



Reference Architecture Foundation for Service Oriented Architecture Version 1.0

Committee Draft 02

14 October 2009

Specification URIs:

This Version:

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

Previous Version:

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

Latest Version:

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

Technical Committee:

OASIS Service Oriented Architecture Reference Model TC

Chair(s):

Ken Laskey, MITRE Corporation

Editor(s):

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

Related work:

This specification is related to:

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

Abstract:

This document specifies the OASIS Reference Architecture for Service Oriented Architecture. It follows from the concepts and relationships defined in the OASIS Reference Model for Service Oriented Architecture. While it remains abstract in nature, the current document describes one possible template upon which a SOA concrete architecture can be built.

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

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

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

Status:

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

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

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

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

¹ American National Standards Institute

² Institute of Electrical and Electronics Engineers

³ International Organization for Standardization

⁴ International Electrotechnical Commission

Notices

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

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

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

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

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

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

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

OASIS takes no position regarding the validity or scope of any intellectual property or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; neither does it represent that it has made any effort to identify any such rights. Information on OASIS' procedures with respect to rights in any document or deliverable produced by an OASIS Technical Committee can be found on the OASIS website. Copies of claims of rights made available for publication and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this OASIS Committee Specification or OASIS Standard, can be obtained from the OASIS TC Administrator. OASIS makes no representation that any information or list of intellectual property rights will at any time be complete, or that any claims in such list are, in fact, Essential Claims.

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

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

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 Foundation?.....	9
1.1.3	Relationship to the OASIS Reference Model for SOA	10
1.1.4	Relationship to other Reference Architectures.....	10
1.1.5	Relationship to other SOA Open Standards	10
1.1.6	Expectations set by this Reference Architecture.....	12
1.2	Service Oriented Architecture – An Ecosystems Perspective	12
1.3	Viewpoints, Views and Models	13
1.3.1	ANSI/IEEE 1471-2000::ISO/IEC 42010-2007	13
1.3.2	UML Modeling Notation.....	14
1.4	Viewpoints of this Reference Architecture.....	15
1.4.1	Service Ecosystem Viewpoint	16
1.4.2	Realizing Service Oriented Architectures Viewpoint.....	16
1.4.3	Owning Service Oriented Architectures Viewpoint.....	17
1.5	Terminology.....	17
1.5.1	Usage of Terms.....	17
1.6	References	17
1.6.1	Normative References	17
1.6.2	Non-Normative References.....	18
2	Architectural Goals and Principles.....	19
2.1	Goals and Critical Success Factors of this Reference Architecture	19
2.1.1	Goals.....	20
2.1.2	Critical Success Factors.....	20
2.2	Principles of this Reference Architecture.....	21
3	Service Ecosystem View	24
3.1	Acting in a SOA Ecosystem Model.....	25
3.1.1	Actors, Delegates and Participants.....	26
3.1.2	Action and Joint Action.....	27
3.1.3	Communication as Joint Action.....	30
3.1.4	Using Communication for Service Action.....	32
3.2	Social Structure Model	33
3.2.1	Roles in Social Structures	35
3.2.2	Shared State and Social Facts.....	36
3.3	Acting in a Social Context Model.....	39
3.3.1	Service Providers and Consumers.....	40
3.3.2	Needs and Capabilities	41
3.3.3	Resources	42
3.3.4	Ownership	43
3.3.5	Trust, Risk and Willingness.....	44
3.3.6	Transactions and Exchanges.....	46
4	Realizing Service Oriented Architectures View	48

4.1	Service Description Model	48
4.1.1	The Model for Service Description	49
4.1.2	Use Of Service Description	57
4.1.3	Relationship to Other Description Models	62
4.1.4	Architectural Implications	63
4.2	Service Visibility Model	64
4.2.1	Visibility to Business	64
4.2.2	Visibility	65
4.2.3	Architectural Implications	70
4.3	Interacting with Services Model	71
4.3.1	Interaction Dependencies	71
4.3.2	Actions and Events	71
4.3.3	Message Exchange	72
4.3.4	Composition of Services	74
4.3.5	Architectural Implications of Interacting with Services	79
4.4	Policies and Contracts Model	80
4.4.1	Policy and Contract Principles	80
4.4.2	Policy Metrics	83
4.4.3	Automating Support for Policies and Contracts	83
4.4.4	Architectural Implications	84
5	Owning Service Oriented Architectures View	85
5.1	Governance Model	85
5.1.1	Understanding Governance	85
5.1.2	A Generic Model for Governance	87
5.1.3	Governance Applied to SOA	91
5.1.4	Architectural Implications of SOA Governance	95
5.2	Security Model	96
5.2.1	Secure Interaction Concepts	96
5.2.2	Where SOA Security is Different	99
5.2.3	Security Threats	99
5.2.4	Security Responses	100
5.2.5	Architectural Implications of SOA Security	101
5.3	Management Model	102
5.3.1	Management and Governance	105
5.3.2	Management Contracts and Policies	105
5.3.3	Management Infrastructure	105
5.3.4	Service Life-cycle	106
5.4	SOA Testing Model	106
5.4.1	Traditional Software Testing as Basis for SOA Testing	106
5.4.2	Testing and the SOA Ecosystem	108
5.4.3	Elements of SOA Testing	109
5.4.4	Testing SOA Services	113
5.4.5	Architectural Implications for SOA Testing	116
6	Conformance	117
A.	Acknowledgements	118

B. Critical Factors Analysis.....	119
B.1 Goals.....	119
Critical Success Factors.....	119
Requirements.....	119
CFA Diagrams.....	119

Table of Figures

Figure 1 SOA Reference Architecture Positioning [ERA].....	12
Figure 2 Example UML class diagram—Resources.....	15
Figure 3 Critical Factors Analysis of the Reference Architecture	19
Figure 4 Model elements described in the Service Ecosystem view	25
Figure 5 Actors, Participants and Delegates	26
Figure 6 Actions, Real World Effect and Events	27
Figure 7 Joint Action	28
Figure 8 Communication as Joint Action	30
Figure 9 Communicative actions as Service Actions	32
Figure 10 Social Structure	33
Figure 11 Enterprise as a Social Structure	34
Figure 12 Roles, Rights and Responsibilities	35
Figure 13 Shared State and Social Facts	37
Figure 14 Propositions	38
Figure 15 Assertions and Promises	39
Figure 16 Acting within Social Structures	40
Figure 17 Service Participants	40
Figure 18 Needs and Capabilities.....	41
Figure 19 Resources	42
Figure 20 Resource Ownership	43
Figure 21 Trusting Actor and Willingness	44
Figure 22 Assessing Trust and Risk	45
Figure 23 Business Transaction	46
Figure 24 Model Elements Described in the Realizing a Service Oriented Architecture View	48
Figure 25 General Description	50
Figure 26 Representation of a Description	51
Figure 27 Service Description.....	53
Figure 28 Service Interface.....	54
Figure 29 Service Functionality	55
Figure 30 Model for Policies and Contracts as related to Service Participants	56
Figure 31 Action-Level and Service-Level Policies.....	56
Figure 32 Policies and Contracts, Metrics, and Compliance Records.....	57
Figure 33 Relationship Between Action and Service Description Components	58
Figure 34 Execution Context	61
Figure 35 Interaction Description	62
Figure 36 Visibility to Business	65
Figure 37 Mediated Service Awareness	67
Figure 38 Awareness In a SOA Ecosystem.....	68
Figure 39 Business, Description and Willingness.....	69
Figure 40 Service Reachability	69

Figure 41 Interaction dependencies.	71
Figure 42 A "message" conveys either an action or an event.	72
Figure 43 Fundamental SOA message exchange patterns (MEPs).....	73
Figure 44 Simple model of service composition ("public" composition).....	75
Figure 45 Abstract example of orchestration of service-oriented business process.....	77
Figure 46 Abstract example of choreography of service-oriented business collaboration.....	78
Figure 47 Policies and Contracts.....	81
Figure 48 Model Elements Described in the Owing Service Oriented Architectures View	85
Figure 49 Motivating governance model.....	87
Figure 50 Setting up governance model.....	88
Figure 51 Carrying out governance model	89
Figure 52 Ensuring governance compliance model.....	90
Figure 53 Relationship among types of governance	92
Figure 54 Confidentiality and Integrity	97
Figure 55 Authentication.....	97
Figure 56 Authorization.....	98
Figure 57 Managing resources in a SOA.....	103

1 Introduction

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

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

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

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

1.1 Context for Reference Architecture for SOA

1.1.1 What is a Reference Architecture?

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

It is possible to define reference architectures at many levels of detail or abstraction, and for many different purposes. A reference architecture need not be a concrete architecture; i.e., depending on the requirements being addressed by the reference architecture, it may not be necessary to completely specify all the technologies, components and their relationships in sufficient detail to enable direct implementation.

1.1.2 What is this Reference Architecture Foundation?

This Reference Architecture Foundation is an abstract realization of SOA, focusing on the elements and their relationships needed to enable SOA-based systems to be used, realized and owned; while avoiding reliance on specific concrete technologies.

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

- resources that are distributed across ownership boundaries;
- people and systems interacting with each other, also across ownership boundaries;

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

- 41 • security, management and governance that are similarly distributed across ownership
42 boundaries; and
- 43 • interaction between people and systems that is primarily through the exchange of messages with
44 reliability that is appropriate for the intended uses and purposes.

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

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

55 **1.1.3 Relationship to the OASIS Reference Model for SOA**

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

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

65 In terms of approach, the primary difference between the Reference Model and this Reference
66 Architecture Foundation is that the former focuses entirely on a common language of the distinguishing
67 features of SOA; whereas this document introduces concepts and architectural elements as needed in
68 order to fulfill the core requirement of using, realizing and owning SOA-based systems.

69 **1.1.4 Relationship to other Reference Architectures**

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

75 As with the Reference Model for SOA, this Reference Architecture Foundation for SOA is primarily
76 focused on large-scale distributed IT systems where the participants may be legally separate entities. It is
77 quite possible for many aspects of 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 concepts on which to build other reference architectures and eventual concrete
80 architectures. The relationship to other industry reference architectures for SOA and related SOA open
81 standards is described below in Section 1.1.5

82 **1.1.5 Relationship to other SOA Open Standards**

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

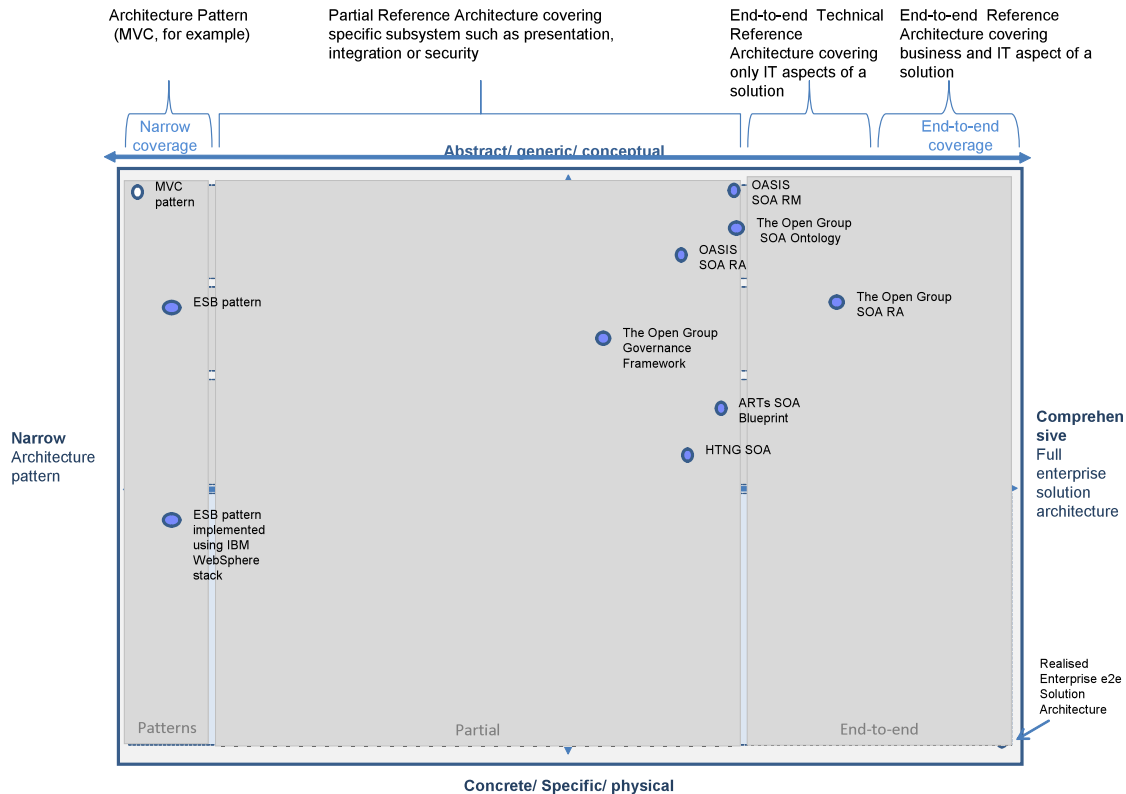
87 This joint white paper explains and positions standards for SOA reference models, ontologies, reference
88 architectures, maturity models, modeling languages, and standards work on SOA governance. It outlines
89 where the works are similar, highlights the strengths of each body of work, and touches on how the work

90 can be used together in complementary ways.. It is also meant as a guide to users of these specifications
91 for selecting the technical products most appropriate for their needs, consistent with where they are today
92 and where they plan to head on their SOA journeys.

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

98 The following is a summary of the positioning and guidance on the specifications:

- 99 • The OASIS Reference Model for SOA (SOA RM) is the most abstract of the specifications
100 positioned. It is used for understanding of core SOA concepts
- 101 • The Open Group SOA Ontology extends, refines, and formalizes some of the core concepts of
102 the SOA RM. It is used for understanding of core SOA concepts and facilitate a model-driven
103 approach to SOA development.
- 104 • The OASIS Reference Architecture Foundation for SOA (this document) is an abstract,
105 foundation reference architecture addressing the ecosystem viewpoint for building and interacting
106 within the SOA paradigm. It is used for understanding different elements of SOA, the
107 completeness of SOA architectures and implementations, and considerations for cross ownership
108 boundaries where there is no single authoritative entity for SOA and SOA governance.
- 109 • The Open Group SOA Reference Architecture is a layered architecture from consumer and
110 provider perspective with cross cutting concerns describing these architectural building blocks
111 and principles that support the realizations of SOA. It is used for understanding the different
112 elements of SOA, deployment of SOA in enterprise, basis for an industry or organizational
113 reference architecture, implication of architectural decisions, and positioning of vendor products in
114 a SOA context
- 115 • The Open Group SOA Governance Framework is a governance domain reference model and
116 method. It is for understanding SOA governance in organizations. The OASIS Reference
117 Architecture for SOA Foundation contains an abstract discussion of governance principles as
118 applied to SOA across boundaries
- 119 • The Open Group SOA Integration Maturity Model (OSIMM) is a means to assess an
120 organization's maturity within a broad SOA spectrum and define a roadmap for incremental
121 adoption. It is used for understanding the level of SOA maturity in an organization
- 122 • The Object Management Group SoaML Specification supports services modeling UML
123 extensions. It can be seen as an instantiation of a subset of the Open Group RA used for
124 representing SOA artifacts in UML.



September 2007, v0.1

SOA-enabled Business Transformation Framework (SBTF)

33

125

126 *Figure 1 SOA Reference Architecture Positioning [ERA].*

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

132 It is anticipated that future work on SOA standards may consider the positioning in this paper to reduce
 133 inconsistencies, overlaps, and gaps between related standards and to ensure that they continue to evolve
 134 in as consistent and complete a manner as possible.

135 **1.1.6 Expectations set by this Reference Architecture**

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

141 **1.2 Service Oriented Architecture – An Ecosystems Perspective**

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

147 From a holistic perspective, a SOA-based system is a network of independent services, machines, the
 148 people who operate, affect, use, and govern those services as well as the suppliers of equipment and
 149 personnel to these people and services. This includes any entity, animate or inanimate, that may affect or

150 be affected by the system. With a system that large, it is clear that nobody is really "in control" or "in
151 charge" of the whole ecosystem; although there are definite stakeholders involved, each of whom has
152 some control and influence over the community.

153 Instead of visualizing a SOA as a single complex machine, this Reference Architecture Foundation views
154 it as an ecosystem: a space people, machines and services inhabit in order to further both their own
155 objectives and the objectives of the larger community.

156 This view of SOA as ecosystem has been a consistent guide to the development of this architecture.

157 Taking an ecosystems perspective often means taking a step back: for example, instead of specifying an
158 application hierarchy, we model the system as a network of peer-like entities; instead of specifying a
159 hierarchy of control, we specify rules for the interactions between participants.

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

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

166 1.3 Viewpoints, Views and Models

167 1.3.1 ANSI/IEEE 1471-2000::ISO/IEC 42010-2007

168 This Reference Architecture Foundation follows the ANSI⁶/IEEE⁷ 1471-2000 and ISO⁸/IEC⁹ 42010-2007
169 standard. Recommended Practice for Architectural Description of Software-Intensive Systems
170 [ANSI/IEEE 1471, ISO/IEC 42010]. An architectural description conforming to the ANSI/IEEE 1471-
171 2000::ISO/IEC 42010-2007 recommended practice is described by a clause that includes the following six
172 (6) elements:

- 173 1. Architectural description identification, version, and overview information
- 174 2. Identification of the system stakeholders and their concerns judged to be relevant to the
175 architecture
- 176 3. Specifications of each viewpoint that has been selected to organize the representation of the
177 architecture and the rationale for those selections
- 178 4. One or more architectural views
- 179 5. A record of all known inconsistencies among the architectural description's required constituents
- 180 6. A rationale for selection of the architecture (in particular, showing how the architecture supports
181 the identified stakeholders' concerns).

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

183 **Architecture**

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

186 **Architectural Description**

187 A collection of products that document the architecture.

⁶ American National Standards Institute

⁷ Institute of Electrical and Electronics Engineers

⁸ International Organization for Standardization

⁹ International Electrotechnical Commission

188 **System**

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

190 **System Stakeholder**

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

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

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

201 **View**

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

203 **Viewpoint**

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

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

209 It is important to note that viewpoints are independent of a particular system. In this way, the architect can
210 select a set of candidate viewpoints first, or create a set of candidate viewpoints, and then use those
211 viewpoints to construct specific views that will be used to organize the architectural description. A view,
212 on the other hand, is specific to a particular system. Therefore, the practice of creating an architectural
213 description involves first selecting the viewpoints and then using those viewpoints to construct specific
214 views for a particular system or subsystem. Note that ANSI/IEEE 1471-2000::ISO/IEC 42010-2007
215 requires that each view corresponds to exactly one viewpoint. This helps maintain consistency among
216 architectural views; a normative requirement of the standard.

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

218 **Model**

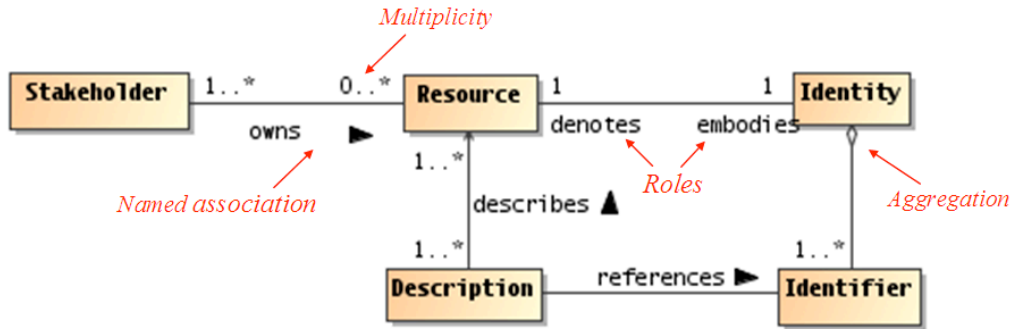
219 An abstraction or representation of some aspect of a thing (in this case, a system)

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

222 **1.3.2 UML Modeling Notation**

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

229 Figure 1 illustrates an annotated example of a UML class diagram that is used to represent a visual
230 model depiction of the Resources Model in the Service Ecosystem View (Section 3). The figure caption
231 describes the UML semantics of this diagram.



232

233 *Figure 2 Example UML class diagram—Resources.*

234 Lines connecting boxes (classifiers) represent associations between things. An association has two roles
 235 (one in each direction). A role can have multiplicity, for example, one or more (“1..*”) **Stakeholders** own
 236 zero or more (“0..*”) **Resources**. The role from classifier A to B is labeled closest to B, and vice versa, for
 237 example, the role between **Resource** to **Identity** can be read a **Resource** embodies **Identity**, and
 238 **Identity** denotes a **Resource**.

239 Mostly, we use named associations, which are denoted with a verb or verb phrase associated with an
 240 arrowhead. A named association reads from classifier A to B, for example, one or more **Stakeholders**
 241 owns zero or more **Resources**. Named associations are a very effective way to model relationships
 242 between concepts.

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

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

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

258

259 1.4 Viewpoints of this Reference Architecture

260 This Reference Architecture Foundation is partitioned into three views that conform to three primary
 261 viewpoints, reflecting the main division of concerns noted above: the Service Oriented Architecture – An
 262 Ecosystems Perspective viewpoint focuses on how people conduct their business using SOA-based
 263 systems; the Realizing Service Oriented Architecture viewpoint focuses on the salient aspects of building
 264 a SOA, and the Owning Service Oriented Architectures viewpoint focuses on those aspects that relate to
 265 owning, managing and controlling a SOA.

266 The viewpoint specifications for each of the primary viewpoints of this Reference Architecture are
 267 summarized in Table 1. Additional detail on each of the three viewpoints is further elaborated in the
 268 following subsections. For this Reference Architecture, there is a one-to-one correspondence between
 269 viewpoints and views.

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

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

271 **1.4.1 Service Ecosystem Viewpoint**

272 The Service Ecosystem viewpoint is intended to capture what using a SOA-based system means for
 273 people using it to conduct their business. We do not limit the applicability of SOA-based systems to
 274 commercial and enterprise systems. We use the term **business** to include any activity of interest to a
 275 user; especially activities shared by multiple users.

276 From this viewpoint, we are concerned with how SOA integrates with and supports the service model
 277 from the perspective of the people who perform their tasks and achieve their goals as mediated by
 278 Service Oriented Architectures. The Service Ecosystem viewpoint also sets the context and background
 279 for the other viewpoints in the Reference Architecture.

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

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

288 **1.4.2 Realizing Service Oriented Architectures Viewpoint**

289 The Realizing Service Oriented Architectures Viewpoint focuses on the infrastructure elements that are
 290 needed to support the construction of SOA-based systems. From this viewpoint, we are concerned with
 291 the application of well-understood technologies available to system architects to realize the vision of a
 292 SOA that may cross ownership boundaries. In particular, we are aware of the importance and relevance
 293 of other standard specifications that may be used to facilitate the building of a SOA.

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

296 1.4.3 Owing Service Oriented Architectures Viewpoint

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

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

306 1.5 Terminology

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

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

311 Terms such as this “Reference Architecture” refer to this document, and “the Reference Model” refer to
312 the OASIS Reference Model for Service Oriented Architecture”. **[SOA-RM]**.

313 1.5.1 Usage of Terms

314 Certain terms are used throughout this document, such as model, action, and rule, which are regular
315 concepts and which have formal definitions within the Reference Architecture. Where a reference to the
316 formally defined concept is intended, we use a capitalized form such as Model, Action and Rule. Where
317 the more colloquial and informal meaning is intended, words are used without special capitalization.

318 1.6 References

319 1.6.1 Normative References

- 320 **[ANSI/IEEE 1471]** *IEEE Recommended Practice for Architectural Description of Software-Intensive*
321 *Systems*, American National Standards Institute/Institute for Electrical and
322 Electronics Engineers, September 21, 2000.
- 323 **[ISO/IEC 42010]** International Organization for Standardization and International Electrotechnical
324 Commission, *System and software engineering — Recommended practice for*
325 *architectural description of software-intensive systems*, July 15, 2007.
- 326 **[RFC2119]** S. Bradner, *Key words for use in RFCs to Indicate Requirement Levels*,
327 <http://www.ietf.org/rfc/rfc2119.txt>, IETF RFC 2119, March 1997.
- 328 **[SOA-RM]** OASIS Standard, "Reference Model for Service Oriented Architecture 1.0, 12
329 October 2006. <http://docs.oasis-open.org/soa-rm/v1.0/soa-rm.pdf>
- 330 **[UML 2]** *Unified Modeling Language: Superstructure*, Ver. 2.1.1, OMG Adopted
331 Specification, OMG document formal/2007-02-05, Object Management Group,
332 Needham, MA, February 5, 2007.
- 333 **[WA]** Architecture of the World Wide Web, W3C, 2004. <http://www.w3.org/TR/webarch>.
- 334 **[WSA]** David Booth, et al., "Web Services Architecture", W3C Working Group Note,
335 World Wide Web Consortium (W3C) (Massachusetts Institute of Technology,
336 European Research Consortium for Informatics and Mathematics, Keio
337 University), February, 2004. [http://www.w3.org/TR/2004/NOTE-ws-arch-
338 20040211/](http://www.w3.org/TR/2004/NOTE-ws-arch-20040211/)

339 **1.6.2 Non-Normative References**

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

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

344 **[ERA]** A. Fattah, "Enterprise Reference Architecture," paper presented at 22nd
345 Enterprise Architecture Practitioners Conference, London, UK, April 2009.

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

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

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

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

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

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

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

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

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

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

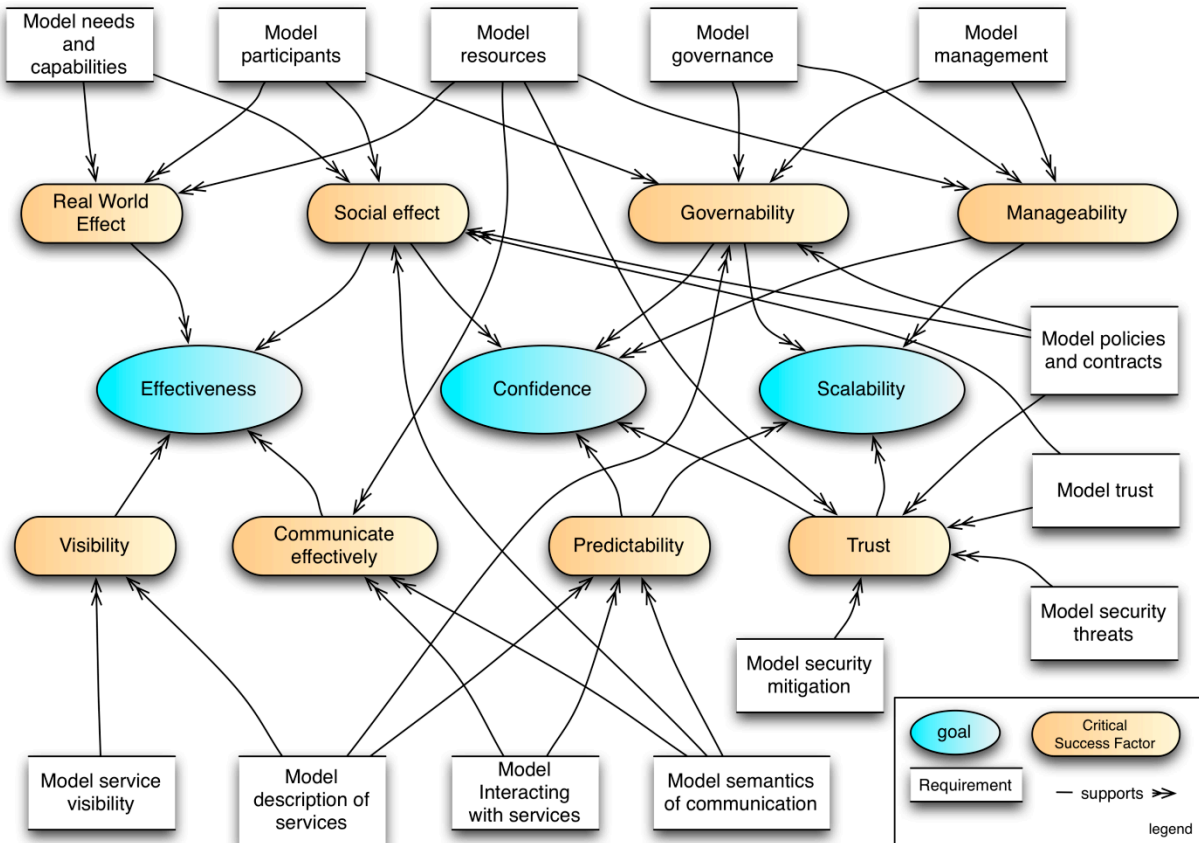
374 **2 Architectural Goals and Principles**

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

377 **2.1 Goals and Critical Success Factors of this Reference Architecture**

378 There are three principal goals of this Reference Architecture:

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



384
 385 *Figure 3 Critical Factors Analysis of the Reference Architecture*

386 Figure 3 represents a Critical Factors Analysis (CFA) diagram demonstrating the relationship between the
 387 primary goals of this reference architecture, critical factors that determine the success of the architecture
 388 and individual elements that need to be modeled.

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

393 **2.1.1 Goals**

394 **2.1.1.1 Effectiveness Goal**

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

399 The critical factors that determine effectiveness are visibility between the participants, that they can
400 communicate effectively, and that actual real world effects and social effects can be realized. In addition,
401 the overall system must be manageable and governable.

402 **2.1.1.2 Confidence Goal**

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

407 In addition to ensuring that social effects are properly captured, other critical factors that are important for
408 ensuring confidence are trust, predictability, manageability and proper governance.

409 **2.1.1.3 Scalability Goal**

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

414 The critical factors that determine scalability, particularly in the context of multiple domains of ownership
415 are predictability, trust, governability and manageability. This is in addition to more traditional measures of
416 scalability such as performance of message exchange.

417 **2.1.2 Critical Success Factors**

418 A critical success factor (CSF) is a property of the intended system, or a sub-goal that directly supports a
419 goal and there is strong belief that without it the goal is unattainable. CSFs are not necessarily
420 measurable in themselves. As illustrated in Figure 3, CSFs can be associated with more than one goal.

421 **2.1.2.1 Real World Effect**

422 It is of the essence that participants can use a SOA-based system to realize actual effects in the world.
423 This implies that the capabilities that are accessed as a result of service interaction are embedded in and
424 backed by the real world.

