and **role** adoption
- 1122 • Social structures overlap and interact, giving rise to situations in which rules of engagement may
1123 conflict. In addition, a given **actor** may be member of multiple **social structures** and the social
1124 structures may be associated with different jurisdictions. This requires mechanisms to:
 - 1125 ○ identify the social structures that are active during a series of joint actions;
 - 1126 ○ identify and resolve conflicts and inconsistencies.

1127 **3.3.2 Resource and Ownership**

1128 Communication about and between, visibility into, and leveraging of resources requires the unambiguous
1129 identification of those resources. Ensuring unambiguous identities implies

- 1130
- 1131
- 1132
- 1133
- 1134
- Mechanism for assigning and guaranteeing uniqueness of globally unique identifiers
 - Identifying the extent of the enterprise over which the identifier needs to be understandable and unique
 - Mechanism and framework for ensuring the long-livedness of identifiers (i.e., they cannot just change arbitrarily)

1135 3.3.3 Policies and Contracts

- 1136
- 1137
- 1138
- 1139
- 1140
- 1141
- 1142
- 1143
- 1144
- 1145
- Policies are constraints
 - Policies **MUST** be expressed
 - Constraints **MUST** be enforceable
 - Management of potentially large numbers of policies **MUST** be achievable
 - Policies have owners
 - Policies **SHOULD** be established by social structures.
 - Policies may not be consistent with one another
 - **Policy** conflict resolution techniques **MUST** exist and be in place
 - Agreements are constraints agreed to
 - Contracts **SHOULD** be enforced by mechanisms of the social structure

1146 3.3.4 Communications as a Means of Mediating Action

1147 Using message exchange for mediating **action** implies

- 1148
- 1149
- 1150
- 1151
- 1152
- 1153
- 1154
- 1155
- 1156
- 1157
- 1158
- 1159
- 1160
- Ensuring correct identification of the structure of messages:
 - Identifying the syntax of the message;
 - Identifying the vocabularies used in the communication
 - Identifying the higher-level structure of the communication, such as policy assertion, contract enforcement, etc.
 - A principal objective of communication is to mediate **action**
 - Messages convey actions and **events**
 - Receiving a message is an **action**, but is not the same **action** as the action conveyed by the message
 - Actions are associated with objectives of the **actors** involved
 - Explicit representation of objectives may facilitate automated processing of messages
 - An **actor** agreeing to adopt an objective becomes responsible for that objective

1161 3.3.5 Semantics

1162 Semantics is pervasive in a SOA ecosystem. There are many forms of utterance that are relevant to the

1163 ecosystem: apart from communicated content there are policy statements, goals, purposes, descriptions,

1164 and agreements which are all forms of utterance.

1165 The operation of the SOA ecosystem is significantly enhanced if

- 1166
- 1167
- 1168
- 1169
- 1170
- 1171
- 1172
- 1173
- 1174
- A careful distinction is made between public semantics and private semantics. In particular, it **MUST** be possible for actors to process content such as communications, descriptions and policies solely on the basis of the public semantics of those utterances.
 - A well founded semantics ensures that any assertions that are essential to the operator of the ecosystem (such as policy statements, and descriptions) have carefully chosen written expressions and associated decision procedures.
 - The role of vocabularies as a focal point for multiple actors to be able to understand each other is critical. While no two actors can fully share their interpretation of elements of vocabularies, ensuring that they do understand the public meaning of vocabularies' elements is essential.

1175 3.3.6 Trust and Risk

1176 In traditional systems, the balance between trust and risk is achieved by severely restricting interactions

1177 and by controlling the **participants** of a system.

1178 It is important that **actors** are able to explicitly reason about both trust and risk in order to effectively
1179 participate in a SOA ecosystem. The more open and public the SOA ecosystem is, the more important it
1180 is for **actors** to be able to reason about their participation.

1181 **3.3.7 Needs, Requirements and Capabilities**

1182 In the process of capturing needs as requirements, and the subsequent requirements decomposition and
1183 allocation processes need to be informed by capabilities that already exist.

- 1184 • Architecture needs to
- 1185 ○ Take into account existing capabilities available as services

1186 **3.3.8 The Importance of Action**

1187 **Participants** participate in a SOA ecosystem in order to get their needs met. This involves **action**; both
1188 individual actions and joint actions.

1189 Any architectural realization of a SOA ecosystem should address:

- 1190 • How actions are modeled:
- 1191 ○ Identifying the performer or agent of the **action**;
- 1192 ○ the target of the **action**; and the
- 1193 ○ verb of the **action**.

1194 Any explicit models of joint action should take into account

- 1195 • The choreography that defines the joint action.
- 1196 • The potential for multiple joint actions to be layered on top of each other

4 Realization of a SOA Ecosystem view

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

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

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

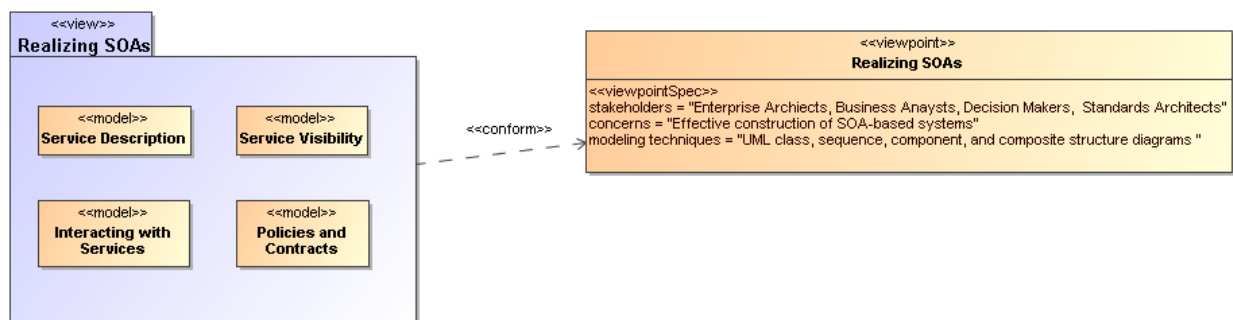


Figure 10 - Model Elements Described in the Realization of a SOA Ecosystem view

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

4.1 Service Description Model

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

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

There are several points to note regarding service description:

- The Reference [Model](#) states that one of the hallmarks of SOA is the large amount of associated description. The model presented below focuses on the description of services but it is equally important to consider the descriptions of the consumer, other participants, and needed resources other than services.
- [Descriptions](#) are inherently incomplete but may be determined as *sufficient* when it is possible for the participants to access and use the described services based only on the descriptions provided. This means that, at one end of the spectrum, a description along the lines of "That service on that machine" may be sufficient for the intended audience. On the other extreme, a service description with a machine-process-able description of the semantics of its operations

- 1234 and real world effects may be required for services accessed via automated service discovery
1235 and planning systems.
- 1236 • Descriptions come with context, i.e. a given description comprises information needed to
1237 adequately support the context. For example, a list of items can define a version of a service, but
1238 for many contexts an indicated version number is sufficient without the detailed list. The current
1239 model focuses on the description needed by a service consumer to understand what the service
1240 does, under what conditions he service will do it, how well the service does it, and what steps are
1241 needed by the consumer to initiate and complete a service interaction. Such information also
1242 enables the service provider to clearly specify what is being provided and the intended conditions
1243 of use.
 - 1244 • Descriptions change over time as, for example, the ingredients and nutrition information for food
1245 labeling continues to evolve. A requirement for transparency of transactions may require
1246 additional description for those associated contexts.
 - 1247 • Description always proceeds from a basis of what is considered "common knowledge". This may
1248 be social conventions that are commonly expected or possibly codified in law. It is impossible to
1249 describe everything and it can be expected that a mechanism as far reaching as SOA will also
1250 connect entities where there is inconsistent "common" knowledge.
 - 1251 • Descriptions become the collection point of information related to a service or any other resource,
1252 but it is not necessarily the originating point or the motivation for generating this information. In
1253 particular, given a SOA service as the access to an underlying capability, the service may point to
1254 some of the capability's previously generated description, e.g. a service providing access to a
1255 data store may also have access to information indicating the freshness of the data.

1256 These points emphasize that there is no one "right" description for all contexts and for all time. Several
1257 descriptions for the same subject may exist at the same time, and this emphasizes the importance of the
1258 description referencing source material maintained by that material's owner rather than having multiple
1259 copies that become out of synch and inconsistent.

1260 It may also prove useful for a description assembled for one context to cross-reference description
1261 assembled for another context as a way of referencing ancillary information without overburdening any
1262 single description. Rather than a single artifact, description can be thought of as a web of documents that
1263 enhance the total available description.

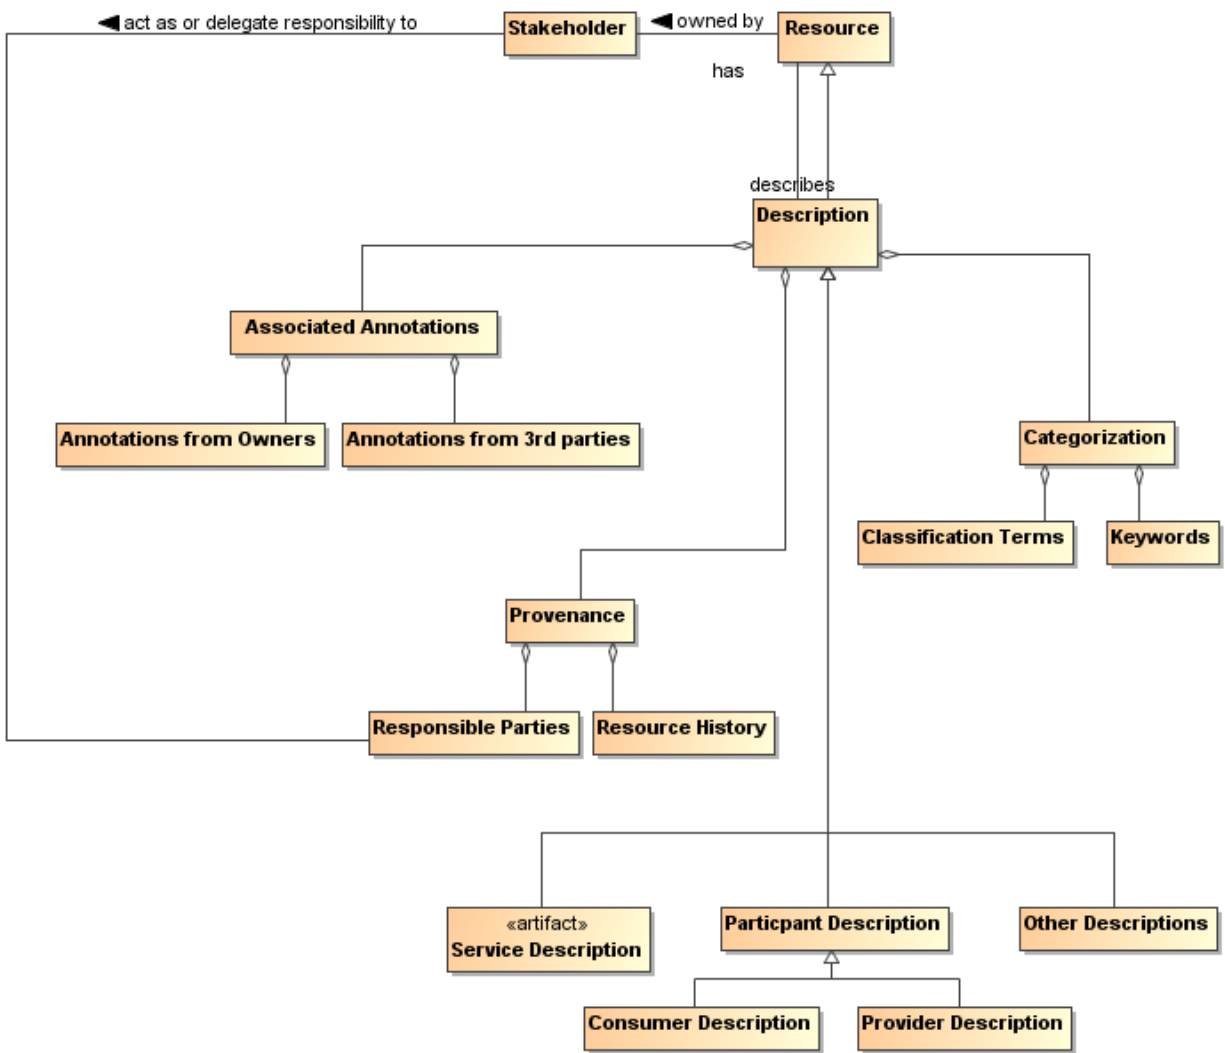
1264 This Reference Architecture Foundation uses the term service description for consistency with the
1265 concept defined in the Reference Model. Some SOA literature treats the idea of a "service contract" as
1266 equivalent to service description. In the SOA-RAF, the term service description is preferred. Replacing
1267 the term "service description" with the term "service contract" implies that just one side of the interaction is
1268 governing and misses the point that a single set of policies identified by a service description may lead to
1269 numerous contracts, i.e. service level agreements, leveraging the same description.

1270 **4.1.1 The Model for Service Description**

1271 Figure 11 shows Service Description as a subclass of the general Description class, where Description is
1272 a subclass of the resource class as defined in Section 3.1.3.1. In addition, each resource is assumed to
1273 have a description. The following section discusses the relationships among elements of general
1274 description and the subsequent sections focus on service description. Other descriptions, such as those
1275 of participants, are important to SOA but are not individually elaborated in this document.

1276 **4.1.1.1 Elements Common to General Description**

1277 The general Description class is composed of a number of elements that are expected to be common
1278 among all descriptions supporting a service oriented architecture. A registry often contains a subset of the
1279 description instance, where the chosen subset is identified as that which facilitates discovery. Additional
1280 information contained in a more complete description may be needed to initiate and continue interaction.



1282
1283 *Figure 11 - General Description*

1284 **4.1.1.1 Provenance**

1285 While the resource Identifier provides the means to know which subject and subject description are being
 1286 considered, Provenance as related to the Description class provides information that reflects on the
 1287 quality or usability of the subject. Provenance specifically identifies the stakeholder (human, defined role,
 1288 organization, ...) that assumes responsibility for the resource being described and tracks historic
 1289 information that establishes a context for understanding what the resource provides and how it has
 1290 changed over time. Responsibilities may be directly assumed by the stakeholder who owns a resource or
 1291 the Owner may designate Responsible Parties for the various aspects of maintaining the resource and
 1292 provisioning it for use by others. There may be more than one stakeholder identified under Responsible
 1293 Parties; for example, one stakeholder may be responsible for code maintenance while another is
 1294 responsible for provisioning of the executable code.

1295 **4.1.1.2 Keywords and Classification Terms**

1296 A traditional element of description has been to associate the resource being described with predefined
 1297 keywords or classification taxonomies that derive from referenceable formal definitions and vocabularies.
 1298 This Reference Architecture Foundation does not prescribe which vocabularies or taxonomies may be
 1299 referenced, nor does it limit the number of keywords or classifications that may be associated with the

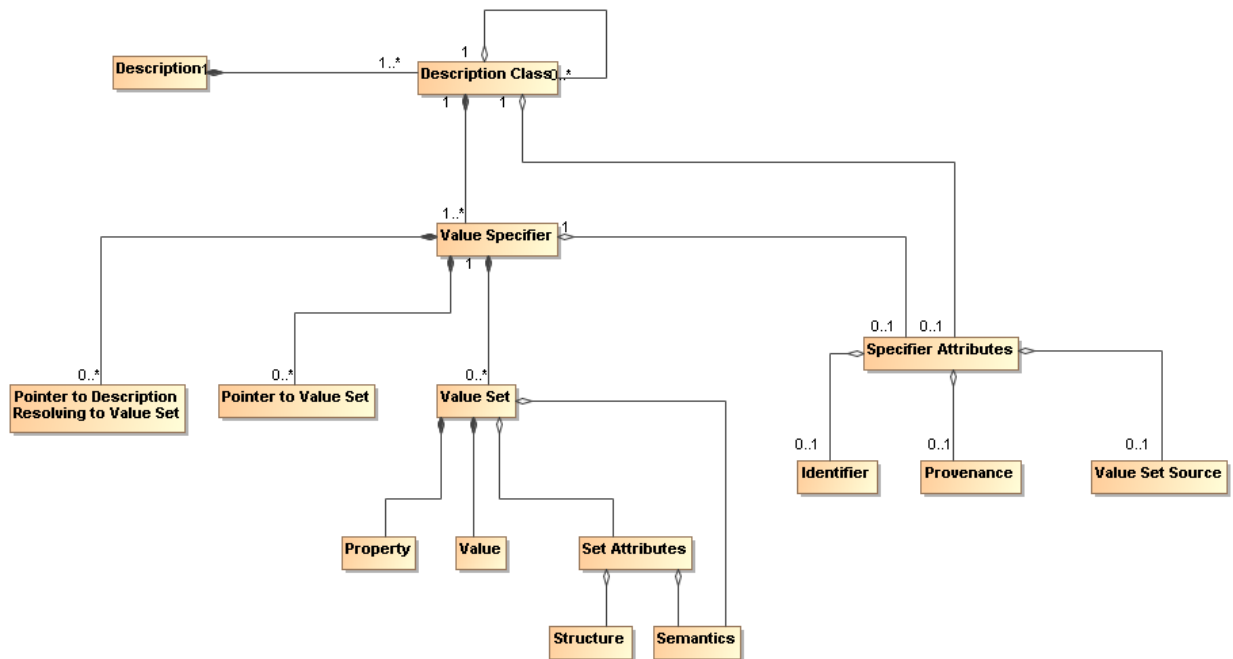
1300 resource. It does, however, state that a normative definition of any terms or keywords SHOULD be
 1301 referenced, whether that be a representation in a formal ontology language, a pointer to an online
 1302 dictionary, or any other accessible source. See Section 4.1.1.2 for further discussion on associating
 1303 semantics with assigned values.

1304 4.1.1.1.3 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 12 - Representation of a Description*

1316 Figure 11 shows the template for a general description, but individual description instances depend on
 1317 the ability to associate meaningful values with the identified elements. Figure 12 shows a model for a
 1318 collection of information that provides for value assignment and traceability for both the meaning and the
 1319 source of a value. The model is not meant to replace existing or future schema or other structures that
 1320 have or will be defined for specific implementations, but it is meant as guidance for the information such
 1321 structures need to capture to generate sufficient description. It is expected that tools will be developed to
 1322 assist the user in populating description and auto-filling many of these fields, and in that context, this
 1323 model provides guidance to the tool developers.

1324 In Figure 12, each class has an associated value specifier or is made up of components that eventually
 1325 resolve to a value specifier. For example, Description has several components, one of which is
 1326 Categorization, which would have an associated 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
 1329 pointers 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
1333 elsewhere;

1334 • provenance information that identifies the party (individual, role, or organization) that has
1335 responsibility for assigning the value sets to any description component;

1336 • an optional source of the value set, if appropriate and meaningful, e.g. if a particular data source
1337 is mandated.

1338 If the value specifier is contained within a higher-level component (such as Service Description containing
1339 Service Functionality), the component may assume values from the attributes of its container.

1340 Note, provenance as a qualifying attribute of a value specifier is different from provenance as part of an
1341 instance of Description. Provenance for a service identifies those who own and are responsible for the
1342 service, as described in Section 3.1.3. Provenance for a value specifier identifies who is responsible for
1343 choosing and assigning values to the value sets that comprise the value specifier. It is assumed that
1344 granularity at the value specifier level is sufficient and provenance is not required for each value set.

1345 The value set also has attributes that define its structure and semantics.

1346 • The semantics of the value set property should be associated with a semantic context conveying
1347 the meaning of the property within the execution context, where the semantic context could vary
1348 from a free text definition to a formal ontology.

1349 • For numeric values, the structure would provide the numeric format of the value and the
1350 “semantics” would be conveyed by a dimensional unit with an identifier to an authoritative source
1351 defining the dimensional unit and preferred mechanisms for its conversion to other dimensional
1352 units of like type.

1353 • For nonnumeric values, the structure would provide the data structure for the value
1354 representation and the semantics would be an associated semantic model.

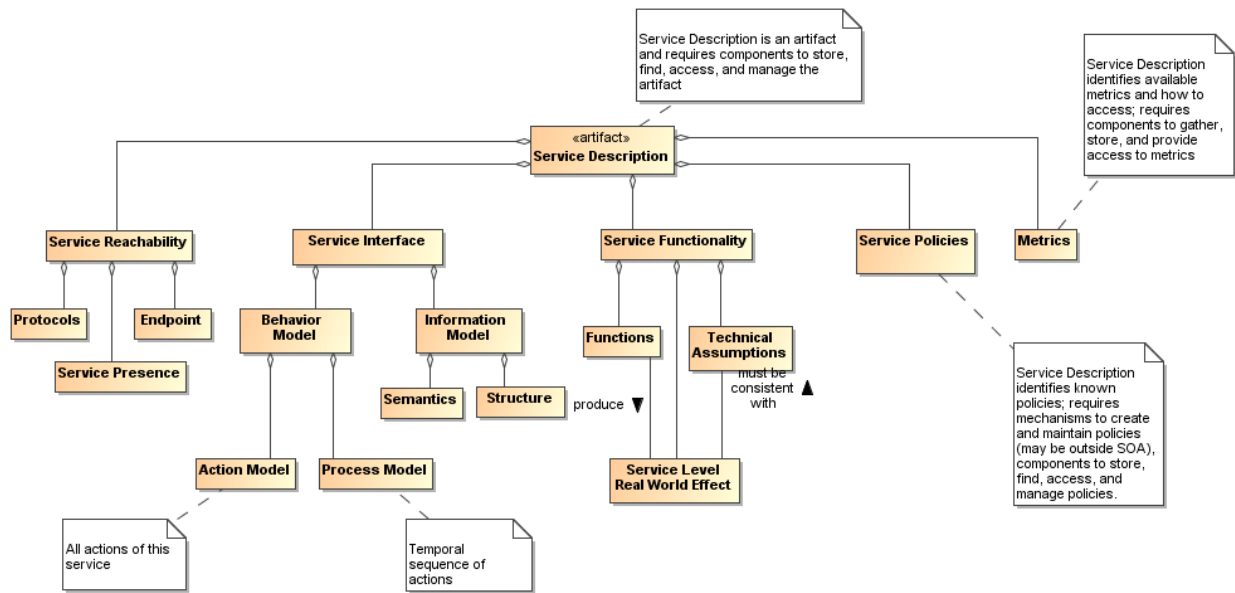
1355 • For pointers, architectural guidelines would define the preferred addressing scheme.

1356 The value specifier may indicate a default semantic model for its component value sets and the individual
1357 value sets may provide an override.

1358 The property-value pair construct is introduced for the value set to emphasize the need to identify
1359 unambiguously both what is being specified and what is a consistent associated value. The further
1360 qualifying of Structure and Semantics in the Set Attributes allows for flexibility in defining the form of the
1361 associated values.

1362 **4.1.1.3 Model Elements Specific to Service Description**

1363



1364

1365 *Figure 13 - Service Description*

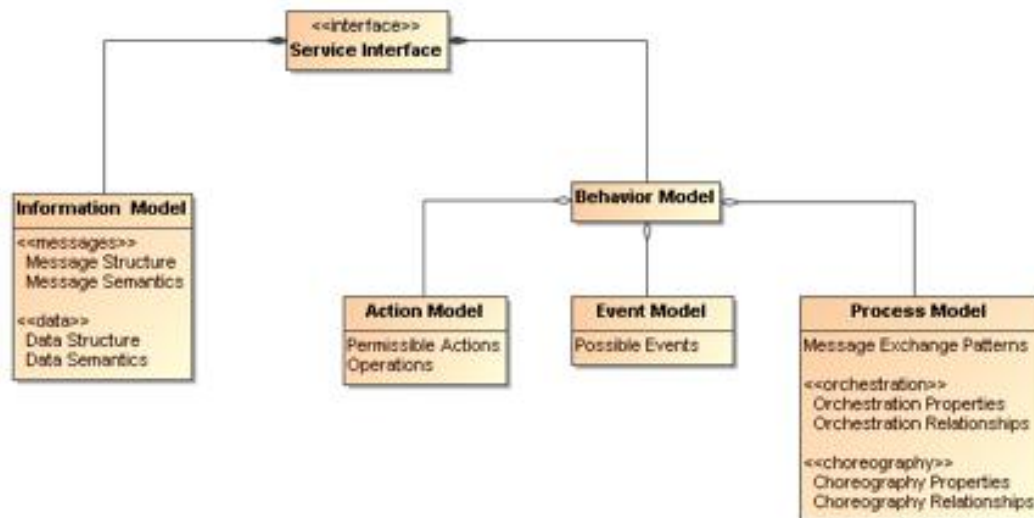
1366 The major elements for the Service Description subclass follow directly from the areas discussed in the
 1367 Reference Model. Here, we discuss the detail shown in Figure 13 and the purpose served by each
 1368 element of service description.

1369 Note, the intent in the subsections that follow is to describe how a particular element, such as the service
 1370 interface, is reflected in the service description, not to elaborate on the details of that element.

1371 **4.1.1.3.1 Service Interface**

1372 As noted in the Reference Model, the service interface is the means for interacting with a service. For the
 1373 SOA-RAF and as shown in Section 4.3 the service interface supports an exchange of messages, where

- 1374 • the message conforms to a referenceable message exchange pattern (MEP),
- 1375 • the message payload conforms to the structure and semantics of the indicated information model,
- 1376 • the messages are used to denote events or actions against the service, where the actions are
- 1377 specified in the action model and any required sequencing of actions is specified in the process
- 1378 model.



1379

1380 *Figure 14 - Service Interface*

1381 Note we distinguish the structure and semantics of the message from that of the underlying protocol that
 1382 conveys the message. The message structure may include nested structures that are independently
 1383 defined, such as an enclosing envelope structure and an enclosed data structure.

1384 These aspects of messages are discussed in more detail in Section 4.3.2.

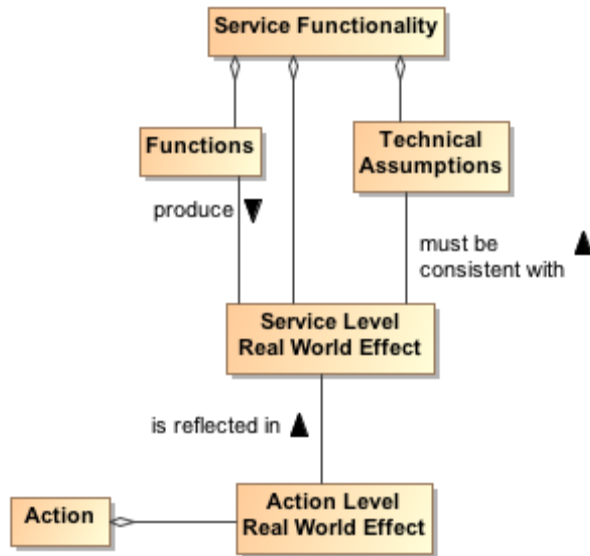
1385 **4.1.1.3.2 Service Reachability**

1386 Service reachability, as modeled in Section 4.2.2.3 enables service participants to locate and interact with
 1387 one another. To support service reachability, the service description should indicate the endpoints to
 1388 which a service consumer can direct messages to invoke actions and the protocol to be used for
 1389 message exchange using that endpoint.

1390 As generally applied to an action, the endpoint is the conceptual location where one applies an action;
 1391 with respect to service description, it is the actual address where a message is sent.

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 as a result of interacting with a service. Service Functionality, shown in Figure 13 as part of the
 1396 overall Service Description model and extended in Figure 15, is a clear 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 *Figure 15 - Service Functionality*

1401 The Service Functionality may also be limited by technical assumptions/constraints that underlie the
1402 effects that can result. Technical constraints are defined as domain specific restrictions and may express
1403 underlying physical limitations, such as flow speeds must be below sonic velocity or disk access that
1404 cannot be faster than the maximum for its host drive. Technical constraints are related to the underlying
1405 capability accessed by the service. In any case, the real world effects must be consistent with the
1406 technical assumptions/constraints.

1407 In Figure 13 and Figure 15, we specifically refer to Service Level and Action Level real world effects.

1408 **Service Level Real World Effect**

1409 A service level real world effect is a specific change in the state or the information returned as a
1410 result of interacting with a service.

1411 **Action Level Real World Effect**

1412 An action level real world effect is a specific change in the state or the information returned as a
1413 result of interacting through a specific action.

1414 Service description describes the service as a whole while the component aspects should contribute to
1415 that whole. Thus, while individual Actions may contribute to the real world effects to be realized from
1416 interaction with the service, there would be a serious disconnect for Actions to contribute real world
1417 effects that could not consistently be reflected in the Service Level Real World Effects and thus the
1418 Service Functionality. The relationship to Action Level Real World Effects and the implications on
1419 defining the scope of a service are discussed in Section 4.1.2.1.

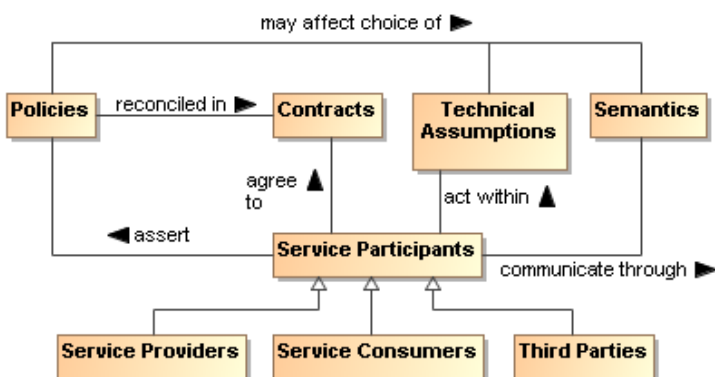
1420 Elements of Service Functionality may be expressed as natural language text, reference an existing
1421 taxonomy of functions or other formal model.

1422 **4.1.1.3.4 Service Policies, Metrics, and Compliance Records**

1423 Policies prescribe the conditions and constraints for interacting with a service and impact the willingness
1424 to continue visibility with the other participants. Whereas technical constraints are statements of “physical”
1425 fact, policies are subjective assertions made by the service provider (sometimes as passed on from
1426 higher authorities).

1427 The service description provides a central location for identifying what policies have been asserted by the
1428 service provider. The specific representation of the policy, e.g. in some formal policy language, is outside
1429 of the service description. The service description would reference the normative definition of the policy.

1430 Policies may also be asserted by other service participants, as illustrated by the model shown in Figure
 1431 16. Policies that are generally applicable to any interaction with the service are asserted by the service
 1432 provider and included in the Service Policies section of the service description.



1433
 1434 *Figure 16 - Model for Policies and Contracts as related to Service Participants*

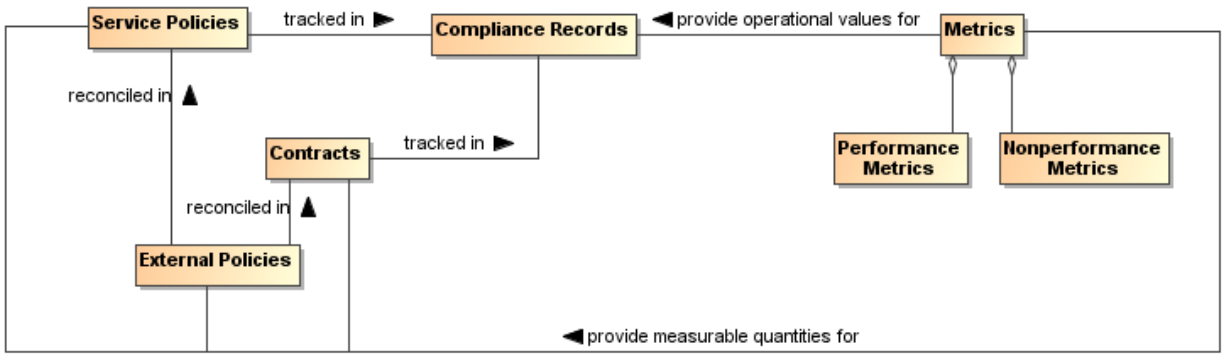
1435 In Figure 16, we specifically refer to policies at the service level. In a similar manner to that discussed for
 1436 Service Level vs. Action Level Real World Effects in Section 4.1.1.3.3, individual Actions may have
 1437 associated policies stating conditions for performing the action, but these must be reflected in and be
 1438 consistent with the policies made visible at the service level and thus the description of the service as a
 1439 whole. The relationship to Action Level Policies and the implications on defining the scope of a service
 1440 are discussed in Section 4.1.2.1.