425 We identify three models that address how service interactions can result in real world effects: a needs
426 and capabilities model, a participants model and a resources model.

427 **2.1.2.2 Social Effect**

428 Effects that are desired in the use of SOA-based systems may be social effects as well as physical
429 effects. For example, opening a bank account is primarily about the relationship between a customer and
430 a bank – the effect of the opened account is a change in the relationship between the customer and the
431 bank.

432 The models that are important in addressing this critical factor are similar to the more general real world
433 effect: the participants model, the needs and capabilities model, the resources model. In addition, the
434 semantics of communication model and the policy and contracts model directly support the objective of
435 realizing the appropriate social effect.

436 **2.1.2.3 Visibility**

437 Ensuring that participants can see each other is clearly also a critical factor in ensuring effectiveness of
438 interaction. Enabling visibility requires addressing the visibility of services and the correct descriptions of
439 services and related artifacts.

440 **2.1.2.4 Communicate effectively**

441 In order for there to be effective uses of capabilities and meeting of needs, it is critical that participants
442 can see and interact with each other. The models that address this are the Interacting with Services
443 model, the Resources model and the Semantics of Communication model.

444 **2.1.2.5 Manageability and Governability**

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

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

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

457 **2.1.2.6 Trust**

458 Trust is a critical factor in ensuring confidence. Trust can be analyzed in terms of trust in infrastructure
459 facilities (otherwise known as reliability), trust in the relationships and effects that are realized by
460 interactions with services, and trust in the integrity and confidentiality of those interactions particularly with
461 respect to external factors (otherwise known as security).

462 The threat model in Section 5.2.3 captures what is meant by trust; the security models capture how
463 external entities might attempt to corrupt that trust and how SOA-based systems can mitigate against
464 those risks.

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

469 **2.1.2.7 Predictability**

470 A factor that engenders confidence in any system is predictability. By predictability, we principally mean
471 that the expectations of participants of SOA-based systems can be tied to the actual performance of
472 those systems (what you see is what you get).

473 The primary means of ensuring predictability is effective descriptions: service descriptions document
474 services, the interacting with services model addresses expectations relating to how services are used
475 and the semantics of communications model addresses how meaning and intent can be exchanged
476 between participants.

477 **2.2 Principles of this Reference Architecture**

478 The following principles serve as core tenets that guide the evolution of this Reference Architecture. The
479 ordered numbering of these principles does not imply priority order.

480 **Principle 1: Technology Neutrality**

481 Statement: Technology neutrality refers to independence from particular technologies.

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

489 Implications: This Reference Architecture must be technology neutral, meaning that we assume that
490 technology will continue to evolve, and that over the lifetime of this architecture that
491 multiple, potentially competing technologies will co-exist. Another immediate implication
492 of technology independence is that greater effort on the part of architects and other
493 decision makers to construct systems based on this architecture is needed.

494 **Principle 2: Parsimony**

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

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

500 Implications: Occam’s (or Ockham’s) Razor applies, which states that the explanation of any
501 phenomenon should make as few assumptions as possible, eliminating those that make
502 no difference in the observable predictions of the explanatory hypothesis or theory. With
503 respect to this Reference Architecture, this is made apparent by avoiding the elaboration
504 of certain details which though that may be required for any particular solution, are likely
505 to vary substantially from application to application. The complement of a parsimonious
506 design is a feature-rich design. Parsimoniously designed systems tend to have fewer
507 features. This, in turn, means that people attempting to use such a system may have to
508 work harder to ensure that their application requirements have been met.

509 **Principle 3: Separation of Concerns**

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

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

521 Implications: This is a core tenet that drives this Reference Architecture to adopt the notion of
522 architectural viewpoints and corresponding views. A *viewpoint* provides the formalization
523 of the groupings of models representing one set of concerns relative to an architecture,
524 while a *view* is the actual representation of a particular system. The ability to leverage an
525 industry standard that formalizes this notion of architectural viewpoints and views helps
526 us better ground these concepts for not only the developers of this Reference
527 Architecture but also for its readers. Fortunately, such a standard exists in the IEEE
528 Recommended Practice for Architectural Description of Software-Intensive Systems
529 **[ANSI/IEEE 1471-2000::ISO/IEC 42010-2007]**; and it is this standard that serves as the
530 basis for the structure and organization of this Reference Architecture.

531 **Principle 4: Applicability**
532 **Statement:** Applicability refers to that which is relevant. Here, an architecture is sought that is
533 relevant to as many facets and applications of SOA-based systems as possible; even
534 those yet unforeseen.

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

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

542
543
544
545
546
547
548
549
550
551
552
553

3 Service 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

554
555
556
557
558
559
560
561
562
563
564
565
566
567
568
569
570
571
572
573
574
575
576
577
578
579
580
581

The *Service Ecosystem View* focuses on what a SOA-based system means for people to participate in it to conduct their business.¹⁰ Business, in general, is characterized in terms of providing services and consuming business services to realize mutually desirable real world effects; in a SOA-based system, the conduct of business involves the effective connectivity of IT-accessible resources as an important element in how these real world effects are realized.

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

However, the primary motivation for participants to interact with each other is to achieve goals – to get things done. While SOA implies the use of IT resources and artifacts, these are merely tools to an end and are usually not the primary interest of the participants. Describing what it means to *act* in the SOA ecosystem when participants may be in different organizations, with different rules and expectations is one of the primary modeling objectives of this section.

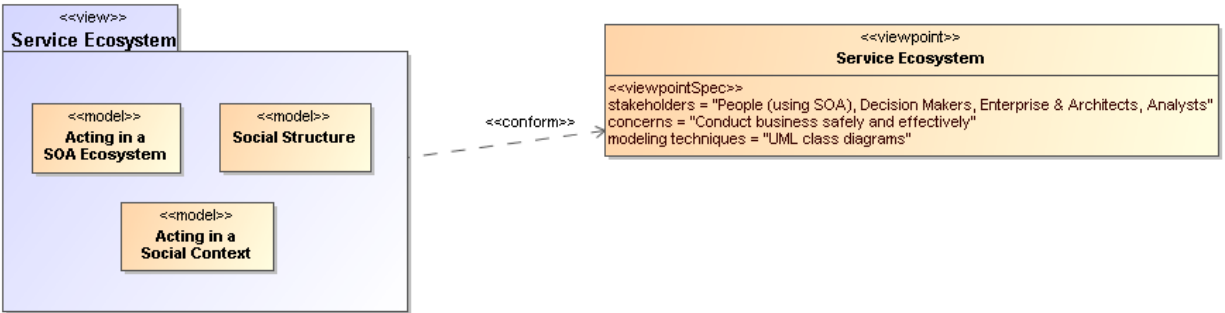
Since there is inherently some *mediation* involved when people interact using electronic means, we lay the foundations for how *communication* can be used to represent *action*. This foundation forms the backdrop for how services are realized – covered in Section 4 – as well as how SOA-based systems are managed as owned entities – covered in Section 5.

Thus, our tasks in this view are to model the people involved—the participants and other stakeholders—their goals and activities and the relevant relationships between people as they affect the utility and safety of actions that are performed.

The models in this view form the basis for many of the activities of SOA participants, especially in areas such as management and security. They lay a groundwork for those areas and will be referenced in the other views to provide a consistent discussion throughout this document.

In particular, the Acting in a SOA ecosystem model introduces the key concepts involved in actions, the Social Structure Model introduces the key elements that underlie the relationships between participants. The Acting in a Social Context model pulls the two together and shows how ownership, risk and transactions are key concepts in the SOA ecosystem.

¹⁰ By *business* we mean to include any activity entered into whose goal is to satisfy some need or desire of the participant.



582
583 *Figure 4 Model elements described in the Service Ecosystem view*

584 **3.1 Acting in a SOA Ecosystem Model**

585 At the core of participants interacting in a SOA ecosystem is the concept of action – participants are
586 acting against services in order to get their needs met. Service providers are acting in order to satisfy
587 those needs and governance parties act in order to ensure the smooth operation of the systems.

588 Actions may also be across one or more ownership boundaries; in particular, participants in a SOA
589 ecosystem may be performing actions involving systems that do *not belong to them*.


590 For example, if a consumer wishes to use a service to book an airline ticket for a journey, the consumer
591 must interact with the airline reservation system in order to achieve the goal of a reservation. The
592 consumer must be able to accomplish her goal using actions recognized by a reservation system that
593 was designed to satisfy the policies and other assumptions about context made by the airline and those
594 implementing the system.

595 When the consumer purchases a ticket, the action is to purchase the ticket but the means of doing so
596 involves an interaction with the airline. However, both the interaction itself and the purchase are actions
597 that must be understood at different levels – at the level of the IT systems through which messages are
598 communicated and at the level of the reservation service through which the effects of the purchase are
599 recorded.

600 There are many parallels between the way that human society is organized, and the way that humans
601 can act using the power of others. There are also parallels in satisfying business needs and satisfying the
602 mechanistic needs of the systems and processes that enable the bringing together of needs and
603 capabilities to satisfy our goals

604 In this section we establish the key principles of action as an abstract concept. We elaborate on action in
605 the context of acting in a social context as *joint action*. And we also establish the connections necessary
606 between the different levels of understanding of action that allow participants to interact as a means of
607 getting things done.

608 **3.1.1 Actors, Delegates and Participants**



609
610 *Figure 5 Actors, Participants and Delegates*

611 **Actor**

612 An **actor** is an entity, human, non-human or organization of entities, that is capable of **action**.¹¹

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

615 **Stakeholder**

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

618 **Participant**

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

620 A participant is a stakeholder whose interests lie in the successful use of and fulfillment of services.
621 However, human participants always require *representation* in an electronic system – they require
622 mechanisms to facilitate their interactions: they require delegates.

623 Note that we admit non-human agents that have no identifiable representative as an extreme case: the
624 normal situation is where participants are either human or organizations.

625 **Non-Participant Stakeholder**

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

627 Stakeholders do not necessarily participate in service interactions. For example, a government may have
628 an interest in the outcomes of commercial services deployed in a SOA ecosystem for the purposes of
629 collecting tax from one or more of the participants. A government may also be interested in regulatory
630 compliance as it affects service interactions.

631 There are two main classes of such non-participatory stakeholders: third parties who are affected by
632 someone's use or provisioning of a service, and regulatory agencies who wish to control the outcome of
633 service interactions in some way (such as by taxation). An example of an affected third party may be
634 someone using the service infrastructure whose activities are impeded because an errant participant is
635 consuming excessive bandwidth in another interaction.

636 **Delegate**

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

638 In order for people to be able to offer, consume and otherwise participate in SOA service interactions,
639 they require the use of an entity capable of directly interacting with electronic communications – we use
640 the term **delegate** to identify that entity. Common examples are software applications that make use of
641 services, hardware devices that embody a particular mission, and enterprise systems that offer services.

642 We do not attempt to characterize **delegates** in terms of their internal architecture, computational
643 requirements or platforms here.

¹¹ Note that there is potential confusion between the concept of Actor in UML2.0 and an actor in an ecosystem. Section 3.2.1 defines the concepts of role and an actor adopting a role.

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

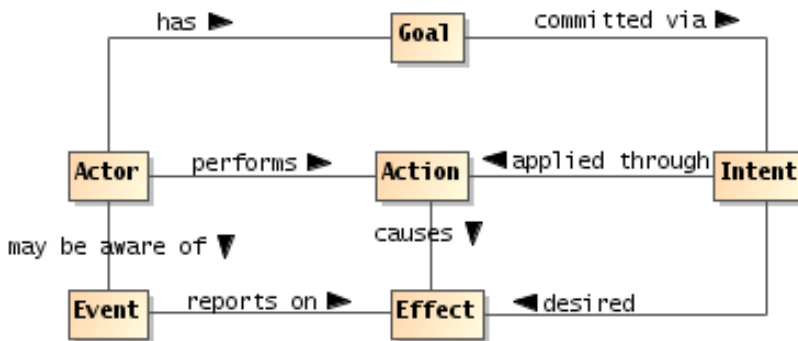
648 In the different models in this architecture we use the **actor** concept when it is not important whether the
649 entity involved is a delegate, participant or some other entity. If the entity is acting on behalf of another,
650 then we use the **delegate** concept. If the entity is a stakeholder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

652 3.1.2 Action and Joint Action

653 Entities act in order to achieve their **goals**. In this model, we look at the most basic form of action – an
654 action performed by a single actor. Figure 6 depicts a model of action showing the relationships between
655 action, goals and effects of action.

656 3.1.2.1 Action and Actors

657 Within this initial model of action, we focus on the actions of individual entities. However, we should
658 remark that for the most part within a SOA ecosystem, the actions we are most interested in are actions
659 involving multiple participants – we address this further in Section 3.1.2.2.



660
661 *Figure 6 Actions, Real World Effect and Events*

662 The most important concept in any model of actions and effects is that of **action** itself:

663 **Action**

664 An **action** is the application of **intent** to achieve an **effect** (within the SOA ecosystem).

665 This concept is simultaneously one of the fulcrums of the Service Oriented Architecture and a touch point
666 for many other aspects of the architecture: such as policies, service descriptions, management, security
667 and so on.

668 The aspect of **action** that distinguishes it from mere force or accident is that someone or something
669 intended the **action** to occur.

670 **Goal**

671 A **goal** is a measurable state of the ecosystem that an actor is seeking to establish.

672 Goals are conditions that people, and more generally actors, are seeking to satisfy. A key aspect of goals
673 is measurability: it should be possible to know if a goal has been satisfied.

674 **Intent**

675 **Intent** is the commitment of an **actor** to achieve a **goal**.

676 An actor's **intent** in performing an **action** is to further one or more of the actor's **goals**.

677 In some situations it may be difficult to determine an **actor's** actual **intent**. This is particularly true for
678 social actions such as those performed within a SOA-based system.

679 However, in most cases, entities in a SOA ecosystem make an assumption of *implied intent*. I.e., if an
680 **actor** performs an **action**, it is assumed that the **actor** also intended to perform the **action** – it was not an
681 accident, or the action of another actor.

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

684 **Effect**

685 An **effect** is a measurable change in the state of the ecosystem.

686 Note the normal **intent** of applying an **action** is to cause an **effect** that reflects the actor's goals.

687 However, there is often the possibility that the actual effects will include unintended consequences that
688 fall outside of, and may run counter to, the intent of the actor.

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

690 **Event**

691 An **event** is the report of an **effect** of which at least one participant has an interest in being
692 aware.

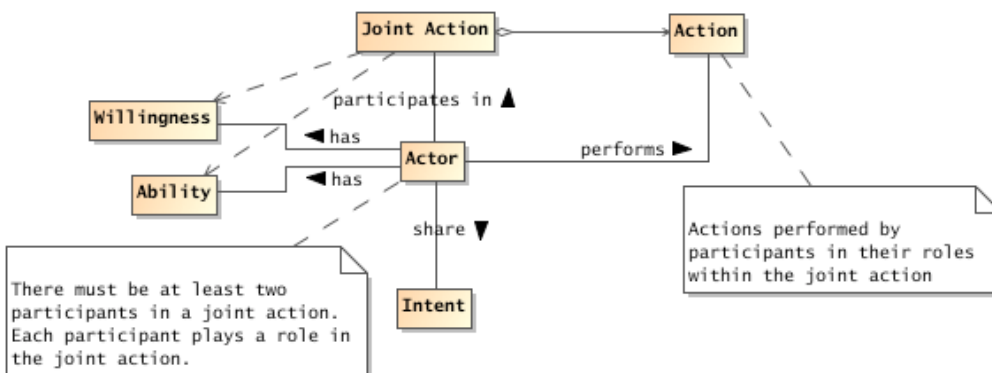
693 In effect (sic), an **event** is a corollary to **action**: in a public arena, actions result in changes to the state of
694 ecosystem (primarily changes to the states of individual **participants**); these changes may be manifested
695 as **events** of which participants in the arena have an awareness.

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

698 **3.1.2.2 Joint Actions**

699 Joint actions are the foundation for understanding interaction between participants in a SOA ecosystem.

700 In this Reference Architecture, we see joint actions at two levels: as communication and as participants
701 using and offering services.



702
703 *Figure 7 Joint Action*

704 **Joint Action**

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

707 In order for multiple actors to participate in a **joint action**, they must each act according to their role within
708 the **joint action**. For example, a common example of a **joint action** is for one **actor** to speak to
709 another.¹² A communication between **actors** cannot take place unless there is both a speaker and a
710 listener – although it is not necessarily required that they both be active simultaneously. The two **actors**
711 involved have different roles – one is a speaker and the other is a listener.

¹² Where speaking and listening includes electronic message sending and receiving.

712 By definition, **joint actions** are **actions** that cannot be performed by single **participants**. Sometimes this
713 is because no single participant has the ability to perform the action on his own; or, in the case of the
714 speaker and listener, the 'joint-ness' of joint actions is inherent.

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

718 There is another sense in which **joint actions** abound: even within a single incident of interaction there
719 are typically several overlapping **joint actions**.

720 For example, when one person says to another: "it is stuffy in here" there is an immediate sense in which
721 there is a **joint action** – a joint communicative action (see below). The intended effect being that the
722 listener believes that the speaker intends him to understand that the speaker believes that the
723 atmosphere is uncomfortable. (The listener may also believe that the atmosphere *is* uncomfortable as a
724 result of the communication.)

725 However, in the right context, there may be another joint action: the apparent declaration may in fact be a
726 command. The intent being that the speaker wishes the listener to understand that the door should be
727 opened, the effect being that of actually opening the door.

728 There may be a further layer to this scenario: the speaker might be aware that there is someone who is
729 waiting to be let in. The command to open the door is actually a command to admit the visitor to the room.

730 Fundamentally all three of these senses of joint action are superimposed on top of each other. However,
731 there is a strong sense in which the different joint actions may be quite interchangeable. For example,
732 instead of declaring that the "room is stuffy", the speaker might have simply said "open the door". Or the
733 speaker might have said "please let John in". In each case the effect would have been the same –
734 modulo the sensitivities of the speaker and listener – the door being open and the visitor admitted to the
735 room.

736 The relationship between the communicative joint action: the utterance of the declaration and the
737 command joint action is a 'uses' relationship. The speaking joint action is used to convey the command
738 joint action; which in turn is used to convey the visitor admittance action.

739 In many situations the best predicate that describes the relationship between these different joint actions
740 is the 'counts as' predicate. The utterance action counts as the command to open the door. The command
741 to open the door counts as the request to admit the visitor.

742 It can be extremely useful to identify and separate the different overlapping senses of joint action. It
743 allows us to separately describe and process the communicative actions from the command joint actions.
744 This, in turn, reflects the fact that each layer has its own logic and ontology.

745 For example, at the utterance level, the issues are to do with the successful understanding of the content
746 of the communication – did the listener hear and understand the words, did the speaker intend to say
747 them, and so on.

748 At the level of the command to open the door, the issues center on whether there is a predisposition on
749 the part of the listener to obey commands given to him by the speaker.

750 In the context of a SOA ecosystem we can separately capture the logic and mechanics of what is
751 involved in electronic communication – the sending of messages, the security of the communication and
752 so on; from the logic and mechanics of command -- does the listener believe that the speaker has the
753 appropriate authority to issue the command.

754 As with human communication, electronic interactions are similarly interchangeable: the commitment to
755 purchase a book requires some form of communication between buyer and seller; but the purchase
756 action itself is unchanged by the use of email or an HTTP post of an XML document.

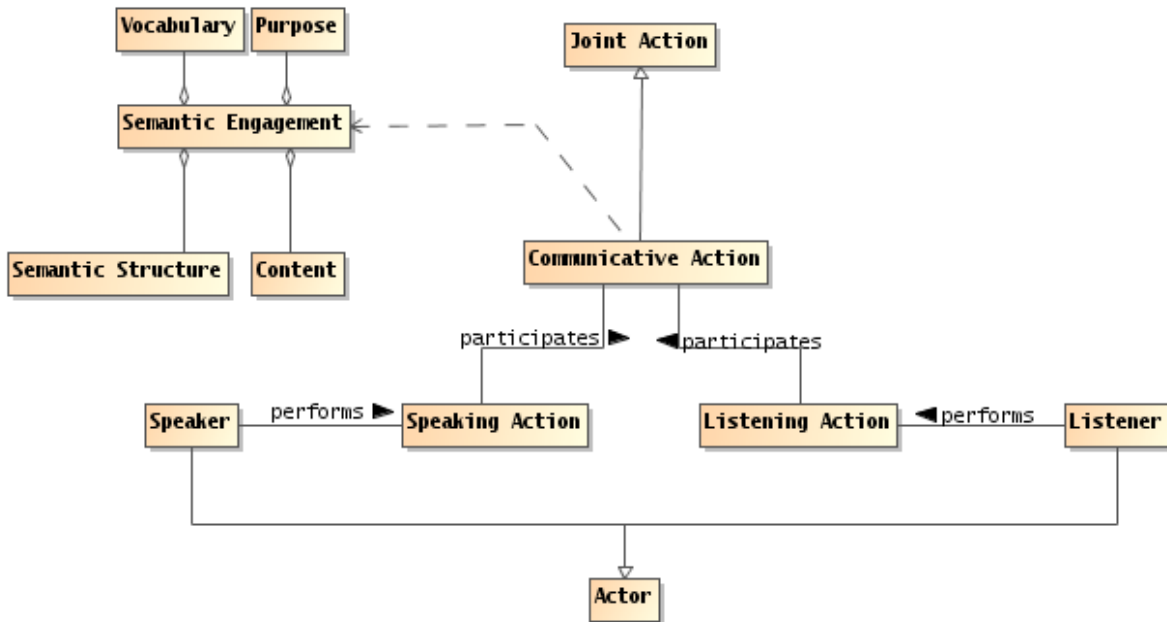
757 In summary, the concept of joint action allows us to honor the fact that both parties in an interaction are
758 required for there to be an actual effect; it allows us to separate out the different levels of the interaction
759 into appropriate semantic layers; and it allows us to recombine those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

761 **3.1.3 Communication as Joint Action**

762 Because there is inherently some separation between **actors** in a SOA ecosystem, they are effectively
 763 driven to use communication techniques to 'get their business done'. When an **actor** sends a message to
 764 another **actor** there are actually two (at least) senses in which the **actor** can be said to be acting: by
 765 communicating with other **actors**; and the purpose of the communication is also **action**.

766 The primary mechanism whereby **actors** interact with each other is through the *exchange of messages*,
 767 where the messages may cross ownership boundaries. Communication and the interpretation of
 768 communicated content is the foundation of all interaction within the SOA ecosystem.

769 However content is actually communicated, communicating is also a form of **action**. We define the
 770 **communicative action** as the **action** of message exchange:



771
 772 *Figure 8 Communication as Joint Action*

773 **Communicative Action**

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

776 A **communicative action** has a speaker and a listener; each of whom must perform their part for the
 777 communicative action to occur.

778 The concept of **communicative action** is important in the explanation of how we can use the exchange
 779 of messages to realize interaction between service participants. The Reference Model defines interaction
 780 as the activity that is involved in making use of a capability offered. A **communicative action** is the **joint**
 781 **action** involved in the exchange of messages.

782 **Speaking Action**

783 A **speaking action** is the **action** required of an **actor** in order to communicate a desired content.

784 **Listening Action**

785 A **listening action** is the **action** required of an **actor** in order to acquire and comprehend
 786 communicated content.

787 Notice that an **actor** listening to a message not only acquires the message but is also able to understand
 788 it. The implications of this are discussed further below.

789 **Speaker & Listener**

790 A **speaker** is an **actor** who performs the speaking action; A **listener** is an **actor** who performs
791 the **listening action**.

792 Speaking and listening are roles that (normally) different actors play in a given communicative action.

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

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

800 **Content**

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

802 Even though communication is effected through **action**, it is not actually effective if the **listener** cannot
803 understand the content of the communication. We can characterize the necessary modes of
804 understanding in terms of a shared *vocabulary* and a shared understanding of the communicated *intent*.

805 The meaning of a communication is typically conveyed as a combination of the syntax of the content, its
806 **semantics** and its **illocutionary force**.

807 Typically, the syntax takes the form of highly regular tree structure, with a well-defined method for
808 interpreting the structure. For example, an invoice will often follow pre-established standards for
809 communicating invoices.

810 **Semantics**

811 The **semantics** of a **communicative action** is the meaning of the content being communicated.

812 The semantics of a fragment of content can be characterized in terms of the **vocabulary** of terms
813 referenced in the content and the relationships between those terms that are represented by the syntactic
814 form of the content.

815 **Vocabulary**

816 A **vocabulary** is a set of terms together with an interpretation that is shared by **actors** involved in
817 a **communicative action**.

818 In order for there to be any communication, there must be sufficient shared understanding of the
819 elements of interaction and of terms used in communication. A shared vocabulary may range from a
820 simple understanding of particular strings as commands to a sophisticated collection of terms that are
821 formalized in shared ontologies.

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

826 **Illocutionary Force**

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


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

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

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

834 3.1.4 Using Communication for Service Action

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



839
840 Figure 9 Communicative actions as Service Actions

841 Service Action

842 A **service action** is an element of the action model of the service.

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

845 Counts as

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

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

850 When we state that a **communicative action counts as** a **service action**, we are relating a system of
851 communication to a system of action against services.¹³ Since a **participant** cannot (normally) act directly
852 on a **service** it must use some means of mediating the **action**. However, from the perspective of all the
853 participants involved, when a participant uses a communicative action appropriately, the participants are
854 *expected* to understand the communication *as though* a **service action** were actually performed.

855 When a customer ‘tells’ an airline service that it ‘confirms’ the purchase of the ticket it is simultaneously a
856 communication and a service action – two ways of understanding the same event, both actions, one
857 layered on top of the other, but with independent semantics.

858 3.1.4.1 Communicating for Action

859 **Actors** participating in a SOA ecosystem are often attempting to get other **actors** to *do* something – the
860 **service action** alluded to above. For example, a customer trying to buy a book has to convince the book
861 selling **service** to deliver the book. Conversely the book selling service has to convince the customer to
862 pay for it.

863 When an actor agrees to a course of **action** as a result of its interactions with other **actors** it is **adopting**
864 an **objective**.

865 Objective

866 An **objective** is a **real world effect** that an **actor** wishes to achieve.

867 **Objectives** refer to **Real World Effects** that **actors** may actively consider achieving.

868 In general, there is a *subsumption* relationship between **actors’ goals** and their **objectives**: an **objective**
869 can be considered to be *consistent* with one of more **goals**. Generally, a **goal** is a long term state of the

¹³ Acting against a service should not be understood to mean acting to foil the effectiveness of the service; but simply as an action involving the normal operation of the service.

870 world that may be, in practice, difficult to measure. On the other hand, an **objective** is a directly
871 measurable and preferably predictable outcome of a particular **action** or set of **actions**.

872 **Objective Adoption**

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

874 A consequence of an **actor** adopting an **objective** on behalf of another **actor** is that the actor becomes
875 **accountable** to the latter for the successful satisfaction of the **objective**.

876 **Accountability**

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

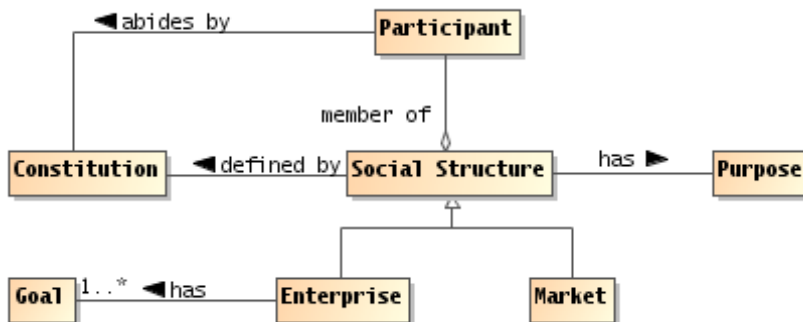
879 An **objective** adopted by one **actor** as a result of an interaction need not be consistent with the
880 **objectives** of the originating **actor**. In many situations, the adopted **objective** is not all the same and may
881 even be contrary to the desires of the original **actor**.

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

884 **3.2 Social Structure Model**

885 The actions undertaken by participants, whether mediated by services or in some other way, are
886 performed in a context that defines the meaning of the actions themselves. That context is fundamentally
887 a *social context* – a context that includes other participants. We can formalize that context as a **social**
888 **structure**: the embodiment of a particular social context.

889 The social structure model is important to defining and understanding the implications of crossing
890 ownership boundaries; it is the foundation for an understanding of security in SOA and also provides the
891 context for determining how SOA-based systems can be effectively managed and governed.



892
893 *Figure 10 Social Structure*

894 **Social Structure**

895 A **social structure**¹⁴ embodies some of the cultural aspects that characterize the relationships
896 and **actions** among a group of **participants**.

897 A **social structure** may have any number of participants, and a given participant can be a member of
898 multiple social structures. Thus, there is frequent interaction among social structures, sometimes resulting
899 in disagreements when the goals of the social structures do not align.

900 In the Reference Architecture, we are concerned primarily with **social structures** that reflect the
901 anticipated participants in SOA-based systems; these are often embodied in legal and quasi-legal
902 frameworks; i.e., they have some rules that are commonly understood. For example, an **enterprise** is a
903 common kind of **social structure**, as is an online chat room. At the other extreme, the legal frameworks
904 of entire countries and regions also count as social structures.

¹⁴ Social structures are sometimes referred to as social institutions.

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

909 **Purpose**

910 A measurable condition ascribed to a thing or action relating it to a goal.
 911 By their nature, purposes are *external* to the purposed entities, whereas goals are *internal* to the entity.
 912 A **social structure** has a *purpose* – the reason for which it exists. All **social structures** have a purpose,
 913 some **social structures** also have **goals**.

914 **Constitution**

915 A **constitution** is an agreement shared by a group of **participants** that defines a **social**
 916 **structure**.

917 Every **social structure** defines the rules by which **participants** interact with each other within the
 918 structure. Whether or not it is explicitly written, the **constitution** is that agreement that identifies the
 919 **social structure** itself.

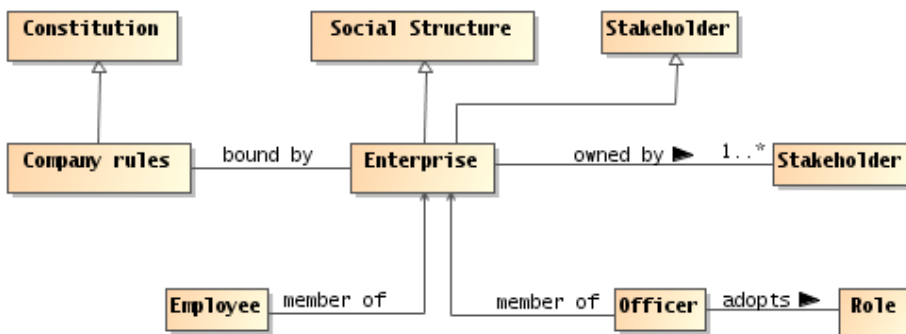
920 A **social structure's** rules are *abided to* by the **participants**. In some cases, this is based on an explicit
 921 agreement, in other cases participants behave as though they agree to the constitution without a formal
 922 agreement. In other cases, participants abide by the rules with some degree of reluctance – this is an
 923 issue raised later on when we discuss governance in SOA-based systems.

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

928 **Enterprise**

929 An **enterprise** is an organization with identifiable officers and with internally established **goals**
 930 that reflect the purpose of the organization.

931 The **enterprise** is marked out as being associated with internal **goals** in a way that a strict market type of
 932 social structure is not. Figure 11 shows a simplified model of enterprises as they relate to social
 933 structures.



934
 935 *Figure 11 Enterprise as a Social Structure*

936 **Market**

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

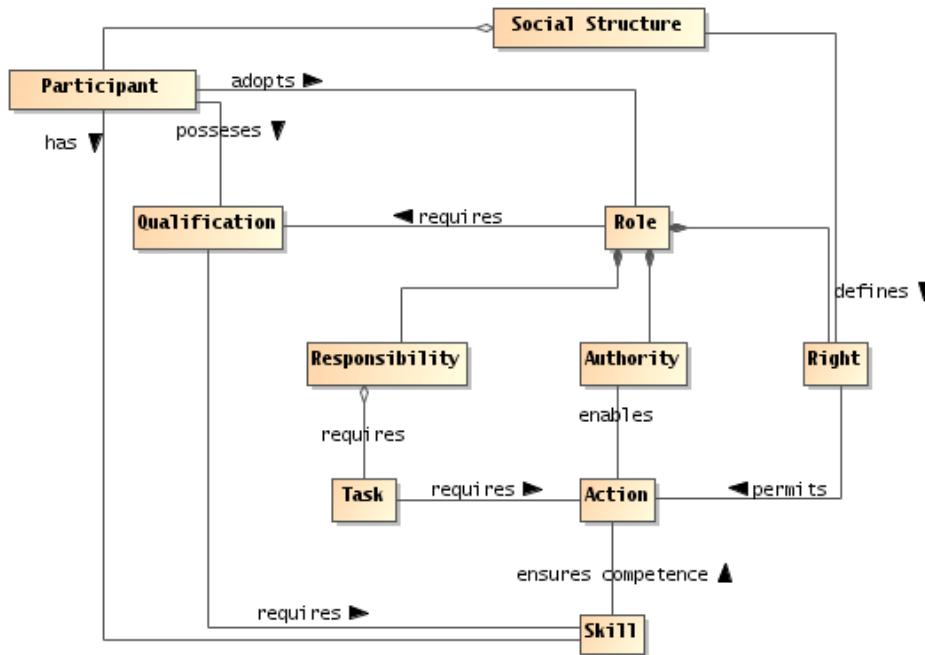
939 If an **enterprise** is often the focus of the differing **roles** and **responsibilities** of members, a **market** or
 940 meeting place is more concerned with the exchange of goods and services for mutual benefit.

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

944 **3.2.1 Roles in Social Structures**

945 One of the primary benefits of formalizing the relationships between people in terms of groups,
 946 corporations, legal entities and so on, is that it allows greater efficiencies in the operation of society.
 947 However, corporations, governments and even society, are abstractions: a government is not a person
 948 that can perform actions – only people or automated processes following the instructions of people can
 949 actually do things.

950 For example, a fishing club is an abstraction that is important to its members. A club, however, is an
 951 abstraction that has no physical ability to act in the world. On the other hand, a person who is
 952 appropriately empowered by the fishing club can act. For example, when that person writes a check and
 953 mails it to the telephone company, that action counts as though the fishing club has paid its bills.



954
 955 *Figure 12 Roles, Rights and Responsibilities*

956 Participants' actions within a social structure are often defined by the roles that they adopt.

957 **Role**

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

961 For many scenarios, the roles of participants are easily identified: for example, a buyer uses the service
 962 offered by the seller to achieve a purchase. However, in particular in situations involving delegation, the
 963 role of a participant may be considerably more complex.

964 A participant can be identified with one or more **roles**. Someone in authority in the social structure may
 965 have formally designated the participant as assuming the role with associated rights and responsibilities.
 966 Qualification and skill describe the expectations of the social structure in who should fill the role, but
 967 formal certifications of those qualifications and skills may or may not be required of the designated
 968 participant.

969 Conversely, someone who exhibits qualification and skill may by consensus assume the role without any
 970 formal designation. Someone with some degree of qualification and skill may become identified with a
 971 role because they perform the associated tasks.

972 Note that, while many roles are clearly identified, with appropriate names and definitions of the
 973 responsibilities, it is also entirely possible to separately bestow rights, responsibilities and so on; usually
 974 in a temporary fashion. For example, when a CEO delegates the responsibility of ensuring that the

975 company accounts are correct to the CTO, this does not imply that the CTO is adopting the full role of
976 CFO.

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

979 **Right**

980 A **right** is a predetermined permission that permits an **actor** to perform some action or adopt a
981 stance in relation to the **social structure** and other **actors**.

982 For example, in most circumstances, sellers have a right to refuse service to potential customers; but may
983 only do so based on certain criteria.

984 **Authority**

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

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

988 An entity may authorize or be assigned another entity to act as its agent. Often the actions that are so
989 authorized are restricted in some sense. In the case of human organizations, the only way that they can
990 act is via an agent.

991 Rights, authorities, responsibilities and roles form the foundation for the security architecture of the
992 Reference Architecture. Rights and responsibilities have similar structure to permissive and obligation
993 policies; except that the focus is from the perspective of the constrained participant rather than the
994 constrained actions.

995 **Responsibility**

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

998 **Skill**

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

1000 Skills are typically associated with **roles** in terms of requirements: a given role description may require
1001 that the role player has a certain skill.

1002 **Qualification**

1003 A **qualification** is a public determination by an issuing authority that an **actor** has achieved some
1004 state.

1005 The issuing authority may require some successful actions on the part of the **actor** (such as
1006 demonstrating some skills). The qualification may have constraints attached to it; for example, the
1007 certification may be time limited.

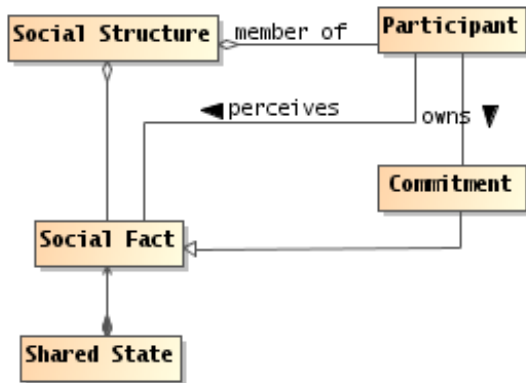
1008 There is a distinction between a **skill** – which is capability that a participant may have to act – and a
1009 publicly accepted right to act. For example, someone may have the skills to fly an airplane but not have a
1010 pilot's license. Conversely, someone may have a pilot license, but because of some temporary cause be
1011 incapable of flying a plane (they may be ill for example).

1012 Qualifications are often used as constraints on roles: any entity adopting a role within an organization (or
1013 other social structure) must have certain qualifications.

1014 **3.2.2 Shared State and Social Facts**

1015 Many of the actions performed by people and most of the important aspects of a person's state are
1016 inherently social in nature. The social context of an action is what gives it much of its meaning. We call
1017 actions in society social actions and, those facts that are understood in a society, social facts. It is often
1018 the case that social actions give rise to social facts.

1019 Compared to facts about the natural world, **social facts** are inherently abstract: they only have meaning
1020 in the context of a social structure.



1021
1022 *Figure 13 Shared State and Social Facts*

1023 **Social Fact**

1024 A **social fact** is an element of the state of a social structure that is defined by that social
1025 structure.

1026 Social structures provide a context in which social facts are given their meaning. For example, the
1027 existence of a valid purchase order with a particular customer has a meaning that is defined primarily by
1028 the company itself, together with the society that the company is part of.

1029 Social facts typically require some kind of ritual to establish the validity of the fact itself. For example, the
1030 existence of an agreed contract typically requires both parties to sign papers and to exchange those
1031 papers. If the signatures are not performed correctly, or if the parties are not properly empowered to
1032 perform the ritual, then it is as though nothing happened.

1033 In the case of agreements reached by electronic means, this involves the exchange of electronic
1034 messages; often with special tokens being exchanged in place of a hand-written signature.

1035 **State**

1036 State is the condition that an entity is in at a particular time.

1037 State is characterized by a set of facts that is true of the entity – in effect we are concerned only with
1038 aspects of an entity that are potentially measurable.

1039 **Private State**

1040 Private state is the set of facts that is known and understood by a participant.

1041 **Shared State**

1042 The set of facts that are knowable by participants as a result of their communicative actions.

1043 Note that shared state *does not* imply the state *is* known to all participants. It simply refers to the
1044 elements of state that *may* be known.

1045 Note that any **participant** has only a partial view of the world. Furthermore, the **participant** will have
1046 internal **private state** that is not accessible to other participants directly. However, elements of the shared
1047 state are in principle accessible to participants even if a given participant does not have access to all
1048 elements at any given time.

1049 **Public Semantics**

1050 The **public semantics** of a **communicative action** is the set of facts that any observer of the
1051 action would be sanctioned to infer by virtue of the observer's situation in a **social structure**.

1052 Of course, the most obvious observer of a communication is the intended recipient of the communication.
1053 However, the key is that the **public semantics** of a communication would enable *any* observer to make
1054 the same inferences.

1055 For example, a standard purchase order denotes a commitment to buy some goods or services. Any
1056 observer of the purchase order would be entitled to interpret it as a purchase order (whether or not the
1057 purchase order was targeted at the observer).

1058 **Public semantics** is often couched in terms of the **shared state** of the various members of the social
1059 structure – a purchase order is interpreted relative to the **social structure** within which it is made.

1060 **Commitment**

1061 A **commitment** is a **social fact** about the future: in the future some fact will be true and a
1062 participant has the current responsibility of ensuring that that fact will indeed be true.

1063 A **commitment** to deliver some good or service is a classic example of a fact about the future.

1064 Other important classes of social facts include the policies adopted by an organization, any agreements
1065 that it is holding for participants, and the assignment of participants to roles within the organization.

1066 Facts have the property of being verifiable (technically, a social fact can be verified to determine if it is
1067 satisfied in the social context). If, as a result of interacting with a service, a buyer incurs the obligation of
1068 paying for some good or service, this obligation (and the discharge of it) is measurable (perhaps by
1069 further interactions with the same or other services).

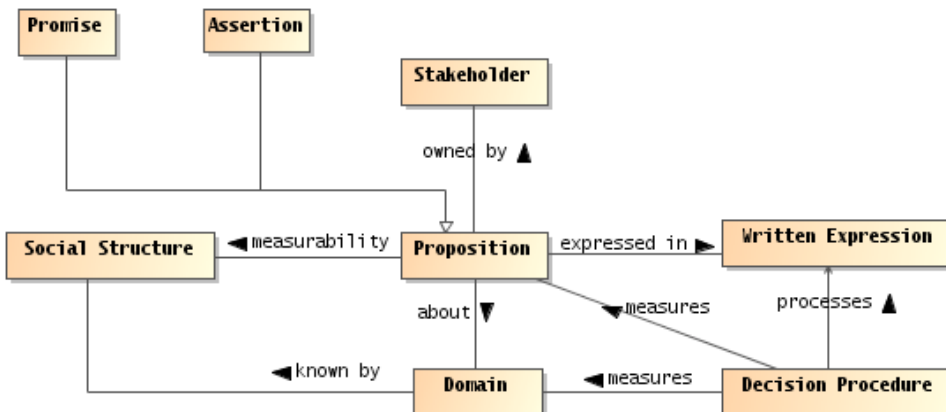
1070 **3.2.2.1 Proposition**

1071 When a participant wishes to share knowledge of a social fact or commitment, it may take the form of one
1072 or more Propositions.

1073 **Proposition**

1074 A proposition is an expression, normally in a language that has a well-defined written form, that
1075 expresses some property of the world from the perspective of a stakeholder.

1076 In principle, the truth of a proposition must be verifiable – using a decision procedure – by examining the
1077 world and checking that the proposition and the world are consistent with each other.¹⁵



1078
1079 *Figure 14 Propositions*

1080 **Decision Procedure**

1081 A decision procedure is a process for determining whether an expression is true, or is satisfied, in
1082 the world.

1083 Decision procedures are algorithms, programs that can measure the world against a formula, expression
1084 or description and answer the question whether the world corresponds to the description. If the truth of a
1085 proposition is indeterminable, then a decision procedure does not exist, and the logic is undecidable.

1086 **Domain**

1087 A domain is a 'world' that is used as the basis for the truth of a proposition.

¹⁵ We exclude here the special case of proposition known as a tautology. Tautologies are important in the study of logic; the kinds of propositions that we are primarily interested in are those which pertain to the world; and as such are only *contingently* true.

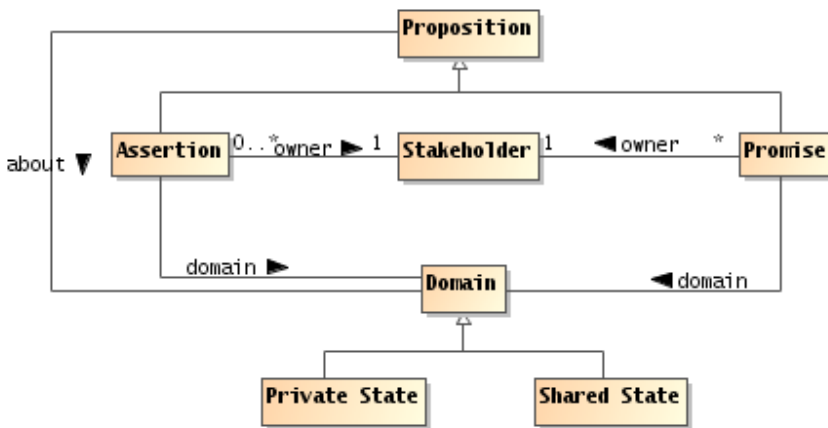
1088 When we say 'world', we are not restricted to the physical world. The criterion is an ability to discover
 1089 facts about it. In our case governmental, commercial and social structures that form the backdrop for
 1090 SOA-based systems are important examples of modeled worlds.

1091 **Written Expression**

1092 The written expression of a proposition is a formula written in a systematic system of marks that
 1093 denotes the proposition.

1094 Note that not all 'systems of marks' have a decision procedure. However, for the uses to which we put the
 1095 concept of proposition: policies, service descriptions, and so on, we require that the language used to
 1096 write policy and other propositions have a decision procedure.

1097 Propositions, as used in reference to needs, policies and contracts can be further analyzed in terms of
 1098 facts that are about the world as it is, will be, or should be. The latter are particularly of concern in policies
 1099 and contracts and other propositions concerning the relationships between people.



1100
 1101 *Figure 15 Assertions and Promises*

1102 **Assertion**

1103 An assertion is a proposition that is held to be true by a stakeholder. It is essentially a claim about
 1104 the state of the world.

1105 **Promise**

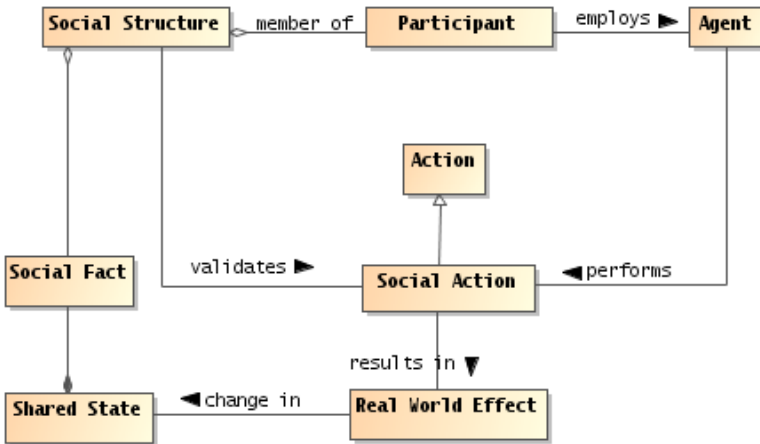
1106 A promise is a proposition regarding the future state of the world by a stakeholder. In particular, it
 1107 represents a commitment by the stakeholder to ensure the truth of the proposition.

1108 For example, an airline may report its record in on-time departures for its various flights. This is a claim
 1109 made by the airline which is, in principle, verifiable. The same airline may promise that some percentage
 1110 of its flights depart within 5 minutes of their scheduled departure. The truth of this promise depends on
 1111 the effectiveness of the airline in meeting its commitments.

1112 Another way of contrasting assertions and promises is to see what happens when the propositions fail: a
 1113 stakeholder that makes a false assertion about the world might be classified as a liar; a stakeholder that
 1114 makes a false promise is said to break its promises.

1115 **3.3 Acting in a Social Context Model**

1116 In the context of SOA, actions are primarily social in nature — one participant is asking another to do
 1117 something that is directly related to the organization(s) that they are part of — and goal oriented — the
 1118 purpose of interacting with a service is to satisfy a need by attempting to ensure that a remote entity
 1119 applies its capabilities to the need.



1120
1121 *Figure 16 Acting within Social Structures*

1122 **Social Action**

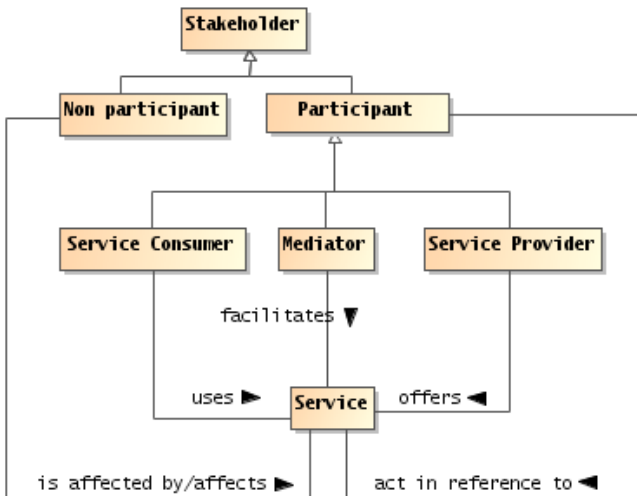
1123 Social actions are actions that are performed in order to achieve some result within a social
1124 structure.

1125 A social action is an action that is defined primarily by the effect it has on the relationship between
1126 participants and state of a social structure by establishing one or more new social facts.

1127 Social actions are always contextualized by a social structure: the organization gives meaning to the
1128 action, and often defines the requirements for an action to be recognized as having an effect within the
1129 organization.

1130 **3.3.1 Service Providers and Consumers**

1131 Section 3.1.1 defines the distinction between participants and nonparticipants. In a SOA social structure,
1132 several types of participants play prominent roles.



1133
1134 *Figure 17 Service Participants*

1135 **Service Provider**

1136 A service provider is a participant that offers a service that enables some capability to be used by
1137 other participants.

1138 Note that several kinds of stakeholders may be involved in provisioning a service. These include but are
1139 not limited to the provider of the capability, an enabler that exposes it as a service, a mediator that

1140 translates and/or manages the relationship between service consumers and the service, a host that offers
1141 support for the service, a government that permits the service and/or collects taxes based on service
1142 interactions.

1143 **Service Consumer**

1144 A service consumer is a participant that interacts with a service in order to realize the real world
1145 effect produced by a capability to address a consumer need.

1146 It is a common understanding that service consumers typically initiate service interactions. Again, this is
1147 not necessarily true in all situations (for example, in publish-and-subscribe scenarios, a service consumer
1148 may initiate an initial subscription, but thereafter, the interactions are initiated by publishers). As with
1149 service providers, several stakeholders may be involved in a service interaction supporting the consumer.

1150 Service providers and service consumers do not represent truly symmetric roles: each participant has
1151 different objectives and often has different capabilities. However, the objectives and the conditions under
1152 which those objectives align are critical for a successful interaction to proceed.

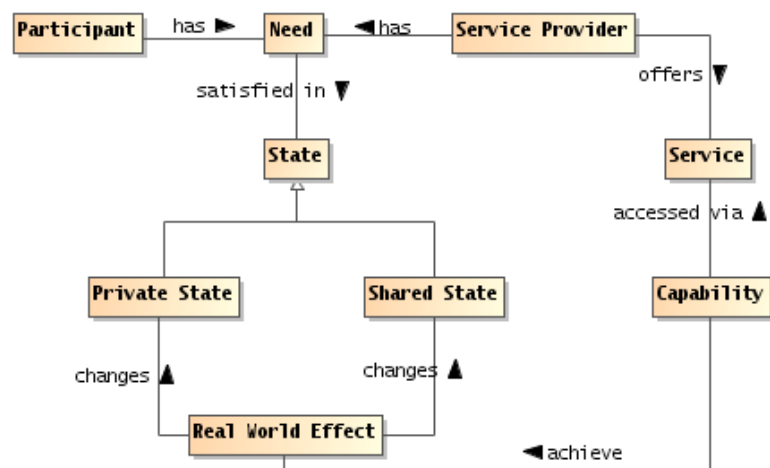
1153 **Service Mediator**

1154 A service mediator is a participant that facilitates the offering or use of services in some way.

1155 There are many kinds of mediator, for example a registry is a kind of mediator that permits providers and
1156 consumers to find each other. Another example might be a filter service that enhances another service by
1157 encrypting and decrypting messages. Yet another example of a mediator is a proxy broker that actively
1158 stands for one or other party in an interaction.

1159 **3.3.2 Needs and Capabilities**

1160 The Reference Model defines SOA in terms of a bringing together of needs and capabilities – the primary
1161 *motivation* for actors to engage with each other. A provider has a capability of generating a set of real
1162 world effects and making that capability available contributes to the satisfaction of some set of provider
1163 needs. The consumer has a need for those real world effects and has the capability of providing monetary
1164 or other return (for example, acknowledgement of effort) to the provider.



1165
1166 *Figure 18 Needs and Capabilities*

1167 **Need**

1168 A need is a measurable requirement that a service participant is actively seeking to satisfy.

1169 A need may or may not be publicly measurable; the needs that this Reference Architecture finds in scope
1170 are those that are publicly measurable. However, the satisfaction of a participant's need can only be
1171 determined by that participant.

1172 The extent to which a need is captured in a formal way is likely to be very different in each situation.

1173 Capability


1174 A capability is an ability to achieve a real world effect.

1175 The model in Section 3.1.3 shows that there is often some indirection between needs and having them
1176 satisfied. Both needs and the effects of using capabilities are expressed in terms of state: a need is
1177 expressed as a condition on the desired state and the Real World Effect of using capabilities is a change
1178 in the state of the world.

1179 By making a capability available for use, the owners aim to address their needs as well as the needs of
1180 other participants who use the service. The extent to which a capability is exposed via a service (or via
1181 multiple services) is controlled by the owner of the capability but may also be limited by the service
1182 provider. As noted in the Reference Model, a given service is not required to provide access to all aspects
1183 of an underlying capability.

1184 3.3.3 Resources

1185 In the SOA-RM and this Reference Architecture, we discuss service, underlying capabilities, and
1186 numerous other entities that are part of the SOA ecosystem. We categorize these as Resources, and
1187 define **Resource** as follows:



1188

1189 *Figure 19 Resources*

1190 Resource

1191 A resource is any entity of some perceived value that has identity.

1192 A resource may have more than one identifier, but any well-formed identifier should unambiguously
1193 resolve to the intended resource.

1194 An important class of resource is the class of capabilities that underlie services. Other examples of
1195 resources are services themselves, descriptions of entities (a kind of meta-resource), IT infrastructure
1196 elements used to deliver services, contracts and policies, and so on.

1197 Identity

1198 Identity is the collection of individual characteristics by which an entity, human or nonhuman, is
1199 recognized or known.

1200 The ability to unambiguously identify a resource in a SOA interaction is critical to determine such things
1201 as authorizations, to understand what functions are being performed and what the results mean, and to
1202 ensure repeatability or characterize differences with future SOA interactions.

1203 Identifier

1204 An identifier is any block of data – such as a string – that unambiguously connects a resource
1205 with a particular identity.

1206 Identifiers typically require a context in order to establish the connection between the identifier and the
1207 resource. A given resource may have multiple identifiers, with different utility for different contexts.

1208 In a SOA eco-system, it is good practice to use globally unique identifiers; for example globally unique
1209 IRIs. An identifier must uniquely disambiguate the indicated resource from other resources but more than
1210 one identifier may uniquely resolve to the same resource.

1211 Description

1212 A description is a structure that may be interpreted as containing assertions about a resource.

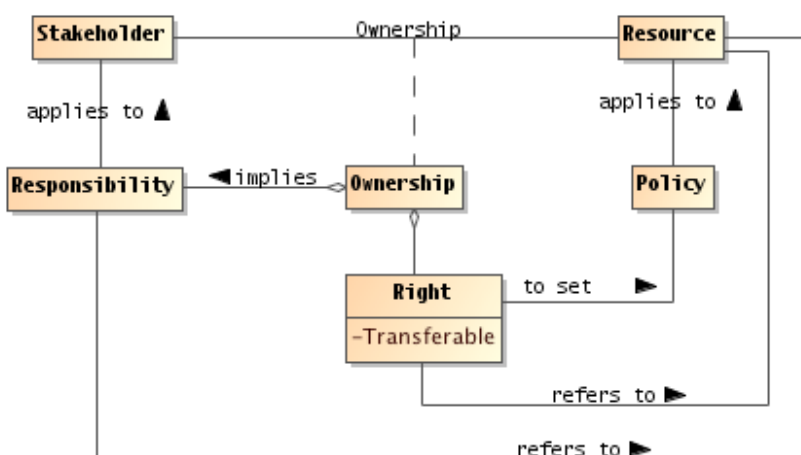
1213 This model of resource is a simplification and an elaboration of the concept that underlies the Web
 1214 Architecture [WA]. Being more abstract, we do not require that the identity of a resource be in any
 1215 particular form (although in practice, many resource identifiers are URIs), nor do we require resources to
 1216 have representations. However, we do require resources to have owners.

1217 3.3.4 Ownership

1218 A fundamental aspect of a resource is that it is owned by a stakeholder. Ownership is also important in
 1219 understanding the various kinds of obligations participants may enter into. Fundamentally, we view
 1220 ownership as a relationship between a stakeholder and a resource, where the owner has certain rights
 1221 over the resource.

1222 Ownership

1223 Ownership is a set of rights and responsibilities that a stakeholder has in relation to a resource;
 1224 including the right to transfer that ownership to another entity.



1225
 1226 *Figure 20 Resource Ownership*

1227 To own a resource implies taking responsibility for creating, maintaining, and if it is to be available to
 1228 others, provisioning the resource. More than one stakeholder may own different rights, such as one
 1229 stakeholder having the right to deploy a capability as a service, another owning the rights to the profits
 1230 that result from using the capability, and yet another owning the rights to use the service.

1231 One who owns a resource may delegate rights and responsibilities to others, but typically retains some
 1232 responsibility to see that the delegated responsibilities are met. There may also be joint ownership of a
 1233 resource, where the responsibility is shared.

1234 A crucial property that distinguishes ownership from a more limited right to use is the right to transfer
 1235 ownership to another person or organization. When a resource is being used without being owned, there
 1236 is an implied requirement that at the end of a period of time the rights and responsibilities relating to the
 1237 resource will be returned to the original owner of the resource.

1238 Ownership is defined in relation to the social structure relative to which rights and responsibilities are
 1239 exercised. In particular, there may be constraints on how ownership may be transferred. For example, a
 1240 government may not permit a corporation to transfer assets to a subsidiary in a different jurisdiction.

1241 Ownership Boundary

1242 An **ownership boundary** is the **social structure** within which the **rights** and **responsibilities**
 1243 associated with a particular ownership may be recognized.

1244 Individual participants are *within* an ownership boundary in relation to a specific owned **resource** if they
 1245 are members of the **social structure** that owns the **resource**.

1246 **3.3.5 Trust, Risk and Willingness**

1247 For interactions to be possible within the SOA ecosystem, each actor must have a sufficient degree of
1248 trust in other actors to form a basis for willingness to engage in the interactions.

1249 **Trust**

1250 **Trust** is a private assessment or internal perception that some entity will perform actions that will
1251 lead to an identifiable set of real world effects.

1252 The reference to real world effects implies the existence of measurements or other observations of
1253 shared state that represent the real world effect.

1254 **Willingness**

1255 **Willingness** is the internal commitment of an actor to carry out its part of an interaction.

1256 As discussed in the Reference Model, willingness on the part of actors to interact is not the same as a
1257 willingness to perform requested actions. A service provider that rejects all attempts to cause it to perform
1258 some action may still be fully willing and engaged in interacting with the consumer.

1259 **Trusting Actor**

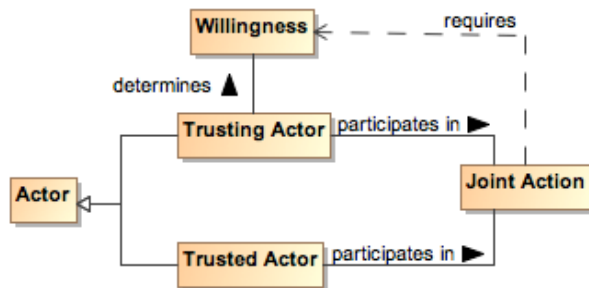
1260 A **Trusting Actor** is an actor who establishes and maintains willingness to proceed with an
1261 interaction based on its trust of other actors.

1262 Typically, it is not important to know the specific actions undertaken by any given actor because these
1263 may be private. Additionally, it is not important to share or even to know the goals of the individual actors
1264 as long as the Trusting Actor believes that individual actions by others will be sufficient to result in
1265 expected real world effects. For example, the Trusting Actor may have a desired real world effect of an
1266 important message being delivered and is willing to pay for this business service; those delivering the
1267 message have no interest in the importance of the message but want to do what is necessary to ensure
1268 payment. Successful completion of the interaction will result in both (and possibly other) real world effects
1269 to be realized.