1441 As noted in Figure 16, the policies asserted may be reflected as Technical Constraints that available
 1442 services or their underlying capabilities must be capable of meeting; it may similarly affect the semantics
 1443 that can be used. For example of the former, there may be a policy that specifies the surge capacity to
 1444 be accommodated by a server, but a service that is not designed to make use of the larger server
 1445 capacity would not satisfy the intent of the policy and would not be appropriate to use. For the latter, a
 1446 policy may require that only services that support interaction via a community-sponsored vocabulary can
 1447 be used.

1448 Contracts are agreements among the service participants. The contract may reconcile inconsistent
 1449 policies asserted by the participants or may specify details of the interaction. Service level agreements
 1450 (SLAs) are one commonly used category of contracts.

1451 The definition and later enforcement of policies and contracts are predicated on the potential for
 1452 measurement; the relationships among the relevant concepts are shown in the model in Figure 17.
 1453 Performance Metrics identify quantities that characterize the speed and quality of realizing the real world
 1454 effects produced using the SOA service; in addition, policies and contracts may depend on
 1455 nonperformance metrics, such as whether a license is in place to use the service. Some of these metrics
 1456 reflect the underlying capability, e.g. a SOA service cannot respond in two seconds if the underlying
 1457 capability is expected to take five seconds to do its processing; some metrics reflect the SOA service, e.g.
 1458 the additional overhead introduced when making data access requests across the network.

1459



1460

1461

Figure 17 - Policies and Contracts, Metrics, and Compliance Records

1462

1463

1464

1465

1466

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

1467

1468

1469

1470

1471

1472

1473

1474

The use of metrics to evaluate compliance is discussed in Section 4.1.1.3.4. The results of compliance evaluation SHOULD be maintained in compliance records and the means to access the compliance records MAY be included in the Service Policies portion of the service description. For example, the description may be in the form of static information (e.g. over the first year of operation, this service had a 91% availability), a link to a dynamically generated metric (e.g. over the past 30 days, the service has had a 93.3% availability), or access to a dynamic means to check the service for current availability (e.g., a ping). The relationship between service presence and the presence of the individual actions that can be invoked is discussed under Reachability in Section 4.2.2.3.

1475

1476

1477

1478

Note, even when policies relate to the perspective of a single participant, policy compliance can be measured and policies may be enforceable without contractual agreement with other participants. While certain elements of contracts and contract compliance are likely private, public aspects of compliance should be reflected in the compliance record information referenced in the service description.

1479

4.1.2 Use of Service Description

1480

4.1.2.1 Service Description in support of Service Interaction

1481

1482

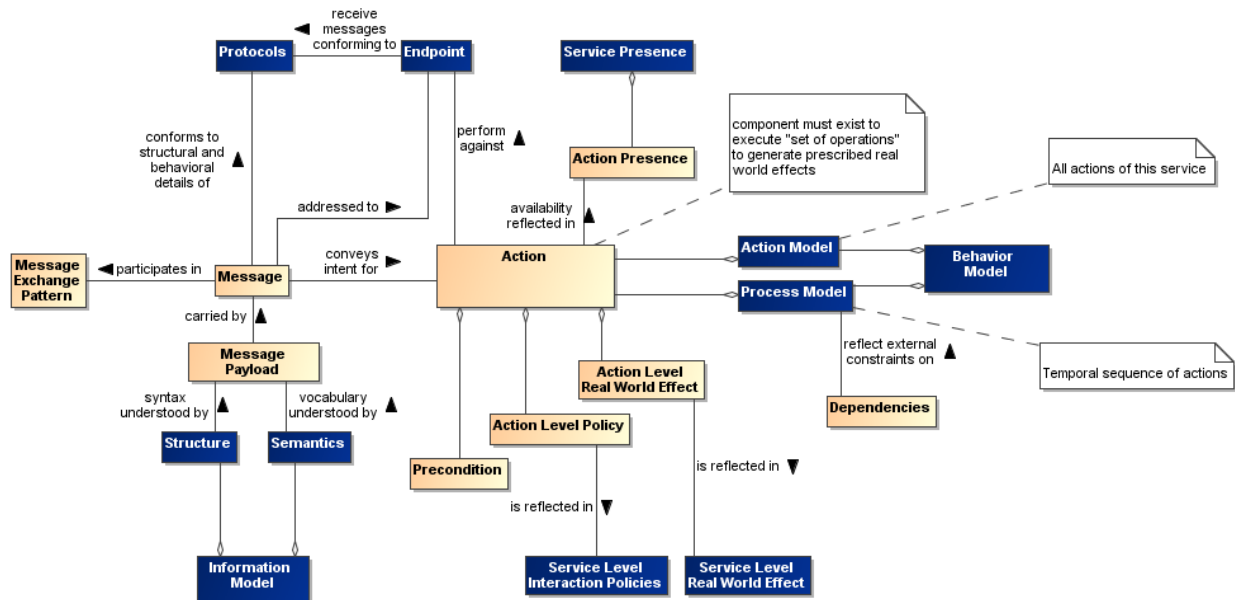
1483

1484

1485

1486

If we assume we have awareness, the service participants must still establish willingness and presence to ensure full visibility (See Section 4.2) and to interact with the service. Service description provides necessary information for many aspects of preparing for and carrying through with interaction. Recall the fundamental definition of service is a mechanism to access an underlying capability; the service description describes this mechanism and its use. It lays the groundwork for what can occur, whereas service interaction comprises the specifics through which real-world effects are realized.



1487

1488 *Figure 18 - Relationship between Action and Components of Service Description Modelx*

1489 Figure 18 combines the models in the subsections of Section 4.1.1 to concisely relate action and the
 1490 relevant components of the Service Description model. The purpose of Figure 18 is to demonstrate that
 1491 the components of service description go beyond arbitrary documentation and form the critical set of
 1492 information needed to define the what and how of action. In Figure 18, the leaf nodes from Figure 13 are
 1493 shown in blue.

1494 Action is typically invoked via a Message where the structure and behavioral details of the message
 1495 conform to an identified Protocol and is directed to the address of the identified endpoint, and the
 1496 message payload conforms to the service Information Model.

1497 The availability of an action is reflected in the Action Presence and each Action Presence contributes to
 1498 the overall Service Presence; this is discussed further in Section 4.2.2.3. Each action has its own
 1499 endpoint and also its own protocols associated with the endpoint⁷ and to some extent, e.g. current or
 1500 average availability, there is presence for the action through that endpoint. The endpoint and service
 1501 presence are also part of the service description.

1502 An action may have preconditions where a Precondition is something that needs to be in place before an
 1503 action can occur, e.g. confirmation of a precursor action. Whether preconditions are satisfied is evaluated
 1504 when an actor tries to perform the action and not before. Presence for an action means an actor can
 1505 initiate it and is independent of whether the preconditions are satisfied. However, the successful
 1506 completion of the action may depend on whether its preconditions were satisfied.

1507 Analogous to the relationship between actions and preconditions, the Process Model may imply
 1508 Dependencies for succeeding steps in a process, e.g. that a previous step has successfully completed, or
 1509 may be isolated to a given step. An example of the latter would be a dependency that the host server has
 1510 scheduled maintenance and access attempts at these times would fail. Dependencies related to the

⁷ 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.

1511 process model do not affect the presence of a service although these may affect whether the business
1512 function successfully completes.

1513 The conditions under which an action can be invoked may depend on policies associated with the action.
1514 The Action Level Policies MUST be reflected in (or subsumed by) the Service Policies because such
1515 policies may be critical to determining whether the conditions for use of the service are consistent with the
1516 policies asserted by the service consumer. The Service Policies are included in the service description.

1517 Similarly, the result of invoking an action is one or more real world effects, and any Action Level Real
1518 World Effects MUST be reflected in the Service Level Real World Effect included in the service
1519 description. The unambiguous expression of action level policies and real world effects as service
1520 counterparts is necessary to adequately describe what constitutes the service interaction.

1521 An adequate service description MUST provide a consumer with information needed to determine if the
1522 service policies, the (business) functions, and service-level real world effects are of interest, and there is
1523 nothing in the technical constraints that preclude use of the service.

1524 Note at the service level, the business functions are not concerned with the action or process models.
1525 These models are detailed separately.

1526 The service description is not intended to be isolated documentation but rather an integral part of service
1527 use. Changes in service description SHOULD immediately be made known to consumers and potential
1528 consumers.

1529 **4.1.2.1.1 Description and Invoking Actions Against a Service**

1530 At this point, let us assume the descriptions were sufficient to establish willingness; see Section 4.2.2.2.
1531 Figure 18 indicates the service endpoint establishes where to actually carry out the interaction. This is
1532 where we start considering the action and process models.

1533 The action model identifies the multiple actions a user can perform against a service and the user would
1534 perform these in the context of the process model as specified or referenced under the Service Interface
1535 portion of Service Description. For a given business function, there is a corresponding process model,
1536 where any process model may involve multiple actions. From the above discussion of model elements of
1537 description we may conclude (1) actions have reachability information, including endpoint and presence,
1538 (2) presence of service is some aggregation of presence of its actions, (3) action preconditions and
1539 service dependencies do not affect presence although these may affect successful completion.

1540 Having established visibility, the interaction can proceed. Given a business function, the consumer knows
1541 what will be accomplished (the service functionality), the conditions under which interaction will proceed
1542 (service policies and contracts), and the process that must be followed (the process model). The
1543 remaining question is how the description information for structure and semantics enable interaction.

1544 We have established the importance of the process model in identifying relevant actions and their
1545 sequence. Interaction proceeds through messages and thus it is the syntax and semantics of the
1546 messages with which we are here concerned. A common approach is to define the structure and
1547 semantics that can appear as part of a message; then assemble the pieces into messages; and,
1548 associate messages with actions. Actions make use of structure and semantics as defined in the
1549 information model to describe its legal messages.

1550 The process model identifies actions to be performed against a service and the sequence for performing
1551 the actions. For a given action, the Reachability portion of description indicates the protocol bindings that
1552 are available, the endpoint corresponding to a binding, and whether there is presence at that endpoint.
1553 An interaction is through the exchange of messages that conform to the structure and semantics defined
1554 in the information model and the message sequence conforming to the action's identified MEP. The
1555 result is some portion of the real world effect that must be assessed and/or processed (e.g. if an error
1556 exists, that part that covers the error processing would be invoked).

1557 **4.1.2.1.2 The Question of Multiple Business Functions**

1558 Action level effects and policies MUST be reflected at the service level for service description to support
1559 visibility.

1560 It is assumed that a SOA service represents an identifiable business function to which policies can be
1561 applied and from which desired business effects can be obtained. While contemporary discussions of
1562 SOA services and supporting standards do not constrain what actions or combinations of actions can or
1563 should be defined for a service, the SOA-RAF considers the implications of service description in defining
1564 the range of actions appropriate for an individual SOA service.

1565 Consider the situation if a given SOA service is the mechanism for access to multiple independent (but
1566 loosely related) business functions. These are not multiple effects from a single function but multiple
1567 functions with potentially different sets of effects for each function. A service can have multiple actions a
1568 user may perform against it, and this does not change with multiple business functions. As an individual
1569 business function corresponds to a process model, so multiple business functions imply multiple process
1570 models. The same action may be used in multiple process models but the aggregated service presence
1571 would be specific to each business function because the components being aggregated may be different
1572 between process models. In summary, for a service with multiple business functions, each function has
1573 (1) its own process model and dependencies, (2) its own aggregated presence, and (3) possibly its own
1574 list of policies and real world effects.

1575 A common variation on this theme is for a single service to have multiple endpoints for different levels of
1576 quality of service (QoS). Different QoS imply separate statements of policy, separate endpoints, possibly
1577 separate dependencies, and so on. One could say the QoS variation does not require this because there
1578 can be a single QoS policy that encompasses the variations, and all other aspects of the service would be
1579 the same except for the endpoint used for each QoS. However, the different aspects of policy at the
1580 service level would need to be mapped to endpoints, and this introduces an undesirable level of coupling
1581 across the elements of description. In addition, it is obvious that description at the service level can
1582 become very complicated if the number of combinations is allowed to grow.

1583 One could imagine a service description that is basically a container for action descriptions, where each
1584 action description is self contained; however, this would lead to duplication of description components
1585 across actions. If common description components are factored, this either is limited to components
1586 common across all actions or requires complicated tagging to capture the components that often but do
1587 not universally apply.

1588 If a provider cannot describe a service as a whole but must describe every action, this leads to the
1589 situation where it may be extremely difficult to construct a clear and concise service description that can
1590 effectively support discovery and use without tedious logic to process the description and assemble the
1591 available permutations. In effect, if adequate description of an action begins to look like description of a
1592 service, it may be best to have it as a separate service.

1593 Recall, more than one service can access the same underlying capability, and this is appropriate if a
1594 different real world effect is to be exposed. Along these lines, one can argue that different QoS are
1595 different services because getting a response in one minute rather than one hour is more than a QoS
1596 difference; it is a fundamental difference in the business function being provided.

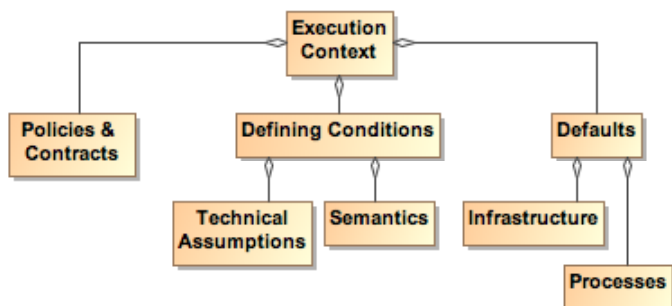
1597 As a best practice, a criteria for whether a service is appropriately scoped may be the ease or difficulty in
1598 creating an unambiguous service description. A consequence of having tightly-scoped services is there
1599 will likely be a greater reliance on combining services, i.e. more fundamental business functions, to create
1600 more advanced business functions. This is consistent with the principles of service oriented architecture
1601 and is the basic position of the Reference Architecture, although not an absolute requirement. Combining
1602 services increases the reliance on understanding and implementing the concepts of orchestration,
1603 choreography, and other approaches yet to be developed; these are discussed in more detail in section
1604 4.4 Interacting with Services.

1605 **4.1.2.1.3 Service Description, Execution Context, and Service Interaction**

1606 The service description MUST provide sufficient information to support service visibility, including the
1607 willingness of service participants to interact. However, the corresponding descriptions for providers and
1608 consumers may both contain policies, technical assumptions, constraints on semantics, and other
1609 technical and procedural conditions that must be aligned to define the terms of willingness. The
1610 agreements which encapsulate the necessary alignment form the basis upon which interactions may
1611 proceed – in the Reference Model, this collection of agreements and the necessary environmental
1612 support establish the execution context.

1613 To illustrate the concept of the execution context, consider a Web-based system for timecard entry. For
 1614 an employee onsite at an employer facility, the execution context requires a computer connected to the
 1615 local network and the employee must enter their network ID and password. Relevant policies include that
 1616 the employee must maintain the most recent anti-virus software and virus definitions for any computer
 1617 connected to the network.

1618 For the same employee connecting from offsite, the execution context specifies the need for a computer
 1619 with installed VPN software and a security token to negotiate the VPN connection. The execution context
 1620 also includes proxy settings as needed to connect to the offsite network. The employee must still comply
 1621 with the requirements for onsite computers and access, but the offsite execution context includes
 1622 additional items before the employee can access the same underlying capability and realize the same
 1623 real world effects, i.e. the timecard entries.

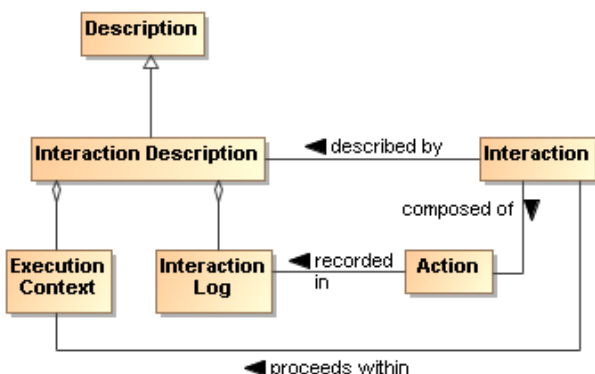


1624
 1625 *Figure 19 - Execution Context*

1626 Figure 19 shows a few broad categories found in execution context. These are not meant to be
 1627 comprehensive. Other items may need to be included to provide a sufficient description of the interaction
 1628 conditions. Any other items not explicitly noted in the model but needed to set the environment SHOULD
 1629 be included in the execution context.

1630 While the execution context captures the conditions under which interaction can occur, it does not capture
 1631 the specific service invocations that do occur in a specific interaction. A service interaction as modeled in
 1632 Figure 20 introduces the concept of an Interaction Description which is composed of both the Execution
 1633 Context and an Interaction Log. The execution context specifies the set of conditions under which the
 1634 interaction occurs and the interaction log captures the sequence of service interactions that occur within
 1635 the execution context. This sequence should follow the Process Model but can include details beyond
 1636 those specified there. For example, the Process Model may specify an action that results in identifying a
 1637 data source, and the identified source is used in a subsequent action. The Interaction Log would record
 1638 the specific data source used.

1639 The execution context can be thought of as a container in which the interaction occurs and the interaction
 1640 log captures what happens inside the container. This combination is needed to support auditability and
 1641 repeatability of the interactions.



1642
 1643 *Figure 20 - Interaction Description*

1644 SOA allows flexibility to accomplish both repeatability and reusability. In facilitating reusability, a service
1645 can be updated without disrupting the user experience of the service. So, Google can improve their
1646 ranking algorithm without notifying the user about the details of the update.

1647 However, it may also be vital for the consumer to be able to recreate past results or to generate
1648 consistent results in the future, and information such as what conditions, which services, and which
1649 versions of those services were used is indispensable in retracing one's path. The interaction log is a
1650 critical part of the resulting real world effects because it defines how the effects were generated and
1651 possibly the meaning of observed effects. This increases in importance as dynamic composability
1652 becomes more feasible. In essence, a result has limited value if one does not know how it was
1653 generated.

1654 The interaction log SHOULD be a detailed trace for a specific interaction, and its reuse is limited to
1655 duplicating that interaction. An execution context can act as a template for identical or similar
1656 interactions. Any given execution context MAY define the conditions of future interactions.

1657 Such uses of execution context imply (1) a standardized format for capturing execution context and (2) a
1658 subclass of general description could be defined to support visibility of saved execution contexts. The
1659 specifics of the relevant formats and descriptions are beyond the scope of this document.

1660 A service description is unlikely to track interaction descriptions or the constituent execution contexts or
1661 interaction logs that include mention of the service. However, as appropriate, linking to specific instances
1662 of either of these could be done through associated annotations.

1663 4.1.3 Relationship to Other Description Models

1664 While the representation shown in Figure 12 is derived from considerations related to service description,
1665 it is acknowledged that other metadata standards are relevant and should, as possible, be incorporated
1666 into this work. Two standards of particular relevance are the Dublin Core Metadata Initiative (DCMI)
1667 **[DCMI]** and ISO 11179 **[ISO 11179]**, especially Part 5.

1668 When the service description (or even the general description class) is considered as the DCMI
1669 "resource", Figure 12 aligns nicely with the DCMI resource model. While some differences exist, these
1670 are mostly in areas where DCMI goes into detail that is considered beyond the scope of the current
1671 Reference Architecture. For example, DCMI defines classes of "shared semantics" whereas this
1672 Reference Architecture Framework considers that an identification of relevant semantic models is
1673 sufficient. Likewise, the DCMI "description model" goes into the details of possible syntax encodings
1674 whereas for the Reference Architecture Framework it is sufficient to identify the relevant formats.

1675 With respect to ISO 11179 Part 5, the metadata fields defined in that reference may be used without
1676 prejudice as the properties in Figure 12. Additionally, other defined metadata sets may be used by the
1677 service provider if the other sets are considered more appropriate, i.e. it is fundamental to this reference
1678 architecture to identify the need and the means to make vocabulary declarations explicit but it is beyond
1679 the scope to specify which vocabularies are to be used. In addition, the identification of domain of the
1680 properties and range of the values has not been included in the current Reference Architecture
1681 discussion, but the text of ISO 11179 Part 5 can be used consistently with the model prescribed in this
1682 document.

1683 Description as defined here considers a wide range of applicability and support of the principles of service
1684 oriented architecture. Other metadata models can be used in concert with the model presented here
1685 because most of these focus on a finer level of detail that is outside the present scope, and so provide a
1686 level of implementation guidance that can be applied as appropriate.

1687 4.1.4 Architectural Implications

1688 The definition of service description indicates numerous architectural implications on the SOA ecosystem:

- 1689 • It changes over time and its contents will reflect changing needs and context. This requires the
1690 existence of:
 - 1691 ○ mechanisms to support the storage, referencing, and access to normative definitions of
1692 one or more versioning schemes that may be applied to identify different aggregations of

- 1693 descriptive information, where the different schemes may be versions of a versioning
 1694 scheme itself;
- 1695 ○ configuration management mechanisms to capture the contents of each aggregation and
 1696 apply a unique identifier in a manner consistent with an identified versioning scheme;
 - 1697 ○ one or more mechanisms to support the storage, referencing, and access to conversion
 1698 relationships between versioning schemes, and the mechanisms to carry out such
 1699 conversions.
- 1700 • Description makes use of defined semantics, where the semantics may be used for
 1701 categorization or providing other property and value information for description classes. This
 1702 requires the existence of:
 - 1703 ○ semantic models that provide normative descriptions of the utilized terms, where the
 1704 models may range from a simple dictionary of terms to an ontology showing complex
 1705 relationships and capable of supporting enhanced reasoning;
 - 1706 ○ mechanisms to support the storage, referencing, and access to these semantic models;
 - 1707 ○ configuration management mechanisms to capture the normative description of each
 1708 semantic model and to apply a unique identifier in a manner consistent with an identified
 1709 versioning scheme;
 - 1710 ○ one or more mechanisms to support the storage, referencing, and access to conversion
 1711 relationships between semantic models, and the mechanisms to carry out such
 1712 conversions.
 - 1713 • Descriptions include reference to policies defining conditions of use. In this sense, policies are
 1714 also resources that need to be visible, discoverable, and accessible. This requires the existence
 1715 of (as also enumerated under governance):
 - 1716 ○ description of policies, including a unique identifier for the policy and a sufficient, and
 1717 preferably a machine processible, representation of the meaning of terms used to
 1718 describe the policy, its functions, and its effects;
 - 1719 ○ one or more discovery mechanisms that enable searching for policies that best meet the
 1720 search criteria specified by the service participant; where the discovery mechanism has
 1721 access to the individual policy descriptions, possibly through some repository
 1722 mechanism;
 - 1723 ○ accessible storage of policies and policy descriptions, so service participants can access,
 1724 examine, and use the policies as defined.
 - 1725 • Descriptions include references to metrics which describe the operational characteristics of the
 1726 subjects being described. This requires the existence of (as partially enumerated under
 1727 governance):
 - 1728 ○ the infrastructure monitoring and reporting information on SOA resources;
 - 1729 ○ possible interface requirements to make accessible metrics information generated;
 - 1730 ○ mechanisms to catalog and enable discovery of which metrics are available for a
 1731 described resources and information on how these metrics can be accessed;
 - 1732 ○ mechanisms to catalog and enable discovery of compliance records associated with
 1733 policies and contracts that are based on these metrics.
 - 1734 • Descriptions of the interactions are important for enabling auditability and repeatability, thereby
 1735 establishing a context for results and support for understanding observed change in performance
 1736 or results. This requires the existence of:
 - 1737 ○ one or more mechanisms to capture, describe, store, discover, and retrieve interaction
 1738 logs, execution contexts, and the combined interaction descriptions;
 - 1739 ○ one or more mechanisms for attaching to any results the means to identify and retrieve
 1740 the interaction description under which the results were generated.
 - 1741 • Descriptions may capture very focused information subsets or can be an aggregate of numerous
 1742 component descriptions. Service description is an example of an aggregate for which manual
 1743 maintenance of the whole would not be feasible. This requires the existence of:
 - 1744 ○ tools to facilitate identifying description elements that are to be aggregated to assemble
 1745 the composite description;
 - 1746 ○ tools to facilitate identifying the sources of information to associate with the description
 1747 elements;
 - 1748 ○ tools to collect the identified description elements and their associated sources into a
 1749 standard, referenceable format that can support general access and understanding;

- 1750 ○ tools to automatically update the composite description as the component sources
1751 change, and to consistently apply versioning schemes to identify the new description
1752 contents and the type and significance of change that occurred.
- 1753 • The description is the source of vital information in establishing willingness to interact with a
1754 resource, reachability to make interaction possible, and compliance with relevant conditions of
1755 use. This requires the existence of:
- 1756 ○ one or more discovery mechanisms that enable searching for described resources that
1757 best meet the criteria specified by a service participant;
- 1758 ○ tools to appropriately track users of the descriptions and notify them when a new version
1759 of the description is available.

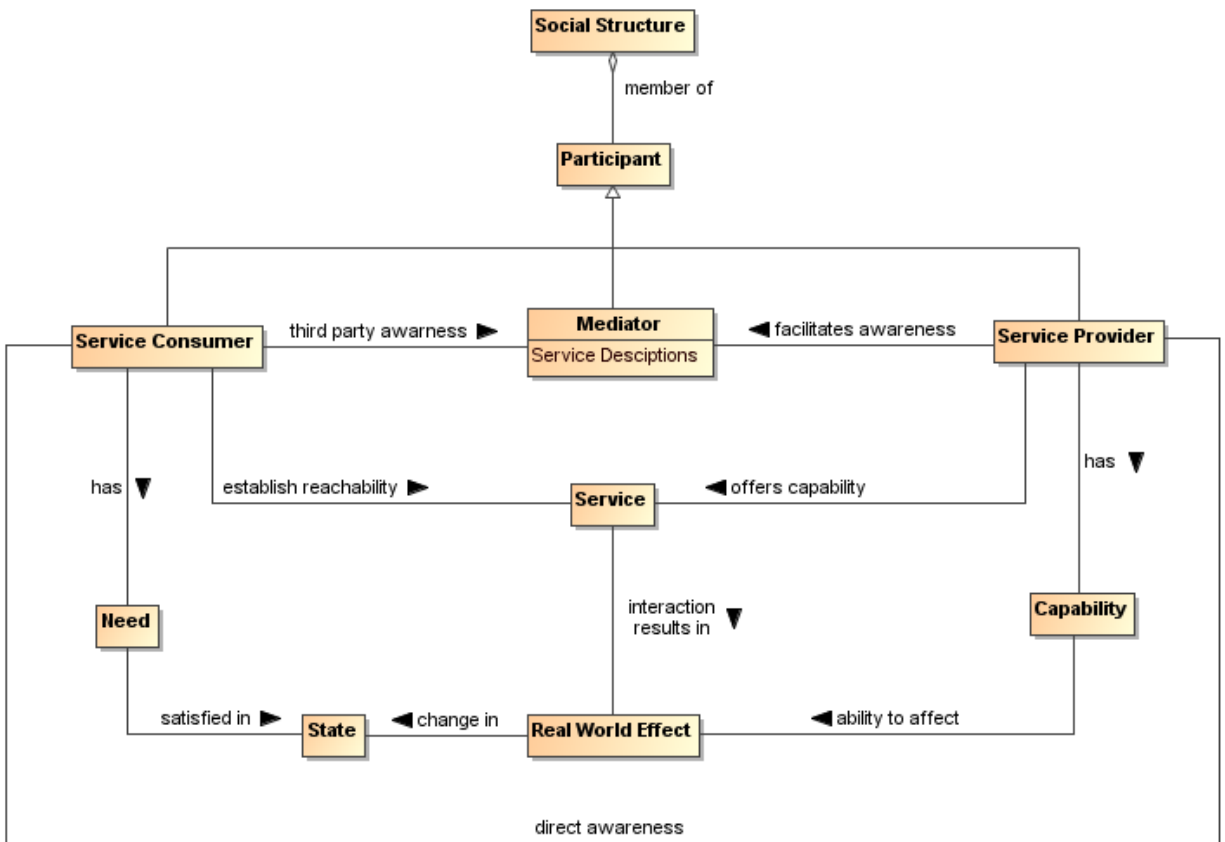
1760 **4.2 Service Visibility Model**

1761 One of the key requirements for participants interacting with each other in the context of a SOA is
1762 achieving visibility: before services can interoperate, the participants have to be visible to each other
1763 using whatever means are appropriate. The Reference Model analyzes visibility in terms of awareness,
1764 willingness, and reachability. In this section, we explore how visibility may be achieved.

1765 **4.2.1 Visibility to Business**

1766 The relationship of visibility to the SOA ecosystem encompasses both human social structures and
1767 automated IT mechanisms. Figure 21 depicts a business setting that is a basis for visibility as related to
1768 the social structure Model in the Participation in a SOA Ecosystem view (see Section 3.1). Service
1769 consumers and service providers may have direct awareness or mediated awareness where mediated
1770 awareness is achieved through some third party. A consumer's willingness to use a service is reflected by
1771 the consumer's presumption of satisfying goals and needs based on the service description. Service
1772 providers offer capabilities that have real world effects that result in a change in state. Reachability of the
1773 service by the consumer may lead to interactions that change the state of the SOA ecosystem. The
1774 consumer can measure the change of state to determine if the claims made by description and the real
1775 world effects of consuming the service meet the consumer's needs.

1776



1778

1779 *Figure 21 - Visibility to Business*

1780 Visibility and interoperability in a SOA ecosystem requires more than location and interface information.
 1781 A meta-model for this broader view of visibility is depicted in Section 4.1. In addition to providing
 1782 improved awareness of service capabilities through description of information such as reachability,
 1783 behavior models, information models, functionality, and metrics, the service description may contain
 1784 policies valuable for determination of willingness to interact.

1785 A mediator using service descriptions may provide event notifications to both consumers and providers
 1786 about information relating to the descriptions. One example of this capability is a publish/subscribe model
 1787 where the mediator allows consumers to subscribe to service description version changes made by the
 1788 provider. Likewise, the mediator may provide notifications to the provider of consumers that have
 1789 subscribed to service description updates.

1790 Another important capability in a SOA environment is the ability to narrow visibility to trusted members
 1791 within a social structure. Mediators for awareness may provide policy based access to service
 1792 descriptions allowing for the dynamic formation of awareness between trusted members.

1793 **4.2.2 Visibility**

1794 Attaining visibility is described in terms of steps that lead to visibility. Different participant communities
 1795 can bring different contexts for visibility within a single social structure, and the same general steps can
 1796 be applied to each of the contexts to accomplish visibility.

1797 Attaining SOA visibility requires

- 1798 • service description creation and maintenance,
- 1799 • processes and mechanisms for achieving awareness of and accessing descriptions,
- 1800 • processes and mechanisms for establishing willingness of participants,
- 1801 • processes and mechanisms to determine reachability.