1270 **Trusted Actor**

1271 A **Trusted Actor** is an actor with which a Trusting Actor has sufficient trust for that Trusting Actor
1272 to be willing to proceed with an interaction.

1273 The relationship of Willingness to the Trusting and Trusted Actors is shown in Figure 21.



1274
1275 *Figure 21 Trusting Actor and Willingness*

1276 **Risk**

1277 **Risk** is a private assessment or internal perception that certain undesirable real world effects may
1278 come into being.

1279 The Actor perceiving risk may take actions to mitigate the risk. For example, the actor may assess a high
1280 degree of risk to clicking on an email link where the actor believes the email to be spam, and the actor
1281 forgoes any possible benefit by not clicking on the link. Alternately, the actor may see a risk in having a
1282 hard drive fail and mitigate the effect of losing files by backing up those files considered important.

1283 **3.3.5.1 Assessing Trust and Risk**

1284 The assessments of trust and risk are based on evidence available to the Trusting Actor.

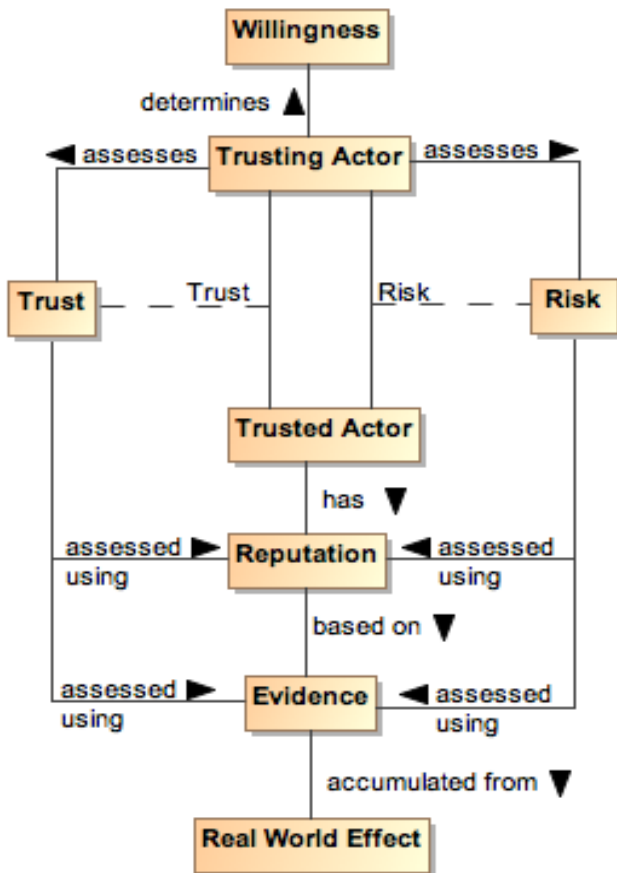
1285 **Evidence**

1286 **Evidence** is the accumulation of real world effects by which a Trusting Actor can assess trust and
1287 risk.

1288 The evidence may be physical artifacts or a set of information from which the Trusting Actor can assess
1289 the degree of trust. The evidence may include a history of previous interaction between the Trusting and
1290 Trusted Actors or previous interactions of the Trusted Actor with other actors for which the real world
1291 effects of their interactions are public. Such an accumulation of real world effects forms the basis of the
1292 Trusted Actor's reputation.

1293 **Reputation**

1294 **Reputation** is the social assessment of an actor with respect to an expectation of behavior or
1295 skill, where the assessment is made on the basis of evidence.



1296
1297 *Figure 22 Assessing Trust and Risk*

1298 Trust is based on the confidence the Trusting Actor has in the accuracy and sufficiency of the gathered
1299 evidence and the degree to which any assessment is appropriate for the situation for which trust is being
1300 assessed. Trust is not binary, i.e. an Actor is not completely trusted or untrusted, because there is
1301 typically some degree of uncertainty in the accuracy or completeness of the evidence or the assessment.
1302 Similarly, there is uncertainty in the amount and consequences of potential risk.

1303 The balance between perceived trust and perceived risk results in a willingness or unwillingness to
1304 proceed. If there is little or no perceived risk, then the degree of trust may not be relevant in assessing
1305 possible actions. For example, most people consider there to be an acceptable level of risk to privacy
1306 when using search engines, and submit queries without any sense of trust being considered.

1307 As perceived risk increases, the issue of trust becomes more of a consideration. There are recognized
1308 risks in providing or accepting credit cards as payment, and standard procedures have been put in place
1309 to increase trust or, at a minimum, bringing trust and risk into balance by mitigating risk. For interactions

1310 with a high degree of risk, the Trusting Actor requires stronger or additional evidence when evaluating the
1311 balance between risk and trust when deciding whether to participate in an interaction.

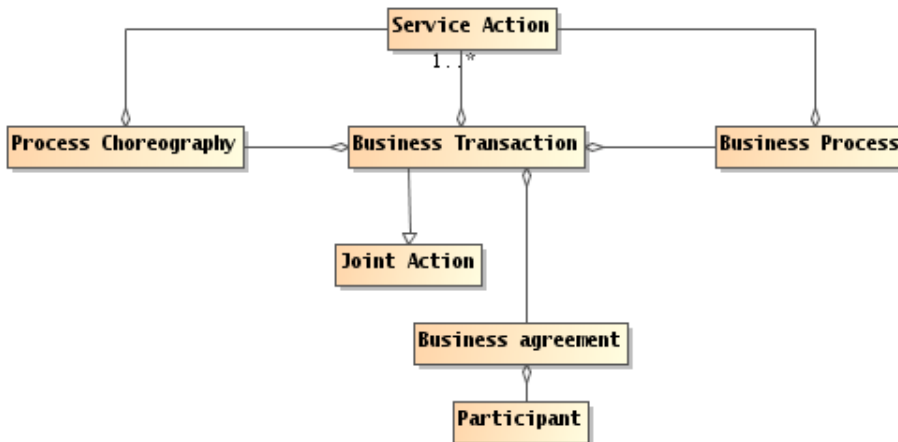
1312 3.3.5.2 Trust and SOA

1313 In traditional systems, the balance between trust and risk was achieved by severely restricting the
1314 interactions and those that could participate: the more trust and the higher the perceived risk, the more
1315 tightly coupled we made the corresponding system.

1316 Realizing many of the perceived benefits of SOA will require a fuller understanding of what trust and risk
1317 mean in the relevant business processes and what is an appropriate balance to be achieved. Actors
1318 need to assess trust and risk and act on those assessments while remaining part of the ecosystem and
1319 not just in a walled garden.

1320 3.3.6 Transactions and Exchanges

1321 An important class of **joint action** is the **business transaction**, or **contract exchange**. Many
1322 interactions between participants in the SOA ecosystem are based around business transactions.



1323
1324 *Figure 23 Business Transaction*

1325 Business Transaction

1326 A **business transaction** is a **joint action** engaged in by two or more **participants** in which the
1327 **ownership** of one of more **resources** is exchanged.

1328 A classic business transaction is buying some good or service, but there is a huge variety of kinds of
1329 possible business transactions.

1330 Key to the concept of business transaction is the contract or agreement to exchange. The form of the
1331 contract can vary from a simple handshake to an elaborately drawn contract with lawyers giving advice
1332 from all sides.

1333 A completed transaction establishes a set of social facts relating to the exchange; typically to the changes
1334 of ownerships of the resources being exchanged.

1335 Business Agreement

1336 A **business agreement** is an agreement entered into by two or more partners that constrains
1337 their future behaviors and permitted states.

1338 A business agreement is typically associated with business transactions: the transaction is guided by the
1339 agreement and an agreement can be the result of a transaction.

1340 Business transactions often have a well defined life-cycle: a negotiation phase in which the terms of the
1341 transaction are discussed, an agreement action which establishes the commitment to the transaction, an
1342 action phase in which the agreed-upon items are exchanged (they may need to be manufactured before
1343 they can be exchanged), and a termination phase in which there may be long-term commitments by both

1344 parties but no particular actions required (e.g., if the exchanged goods are found to be defective, then
1345 there is likely a commitment to repair or replace them).

1346 From an architectural perspective, the business transaction often represents the top-most mode of
1347 interpretation of service interactions. When participants interact in a service, they exchange information
1348 and perform actions that have an effect in the world. These exchanges can be interpreted as realizing
1349 part of, and in support of, business transactions.

1350 **Business Process**

1351 A **business process** is a description of the tasks, participants' roles and information needed to
1352 fulfill a business objective.

1353 Business processes are often used to describe the actions and interactions that form business
1354 transactions. This is most clear when the business process defines an activity involving parties external to
1355 the organization; however, even within an enterprise, a business process typically involves multiple
1356 participants and stakeholders.

1357 In the context of transactions mediated and supported by electronic means, business processes are often
1358 required to be defined well enough to permit automation. The forms of such definitions are often referred
1359 to as choreographies:

1360 **Process Choreography**

1361 A process choreography is a description of the possible interactions that may take place between
1362 two or more participants to fulfill an objective.

1363 A choreography is, in effect, a description of what the forms of permitted joint actions are when trying to
1364 achieve a particular result. Joint actions are by nature formed out of the individual actions of the
1365 participants; a choreography can be used to describe those interlocking actions that make up the joint
1366 action itself.

4 Realizing Service Oriented Architectures View

Make everything as simple as possible but no simpler.

Albert Einstein

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

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

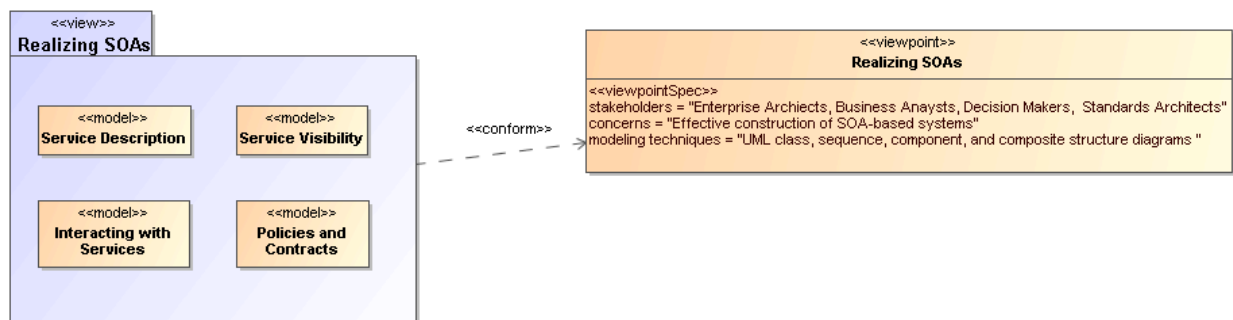


Figure 24 Model Elements Described in the Realizing a Service Oriented Architecture View

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

4.1 Service Description Model

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

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

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

- 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 and real world effects may be required for services accessed via automated service discovery and planning systems.

- 1405 • Descriptions come with context, i.e. a given description comprises information needed to adequately
1406 support the context. For example, a list of items can define a version of a service, but for many
1407 contexts an indicated version number is sufficient without the detailed list. The current model focuses
1408 on the description needed by a service consumer to understand what the service does, under what
1409 conditions will the service do it, how well does the service do it, and what steps are needed by the
1410 consumer to initiate and complete a service interaction. Such information also enables the service
1411 provider to clearly specify what is being provided and the intended conditions of use.
- 1412 • Descriptions will change over time as, for example, the ingredients and nutrition information for food
1413 labeling continues to evolve. A requirement for transparency of transactions may require additional
1414 description for those associated contexts.
- 1415 • Description always proceeds from a basis of what is considered "common knowledge". This may be
1416 social conventions that are commonly expected or possibly codified in law. It is impossible to describe
1417 everything and it can be expected that a mechanism as far reaching as SOA will also connect entities
1418 where there is inconsistent "common" knowledge.
- 1419 • Descriptions will become the collection point of information related to a service or any other resource,
1420 but it will not necessarily be the originating point or the motivation for generating this information. In
1421 particular, given a SOA service as the access to an underlying capability, the service may point to
1422 some of the capability's previously generated description, e.g. a service providing access to a data
1423 store may reference update records that indicate the freshness of the data.
- 1424 • Descriptions of the provider and consumer are the essential building blocks for establishing the
1425 execution context of an interaction.

1426 These points emphasize that there is no one "right" description for all contexts and for all time. Several
1427 descriptions for the same subject may exist at the same time, and this emphasizes the importance of the
1428 description referencing source material maintained by that material's owner rather than having multiple
1429 copies that become out of synch and inconsistent.

1430 It may also prove useful for a description assembled for one context to cross-reference description
1431 assembled for another context as a way of referencing ancillary information without overburdening any
1432 single description. Rather than a single artifact, description can be thought of as a web of documents that
1433 enhance the total available description.

1434 This Reference Architecture uses the term service description for consistency with the concept defined in
1435 the Reference Model. Some SOA literature treats the idea of a "service contract" as equivalent to service
1436 description. In this Reference Architecture, the term service description is preferred. Replacing service
1437 description with service contract implies just one side of the interaction is governing and misses the point
1438 that a single set of policies identified by a service description may lead to numerous contracts, i.e. service
1439 level agreements, leveraging the same description.

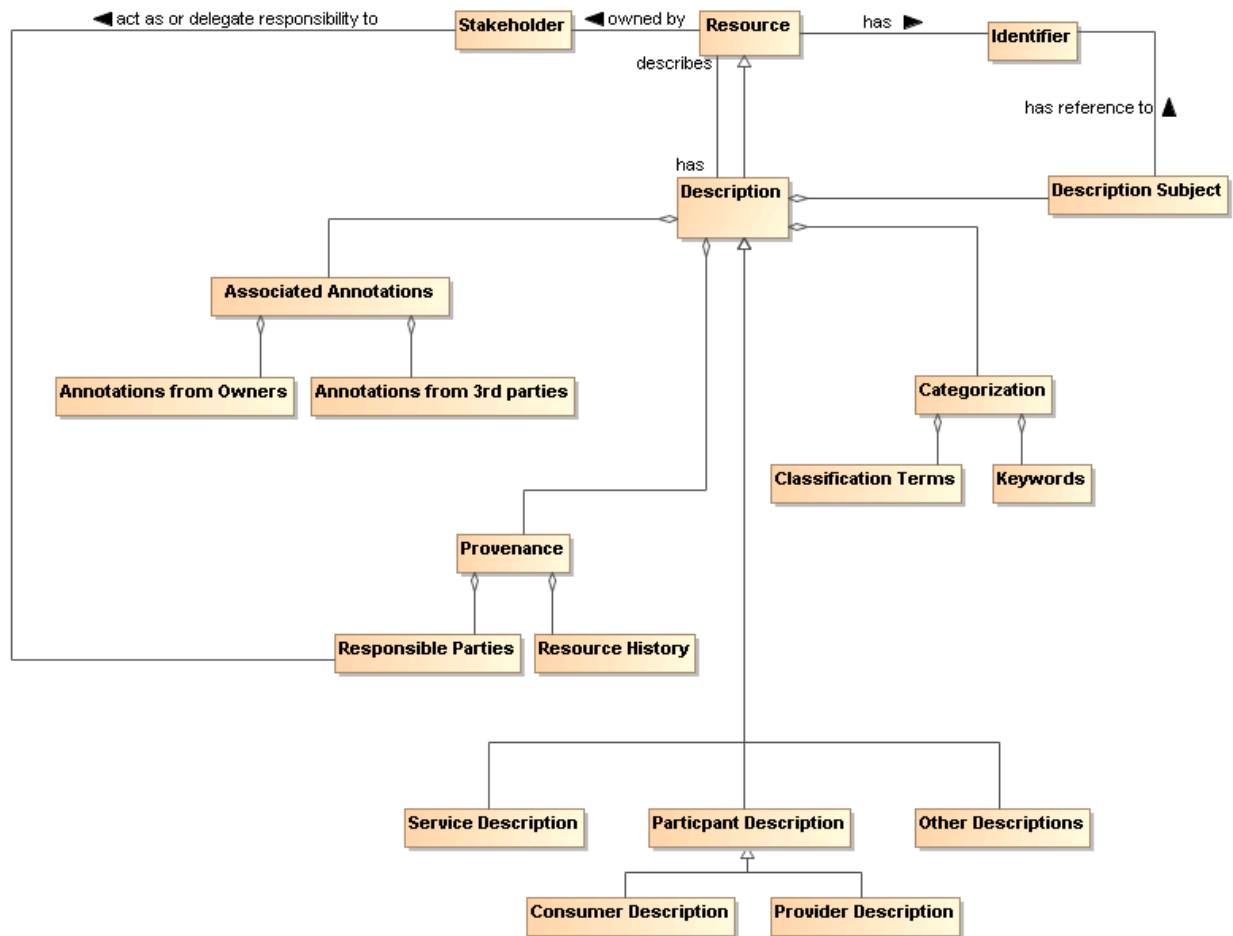
1440 **4.1.1 The Model for Service Description**

1441 *Figure 25* shows Service Description as a subclass of the general Description class, where Description is
1442 a subclass of the Resource class as defined in Section 3.3.3. In addition, each Resource is assumed to
1443 have a description. The following section discusses the relationships among elements of general
1444 description and the subsequent sections focus on service description itself. Note, other descriptions, such
1445 as those of participants, are important to SOA but are not individually elaborated in this document.

1446

1447 **4.1.1.1 Elements Common to General Description**

1448 The general Description class is composed of a number of elements that are expected to be common
1449 among all specialized descriptions supporting a service oriented architecture. A registry often contains a
1450 subset of the description instance, where the chosen subset is identified as that which facilitates mediated
1451 discovery. Additional information contained in a more complete description may be needed to initiate and
1452 continue interaction.



1453
1454 *Figure 25 General Description*

1455 **4.1.1.1.1 Description Subject**

1456 The subject of a description is a Resource. The value assigned to the Description Subject class may be
1457 of any form that provides understanding of what constitutes the Resource, but it is often in human-
1458 readable text. The Description Subject MUST also reference the Identifier of the resource it describes so
1459 it can unambiguously identify the subject of each description instance.

1460 As a Resource, Description also has an identifier with a unique value for each description instance. The
1461 description instance provides vital information needed to both establish visibility of the resource and to
1462 support its use in the execution context for the associated interaction. The identifier of the description
1463 instance allows the description itself to be referenced for discussion, access, or reuse of its content.

1464

1465 **4.1.1.1.2 Provenance**

1466 While the Resource Identifier provides the means to know which subject and subject description are
1467 being considered, Provenance as related to the Description class provides information that reflects on the
1468 quality or usability of the subject. Provenance specifically identifies the entity (human, defined role,
1469 organization, ...) that assumes responsibility for the resource being described and tracks historic
1470 information that establishes a context for understanding what the resource provides and how it has
1471 changed over time. Responsibilities may be directly assumed by the Stakeholder who owns a Resource
1472 or the Owner may designate Responsible Parties for the various aspects of maintaining the resource and
1473 provisioning it for use by others. There may be more than one entity identified under Responsible Parties;
1474 for example, one entity may be responsible for code maintenance while another is responsible for

1475 provisioning of the executable code. The historical aspects may also have multiple entries, such as when
 1476 and how data was collected and when and how it was subsequently processed, and as with other
 1477 elements of description, may provide links to other assets maintained by the Resource owner.

1478 **4.1.1.1.3 Keywords and Classification Terms**

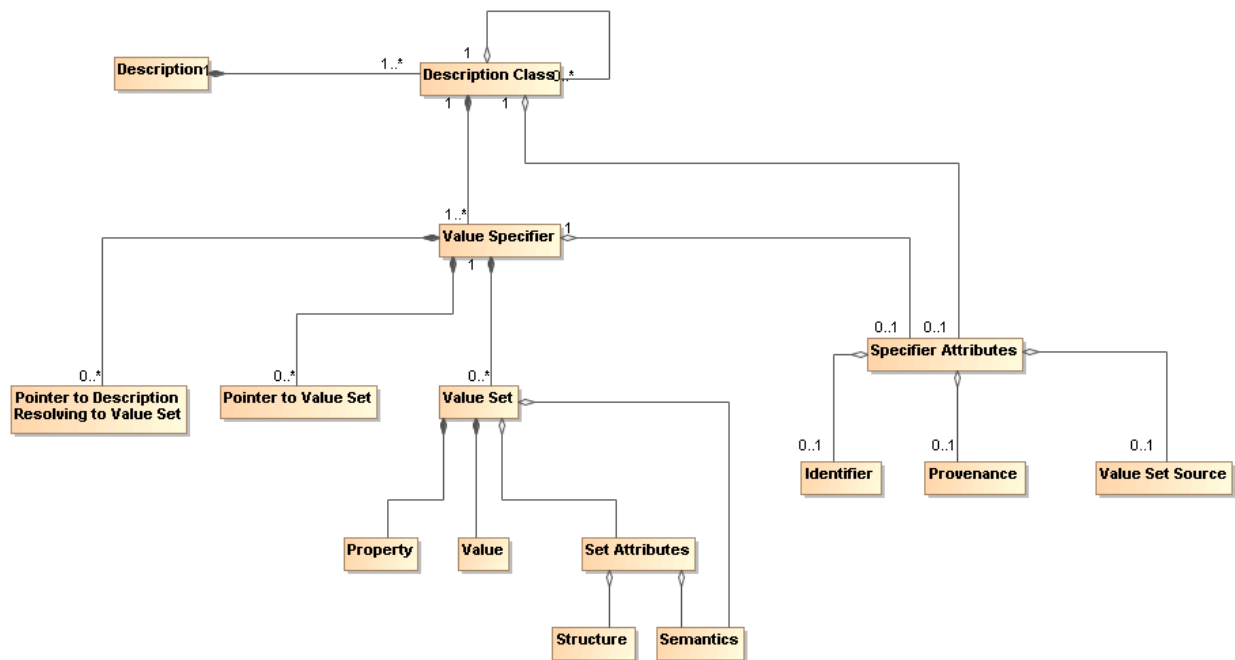
1479 A traditional element of description has been to associate the resource being described with predefined
 1480 keywords or classification taxonomies that derive from referenceable formal definitions and vocabularies.
 1481 This Reference Architecture does not prescribe which vocabularies or taxonomies may be referenced,
 1482 nor does it limit the number of keywords or classifications that may be associated with the resource. It
 1483 does, however, state that a normative definition SHOULD be referenced, whether that be a
 1484 representation in a formal ontology language, a pointer to an online dictionary, or any other accessible
 1485 source. See Section 4.1.1.2 for further discussion on associating semantics with assigned values.

1486 **4.1.1.1.4 Associated Annotations**

1487 The general description instance may also reference associated documentation that is in addition to that
 1488 considered necessary in this model. For example, the owner of a service may have documentation on
 1489 best practices for using the service. Alternately, a third party may certify a service based on their own
 1490 criteria and certification process; this may be vital information to other prospective consumers if they were
 1491 willing to accept the certification in lieu of having to perform another certification themselves. Note, while
 1492 the examples of Associated Documentation presented here are related to services, the concept applies
 1493 equally to description of other entities.
 1494

1495 **4.1.1.2 Assigning Values to Description Instances**

1496



1497
 1498 *Figure 26 Representation of a Description*

1499 Figure 25 shows the template for a general description but individual description instances depend on the
 1500 ability to associate meaningful values with the identified elements. Figure 26 shows a model for a
 1501 collection of information that provides for value assignment and traceability for both the value meaning
 1502 and the source of a value. The model is not meant to replace existing or future schema or other
 1503 structures that have or will be defined for specific implementations, but it is meant as guidance for the

1504 information such structures need to capture to generate sufficient description. It is expected that tools will
1505 be developed to assist the user in populating description and auto-filling many of these fields, and in that
1506 context, this model provides guidance to the tool developers.

1507 In Figure 26 each class has an associated value specifier or is made up of components that will
1508 eventually resolve to a value specifier. For example, Description has several components, one of which is
1509 Categorization, which would have an associated a value specifier.

1510 A value specifier consists of

- 1511 • a collection of value sets with associated property-value pairs, pointers to such value sets, or pointers
1512 to descriptions that eventually resolve to value sets that describe the component; and
- 1513 • attributes that qualify the value specifier and the value sets it contains.

1514 The qualifying attributes for the value specifier include

- 1515 • an optional identifier that would allow the value set to be defined, accessed, and reused elsewhere;
- 1516 • provenance information that identifies the party (individual, role, or organization) that has
1517 responsibility for assigning the value sets to any description component;
- 1518 • an optional source of the value set, if appropriate and meaningful, e.g. if a particular data source is
1519 mandated.

1520 If the value specifier is contained within a higher-level component, (such as Service Description
1521 containing Service Functionality), the component may inherit values from the attributes from its container.

1522 Note, provenance as a qualifying attribute of a value specifier is different from provenance as part of an
1523 instance of Description. Provenance for a service identifies those who own and are responsible for the
1524 service, as described in Section 3.3.4. Provenance for a value specifier identifies who is responsible for
1525 choosing and assigning values to the value sets that comprise the value specifier. It is assumed that
1526 granularity at the value specifier level is sufficient and provenance is not required for each value set.

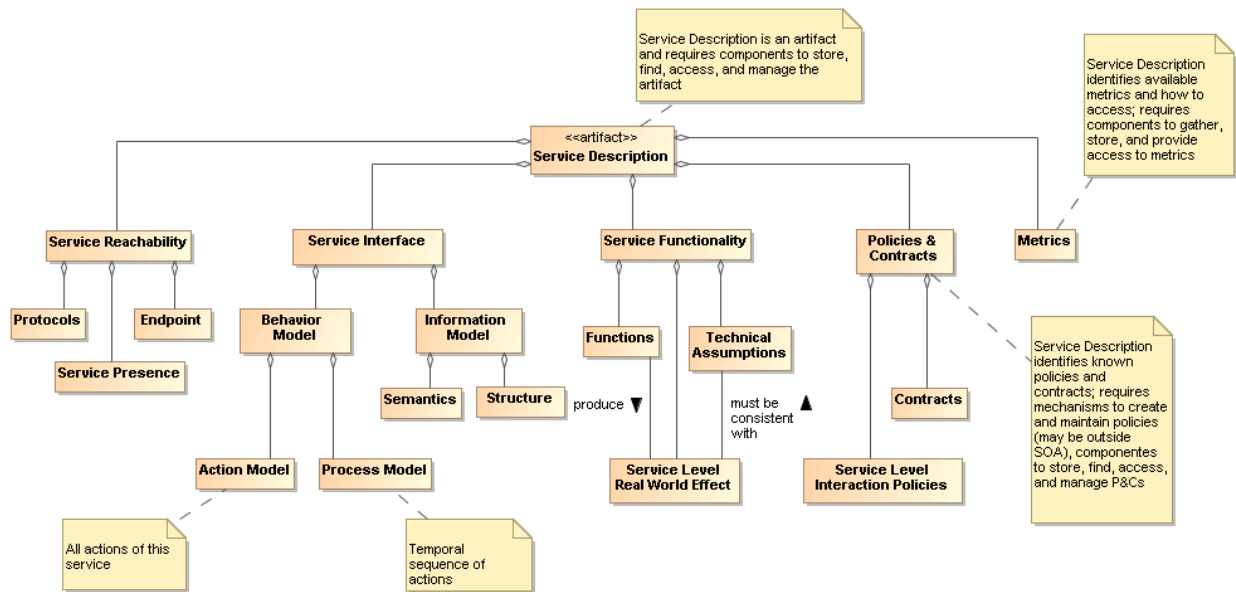
1527 The value set also has attributes that define its structure and semantics.

- 1528 • The semantics of the value set property should be associated with a semantic context conveying the
1529 meaning of the property within the execution context, where the semantic context could vary from a
1530 free text definition to a formal ontology.
- 1531 • For numeric values, the structure would provide the numeric format of the value and the “semantics”
1532 would be conveyed by a dimensional unit with an identifier to an authoritative source defining the
1533 dimensional unit and preferred mechanisms for its conversion to other dimensional units of like type.
- 1534 • For nonnumeric values, the structure would provide the data structure for the value representation
1535 and the semantics would be an associated semantic model.
- 1536 • For pointers, architectural guidelines would define the preferred addressing scheme.

1537 The value specifier may indicate a default semantic model for its component value sets and the individual
1538 value sets may provide an override.

1539 The property-value pair construct is introduced for the value set to emphasize the need to identify
1540 unambiguously both what is being specified and what is a consistent associated value. The further
1541 qualifying of Structure and Semantics in the Set Attributes allows for flexibility in defining the form of the
1542 associated values.

1543 **4.1.1.3 Model Elements Specific to Service Description**



1544
1545 *Figure 27 Service Description*

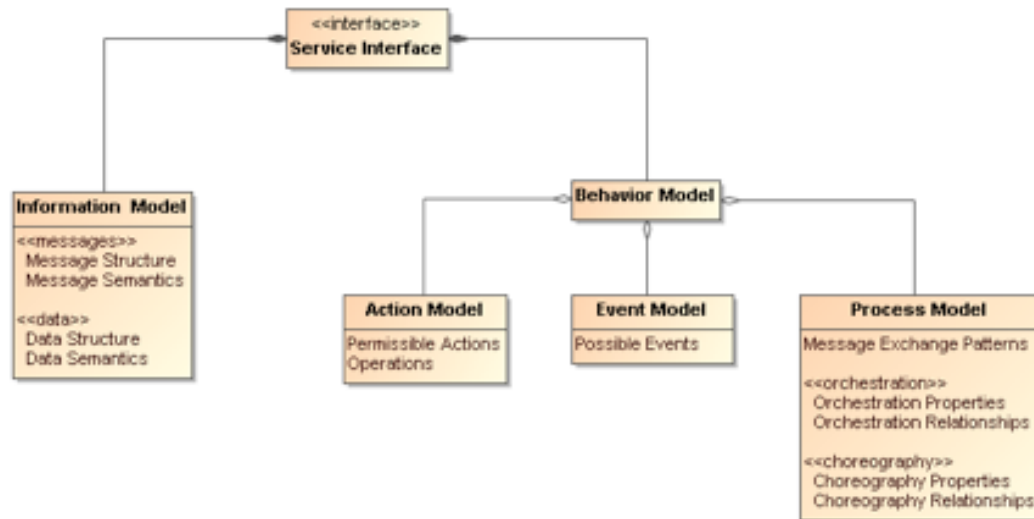
1546 The major elements for the Service Description subclass follow directly from the areas discussed in the
1547 Reference Model. Here, we discuss the detail shown in *Figure 27* and the purpose served by each
1548 element of service description.