1802 Visibility may occur in stages, i.e. a participant can become aware enough to look or ask for further
1803 description, and with this description, the participant can decide on willingness, possibly requiring
1804 additional description. For example, if a potential consumer has a need for a tree cutting (business)
1805 service, the consumer can use a web search engine to find web sites of providers. The web search
1806 engine (a mediator) gives the consumer links to relevant web pages and the consumer can access those
1807 descriptions. For those prospective providers that satisfy the consumer's criteria, the consumer's
1808 willingness to interact increases. The consumer may contact several tree services to get detailed cost
1809 information (or arrange for an estimate) and may ask for references (further description). The consumer is
1810 likely to establish full visibility and proceed with interaction with the tree service who mutually establishes
1811 visibility.

1812 **4.2.2.1 Awareness**

1813 An important means for a service participant to be aware of another participant is to have access to a
1814 description of that participant and for the description to have sufficient completeness to establish the other
1815 requirements of visibility.

1816 Awareness is inherently a function of a participant; awareness can be established without any action on
1817 the part of the target participant other than the target providing appropriate descriptions. Awareness is
1818 often discussed in terms of consumer awareness of providers but the concepts are equally valid for
1819 provider awareness of consumers.

1820 Awareness can be decomposed into: creating the descriptions, making them available, and discovering
1821 the descriptions. Discovery can be initiated or it can be by notification. Initiated discovery for business
1822 may require formalization of the required capabilities and resources to achieve business goals.

1823 Achieving awareness in a SOA can range from word of mouth to formal service descriptions in a
1824 standards-based registry-repository. Some other examples of achieving awareness in a SOA are the
1825 use of a web page containing description information, email notifications of descriptions, and document
1826 based descriptions.

1827 A mediator for awareness is a third party participant that provides awareness to one or more consumers
1828 of one or more services. Direct awareness is awareness between a consumer and provider without the
1829 use of a third party. A registry/repository can act as a mediator; a Web page displaying similar
1830 information can also be considered a mediator.

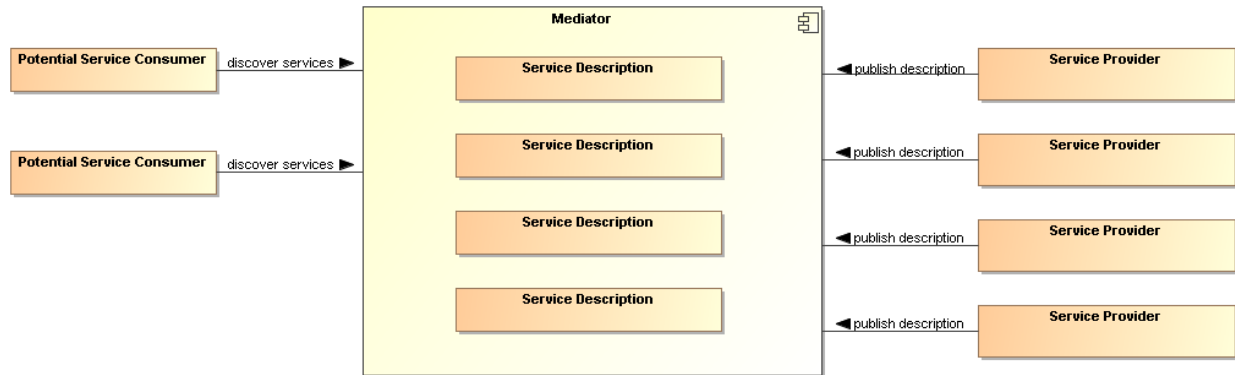
1831 Direct awareness may be the result of having previously established an execution context, or direct
1832 awareness may include determining the presence of services and then querying the service directly for
1833 description. As an example, a priori visibility of some sensor device may provide the means for interaction
1834 or a query for standardized sensor device metadata may be broadcast to multiple locations. If
1835 acknowledged, the service interface for the device may directly provide description to a consumer so the
1836 consumer can determine willingness to interact.

1837 The same medium for awareness may be direct in one context and may be mediated in another context.
1838 For example, a service provider may maintain a web site with links to the provider's descriptions of
1839 services giving the consumers direct awareness to the provider's services. Alternatively, a community
1840 may maintain a mediated web site with links to various provider descriptions of services for any number of
1841 consumers. More than one mediator may be involved, as different mediators may specialize in different
1842 mediation functions.

1843 Descriptions may be formal or informal. Section 4.1, provides a comprehensive model for service
1844 description that can be used to mediate visibility. Using consistent description taxonomies and standards
1845 based mediated awareness helps provide more effective awareness.

1846 **4.2.2.1.1 Mediated Awareness**

1847 Mediated awareness promotes loose coupling by keeping the consumers and services from explicitly
1848 referring to each other. Mediation lets interaction vary independently. Rather than all potential service
1849 consumers being informed on a continual basis about all services, there is a known or agreed upon
1850 facility or location that stores and supports discovery and/or notification related to the service description.



1851

1852 *Figure 22 - Mediated Service Awareness*

1853 In Figure 22, the potential service consumers perform queries or are notified in order to locate those
 1854 services that satisfy their needs. As an example, the telephone book is a mediating registry where
 1855 individuals perform manual searches to locate services (i.e. the yellow pages). The telephone book is
 1856 also a mediated registry for solicitors to find and notify potential customers (i.e. the white pages).

1857 In mediated service awareness for large and dynamic numbers of service consumers and service
 1858 providers, the benefits of utilizing the mediator typically far outweigh the management issues associated
 1859 with it. Some of the benefits of mediated service awareness are

- 1860 • Potential service consumers have a known location for searching thereby eliminating needless
 1861 and random searches
- 1862 • Typically a consortium of interested parties (or a sufficiently large corporation) signs up to host
 1863 the mediation facility
- 1864 • Standardized tools and methods can be developed and promulgated to promote interoperability
 1865 and ease of use.

1866 However, mediated awareness can have some risks associated with it:

- 1867 • A single point of failure. If the mediation service fails then a large number of service providers and
 1868 consumers are potentially adversely affected.
- 1869 • A single point of control. If the central mediation service is owned by, or controlled by, someone
 1870 other than the service consumers and/or providers then the latter may be put at a competitive
 1871 disadvantage based on policies of the discovery provider.

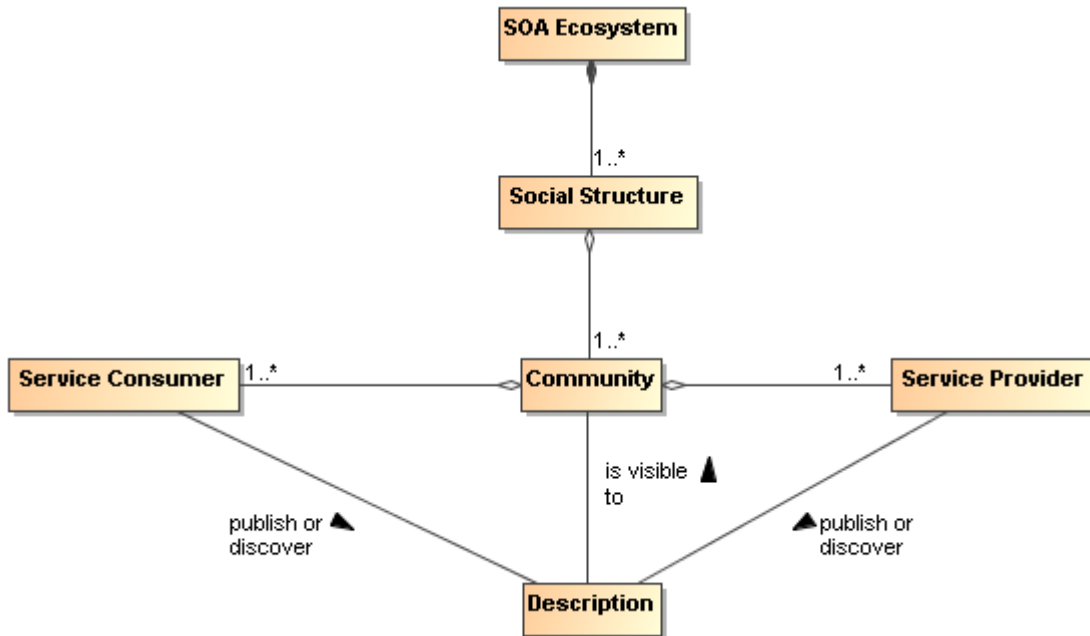
1872 A common mechanism for mediated awareness is a registry/repository. The registry stores links or
 1873 pointers to service description artifacts. The repository in this example is the storage location for the
 1874 service description artifacts. Service descriptions can be pushed (publish/subscribe for example) or pulled
 1875 from the registry/repository mediator.

1876 Registries/repositories may be referred to as federated when supported functions, such as responding to
 1877 discovery requests, are distributed across multiple registry/repository instances.

1878 **4.2.2.1.2 Awareness in Complex Social Structures**

1879 Awareness applies to one or more communities within one or more social structures where a community
 1880 consists of at least one description provider and one description consumer. These communities may be
 1881 part of the same social structure or be part of different ones.

1882 In Figure 23, awareness can be between consumers and providers within a single community, multiple
 1883 communities, or all communities in the social structure. The social structure can encourage or restrict
 1884 awareness through its policies, and these policies can affect participant willingness. The information
 1885 about policies should be incorporated in the relevant descriptions. The social structure also governs the
 1886 conditions for establishing contracts, the results of which are reflected in the execution context if
 1887 interaction is to proceed.



1888

1889 *Figure 23 - Awareness in a SOA Ecosystem*

1890 IT policy/contract mechanisms can be used by visibility mechanisms to provide awareness between
 1891 communities. The IT mechanisms for awareness may incorporate trust mechanisms to enable
 1892 awareness between trusted communities. For example, government organizations may want to limit
 1893 awareness of an organization's services to specific communities of interest.

1894 Another common business model for awareness is maximizing awareness to communities within the
 1895 social structure, the traditional market place business model. A centralized mediator often arises as a
 1896 provider for this global visibility, a gatekeeper of visibility so to speak. For example, Google is a
 1897 centralized mediator for accessing information on the web. As another example, television networks have
 1898 centralized entities providing a level of awareness to communities that otherwise could not be achieved
 1899 without going through the television network.

1900 However, mediators have motivations, and they may be selective in which information they choose to
 1901 make available to potential consumers. For example, in a secure environment, the mediator may enforce
 1902 security policies and make information selectively available depending on the security clearance of the
 1903 consumers.

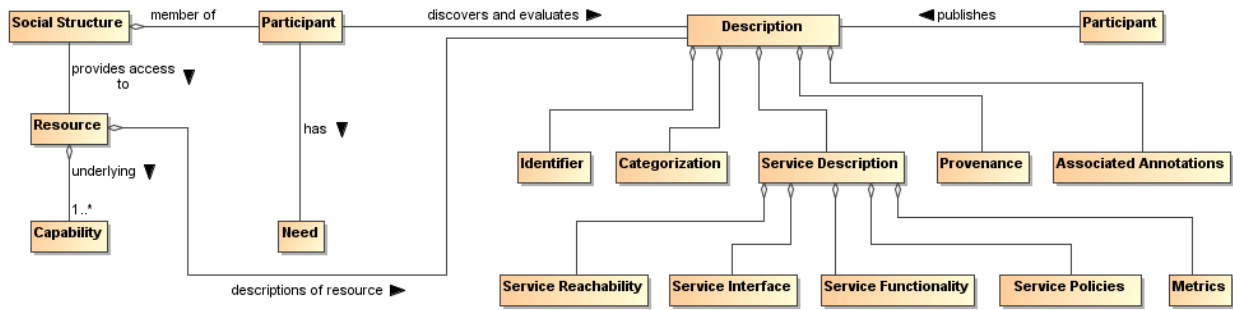
1904 **4.2.2.2 Willingness**

1905 Having achieved awareness, participants use descriptions to help determine their willingness to interact
 1906 with another participant. Both awareness and willingness are determined prior to consumer/provider
 1907 interaction.

1908

1909

1910



1911

1912 *Figure 24 - Business, Description and Willingness*

1913 Figure 24 relates elements of the *Participation in a SOA Ecosystem* view, and elements from the Service
 1914 Description Model to willingness. By having a willingness to interact within a particular social structure,
 1915 the social structure provides the participant access to capabilities based on conditions the social structure
 1916 finds appropriate for its context. The participant can use these capabilities to satisfy goals and objectives
 1917 as specified by the participant's needs.

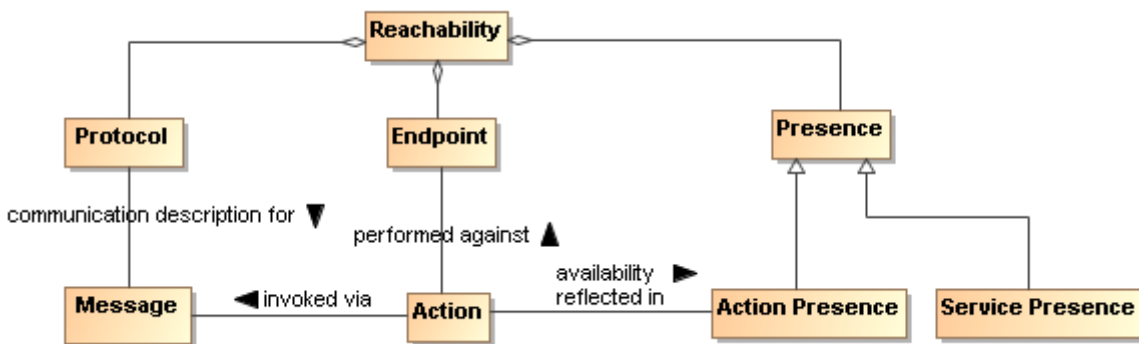
1918 In Figure 24, information used to determine willingness is defined by Description. Information referenced
 1919 by Description may come from many sources. For example, a mediator for descriptions may provide 3rd
 1920 party annotations for reputation. Another source for reputation may be a participant's own history of
 1921 interactions with another participant.

1922 A participant inspects functionality for potential satisfaction of needs. Identity is associated with any
 1923 participant, however, identity may or may not be verified. If available, participant reputation may be a
 1924 deciding factor for willingness to interact. Policies and contracts referenced by the description may be
 1925 particularly important to determine the agreements and commitments required for business interactions.
 1926 Provenance may be used for verification of authenticity of a resource.

1927 Mechanisms that aid in determining willingness make use of the artifacts referenced by descriptions of
 1928 services. Mechanisms for establishing willingness could be as simple as rendering service description
 1929 information for human consumption to automated evaluation of functionality, policies, and contracts by a
 1930 rules engine. The rules engine for determining willingness could operate as a policy decision procedure
 1931 as defined in Section 4.4.

1932 4.2.2.3 Reachability

1933 Reachability involves knowing the endpoint, protocol, and presence of a service. At a minimum,
 1934 reachability requires information about the location of the service and the protocol describing the means
 1935 of communication.



1936

1937 *Figure 25 - Service Reachability*

1938 Endpoint

1939 An endpoint is a reference-able entity, processor or resource against which an action can be
1940 performed.

1941 **Protocol**

1942 A protocol is a structured means by which details of a service interaction mechanism are defined.

1943 **Presence**

1944 Presence is the measurement of reachability of a service at a particular point in time.

1945 A protocol defines a structured method of communication. Presence is determined by interaction through
1946 a communication protocol. Presence may not be known in many cases until the interaction begins. To
1947 overcome this problem, IT mechanisms may make use of presence protocols to provide the current
1948 up/down status of a service.

1949 Service reachability enables service participants to locate and interact with one another. Each action may
1950 have its own endpoint and also its own protocols associated with the endpoint and whether there is
1951 presence for the action through that endpoint. Presence of a service is an aggregation of the presence of
1952 the service's actions, and the service level may aggregate to some degraded or restricted presence if
1953 some action presence is not confirmed. For example, if error processing actions are not available, the
1954 service can still provide required functionality if no error processing is needed. This implies reachability
1955 relates to each action as well as applying to the service/business as a whole.

1956 **4.2.3 Architectural Implications**

1957 Visibility in a SOA ecosystem has the following architectural implications on mechanisms providing
1958 support for awareness, willingness, and reachability:

- 1959 • Mechanisms providing support for awareness have the following minimum capabilities:
 - 1960 ○ creation of Description, preferably conforming to a standard Description format and
 - 1961 structure;
 - 1962 ○ publishing of Description directly to a consumer or through a third party mediator;
 - 1963 ○ discovery of Description, preferably conforming to a standard for Description discovery;
 - 1964 ○ notification of Description updates or notification of the addition of new and relevant
 - 1965 Descriptions;
 - 1966 ○ classification of Description elements according to standardized classification schemes.
- 1967 • In a SOA ecosystem with complex social structures, awareness may be provided for specific
1968 communities of interest. The architectural mechanisms for providing awareness to communities
1969 of interest require support for:
 - 1970 ○ policies that allow dynamic formation of communities of interest;
 - 1971 ○ trust that awareness can be provided for and only for specific communities of interest, the
 - 1972 bases of which is typically built on keying and encryption technology.
- 1973 • The architectural mechanisms for determining willingness to interact require support for:
 - 1974 ○ verification of identity and credentials of the provider and/or consumer;
 - 1975 ○ access to and understanding of description;
 - 1976 ○ inspection of functionality and capabilities;
 - 1977 ○ inspection of policies and/or contracts.
- 1978 • The architectural mechanisms for establishing reachability require support for:
 - 1979 ○ the location or address of an endpoint;
 - 1980 ○ verification and use of a service interface by means of a communication protocol;
 - 1981 ○ determination of presence with an endpoint which may only be determined at the point of
 - 1982 interaction but may be further aided by the use of a presence protocol for which the
 - 1983 endpoints actively participate.

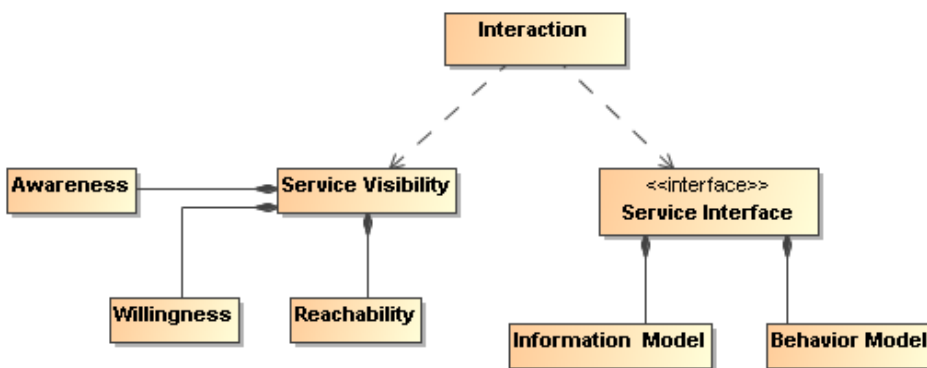
1984 **4.3 Interacting with Services Model**

1985 Interaction is the activity involved in using a service to access capability in order to achieve a particular
1986 desired real world effect, where real world effect is the actual result of using a service. An interaction can
1987 be characterized by a sequence of communicative actions. Consequently, interacting with a service, i.e.
1988 participating in joint action with the service—usually mediated by a series of message exchanges—

1989 involves individual actions performed by both the service and the consumer.⁸ Note that a participant (or
1990 delegate acting on behalf of the participant) can be the sender of a message, the receiver of a message,
1991 or both.

1992 4.3.1 Interaction Dependencies

1993 Recall from the Reference Model that service visibility is the capacity for those with needs and those with
1994 capabilities to be able to interact with each other, and that the service interface is the means by which the
1995 underlying capabilities of a service are accessed. Ideally, the details of the underlying service
1996 implementation are abstracted away by the service interface. [Service] interaction therefore has a direct
1997 dependency on the visibility of the service as well as its implementation-neutral interface (see Figure 26).
1998 Service visibility is composed of awareness, willingness, and reachability and service interface is
1999 composed of the information and behavior models. Service visibility is modeled in Section 4.2 while
2000 service interface is modeled in Section 4.1.

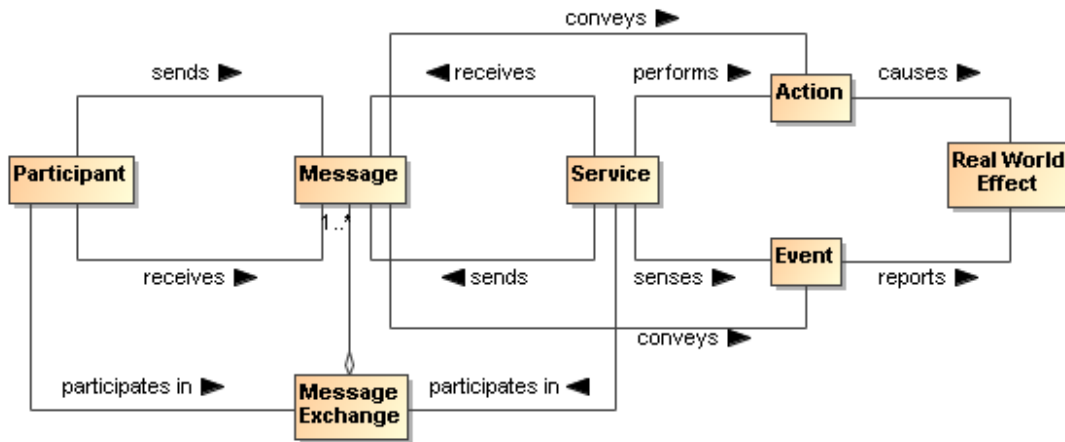


2001
2002 *Figure 26 - Interaction dependencies*

2003 4.3.2 Actions and Events

2004 The SOA-RAF uses message exchange between service participants to denote actions performed
2005 against and by the service, and to denote events that report on real world effects that are caused by the
2006 service actions. A visual model of the relationship between these concepts is shown in Figure 27.

⁸ In order for multiple actors to participate in a joint action, they must each act according to their role within the joint action. For SOA-based systems, this is achieved through a message exchange style of communication. The concept of “joint action” is further described in Section **Error! Reference source not found.**



2007
2008 *Figure 27 - A "message" denotes either an action or an event*

2009 Both actions and **events**, realized by the SOA services, are denoted by the messages. The Reference
2010 Model states that the action model characterizes the “permissible set of actions that may be invoked
2011 against a service.” We extend that notion here to include **events** as part of the event model and that
2012 messages are intended for invoking actions or for notification of **events**.

2013 In Section 3.2.3 we saw that **participants** interact with each other in order to participate in joint actions. A
2014 joint **action** is not itself the same thing as the result of the joint **action**. When a joint **action** is participated in
2015 with a service, the **real world effect** that results may be reported in the form of an event notification.

2016 **4.3.3 Message Exchange**

2017 *Message exchange* is the means by which service **participants** (or their **delegates**) interact with each
2018 other. There are two primary modes of interaction: joint actions that cause **real world effects** and
2019 notification of **events** that report real world effects.⁹

2020 A message exchange is used to affect an **action** when the messages contain the appropriately formatted
2021 content, are directed towards a particular action in accordance with the action model, and the **delegates**
2022 involved interpret the message appropriately.

2023 A message exchange is also used to communicate **event** notifications. An **event** is an occurrence that is
2024 of interest to some **participant**; in our case when some **real world effect** has occurred. Just as action
2025 messages have formatting requirements, so do event notification messages. In this way, the Information
2026 Model of a service must specify the syntax (structure), and semantics (meaning) of the action messages
2027 and event notification messages as part of a service interface. It must also specify the syntax and
2028 semantics of any data that is carried as part of a payload of the action or event notification message. The
2029 Information Model is described in greater detail in the Service Description Model (see Section 4.1).

2030 In addition to the Information Model that describes the syntax and semantics of the messages and data
2031 payloads, exception conditions and error handling in the event of faults (e.g., network outages, improper
2032 message formats, etc.) must be specified or referenced as part of the Service Description.

⁹ The notion of “joint” in joint action implies that you have to have a speaker *and* a listener in order to interact.

2033 When a message is used to invoke an **action**, the correct interpretation typically requires the receiver to
 2034 perform an operation, which itself invokes a set of private, internal actions. These *operations* represent
 2035 the sequence of (private) actions a service must perform in order to validly participate in a given joint
 2036 action.

2037 Similarly, the correct consequence of realizing a **real world effect** may be to initiate the reporting of that
 2038 real world effect via an event notification.

2039 **Message Exchange**

2040 The means by which joint action and event notifications are coordinated by service **participants**
 2041 (or **delegates**).

2042 **Operations**


2043 The sequence of actions a service must perform in order to validly participate in a given joint
 2044 action.

2045 **4.3.3.1 Message Exchange Patterns (MEPs)**

2046 The basic temporal aspect of service interaction can be characterized by two fundamental message
 2047 exchange patterns (MEPs):

- 2048 • Request/response to represent how actions cause a **real world effect**
- 2049 • Event notification to represent how **events** report a **real world effect**

2050 This is by no means a complete list of all possible MEPs used for inter- or intra-enterprise messaging but
 2051 it does represent those that are most commonly used in exchange of information and reporting changes
 2052 in state both within organizations and across organizational boundaries.



2053
 2054 *Figure 28 - Fundamental SOA message exchange patterns (MEPs)*

2055 Recall from the Reference Model that the Process Model characterizes “the temporal relationships
2056 between and temporal properties of actions and **events** associated with interacting with the service.”
2057 Thus, MEPs are a key element of the Process Model. The meta-level aspects of the Process Model (just
2058 as with the Action Model) are provided as part of the Service Description Model (see Section 4.1).

2059 In the UML sequence diagram shown in Figure 28 it is assumed that the service **participants** (consumer
2060 and provider) have delegated message handling to hardware or software delegates acting on their behalf.
2061 In the case of the **service consumer**, this is represented by the *Consumer Delegate* component. In the
2062 case of the service provider, the **delegate** is represented by the *Service* component. The message
2063 interchange model illustrated represents a logical view of the MEPs and not a physical view. In other
2064 words, specific hosts, network protocols, and underlying messaging system are not shown as these tend
2065 to be implementation specific. Although such implementation-specific elements are considered outside
2066 the scope of this document, they are important considerations in modeling the SOA execution context.
2067 Recall from the Reference Model that the *execution context* of a service interaction is “the set of
2068 infrastructure elements, process entities, policy assertions and agreements that are identified as part of
2069 an instantiated service interaction, and thus forms a path between those with needs and those with
2070 capabilities.”

2071 **4.3.3.2 Request/Response MEP**

2072 In a request/response MEP, the Consumer Delegate component sends a request message to the Service
2073 component. The Service component then processes the request message. Based on the content of the
2074 message, the Service component performs the service operation and the associated private actions.
2075 Following the completion of these operations, a response message is returned to the Consumer Delegate
2076 component. The response could be that a step in a process is complete, the initiation of a follow-on
2077 operation, or the return of requested information.¹⁰

2078 Although the sequence diagram shows a *synchronous* interaction (because the sender of the request
2079 message, i.e., Consumer Delegate, is blocked from continued processing until a response is returned
2080 from the Service) other variations of request/response are valid, including *asynchronous* (non-blocking)
2081 interaction through use of queues, channels, or other messaging techniques.

2082 What is important to convey here is that the request/response MEP represents **action**, which causes a
2083 **real world effect**, irrespective of the underlying messaging techniques and messaging infrastructure used
2084 to implement the request/response MEP.

2085 **4.3.3.3 Event Notification MEP**

2086 An **event** is made visible to interested consumers by means of an event notification message exchange
2087 that reports a **real world effect**; specifically, a change in shared state between service **participants**. The
2088 basic event notification MEP takes the form of a one-way message sent by a notifier component (in this
2089 case, the Service component) and received by components with an interest in the **event** (here, the
2090 Consumer Delegate component).

2091 Often the sending component may not be fully aware of all the components that wish to receive the
2092 notification; particularly in so-called publish/subscribe (“pub/sub”) situations. In event notification


¹⁰ 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.

2093 message exchanges, it is rare to have a tightly-coupled link between the sending and the receiving
2094 component(s) for a number of practical reasons. One of the most common needs for pub/sub messaging
2095 is the potential for network outages or communication interrupts that can result in loss of notification of
2096 events. Therefore, a third-party mediator component is often used to decouple the sending and receiving
2097 components.

2098 Although this is typically an implementation issue, because this type of third-party decoupling is so
2099 common in event-driven systems, it is warranted for use in modeling this type of message exchange in
2100 the SOA-RAF. This third-party intermediary is shown in Figure 28 as an Event Broker mediator. As with
2101 the request/response MEP, no distinction is made between synchronous versus asynchronous
2102 communication, although asynchronous message exchange is illustrated in the UML sequence diagram
2103 depicted in Figure 28 .

2104 4.3.4 Composition of Services

2105 Composition of services is the act of aggregating or “composing” a single service from one or more other
2106 services. A simple model of service composition is illustrated in Figure 29.



2107
2108 *Figure 29 - Simple model of service composition*

2109 Here, Service A is a service that has an exposed interface IServiceA, which is available to the Consumer
2110 Delegate and relies on two other services in its implementation. The Consumer Delegate does not know
2111 that Services B and C are used by Service A, or whether they are used in serial or parallel, or if their
2112 operations succeed or fail. The Consumer Delegate only cares about the success or failure of Service A.
2113 The exposed interfaces of Services B and C (IService B and IServiceC) are not necessarily hidden from
2114 the Consumer Delegate; only the fact that these services are used as part of the composition of Service
2115 A. In this example, there is no practical reason the Consumer Delegate could not interact with Service B
2116 or Service C in some other interaction scenario.

2117 It is possible for a service composition to be opaque from one perspective and transparent from another.
2118 For example, a service may appear to be a single service from the Consumer’s Delegate’s perspective,
2119 but is transparently composed of one or more services from a service management perspective. A
2120 Service Management capability needs to be able to have visibility into the composition in order to properly
2121 manage the dependencies between the services used in constructing the composite service—including
2122 managing the service’s lifecycle. The subject of services as management entities is described and
2123 modeled in the *Ownership in a SOA Ecosystem* View of the SOA-RAF and is not further elaborated in this
2124 section. The point to be made here is that there can be different levels of opaqueness or transparency
2125 when it comes to visibility of service composition.

2126 Services can be composed in a variety of ways including direct consumer-to-service interaction by using
2127 programming techniques, or they can be aggregated by means of an aggregation engine approach that
2128 leverages a service composition scripting language. Such approaches are further elaborated in the
2129 following sub-sections on service-oriented business processes and collaborations.

2130 4.3.4.1 Service-Oriented Business Processes

2131 The concepts of business processes and collaborations in the context of transactions and exchanges
2132 across organizational boundaries are described and modeled as part of the *Participation in a SOA*
2133 *Ecosystem* view of this reference architecture (see Section 3). Here, we focus on the belief that the
2134 principle of composition of services can be applied to business processes and collaborations. Of course,

2135 business processes and collaborations traditionally represent complex, multi-step business functions that
2136 may involve multiple [participants](#), including internal users, external customers, and trading partners.
2137 Therefore, such complexities cannot simply be ignored when transforming traditional business processes
2138 and collaborations to their service-oriented variants.

2139 **Business Processes**

2140 Business processes are a set of one or more linked activities that are performed to achieve a
2141 certain business outcome.

2142 Service orientation as applied to business processes (i.e., “service-oriented business processes”) means
2143 that the aggregation or composition of all of the abstracted activities, flows, and rules that govern a
2144 business process can themselves be abstracted as a service **[BLOOMBERG/SCHMELZER]**.

2145 When business processes are abstracted in this manner and accessed through SOA services, all of the
2146 concepts used to describe and model composition of services that were articulated in Section 4.3.4 apply.
2147 There are some important differences between a composite service that represents an abstraction of a
2148 business process and a composite service that represents a single-step business interaction. Business
2149 processes have temporal properties and can range from short-lived processes that execute on the order
2150 of minutes or hours to long-lived processes that can execute for weeks, months, or even years. Further,
2151 these processes may involve many [participants](#). These are important considerations for the consumer of
2152 a service-oriented business process and these temporal properties must be articulated as part of the
2153 meta-level aspects of the service-oriented business process in its Service Description, along with the
2154 meta-level aspects of any sub-processes that may be of use or need to be visible to the [service](#)
2155 [consumer](#).

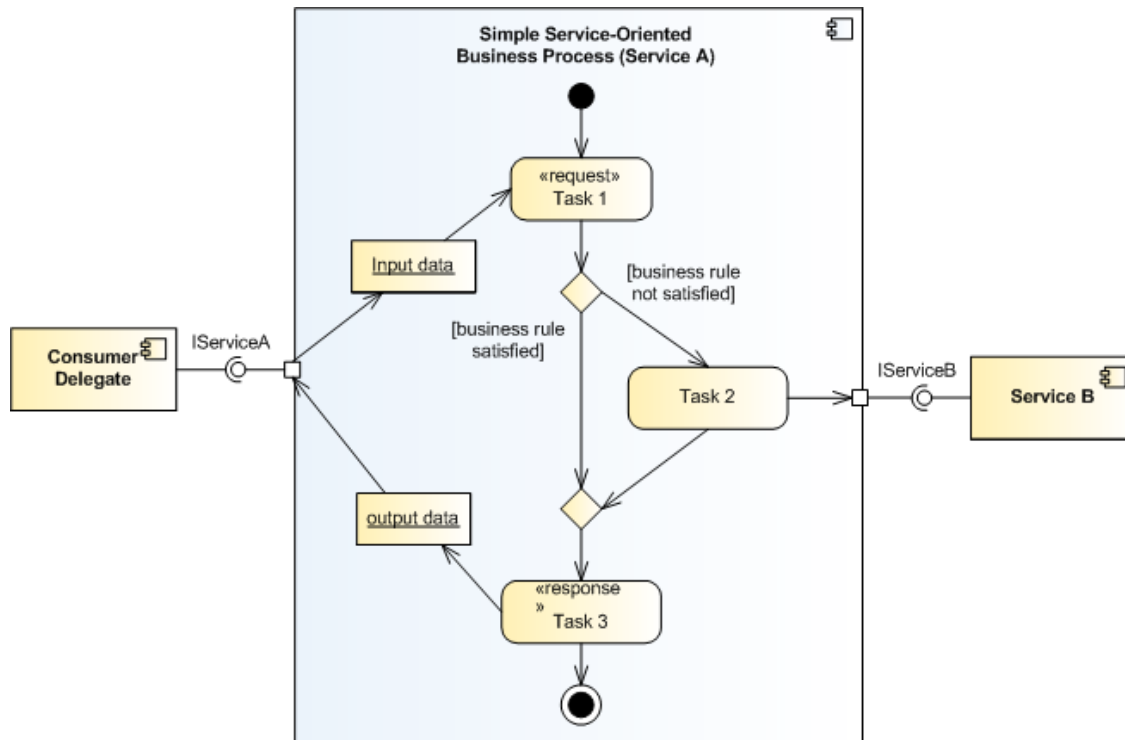
2156 In addition, a workflow activity represents a unit of work that some actor acting in a described [role](#) (i.e.,
2157 [role player](#)) is asked to perform. Activities can be broken down into steps with each step representing a
2158 task for the [role player](#) to perform. A technique that is used to compose service-oriented business
2159 processes that are hierarchical (top-down) and self-contained in nature is known as *orchestration*.

2160 **Orchestration**

2161 A technique used to compose service-oriented business processes that are executed and
2162 coordinated by an [actor](#) acting as “conductor.”

2163 An orchestration is typically implemented using a scripting approach to compose service-oriented
2164 business processes. This typically involves use of a standards-based orchestration scripting language.
2165 In terms of automation, an orchestration can be mechanized using a business process orchestration
2166 engine, which is a hardware or software component ([delegate](#)) responsible for acting in the role of central
2167 conductor/coordinator responsible for executing the flows that comprise the orchestration.

2168 A simple generic example of such an orchestration is illustrated in Figure 30.



2169

2170 *Figure 30 - Abstract example of orchestration of service-oriented business process*

2171 Here, we use a UML activity diagram to model the simple service-oriented business process as it allows
 2172 us to capture the major elements of business processes such as the set of related tasks to be performed,
 2173 linking between tasks in a logical flow, data that is passed between tasks, and any relevant business
 2174 rules that govern the transitions between tasks. A task is a unit of work that an individual, system, or
 2175 organization performs and can be accomplished in one or more steps or subtasks. While subtasks can
 2176 be readily modeled, they are not illustrated in the orchestration model in Figure 30.

2177 This particular example is based on a request/response MEP and captures how one particular task (Task
 2178 2) actually utilizes an externally-provided service, Service B. The entire service-oriented business
 2179 process is exposed as Service A that is accessible via its externally visible interface, IServiceA.

2180 Although not explicitly shown in the orchestration model above, it is assumed that there exists a software
 2181 or hardware component, i.e., orchestration engine that executes the process flow. Recall that a central
 2182 concept to orchestration is that process flow is coordinated and executed by a single conductor delegate;
 2183 hence the name “orchestration.”

2184 **4.3.4.2 Service-Oriented Business Collaborations**

2185 Business collaborations typically represent the interaction involved in executing [business transactions](#).

2186 It is important to note that business collaborations represent “peer”-style interactions; in other words,
 2187 [peers](#) in a business collaboration act as equals. This means that unlike the orchestration of business
 2188 processes, there is no single or central entity that coordinates or “conducts” a business collaboration.
 2189 These peer styles of interactions typically occur between trading partners that span organizational
 2190 boundaries.

2191 Business collaborations can also be service-enabled. For purposes of this Reference Architecture
 2192 Foundation, we refer to these as “service-oriented business collaborations.” Service-oriented business
 2193 collaborations do not necessarily imply exposing the entire peer-style business collaboration as a service
 2194 itself but rather the collaboration uses service-based interchanges.

2195 The technique that is used to compose service-oriented business collaborations in which multiple parties
 2196 collaborate in a peer-style as part of some larger [business transaction](#) by exchanging messages with

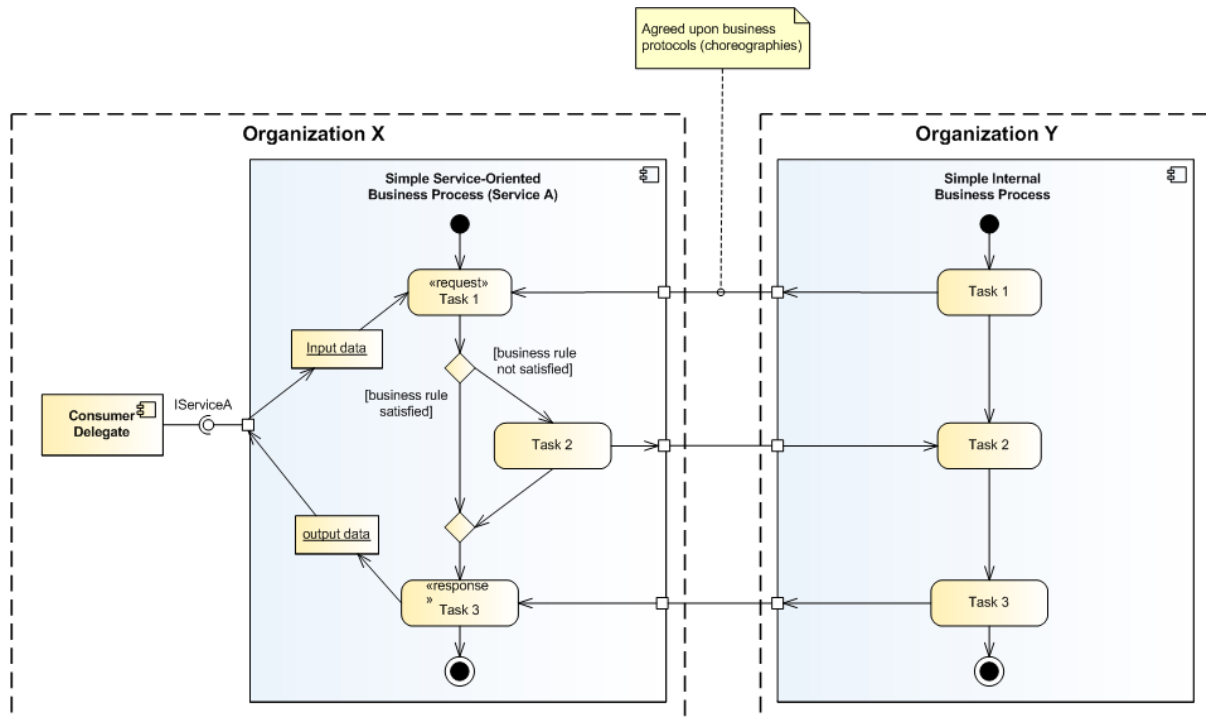
2197 trading partners and external organizations (e.g., suppliers) is known as *choreography*
2198 **[NEWCOMER/LOMOW]**.

2199 **Choreography**

2200 A technique used to characterize service-oriented business collaborations based on ordered
2201 message exchanges between *peer* entities in order to achieve a common business goal.

2202 Choreography differs from orchestration primarily in that each party in a business collaboration describes
2203 its part in the service interaction. Note that choreography as we have defined it here should not be
2204 confused with the term *process choreography*, which is defined in the *Participation in a SOA Ecosystem*
2205 view as “the description of the possible interactions that may take place between two or more *participants*
2206 to fulfill an objective.” This is an example of domain-specific nomenclature that often leads to confusion
2207 and why we are making note of it here.

2208 A simple generic example of a choreography is illustrated in Figure 31.



2209
2210 *Figure 31 - Abstract example of choreography of service-oriented business collaboration*

2211 This example, which is a variant of the orchestration example illustrated earlier in Figure 30 adds trust
2212 boundaries between two organizations; namely, Organization X and Organization Y. It is assumed that
2213 these two organizations are peer entities that have an interest in a business collaboration, for example,
2214 Organization X and Organization Y could be trading partners. Organization X retains the service-oriented
2215 business process Service A, which is exposed to internal consumers via its provided service interface,
2216 IServiceA. Organization Y also has a business process that is involved in the business collaboration;
2217 however, for this example, it is an internal business process that is not exposed to potential consumers
2218 either within or outside its organizational boundary.

2219 The scripting language that is used for the choreography needs to define how and when to pass control
2220 from one trading partner to another, i.e., between Organization X and Organization Y. Defining the
2221 business protocols used in the business collaboration involves precisely specifying the visible message
2222 exchange behavior of each of the parties involved in the protocol, without revealing internal
2223 implementation details **[NEWCOMER/LOMOW]**.

2224 In a peer-style business collaboration, a choreography scripting language must be capable of describing
2225 the coordination of those service-oriented processes that cross organizational boundaries.

2226 4.3.5 Architectural Implications of Interacting with Services

2227 Interacting with Services has the following architectural implications on mechanisms that facilitate service
2228 interaction:

- 2229 • A well-defined service Information Model that:
 - 2230 ○ describes the syntax and semantics of the messages used to denote actions and **events**;
 - 2231 ○ describes the syntax and semantics of the data payload(s) contained within messages;
 - 2232 ○ documents exception conditions in the event of faults due to network outages, improper
2233 message/data formats, etc.;
 - 2234 ○ is both human readable and machine processable;
 - 2235 ○ is referenceable from the Service Description artifact.
- 2236 • A well-defined service Behavior Model that:
 - 2237 ○ characterizes the knowledge of the actions invokes against the service and **events** that
2238 report **real world effects** as a result of those actions;
 - 2239 ○ characterizes the temporal relationships and temporal properties of actions and **events**
2240 associated in a service interaction;
 - 2241 ○ describe activities involved in a workflow activity that represents a unit of work;
 - 2242 ○ describes the **role** (s) that a **role player** performs in a service-oriented business process
2243 or service-oriented business collaboration;
 - 2244 ○ is both human readable and machine processable;
 - 2245 ○ is referenceable from the Service Description artifact.
- 2246 • Service composition mechanisms to support orchestration of service-oriented business processes and
2247 choreography of service-oriented business collaborations such as:
 - 2248 ○ Declarative and programmatic compositional languages;
 - 2249 ○ Orchestration and/or choreography engines that support multi-step processes as part of a
2250 short-lived or long-lived **business transaction**;
 - 2251 ○ Orchestration and/or choreography engines that support compensating transactions in
2252 the presences of exception and fault conditions.
- 2253 • Infrastructure services that provides mechanisms to support service interaction, including but not
2254 limited to:
 - 2255 ○ mediation services such as message and event brokers, providers, and/or buses that
2256 provide message translation/transformation, gateway capability, message persistence,
2257 reliable message delivery, and/or intelligent routing semantics;
 - 2258 ○ binding services that support translation and transformation of multiple application-level
2259 protocols to standard network transport protocols;
 - 2260 ○ auditing and logging services that provide a data store and mechanism to record
2261 information related to service interaction activity such as message traffic patterns,
2262 security violations, and service contract and **policy** violations
 - 2263 ○ security services that provide centralized authorization and authentication support, etc.,
2264 which provide protection against common security threats in a SOA ecosystem;
 - 2265 ○ monitoring services such as hardware and software mechanisms that both monitor the
2266 performance of systems that host services and network traffic during service interaction,
2267 and are capable of generating regular monitoring reports.
- 2268 • A layered and tiered service component architecture that supports multiple message exchange
2269 patterns (MEPs) in order to:
 - 2270 ○ promote the industry best practice of separation of concerns that facilitates flexibility in
2271 the presence of changing business requirements;
 - 2272 ○ promote the industry best practice of separation of **roles** in a service development
2273 lifecycle such that subject matter experts and teams are structured along areas of
2274 expertise;
 - 2275 ○ support numerous standard interaction patterns, peer-to-peer interaction patterns,
2276 enterprise integration patterns, and business-to-business integration patterns.

2277 **4.4 Policies and Contracts Model**

2278 A common phenomenon of many machines and systems is that the scope of potential behavior is much
2279 broader than is actually needed for a particular circumstance. This is especially true of a system as
2280 powerful as a SOA ecosystem. As a result, the behavior and performance of the system tend to be
2281 under-constrained by the implementation; instead, the actual behavior is expressed by means of policies
2282 of some form. Policies define the choices that **stakeholders** make; these choices are used to guide the
2283 actual behavior of the system to the desired behavior and performance.

2284 As noted in Section 3.1.5, a **policy** is a constraint of some form that is promulgated by a **stakeholder** who
2285 has the responsibility of ensuring that the constraint is enforced. In contrast, **contracts** are **agreements**
2286 between **participants**. However, like policies, it is a necessary part of **contracts** that they are enforceable.

2287 While responsibility for enforcement may differ, both **contracts** and policies share a common characteristic
2288 – there is a **constraint** that must be enforced. In both cases the mechanisms needed to enforce
2289 constraints are likely to be identical; in this model we focus on the issues involved in representing policies
2290 and **contracts** and on some of the principles behind their enforcement.

2291 **4.4.1 Policy and Contract Representation**

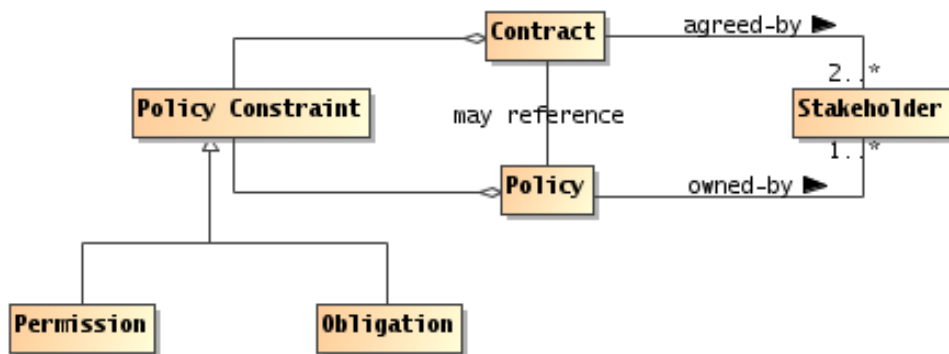
2292 A **policy constraint** is a specific kind of constraint: the ontology of policies and **contracts** includes the
2293 core concepts of **permission**, **obligation**, owner, subject. In addition, it may be necessary to be able
2294 combine policy constraints and to be able to resolve policy conflicts.

2295 **4.4.1.1 Policy Framework**

2296 **Policy Framework**

2297 A policy framework is a language in which policy constraints may be expressed.

2298 A policy framework combines a syntax for expressing policy constraints together with a **decision**
2299 **procedure** for determining if a policy constraint is satisfied.



2300
2301 *Figure 32 - Policies and Contracts*

2302 We can characterize (caricature) a policy framework in terms of a logical framework and an ontology of
2303 policies. The policy ontology details specific kinds of policy constraints that can be expressed; and the
2304 logical framework is a 'glue' that allows us to express combinations of policies.

2305 **Logical Framework**

2306 A logical framework is a linguistic framework consisting of a syntax – a way of writing expressions
2307 – and a semantics – a way of interpreting the expressions.

2308 **Policy Ontology**

2309 A policy ontology is a formalization of a set of concepts that are relevant to forming policy
2310 expressions.

2311 For example, a policy ontology that allows to identify simple constraints – such as the existence of a
2312 property, or that a value of a property should be compared to a fixed value – is often enough to express
2313 many basic constraints.

2314 Included in many policy ontologies are the basic signals of [permissions](#) and [obligations](#). Some policy
2315 frameworks are sufficiently constrained that there is not possibility of representing an [obligation](#); in which
2316 case there is often no need to ‘call out’ the distinction between [permissions](#) and [obligations](#).

2317 The logical framework is also a strong determiner of the expressivity of the policy framework: the richer
2318 the logical framework, the richer the set of policy constraints that can be expressed. However, there is a
2319 strong inverse correlation between expressivity and ease and efficiency of implementation.

2320 In the discussion that follows we assume the following basic policy ontology:

2321 **Policy Owner**

2322 A policy owner is a [stakeholder](#) that asserts and enforces the policy.

2323 **Policy Subject**

2324 A policy subject is an [actor](#) who is subject to the constraints of a policy or contract.

2325 **Policy Constraint**

2326 A policy constraint is a measurable and enforceable proposition that characterizes the constraint
2327 that the policy is about.

2328 **Policy Object**

2329 A policy object is an identifiable state, [action](#) or [resource](#) that is potentially constrained by the
2330 policy.

2331 **4.4.2 Policy and Contract Enforcement**

2332 The enforcement of policy constraints has to address two core problems: how to enforce the atomic policy
2333 constraints, and how to enforce combinations of policy constraints. In addition, it is necessary to address
2334 the resolution of policy conflicts.

2335 **4.4.2.1 Enforcing Simple Policy Constraints**

2336 The two primary kinds of policy constraint – [permission](#) and obligation – naturally lead to different styles
2337 of enforcement. A [permission](#) constraint must typically be enforced *prior* to the policy subject invoking the
2338 **policy object**. On the hand, an obligation constraint must typically be enforced post-facto through some
2339 form of auditing process and remedial action.

2340 For example, if a communications policy required that all communication be encrypted, this is enforceable
2341 at the point of communication: any attempt to communicate a message that is not encrypted can be
2342 blocked.

2343 Similarly, an obligation to pay for services rendered is enforced by ensuring that payment arrives within a
2344 reasonable period of time. Invoices are monitored for prompt (or lack of) payment.

2345 The key concepts in enforcing both forms of policy constraint are the policy decision and the policy
2346 enforcement.

2347 **Policy Decision**

2348 A policy decision is a determination as to whether a given policy constraint is satisfied or not.

2349 A policy decision is effectively a measurement of some state – typically a portion of the SOA ecosystem’s
2350 **shared state**. This implies a certain *timeliness* in the measuring: a measurement that is too early or is too
2351 late does not actually help in determining if the policy constraint is satisfied appropriately.

2352 **Policy Enforcement**

2353 A policy enforcement is the use of a mechanism which limits the behavior and/or state of policy
2354 subjects to comply with a policy decision.

2355 A policy enforcement implies the use of some mechanism to ensure compliance with a policy decision.
2356 The range of mechanisms is completely dependent on the kinds of atomic policy constraints that the

2357 policy framework may support. As noted above, the two primary styles of constraint – **permission** and
2358 **obligation** –lead to different styles of enforcement.

2359 4.4.2.2 Conflict Resolution

2360 Whenever it is possible that more than one policy constraint applies in a given situation, there is the
2361 potential that the policy constraints themselves are not mutually consistent. For example, a policy
2362 constraint that requires communication to be encrypted and a policy constraint that requires an
2363 administrator to read every communication conflict with each other – the two policy constraints cannot
2364 both be satisfied concurrently.

2365 In general, with sufficiently rich policy frameworks, it is not possible to always resolve policy conflicts
2366 automatically. However, a reasonable approach is to augment the policy decision process with simple
2367 policy conflict resolution rules; with the potential for *escalating* a policy conflict to human adjudication.

2368 Policy Conflict

2369 A policy conflict exists between two or more policy constraints in a policy decision process if the
2370 satisfaction of one or more policy constraints leads directly to the violation of one or more other
2371 policy constraints.

2372 Policy Conflict Resolution

2373 A policy conflict resolution rule is a way of determining which policy constraints should prevail if a
2374 policy conflict occurs.

2375 The inevitable consequence of policy conflicts is that it is not possible to guarantee that all policy
2376 constraints are satisfied at all times. This, in turn, implies a certain *flexibility* in the application of policy
2377 constraints: each individual constraint may not always be honored.

2378 4.4.3 Architectural Implications

2379 The key choices that must be made in a system of policies center on the policy framework, policy
2380 enforcement, and conflict resolution

- 2381 • There SHOULD be a standard policy framework that is adopted across ownership domains within the
2382 SOA ecosystem:
 - 2383 ○ This framework MUST permit the expression of simple policy constraints
 - 2384 ○ The framework MAY allow (to a varying extent) the combination of policy constraints,
2385 including
 - 2386 • Both positive and negative constraints
 - 2387 • Conjunctions and disjunctions of constraints
 - 2388 • The quantification of constraints
 - 2389 ○ The framework MUST at least allow the policy subject and the policy object to be identified as
2390 well as the policy constraint.
 - 2391 ○ The framework MAY allow further structuring of policies into modules, inheritance between
2392 policies and so on.
- 2393 • There SHOULD be mechanisms that facilitate the application of policies:
 - 2394 ○ There SHOULD be mechanisms that allow policy decisions to be made, consistent with the
2395 policy frameworks.
 - 2396 ○ There SHOULD be mechanisms to enforce policy decisions
 - 2397 • There SHOULD be mechanisms to support the measurement of whether certain
2398 policy constraints are satisfied or not, or to what degree they are satisfied.
 - 2399 • Such enforcement mechanisms MAY include support for both **permission**-style
2400 constraints and obligation-style constraints.
 - 2401 • Enforcement mechanisms MAY support the simultaneous enforcement of multiple
2402 policy constraints across multiple points in the SOA ecosystem.
 - 2403 ○ There SHOULD be mechanisms to resolve policy conflicts
 - 2404 • This MAY involve escalating policy conflicts to human adjudication.
 - 2405 ○ There SHOULD be mechanisms that support the management and promulgation of policies.

5 Ownership in a SOA Ecosystem View

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

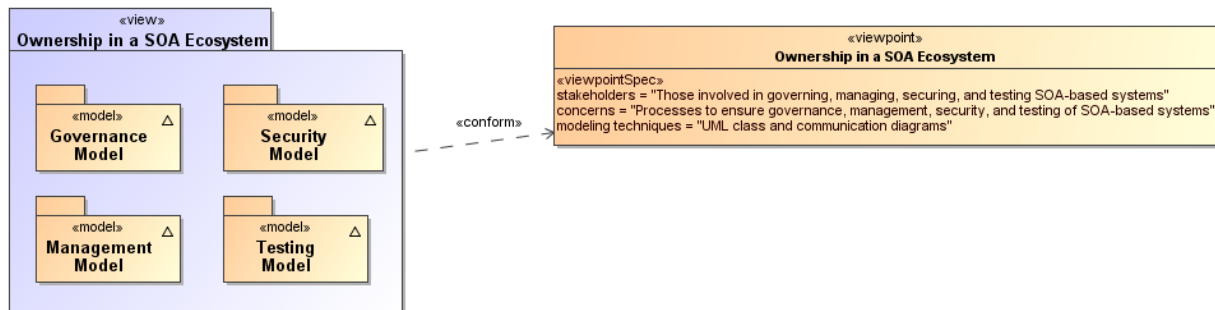
2406
2407
2408
2409
2410
2411
2412
2413
2414
2415
2416
2417
2418
2419
2420
2421
2422

The *Ownership in a SOA Ecosystem View* focuses on the issues, requirements and responsibilities involved in owning a SOA-based system.

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

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

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



2423
2424
2425

Figure 33 - Model Elements Described in the Ownership in a SOA Ecosystem View

The following subsections present models of these functions.

5.1 Governance Model

2426
2427
2428
2429
2430
2431

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

5.1.1 Understanding Governance

2432
2433

5.1.1.1 Terminology

2434
2435
2436
2437
2438
2439

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

2440 To accomplish this, governance requires organizational structure and processes and must identify who
2441 has authority to define and carry out its mandates. It must address the following questions:

- 2442 1. what decisions must be made to ensure effective management and use?,
- 2443 2. who should make these decisions?,
- 2444 3. how will these decisions be made and monitored? , and
- 2445 4. how will these decisions be communicated?

2446 The intent is to achieve goals, add value, and reduce risk.

2447 Within a single ownership domain such as an enterprise, generally there is a hierarchy of governance
2448 structures. Some of the more common enterprise governance structures include corporate governance,
2449 technology governance, IT governance, and architecture governance [TOGAF v8.1]. These governance
2450 structures can exist at multiple levels (global, regional, and local) within the overall enterprise.

2451 It is often asserted that SOA governance is a specialization of IT governance as there is a natural
2452 hierarchy of these types of governance structures; however, the focus of SOA governance is less on
2453 decisions to ensure effective management and use of IT as it is to ensure effective management and use
2454 of SOA-based systems. Certainly, SOA governance must still answer the basic questions also
2455 associated with IT governance, i.e., who should make the decisions, and how these decisions will be
2456 made and monitored.

2457 5.1.1.2 Relationship to Management

2458 There is often confusion centered on the relationship between governance and management. As
2459 described earlier, governance is concerned with decision making. Management, on the other hand, is
2460 concerned with execution. Put another way, governance describes the world as leadership wants it to be;
2461 management executes activities that intends to make the leadership's desired world a reality. Where
2462 governance determines who has the authority and responsibility for making decisions and the
2463 establishment of guidelines for how those decisions should be made, management is the actual process
2464 of making, implementing, and measuring the impact of those decisions [Loeb]. Consequently,
2465 governance and management work in concert to ensure a well-balanced and functioning organization as
2466 well as an ecosystem of inter-related organizations. In the sections that follow, we elaborate further on
2467 the relationship between governance and management in terms of setting and enforcing service policies,
2468 [contracts](#), and standards as well as addressing issues surrounding regulatory compliance.

2469 5.1.1.3 Why is SOA Governance Important?

2470 One of the hallmarks of SOA that distinguishes it from other architectural paradigms for distributed
2471 computing is the ability to provide a uniform means to offer, discover, interact with and use capabilities
2472 (as well the ability to compose new capabilities from existing ones) all in an environment that transcends
2473 domains of [ownership](#). Consequently, [ownership](#), and issues surrounding it, such as obtaining
2474 acceptable terms and conditions (T&Cs) in a contract, is one of the primary topics for SOA governance.
2475 Generally, IT governance does not include T&Cs, for example, as a condition of use as its primary
2476 concern.

2477 Just as other architectural paradigms, technologies, and approaches to IT are subject to change and
2478 evolution, so too is SOA. Setting policies that allow change management and evolution, establishing
2479 strategies for change, resolving disputes that arise, and ensuring that SOA-based systems continue to
2480 fulfill the goals of the business are all reasons why governance is important to SOA.

2481 5.1.1.4 Governance Stakeholders and Concerns

2482 As noted in Section 3.1.1 the [participants](#) in a service interaction include the service provider, the [service](#)
2483 [consumer](#), and other interested or unintentional third parties. Depending on the circumstances, it may
2484 also include the owners of the underlying capabilities that the SOA services access. Governance must
2485 establish the policies and rules under which duties and [responsibilities](#) are defined and the expectations
2486 of [participants](#) are grounded. The expectations include transparency in aspects where transparency is
2487 mandated, trust in the impartial and consistent application of governance, and assurance of reliable and
2488 robust behavior throughout the SOA ecosystem.

2489 **5.1.2 A Generic Model for Governance**

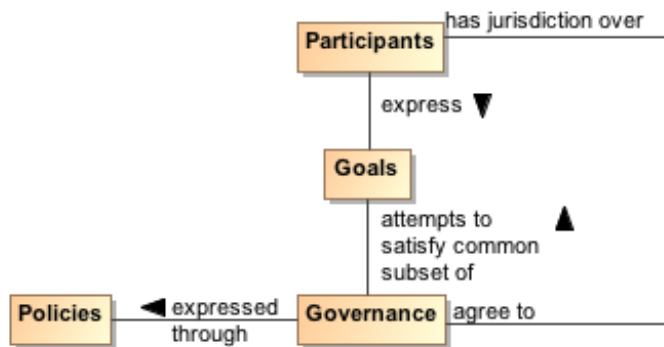
2490 **Governance**

2491 Governance is the prescribing of conditions and constraints consistent with satisfying common
2492 goals and the structures and processes needed to define and respond to actions taken towards
2493 realizing those goals.

2494 The following is a generic model of governance represented by segmented models that begin with
2495 motivation and proceed through measuring compliance. It is not all-encompassing but a focused subset
2496 that captures the aspects necessary to describe governance for SOA. It does not imply that practical
2497 application of governance is a single, isolated instance of these models; in reality, there may be
2498 hierarchical and parallel chains of governance that deal with different aspects or focus on different goals.
2499 This is discussed further in section 5.1.2.5. The defined models are simultaneously applicable to each of
2500 the overlapping instances.

2501 A given [enterprise](#) may already have portions of these models in place. To a large extent, the models
2502 shown here are not specific to SOA; discussions on direct applicability begin in section 5.1.3.

2503 **5.1.2.1 Motivating Governance**



2504
2505 *Figure 34 - Motivating Governance*


2506 An organizational domain such as an [enterprise](#) is made up of [participants](#) who may be individuals or
2507 groups of individuals forming smaller organizational units within the [enterprise](#). The overall business
2508 strategy should be consistent with the Goals of the [participants](#); otherwise, the business strategy would
2509 not provide value to the [participants](#) and governance towards those ends becomes difficult if not
2510 impossible. This is not to say that an instance of governance simultaneously satisfies all the goals of all
2511 the [participants](#); rather, the goals of any governance instance must sufficiently satisfy a useful subset of
2512 each [participant's](#) goals so as to provide value and ensure the cooperation of all the [participants](#).

2513 A policy is the formal characterization of the conditions and constraints that governance deems as
2514 necessary to realize the goals which it is attempting to satisfy. Policy may identify required conditions or
2515 actions or may prescribe limitations or other constraints on permitted conditions or actions. For example,
2516 a policy may prescribe that safeguards must be in place to prevent unauthorized access to sensitive
2517 material. It may also prohibit use of computers for activities unrelated to the specified work assignment.
2518 Policy is made operational through the promulgation and implementation of Rules and Regulations (as
2519 defined in section 5.1.2.3).

2520 As noted in section 4.4.2, policy may be asserted by any [participant](#) or on behalf of the [participant](#) by its
2521 organization. Part of the [purpose](#) of governance is to arbitrate among diverse goals of [participants](#) and
2522 the diverse policies articulated to realize those goals. The intent is to form a consistent whole that allows
2523 governance to minimize ambiguity about its [purpose](#). While resolving all ambiguity would be an ideal, it is
2524 unlikely that all inconsistencies will be identified and resolved before governance becomes operational.