1549 Note, the intent in the subsections that follow is to describe how a particular element, such as the service
1550 interface, is reflected in the service description, not to elaborate on the details of that element. Other
1551 sections of the Reference Model and this Reference Architecture describe the “physics” of each element
1552 whereas the service description subsections will only touch on the meta aspects.

1553 **4.1.1.3.1 Service Interface**

1554 As noted in the Reference Model, the service interface is the means for interacting with a service. For
1555 this Reference Architecture and as shown in Section 4.3 the service interface will support an exchange of
1556 messages, where

- 1557 • the message conforms to a referenceable message exchange pattern (MEP),
- 1558 • the message payload conforms to the structure and semantics of the indicated information model,
- 1559 • the messages are used to denote events or actions against the service, where the actions are
1560 specified in the action model and any required sequencing of actions is specified in the process
1561 model.



1562

1563 *Figure 28 Service Interface*

1564 Note we distinguish the structure and semantics of the message from that of the underlying protocol that
 1565 conveys the message. The message structure may include nested structures that are independently
 1566 defined, such as an enclosing envelope structure and an enclosed data structure.

1567 These aspects of messages are discussed in more detail in Section 4.3

1568 **4.1.1.3.2 Service Reachability**

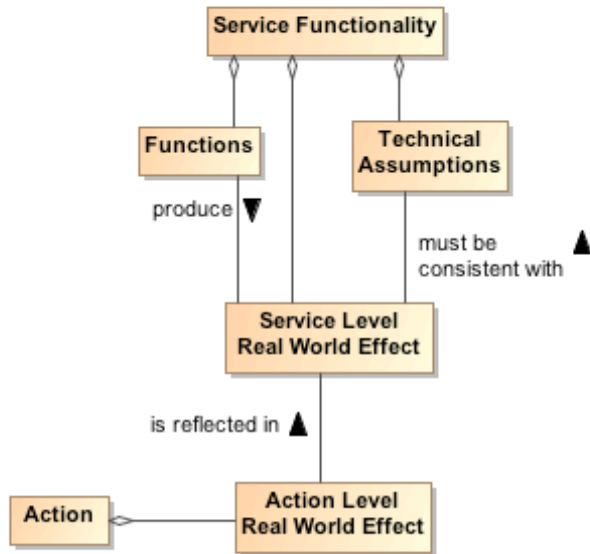
1569 Service reachability, as modeled in Section 4.2.2.3 enables service participants to locate and interact with
 1570 one another. To support service reachability, the service description should indicate the endpoints to
 1571 which a service consumer can direct messages to invoke actions and the protocol to be used for
 1572 message exchange using that endpoint.

1573 As applied in general to an action, the endpoint is the conceptual location where one applies an action;
 1574 with respect to service description, it is the actual address where a message is sent.

1575 In addition, the service description should provide information on collected metrics for service presence;
 1576 see Section 4.1.1.3.4 for the discussion of metrics as part of service description.

1577 **4.1.1.3.3 Service Functionality**

1578 While the service interface and service reachability are concerned with the mechanics of using a service,
 1579 service functionality and performance metrics (discussed in Section 4.1.1.3.4) describe what can be
 1580 expected when interacting with a service. Service Functionality, shown in *Figure 27* as part of the overall
 1581 Service Description model and extended in *Figure 29*, is an unambiguous expression of service function(s)
 1582 and the real world effects of invoking the function. The Functions likely represent business activities in
 1583 some domain that produce the desired Real World Effects.



1584
1585 *Figure 29 Service Functionality*

1586 The Service Functionality may also be constrained by Technical Assumptions that underlie the effects
1587 that can result. Technical assumptions are defined as domain specific restrictions and may express
1588 underlying physical limitations, such as flow speeds must be below sonic velocity or disk access that
1589 cannot be faster than the maximum for its host drive. Technical assumptions are likely related to the
1590 underlying capability accessed by the service. In any case, the Real World Effects must be consistent
1591 with the Technical Assumptions.

1592 In *Figure 27* and *Figure 29*, we specifically refer to Service Level and Action Level Real World Effects.

1593 **Service Level Real World Effect**

1594 A service level real world effect is a specific change in shared state or information returned as a
1595 result of interacting with a service.

1596 **Action Level Real World Effect**

1597 An action level real world effect is a specific change in shared state or information returned as a
1598 result of performing a specific action against a service.

1599 Service description describes the service as a whole while the component aspects should contribute to
1600 that whole. Thus, while individual Actions may contribute to the real world effects to be realized from
1601 interaction with the service, there would be a serious disconnect for Actions to contribute real world
1602 effects that could not consistently be reflected in the Service Level Real World Effects and thus the
1603 Service Functionality. The relationship to Action Level Real World Effects and the implications on
1604 defining the scope of a service are discussed in Section 4.1.2.1.

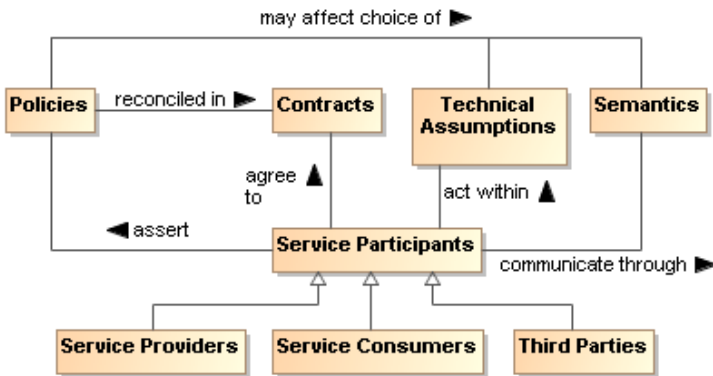
1605 Elements of Service Functionality may be expressed as natural language text, reference to an existing
1606 taxonomy of functions, or reference to a more formal knowledge capture providing richer description and
1607 context.

1608 **4.1.1.3.4 Policies and Contracts, Metrics, and Compliance Records**

1609 Policies prescribe the conditions and constraints for interacting with a service and impact the willingness
1610 to continue visibility with the other participants. Whereas technical assumptions are statements of
1611 “physical” fact, policies are subjective assertions made by the service provider (sometimes as passed on
1612 from higher authorities).


1613 The service description provides a central location for identifying what policies have been asserted by the
1614 service provider. The specific representation of the policy, e.g. in some formal policy language, is likely
1615 done outside of the service description and the service description would reference the normative
1616 definition of the policy.

1617 Policies may also be asserted by other service participants, as illustrated by the model shown in Figure
 1618 30. Policies that are generally applicable to any interaction with the service are likely to be asserted by
 1619 the service provider and included in the Policies and Contracts section of the service description.
 1620 Conversely, policies that are asserted by specific consumers or consumer communities would likely be
 1621 identified as part of a description's Annotations from 3rd parties (see Section 4.1.1.1.4) because these
 1622 would be specific to those parties and not a general aspect of the service being described.



1623
 1624 *Figure 30 Model for Policies and Contracts as related to Service Participants*

1625 In Figure 27 and Figure 31, we specifically refer to Service Level Interaction Policies. In a similar manner
 1626 to that discussed for Service Level vs. Action Level Real World Effects in Section 4.1.1.3.3, individual
 1627 Actions may have associated policies stating conditions for performing the action, but these must be
 1628 reflected in and be consistent with the policies made visible at the service level and thus the description of
 1629 the service as a whole. The relationship to Action Level Policies and the implications on defining the
 1630 scope of a service are discussed in Section 4.1.2.1.



1631
 1632 *Figure 31 Action-Level and Service-Level Policies*

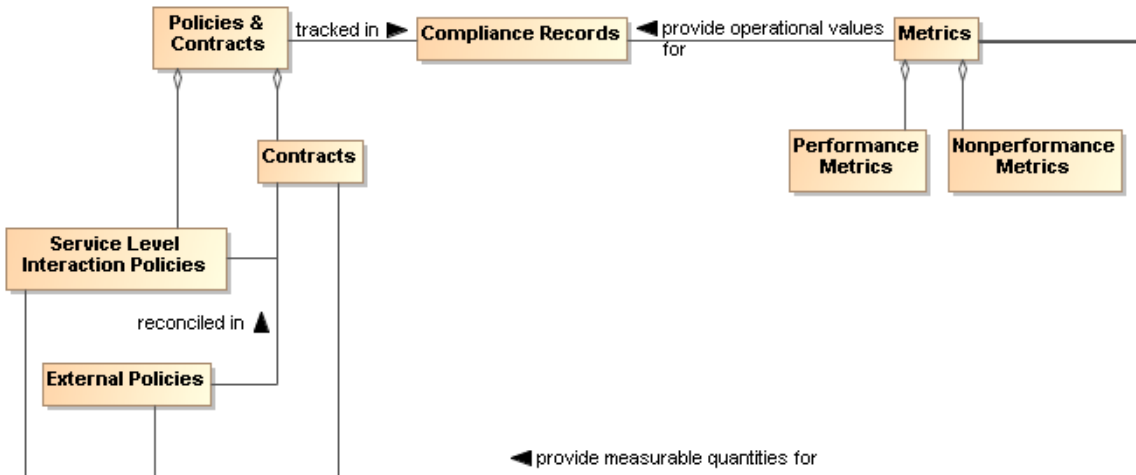
1633 As noted in Figure 30, the policies asserted may affect the allowable Technical Assumptions that can be
 1634 embodied in services or their underlying capabilities and may affect the semantics that can be used. For
 1635 example of the former, there may be a policy that specifies the surge capacity to be accommodated by a
 1636 server, and a service that designs for a smaller capacity would not be appropriate to use. For the latter, a
 1637 policy may require that only services using a community-sponsored vocabulary can be used.

1638 Contracts are agreements among the service participants. The contract may reconcile inconsistent
 1639 policies asserted by the participants or may specify details of the interaction. Service level agreements
 1640 (SLAs) are one commonly used category of contracts.

1641 References to contracts under which the service can be used may also be included in the service
 1642 description. As with policies, the specific representation of the contract, e.g. in some formal contract

1643 language, is likely done outside of the service description and the service description would reference the
1644 normative definition of the contract. Policies and contracts are discussed further in Section 4.4.

1645 The definition and later enforcement of policies and contracts are predicated on the existence of metrics;
1646 the relationships among the relevant concepts are shown in the model in Figure 32. Performance Metrics
1647 identify quantities that characterize the speed and quality of realizing the real world effects produced via
1648 the SOA service; in addition, policies and contracts may depend on nonperformance metrics, such as
1649 whether a license is in place to use the service. Some of these metrics reflect the underlying capability,
1650 e.g. a SOA service cannot respond in two seconds if the underlying capability is expected to take five
1651 seconds to do its processing; some metrics reflect the implementation of the SOA service, e.g. what level
1652 of caching is present to minimize data access requests across the network.



1653
1654 **Figure 32 Policies and Contracts, Metrics, and Compliance Records**

1655 As with many quantities, the metrics associated with a service are not themselves defined by this Service
1656 Description because it is not known *a priori* which metrics are being collected or otherwise checked by the
1657 services, the SOA infrastructure, or other resources that participate in the SOA interactions. However,
1658 the service description SHOULD provide a placeholder (possibly through a link to an externally compiled
1659 list) for identifying which metrics are available and how these can be accessed.

1660 The use of metrics to evaluate compliance is discussed in Section 4.4.2. The results of compliance
1661 evaluation SHOULD be maintained in compliance records and the means to access the compliance
1662 records SHOULD be included in the Policies and Contracts portion of the service description. For
1663 example, the description may be in the form of static information (e.g. over the first year of operation, this
1664 service had a 91% availability), a link to a dynamically generated metric (e.g. over the past 30 days, the
1665 service has had a 93.3% availability), or access to a dynamic means to check the service for current
1666 availability (e.g. a ping). The relationship between service presence and the presence of the individual
1667 actions that can be invoked is discussed under Reachability in Section 4.2.2.3.

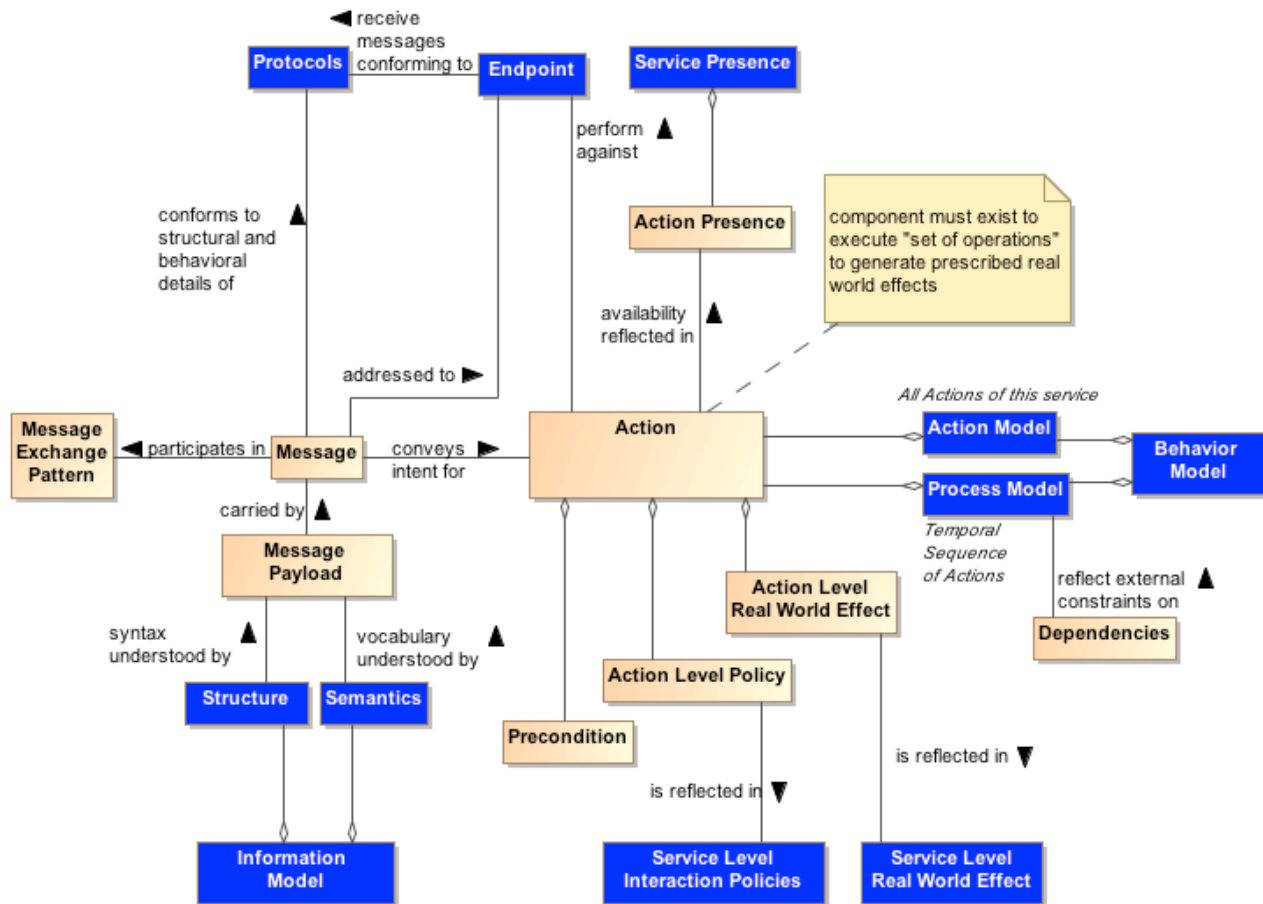
1668 Note, even when policies relate the perspective of a single participant, policy compliance can be
1669 measured and policies may be enforceable without contractual agreement with other participants. This
1670 should be reflected in the policy, contract, and compliance record information maintained in the service
1671 description.

1672 **4.1.2 Use Of Service Description**

1673 **4.1.2.1 Service Description in support of Service Interaction**

1674 If we assume we have awareness, i.e. access to relevant descriptions, the service participants must still
1675 establish willingness and presence to ensure full visibility (See Section 4.2) and to interact with the
1676 service. Service description provides necessary information for many aspects of preparing for and
1677 carrying through with interaction. Recall the fundamental definition of service is a mechanism to access
1678 an underlying capability; the service description describes this mechanism and its use. It lays the

1679 groundwork for what can occur, whereas service interaction defines the specifics through which
 1680 occurrences are realized.



1681
 1682 *Figure 33 Relationship Between Action and Service Description Components*

1683 Figure 33 combines the models in the subsections of Section 4.1.1 to concisely relate Action and the
 1684 relevant components of Service Description. The purpose of Figure 33 is to demonstrate that the
 1685 components of service description go beyond arbitrary documentation and form the critical set of
 1686 information needed to define the what and how of Action. In Figure 33, the leaf nodes from Figure 27 are
 1687 shown in blue.

1688 Action is invoked via a Message where the structure and behavioral details of the message conform to an
 1689 identified Protocol and is directed to the address of the identified endpoint, and the message payload
 1690 conforms to the service Information Model.

1691 The availability of an action is reflected in the Action Presence and each Action Presence contributes to
 1692 the overall Service Presence; this is discussed further in Section 4.2.2.3. Each action has its own
 1693 endpoint and also its own protocols associated with the endpoint¹⁶ and to what extent, e.g. current or
 1694 average availability, there is presence for the action through that endpoint. The endpoint and service
 1695 presence are also part of the service description.

1696 An action may have preconditions where a Precondition is something that needs to be in place before an
 1697 action can occur, e.g. confirmation of a precursor action. Whether preconditions are satisfied is evaluated
 1698 when someone tries to perform the action and not before. Presence for an action means someone can

¹⁶ 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.

1699 initiate it and is independent of whether the preconditions are satisfied. However, the successful
1700 completion of the action may depend on whether its preconditions were satisfied.

1701 Analogous to the relationship between actions and preconditions, the Process Model may imply
1702 Dependencies for succeeding steps in a process, e.g. that a previous step has successfully completed, or
1703 may be isolated to a given step. An example of the latter would be a dependency that the host server has
1704 scheduled maintenance and access attempts at these times would fail. Dependencies related to the
1705 process model do not affect the presence of a service although these may affect whether the business
1706 function successfully completes.

1707 The conditions under which an action can be invoked may depend on policies associated with the action.
1708 The Action Level Policies MUST be reflected in the Service Level Interaction Policies because such
1709 policies may be critical to determining whether the conditions for use of the service are consistent with the
1710 policies asserted by the service consumer. The service level interaction policies are included in the
1711 service description.

1712 Similarly, the result of invoking an action is one or more real world effects, and the Action Level Real
1713 World Effects MUST be reflected in the Service Level Real World Effect included in the service
1714 description. The unambiguous expression of action level policies and real world effects as service
1715 counterparts is necessary to adequately understand what constitutes the service interaction.

1716 An adequate service description MUST provide a consumer with information needed to determine if the
1717 service policies and the (business) functions and service-level real world effects are of interest and there
1718 is nothing in the technical assumptions that preclude use of the service.

1719 Note at this level, the business functions are not concerned with the action or process models. These
1720 models are detailed separately.

1721 The service description is not intended to be isolated documentation but rather an integral part of service
1722 use. Changes in service description SHOULD immediately be made known to consumers and potential
1723 consumers.

1724 **4.1.2.1.1 Description and Invoking Actions Against a Service**

1725 At this point, let us assume the descriptions were sufficient to establish willingness; see Section 4.2.2.2.
1726 Figure 33 indicates the service endpoint establishes where to actually carry out the interaction. This is
1727 where we start considering the action and process models.

1728 The action model identifies the multiple actions a user can perform against a service and the user would
1729 perform these in the context of the process model as specified or referenced under the Service Interface
1730 portion of Service Description. For a given business function, there is a corresponding process model,
1731 where any process model may involve multiple actions. From the above discussion of model elements of
1732 description we may conclude (1) actions have reachability information, including endpoint and presence,
1733 (2) presence of service is some aggregation of presence of its actions, (3) action preconditions and
1734 service dependencies do not affect presence although these may affect successful completion.

1735 Having established visibility, the interaction can proceed. Given a business function, the consumer knows
1736 what will be accomplished (the service functionality), the conditions under which interaction will proceed
1737 (service policies and contracts), and the process that must be followed (the process model). The
1738 remaining question is how does the description information for structure and semantics enable
1739 interaction.

1740 We have established the importance of the process model in identifying relevant actions and their
1741 sequence. Interaction proceeds through messages and thus it is the syntax and semantics of the
1742 messages with which we are here concerned. A common approach is to define the structure and
1743 semantics that can appear as part of a message; then assemble the pieces into messages; and,
1744 associate messages with actions. Actions make use of structure and semantics as defined in the
1745 information model to describe its legal messages.

1746 The process model identifies actions to be performed against a service and the sequence for performing
1747 the actions. For a given action, the Reachability portion of description indicates the protocol bindings that
1748 are available, the endpoint corresponding to a binding, and whether there is presence at that endpoint.
1749 The interaction with actions is through messages that conform to the structure and semantics defined in
1750 the information model and the message sequence conforming to the action's identified MEP. The result

1751 is some portion of the real world effect that will need to be assessed and/or processed (e.g. if an error
1752 exists, that part that covers the error processing would be invoked).

1753 **4.1.2.1.2 The Question of Multiple Business Functions**

1754 Action level effects and policies MUST be reflected at the service level for service description to support
1755 visibility.

1756 It is assumed that a SOA service represents an identifiable business function to which policies can be
1757 applied and from which desired business effects can be obtained. While contemporary discussions of
1758 SOA services and supporting standards do not constrain what actions or combinations of actions can or
1759 should be defined for a service, this Reference Architecture considers the implications of service
1760 description in defining the range of actions appropriate for an individual SOA service.

1761 Consider the situation if a given SOA service is the container for multiple independent (but loosely
1762 related) business functions. These are not multiple effects from a single function but multiple functions
1763 with potentially different sets of effects for each function. A service can have multiple actions a user may
1764 perform against it, and this does not change with multiple business functions. As an individual business
1765 function corresponds to a process model, so multiple business functions imply multiple process models.
1766 The same action may be used in multiple process models but the aggregated service presence would be
1767 specific to each business function because the components being aggregated will likely be different
1768 between process models. In summary, for a service with multiple business functions, each function has
1769 (1) its own process model and dependencies, (2) its own aggregated presence, and (3) possibly its own
1770 list of policies and real world effects.

1771 A common variation on this theme is for a single service to have multiple endpoints for different levels of
1772 quality of service (QoS). Different QoS imply separate statements of policy, separate endpoints, possibly
1773 separate dependencies, and so on. One could say the QoS variation does not require this because there
1774 can be a single QoS policy that encompasses the variations. and all other aspects of the service would be
1775 the same except for the endpoint used for each QoS. However, the different aspects of policy at the
1776 service level would need to be mapped to endpoints, and this introduces an undesirable level of coupling
1777 across the elements of description. In addition, it is obvious that description at the service level can
1778 become very complicated if the number of combinations is allowed to grow.

1779 One could imagine a service description that is basically a container for action descriptions, where each
1780 action description is self contained; however, this would lead to duplication of description components
1781 across actions. If common description components are factored, this either is limited to components
1782 common across all actions or requires complicated tagging to capture the components that often but do
1783 not universally apply.

1784 If a provider cannot describe a service as a whole but must describe every action, this leads to the
1785 situation where it may be extremely difficult to construct a clear and concise service description that can
1786 effectively support discovery and use without tedious logic to process the description and assemble the
1787 available permutations. In effect, if adequate description of an action begins to look like description of a
1788 service, it may be best to have it as a separate service.

1789 Recall, more than one service can access the same underlying capability, and this is appropriate if a
1790 different real world effect is to be exposed. Along these lines, one can argue that different QoS are
1791 different services because getting a response in one minute rather than one hour is more than a QoS
1792 difference; it is a fundamental difference in the business function being provided.

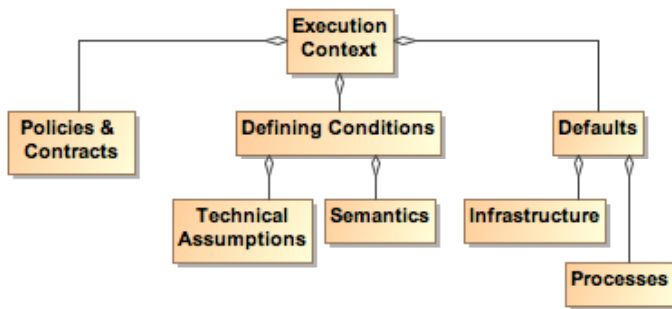
1793 As a best practice, a criteria for whether a service is appropriately scoped may be the ease or difficulty in
1794 creating an unambiguous service description. A consequence of having tightly-scoped services is there
1795 will be a greater reliance on combining services, i.e. more fundamental business functions, to create more
1796 advanced business functions. This is consistent with the principles of service oriented architecture and is
1797 the basic position of the Reference Architecture, although not an absolute requirement. Combining
1798 services increases the reliance on understanding and implementing the concepts of orchestration,
1799 choreography, and other approaches yet to be developed; these are discussed in more detail in section
1800 4.4 Interacting with Services.

1801 **4.1.2.1.3 Service Description, Execution Context, and Service Interaction**

1802 The service description MUST provide sufficient information to support service visibility, including the
1803 willingness of service participants to interact. However, the corresponding descriptions for providers and
1804 consumers may both contain policies, technical assumptions, constraints on semantics, and other
1805 technical and procedural conditions that must be aligned to define the terms of willingness. The
1806 agreements which encapsulate the necessary alignment form the basis upon which interactions may
1807 proceed – in the Reference Model, this collection of agreements and the necessary environmental
1808 support establish the execution context.

1809 To illustrate the concept of the execution context, consider a Web-based system for timecard entry. For
1810 an employee onsite at an employer facility, the execution context requires a computer connected to the
1811 local network and the employee must enter their network ID and password. Relevant policies include that
1812 the employee must maintain the most recent anti-virus software and virus definitions for any computer
1813 connected to the network.

1814 For the same employee connecting from offsite, the execution context specifies the need for a computer
1815 with installed VPN software and a security token to negotiate the VPN connection. The execution context
1816 also includes proxy settings as needed to connect to the offsite network. The employee must still comply
1817 with the requirements for onsite computers and access, but the offsite execution context includes
1818 additional items before the employee can access the same underlying capability and realize the same
1819 real world effects, i.e. the timecard entries.

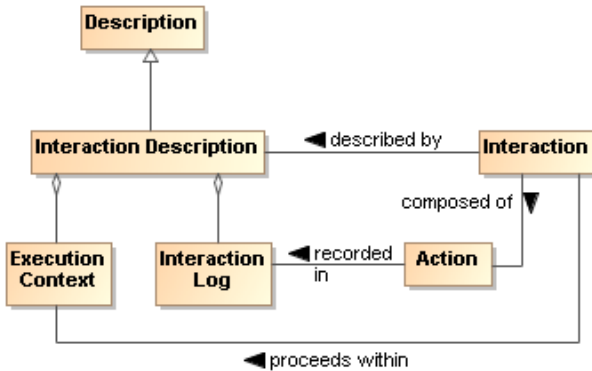


1820
1821 *Figure 34 Execution Context*

1822 Figure 34 shows a few broad categories found in execution context. These are not meant to be
1823 comprehensive. Other items may need to be included to collect a sufficient description of the interaction
1824 conditions. Any other items not explicitly noted in the model but needed to set the environment SHOULD
1825 be included in the execution context.

1826 While the execution context captures the conditions under which interaction can occur, it does not capture
1827 the specific service invocations that do occur in a specific interaction. A service interaction as modeled in
1828 Figure 33 introduces the concept of an Interaction Description which is composed of both the Execution
1829 Context and an Interaction Log. The execution context specifies the set of conditions under which the
1830 interaction occurs and the interaction log captures the sequence of service interactions that occur within
1831 the execution context. This sequence should follow the Process Model but can include details beyond
1832 those specified there. For example, the Process Model may specify an action that results in identifying a
1833 data source, and the identified source is used in a subsequent action. The Interaction Log would record
1834 the specific data source used.

1835 The execution context can be thought of as the container in which the interaction occurs and the
1836 interaction log captures what happens inside the container. This combination is needed to support
1837 auditability and repeatability of the interactions.



1838

1839 *Figure 35 Interaction Description*

1840 SOA allows flexibility to accomplish repeatability or reusability. One benefit of this is that a service can be
 1841 updated without disrupting the user experience of the service. So, Google can improve their ranking
 1842 algorithm without notifying the user about the details of the update.

1843 However, it may also be vital for the consumer to be able to recreate past results or to generate
 1844 consistent results in the future, and information such as what conditions, which services, and which
 1845 versions of those services are used is indispensable in retracing one's path. The interaction log is a
 1846 critical part of the resulting real world effects because it defines how the effects were generated and
 1847 possibly the meaning of observed effects. This increases in importance as dynamic composability
 1848 becomes more feasible. In essence, a result has limited value if one does not know how it was
 1849 generated.

1850 The interaction log SHOULD be a detailed trace for a specific interaction, and its reuse is limited to
 1851 duplicating that interaction. An execution context can act as a template for identical or similar
 1852 interactions. Any given execution context MAY define the conditions of future interactions.

1853 Such uses of execution context imply (1) a standardized format for capturing execution context and (2) a
 1854 subclass of general description could be defined to support visibility of saved execution contexts. The
 1855 specifics of the relevant formats and descriptions are beyond the scope of this Reference Architecture.

1856 A service description is unlikely to track interaction descriptions or the constituent execution contexts or
 1857 interaction logs that include mention of the service. However, as appropriate, linking to specific instances
 1858 of either of these could be done through associated annotations.

1859 **4.1.3 Relationship to Other Description Models**

1860 While the representation shown in Figure 26 is derived from considerations related to service description,
 1861 it is acknowledged that other metadata standards are relevant and should, as possible, be incorporated
 1862 into this work. Two standards of particular relevance are the Dublin Core Metadata Initiative (DCMI) and
 1863 ISO 11179, especially Part 5.

1864 When the service description (or even the general description class) is considered as the DCMI
 1865 "resource", Figure 26 aligns nicely with the DCMI resource model. While some differences exist, these
 1866 are mostly in areas where DCMI goes into detail that is considered beyond the scope of the current
 1867 Reference Architecture. For example, DCMI defines classes of "shared semantics" whereas this
 1868 Reference Architecture considers that an identification of relevant semantic models is sufficient.
 1869 Likewise, the DCMI "description model" goes into the details of possible syntax encodings whereas for
 1870 the Reference Architecture it is sufficient to identify the relevant formats.