2525 For governance to have effective jurisdiction over [participants](#), there must be some degree of agreement
2526 by all [participants](#) that they will abide by the governance mandates. A minimal degree of agreement often
2527 presages [participants](#) who “slow-roll” if not actively rejecting compliance with Policies that express the
2528 specifics of governance.

2529 **5.1.2.2 Setting Up Governance**



2530
2531 *Figure 35 - Setting Up Governance*

2532 **Leadership**

2533 Leadership is the entity who has the responsibility and authority to generate consistent policies
2534 through which the goals of governance can be expressed and to define and champion the
2535 structures and processes through which governance is realized.

2536 **Governance Framework**

2537 The Governance Framework is a set of organizational structures that enable governance to be
2538 consistently defined, clarified, and as needed, modified to respond to changes in its domain of
2539 concern.

2540 **Governance Processes**

2541 Governance Processes are the defined set of activities that are performed within the Governance
2542 Framework to enable the consistent definition, application, and as needed, modification of Rules
2543 that organize and regulate the activities of participants for the fulfillment of expressed policies.
2544 (See section 5.1.2.3 for elaboration on the relationship of Governance Processes and Rules.)

2545 As noted earlier, governance requires an appropriate organizational structure and identification of who
2546 has authority to make governance decisions. In Figure 35, the entity with governance authority is
2547 designated the Leadership. This is someone, possibly one or more of the participants, which participants
2548 recognize as having authority for a given purpose or over a given set of issues or concerns.

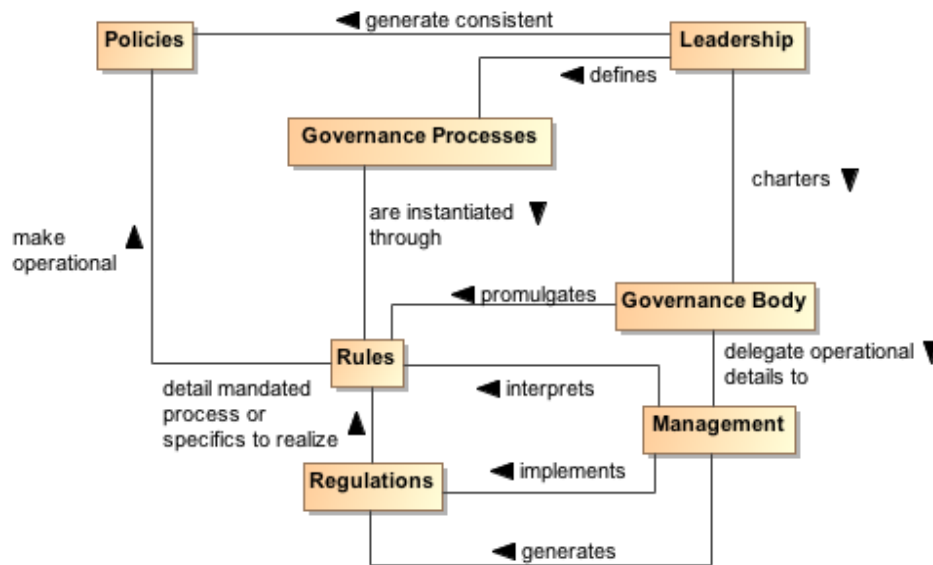
2549 The Leadership is responsible for prescribing or delegating a working group to prescribe the Governance
2550 Framework that forms the structure for Governance Processes which define how governance is to be
2551 carried out. This does not itself define the specifics of how governance is to be applied, but it does
2552 provide an unambiguous set of procedures that should ensure consistent actions which participants agree
2553 are fair and account for sufficient input on the subjects to which governance is applied.

2554 The participants may be part of the working group that codifies the Governance Framework and
2555 Processes. When complete, the participants must acknowledge and agree to abide by the products
2556 generated through application of this structure.

2557 The Governance Framework and Processes are often documented in the charter of a body created or
2558 designated to oversee governance. This is discussed further in the next section. Note that the
2559 Governance Processes should also include those necessary to modify the Governance Framework itself.

2560 An important function of Leadership is not only to initiate but also be the consistent champion of
 2561 governance. Those responsible for carrying out governance mandates must have Leadership who make
 2562 it clear to [participants](#) that expressed Policies are seen as a means to realizing established goals and that
 2563 compliance with governance is required.

2564 **5.1.2.3 Carrying Out Governance**



2565
 2566 *Figure 36 - Carrying Out Governance*

2567 **Rule**

2568 A Rule is a prescribed guide for carrying out activities and processes leading to desired results,
 2569 e.g. the operational realization of policies.

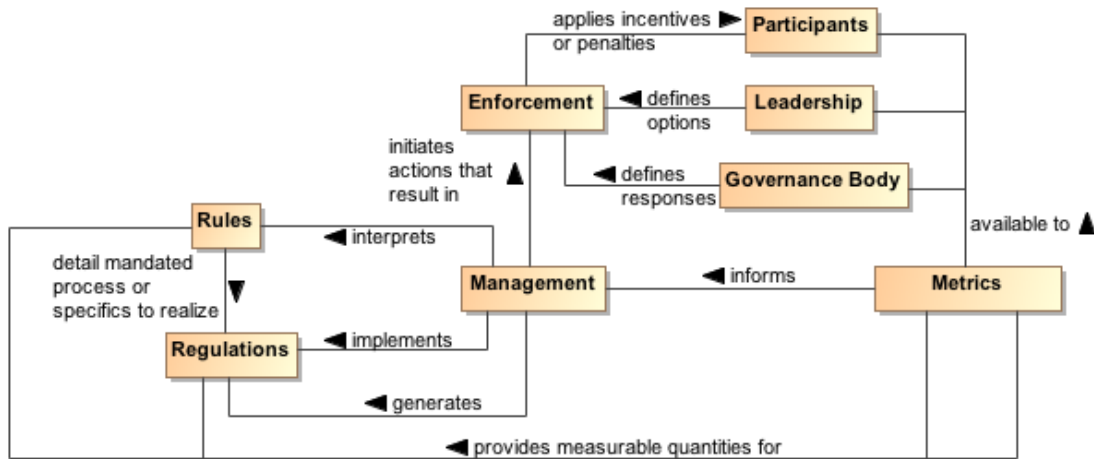
2570 **Regulation**

2571 A Regulation is a mandated process or the specific details that derive from the interpretation of
 2572 Rules and lead to measurable quantities against which compliance can be measured.

2573 To carry out governance, Leadership charts a Governance Body to promulgate the Rules needed to
 2574 make the Policies operational. The Governance Body acts in line with Governance Processes for its rule-
 2575 making process and other functions. Whereas Governance is the setting of Policies and defining the
 2576 Rules that provide an operational context for Policies, the operational details of governance may be
 2577 delegated by the Governance Body to Management. Management generates Regulations that specify
 2578 details for Rules and other procedures to implement both Rules and Regulations. For example,
 2579 Leadership could set a Policy that all authorized parties should have access to data, the Governance
 2580 Body would promulgate a Rule that PKI certificates are required to establish identity of authorized parties,
 2581 and Management can specify a Regulation of who it deems to be a recognized PKI issuing body. In
 2582 summary, Policy is a predicate to be satisfied and Rules prescribe the activities by which that satisfying
 2583 occurs. A number of rules may be required to satisfy a given policy; the carrying out of a rule may
 2584 contribute to several policies being realized.

2585 Whereas the Governance Framework and Processes are fundamental for having [participants](#)
 2586 acknowledge and commit to compliance with governance, the Rules and Regulations provide operational
 2587 constraints which may require resource [commitments](#) or other levies on the [participants](#). It is important
 2588 for [participants](#) to consider the framework and processes to be fair, unambiguous, and capable of being
 2589 carried out in a consistent manner and to have an opportunity to formally accept or ratify this situation.
 2590 Rules and Regulations, however, do not require individual acceptance by any given [participant](#) although
 2591 some level of community comment may be part of the Governance Processes. Having agreed to
 2592 governance, the [participants](#) are bound to comply or be subject to prescribed mechanisms for
 2593 enforcement.

2594 **5.1.2.4 Ensuring Governance Compliance**



2595
2596 *Figure 37 - Ensuring Governance Compliance*

2597 Setting Rules and Regulations does not ensure effective governance unless compliance can be
 2598 measured and Rules and Regulations can be enforced. Metrics are those conditions and quantities that
 2599 can be measured to characterize actions and results. Rules and Regulations MUST be based on
 2600 collected Metrics or there is no means for Management to assess compliance. The Metrics are available
 2601 to the participants, the Leadership, and the Governance Body so what is measured and the results of
 2602 measurement are clear to everyone.

2603 The Leadership in its relationship with participants has certain options that can be used for Enforcement.
 2604 A common option may be to affect future funding. The Governance Body defines specific enforcement
 2605 responses, such as what degree of compliance is necessary for full funding to be restored. It is up to
 2606 Management to identify compliance shortfalls and to initiate the Enforcement process.

2607 Note, enforcement does not strictly need to be negative consequences. Management can use Metrics to
 2608 identify exemplars of compliance and Leadership can provide options for rewarding the participants. The
 2609 Governance Body defines awards or other incentives.

2610 **5.1.2.5 Considerations for Multiple Governance Chains**

2611 As noted in section 5.1.2, instances of the governance model often occur as a tiered arrangement, with
 2612 governance at some level delegating specific authority and responsibility to accomplish a focused portion
 2613 of the original level's mandate. For example, a corporation may encompass several lines of business and
 2614 each line of business governs its own affairs in a manner that is consistent with and contributes to the
 2615 goals of the parent organization. Within the line of business, an IT group may be given the mandate to
 2616 provide and maintain IT resources, giving rise to IT governance.

2617 In addition to tiered governance, there may be multiple governance chains working in parallel. For
 2618 example, a company making widgets has policies intended to ensure they make high quality widgets and
 2619 make an impressive profit for their shareholders. On the other hand, Sarbanes-Oxley is a parallel
 2620 governance chain in the United States that specifies how the management must handle its accounting
 2621 and information that needs to be given to its shareholders. The parallel chains may just be additive or
 2622 may be in conflict and require some harmonization.

2623 Being distributed and representing different ownership domains, a SOA participant falls under the
 2624 jurisdiction of multiple governance domains simultaneously and may individually need to resolve
 2625 consequent conflicts. The governance domains may specify precedence for governance conformance or
 2626 it may fall to the discretion of the participant to decide on the course of actions they believe appropriate.

2627 5.1.3 Governance Applied to SOA

2628 5.1.3.1 Where SOA Governance is Different

2629 SOA governance is often discussed in terms of IT governance, but rather than a parent-child relationship,
2630 Figure 38 shows the two as siblings of the general governance described in section 5.1.2. There are
2631 obvious dependencies and a need for coordination between the two, but the idea of aligning IT with
2632 business already demonstrates that resource providers and resource consumers must be working
2633 towards common goals if they are to be productive and efficient. While SOA governance is shown to be
2634 active in the area of infrastructure, it is a specialized concern for having a dependable platform to support
2635 service interaction; a range of traditional IT issues is therefore out of scope of this document. A SOA
2636 governance plan for an [enterprise](#) will not of itself resolve shortcomings with the enterprise's IT
2637 governance.

2638 Governance in the context of SOA is that organization of services: that promotes their visibility; that
2639 facilitates interaction among service [participants](#); and that directs that the results of service interactions
2640 are those [real world effects](#) as described within the service description and constrained by policies and
2641 [contracts](#) as assembled in the execution context.

2642 SOA governance must specifically account for control across different ownership domains, i.e. all the
2643 [participants](#) may not be under the jurisdiction of a single governance authority. However, for governance
2644 to be effective, the [participants](#) must agree to recognize the [authority](#) of the Governance Body and must
2645 operate within the Governance Framework and through the Governance Processes so defined.

2646 SOA governance must account for interactions across [ownership boundaries](#), which may also imply
2647 across enterprise governance boundaries. For such situations, governance emphasizes the need for
2648 agreement that some Governance Framework and Governance Processes have jurisdiction, and the
2649 governance defined must satisfy the Goals of the [participants](#) for cooperation to continue. A standards
2650 development organization such as OASIS is an example of voluntary agreement to governance over a
2651 limited domain to satisfy common goals.


2652 The specifics discussed in the figures in the previous sections are equally applicable to governance
2653 across [ownership boundaries](#) as it is within a single boundary. There is a charter agreed to when
2654 [participants](#) become members of the organization, and this charter sets up the structures and processes
2655 to be followed. Leadership may be shared by the leadership of the overall organization and the
2656 leadership of individual groups themselves chartered per the Governance Processes. There are
2657 Rules/Regulations specific to individual efforts for which [participants](#) agree to local goals, and
2658 Enforcement can be loss of voting rights or under extreme circumstances, expulsion from the group.

2659 Thus, the major difference for SOA governance is an appreciation for the cooperative nature of the
2660 enterprise and its reliance on furthering common goals if productive participation is to continue.

2661 5.1.3.2 What Must be Governed

2662 An expected benefit of employing SOA principles is the ability to quickly bring [resources](#) to bear to deal
2663 with unexpected and evolving situations. This requires a great deal of confidence in the underlying
2664 capabilities that can be accessed and in the services that enable the access. It also requires
2665 considerable flexibility in the ways these [resources](#) can be employed. Thus, SOA governance requires
2666 establishing confidence and trust while instituting a solid framework that enables flexibility, indicating a
2667 combination of strict control over a limited set of foundational aspects but minimum constraints beyond
2668 those bounds.

2669



2670
2671 *Figure 38 - Relationship Among Types of Governance*

2672 SOA governance applies to three aspects of service definition and use:

- 2673 • SOA infrastructure – the “plumbing” that provides utility functions that enable and support the use
- 2674 of the service
- 2675 • Service inventory – the requirements on a service to permit it to be accessed within the
- 2676 infrastructure
- 2677 • Participant interaction – the consistent expectations with which all [participants](#) are expected to
- 2678 comply

2679 **5.1.3.2.1 Governance of SOA Infrastructure**

2680 The SOA infrastructure is likely composed of several families of SOA services that provide access to
2681 fundamental computing business services. These include, among many others, services such as
2682 messaging, security, storage, discovery, and mediation. The provisioning of an infrastructure on which
2683 these services may be accessed and the general realm of those contributing as utility functions of the
2684 infrastructure are a traditional IT governance concern. In contrast, the focus of SOA governance is how
2685 the existence and use of the services enables the SOA ecosystem.

2686 By characterizing the environment as containing families of SOA services, the assumption is that there
2687 may be multiple approaches to providing the business services or variations in the actual business
2688 services provided. For example, discovery could be based on text search, on metadata search, on
2689 approximate matches when exact matches are not available, and numerous other variations. The
2690 underlying implementation of search algorithms are not the purview of SOA governance, but the access
2691 to the resulting service infrastructure enabling discovery must be stable, reliable, and extremely robust to
2692 all operating conditions. Such access enables other specialized SOA services to use the infrastructure in
2693 dependable and predictable ways, and is where governance is important.

2694 **5.1.3.2.2 Governance of the Service Inventory**

2695 Given an infrastructure in which other SOA services can operate, a key governance issue is which SOA
2696 services to allow in the ecosystem. The major concern SHOULD be a definition of well-behaved services,
2697 where the required behavior will inherit their characteristics from experiences with distributed computing
2698 but also evolve with SOA experience. A major requirement for ensuring well-behaved services is
2699 collecting sufficient metrics to know how the service affects the SOA infrastructure and whether it
2700 complies with established infrastructure policies.

2701 Another common concern of service approval is whether there is a possibility of duplication of function by
2702 multiple services. Some governance models talk to a tightly controlled environment where a primary
2703 concern is to avoid any service duplication. Other governance models talk to a market of services where
2704 the consumers have wide choices. For the latter, it is anticipated that the better services will emerge from
2705 market consensus and the availability of alternatives will drive innovation.

2706 Some combination of control and openness will emerge, possibly with a different appropriate balance for
2707 different categories of use. For SOA governance, the issue is less which services are approved but rather
2708 ensuring that sufficient description is available to support informed decisions for appropriate use. Thus,
2709 SOA governance SHOULD concentrate on identifying the required attributes to adequately describe a
2710 service, the required target values of the attributes, and the standards for defining the meaning of the
2711 attributes and their target values. Governance may also specify the processes by which the attribute
2712 values are measured and the corresponding certification that some realized attribute set may imply.

2713 For example, unlimited access for using a service may require a degree of life cycle maturity that has
2714 demonstrated sufficient testing over a certain size community. Alternately, the policy may specify that a
2715 service in an earlier phase of its life cycle may be made available to a smaller, more technically
2716 sophisticated group in order to collect the metrics that would eventually allow the service to advance its
2717 life cycle status.

2718 This aspect of governance is tightly connected to description because, given a well-behaved set of
2719 services, it is the responsibility of the consumer (or policies promulgated by the consumer's organization)
2720 to decide whether a service is sufficient for that consumer's intended use. The goal is to avoid global
2721 governance specifying criteria that are too restrictive or too lax for the local needs of which global
2722 governance has little insight.

2723 Such an approach to specifying governance allows independent domains to describe services in local
2724 terms while still having the services available for informed use across domains. In addition, changes to
2725 the attribute sets within a domain can be similarly described, thus supporting the use of newly described
2726 [resources](#) with the existing ones without having to update the description of all the legacy content.

2727 **5.1.3.2.3 Governance of Participant Interaction**

2728 Finally, given a reliable services infrastructure and a predictable set of services, the third aspect of
2729 governance is prescribing what is required during a service interaction.

2730 Governance would specify adherence to service interface and service reachability parameters and would
2731 require that the result of an interaction MUST correspond to the [real world effects](#) as contained in the
2732 service description. Governance would ensure preconditions for service use are satisfied, in particular
2733 those related to security aspects such as user authentication, authorization, and non-repudiation. If
2734 conflicts arise, governance would specify resolution processes to ensure appropriate agreements,
2735 policies, and conditions are met.

2736 It would also rely on sufficient monitoring by the SOA infrastructure to ensure services remain well-
2737 behaved during interactions, e.g. do not use excessive resources or exhibit other prohibited behavior.
2738 Governance would also require that policy agreements as documented in the execution context for the
2739 interaction are observed and that the results and any after effects are consistent with the agreed policies.
2740 Governance focuses on more contractual and legal aspects rather than the precursor descriptive aspects.
2741 SOA governance may prescribe the processes by which SOA-specific policies are allowed to change, but
2742 there are probably more business-specific policies that will be governed by processes outside SOA
2743 governance.

2744 **5.1.3.3 Overarching Governance Concerns**

2745 There are numerous governance related concerns whose effects span the three areas just discussed.
2746 One is the area of standards, how these are mandated, and how the mandates may change. The Web
2747 Services standards stack is an example of relevant standards where a significant number are still under
2748 development. In addition, while there are notional scenarios that guide what standards are being
2749 developed, the fact that many of these standards do not yet exist precludes operational testing of their
2750 adequacy or effectiveness as a necessary and sufficient set.

2751 That said, standards are critical to creating a SOA ecosystem where SOA services can be introduced,
2752 used singularly, and combined with other services to deliver complex business functionality. As with
2753 other aspects of SOA governance, the Governance Body should identify the minimum set felt to be
2754 needed and rigorously enforce that that set be used where appropriate. The Governance Body takes
2755 care to expand and evolve the mandated standards in a predictable manner and with sufficient technical
2756 guidance that new services are able to coexist as much as possible with the old, and changes to
2757 standards do not cause major disruptions.

2758 Another area that may see increasing activity as SOA expands is additional regulation by governments
2759 and associated legal institutions. New laws may deal with transactions which are service based, possibly
2760 including taxes on the transactions. Disclosure laws may mandate certain elements of description so
2761 both the consumer and provider act in a predictable environment and are protected from ambiguity in
2762 intent or [action](#). Such laws spawn rules and regulations that will influence the metrics collected for
2763 evaluation of compliance.

2764 **5.1.3.4 Considerations for SOA Governance**

2765 The Reference Architecture definition of a loosely coupled system is one in which the constraints on the
2766 interactions between components is minimal: sufficient to permit interoperation without additional
2767 constraints that may be an artifact of implementation technology. While governance experience for
2768 standalone systems provides useful guides, we must be careful not to apply constraints that would
2769 preclude the flexibility, agility, and adaptability we expect to realize from a SOA ecosystem.

2770 One of the strengths of SOA is it can make effective use of diversity rather than requiring monolithic
2771 solutions. Heterogeneous organizations can interact without requiring each conforms to uniform tools,
2772 representation, and processes. However, with this diversity comes the need to adequately define those
2773 elements necessary for consistent interaction among systems and [participants](#), such as which
2774 communication protocol, what level of security, which vocabulary for payload content of messages. The
2775 solution is not always to lock down these choices but to standardize alternatives and standardize the
2776 representations through which an unambiguous identification of the alternative chosen can be conveyed.
2777 For example, the URI standard specifies the URI string, including what protocol is being used, what is the
2778 target of the message, and how parameters may be attached. It does not limit the available protocols, the
2779 semantics of the target address, or the parameters that can be transferred. Thus, as with our definition of
2780 loose coupling, it provides absolute constraints but minimizes which constraints it imposes.

2781 There is not a one-size-fits-all governance but a need to understand the types of things governance is
2782 called upon to do in the context of the goals of SOA. Some communities may initially desire and require
2783 very stringent governance policies and procedures while others see need for very little. Over time, best
2784 practices will evolve, resulting in some consensus on a sensible minimum and, except in extreme cases
2785 where it is demonstrated to be necessary, a loosening of strict governance toward the best practice
2786 mean.

2787 A question of how much governance may center on how much time governance activities require versus
2788 how quickly is the system being governed expected to respond to changing conditions. For large single
2789 systems that take years to develop, the governance process could move slowly without having a serious
2790 negative impact. For example, if something takes two years to develop and the steps involved in
2791 governance take two months to navigate, then the governance can go along in parallel and may not have
2792 a significant impact on system response to changes. Situations where it takes as long to navigate
2793 governance requirements as it does to develop a response are examples where governance may need to
2794 be reevaluated as to whether it facilitates or inhibits the desired results. Thus, the speed at which
2795 services are expected to appear and evolve needs to be considered when deciding the processes for
2796 control. The added weight of governance should be appropriate for overall goals of the application
2797 domain and the service environment.

2798 Governance, as with other aspects of any SOA implementation, should start small and be conceptualized
2799 in a way that keeps it flexible, scalable, and realistic. A set of useful guidelines would include:

- 2800 • Do not hardwire things that will inevitably change. For example, develop a system that uses the
2801 representation of policies rather than code the policies into the implementations.
- 2802 • Avoid setting up processes that demo well for three services without considering how they may
2803 work for 300. Similarly, consider whether the display of status and activity for a small number of
2804 services will also be effective for an operator in a crisis situation looking at dozens of services,
2805 each with numerous, sometimes overlapping and sometimes differing activities.
- 2806 • Maintain consistency and realism. A service solution responding to a natural disaster cannot be
2807 expected to complete a 6-week review cycle but be effective in a matter of hours.

2808 **5.1.4 Architectural Implications of SOA Governance**

2809 The description of SOA governance indicates numerous architectural requirements on the SOA
2810 ecosystem:

- 2811 • Governance is expressed through policies and assumes multiple use of focused policy modules
2812 that can be employed across many common circumstances. This requires the existence of:
 - 2813 ○ descriptions to enable the policy modules to be visible, where the description includes a
2814 unique identifier for the policy and a sufficient, and preferably a machine process-able,
2815 representation of the meaning of terms used to describe the policy, its functions, and its
2816 effects;
 - 2817 ○ one or more discovery mechanisms that enable searching for policies that best meet the
2818 search criteria specified by the service [participant](#); where the discovery mechanism will
2819 have access to the individual policy descriptions, possibly through some repository
2820 mechanism;
 - 2821 ○ accessible storage of policies and policy descriptions, so service [participants](#) can access,
2822 examine, and use the policies as defined.
- 2823 • Governance requires that the [participants](#) understand the intent of governance, the structures
2824 created to define and implement governance, and the processes to be followed to make
2825 governance operational. This requires the existence of:
 - 2826 ○ an information collection site, such as a Web page or portal, where governance
2827 information is stored and from which the information is always available for access;
 - 2828 ○ a mechanism to inform [participants](#) of significant governance [events](#), such as changes in
2829 policies, rules, or regulations;
 - 2830 ○ accessible storage of the specifics of Governance Processes;
 - 2831 ○ SOA services to access automated implementations of the Governance Processes
- 2832 • Governance policies are made operational through rules and regulations. This requires the
2833 existence of:
 - 2834 ○ descriptions to enable the rules and regulations to be visible, where the description
2835 includes a unique identifier and a sufficient, and preferably a machine process-able,
2836 representation of the meaning of terms used to describe the rules and regulations;
 - 2837 ○ one or more discovery mechanisms that enable searching for rules and regulations that
2838 may apply to situations corresponding to the search criteria specified by the service
2839 [participant](#); where the discovery mechanism will have access to the individual
2840 descriptions of rules and regulations, possibly through some repository mechanism;
 - 2841 ○ accessible storage of rules and regulations and their respective descriptions, so service
2842 [participants](#) can understand and prepare for compliance, as defined.
 - 2843 ○ SOA services to access automated implementations of the Governance Processes.
- 2844 • Governance implies management to define and enforce rules and regulations. Management is
2845 discussed more specifically in section 5.3, but in a parallel to governance, management requires
2846 the existence of:
 - 2847 ○ an information collection site, such as a Web page or portal, where management
2848 information is stored and from which the information is always available for access;
 - 2849 ○ a mechanism to inform [participants](#) of significant management [events](#), such as changes
2850 in rules or regulations;
 - 2851 ○ accessible storage of the specifics of processes followed by management.
- 2852 • Governance relies on metrics to define and measure compliance. This requires the existence of:
 - 2853 ○ the infrastructure monitoring and reporting information on SOA resources;
 - 2854 ○ possible interface requirements to make accessible metrics information generated or
2855 most easily accessed by the service itself.

2856 **5.2 Security Model**

2857 Security is one aspect of confidence – the confidence in the integrity, reliability, and confidentiality of the
2858 system. In particular, security focuses on those aspects of assurance that involve the accidental or malign
2859 intent of other people to damage or compromise trust in the system and on the availability of SOA-based
2860 systems to perform desired capability.

2861 **Security**

2862 Security concerns the set of mechanisms for ensuring and enhancing trust and confidence in the
2863 SOA ecosystem.

2864 Providing for security for Service Oriented Architecture is somewhat different than for other contexts;
2865 although many of the same principles apply equally to SOA and to other systems. The fact that SOA
2866 embraces crossing [ownership boundaries](#) makes the issues involved with moving data more visible.

2867 As well as securing the movement of data within and across [ownership boundaries](#), security often
2868 revolves around [resources](#): the need to guard certain resources against inappropriate access – whether
2869 reading, writing or otherwise manipulating those resources.

2870 Any comprehensive security solution must take into account the people that are using, maintaining and
2871 managing the SOA. Furthermore, the relationships between them must also be incorporated: any security
2872 assertions that may be associated with particular interactions originate in the people that are behind the
2873 interaction.

2874 We analyze security in terms of the [social structures](#) that define the legitimate [permissions](#), [obligations](#)
2875 and [roles](#) of people in relation to the system, and mechanisms that must be put into place to realize a
2876 secure system. The former are typically captured in a series of security policy statements; the latter in
2877 terms of security *guards* that ensure that policies are enforced.

2878 How and when to apply these derived security policy mechanisms is directly associated with the
2879 assessment of the *threat model* and a *security response model*. The threat model identifies the kinds of
2880 threats that directly impact the message and/or application of constraints, and the response model is the
2881 proposed mitigation to those threats. Properly implemented, the result can be an acceptable level of risk
2882 to the safety and integrity of the system.

2883 **5.2.1 Secure Interaction Concepts**

2884 We can characterize secure interactions in terms of key security concepts **[ISO/IEC 27002]**:
2885 confidentiality, integrity, authentication, authorization, non-repudiation, and availability. The concepts for
2886 secure interactions are well defined in other standards and publications. The security concepts here are
2887 not defined but rather related to the SOA ecosystem perspective of the SOA-RAF.

2888 **5.2.1.1 Confidentiality**

2889 Confidentiality concerns the protection of privacy of [participants](#) in their interactions. Confidentiality refers
2890 to the assurance that unauthorized entities are not able to read messages or parts of messages that are
2891 transmitted.

2892 Note that confidentiality has degrees: in a completely confidential exchange, third parties would not even
2893 be aware that a confidential exchange has occurred. In a partially confidential exchange, the identities of
2894 the participants may be known but the content of the exchange obscured.

2895 **5.2.1.2 Integrity**

2896 Integrity concerns the protection of information that is exchanged – either from unauthorized writing or
2897 inadvertent corruption. Integrity refers to the assurance that information that has been exchanged has not
2898 been altered.

2899 Integrity is different from confidentiality in that messages that are sent from one participant to another
2900 may be obscured to a third party, but the third party may still be able to introduce his own content into the
2901 exchange without the knowledge of the participants.

2902 Section 5.2.4 describes common computing techniques for providing confidentiality and integrity during
2903 message exchanges.


2904 **5.2.1.3 Authentication**

2905 Authentication concerns the identity of the participants in an exchange. Authentication refers to the
2906 means by which one participant can be assured of the identity of other participants.

2907

2908 *Figure 39 - Authentication*

2909 applies authentication to the identity of participants.

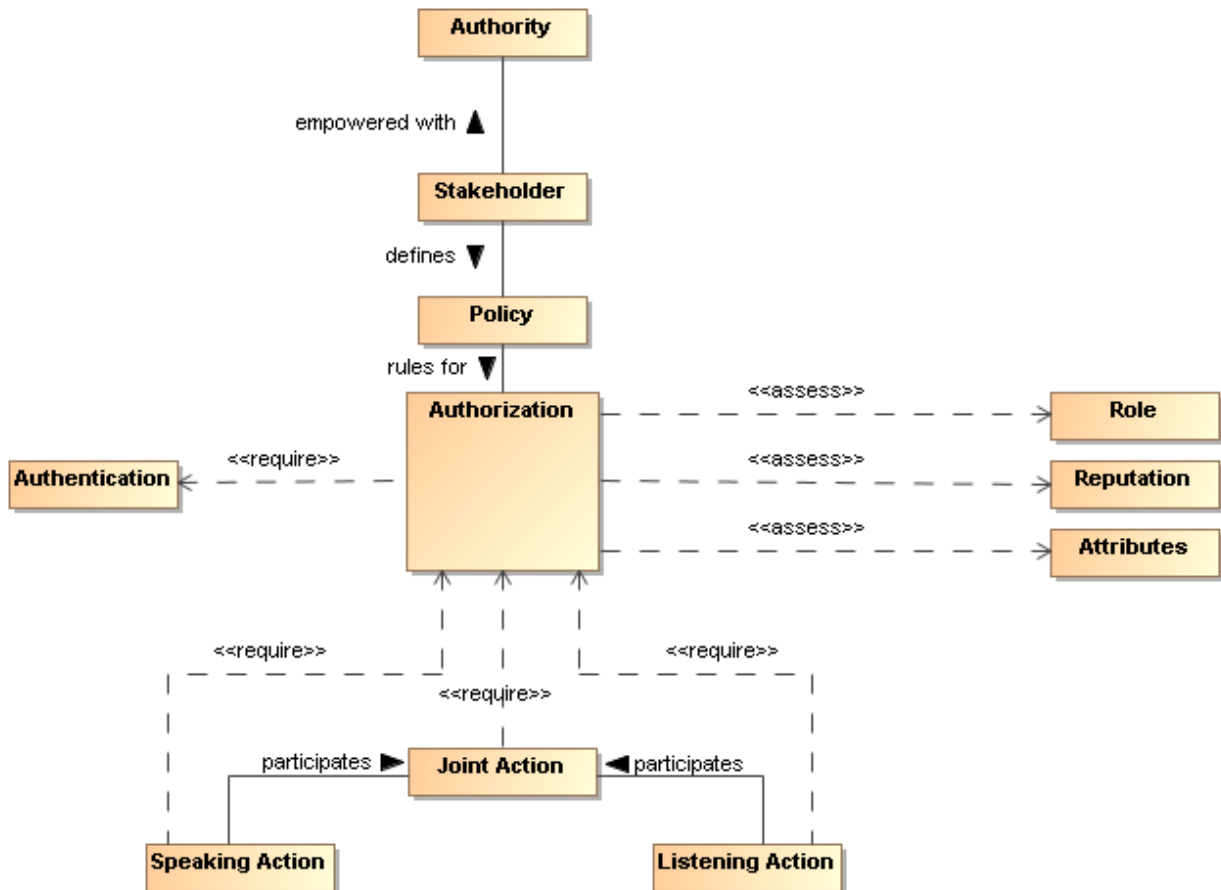


2910

2911 *Figure 39 - Authentication*

2912 5.2.1.4 Authentication

2913 Authorization concerns the legitimacy of the interaction. Authorization refers to the means by which a
2914 stakeholder may be assured that the information and actions that are exchanged are either explicitly or
2915 implicitly approved.



2916

2917 *Figure 40 - Authorization*

2918 The **roles** and attributes which provide a **participant's** credentials are expanded to include reputation.
2919 Reputation often helps determine willingness to interact, for example, reviews of a service provider will
2920 influence the decision to interact with the service provider. The **roles**, reputation, and attributes are
2921 represented as assertions measured by authorization decision points.

2922 The role of policy for security is to permit stakeholders to express their choices. In Figure 40, a policy is a
2923 written constraint and the role, reputation, and attribute assertions are evaluated according to the
2924 constraints in the authorization policy. A combination of security mechanisms and their control via
2925 explicit policies can form the basis of an authorization solution.

2926 **5.2.1.5 Non-repudiation**

2927 Non-repudiation concerns the accountability of [participants](#). To foster trust in the performance of a system
2928 used to conduct shared activities it is important that the [participants](#) are not able to later deny their
2929 actions: to repudiate them. Non-repudiation refers to the means by which a [participant](#) may not, at a later
2930 time, successfully deny having participated in the interaction or having performed the actions as reported
2931 by other [participants](#).

2932 **5.2.1.6 Availability**

2933 Availability concerns the ability of systems to use and offer the services for which they were designed.
2934 One of the threats against availability is the so-called denial of service attack in which attackers attempt to
2935 prevent legitimate access to the system.

2936 We differentiate here between general availability – which includes aspects such as systems reliability –
2937 and availability as a security concept where we need to respond to active threats to the system.

2938 **5.2.2 Where SOA Security is Different**

2939 The core security concepts are fundamental to all social interactions. The evolution of sharing
2940 information using a SOA requires the flexibility to dynamically secure computing interactions in a
2941 computing ecosystem where the owning social groups, [roles](#), and [authority](#) are constantly changing as
2942 described in section 5.1.3.1.

2943 SOA policy-based security can be more adaptive for a computing ecosystem than previous computing
2944 technologies allow for, and typically involves a greater degree of distributed mechanisms.

2945 Standards for security, as is the case with all aspects of SOA, play a large role in flexible security on a
2946 global scale. SOA security may also involve greater auditing and reporting to adhere to regulatory
2947 compliance established by governance structures.

2948 **5.2.3 Security Threats**

2949 There are a number of ways in which an attacker may attempt to compromise the security of a system.
2950 The two primary sources of attack are third parties attempting to subvert interactions between legitimate
2951 [participants](#) and an entity that is participating but attempting to subvert its partner(s). The latter is
2952 particularly important in a SOA where there may be multiple [ownership boundaries](#) and trust boundaries.

2953 The threat model lists some common threats that relate to the core security concepts listed in Section
2954 5.2.1. Each technology choice in the realization of a SOA can potentially have many threats to consider.

2955 **Message alteration**

2956 If an attacker is able to modify the content (or even the order) of messages that are exchanged
2957 without the legitimate [participants](#) being aware of it then the attacker has successfully
2958 compromised the security of the system. In effect, the [participants](#) may unwittingly serve the
2959 needs of the attacker rather than their own.

2960 An attacker may not need to completely replace a message with his own to achieve his objective:
2961 replacing the identity of the beneficiary of a transaction may be enough.

2962 **Message interception**

2963 If an attacker is able to intercept and understand messages exchanged between [participants](#),
2964 then the attacker may be able to gain advantage. This is probably the most commonly understood
2965 security threat.

2966 **Man in the middle**

2967 In a man-in-the-middle attack, the legitimate [participants](#) believe that they are interacting with
2968 each other; but are in fact interacting with the attacker. The attacker attempts to convince each
2969 [participant](#) that he is their correspondent; whereas in fact he is not.

2970 In a successful man-in-the-middle attack, legitimate [participants](#) do not have an accurate
2971 understanding of the state of the other [participants](#). The attacker can use this to subvert the
2972 intentions of the [participants](#).

2973 **Spoofing**

2974 In a spoofing attack, the attacker convinces a [participant](#) that he is really someone else –
2975 someone that the [participant](#) would normally trust.

2976 **Denial of service attack**

2977 In a denial of service (DoS) attack, the attacker attempts to prevent legitimate users from making
2978 use of the service. A DoS attack is easy to mount and can cause considerable harm: by
2979 preventing legitimate interactions, or by slowing them down enough, the attacker may be able to
2980 simultaneously prevent legitimate access to a service and to attack the service by another
2981 means.

2982 A variation of the DoS attack is the Distributed Denial of Service attack. In a DDoS attack the
2983 attacker uses multiple agents to attack the target. In some circumstances this can be
2984 extremely difficult to counteract effectively.

2985 One of the features of a DoS attack is that it does not require valid interactions to be effective:
2986 responding to invalid messages also takes resources and that may be sufficient to cripple the
2987 target.

2988 **Replay attack**

2989 In a replay attack, the attacker captures the message traffic during a legitimate interaction and
2990 then replays part of it to the target. The target is persuaded that a similar transaction to the
2991 previous one is being repeated and it responds as though it were a legitimate interaction.

2992 A replay attack may not require that the attacker understand any of the individual
2993 communications; the attacker may have different objectives (for example attempting to predict
2994 how the target would react to a particular request).

2995 **False repudiation**

2996 In false repudiation, a user completes a normal transaction and then later attempts to deny that
2997 the transaction occurred. For example, a customer may use a service to buy a book using a credit
2998 card; then, when the book is delivered, refuse to pay the credit card bill claiming that *someone*
2999 *else* must have ordered the book.

3000 **5.2.4 Security Responses**

3001 Security goals are never absolute: it is not possible to guarantee 100% confidentiality, non-repudiation,
3002 etc. However, a well designed and implemented security response model can ensure acceptable levels of
3003 security risk. For example, using a well-designed cipher to encrypt messages may make the cost of
3004 breaking communications so great and so lengthy that the information obtained is valueless.

3005 Performing threat assessments, devising mitigation strategies, and determining acceptable levels of risk
3006 are the foundation for an effective process to mitigating threats in a cost-effective way.¹¹ The choice in
3007 hardware and software to realize a SOA will be a basis for threat assessments and mitigation strategies.
3008 The [stakeholders](#) of a specific SOA implementation should determine acceptable levels of risk based on
3009 threat assessments and the cost of mitigating those threats.

3010 **5.2.4.1 Privacy Enforcement**

3011 The most efficient mechanism to assure confidentiality is the encryption of information. Encryption is
3012 particularly important when messages must cross trust boundaries; especially over the Internet. Note that
3013 encryption need not be limited to the content of messages: it is possible to obscure even the existence of
3014 messages themselves through encryption and ‘white noise’ generation in the communications channel.

3015 The specifics of encryption are beyond the scope of this architecture. However, we are concerned about
3016 how the connection between privacy-related policies and their enforcement is made.

3017 A policy enforcement point for enforcing privacy may take the form of an automatic function to encrypt
3018 messages as they leave a trust boundary; or perhaps simply ensuring that such messages are suitably
3019 encrypted.

3020 Any policies relating to the level of encryption being used would then apply to these centralized
3021 messaging functions.

3022 **5.2.4.2 Integrity Protection**

3023 To protect against message tampering or inadvertent message alteration, and to allow the receiver of a
3024 message to authenticate the sender, messages may be accompanied by a digital signature. Digital
3025 signatures provide a means to detect if signed data has been altered. This protection can also extend to
3026 authentication and non-repudiation of a sender.

3027 A common way a digital signature is generated is with the use of a private key that is associated with a
3028 public key and a digital certificate. The private key of some entity in the system is used to create a digital
3029 signature for some set of data. Other entities in the system can check the integrity of the signed data set
3030 via signature verification algorithms. Any changes to the data that was signed will cause signature
3031 verification to fail, which indicates that integrity of the data set has been compromised.

3032 A party verifying a digital signature must have access to the public key that corresponds to the private key
3033 used to generate the signature. A digital certificate contains the public key of the owner, and is itself
3034 protected by a digital signature created using the private key of the issuing Certificate Authority (CA).

3035 **5.2.4.3 Message Replay Protection**

3036 To protect against replay attacks, messages may contain information that can be used to detect replayed
3037 messages. The simplest requirement to prevent replay attacks is that each message that is ever sent is
3038 unique. For example, a message may contain a message ID, a timestamp, and the intended destination.

3039 By storing message IDs, and comparing each new message with the store, it becomes possible to verify
3040 whether a given message has been received before (and therefore should be discarded).

¹¹ 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.

3041 The timestamp may be included in the message to help check for message freshness. Messages that
3042 arrive after their message ID could have been cleared (after receiving the same message some time
3043 previously) may also have been replayed. A common means for representing timestamps is a useful part
3044 of an interoperable replay detection mechanism.

3045 The destination information is used to determine if the message was misdirected or replayed. If the
3046 replayed message is sent to a different endpoint than the destination of the original message, the replay
3047 could go undetected if the message does not contain information about the intended destination.

3048 In the case of messages that are replies to prior messages, it is also possible to include seed information
3049 in the prior messages that is randomly and uniquely generated for each message that is sent out. A
3050 replay attack can then be detected if the reply does not embed the random number that corresponds to
3051 the original message.

3052 **5.2.4.4 Auditing and Logging**

3053 False repudiation involves a [participant](#) denying that it authorized a previous interaction. An effective
3054 strategy for responding to such a denial is to maintain careful and complete logs of interactions which can
3055 be used for auditing [purposes](#). The more detailed and comprehensive an audit trail is, the less likely it is
3056 that a false repudiation would be successful.

3057 The countermeasures assume that the non-repudiation tactic (e.g. digital signatures) is not undermined
3058 itself. For example, if private key is stolen and used by an adversary, even extensive logging cannot
3059 assist in rejecting a false repudiation.

3060 Unlike many of the security responses discussed here, it is likely that the scope for automation in
3061 rejecting a repudiation attempt is limited to careful logging.

3062 **5.2.4.5 Graduated engagement**

3063 The key to managing and responding to DoS attacks is to be careful in the use of [resources](#) when
3064 responding to interaction. Put simply, a system has a choice to respond to a communication or to ignore
3065 it. In order to avoid vulnerability to DoS attacks a service provider should be careful not to commit
3066 [resources](#) beyond those implied by the current state of interactions; this permits a graduation in
3067 commitment by the service provider that mirrors any [commitment](#) on the part of [service consumers](#) and
3068 attackers alike.

3069 **5.2.5 Architectural Implications of SOA Security**

3070 Providing SOA security in an ecosystem of governed services has the following implications on the policy
3071 support and the distributed nature of mechanisms used to assure SOA security:

- 3072 • Security expressed through policies have the same architectural implications as described in
3073 Section 4.4.3 for policies and [contracts](#) architectural implications.
- 3074 • Security policies require mechanisms to support security description administration, storage, and
3075 distribution.
- 3076 • Service descriptions supporting security policies should:
 - 3077 ○ have a meta-structure sufficiently rich to support security policies;
 - 3078 ○ be able to reference one or more security policy artifacts;
 - 3079 ○ have a framework for resolving conflicts between security policies.
- 3080 • The mechanisms that make-up the execution context in secure SOA-based systems should:
 - 3081 ○ provide protection of the confidentiality and integrity of message exchanges;
 - 3082 ○ be distributed so as to provide centralized or decentralized policy-based identification,
3083 authentication, and authorization;
 - 3084 ○ ensure service availability to consumers;
 - 3085 ○ be able to scale to support security for a growing ecosystem of services;
 - 3086 ○ be able to support security between different communication technologies;
- 3087 • Common security services include:
 - 3088 ○ services that abstract encryption techniques;
 - 3089 ○ services for auditing and logging interactions and security violations;
 - 3090 ○ services for identification;

- 3091 ○ services for authentication;
- 3092 ○ services for authorization;
- 3093 ○ services for intrusion detection and prevention;
- 3094 ○ services for availability including support for quality of service specifications and metrics.

3095 **5.3 Management Model**

3096 **5.3.1 Management**

3097 Management is a process of controlling resources in accordance with the policies and principles defined
3098 by Governance.

3099 There are three separate but linked domains of interest within the management of SOA ecosystem:

- 3100 1. the management and support of the resources that are involved in any complex structures – of
3101 which SOA ecosystems are excellent examples;
- 3102 2. the promulgation and enforcement of the policies and service contracts agreed to by the
3103 stakeholders in the SOA ecosystem;
- 3104 3. the management of the relationships of the participants – both to each other and to the services
3105 that they use and offer.

3106 There are many artifacts related to management. Historically, systems management capabilities have
3107 been organized by the “FCAPS” functions (based on ITU-T Rec. M.3400 (02/2000), “TMN Management
3108 Functions”):

- 3109 • fault management,
- 3110 • configuration management,
- 3111 • account management,
- 3112 • performance and security management.

3113 The primary task of the functional groups is to concentrate on maintaining systems in a trusted, active,
3114 and accessible state.

3115 In the context of the SOA ecosystem, we see many possible resources that may require management
3116 such as services, service descriptions, service contracts, policies, roles, relationships, security, people
3117 and systems that implement services and infrastructure elements. In addition, given the ecosystem
3118 nature, it is also potentially necessary to manage the business relationships between participants.

3119 Successful operation of a SOA ecosystem requires trust among the stakeholders and between them and
3120 the SOA-based system elements. In contrast, regular systems in technology are not necessarily operated
3121 or used in an environment requiring trust before the stakeholders make use of the system. Indeed, many
3122 of these systems exist in hierarchical management structures, within which use may be mandated by
3123 legal requirement, executive decision, or good business practice in furthering the business’ strategy. The
3124 pre-condition of trust in the SOA ecosystem is rooted both in the principles of service orientation and in
3125 the distributed, authoritative ownership of independent services. Even for hierarchical management
3126 structures applied to a SOA ecosystem, the service in use should have a contractual basis rather than
3127 being mandated.

3128 Trust may be established through agreements/contracts, policies, or implicitly through observation of
3129 repeated interactions with others. Explicit trust is usually accompanied by formalized documents suitable
3130 for management. Implicit trust adds fragility to the management of a SOA ecosystem because failure to
3131 maintain consistent and predictable interactions will undermine the trust between participants and within
3132 the ecosystem as a whole.

3133 Management in a SOA ecosystem is thus concerned with management taking actions that will establish
3134 the condition of trust that must be present before engaging in service interactions. These concerns should
3135 largely be handled within the governance of the ecosystem. The policies, agreements, and practices
3136 defined through governance provide the boundaries within which management operates and for which
3137 management must provide enforcement and feedback. However, governance alone cannot foresee all
3138 circumstances but must offer sufficient guidance where agreement between all stakeholders cannot be
3139 reached. Management in these cases must be flexible and adaptable to handle unanticipated conditions
3140 without unnecessarily breaking trust relationships.

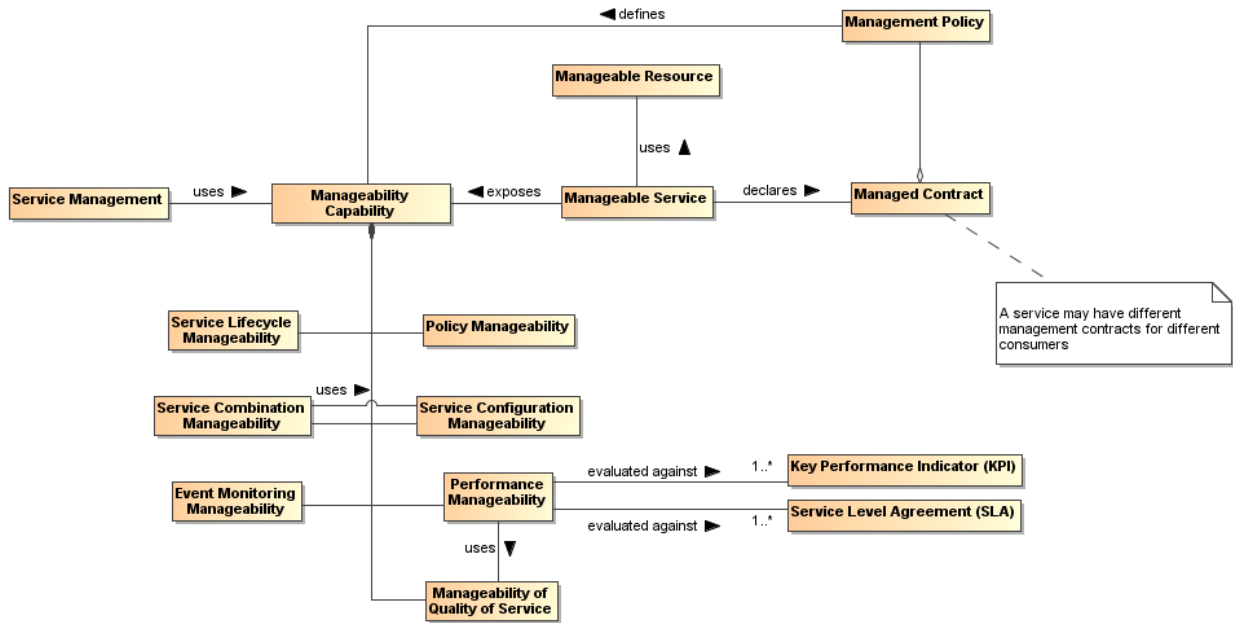
3141 Service management is the process – manual, automated, or a combination – of proactively monitoring
 3142 and controlling the behavior of a service or a set of services. Service management operates under
 3143 constraints attributed to the business and social context. Specific policies may be used to govern cross-
 3144 boundary relationships. Managing solutions based on such policies (and that may be used across
 3145 ownership boundaries) raises issues that are not typically present when managing a service within a
 3146 single ownership domain. Care is therefore required in managing a service when the owner of the
 3147 service, the provider of the service, the host of the service and mediators to the service may all belong to
 3148 different stakeholders.

3149 Cross-boundary service management takes place in, at least, the following situations:

- 3150 • using combinations of services that belong to different ownership domains
- 3151 • using of services that mediate between ownership domains
- 3152 • sharing monitoring and reporting means and results.

3153 These situations are particularly important in ecosystems that are highly decentralized, in which the
 3154 participants interact as peers as well as in the “master-servant” mode.

3155 The management model shown in Figure 41 conveys how the SOA applies to managing services.
 3156 Services management operates via service metadata, such as service lifecycles and attributes associated
 3157 with service use, which are typically collected in or accessed through the service description.
 3158



3159
 3160 *Figure 41 - Manageability capabilities in SOA ecosystem*

3161 The service metadata of interest is that set of service properties that is manageable. These manageability
 3162 properties are generally identifiable for any service consumed or supplied within the ecosystem. The
 3163 necessary existence of these properties within the SOA ecosystem motivates the following definitions:

3164 **Manageability** of a resource is the capability that allows it to be controlled, monitored, and reported on
 3165 with respect to some property. Note that manageability is not necessarily a part of the managed entities
 3166 themselves and are generally considered to be external to the managed entities.

3167 Each resource may be managed through a number of aspects of management, and the resources may
 3168 be grouped based on similar aspects. For example, resources may be grouped according to the aspect
 3169 referred to as “Configuration Manageability” for the collection of services. Some resources may not be
 3170 managed under a particular capability if there are no manageability aspects with a clear meaning or use.
 3171 As an example, all resources within a SOA ecosystem have a lifecycle that is meaningful within the
 3172 ecosystem. Thus, all resources are manageable under Lifecycle Manageability. In contrast, not all

3173 resources report or handle events. Thus, Event Manageability is only concerned with those resources for
3174 which events are meaningful.

3175 **Life-cycle Manageability** of a service typically refers to how the service is created, how it is destroyed
3176 and how service versions must be managed. This manageability is a feature of the SOA ecosystem
3177 because the service cannot manage its own life cycle.

3178 Another important consideration is that services may have resource requirements that must be
3179 established at various points in the services' life cycles. However actual providers of these resources may
3180 not be known at the time of the service creation and, thus, have to be managed at service run-time.

3181 **Combination Manageability** of a service addresses management of service characteristics that allow for
3182 creating and changing combinations in which the service participates or that the service combines itself.
3183 Known models of such combinations are aggregations and compositions. Examples of patterns of
3184 combinations are choreography and orchestration. Combination Manageability drives implementation of
3185 the Service Composability Principle of service orientation.

3186 Service combination manageability resonates with the methodology of process management.
3187 Combination Manageability may be applied at different phases of service creation and execution and, in
3188 some cases, can utilize Configuration Manageability.

3189 Service combinations typically contribute the most in delivering business values to the stakeholders.
3190 Managing service combinations is the one of the most important tasks and features of the SOA
3191 ecosystem.

3192 **Configuration Manageability** of a service allows managing the identity of and the interactions among
3193 internal elements of the service. Also, Configuration Manageability correlates with the management of
3194 service versions and configuration of the deployment of new services into the ecosystem. Configuration
3195 Management differs from the Combination Manageability in the scope and scale of manageability, and
3196 addresses lower level concerns than the architectural combination of services.

3197 **Event Monitoring Manageability** allows managing the categories of events of interest related to services
3198 and reporting recognized events to the interested stakeholders. Such events may be the ones that trigger
3199 service invocations as well as execution of particular functionality provided by the service.

3200 Event Monitoring Manageability is a key lower-level manageability aspect that the service provider and
3201 associated stakeholders are interested. Monitored events may be internal or external to the SOA
3202 ecosystem. For example, a disaster in the oil industry, which is outside the SOA ecosystem of the Insurer,
3203 can trigger the service's functionality that is responsible for immediate or constant monitoring of oil prices
3204 in the oil trading exchanges and, respectively, modify the premium paid by the insured oil companies.

3205 **Performance Manageability** of a service allows controlling the service results, shared and sharable real
3206 world effects against the business goals and objectives of the service. This manageability assumes
3207 monitoring of the business performance as well as the management of this monitoring itself. Performance
3208 Manageability includes business and technical performance manageability through a performance criteria
3209 set, such as business key performance indicators (KPI) and service-level agreements (SLA).

3210 The performance business- and technical-level characteristics of the service should be known from the
3211 service contract. The service provider and consumer must be able to monitor and measure these
3212 characteristics or be informed about the results measured by a third party.

3213 Performance Manageability is the instrument for providing compliance of the service with its service
3214 contracts. Performance Manageability utilizes Manageability of Quality of Service.

3215 **Manageability of Quality of Service** deals with management of service non-functional characteristics
3216 that may be of significant value to the service consumers and other stakeholders in the SOA ecosystem.
3217 Classic examples of this include bandwidth offerings associated with a service.

3218 Manageability of quality of service assumes that the properties associated with service qualities are
3219 monitored during the service execution. Results of monitoring may be challenged against SLA and even
3220 KPI, which results in the continuous validation of how the service contract is preserved by the service
3221 provider.

3222 **Policy Manageability** allows additions, changes and replacements of the policies associated with a
3223 resource in the SOA ecosystem. The ability to manage those policies (such as promulgating policies,

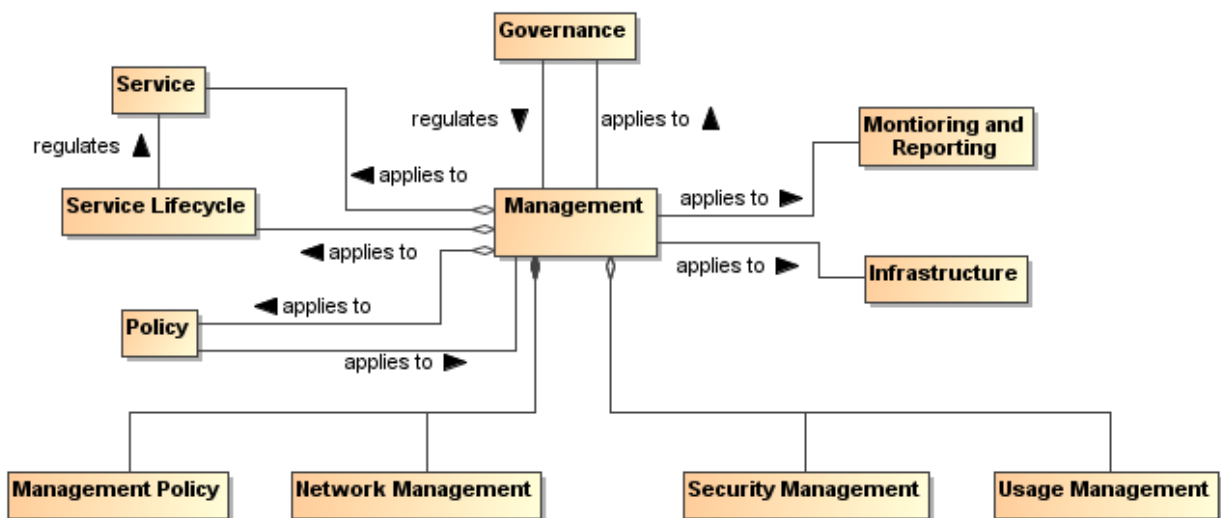
3224 retiring policies and ensuring that policy decision points and enforcement points are current) enables the
3225 ecosystem to apply policies and *evaluate* the results.

3226 The capability to manage, i.e. use a particular manageability, requires policies from governance to be
3227 translated into detailed rules and regulations which are measured and monitored providing corresponding
3228 feedback for enforcement.

3229 Management of SOA ecosystem recognises the manageability challenge and requires manageability
3230 properties to be considered for all aforementioned manageability cases. In the following sub-sections, we
3231 describe how these properties are used in the management. Also we describe some relationships
3232 between management and other components of SOA ecosystem.

3233 5.3.2 Management Means and Relationships

3234 A minimal set of management issues for the SOA ecosystem is shown on Figure 42 and elaborated in the
3235 following sections.



3236
3237 Figure 42 - Management Means and Relationships in SOA ecosystem

3238 5.3.2.1 Management Policy

3239 The management of resources within the SOA ecosystem may be governed by management policies. In
3240 a deployed SOA-based solution, it may well be that different aspects of the management of a given
3241 service are managed by different management services. For example, the life-cycle management of
3242 services often involves managing service versions. Managing quality of service is often very specific to
3243 the service itself; for example, quality of service attributes for a video streaming service are quite different
3244 to those for a banking system.

3245 Additional concepts of management also apply to IT management.

3246 5.3.2.2 Network Management

3247 Network management deals with the maintenance and administration of large scale physical networks
3248 such as computer networks and telecommunication networks. Specifics of the networks may affect
3249 service interactions from performance and operational perspectives.

3250 Network and related system management executes a set of functions required for controlling, planning,
3251 deploying, coordinating, and monitoring the distributed services in the SOA ecosystem. However, while
3252 recognizing their importance, the specifics of systems management or network management are out of
3253 scope for this Reference Architecture Foundation.

3254 **5.3.2.3 Security Management**

3255 Security Management includes identification of roles, permissions, access rights, and policy attributes
3256 defining security boundaries and events that may trigger a security response.

3257 Security management within a SOA ecosystem is essential to maintaining the trust relationships between
3258 participants residing in different ownership domains. Security management must consider not just the
3259 internal properties related to interactions between participants but ecosystem properties that preserve the
3260 integrity of the ecosystem from external threats.

3261 **5.3.2.4 Usage Management**

3262 Usage Management is concerned with how resources are used, including:

- 3263 • how the resource is accessed, who is using the resource, and the state of the resource (access
3264 properties);
 - 3265 • controlling or shaping demand for resources to optimize the overall operation of the ecosystem
3266 (demand properties);
 - 3267 • with assigning costs to the use of resources and distributing those cost assignments to the
3268 participants in an appropriate manner (financial properties).
- 3269

3270 **5.3.3 Management and Governance**

3271 The primary role of governance in the context of a SOA ecosystem is to foster an atmosphere of
3272 predictability, trust, and efficiency, and it accomplishes this by allowing the stakeholders to negotiate and
3273 set the key policies that govern the running of the SOA-based solution. Recall that in an ecosystems
3274 perspective, the goal of governance is less to have complete fine-grained control but more to enable the
3275 individual participants to work together.

3276 Policies for a SOA ecosystem will tend to focus on the rules of engagement between participants; for
3277 example, what kinds of interactions are permissible, how disputes are resolved, etc. While governance
3278 may primarily focus on setting policies, management will focus on the realization and enforcement of
3279 policies. Effective management in the SOA ecosystem requires an ability for governance to understand
3280 the consequences of its policies, guidelines, and principles, and to adjust those as needed when
3281 inconsistencies or ambiguity become evident from the operation of the management functions. This
3282 understanding and adjustment must be facilitated by the results of management and so the mechanisms
3283 for providing feedback from management into governance must exist.

3284 Governance operates via specialized activities and, thus, should be managed itself. Governance policies
3285 are included in the Governance Framework and Processes, and driven by the enterprise business model,
3286 business objectives and strategies. Where corporate management policies exist, these are usually guided
3287 and directed by the corporate executives. In peer relationships, governance policies are set by either an
3288 external entity and accepted by the peers or by the peers themselves. This creates the appropriate
3289 authoritative level for the policies used for the management of the Governance Framework and
3290 Processes. Management to operationalize governance controls the life-cycle of the governing policies,
3291 including procedures and processes, for modifying the Governance Framework and Processes.

3292 **5.3.4 Management and Contracts**

3293 **5.3.4.1 Management for Contracts and Policies**

3294 As we noted above, management can often be viewed as the application of contracts and individual
3295 policies to ensure the smooth running of the SOA ecosystem. Policies and service contracts specify the
3296 service characteristics that have to be monitored, analysed and managed and play an important role as
3297 the guiding constraints for management, as well as being artifacts (e.g., policy and contractual
3298 documents) that also need to be managed.

3299 **5.3.4.2 Contracts**

3300 As described in sections “Participation in a SOA Ecosystem view” and “Realization of a SOA Ecosystem
3301 view”, there are several types of contractual information in the SOA ecosystem. From the management
3302 perspective, three basic types of the contractual information relate to:

- 3303 • relationship between service provider and consumer;
- 3304 • communication with the service;
- 3305 • control of the quality of the service execution.

3306 When a consumer prepares to interact with a service, the consumer and the service provider must come
3307 to an agreement on the service features and characteristics that will be provided by the service and made
3308 available to the consumer. This agreement is known as a service contract.