1871 With respect to ISO 11179 Part 5, the metadata fields defined in that reference may be used without
 1872 prejudice as the properties in Figure 26. Additionally, other defined metadata sets may be used by the
 1873 service provider if the other sets are considered more appropriate, i.e. it is fundamental to this Reference
 1874 Architecture to identify the need and the means to make vocabulary declarations explicit but it is beyond
 1875 the scope to specify which vocabularies are to be used. In addition, the identification of domain of the
 1876 properties and range of the values has not been included in the current Reference Architecture

1877 discussion, but the text of ISO 11179 Part 5 can be used consistently with the model prescribed in this
1878 document.

1879 Description as defined in the context of this Reference Architecture considers a wide range of applicability
1880 and support of the principles of service oriented architecture. Other metadata models can be used in
1881 concert with the model presented here because most of these focus on a finer level of detail that is
1882 outside the present scope, and so provide a level of implementation guidance that can be applied as
1883 appropriate.

1884 **4.1.4 Architectural Implications**

1885 The description of service description indicates numerous architectural implications on the SOA
1886 ecosystem:

- 1887 • Description will change over time and its contents will reflect changing needs and context. This
1888 requires the existence of:
 - 1889 ○ mechanisms to support the storage, referencing, and access to normative definitions of
1890 one or more versioning schemes that may be applied to identify different aggregations of
1891 descriptive information, where the different schemes may be versions of a versioning
1892 scheme itself;
 - 1893 ○ configuration management mechanisms to capture the contents of the each aggregation
1894 and apply a unique identifier in a manner consistent with an identified versioning scheme;
 - 1895 ○ one or more mechanisms to support the storage, referencing, and access to conversion
1896 relationships between versioning schemes, and the mechanisms to carry out such
1897 conversions.
- 1898 • Description makes use of defined semantics, where the semantics may be used for
1899 categorization or providing other property and value information for description classes. This
1900 requires the existence of:
 - 1901 ○ semantic models that provide normative descriptions of the utilized terms, where the
1902 models may range from a simple dictionary of terms to an ontology showing complex
1903 relationships and capable of supporting enhanced reasoning;
 - 1904 ○ mechanisms to support the storage, referencing, and access to these semantic models;
 - 1905 ○ configuration management mechanisms to capture the normative description of each
1906 semantic model and to apply a unique identifier in a manner consistent with an identified
1907 versioning scheme;
 - 1908 ○ one or more mechanisms to support the storage, referencing, and access to conversion
1909 relationships between semantic models, and the mechanisms to carry out such
1910 conversions.
- 1911 • Descriptions include reference to policies defining conditions of use and optionally contracts
1912 representing agreement on policies and other conditions. This requires the existence of (as also
1913 enumerated under governance):
 - 1914 ○ descriptions to enable the policy modules to be visible, where the description includes a
1915 unique identifier for the policy and a sufficient, and preferably a machine processible,
1916 representation of the meaning of terms used to describe the policy, its functions, and its
1917 effects;
 - 1918 ○ one or more discovery mechanisms that enable searching for policies that best meet the
1919 search criteria specified by the service participant; where the discovery mechanism will
1920 have access to the individual policy descriptions, possibly through some repository
1921 mechanism;
 - 1922 ○ accessible storage of policies and policy descriptions, so service participants can access,
1923 examine, and use the policies as defined.
- 1924 • Descriptions include references to metrics which describe the operational characteristics of the
1925 subjects being described. This requires the existence of (as partially enumerated under
1926 governance):
 - 1927 ○ the infrastructure monitoring and reporting information on SOA resources;
 - 1928 ○ possible interface requirements to make accessible metrics information generated or
1929 most easily accessed by the service itself;

- 1930 ○ mechanisms to catalog and enable discovery of which metrics are available for a
- 1931 described resources and information on how these metrics can be accessed;
- 1932 ○ mechanisms to catalog and enable discovery of compliance records associated with
- 1933 policies and contracts that are based on these metrics.
- 1934 • Descriptions of the interactions are important for enabling auditability and repeatability, thereby
- 1935 establishing a context for results and support for understanding observed change in performance
- 1936 or results. This requires the existence of:
- 1937 ○ one or more mechanisms to capture, describe, store, discover, and retrieve interaction
- 1938 logs, execution contexts, and the combined interaction descriptions;
- 1939 ○ one or more mechanisms for attaching to any results the means to identify and retrieve
- 1940 the interaction description under which the results were generated.
- 1941 • Descriptions may capture very focused information subsets or can be an aggregate of numerous
- 1942 component descriptions. Service description is an example of a likely aggregate for which
- 1943 manual maintenance of all aspects would not be feasible. This requires the existence of:
- 1944 ○ tools to facilitate identifying description elements that are to be aggregated to assemble
- 1945 the composite description;
- 1946 ○ tools to facilitate identifying the sources of information to associate with the description
- 1947 elements;
- 1948 ○ tools to collect the identified description elements and their associated sources into a
- 1949 standard, referenceable format that can support general access and understanding;
- 1950 ○ tools to automatically update the composite description as the component sources
- 1951 change, and to consistently apply versioning schemes to identify the new description
- 1952 contents and the type and significance of change that occurred.
- 1953 • Descriptions provide up-to-date information on what a resource is, the conditions for interacting
- 1954 with the resource, and the results of such interactions. As such, the description is the source of
- 1955 vital information in establishing willingness to interact with a resource, reachability to make
- 1956 interaction possible, and compliance with relevant conditions of use. This requires the existence
- 1957 of:
- 1958 ○ one or more discovery mechanisms that enable searching for described resources that
- 1959 best meet the criteria specified by a service participant, where the discovery mechanism
- 1960 will have access to individual descriptions, possibly through some repository mechanism;
- 1961 ○ tools to appropriately track users of the descriptions and notify them when a new version
- 1962 of the description is available.

1963 4.2 Service Visibility Model

1964 One of the key requirements for participants interacting with each other in the context of a SOA is

1965 achieving visibility: before services can interoperate, the participants have to be visible to each other

1966 using whatever means are appropriate. The Reference Model analyzes visibility in terms of awareness,

1967 willingness, and reachability. In this section, we explore how visibility may be achieved.

1968 4.2.1 Visibility to Business

1969 The relationship of visibility to the SOA ecosystem encompasses both human social structures and

1970 automated IT mechanisms. Figure 36 depicts a business setting that is a basis for visibility as related to

1971 the Social Structure Model in the Service Ecosystem View (see Section 3.2). Service consumers and

1972 service providers may have direct awareness or mediated awareness where mediated awareness is

1973 achieved through some third party. A consumer's willingness to use a service is reflected by the

1974 consumer's presumption of satisfying goals and needs based on the description of the service. Service

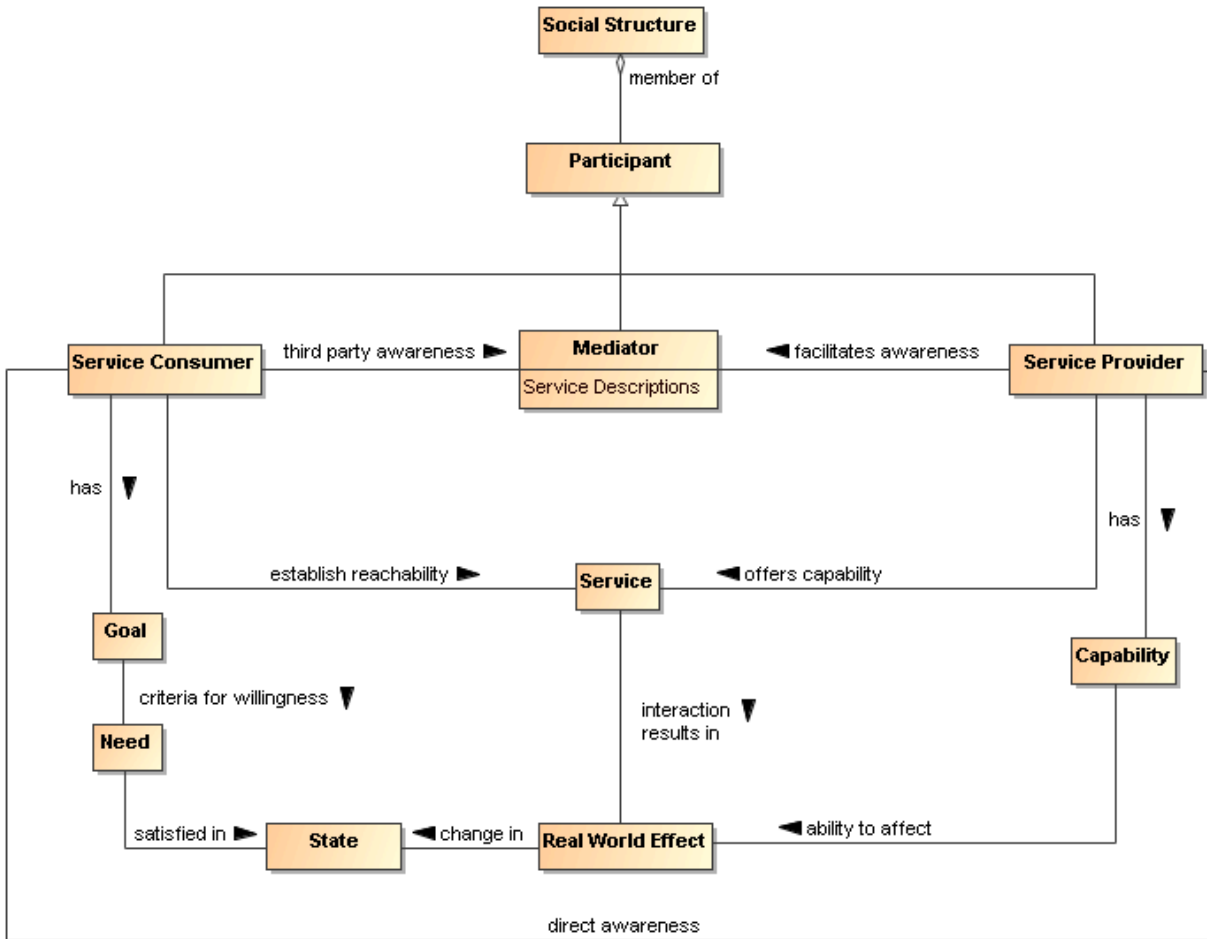
1975 providers offer capabilities that have real world effects that result in a change in state of the consumer.

1976 Reachability of the service by the consumer leads to interactions that change the state of the consumer.

1977 The consumer can measure the change of state to determine if the claims made by description and the

1978 real world effects of consuming the service meet the consumer's needs.

1979



1980

1981 *Figure 36 Visibility to Business*

1982 Visibility and interoperability in a SOA ecosystem requires more than location and interface information.
 1983 A meta-model for this broader view of visibility is depicted in Section 4.1. In addition to providing
 1984 improved awareness of service capabilities through description of information such as reachability,
 1985 behavior models, information models, functionality, and metrics, the service description may contain
 1986 policies valuable for determination of willingness to interact.

1987 A mediator of service descriptions may provide event notifications to both consumers and providers about
 1988 information relating to service descriptions. One example of this capability is a publish/subscribe model
 1989 where the mediator allows consumers to subscribe to service description version changes made by the
 1990 provider. Likewise, the mediator may provide notifications to the provider of consumers that have
 1991 subscribed to service description updates.

1992 Another important business capability in a SOA environment is the ability to narrow visibility to trusted
 1993 members within a social structure. Mediators for awareness may provide policy based access to service
 1994 descriptions allowing for the dynamic formation of awareness between trusted members.

1995 **4.2.2 Visibility**

1996 Attaining visibility is described in terms of steps that lead to visibility. While there can be many contexts
 1997 for visibility within a single social structure, the same general steps can be applied to each of the contexts
 1998 to accomplish visibility.

1999 Attaining SOA visibility requires

- 2000 • service description creation and maintenance,
 2001 • processes and mechanisms for achieving awareness of and accessing descriptions,

- 2002 • processes and mechanisms for establishing willingness of participants,
- 2003 • processes and mechanisms to determine reachability.

2004 Visibility may occur in stages, i.e. a participant can become aware enough to look or ask for further
2005 description, and with this description, the participant can decide on willingness, possibly requiring
2006 additional description. For example, if a potential consumer has a need for a tree cutting (business)
2007 service, the consumer can use a web search engine to find web sites of providers. The web search
2008 engine (a mediator) gives the consumer links to relevant web pages and the consumer can access those
2009 descriptions. For those prospective providers that satisfy the consumer's criteria, the consumer's
2010 willingness to interact increases. The consumer likely contacts several tree services to get detailed cost
2011 information (or arrange for an estimate) and may ask for references (further description). Likely, the
2012 consumer will establish full visibility and proceed with the interaction with a tree service who mutually
2013 establishes visibility.

2014 **4.2.2.1 Awareness**

2015 A service participant is aware of another participant if it has access to a description of that participant with
2016 sufficient completeness to establish the other requirements of visibility.

2017 Awareness is inherently a function of a participant; awareness can be established without any action on
2018 the part of the target participant other than the target providing appropriate descriptions. Awareness is
2019 often discussed in terms of consumer awareness of providers but the concepts are equally valid for
2020 provider awareness of consumers.

2021 Awareness can be decomposed into the creation of descriptions, making them available, and discovering
2022 the descriptions. Discovery can be initiated or it can be by notification. Initiated discovery for business
2023 may require formalization of the required capabilities and resources to achieve business goals.

2024 Achieving awareness in a SOA can range from word of mouth to formal service descriptions in a
2025 standards-based registry-repository. Some other examples of achieving awareness in a SOA are the
2026 use of a web page containing description information, email notifications of descriptions, and document
2027 based descriptions.

2028 A mediator as discussed for awareness is a third party participant that provides awareness to one or
2029 more consumers of one or more services. Direct awareness is awareness between a consumer and
2030 provider without the use of a third party.

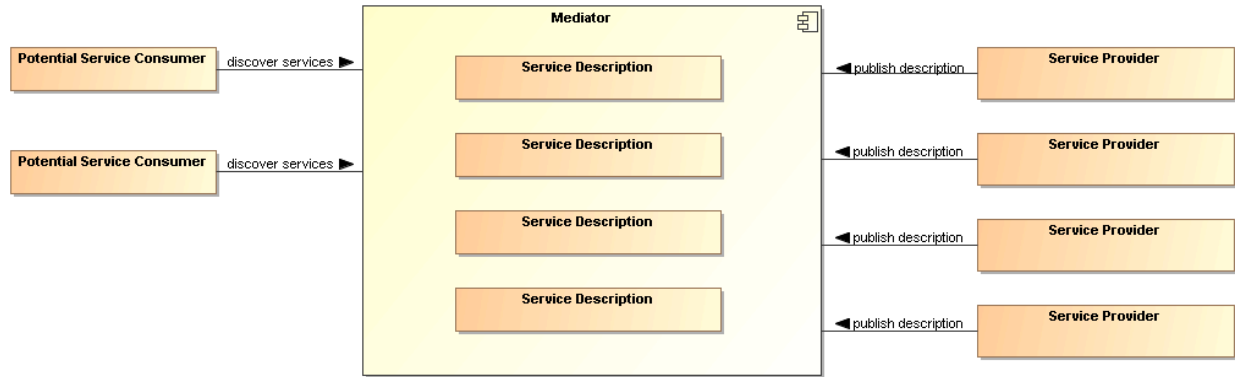
2031 Direct awareness may be the result of having previously established an execution context, or direct
2032 awareness may include determining the presence of services and then querying the service directly for
2033 description. As an example, a priori visibility of some sensor device may provide the means for interaction
2034 or a query for standardized sensor device metadata may be broadcast to multiple locations. If
2035 acknowledged, the service interface for the device may directly provide description to a consumer so the
2036 consumer can determine willingness to interact.

2037 The same medium for awareness may be direct in one context and may be mediated in another context.
2038 For example, a service provider may maintain a web site with links to the provider's descriptions of
2039 services giving the consumers direct awareness to the provider's services. Alternatively, a community
2040 may maintain a mediated web site with links to various provider descriptions of services for any number of
2041 consumers. More than one mediator may be involved, as different mediators may specialize in different
2042 mediation functions.

2043 Descriptions may be formal or informal. Section 4.1, provides a comprehensive model for service
2044 description that can be applied to formal registry/repositories used to mediate visibility. Using consistent
2045 description taxonomies and standards based mediated awareness helps provide more effective
2046 awareness.

2047 **4.2.2.1.1 Mediated Awareness**

2048 Mediated awareness promotes loose coupling by keeping the consumers and services from explicitly
2049 referring to each other and the descriptions. Mediation lets interaction vary independently. Rather than all
2050 potential service consumers being informed on a continual basis about all services, there is a known or
2051 agreed upon facility or location that houses the service description.



2052

2053 *Figure 37 Mediated Service Awareness*

2054 In Figure 37, the potential service consumers perform queries or are notified in order to locate those
 2055 services that satisfy their needs. As an example, the telephone book is a mediated registry where
 2056 individuals perform manual searches to locate services (i.e. the yellow pages). The telephone book is
 2057 also a mediated registry for solicitors to find and notify potential customers (i.e. the white pages).

2058 In mediated service awareness for large and dynamic numbers of service consumers and service
 2059 providers, the benefits typically far outweigh the management issues associated with it. Some of the
 2060 benefits of mediated service awareness are

- 2061 • Potential service consumers have a known location for searching thereby eliminating needless and
 2062 random searches
- 2063 • Typically a consortium of interested parties (or a sufficiently large corporation) signs up to host the
 2064 mediation facility
- 2065 • Standardized tools and methods can be developed and promulgated to promote interoperability and
 2066 ease of use.

2067 However, mediated awareness can have some risks associated with it:

- 2068 • A single point of failure. If the central mediation service fails then a potentially large number of service
 2069 providers and consumers will be adversely affected.
- 2070 • A single point of control. If the central mediation service is owned by, or controlled by, someone other
 2071 than the service consumers and/or providers then the latter may be put at a competitive disadvantage
 2072 based on policies of the discovery provider.

2073 A common mechanism for mediated awareness is a registry-repository. The registry stores links or
 2074 pointers to service description artifacts. The repository in this example is the storage location for the
 2075 service description artifacts. Service descriptions can be pushed (publish/subscribe for example) or pulled
 2076 from the register-repository mediator.

2077 The registry is like a card catalog at the library and a repository is like the shelves for the books.
 2078 Standardized metadata describing repository content can be stored as registry objects in a registry and
 2079 any type of content can be stored as repository items in a repository. The registry may be constructed
 2080 such that description items stored within the mediation facility repository will have intrinsic links in the
 2081 registry while description items stored outside the mediation facility will have extrinsic links in the registry.

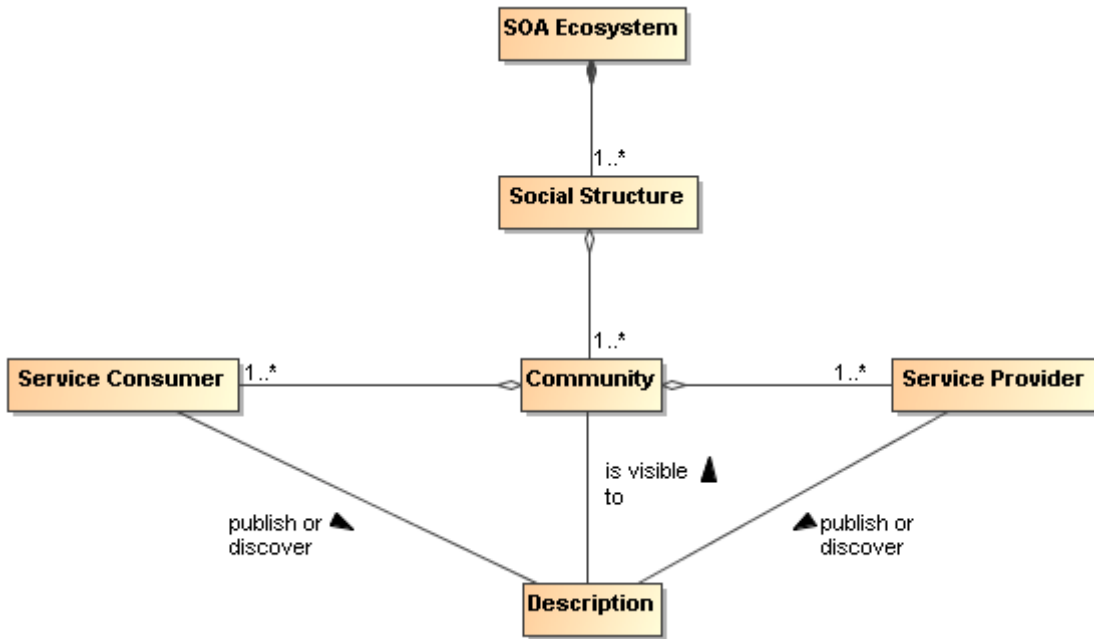
2082 When independent but like SOA IT mechanisms interoperate with one another, the IT mechanisms may
 2083 be referred to as federated.

2084 **4.2.2.1.2 Awareness in Complex Social Structures**

2085 Awareness applies to one or more communities within one or more social structures where a community
 2086 consists of at least one description provider and one description consumer. These communities may be
 2087 part of the same social structure or be part of different ones.

2088 In Figure 38, awareness can be within a single community, multiple communities, or all communities in
 2089 the social structure. The social structure can encourage or restrict awareness through its policies, and
 2090 these policies can affect participant willingness. The information about policies should be incorporated in

2091 the relevant descriptions. The social structure also governs the conditions for establishing contracts, the
2092 results of which will be reflected in the execution context if interaction is to proceed.



2093
2094 *Figure 38 Awareness In a SOA Ecosystem*

2095 IT policy/contract mechanisms can be used by visibility mechanisms to provide awareness between
2096 communities. The IT mechanisms for awareness may incorporate trust mechanisms to assure
2097 awareness between trusted communities. For example, government organizations will often want to limit
2098 awareness of an organization's services to specific communities of interest.

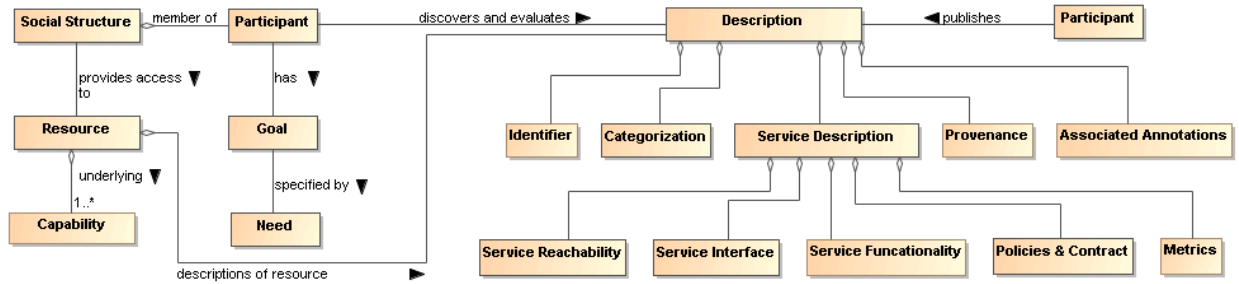
2099 Another common business model for awareness is maximizing awareness to communities within the
2100 social structure, the traditional market place business model. A centralized mediator often arises as a
2101 provider for this global visibility, a gatekeeper of visibility so to speak. For example, Google is a
2102 centralized mediator for accessing information on the web. As another example, television networks have
2103 centralized entities providing a level of awareness to communities that otherwise could not be achieved
2104 without going through the television network.

2105 However, mediators have motivations, and they may be selective in which information they choose to
2106 make available to potential consumers. For example, in a secure environment, the mediator may enforce
2107 security policies and make information selectively available depending on the security clearance of the
2108 consumers.

2109 **4.2.2.2 Willingness**

2110 Having achieved awareness, participants use descriptions to help determine their willingness to interact
2111 with another participant. Both awareness and willingness are determined prior to consumer/provider
2112 interaction.

2113
2114



2115
2116 *Figure 39 Business, Description and Willingness*

2117 Figure 39 relates elements of the Service Ecosystem View, and elements from the Service Description
2118 Model to willingness. By having a willingness to interact within a particular social structure, the social
2119 structure provides the participant access to capabilities based on conditions the social structure finds
2120 appropriate for its context. The participant can use these capabilities to satisfy goals and objectives as
2121 specified by the participant's needs.

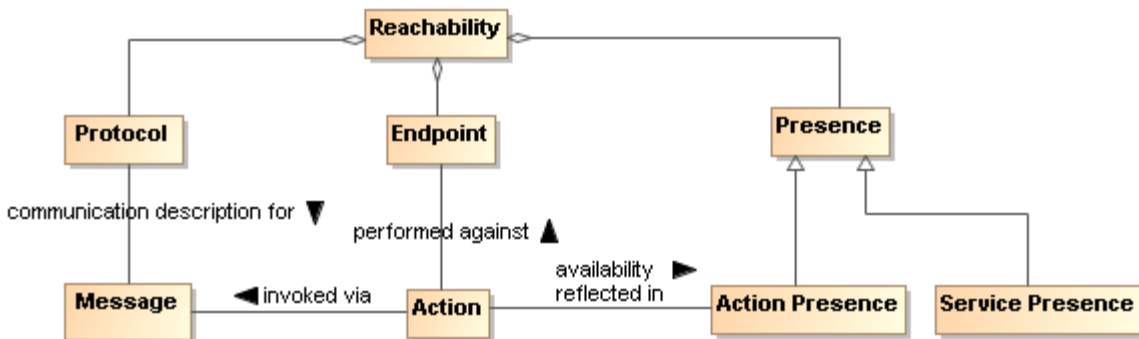
2122 In Figure 39, information used to determine willingness is defined by Description. Information referenced
2123 by Description may come from many sources. For example, a mediator for descriptions may provide 3rd
2124 party annotations for reputation. Another source for reputation may be a participant's own history of
2125 interactions with another participant.

2126 A participant will inspect functionality for potential satisfaction of needs. Identity is associated with any
2127 participant, however, identity may or may not be verified. If available, participant reputation may be a
2128 deciding factor for willingness to interact. Policies and contracts referenced by the description may be
2129 particularly important to determine the agreements and commitments required for business interactions.
2130 Provenance may be used for verification of authenticity of a resource.

2131 Mechanisms that aid in determining willingness will likely make use of the artifacts referenced by
2132 descriptions of services. Mechanisms for establishing willingness could be as simple as rendering
2133 service description information for human consumption to automated evaluation of functionality, policies,
2134 and contracts by a rules engine. The rules engine for determining willingness could operate as a policy
2135 decision procedure as defined in Section 4.4.

2136 **4.2.2.3 Reachability**

2137 Reachability involves knowing the endpoint, protocol, and presence of a service. At a minimum,
2138 reachability requires information about the location of the service and the protocol describing the means
2139 of communication.



2140
2141 *Figure 40 Service Reachability*

2142

2143 **Endpoint**

2144 An endpoint is a reference-able entity, processor or resource against which an action can be
2145 performed.

2146 **Protocol**

2147 A protocol is a structured means by which service interaction is regulated.
2148

2149 **Presence**

2150 Presence is the measurement of reachability of a service at a particular point in time.

2151 A protocol defines a structured method of communication with a service. Presence is determined by
2152 interaction through a communication protocol. Presence may not be known in many cases until the act of
2153 interaction begins. To overcome this problem, IT mechanisms may make use of presence protocols to
2154 provide the current up/down status of a service.

2155 Service reachability enables service participants to locate and interact with one another. Each action may
2156 have its own endpoint and also its own protocols associated with the endpoint and whether there is
2157 presence for the action through that endpoint. Presence of a service is an aggregation of the presence of
2158 the service's actions, and the service level may aggregate to some degraded or restricted presence if
2159 some action presence is not confirmed. For example, if error processing actions are not available, the
2160 service can still provide required functionality if no error processing is needed. This implies reachability
2161 relates to each action as well as applying to the service/business as a whole.
2162

2163 **4.2.3 Architectural Implications**

2164 Visibility in a SOA ecosystem has the following architectural implications on mechanisms providing
2165 support for awareness, willingness, and reachability:

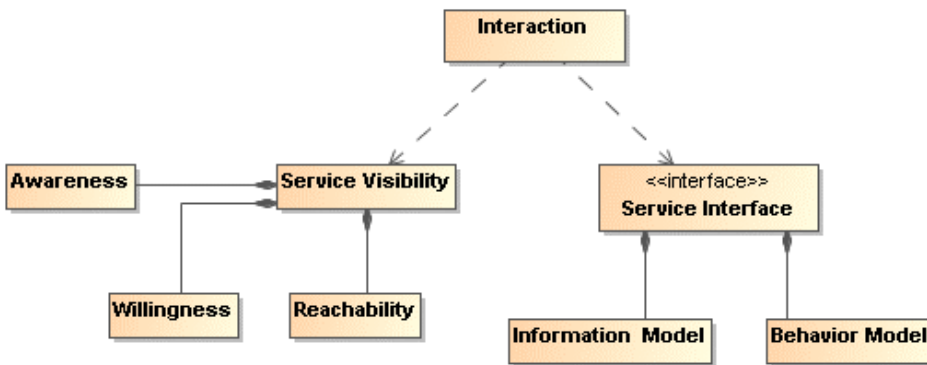
- 2166 • Mechanisms providing support for awareness will likely have the following minimum capabilities:
 - 2167 ○ creation of Description, preferably conforming to a standard Description format and structure;
 - 2168 ○ publishing of Description directly to a consumer or through a third party mediator;
 - 2169 ○ discovery of Description, preferably conforming to a standard for Description discovery;
 - 2170 ○ notification of Description updates or notification of the addition of new and relevant
2171 Descriptions;
 - 2172 ○ classification of Description elements according to standardized classification schemes.
- 2173 • In a SOA ecosystem with complex social structures, awareness may be provided for specific
2174 communities of interest. The architectural mechanisms for providing awareness to communities of
2175 interest will require support for:
 - 2176 ○ policies that allow dynamic formation of communities of interest;
 - 2177 ○ trust that awareness can be provided for and only for specific communities of interest, the
2178 bases of which is typically built on keying and encryption technology.
- 2179 • The architectural mechanisms for determining willingness to interact will require support for:
 - 2180 ○ verification of identity and credentials of the provider and/or consumer;
 - 2181 ○ access to and understanding of description;
 - 2182 ○ inspection of functionality and capabilities;
 - 2183 ○ inspection of policies and/or contracts.
- 2184 • The architectural mechanisms for establishing reachability will require support for:
 - 2185 ○ the location or address of an endpoint;
 - 2186 ○ verification and use of a service interface by means of a communication protocol;
 - 2187 ○ determination of presence with an endpoint which may only be determined at the point of
2188 interaction but may be further aided by the use of a presence protocol for which the endpoints
2189 actively participate.