3309 **Service Contract**

3310 An implicit or explicit documented agreement between the service consumer and service provider
3311 about the use of the service based on

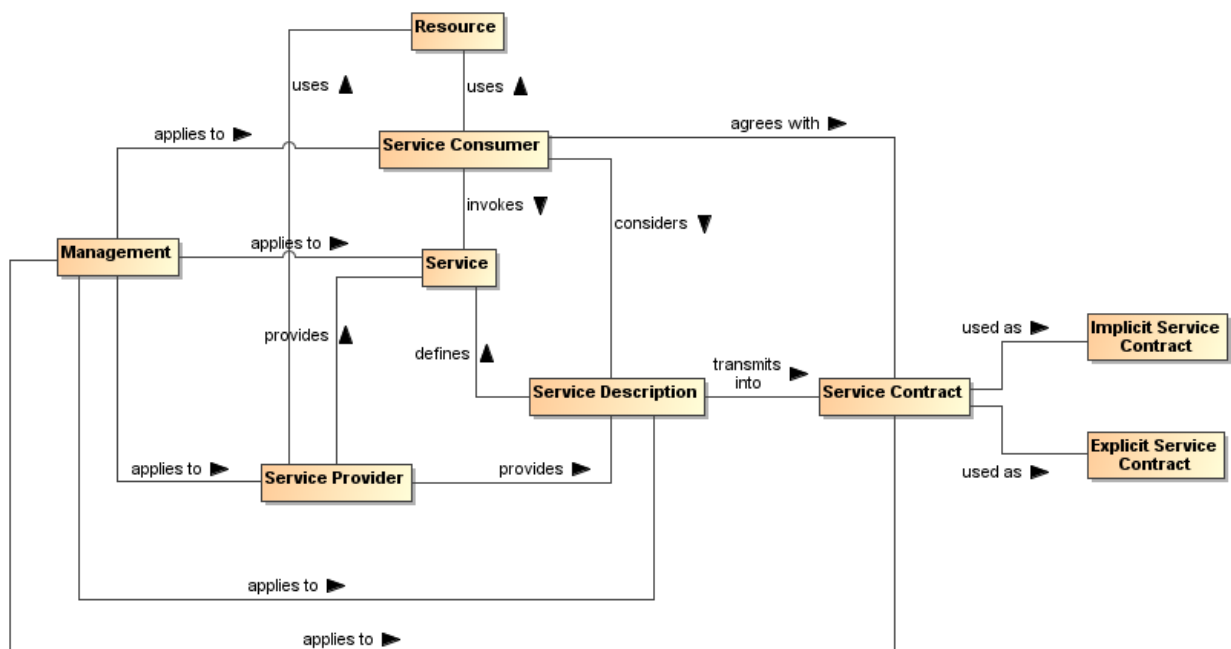
- 3312 • the commitment by a service provider to provide service functionality and results consistent
3313 with identified real world effects and
- 3314 • the commitment by a service consumer to interact with the service per specific means and
3315 per specified policies,

3316 where both consumer and provider actions are in the manner described in the service description.

3317 The service description provides the basis for the service contract and, in some situations, may be used
3318 as an implicit default service contract. In addition, the service description may set mandatory aspects of a
3319 service contract, e.g. for security services, or may specify acceptable alternatives. As an example of
3320 alternatives, the service description may identify which versions of a vocabulary will be recognized, and
3321 the specifics of the contract are satisfied when the consumer uses one of the alternatives. Another
3322 alternative could have a consumer identify a policy they require be satisfied, e.g. a standard privacy policy
3323 on handling personal information, and a provider that is prepared to accept a policy request would
3324 indicate acceptance as part of the service contract by continuing with the interaction. In each of these
3325 cases, the actions of the participants are consistent with an implicit service contract without the existence
3326 of a formal agreement between the participants.

3327 In the case of business services, it is anticipated that the service contract may take an explicit form and
3328 the agreement between business consumer and business service provider is formalized. Formalization
3329 requires up-front interactions between service consumer and service provider. In many business
3330 interactions, especially between business organizations within or across corporate boundaries, a
3331 consumer needs a contractual assurance from the provider or wants to explicitly indicate choices among
3332 alternatives, e.g., only use a subset of the business functionality offered by the service and pay a

3333 prorated cost.



3334
3335 *Figure 43 - Management of the service interaction*

3336 Consequently, an implicit service contract is an agreement (1) on the consumer side with the terms,
3337 conditions, features and interaction means specified in the service description "as is" or (2) a selection
3338 from alternatives that are made available through mechanisms included in the service description, and
3339 neither of these require any a priori interactions between the service consumer and the service provider.
3340 An explicit service contract always requires a form of interaction between the service consumer and the
3341 service provider prior to the service invocation. This interaction may regard the choice or selection of the
3342 subset of the elements of the service description or other alternatives introduced through the formal
3343 agreement process that would be applicable to the interaction with the service and affect related joint
3344 action.

3345 Any form of explicit contract couples the service consumer and provider. While explicit contracts may be
3346 necessary or desirable in some cases, such as in supply chain management, commerce often uses a mix
3347 of implicit and explicit contracts, and a service provider may offer (via service description) a conditional
3348 shift from implicit to explicit contract. For example, Twitter offers an implicit contract on the use of its APIs
3349 to any application with the limit on the amount of service invocations; if the application needs to use more
3350 invocations, one has to enter into the explicit fee-based contract with the provider. A case where an
3351 implicit contract transforms into an explicit contract may be illustrated when one buys a new computer and
3352 it does not work. The buyer returns the computer for repair under the manufacturer's warranty as stated
3353 by an implicit purchase contract. However, if the repair does not fix the problem and the seller offers an
3354 upgraded model in replacement, the buyer may agree to an explicit contract that limits the rights of the
3355 buyer to make the explicit agreement public.

3356 Control of the quality of the service execution, often represented as a service level agreement (SLA), is
3357 performed by service monitoring systems and includes both technical and operational business controls.
3358 SLA is a part of the service contract and, because of the individual nature of such contracts, may vary
3359 from one service contract to another, even for the same consumer. Typically, a particular SLA in the
3360 service contract is a concrete instance of the SLA declared in the service description.

3361 Management of the service contracts is based on management policies that may be mentioned in the
3362 service description and in the service contracts. Management of the service contracts is mandatory for
3363 consumer relationship management. In the case of explicit service contracts, the contracts have to be
3364 created, stored, maintained, reviewed/controlled and archived/destroyed as needed. All the activities are
3365 management concerns. Explicit service contracts may be stored in specialized repositories that provide
3366 appropriate level of security.

3367 Management of the service interfaces is based on several management policies that regulate
3368 • availability of interfaces specified in the service contracts,
3369 • accessibility of interfaces,
3370 • procedures for interface changes,
3371 • interface versions as well as the versions of all parts of the interfaces,
3372 • traceability of the interfaces and their versions back to the service description document.

3373 Management of the SLA is integral to the management of service monitoring and operational service
3374 behavior at run-time. An SLA usually enumerates service characteristics and expected performances of
3375 the service. Since an SLA carries the connotation of a “promise”, monitoring is needed to know if the
3376 promise is being kept. Existence of an SLA itself does not guarantee that the consumer will be provided
3377 with the service level specified in the service contract.

3378 The use of an SLA in a SOA ecosystem can be wider than just an agreement on technical performances.
3379 An SLA may contain remedies for situations where the promised service cannot be maintained, or the
3380 real world effect cannot be achieved due to developments subsequent to the agreement. A service
3381 consumer that acts accordingly to realize the real world effect may be compensated for the breach of the
3382 SLA if the effect is not realized.

3383 Management of the SLA includes, among others, policies to change, update, and replace the SLA. This
3384 aspect concerns service Execution Context because the business logic associated with a defined
3385 interface may differ in different Execution Contexts and affect the overall performance of the service.

3386 **5.3.4.3 Policies**

3387 "Although provision of management capabilities enables a service to become manageable, the extent and
3388 degree of permissible management are defined in management policies that are associated with the
3389 services. Management policies are used to define the obligations for, and permissions to, managing the
3390 service" [WSA]. Management policies, in essence, are the realization of governing rules and regulations.
3391 As such, some management policies may target services while other policies may target the management
3392 of the services.

3393 In practice, a policy without any means of enforcing it is vacuous. In the case of management policy, we
3394 rely on a management infrastructure to realize and enforce management policy.

3395 **5.3.4.4 Service Description and Management**

3396 The service description identifies several management objects such as a set of service interfaces and
3397 related set of SLAs. Service behavioral characteristics and performances specified in the SLA depend on
3398 the interface type and its Execution Context. In the service description, a service consumer can find
3399 references to management policies, SLA metrics, and the means of accessing measured values that
3400 together increase assurance in the service quality. At the same time, service description is an artifact that
3401 needs to be managed.

3402 In the SOA ecosystem, the service description is the assembled information that describes the service but
3403 it may be reported or displayed in different presentations. While each separate version of the service has
3404 one and only one service description, different categories of service consumers may focus their interests
3405 on different aspects of the service description. Thus, the same service description may be displayed not
3406 only in different languages but also with different cultural and professional accents in the content.

3407 New service description may be issued to reflect changes and update in the service. If the change in the
3408 service does not affect its service description, the new service version may have the same service
3409 description as the previous version except for the updated version identifier. For example, a service
3410 description may stay the same if bugs were fixed in the service. However, if a change in the service
3411 influences any aspects of the service quality that can affect the real world effect resulting from
3412 interactions with the service, the service description must reflect this change even if there are no changes
3413 to the service interface.

3414 Management of service description and related explicit service contracts is essential for delivery of the
3415 service to the consumer satisfaction. This management can also prevent business problems rooted in
3416 poor communication between the service consumers and the service providers.

3417 Thus, management of service description contains, among others, management of the service description
3418 presentations, the life-cycles of the service descriptions, service description distribution practices and
3419 storage of the service descriptions and related service contracts. Collections of service descriptions in
3420 the enterprise may manifest a need for specialised registries and/or repositories. Depending on the
3421 enterprise policies, an allocation of purposes and duties of registries and repositories may vary but this
3422 topic is beyond the current scope.

3423 **5.3.5 Management for Monitoring and Reporting**

3424 The successful application of management relies on the monitoring and reporting aspects of management
3425 to enable the control aspect. Monitoring in the context of management consists of measuring values of
3426 managed aspects and evaluating that measurement in relationship to some expectation. Monitoring in a
3427 SOA ecosystem is enabled through the use of mechanisms by resources for exposing managed aspects.
3428 In the SOA framework, this mechanism may be a service for obtaining the measurement. Alternatively,
3429 the measurement may be monitored by means of event generation containing updated values of the
3430 managed aspect.

3431 Approaches to monitoring may use a polling strategy in which the measurements are requested from
3432 resources in periodic intervals, in a pull strategy in which the measurements are requested from
3433 resources at random times, or in a push strategy in which the measurements are supplied by the resource
3434 without request. The push strategy can be used in a periodic update approach or in an “update on
3435 change” approach. Management services must be capable of handling these different approaches to
3436 monitoring.

3437 Reporting is the complement to monitoring. Where monitoring is responsible for obtaining measurements,
3438 reporting is responsible for distributing those measurements to interested stakeholders. The separation
3439 between monitoring and reporting is made to include the possibility that data obtained through monitoring
3440 might not be used until an event impacting the ecosystem occurs or the measurement requires further
3441 processing to be useful. In the SOA framework, reporting is provided using services for requesting
3442 measurement reports. These reports may consist of raw measurement data, formatted collections of data,
3443 or the results of analysis performed on measurement data from collections of different managed aspects.
3444 Reporting is also used to support logging and auditing capabilities, where the reporting mechanisms
3445 create log or audit entries.

3446 **5.3.6 Management for Infrastructure**

3447 All of the properties, policies, interactions, resources, and management are only possible if a SOA
3448 ecosystem infrastructure provides support for managed capabilities. Each managed capability imposes
3449 different requirements on the capabilities supplied by the infrastructure in SOA ecosystem and requires
3450 that those capabilities be usable as services or at the very least be interoperable.

3451 While not providing a full list of infrastructural elements of a SOA ecosystem, we list some examples here:

- 3452 1. Registries and repositories for services, policies, and related descriptions and contracts
- 3453 2. Synchronous and asynchronous communication channels for service interactions (e.g., network,
3454 e-mail, message routing with ability of mediating transport protocols, etc.)
- 3455 3. Recovery capabilities
- 3456 4. Security controls

3457 A SOA ecosystem infrastructure, enabling service management, should also support:

- 3458 1. Management enforcement and control means
- 3459 2. Monitoring and SLA validation controls
- 3460 3. Testing and Reporting capabilities

3461 The combination of manageability capabilities and infrastructure elements constitutes a certain level of
3462 SOA management maturity. While several maturity models exist, this topic is out of the scope of the
3463 current document.

3464 5.3.7 Architectural Implication of the Management Model

3465 5.4 SOA Testing Model

3466 Testing for SOA combines the typical challenges of software testing and certification with the additional
3467 needs of accommodating the distributed nature of the **resources**, the greater access of a more
3468 unbounded consumer population, and the desired flexibility to create new solutions from existing
3469 components over which the solution developer has little if any control. The **purpose** of testing is to
3470 demonstrate a required level of reliability, correctness, and effectiveness that enable prospective
3471 consumers to have adequate confidence in using a service. Adequacy is defined by the consumer based
3472 on the consumer's needs and context of use. Absolute correctness and completeness cannot be proven
3473 by testing; however, for SOA, it is critical for the prospective consumer to know what testing has been
3474 performed, how it has been performed, and what were the results.

3475 5.4.1 Traditional Software Testing as Basis for SOA Testing

3476 SOA services are largely software artifacts and can leverage the body of experience that has evolved
3477 around software testing. IEEE-829 [IEEE-829] specifies the basic set of software test documents while
3478 allowing flexibility for tailored use. As such, the document structure can also provide guidance to SOA
3479 testing.

3480 IEEE-829 covers test specification and test reporting through use of the following document types:

- 3481 • *Test plan* documenting the scope (what is to be tested, both which entity and what features of the
3482 entity), the approach (how it is tested), and the needed **resources** (who does the testing, for how
3483 long), with details contained in the:
- 3484 • *Test design specification*: features to be tested, test conditions (e.g. test cases, test procedures
3485 needed) and expected results (criteria for passing test); entrance and exit criteria
- 3486 • *Test case specification*: test data used for input and expected output
- 3487 • *Test procedure specification*: steps required to run the test, including any set-up preconditions
- 3488 • *Test item transmittal* to identify the test items being transmitted for testing
- 3489 • *Test log* to record what occurred during test, i.e. which tests run, who ran, what order, what
3490 happened
- 3491 • *Test incident report* to capture any **event** that happened during test which requires further
3492 investigation
- 3493 • *Test summary* as a management report summarizing test run and results, conclusions

3494 In summary, IEEE-829 captures (1) what was tested, (2) how it was tested, e.g. the test procedure used,
3495 and (3) the results of the test.

3496 5.4.1.1 Types of Testing

3497 There are numerous aspects of testing that, in total, work to establish that an entity is (1) built as required
3498 per policies and related specifications prescribed by the entity's owner, and (2) delivers the functionality
3499 required by its intended users. This is often referred to as verification and validation.

3500 Policies, as described in Section 4.4, that are related to testing may prescribe but are not limited to the
3501 business processes to be followed, the standards with which an implementation must comply, and the
3502 **qualifications** of and restrictions on the users. In addition to the functional requirements prescribing what
3503 an entity does, there may also be non-functional performance and/or quality metrics that state how well
3504 the entity does it. The relation of these policies to SOA testing is discussed further below.

3505 The identification of policies is the purview of governance (section 5.1) and the assuring of compliance
3506 (including response to noncompliance) with policies is a matter for management (section 5.3).

3507 5.4.1.2 Range of Test Conditions

3508 Test conditions and expected responses are detailed in the test case specification. The test conditions
3509 should be designed to cover the areas for which the entity's response must be documented and may
3510 include:

- 3511 • nominal conditions
- 3512 • boundaries and extremes of expected conditions
- 3513 • breaking point where the entity has degraded below a certain level or has otherwise ceased
- 3514 effective functioning
- 3515 • random conditions to investigate unidentified dependencies among combinations of conditions
- 3516 • errors conditions to test error handling

3517 The specification of how each of these conditions should be tested for SOA resources, including the
3518 infrastructure elements of the SOA ecosystem, is beyond the scope of this document but is an area that
3519 evolves along with operational SOA experience.

3520 **5.4.1.3 Configuration Management of Test Artifacts**

3521 The test item transmittal provides an unambiguous identification of the entity being tested, thus
3522 REQUIRING that the configuration of the entity is appropriately tracked and documented. In addition, the
3523 test documents (such as those specified by IEEE-829) MUST also be under a documented and
3524 appropriately audited configuration management process, as should other [resources](#) used for testing.
3525 The description of each artifact would follow the general description model as discussed in section
3526 4.1.1.1; in particular, it would include a version number for the artifact and reference to the documentation
3527 describing the versioning scheme from which the version number is derived.

3528 **5.4.2 Testing and the SOA Ecosystem**

3529 Testing of SOA artifacts for use in the SOA ecosystem differs from traditional software testing for several
3530 reasons. First, a highly touted benefit of SOA is to enable unanticipated consumers to make use of
3531 services for unanticipated purposes. Examples of this could include the consumer using a service for a
3532 result that was not considered the primary one by the provider, or the service may be used in combination
3533 with other services in a scenario that is different from the one considered when designing for the initial
3534 target consumer community. It is unlikely that a new consumer will push the services back to anything
3535 resembling the initial test phase to test the new use, and thus additional paradigms for testing are
3536 necessary. Some testing may depend on the availability of test resources made available as a service
3537 outside the initial test community, while some testing is likely to be done as part of limited use in the
3538 operational setting. The potential [responsibilities](#) related to such "consumer testing" is discussed further
3539 below.

3540 Secondly, in addition to consumers who interact with a service to realize the described [real world effects](#),
3541 the developer community is also intended to be a consumer. In the SOA vision of reuse, the developer
3542 composes new solutions using existing services, where the existing services provides access to some
3543 desired [real world effects](#) that are needed by the new solution. The new solution is a consumer of the
3544 existing services, enabling repeated interactions with the existing services playing the role of reusable
3545 components. Note, those components are used at the locations where they individually reside and are not
3546 typically duplicated for the new solution. The new solution may itself be offered as a SOA service, and a
3547 consumer of the service composition representing the new solution may be totally unaware of the
3548 component services being used. (See section 4.3.4 for further discussion on service compositions.)

3549 Another difference from traditional testing is that the distributed, unbounded nature of the SOA ecosystem
3550 makes it unlikely to have an isolated test environment that duplicates the operational environment. A
3551 traditional testing approach often makes use of a test system that is identical to the eventual operational
3552 system but isolated for testing. After testing is successfully completed, the tested entity would be
3553 migrated to the operational environment, or the test environment may be delivered as part of the system
3554 to become operational. This is not feasible for the SOA ecosystem as a whole.

3555 SOA services must be testable in the environment and under the conditions that can be encountered in
3556 the operational SOA ecosystem. As the ecosystem is in a state of constant change, so some level of
3557 testing is continuous through the lifetime of the service, leveraging utility services used by the ecosystem
3558 infrastructure to monitor its own health and respond to situations that could lead to degraded
3559 performance. This implies the test resources must incorporate aspects of the SOA paradigm, and a
3560 category of services may be created to specifically support and enable effective monitoring and
3561 continuous testing for [resources](#) participating in the SOA ecosystem.

3562 While SOA within an enterprise may represent a more constrained and predictable operational
3563 environment, the composability and unanticipated use aspects are highly touted within the enterprise.
3564 The expanded perspective on testing may not be as demanding within an enterprise but fuller
3565 consideration of the ecosystem enables the enterprise to be more responsive should conditions change.

3566 **5.4.3 Elements of SOA Testing**

3567 IEEE-829 identifies fundamental aspects of testing, and many of these should carry over to SOA testing:
3568 in particular, the identification of what is to be tested, how it is to be tested, and by whom the testing is to
3569 be done. While IEEE-829 identifies a suggested document tree, the availability of these documents in the
3570 SOA ecosystem is discussed below.

3571 **5.4.3.1 What is to be Tested**

3572 The focus of this discussion is the SOA service. It is recognized that the infrastructure components of
3573 any SOA environment are likely to also be SOA services and, as such, falls under the same testing
3574 guidance. Other resources that contribute to a SOA environment may not be SOA services, but are
3575 expected to satisfy the intent if not the letter of guidance presented here. Specific differences for such
3576 [resources](#) are as yet largely undefined and further elaboration is beyond the scope of the SOA-RAF.

3577 The following discussion often focuses on a singular SOA service but it is implicit that any service may be
3578 a composite of other services. As such, testing the functionality of a composite service may effectively be
3579 testing an end-to-end business process that is being provided by the composite service. If new versions
3580 are available for the component services, appropriate end-to-end testing of the composite may be
3581 required in order to verify that the composite functionality is still adequately provided. The level of
3582 required testing of an updated composite depends on policies of those providing the service, policies of
3583 those using the service, and mission criticality of those depending on the service results.

3584 The SOA service to be tested MUST be unambiguously identified as specified by its applicable
3585 configuration management scheme. Specifying such a scheme is beyond the scope of the SOA-RAF
3586 other than to say the scheme should be documented and itself under configuration management.

3587 **5.4.3.1.1 Origin of Test Requirements**

3588 In the Service Description model (Figure 13), the aspects of a service that need to be described are:

- 3589 • the service functionality and technical assumptions that underlie the functionality;
- 3590 • the policies that describe conditions of use;
- 3591 • the service interface that defines information exchange with the service;
- 3592 • service reachability that identifies how and where message exchange is to occur; and
- 3593 • metrics access for any [participant](#) to have information on how a service is performing.

3594 Service testing must provide adequate assurance that each of these aspects is operational as defined.

3595 The information in the service description comes from different sources. The functionality is defined
3596 through whatever process identifies needs and the community for which these needs are addressed. The
3597 process may be ad hoc as serves the prospective service owner or strictly governed, but defining the
3598 functionality is an essential first step in development. It is also an early and ongoing focus of testing to
3599 ensure the service accurately reflects the described functionality and the described functionality
3600 accurately addresses the consumer needs.

3601 Policies define the conditions of development and conditions of use for a service and are typically
3602 specified as part of the governance process. Policies constraining service development, such as coding
3603 standards and best practices, require appropriate testing and auditing during development to ensure
3604 compliance. While the governance process identifies development policies, these are likely to originate
3605 from the technical community responsible for development activities. Policies that define conditions of
3606 use often define business practices that service owners and providers or those responsible for the SOA
3607 infrastructure want followed. These policies are initially tested during service development and are
3608 continuously monitored during the operational lifetime of the service.

3609 The testing of the service interface and service reachability are often related but essentially reflect
3610 different motivations and needs. The service interface is specified as a joint product of the service

3611 owners and providers who define service functionality, the prospective consumer community, the service
3612 developer, and the governance process. The semantics of the information model must align with the
3613 semantics of those who consume the service in order for there to be meaningful exchange of information.
3614 The structure of the information is influenced by the consumer semantics and the requirements and
3615 constraints of the representation as interpreted by the service developer. The service process model that
3616 defines actions which can be performed against a service and any temporal dependencies derive from
3617 the defined functionality and may be influenced by the development process. Any of these constraints
3618 may be identified and expressed as policy through the governance process.

3619 Service reachability conditions are the purview of the service provider who identifies the service endpoint
3620 and the protocols recognized at the endpoint. These may be constrained by governance decisions on
3621 how endpoint addresses may be allocated and what protocols should be used.

3622 While the considerations for defining the service interface derive from several sources, testing of the
3623 service interface is more straightforward and isolated in the testing process. At any point where the
3624 interface is modified or exposes a new [resource](#), the message exchange should be monitored both to
3625 ensure the message reaches its intended destination and it is parsed correctly once received. Once an
3626 interface has been shown to function properly, it is unlikely to fail later unless something fundamental to
3627 the service changes.

3628 The service interface is also tested when the service endpoint changes. Testing of the endpoint ensures
3629 message exchange can occur at the time of testing and the initial testing shows the interface is being
3630 processed properly at the new endpoint. Functioning of a service endpoint at one time does not
3631 guarantee it is functioning at another time, e.g. the server with the endpoint address may be down,
3632 making testing of service reachability a continual monitoring function through the life of the service's use
3633 of the endpoint. Also, while testing of the service endpoint is a necessary and most commonly noted part
3634 of the test regiment, it is not in itself sufficient to ensure the other aspects of testing discussed in this
3635 section.

3636 Finally, governance is impossible without the collection of metrics against which service behavior can be
3637 assessed. Metrics are also a key indicator for consumers to decide if a service is adequate for their
3638 needs. For instance, the average response time or the recent availability can be determining factors even
3639 if there are no rules or regulations promulgated through the governance process against which these
3640 metrics are assessed. The available metrics are a combination of those expected by the consumer
3641 community and those mandated through the governance process. The total set of metrics will evolve
3642 over time with SOA experience. Testing of the services that gather and provide access to the metrics will
3643 follow testing as described in this section, but for an individual service, testing will ensure that the metrics
3644 access indicated in the service description is accurate.

3645 The individual test requirements highlight aspects of the service that testing must consider but testing
3646 must establish more than isolated behavior. The emphasis is the holistic results of interacting with the
3647 service in the SOA environment. Recall that the execution context is the set of agreements between a
3648 consumer and a provider that define the conditions under which service interaction occurs. The
3649 agreements are expected to be predominantly the acceptance of the standard conditions as enumerated
3650 by the service provider, but it may include the identification of alternate conditions that will govern the
3651 interaction.

3652 For example, the provider may prefer a policy where it can sell the contact information of its consumers
3653 but will honor the request of a consumer to keep such information private. The identification of the
3654 alternate privacy policy is part of the execution context, and it is the application of and compliance with
3655 this policy that operational monitoring will attempt to measure. The collection of metrics showing this
3656 condition is indeed met when chosen is considered part of the ongoing testing of the service.

3657 Other variations in the execution context also require monitoring to ensure that different combinations of
3658 conditions perform together as desired. For example, if a new privacy policy takes additional [resources](#) to
3659 apply, this may affect quality of service and propagate other effects. These could not be tested during the
3660 original testing if the alternate policy did not exist at that time.

3661 **5.4.3.1.2 Testing Against Non-Functional Requirements**

3662 Testing against non-functional requirements constitutes testing of business usability of the service. In a
3663 marketplace of services, non-functional characteristics may be the primary differentiator between services
3664 that produce essentially the same [real world effects](#).

3665 As noted in the previous section, non-functional characteristics are often associated with policies or other
3666 terms of use and may be collected in service level contracts offered by the service providers. Non-
3667 functional requirements may also reflect the network and hardware infrastructure that support
3668 communication with the service, and changes may impact quality of service. The [service consumer](#) and
3669 even the service provider may not be aware of all such infrastructure changes but the changes may
3670 manifest in shared states that impact the usability of the service.

3671 In general, a change in the non-functional requirements results in a change to the execution context, but
3672 as with any collection of information that constitutes a description, the execution context is unable to
3673 explicitly capture all non-functional requirements that may apply. A change in non-functional
3674 requirements, whether explicitly part of the execution context or an implicit contributor, may require
3675 retesting of the service even if its functionality and the implementation of the functionality has not
3676 changed. Depending on the circumstances, retesting may require a formal recertifying of end-to-end
3677 behavior or more likely will be part of the continuous monitoring that applies throughout the service
3678 lifetime.

3679 **5.4.3.1.3 Testing Content and the Interests of Consumers**

3680 As noted in section 5.4.1.1, testing may involve verification of conformance with respect to policies and
3681 technical specifications and validation with respect to sufficiency of functionality to meet some prescribed
3682 use. It may also include demonstration of performance and quality aspects. For some of these items,
3683 such as demonstrating the business processes followed in developing the service or the use of standards
3684 in implementing the service, the testing or relevant auditing is done internal to the service development
3685 process and follows traditional software testing and quality assurance. If it is believed of value to
3686 potential consumers, information about such testing could be included in the service description.
3687 However, it is not required that all test or compliance artifacts be available to consumers, as many of the
3688 details tested may be part of the opacity of the service implementation.

3689 Some aspects of the service being tested will reflect directly on the [real world effects](#) realized through
3690 interaction with the service. In these cases, it is more likely that testing results will be directly relevant to
3691 potential consumers. For example, if the service was designed to correspond to certain elements of a
3692 business process or that a certain workflow is followed, testing should verify that the [real world effects](#)
3693 reflect that the business process or workflow were satisfactorily captured.

3694 The testing may also need to demonstrate that specified conditions of use are satisfied. For example,
3695 policies may be asserted that require certain [qualifications](#) of or impose restrictions on the consumers
3696 who may interact with the service. The service testing must demonstrate that the service independently
3697 enforces the policies or it provides the required information exchanges with the SOA ecosystem so other
3698 [resources](#) can ensure the specified conditions.

3699 The completeness of the testing, both in terms of the features tested and the range of parameters for
3700 which response is tested, depends on the context of expected use: the more critical the use, the more
3701 complete the testing. There are always limits on the [resources](#) available for testing, if nothing else than
3702 the service must be available for use in a finite amount of time.

3703 This again emphasizes the need for adequate documentation to be available. If the original testing is
3704 very thorough, it may be adequate for less demanding uses in the future. If the original testing was more
3705 constrained, then well-documented test results establish the foundation on which further testing can be
3706 defined and executed.

3707 **5.4.3.2 How Testing is to be Done**

3708 Testing should follow well-defined methodologies and, if possible, should reuse test artifacts that have
3709 proven generally useful for past testing. For example, IEEE-829 notes that test cases are separated from
3710 test designs to allow for use in more than one design and to allow for reuse in other situations. In the

3711 SOA ecosystem, description of such artifacts, as with description of a service, enables awareness of the
3712 item and describes how the artifact may be accessed or used.

3713 As with traditional testing, the specific test procedures and test case inputs are important so the tests are
3714 unambiguously defined and entities can be retested in the future. Automated testing and regression
3715 testing may be more important in the SOA ecosystem in order to re-verify a service is still acceptable
3716 when incorporated in a new use. For example, if a new use requires the services to deal with input
3717 parameters outside the range of initial testing, the tests could be rerun with the new parameters. If the
3718 testing resources are available to consumers within the SOA ecosystem, the testing as designed by test
3719 professionals could be consumed through a service accessed by consumers, and their results could
3720 augment those already in place. This is discussed further in the next section.

3721 5.4.3.3 Who Performs the Testing

3722 As with any software, the first line of testing is unit testing done by software developers. It is likely that
3723 initial testing will be done by those developing the software but may also be done independently by other
3724 developers. For SOA development, unit testing is likely confined to a development sandbox isolated from
3725 the SOA ecosystem.

3726 SOA testing will differ from traditional software testing in that testing beyond the development sandbox
3727 must incorporate aspects of the SOA ecosystem, and those doing the testing must be familiar with both
3728 the characteristics and responses of the ecosystem and the tools, especially those available as services,
3729 to facilitate and standardize testing. Test professionals will know what level of assurance must be
3730 established as the exposure of the service to the ecosystem and ecosystem to the service increases
3731 towards operational status. These test professionals may be internal resources to an organization or may
3732 evolve as a separate discipline provided through external contracting.

3733 As noted above, it is unlikely that a complete duplicate of the SOA ecosystem will be available for isolated
3734 testing, and thus use of ecosystem [resources](#) will manifest as a transition process rather than a step
3735 change from a test environment to an operational one. This is especially true for new composite services
3736 that incorporate existing operational services to achieve the new functionality. The test professionals will
3737 need to understand the available resources and the ramifications of this transition.

3738 As with current software development, a stage beyond work by test professionals will make use of a
3739 select group of typical users, commonly referred to as beta testers, to report on service response during
3740 typical intended use. This establishes fitness by the consumers, providing final validation of previously
3741 verified processes, requirements, and final implementation.

3742 In traditional software development, beta testing is the end of testing for a given version of the software.
3743 However, although the initial test phase can establish an appropriate level of confidence consistent with
3744 the designed use for the initial target consumer community, the operational service will exist in an
3745 evolving ecosystem, and later conditions of use may differ from those thought to be sufficient during the
3746 initial testing. Thus, operational monitoring becomes an extension of testing through the service lifetime.
3747 This continuous testing will attempt to ensure that a service does not consume an inordinate amount of
3748 ecosystem resources or display other behavior that degrades the ecosystem, but it will not undercover
3749 functional errors that may surface over time.