2190 **4.3 Interacting with Services Model**

2191 Interaction is the activity involved in using a service to access capability in order to achieve a particular
 2192 desired real world effect, where real world effect is the actual *result* of using a service. An interaction can
 2193 be characterized by a sequence of actions. Consequently, interacting with a service, i.e. performing
 2194 actions against the service—usually mediated by a series of message exchanges—involves actions
 2195 performed by the service. Different modes of interaction are possible such as modifying the shared state
 2196 of a resource. Note that a participant (or agent acting on behalf of the participant) can be the sender of a
 2197 message, the receiver of a message, or both.

2198 **4.3.1 Interaction Dependencies**

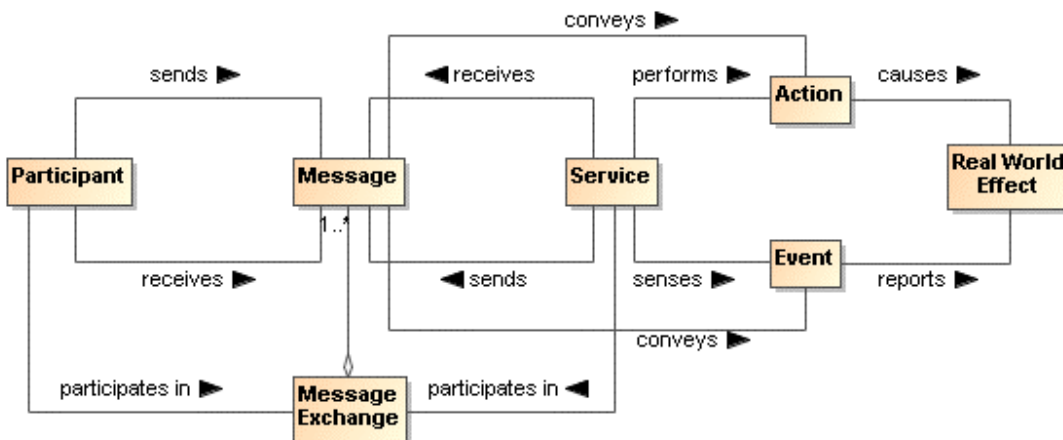
2199 Recall from the Reference Model that service visibility is the capacity for those with needs and those with
 2200 capabilities to be able to interact with each other, and that the service interface is the means by which the
 2201 underlying capabilities of a service are accessed. Ideally, the details of the underlying service
 2202 implementation are abstracted away by the service interface. [Service] interaction therefore has a direct
 2203 dependency on the visibility of the service as well as its implementation-neutral interface (see Figure x).
 2204 Service visibility is composed of awareness, willingness, and reachability and service interface is
 2205 composed of the information and behavior models. Service visibility is modeled in Section 4.2 while
 2206 service interface is modeled in Section x.



2207
 2208 *Figure 41 Interaction dependencies.*

2209 **4.3.2 Actions and Events**

2210 For purposes of this Reference Architecture, the authors have committed to the use of message
 2211 exchange between service participants to denote actions performed against and by the service, and to
 2212 denote events that report on real world effects that are caused by the service actions. A visual model of
 2213 the relationship between these concepts is shown in Figure 42.



2215 *Figure 42 A "message" conveys either an action or an event.*

2216 A *message* conveys either an action or an event. In other words, both actions and events, realized by the
2217 SOA services, are denoted by the messages. The Reference Model states that the action model
2218 characterizes the "permissible set of actions that may be invoked against a service." We extend that
2219 notion here to include events as part of the event model and that messages denote either actions or
2220 events.

2221 In Section 3.3, we saw that participants interact with each other in order to perform actions. An action is
2222 not itself the same thing as the result of performing the action. When an action is performed against a
2223 service, the real world effect that results is reported in the form of events

2224 **4.3.3 Message Exchange**

2225 *Message exchange* is the means by which service participants (or their agents) interact with each other.
2226 There are two primary modes of interaction: joint actions that cause real world effects, and notification of
2227 events that report real world effects.¹⁷

2228 A message exchange is used to affect an action when the messages contain the appropriately formatted
2229 content that should be interpreted as joint action and the agents involved interpret the message
2230 appropriately.

2231 A message exchange is also used to communicate event notifications. An event is a report of an
2232 occurrence that is of interest to some participant; in our case when some real world effect has occurred.
2233 Just as action messages will have formatting requirements, so will event notification messages. In this
2234 way, the Information Model of a service must specify the syntax (structure), and semantics (meaning) of
2235 the action messages and event notification messages as part of a service interface. It must also specify
2236 the syntax and semantics of any data that is carried as part of a payload of the action or event notification
2237 message. The Information Model is described in greater detail in the Service Description Model (see
2238 Section 4.1).

2239 In addition to the Information Model that describes the syntax and semantics of the messages and data
2240 payloads, exception conditions and error handling in the event of faults (e.g., network outages, improper
2241 message formats, etc.) must be specified or referenced as part of the Service Description.

2242 When a message is interpreted as an action, the correct interpretation typically requires the receiver to
2243 perform a set of operations. These *operations* represent the sequence of actions (often private) a service
2244 must perform in order to validly participate in a given joint action.

2245 Similarly, the correct consequence of realizing a real world effect may be to initiate the reporting of that
2246 real world effect via an event notification.

2247 **Message Exchange**

2248 The means by which joint actions and event notifications are coordinated by service participants
2249 (or agents).

2250 **Operations**

2251 The sequence of actions a service must perform in order to validly participate in a given joint
2252 action.

2253 **4.3.3.1 Message Exchange Patterns (MEPs)**

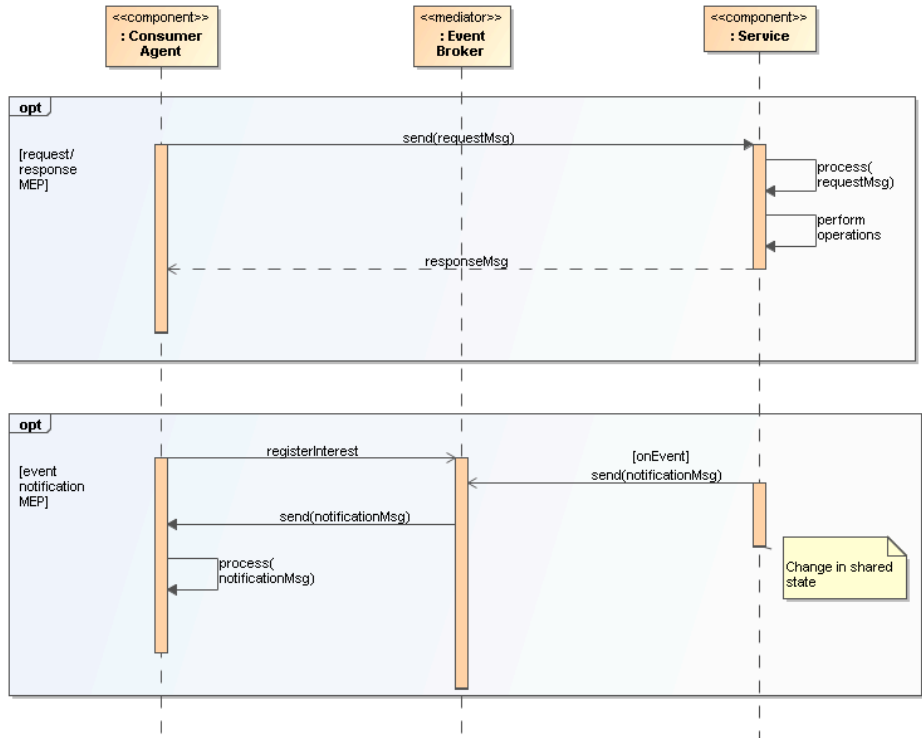
2254 As stated earlier, this Reference Architecture commits to the use of message exchange to denote actions
2255 against the services, and to denote events that report on real world effects that arise from those actions.

2256 Based on these assumptions, the basic temporal aspect of service interaction can be characterized by
2257 two fundamental message exchange patterns (MEPs):

- 2258 • Request/response to represent how actions cause a real world effect

¹⁷ The notion of "joint" in joint action implies that you have to have a speaker *and* a listener in order to interact.

2259 • Event notification to represent how events report a real world effect
 2260 This is by no means a complete list of all possible MEPs used for inter- or intra-enterprise messaging but
 2261 it does represent those that are most commonly used in exchange of information and reporting changes
 2262 in state both within organizations and across organizational boundaries, a hallmark of a SOA.
 2263 Recall from the OASIS Reference Model that the Process Model characterizes “the temporal relationships
 2264 between and temporal properties of actions and events associated with interacting with the service.”
 2265 Thus, MEPs are a key element of the Process Model. The meta-level aspects of the Process Model (just
 2266 as with the Action Model) are provided as part of the Service Description Model (see Section 4.1).



2267
 2268 *Figure 43 Fundamental SOA message exchange patterns (MEPs)*

2269 In the UML sequence diagram shown in Figure 43 it is assumed that the service participants (consumer
 2270 and provider) have delegated message handling to hardware or software agents acting on their behalf. In
 2271 the case of the service consumer, this is represented by the *Consumer Agent* component. In the case of
 2272 the service provider, the agent is represented by the *Service* component. The message interchange
 2273 model illustrated represents a logical view of the MEPs and not a physical view. In other words, specific
 2274 hosts, network protocols, and underlying messaging system are not shown as these tend to be
 2275 implementation specific. Although such implementation-specific elements are considered outside the
 2276 scope of this Reference Architecture, they are important considerations in modeling the SOA execution
 2277 context. Recall from the Reference Model that the *execution context* of a service interaction is “the set of
 2278 infrastructure elements, process entities, policy assertions and agreements that are identified as part of
 2279 an instantiated service interaction, and thus forms a path between those with needs and those with
 2280 capabilities.”

2281 **4.3.3.2 Request/Response MEP**

2282 In a request/response MEP, the Consumer Agent component sends a request message to the Service
 2283 component. The Service component then processes the request message. Based on the content of the
 2284 message, the Service component performs the service operations. Following the completion of these
 2285 operations, a response message is returned to the Consumer Agent component. The response could be

2286 that a step in a process is complete, the initiation of a follow-on operation, or the return of requested
2287 information.¹⁸

2288 Although the sequence diagram shows a *synchronous* interaction (because the sender of the request
2289 message, i.e., Consumer Agent, is blocked from continued processing until a response is returned from
2290 the Service) other variations of request/response are valid, including *asynchronous* (non-blocking)
2291 interaction through use of queues, channels, or other messaging techniques.

2292 What is important to convey here is that the request/response MEP represents *action*, which causes a
2293 real world effect, irrespective of the underlying messaging techniques and messaging infrastructure used
2294 to implement the request/response MEP.

2295 **4.3.3.3 Event Notification MEP**

2296 An event is realized by means of an event notification message exchange that reports a real world effect;
2297 specifically, a change in shared state between service participants. The basic event notification MEP
2298 takes the form of a one-way message sent by a notifier agent (in this case, the Service component) and
2299 received by agents with an interest in the event (here, the Consumer Agent component).

2300 Often the sending agent may not be fully aware of all the agents that will receive the notification;
2301 particularly in so-called publish/subscribe (“pub/sub”) situations. In event notification message
2302 exchanges, it is rare to have a tightly-coupled link between the sending and the receiving agent(s) for a
2303 number of practical reasons. One of the most common is the potential for network outages or
2304 communication interrupts that can result in loss of notification of events. Therefore, a third-party agent or
2305 service is often used to decouple the sending and receiving agents .

2306 Although this is typically an implementation issue, because this type of third-party decoupling is so
2307 common in event-driven systems, we felt that for this Reference Architecture, it was warranted for use in
2308 modeling this type of message exchange. This third-party intermediary is shown in Figure 43 as an Event
2309 Broker mediator. As with the request/response MEP, no distinction is made between synchronous versus
2310 asynchronous communication, although asynchronous message exchange is illustrated in Figure 43 .

2311 **4.3.4 Composition of Services**

2312 Composition of services is the act of aggregating or “composing” a single service from one or more other
2313 services. Before we provide an architectural model of service composition, it is important that we
2314 distinguish two fundamentally different types of services, *atomic services* and *composite services*.

2315 **Atomic Service**

2316 A service visible to a service consumer (or agent) via a single interface and described via a single
2317 service description that does not use or interact with other services.

2318 **Composite Service**

2319 A service visible to a service consumer (or agent) via a single interface and described via a single
2320 service description that is the aggregation or composition of one or more other services. These
2321 other services can be atomic services, other composite services, or a combination of both.¹⁹

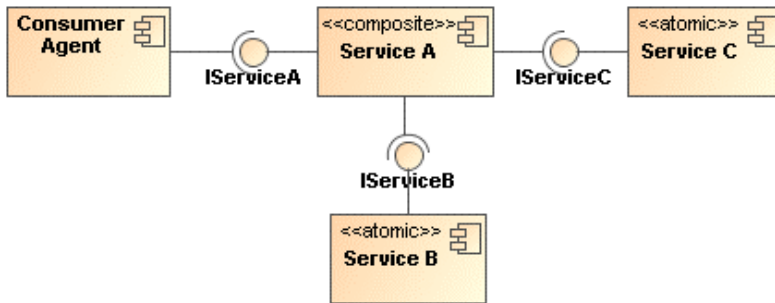
2322 From the consumer’s point of view, the distinction is, of course, mostly irrelevant. The consumer still
2323 interacts with a composite service via a single interface and utilizes the meta-level information about the
2324 composite service provided by a single Service Description. Nevertheless, there are important

¹⁸ 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.

¹⁹ The term *composition* as used herein does not embrace the semantics of a UML composition binary relationship. Here we are referring to the relationship between services.

2325 dependencies that need to be considered in services that utilize other services such as propagation of
2326 policy constraints, security profiles, etc.

2327 A simple model of service composition is illustrated in Figure 44.



2328

2329 *Figure 44 Simple model of service composition ("public" composition).*

2330 Here, Service A is a composite service that has an exposed interface IServiceA that is available to the
2331 Consumer Agent component and relies on two other service components in its implementation. The
2332 Consumer Agent does not know that atomic Services B and C are used by Service A, or whether they are
2333 used in serial or parallel, or if their operations succeed or fail. The Consumer Agent only cares about the
2334 success or failure of Service A. The exposed interfaces of Services B and C (IService B and IServiceC)
2335 are not necessarily hidden from the Consumer Agent; only the fact that these services are used as part of
2336 the composition of Service A. In this example, there is no practical reason the Consumer Agent could not
2337 interact with Service B or Service C in some other interaction scenario.

2338 It is possible for a service composition to be opaque from one perspective and transparent from another.
2339 For example, a service may appear to be a single service from the Consumer Agent's perspective, but is
2340 transparently composed of one or more services from a service management perspective. A Service
2341 Management Service needs to be able to have visibility into the composition in order to properly manage
2342 the dependencies between the services used in constructing the composite service—including managing
2343 the service's lifecycle. The subject of services as management entities is described and modeled in the
2344 Owing Service Oriented Architectures View of this Reference Architecture and will not be further
2345 elaborated here. The point to be made here is that there can be different levels of opaqueness or
2346 transparency when it comes to visibility of service composition.

2347 Services can be composed in a variety of ways including direct service-to-service interaction by using
2348 programming techniques, or they can be aggregated by means of a scripting approach that leverages a
2349 service composition scripting language. Such scripting approaches are further elaborated in the following
2350 sub-sections on service-oriented business processes and collaborations.

2351 4.3.4.1 Service-Oriented Business Processes

2352 The concepts of business processes and collaborations in the context of transactions and exchanges
2353 across organizational boundaries are described and modeled as part of the Service Ecosystem View of
2354 this Reference Architecture (see Section 3). Here, we focus on the belief that the principle of composition
2355 of services can be applied to business processes and collaborations. Of course, business processes and
2356 collaborations traditionally represent complex, multi-step business functions that may involve multiple
2357 participants, including internal users, external customers, and trading partners. Therefore, such
2358 complexities cannot simply be ignored when transforming traditional business processes and
2359 collaborations to their service-oriented variants.

2360 Business processes are comprised of a set of coherent activities that, when performed in a logical
2361 sequence over a period of time and with appropriate rules applied, result in a certain business outcome.
2362 Service orientation as applied to business processes (i.e., "service-oriented business processes") means
2363 that the aggregation or composition of all of the abstracted activities, flows, and rules that govern a
2364 business process can themselves be abstracted as a service [BLOOMBERG/SCHMELZER].

2365 When business processes are abstracted in this manner and accessed through SOA services, all of the
2366 concepts used to describe and model composition of services that were articulated in Section 4.3.4 apply.

2367 There are some important differences from a composite service that represents an abstraction of a
2368 business process from a composite service that represents a single-step business interaction. As stated
2369 earlier, business processes have temporal properties and can range from short-lived processes that
2370 execute on the order of minutes or hours to long-lived processes that can execute for weeks, months, or
2371 even years. Further, these processes may involve many participants. These are important
2372 considerations for the consumer of a service-oriented business process and these temporal properties
2373 must be articulated as part of the meta-level aspects of the service-oriented business process in its
2374 Service Description, along with the meta-level aspects of any sub-processes that may be of use or need
2375 to be visible to the Service Consumer.

2376 In addition, a workflow activity represents a unit of work that some entity acting in a described role (i.e.,
2377 role player) is asked to perform. Activities can be broken down into steps with each step representing a
2378 task for the role player to perform. Based on our earlier assertion that messages denote joint action
2379 between service participants, we model these tasks as actions, i.e., message exchanges, which model
2380 activities as a collection of action-specific message exchanges. The role player performing a task or sub-
2381 task of a particular activity in an overall process flow may actually be a human entity and not a software or
2382 hardware agent.

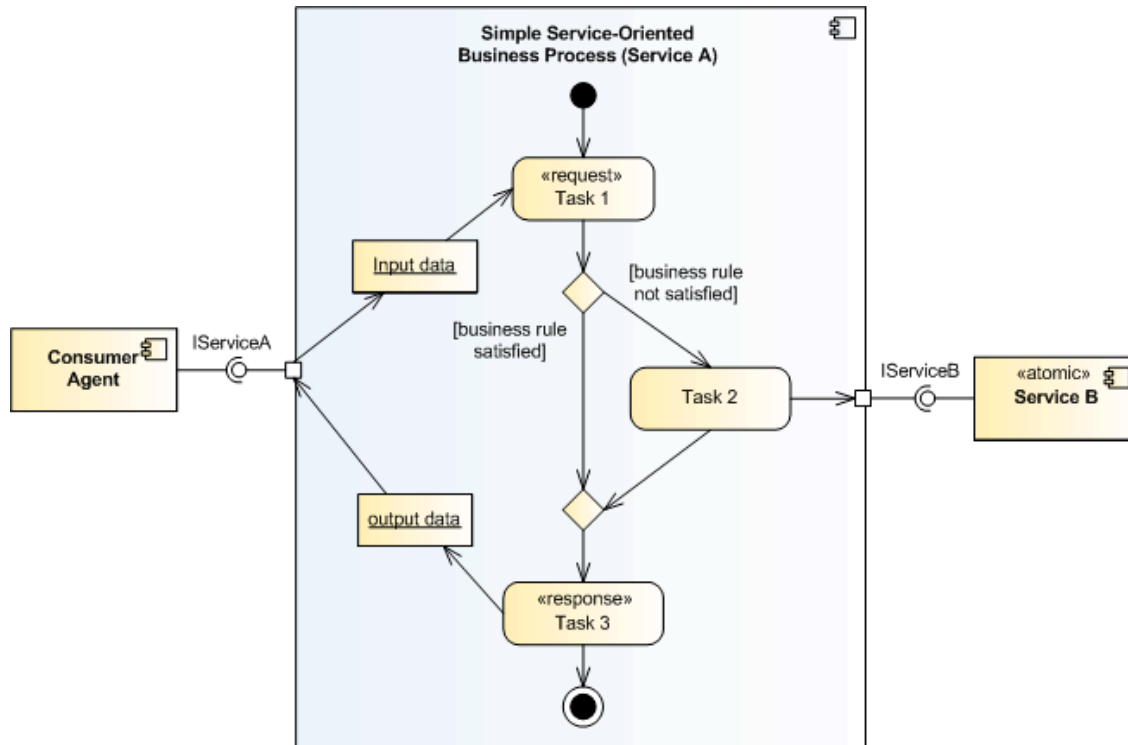
2383 A technique that is used to compose service-oriented business processes that are hierarchical (top-down)
2384 and self-contained in nature is known as *orchestration*.

2385 **Orchestration**

2386 A technique used to compose hierarchical and self-contained service-oriented business
2387 processes that are executed and coordinated by a single agent acting in a “conductor” role.

2388 An orchestration is typically implemented using a scripting approach to compose service-oriented
2389 business processes. This typically involves use of a standards-based orchestration scripting language.
2390 In terms of automation, an orchestration can be mechanized using a business process orchestration
2391 engine, which is a hardware or software component (agent) responsible for acting in the role of central
2392 conductor/coordinator responsible for executing the flows that comprise the orchestration.

2393 A simple generic example of such an orchestration is illustrated in Figure 45.



2394
 2395 *Figure 45 Abstract example of orchestration of service-oriented business process.*

2396 Here, we use a UML activity diagram to model the simple service-oriented business process as it allows
 2397 us to capture the major elements of business processes such as the set of related tasks to be performed,
 2398 linking between tasks in a logical flow, data that is passed between tasks, and any relevant business
 2399 rules that govern the transitions between tasks. A task is a unit of work that an individual, system, or
 2400 organization performs and can be accomplished in one or more steps or subtasks. While subtasks can
 2401 be readily modeled, they are not illustrated in the orchestration model In Figure 45..

2402 This particular example is based on a request/response MEP and captures how one particular task (Task
 2403 2) actually utilizes an externally-provided service, Service B. The entire service-oriented business
 2404 process is exposed as Service A that is accessible via its externally visible interface, IServiceA.

2405 Although not explicitly shown in the orchestration model above, it is assumed that there exists a software
 2406 or hardware component, i.e., orchestration engine that executes the process flow. Recall that a central
 2407 concept to orchestration is that process flow is coordinated and executed by a single conductor agent;
 2408 hence the name “orchestration.”

2409 **4.3.4.2 Service-Oriented Business Collaborations**

2410 Business collaborations typically represent the interaction involved in executing business transactions,
 2411 where a *business transaction* is defined in the Service Ecosystem View as “a joint action engaged in by
 2412 two or more participants in which resources are exchanged” (see Section 3.3.5).

2413 It is important to note that business collaborations represent “peer”-style interactions; in other words,
 2414 peers in a business collaboration act as equals. This means that unlike the orchestration of business
 2415 processes, there is no single or central entity that coordinates or “conducts” a business collaboration.
 2416 These peer styles of interactions typically occur between trading partners that span organizational
 2417 boundaries.

2418 Business collaborations can also be service-enabled. For purposes of this Reference Architecture, we
 2419 refer to these as “service-oriented business collaborations.” Service-oriented business collaborations do
 2420 not necessarily imply exposing the entire peer-style business collaboration as a service itself but rather
 2421 the collaboration uses service-based interchanges.

2422 The technique that is used to compose service-oriented business collaborations in which multiple parties
 2423 collaborate in a peer-style as part of some larger business transaction by exchanging messages with
 2424 trading partners and external organizations (e.g., suppliers) is known as *choreography*
 2425 **[NEWCOMER/LOMOW]**.

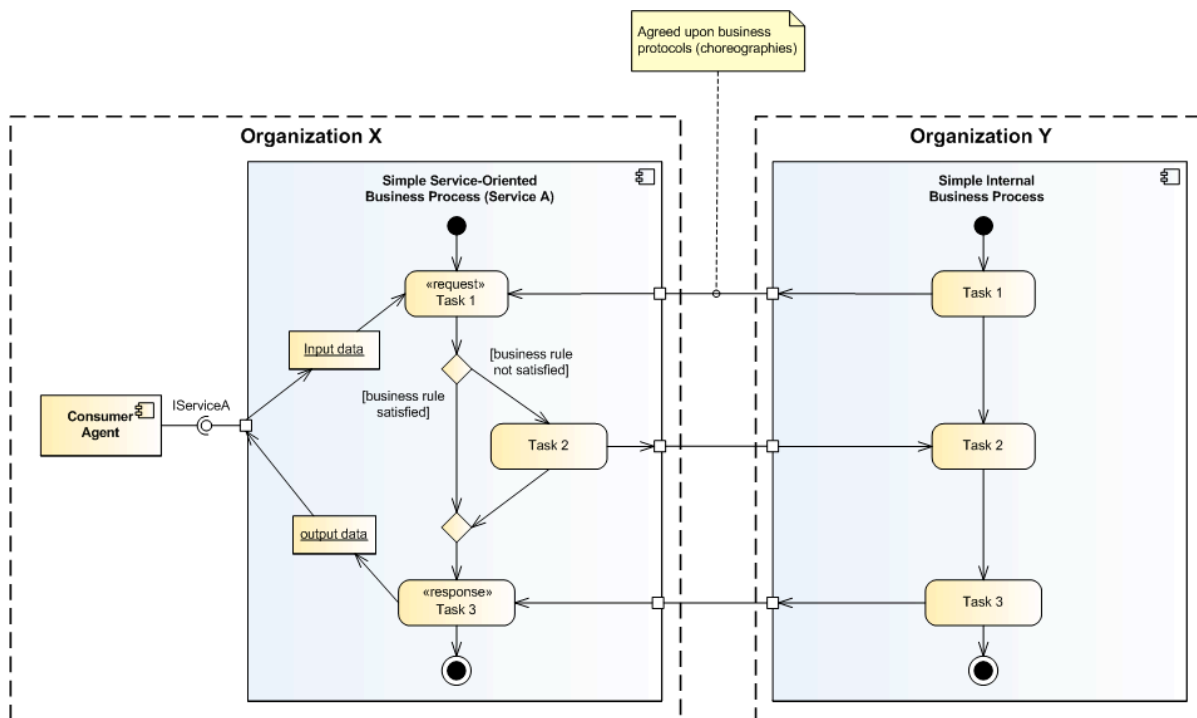
2426 **Choreography**

2427 A technique used to characterize and to compose service-oriented business collaborations based
 2428 on ordered message exchanges between peer entities in order to achieve a common business
 2429 goal.

2430 Choreography differs from orchestration primarily in that each party in a business collaboration describes
 2431 its part in the service interaction in terms of public message exchanges that occur between the multiple
 2432 parties as standard atomic or composite services, rather than as specific service-oriented business
 2433 processes that a single conductor/coordinator (e.g., orchestration engine) executes. Note that
 2434 choreography as we have defined it here should not be confused with the term *process choreography*,
 2435 which is defined in the Service Ecosystem View as “the description of the possible interactions that may
 2436 take place between two or more participants to fulfill an objective.” This is an example of domain-specific
 2437 nomenclature that often leads to confusion and why we are making note of it here.

2438 As is the case of an orchestration, a choreography is typically implemented by using a scripting approach
 2439 to composing service-oriented business collaborations. This typically involves use of a standards-based
 2440 choreography scripting language.

2441 A simple generic example of a choreography is illustrated in Figure 46



2442
 2443 *Figure 46 Abstract example of choreography of service-oriented business collaboration.*

2444 This example, which is a variant of the orchestration example illustrated earlier in Figure 45 adds trust
 2445 boundaries between two organizations; namely, Organization X and Organization Y. It is assumed that
 2446 these two organizations are peer entities that have an interest in a business collaboration, for example,
 2447 Organization X and Organization Y could be trading partners. Organization X retains the service-oriented
 2448 business process Service A, which is exposed to internal consumers via its provided service interface,
 2449 IServiceA. Organization Y also has a business process that is involved in the business collaboration;
 2450 however, for this example, it is an internal business process that is not exposed to potential consumers
 2451 either within or outside its organizational boundary.

2452 The scripting language that is used for the choreography needs to define how and when to pass control
2453 from one trading partner to another, i.e., Organization X and Organization Y. Defining the business
2454 protocols used in the business collaboration involves precisely specifying the visible message exchange
2455 behavior of each of the parties involved in the protocol, without revealing internal implementation details
2456 **[NEWCOMER/LOMOW]**.

2457 If, a peer-style business collaboration in which visibility into and use of each participating organization's
2458 internal service-oriented business processes was necessary as part of an end-to-end business
2459 transaction, then it would be desirable to select a choreography scripting language that would support
2460 interaction between different orchestration engines that spans organizational boundaries. WS-CDL is an
2461 example of such a language.

2462 **4.3.5 Architectural Implications of Interacting with Services**

2463 Interacting with Services has the following architectural implications on mechanisms that facilitate service
2464 interaction:

- 2465 • A well-defined service Information Model that:
 - 2466 ○ describes the syntax and semantics of the messages used to denote actions and events;
 - 2467 ○ describes the syntax and semantics of the data payload(s) contained within messages;
 - 2468 ○ documents exception conditions in the event of faults due to network outages, improper
2469 message/data formats, etc.;
 - 2470 ○ is both human readable and machine processable;
 - 2471 ○ is referenceable from the Service Description artifact.
- 2472 • A well-defined service Behavior Model that:
 - 2473 ○ characterizes the knowledge of the actions invokes against the service and events that
2474 report real world effects as a result of those actions;
 - 2475 ○ characterizes the temporal relationships and temporal properties of actions and events
2476 associated in a service interaction;
 - 2477 ○ describe activities involved in a workflow activity that represents a unit of work;
 - 2478 ○ describes the role(s) that a role player performs in a service-oriented business process or
2479 service-oriented business collaboration;
 - 2480 ○ is both human readable and machine processable;
 - 2481 ○ is referenceable from the Service Description artifact.
- 2482 • Service composition mechanisms to support orchestration of service-oriented business processes and
2483 choreography of service-oriented business collaborations such as:
 - 2484 ○ Declarative and programmatic compositional languages;
 - 2485 ○ Orchestration and/or choreography engines that support multi-step processes as part of a
2486 short-lived or long-lived business transaction;
 - 2487 ○ Orchestration and/or choreography engines that support compensating transactions in
2488 the presences of exception and fault conditions.
- 2489 • Infrastructure services that provides mechanisms to support service interaction, including but not
2490 limited to:
 - 2491 ○ mediation services such as message and event brokers, providers, and/or buses that
2492 provide message translation/transformation, gateway capability, message persistence,
2493 reliable message delivery, and/or intelligent routing semantics;
 - 2494 ○ binding services that support translation and transformation of multiple application-level
2495 protocols to standard network transport protocols;
 - 2496 ○ auditing and logging services that provide a data store and mechanism to record
2497 information related to service interaction activity such as message traffic patterns,
2498 security violations, and service contract and policy violations
 - 2499 ○ security services that abstract techniques such as public key cryptography, secure
2500 networks, virus protection, etc., which provide protection against common security threats
2501 in a SOA ecosystem;
 - 2502 ○ monitoring services such as hardware and software mechanisms that both monitor the
2503 performance of systems that host services and network traffic during service interaction,
2504 and are capable of generating regular monitoring reports.