3750 As with any software, it is the responsibility of the consumers to consider the reasonableness of solutions
3751 in order to spot errors in either the software or the way the software is being used. This is especially
3752 important for consumers with unanticipated uses that may go beyond the original test conditions. It is
3753 unlikely the consumers will initiate a new round of formal testing unless the new use requires a
3754 significantly higher level of confidence in the service. Rather the consumer becomes a new extension to
3755 the testing regiment. Obvious testing would include a sanity check of results during the new use.
3756 However, if the details of legacy testing are associated with the service through the service description
3757 and if testing resources are available through automated testing services, then the new consumers can
3758 rerun and extend previous testing to include the extended test conditions. If the test results are
3759 acceptable, these can be added to the documentation of previous results and become the extended basis
3760 for future decisions by prospective consumers on the appropriateness of the service. If the results are not
3761 acceptable or in some way questionable, the responsible party for the service or testing professionals can
3762 be brought in to decide if remedial [action](#) is necessary.

3763 **5.4.3.4 How Testing Results are Reported**

3764 For any SOA service, an accurate reporting of the testing a service has undergone and the results of the
3765 testing is vital to consumers deciding whether a service is appropriate for intended use. Appropriateness
3766 may be defined by a consumer organization and require specific test regiments culminating in a
3767 certification; appropriateness could be established by accepting testing and certifications that have been
3768 conferred by others.

3769 The testing and certification information should be identified in the service description. Referring to the
3770 general description model of Figure 11, tests conducted by or under a request from the service owner
3771 (see [ownership](#) in section 3.1.3) would be captured under Annotations from Owners. Testing done by
3772 others, such as consumers with unanticipated uses, could be associated through Annotations from 3rd
3773 Parties. The annotations should clearly indicate what was tested, how the testing was done, who did the
3774 testing, and the testing results. The clear description of each of these artifacts and of standardized
3775 testing protocols for various levels of sophistication and completeness of testing would enable a common
3776 understanding and comparison of test coverage. It will also make it more straightforward to conduct and
3777 report on future testing, facilitating the maintenance of the service description.

3778 Consumer testing and the reporting of results raises additional issues. While stating who did the testing
3779 is mandatory, there may be formal requirements for authentication of the tester to ensure traceability of
3780 the testing claims. In some circumstances, persons or organizations would not be allowed to state testing
3781 claims unless the tester was an approved entity. In other cases, ensuring the tester had a valid email
3782 may be sufficient. In either case, it would be at the discretion of the potential consumer to decide what
3783 level of authentication was acceptable and which testers are considered authoritative in the context of
3784 their anticipated use.

3785 Finally, in a world of openly shared information, we would see an ever-expanding set of testing
3786 information as new uses and new consumers interact with a service. In reality, these new uses may
3787 represent proprietary processes or classified use that should only be available to authorized parties.
3788 Testing information, as with other elements of description, may require special access controls to ensure
3789 appropriate access and use.

3790 **5.4.4 Testing SOA Services**

3791 Testing of SOA services should be consistent with the SOA paradigm. In particular, testing resources
3792 and artifacts should be visible in support of service interaction between providers and consumers, where
3793 here the interaction is between the testing resource and the tester. In addition, the idea of opacity of the
3794 implementation should limit the details that need to be available for effective use of the test resources.
3795 Testing that requires knowledge of the internal structure of the service or its underlying capability should
3796 be performed as part of unit testing in the development sandbox, and should represent a minimum level
3797 of confidence before the service begins its transition to further testing and eventual operation in the SOA
3798 ecosystem.

3799 **5.4.4.1 Progression of SOA Testing**

3800 Software testing is a gradual exercise going from micro inspection to testing macro effects. The first step
3801 in testing is likely the traditional code reviews. SOA considerations would account for the distributed
3802 nature of SOA, including issues of distributed security and best practices to ensure secure resources. It
3803 would also set the groundwork for opacity of implementation, hiding programming details and simplifying
3804 the use of the service.

3805 Code review is likely followed by unit testing in a development sandbox isolated from the operational
3806 environment. The unit testing is done with full knowledge of the service internal structure and knowledge
3807 of resources representing underlying capabilities. It tests the interface to ensure exchanged messages
3808 are as specified in the service description and the messages can be parsed and interpreted as intended.
3809 Unit testing also verifies intended functionality and that the software has dealt correctly with internal
3810 dependencies, such as structure of a file system or access to other dedicated resources.

3811 Some aspects of unit testing require external dependencies be satisfied, and this is often done using
3812 mock objects to substitute for the external resources. In particular, it will likely be necessary to include

3813 mocks of existing operational services, both those provided as part of the SOA infrastructure and services
3814 from other providers.

3815 **Service Mock**

3816 A service mock is an entity that mimics some aspect of the performance of an operational service
3817 without committing to the [real world effects](#) that the operational service would produce.

3818 Mocks are discussed in detail in sections 5.4.4.3 and 5.4.4.4.

3819 After unit testing has demonstrated an adequate level of confidence in the service, the testing must
3820 transition from the tightly controlled environment of the development sandbox to an environment that
3821 more clearly resembles the operational SOA ecosystem or, at a minimum, the intended enterprise. While
3822 sandbox testing will use simple mocks of some aspects of the SOA environment, such as an interface to
3823 a security service without the security service functionality, the dynamic nature of SOA makes a full
3824 simulation infeasible to create or maintain. This is especially true when a new composite service makes
3825 use of operational services provided by others. Thus, at some point before testing is complete, the
3826 service will need to demonstrate its functionality by using resources and dealing with conditions that only
3827 exist in the full ecosystem or the intended enterprise. Some of these resources may still provide test
3828 interfaces -- more on this below -- but the interfaces will be accessible using the SOA environment and
3829 not just implemented for the sandbox.

3830 At this stage, the opacity of the service becomes important as the details of interacting with the service
3831 now rely on correct use of the service interface and not knowledge of the service internals. The workings
3832 of the service will only be observable through the [real world effects](#) realized through service interactions
3833 and external indications that conditions of use, such as user authentication, are satisfied. Monitoring the
3834 behavior of the service will depend on service interfaces that expose internal monitoring or provide
3835 required information to the SOA infrastructure monitoring function. The monitoring required to test a new
3836 service is likely to have significant overlap with the monitoring the SOA infrastructure includes to monitor
3837 its own health and to identify and isolate behavior outside of acceptable bounds. This is exactly what is
3838 needed as part of service testing, and it is reasonable to assume that the ecosystem transition includes
3839 use of operational monitoring rather than solely dedicated monitoring for each service being tested.

3840 Use of SOA monitoring resources during the explicit testing phase sets the stage for monitoring and a
3841 level of continual testing throughout the service lifetime.

3842 **5.4.4.2 Testing Traditional Dependencies vs. Service Interactions**

3843 A SOA service is not required to make use of other operational services beyond what may be required for
3844 monitoring by the ecosystem infrastructure. The service can implement hardcoded dependencies which
3845 have been tested in the development sandbox through the use of dedicated mocks. While coordination
3846 may be required with real data sources during integration testing, the dependencies can be constrained to
3847 things that can be tested in a more traditional manner. Policies can also be set to restrict access to pre-
3848 approved users, and thus the question of unanticipated users and unanticipated uses can be eliminated.
3849 Operational readiness can be defined in terms of what can be proven in isolated testing. While all this
3850 may provide more confidence in the service for its designed [purpose](#), such a service will not fully
3851 participate in the benefits or challenges of the ecosystem. This is akin to filling a swimming pool with sea
3852 water and having someone in the pool say they are swimming in the ocean.

3853 In considering the testing needed for a fully participating service, consider the example of a new
3854 composite service that combines the [real world effects](#) and complies with the conditions of use of five
3855 existing operational services. The developer of the composite service does not own any of the
3856 component services and has limited, if any, ability to get the distributed owners to do any customization.
3857 The developer also is limited by the principle of opacity to information comprising the service description,
3858 and does not know internal details of the component services. The developer of the composite service
3859 must use the component services as they exist as part of the SOA environment, including what is
3860 provided to support testing by new users. This introduces requirements for what is needed in the way of
3861 service mocks.

3862 5.4.4.3 Use of Service Mocks

3863 Service mocks enables the tested service to respond to specific features of an operational service that is
3864 being used as a component. It allows service testing to proceed without needing access to or with only
3865 limited engagement with the component service. Mocks can also mimic difficult to create situations for
3866 which it is desired to test the new service response. For composite services using multiple component
3867 services, mocks may be used in combination to function for any number of the components. Note, when
3868 using service mocks, it is important to remember that it is not the component service that is being tested
3869 (although anomalous behavior may be uncovered during testing) but the use of the component in the new
3870 composite.

3871 Individual service mocks can emphasize different features of the component service they represent but
3872 any given mock does not have to mimic all features. For example, a mock of the service interface can
3873 echo a sent message and demonstrate the message is reaching its intended destination. A mock could
3874 go further and parse the sent message to demonstrate the message not only reached its destination but
3875 was understood. As a final step, the mock could report back what actions would have been taken by the
3876 component service and what [real world effects](#) would result. If the response mimicked the operational
3877 response, functional testing could proceed as if the [real world effect](#) actually occurred.

3878 There are numerous ways to provide mock functionality. The service mock could be a simulation of the
3879 operational service and return simulated results in a realistic response message or event notification. It is
3880 also possible for the operational service to act as its own mock and simply not execute the commit stage
3881 of its functionality. The service mock could use a combination of simulation and service action without
3882 commit to generate a report of what would have occurred during the defined interaction with the
3883 operational service.

3884 As the service proceeds through testing, mocks should be systematically replaced by the component
3885 resources accessed through their operational interfaces. Before beta testing begins, end-to-end testing,
3886 i.e. proceeding from the beginning of the service interaction to the resulting real world results, should be
3887 accomplished using component resources via their operational interfaces.

3888 5.4.4.4 Providers of Service Mocks

3889 In traditional testing, it is often the test professionals who design and develop the mocks, but in the
3890 distributed world of SOA, this may not be efficient or desirable.

3891 In the development sandbox, it is likely the new service developer or test professionals working with the
3892 developer will create mocks adequate for unit testing. Given that most of this testing is to verify the new
3893 service is performing as designed, it is not necessary to have high fidelity models of other resources
3894 being accessed. In addition, given opacity of SOA implementation, the developer of the new service may
3895 not have sufficient detailed knowledge of a component service to build a detailed mock of the component
3896 service functionality. Sharing existing mocks at this stage may be possible but the mocks would need to
3897 be implemented in the sandbox, and for simple models it is likely easier to build the mock from scratch.

3898 As testing begins its transition to the wider SOA environment, mocks may be available as services. For
3899 existing resources, it is possible that an Open Source model could evolve where service mocks of
3900 available functions can be catalogued and used during initial interaction of the tested service and the
3901 operational environment. Widely used functions may have numerous service mocks, some mimicking
3902 detailed conditions within the SOA infrastructure. However, the Open Source model is less likely to be
3903 sufficient for specialty services that are not widely used by a large consumer community.

3904 The service developer is probably best qualified for also developing more detailed service mocks or for
3905 mock modes of operational services. This implies that in addition to their operational interfaces, services
3906 will routinely provide test interfaces to enable service mocks to be used as services. As noted above, a
3907 new service developer wanting to build a mock of component services is limited to the description
3908 provided by the component service developer or owner. The description typically will detail [real world](#)
3909 [effects](#) and conditions of use but will not provide implementation details, some of which may be
3910 proprietary. Just as important in the SOA ecosystem, if it becomes standard protocol for developers to
3911 create service mocks of their own services, a new service developer is only responsible for building his
3912 own mocks and can expect other mocks to be available from other developers. This reduces duplication
3913 of effort where multiple developers would be trying to build the same mocks from the same insufficient

3914 information. Finally, a service developer is probably best qualified to know when and how a service mock
3915 should be updated to reflect modified functionality or message exchange.

3916 It is also possible that testing organizations will evolve to provide high-fidelity test harnesses for new
3917 services. The harnesses would allow new services to plug into a test environment and would facilitate
3918 accessing mocks of component services. However, it will remain a constant challenge for such
3919 organizations to capture evolving uses and characteristics of service interactions in the real SOA
3920 environment and maintain the fidelity and accuracy of the test systems.

3921 **5.4.4.5 Fundamental Questions for SOA Testing**

3922 In order for the transition to the SOA operational environment to proceed, it is necessary to answer two
3923 fundamental questions:

- 3924 • Who provides what testing resources for the SOA operational environment, e.g. mocks of
3925 interfaces, mocks of functionality, monitoring tools?
- 3926 • What testing needs to be accomplished before operational environment resources can be
3927 accessed for further testing?

3928 The discussion in section 5.4.4.4 notes various levels of sophistication of service mocks and different
3929 communities are likely to be responsible for different levels. Section 5.4.4.4 advocates a significant role
3930 for service developers, but there needs to be community consensus that such mocks are needed and that
3931 service developers will agree to fulfilling this [role](#). There is also a need for consensus as to what tools
3932 should be available as services from the SOA infrastructure.

3933 As for use of the service mocks and SOA environment monitoring services, practical experience is
3934 needed upon which guidelines can be established for when a new service has been adequately tested to
3935 proceed with a greater level of exposure with the SOA environment. Malfunctioning services could cause
3936 serious problems if they cannot be identified and isolated. On the other hand, without adequate testing
3937 under SOA operational conditions, it is unlikely that problems can be uncovered and corrected before
3938 they reach an operational stage.

3939 As noted in section 5.4.4.2, some of these questions can be avoided by restricting services to more
3940 traditional use scenarios. However, such restriction will limit the effectiveness of SOA use and the result
3941 will resemble the constraints of traditional integration activities we are trying to move beyond.

3942 **5.4.5 Architectural Implications for SOA Testing**

3943 The discussion of SOA Testing indicates numerous architectural implications on the SOA ecosystem:

- 3944 • The distributed, boundary-less nature of the SOA ecosystem makes it infeasible to create and
3945 maintain a single mock of the entire ecosystem to support testing activities.
- 3946 • A standard suite of monitoring services needs to be defined, developed, and maintained. This
3947 should be done in a manner consistent with the evolving nature of the ecosystem.
- 3948 • Services should provide interfaces that support access in a test mode.
- 3949 • Testing resources must be described and their descriptions must be catalogued in a manner that
3950 enables their discovery and access.
- 3951 • Guidelines for testing and ecosystem access need to be established and the ecosystem must be
3952 able to enforce those guidelines asserted as policies.
- 3953 • Services should be available to support automated testing and regression testing.
- 3954 • Services should be available to facilitate updating service description by anyone who has
3955 performed testing of a service.

3956 6 Conformance

3957 This Reference Architecture Framework is an abstract architectural description of Service Oriented
3958 Architecture, which means that it is especially difficult to construct tests for conformance to the
3959 architecture. In addition, conformance to an architectural specification does not, by itself, guarantee any
3960 form of interoperability between multiple implementations.

3961 However, it *is* possible to decide whether or not a given architecture is conformant to an architectural
3962 description such as this one. In discussions of conformance we use the term **target architecture** to
3963 identify the (typically concrete) architecture that may be viewable as conforming to the abstract principles
3964 outlined in this document.

3965 Target Architecture

3966 A target architecture is an architectural description of a system that is intended to be viewed as
3967 conforming to the SOA-RAF.

3968 While we cannot guarantee interoperability between target architectures (or more specifically between
3969 applications and systems residing within the ecosystems of those target architectures), interoperability
3970 between target architectures is promoted by conformance to this Reference Architecture Framework as it
3971 reduces the semantic impedance mismatch between the different ecosystems.

3972 The primary measure of conformance is whether given concepts as described in document have
3973 corresponding concepts in the target architecture. Such a correspondence **MUST** honor the relationships
3974 identified within this document for the target architecture to be considered conforming.

3975 For example, in Section 3.1.3.1 we identify [resource](#) as a key concept. A [resource](#) is associated with an
3976 owner and a number of identifiers. For a target architecture to conform to the SOA-RAF, it must be
3977 possible to find corresponding concepts of [resource](#), identifier and owner within the target architecture:
3978 say *entity*, *token* and *user*. Furthermore, the relationships between *entity*, *token* and *user* **MUST** mirror
3979 the relationships between [resource](#), identifier and owner appropriately.

3980 Clearly, such correspondence is simpler if the terminology within the target architecture is identical to that
3981 in the SOA-RAF. But so long as the 'graph' of concepts and relationships is consistent, that is all that is
3982 required for the target architecture to conform to this Reference Architecture Framework.

3983 [EDITOR'S NOTE: The conformance section is not complete]

3984

3985

A. Acknowledgements

3986 The following individuals have participated in the work of the technical committee responsible for creation
3987 of this specification and are gratefully acknowledged:

3988 **Participants:**

- 3989 Chris Bashioum, MITRE Corporation
- 3990 Rex Brooks, Individual Member
- 3991 Peter Brown, Individual Member
- 3992 Scott Came, Search Group Inc.
- 3993 Joseph Chiusano, Booz Allen Hamilton
- 3994 Robert Ellinger, Northrop Grumman Corporation
- 3995 David Ellis, Sandia National Laboratories
- 3996 Jeff A. Estefan, Jet Propulsion Laboratory
- 3997 Don Flinn, Individual Member
- 3998 Anil John, Johns Hopkins University
- 3999 Ken Laskey, MITRE Corporation
- 4000 Boris Lublinsky, Nokia Corporation
- 4001 Francis G. McCabe, Individual Member
- 4002 Christopher McDaniels, USSTRATCOM
- 4003 Tom Merkle, Lockheed Martin Corporation
- 4004 Jyoti Namjoshi, Patni Computer Systems Ltd.
- 4005 Duane Nickull, Adobe Inc.
- 4006 James Odell, Associate
- 4007 Michael Poulin, Fidelity Investments
- 4008 Michael Stiefel, Associate
- 4009 Danny Thornton, Northrop Grumman
- 4010 Timothy Vibbert, Lockheed Martin Corporation
- 4011 Robert Vitello, New York Dept. of Labor

4012 The committee would particularly like to underline the significant contributions made by Rex Brooks, Jeff
4013 Estefan, Ken Laskey, Boris Lublinsky, Frank McCabe, Michael Poulin and Danny Thornton

4014 **B. Index of Defined Terms**

4015 The first page number refers to the first use of the term. The second, where necessary, refers to the page
4016 where the term is formally defined.

4017 Action

4018 Action Level Real World Effect

4019 Actor

4020 Architecture

4021 Architectural Description

4022 Authority

4023 Business Processes

4024 Capability

4025 Choreography

4026 Commitment

4027 Communicative Action

4028 Constitution

4029 Contract

4030 Delegate

4031 Description

4032 Endpoint

4033 Enterprise

4034 Governance

4035 Governance Framework

4036 Governance Processes

4037 Identifier

4038 Identity

4039 Joint Action

4040 Leadership

4041 Life-cycle manageability

4042 Logical Framework

4043 Management

4044 Management Policy

4045 Management Service

4046 Manageability Capability

4047 Message Exchange

4048 Model

4049 Obligation

4050 Objective

4051 Operations

4052 Orchestration

4053 Ownership

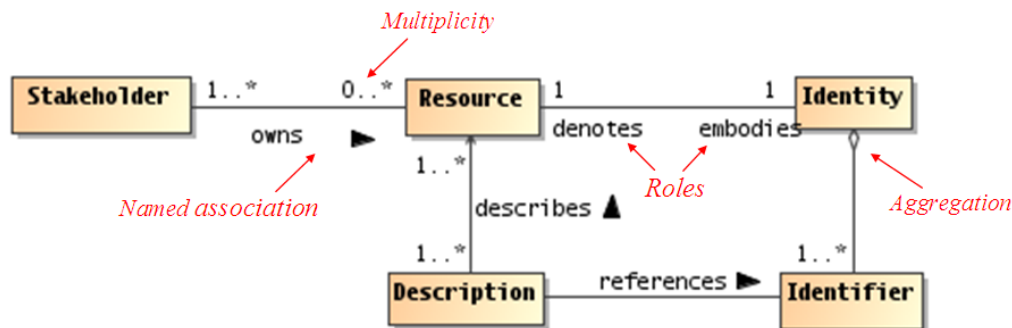
4054 Ownership Boundary
4055 Participant
4056 Peer
4057 Permission
4058 Policy
4059 Policy Conflict
4060 Policy Conflict Resolution
4061 Policy Constraint
4062 Policy Decision
4063 Policy Enforcement
4064 Policy Framework
4065 Policy Object
4066 Policy Ontology
4067 Policy Owner
4068 Policy Subject
4069 Presence
4070 Private State
4071 Protocol
4072 Public Semantics
4073 Qualification
4074 Real World Effect
4075 Regulation
4076 Resource
4077 Responsibility
4078 Right
4079 Risk
4080 Role
4081 Rule
4082 Security
4083 Semantic Engagement
4084 Service Action
4085 Service Consumer
4086 Service Level Real World Effect
4087 Service Mediator
4088 Service Provider
4089 Shared State
4090 Skill
4091 Social Structure
4092 Stakeholder
4093 State
4094 System
4095 System Stakeholder

- 4096 Trust
- 4097 View
- 4098 Viewpoint

4099

C. The Unified Modeling Language, UML

4100 Figure 44 illustrates an annotated example of a UML class diagram that is used to represent a visual
4101 model depiction of the Resources Model in the *Participation in a SOA Ecosystem* view.



4102

4103 *Figure 44 - Example UML class diagram—Resources*

4104 Lines connecting boxes (classifiers) represent associations between things. An association has two roles
4105 (one in each direction). A role can have cardinality, for example, one or more (“1..*”) **stakeholders** own
4106 zero or more (“0..*”) **resources**. The role from classifier A to B is labeled closest to B, and vice versa, for
4107 example, the role between **resource** to Identity can be read as **resource** embodies Identity, and Identity
4108 denotes a **resource**.

4109 Mostly, we use named associations, which are denoted with a verb or verb phrase associated with an
4110 arrowhead. A named association reads from classifier A to B, for example, one or more **stakeholders**
4111 owns zero or more **resources**. Named associations are a very effective way to model relationships
4112 between concepts.

4113 An open diamond (at the end of an association line) denotes an aggregation, which is a part-of
4114 relationship, for example, Identifiers are part of Identity (or conversely, Identity is made up of Identifiers).

4115 A stronger form of aggregation is known as composition, which involves using a filled-in diamond at the
4116 end of an association line (not shown in above diagram). For example, if the association between Identity
4117 and Identifier were a composition rather than an aggregation as shown, deleting Identity would also
4118 delete any owned Identifiers. There is also an element of exclusive ownership in a composition
4119 relationship between classifiers, but this usually refers to specific instances of the owned classes
4120 (objects).

4121 This is by no means a complete description of the semantics of all diagram elements that comprise a
4122 UML class diagram, but rather is intended to serve as an illustrative example for the reader. It should be
4123 noted that the SOA-RAF utilizes additional class diagram elements as well as other UML diagram types
4124 such as sequence diagrams and component diagrams. The reader who is unfamiliar with the UML is
4125 encouraged to review one or more of the many useful online resources and book publications available
4126 describing UML (see, for example, www.uml.org).

4127

D. Critical Factors Analysis

4128
4129
4130
4131
4132
4133

A critical factors analysis (CFA) is an analysis of the key properties of a project. A CFA is analyzed in terms of the goals of the project, the critical factors that will lead to its success and the measurable requirements of the project implementation that support the goals of the project. CFA is particularly suitable for capturing quality attributes of a project, often referred to as “non-functional” or “other-than-functional” requirements: for example, security, scalability, wide-spread adoption, and so on. As such, CFA complements rather than attempts to replace other requirements capture techniques.

4134

D.1 Goals

4135
4136
4137

A goal is an overall target that you are trying to reach with the project. Typically, goals are hard to measure by themselves. Goals are often directed at the potential consumer of the product rather than the technology developer.

4138

D.2 Critical Success Factors

4139
4140
4141

A critical success factor (CSF) is a property, sub-goal that directly supports a goal and there is strong belief that without it the goal is unattainable. CSFs themselves are not necessarily measurable in themselves.

4142

D.3 Requirements

4143
4144
4145
4146

A requirement is a specific measurable property that directly supports a CSF. The key here is measurability: it should be possible to unambiguously determine if a requirement has been met. While goals are typically directed at consumers of the specification, requirements are focused on technical aspects of the specification.

4147

D.4 CFA Diagrams

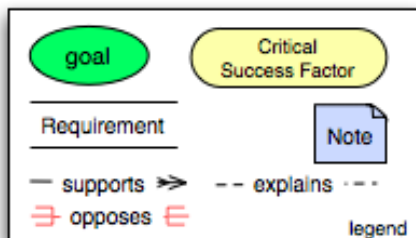
4148
4149
4150

It can often be helpful to illustrate graphically the key concepts and relationships between them. Such diagrams can act as effective indices into the written descriptions of goals etc., but is not intended to replace the text.

4151

The legend:

4152



4153

4154

4155

4156

4157

4158

illustrates the key elements of the graphical notation. Goals are written in round ovals, critical success factors are written in round-ended rectangles and requirements are written using open-ended rectangles. The arrows show whether a CSF/goal/requirement is supported by another element or opposed by it. This highlights the potential for conflict in requirements.

4159

E. Relationship to other SOA Open Standards

4160 The white paper “Navigating the SOA Open Standards Landscape Around Architecture” issued jointly by
4161 OASIS, OMG, and The Open Group **[SOA-NAV]** was written to help the SOA community at large
4162 navigate the myriad of overlapping technical products produced by these organizations with specific
4163 emphasis on the “A” in SOA, i.e., Architecture.

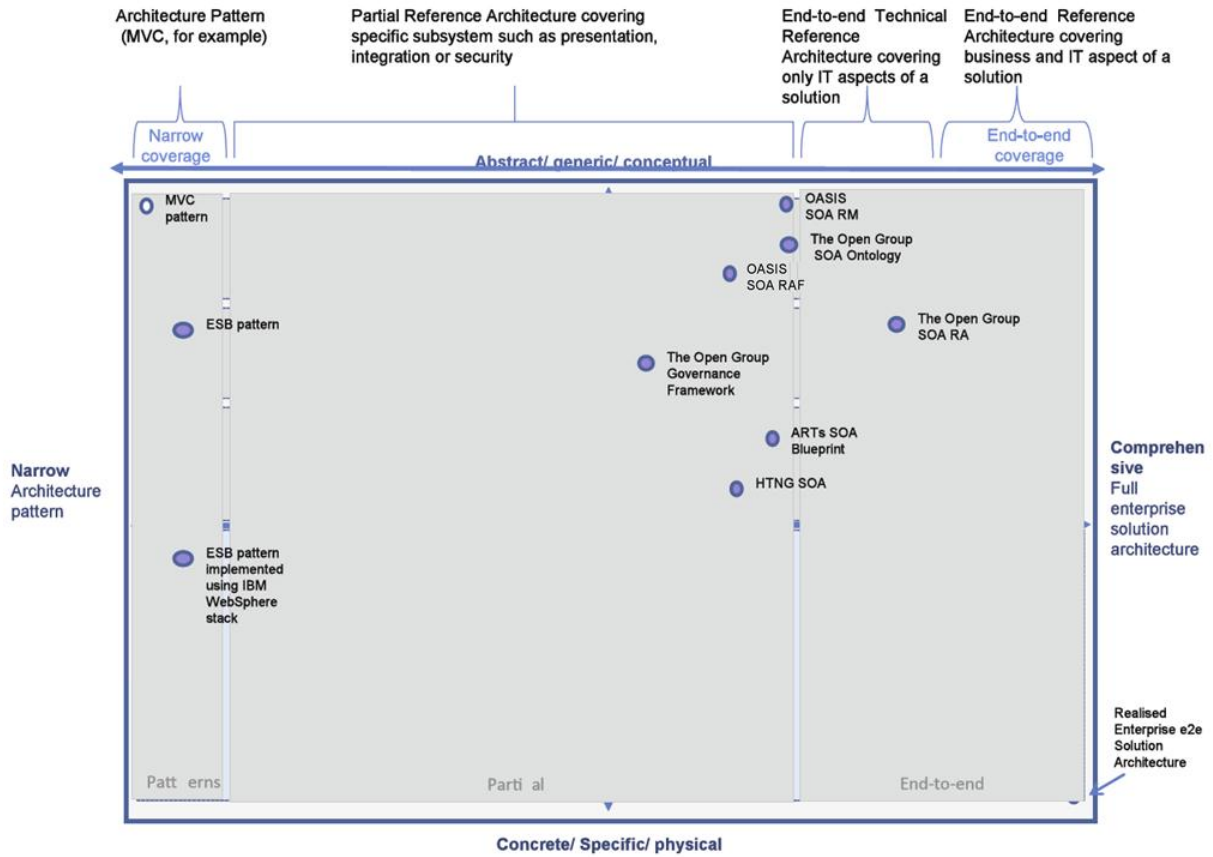
4164 The white paper explains and positions standards for SOA reference models, ontologies, reference
4165 architectures, maturity models, modeling languages, and standards work on SOA governance. It outlines
4166 where the works are similar, highlights the strengths of each body of work, and touches on how the work
4167 can be used together in complementary ways. It is also meant as a guide to users for selecting those
4168 specifications most appropriate for their needs.

4169 While the understanding of SOA and SOA Governance concepts provided by these works is similar, the
4170 evolving standards are written from different perspectives. Each specification supports a similar range of
4171 opportunity, but has provided different depths of detail for the perspectives on which they focus. Although
4172 the definitions and expressions may differ, there is agreement on the fundamental concepts of SOA and
4173 SOA Governance.

4174 The following is a summary taken from **[SOA-NAV]** of the positioning and guidance on the specifications:

- 4175 • The OASIS Reference Model for SOA (SOA RM) is the most abstract of the specifications
4176 positioned. It is used for understanding core SOA concepts
- 4177 • The Open Group SOA Ontology extends, refines, and formalizes some of the core concepts of
4178 the SOA RM. It is used for understanding core SOA concepts and facilitates a model-driven
4179 approach to SOA development.
- 4180 • The OASIS Reference Architecture Foundation for SOA (this document) is an abstract,
4181 foundational reference architecture addressing a broader ecosystem viewpoint for building and
4182 interacting within the SOA paradigm. It is used for understanding different elements of SOA, the
4183 completeness of SOA architectures and implementations, and considerations for reaching across
4184 ownership boundaries where there is no single authoritative entity for SOA and SOA governance.
- 4185 • The Open Group SOA Reference Architecture is a layered architecture from consumer and
4186 provider perspective with cross cutting concerns describing these architectural building blocks
4187 and principles that support the realizations of SOA. It is used for understanding the different
4188 elements of SOA, deployment of SOA in enterprise, basis for an industry or organizational
4189 reference architecture, implication of architectural decisions, and positioning of vendor products in
4190 a SOA context.
- 4191 • The Open Group SOA Governance Framework is a governance domain reference model and
4192 method. It is for understanding SOA governance in organizations. The OASIS Reference
4193 Architecture for SOA Foundation contains an abstract discussion of governance principles as
4194 applied to SOA across boundaries
- 4195 • The Open Group SOA Integration Maturity Model (OSIMM) is a means to assess an
4196 organization’s maturity within a broad SOA spectrum and define a roadmap for incremental
4197 adoption. It is used for understanding the level of SOA maturity in an organization
- 4198 • The Object Management Group SoaML Specification supports services modeling UML
4199 extensions. It can be seen as an instantiation of a subset of the Open Group RA used for
4200 representing SOA artifacts in UML.

4201 Fortunately, there is a great deal of agreement on the foundational core concepts across the many
4202 independent specifications and standards for SOA. This could be best explained by broad and common
4203 experience of users of SOA and its maturity in the marketplace. It also provides assurance that investing
4204 in SOA-based business and IT transformation initiatives that incorporate and use these specifications and
4205 standards helps to mitigate risks that might compromise a successful SOA solution.



4206

4207 *Figure 45 - SOA Reference Architecture Positioning (from "Navigating the SOA Open Standards Landscape Around*
 4208 *Architecture, © OASIS, OMG, The Open Group)*