- 2505 • A layered and tiered service component architecture that supports multiple message exchange
2506 patterns (MEPs) in order to:
- 2507 ○ promote the industry best practice of separation of concerns that facilitates flexibility in
2508 the presence of changing business requirements;
 - 2509 ○ promote the industry best practice of separation of roles in a service development
2510 lifecycle such that subject matter experts and teams are structured along areas of
2511 expertise;
 - 2512 ○ support numerous standard interaction patterns, peer-to-peer interaction patterns,
2513 enterprise integration patterns, and business-to-business integration patterns.

2514 **4.4 Policies and Contracts Model**

2515 As described in the Reference Model, a policy is an enforceable constraint or condition on the use,
2516 deployment, or description of an owned entity as defined by any participant. A contract is a constraint
2517 that has the agreement of the constrained participants.

2518 This Reference Architecture reflects a common separation between *mechanism* and *policy*. In many
2519 situations it is often simpler to build mechanisms that address a more general problem than the one at
2520 hand, and then to use that general mechanism to solve the particular problem. In the case of a SOA
2521 ecosystem, the mechanisms focus on the ability to match participants' needs with service capabilities.
2522 However, each particular combination of need and capability is very likely to be distinct; resulting in a
2523 large number of circumstances. Policies can be used as a framework for managing this combinatorial
2524 explosion.

2525 Policies and contracts have wide applicability within this Reference Architecture. They are used to
2526 express security policies, service policies, relationships and constraints within the social structures that
2527 encapsulate service participants, management of services and many other instances. The enforcement of
2528 a policy or contract may be a part of the SOA-based computing environment or it may be handled outside
2529 of the SOA-based computing environment.

2530 **4.4.1 Policy and Contract Principles**

2531 In the realization of policies and contracts for a SOA, there are common policy principles that will be
2532 encountered in many of the standards and/or technology choices used for the realization. Some of these
2533 common principles are covered in this section.

2534 **4.4.1.1 Goals of Policies and Contracts**

2535 Policies SHOULD reflect the goals of governance or management processes, see Section [5.1](#)
2536 [Governance of Service Oriented Architectures](#) and section [5.3 Services as Managed Entities Model](#). The
2537 governance and management processes SHOULD use formal and standardized policy languages to
2538 enable the widest possible understanding and use of stated policies and contracts, and architecture
2539 components SHOULD be available to enable compliance.

2540 **4.4.1.2 Policy and Contract Specification**

2541 The language used to describe policies and contracts inevitably constrains the forms and types of policies
2542 and contracts expressible in the description. Formal policy language definitions are outside the scope of
2543 this specification. For formal policy languages, standard specifications such as XACML and WS-Policy
2544 may be referenced. Policy/Contract descriptions may be associated with a service through the Service
2545 Description as defined in Section [4.1 Service Description Model](#).

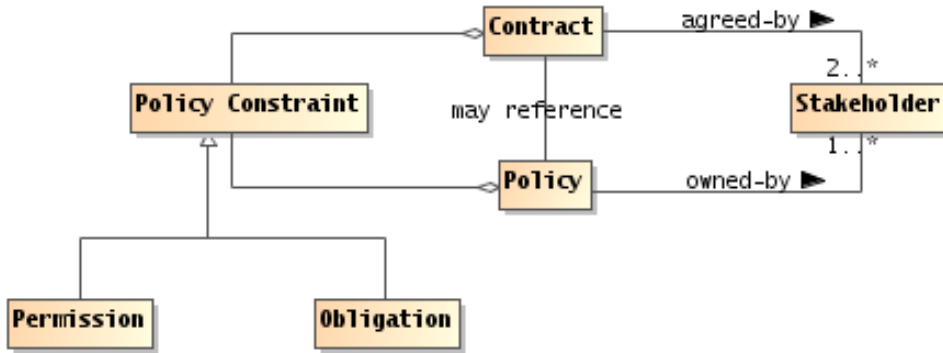
2546 Regardless of the language used to describe policies and contracts, there are certain aspects to capture
2547 in any system for the representation of policies and contracts such as:

- 2548 • how to describe atomic policy constraints
- 2549 • how to nest policy constraints allowing for abstractions and refinements of a policy constraint
- 2550 • how to reference policy constraints allowing for the reuse of a policy constraint

- 2551 • how to define alternative policy constraints for the selection of compatible policy constraints
- 2552 between the consumer and provider
- 2553 • policy versioning
- 2554 • policy modules

2555 **4.4.1.3 Policy Constraints**

2556 Policies are often characterized in terms of permissions or about obligations. An overriding meta-
 2557 constraint is that policy constraints (likewise contract constraints) **MUST** be enforceable – a constraint
 2558 that is not enforceable is not a legitimate element of a system of policies and contracts in the SOA
 2559 ecosystem.



2560
 2561 *Figure 47 Policies and Contracts*

2562 **Policy**

2563 A policy represents some constraint or condition on the use, deployment or description of a
 2564 resource as defined by a participant or, more generally, a stakeholder.

2565 **Contract**

2566 A contract represents an agreement by two or more participants to constrain their behavior and
 2567 state.

2568 **Policy Constraint**

2569 A policy constraint is a measurable proposition that characterizes the constraint that the policy is
 2570 about.

2571 **Permission**

2572 A permission constraint governs the ability of a participant or other actor to perform an action or
 2573 enter some specified state.

2574 Permissions may apply to any action that any actor may be able to perform. Note that permissions are
 2575 distinct from ability and from authority. Authority refers to the legitimate nature of an action, whereas
 2576 permission refers to the right to perform the action.

2577 **Obligation**

2578 An obligation constraint governs the requirement that a participant or other actor should perform
 2579 an action or maintain some specified state.

2580 For example, once the service consumer and provider have entered into an agreement to provide and
 2581 consume a service, both participants incur obligations: the consumer is obligated to pay for the service
 2582 and the provider is obligated to provide the service.

2583 Obligations to maintain state may range from a requirement to maintain a minimum balance on an
 2584 account through a requirement that a service provider ‘remember’ that a particular service consumer is
 2585 logged in.

2586 A permission-style constraint is about the right to access some resource or perform some action; an
2587 obligation-style constraint is about the requirement to perform some action or maintain the state of a
2588 resource.

2589 Obligations and Permissions have a positive form and a negative form. A positive permission refers to
2590 something that you may do, a negative permission refers to something you should not do.

2591 These are combinable, in the sense that you may have a positive permission constraint (for example, you
2592 may use encryption in your messages), whereas a negative permission constraint indicates that there is
2593 something you may not do. Similarly, a positive obligation may be something like you must keep the
2594 balance of your account positive; whereas an example of a negative obligation may be that the bank will
2595 not cover a check for more than the balance in your account.

2596 Permission-style constraints are often checkable a-priori: before the intended action or access is
2597 completed the current permission constraints may be applied to deny the access if necessary. However,
2598 obligation-style constraints can normally only be verified post-priori.

2599 Policies and contracts can contain a mix of permissions and obligations. The degree with which these can
2600 be effectively combined depends on the specific policy management framework in use.

2601 In a business context, contracts are legally binding agreements between two or more parties. A contract
2602 is formed when there is an offer that is duly made and the offer is accepted and there is evidence that
2603 indicates there was a tangible exchange of value between the two parties. While this Reference
2604 Architecture is inclusive of legally binding contracts for a SOA, contracts do not always have to be legally
2605 binding agreements.

2606 A contract may include references to policies and other contracts while a policy may include references to
2607 contracts and other policies. For example, a contract may reference a set of policies and a policy may
2608 prioritize certain contracts over others.

2609 **4.4.1.4 Policy Composition**

2610 Multiple policies may be defined for one or more services in one or more ownership domains. The
2611 application of policies and contracts over distributed services requires the ability to compose one or more
2612 policies into an overarching policy. The composition of policies may be implemented as a hierarchy
2613 and/or it can be implemented as intersections and unions of sets.

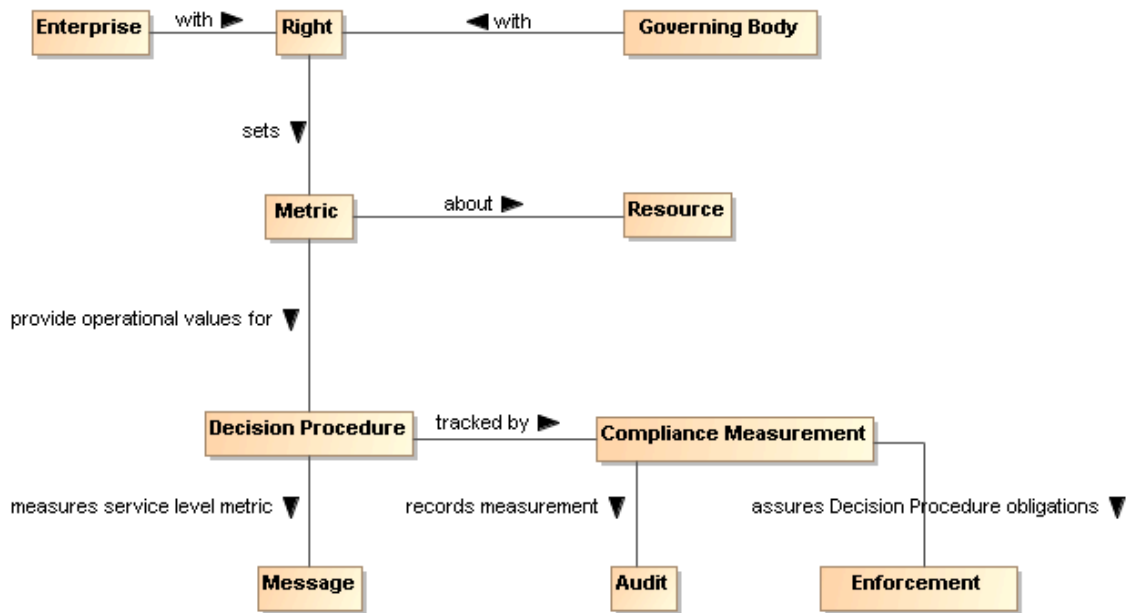
2614 **4.4.1.5 Conflict Resolution**

2615 The analysis of policy rules may result in conflicts between the policy rules. There can be many causes
2616 for policy conflicts such as conflicting policy rules between ownership domains and policy language
2617 specifications that do not convert to first order predicate logic for IT policy mechanisms. This can cause
2618 policy decision results to be indeterminate. Policy administration mechanisms may provide conflict
2619 resolution capabilities prior to the storage/distribution of policies. At run time, conflicts may propagate to
2620 higher authorities inside or outside the SOA-based IT mechanisms.

2621 **4.4.1.6 Delegation of Policy**

2622 Policy authorization may be delegated to agents acting on behalf of a client to enable decentralized policy
2623 administration and/or policy enforcement. This allows policies to be administered and/or enforced in a
2624 hierarchical fashion. Policies may also be transferred to an agent or resource to effectively allow that
2625 agent or resource to separate from an ownership domain. The agent or resource may join another
2626 ownership domain or rejoin the same ownership domain at a later time.

2627 **4.4.2 Policy Metrics**



2628
2629 Figure 50 Policy Metrics

2630
2631 **Metric**

2632 A metric is a policy constraint used to measure compliance.

2633
2634 Metrics are often expressed to measure service performance compliance and regulatory compliance.
2635 Service Level Agreements (SLAs) are one commonly used category of contractual compliance. The
2636 metrics that comprise a SLA often consist of (Service Level) constraints such as service warranties and
2637 remedies or business level constraints such as <business terms and conditions>.

2638 **4.4.3 Automating Support for Policies and Contracts**

2639 There are many functional `control points' in a SOA ecosystem; for example, how messages are
2640 exchanged, what descriptions should be made available to which participant, what services should be
2641 offered and so on.

2642 An effective technique for managing these control points is via policies; together with mechanisms for
2643 distributing and applying policies appropriately.

2644 From the IT perspective, high level policies and contracts need to be translated into low level rules and
2645 measurable properties that programmatic elements can enforce. For low level rules and measurable
2646 properties, both contracts and policies are likely to be enforced by the same type of IT policy
2647 mechanisms.

2648 **4.4.3.1 IT Mechanisms Supporting Policies and Contracts**

2649 The mechanism for enforcing a permission-oriented constraint is typically prevention at the point of
2650 action. The mechanisms for enforcing obligation constraints are typically achieved by a combination of
2651 auditing and remedial action.

2652 A common phenomenon of many machines and systems is that they are much broader in their potential
2653 than is actually needed for a particular circumstance. As a result, the behavior and performance of the
2654 system tend to be under-constrained by the implementation. Policy statements define the choices that a
2655 service provider and/or service consumer (or other stakeholder) makes; these choices are used to guide
2656 the actual behavior of the system to the desired behavior and performance.

2657 While there are many possible approaches to the realization of policy/contracts for a SOA, one approach
2658 based on current policy standardization efforts is depicted in this section. The common policy
2659 architectural elements that are provided in this section are based on the minimal mechanisms required to
2660 provide policy guided delivery of services within an ownership domain and across ownership domains.

2661 **4.4.4 Architectural Implications**

2662 While policy and contract descriptions have much of the same architectural implications as described in
2663 Service Description, languages and mechanisms supporting policies and contracts also have the
2664 following architectural implications:

- 2665 • Policy and Contract language specifications will typically provide support for the following capabilities:
 - 2666 ○ expression of assertion and commitment policy constraints;
 - 2667 ○ expression of positive and negative policy constraints;
 - 2668 ○ expression of permission and obligation policy constraints;
 - 2669 ○ nesting of policy constraints allowing for abstractions and refinements of a policy constraint;
 - 2670 ○ definition of alternative policy constraints to allow for the selection of compatible policy
 - 2671 constraints for a consumer and provider;
 - 2672 ○ composition of policies to combine one or more policies.
- 2673 • Policy and contract mechanisms in a SOA ecosystem will require the following capabilities:
 - 2674 ○ decision procedures which must be able to measure and render decisions on constraints;
 - 2675 ○ enforcement of decisions;
 - 2676 ○ measurement and notification of obligation constraints;
 - 2677 ○ auditability of decisions, enforcement, and obligation measurements;
 - 2678 ○ administration of policy and contract language artifacts;
 - 2679 ○ storage of policies and contracts;
 - 2680 ○ distribution of policies/contracts;
 - 2681 ○ conflict resolution or elevation of conflicts in policy rules;
 - 2682 ○ delegation of policy authority to agents acting on behalf of a client;
 - 2683 ○ decision procedures capable of incorporating roles and/or attributes for rendered decisions.

5 Owning Service Oriented Architectures View

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

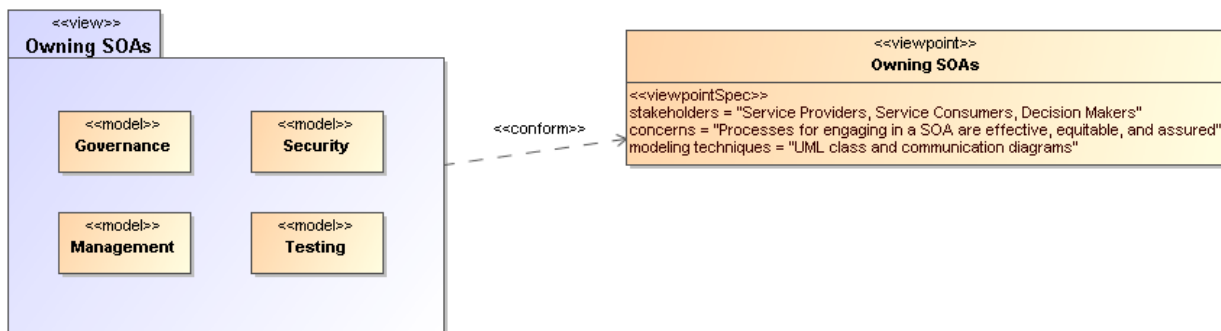
2684
2685
2686
2687
2688

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

2691 Owning a SOA-based system raises significantly different challenges to owning other complex systems --
2692 such as Enterprise suites -- because there are strong limits on the control and authority of any one party
2693 when a system spans multiple ownership domains.

2694 Even when a SOA-based system is deployed internally within an organization, there are multiple internal
2695 stakeholders involved and there may not be a simple hierarchy of control and management. Thus, an
2696 early consideration of how multiple boundaries affect SOA-based systems will provide a firm foundation
2697 for dealing with them in whatever form they are found rather than debating whether the boundaries should
2698 exist.

2699 This view focuses on the Governance of SOA-based systems, on the security challenges involved in
2700 running a SOA-based system and the management challenges.



2701
2702 *Figure 48 Model Elements Described in the Owning Service Oriented Architectures View*

2703 The following subsections present models of these functions.

5.1 Governance Model

2705 The Reference Model defines Service Oriented Architecture as an architectural paradigm for organizing
2706 and utilizing distributed capabilities that may be under the control of different ownership domains [**SOA-**
2707 **RM**]. Consequently, it is important that organizations that plan to engage in service interactions adopt
2708 governance policies and procedures sufficient to ensure that there is standardization across both internal
2709 and external organizational boundaries to promote the effective creation and use of SOA-based services.

5.1.1 Understanding Governance

5.1.1.1 Terminology

2712 Governance is about making decisions that are aligned with the overall organizational strategy and
2713 culture of the enterprise. [**Gartner**] It specifies the decision rights and accountability framework to
2714 encourage desirable behaviors [**Weill/Ross-MIT Sloan School**] towards realizing the strategy and
2715 defines incentives (positive or negative) towards that end. It is less about overt control and strict
2716 adherence to rules, and more about guidance and effective and equitable usage of resources to ensure
2717 sustainability of an organization's strategic objectives. [**TOGAF v8.1**]

2718 To accomplish this, governance requires organizational structure and processes and must identify who
2719 has authority to define and carry out its mandates. It must address the following questions: 1) what
2720 decisions must be made to ensure effective management and use?, 2) who should make these
2721 decisions?, and 3) how will these decisions be made and monitored? , and (4) how will these decisions
2722 be communicated? The intent is to achieve goals, add value, and reduce risk.

2723 Within a single ownership domain such as an enterprise, generally there is a hierarchy of governance
2724 structures. Some of the more common enterprise governance structures include corporate governance,
2725 technology governance, IT governance, and architecture governance **[TOGAF v8.1]**. These governance
2726 structures can exist at multiple levels (global, regional, and local) within the overall enterprise.

2727 It is often asserted that SOA governance is a specialization of IT governance as there is a natural
2728 hierarchy of these types of governance structures; however, the focus of SOA governance is less on
2729 decisions to ensure effective management and use of IT as it is to ensure effective management and use
2730 of SOA-based systems. Certainly, SOA governance must still answer the basic questions also
2731 associated with IT governance, i.e., who should make the decisions, and how these decisions will be
2732 made and monitored.

2733 **5.1.1.2 Relationship to Management**

2734 There is often confusion centered on the relationship between governance and management. As
2735 described earlier, governance is concerned with decision making. Management, on the other hand, is
2736 concerned with execution. Put another way, governance describes the world as leadership wants it to be;
2737 management executes activities that intends to make the leadership's desired world a reality. Where
2738 governance determines who has the authority and responsibility for making decisions and the
2739 establishment of guidelines for how those decisions should be made, management is the actual process
2740 of making, implementing, and measuring the impact of those decisions **[Loeb]**. Consequently,
2741 governance and management work in concert to ensure a well-balanced and functioning organization as
2742 well as an ecosystem of inter-related organizations. In the sections that follow, we elaborate further on
2743 the relationship between governance and management in terms of setting and enforcing service policies,
2744 contracts, and standards as well as addressing issues surrounding regulatory compliance.

2745 **5.1.1.3 Why is SOA Governance Important?**

2746 One of the hallmarks of SOA that distinguishes it from other architectural paradigms for distributed
2747 computing is the ability to provide a uniform means to offer, discover, interact with and use capabilities
2748 (as well the ability to compose new capabilities from existing ones) all in an environment that transcends
2749 domains of ownership. Consequently, ownership, and issues surrounding it, such as obtaining
2750 acceptable terms and conditions (T&Cs) in a contract, is one of the primary topics for SOA governance.
2751 Generally, IT governance does not include T&Cs, for example, as a condition of use as its primary
2752 concern.

2753 Just as other architectural paradigms, technologies, and approaches to IT are subject to change and
2754 evolution, so too is SOA. Setting policies that allow change management and evolution, establishing
2755 strategies for change, resolving disputes that arise, and ensuring that SOA-based systems continue to
2756 fulfill the goals of the business are all reasons why governance is important to SOA.

2757 **5.1.1.4 Governance Stakeholders and Concerns**

2758 As noted in Section 3.3.1 the participants in a service interaction include the service provider, the service
2759 consumer, and other interested or unintentional third parties. Depending on the circumstances, it may
2760 also include the owners of the underlying capabilities that the SOA services access. Governance must
2761 establish the policies and rules under which duties and responsibilities are defined and the expectations
2762 of participants are grounded. The expectations include transparency in aspects where transparency is
2763 mandated, trust in the impartial and consistent application of governance, and assurance of reliable and
2764 robust behavior throughout the SOA ecosystem.

2765 **5.1.2 A Generic Model for Governance**

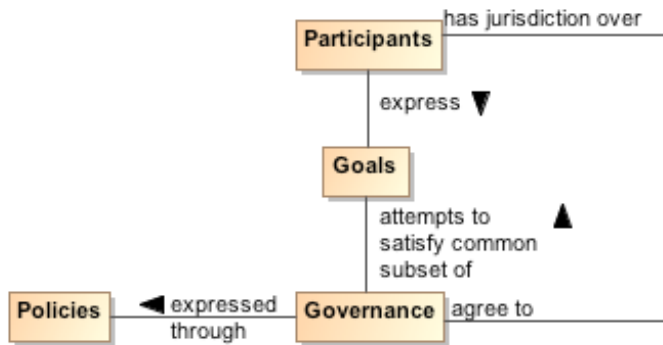
2766 **Governance**

2767 Governance is the prescribing of conditions and constraints consistent with satisfying common
2768 goals and the structures and processes needed to define and respond to actions taken towards
2769 realizing those goals.

2770 The following is a generic model of governance represented by segmented models that begin with
2771 motivation and proceed through measuring compliance. It is not meant to be an all-encompassing treatise
2772 on governance but a focused subset that captures the aspects necessary to describe governance for
2773 SOA. It is also not meant to imply that practical application of governance is a single, isolated instance of
2774 these models; in fact, there are likely hierarchical chains of governance that apply and possibly parallel
2775 chains that govern different aspects or focus on different goals. This is discussed further in section
2776 5.1.2.5. The defined models are simultaneously applicable to each of the overlapping instances.

2777 A given enterprise may already have portions of these models in place. To a large extent, the models
2778 shown here are not specific to SOA; discussions on direct applicability begin in section 5.1.3.

2779 **5.1.2.1 Motivating Governance**



2780
2781 *Figure 49 Motivating governance model*

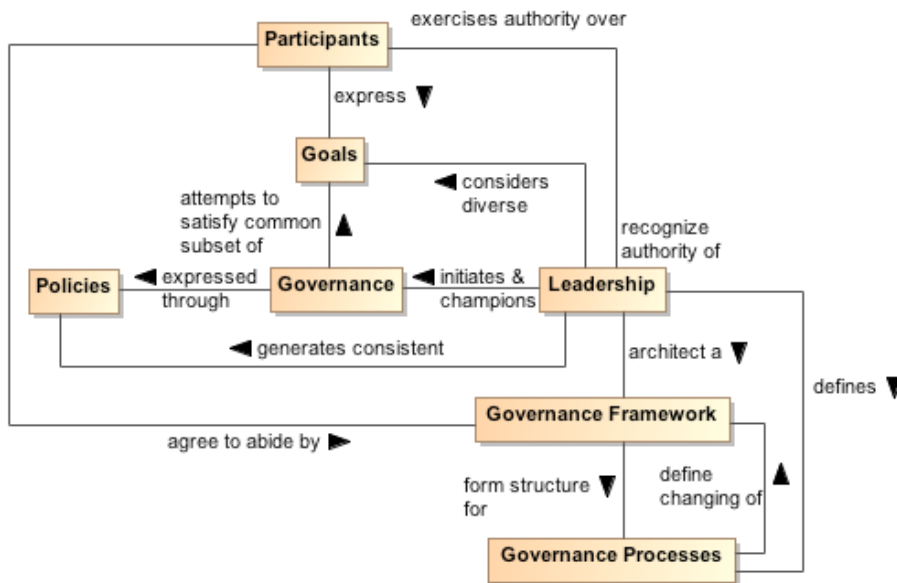
2782 An organizational domain such as an enterprise is made up of Participants who may be individuals or
2783 groups of individuals forming smaller organizational units within the enterprise. The overall business
2784 strategy should be consistent with the Goals of the participants; otherwise, the business strategy would
2785 not provide value to the participants and governance towards those ends becomes difficult if not
2786 impossible. This is not to say that an instance of governance will simultaneously satisfy all the goals of all
2787 the participants; rather, the goals of any governance instance must sufficiently satisfy a useful subset of
2788 each participant's goals so as to provide value and ensure the cooperation of all the participants.

2789 A policy is the formal characterization of the conditions and constraints that governance deems as
2790 necessary to realize the goals which it is attempting to satisfy. Policy may identify required conditions or
2791 actions or may prescribe limitations or other constraints on permitted conditions or actions. For example,
2792 a policy may prescribe that safeguards must be in place to prevent unauthorized access to sensitive
2793 material. It may also prohibit use of computers for activities unrelated to the specified work assignment.
2794 Policy is made operational through the promulgating and implementing of Rules and Regulations (as
2795 defined in section 5.1.2.3).

2796 As noted in section 4.4.2, policy may be asserted by any participant or on behalf of the participant by its
2797 organization. Part of the purpose of governance is to arbitrate among diverse goals of participants and
2798 diverse policies articulated to realize those goals. The intent is to form a consistent whole that allows
2799 governance to minimize ambiguity about its purpose. While resolving all ambiguity would be an ideal, it is
2800 unlikely that all inconsistencies will be identified and resolved before governance becomes operational.

2801 For governance to have effective jurisdiction over participants, there must be some degree of agreement
2802 by each participant that it will abide by the governance mandates. A minimal degree of agreement often
2803 presages participants who "slow-roll" if not actively reject complying with Policies that express the
2804 specifics of governance.

2805 **5.1.2.2 Setting Up Governance**



2806
2807 *Figure 50 Setting up governance model*

2808 **Leadership**

2809 Leadership is the entity who has the responsibility and authority to generate consistent policies
2810 through which the goals of governance can be expressed and to define and champion the
2811 structures and processes through which governance is realized.

2812 **Governance Framework**

2813 The Governance Framework is a set of organizational structures that enable governance to be
2814 consistently defined, clarified, and as needed, modified to respond to changes in its domain of
2815 concern.

2816 **Governance Processes**

2817 Governance Processes are the defined set of activities that are performed within the Governance
2818 Framework to enable the consistent definition, application, and as needed, modification of Rules
2819 that organize and regulate the activities of Participants for the fulfillment of expressed policies.
2820 (See section 5.1.2.3 for elaboration on the relationship of Governance Processes and Rules.)

2821 As noted earlier, governance requires an appropriate organizational structure and identification of who
2822 has authority to make governance decisions. In Figure 50, the entity with governance authority is
2823 designated the Leadership. This is someone, possibly one or more of the Participants, that Participants
2824 recognize as having authority for a given purpose or over a given set of issues or concerns.

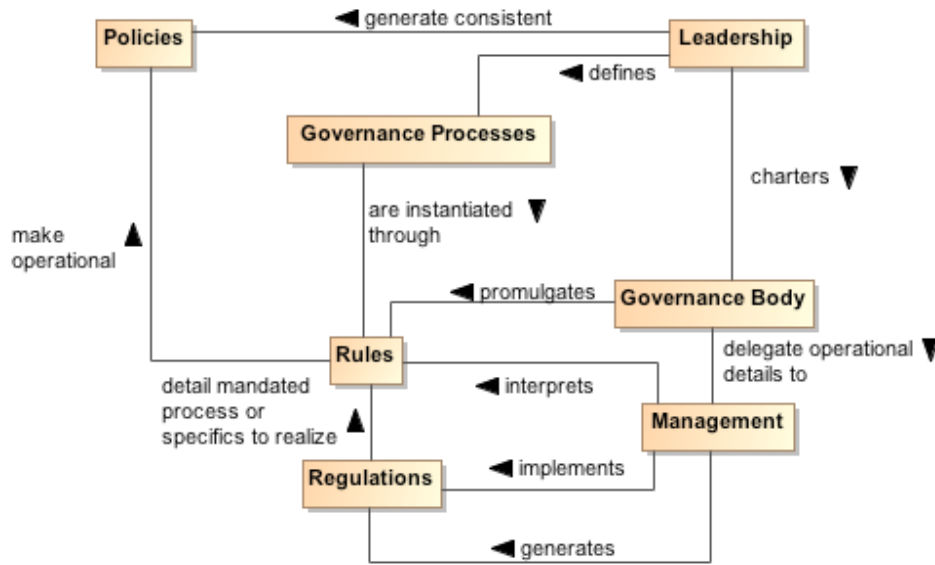
2825 The Leadership is responsible for prescribing or delegating a working group to prescribe the Governance
2826 Framework that forms the structure for Governance Processes which define how governance is to be
2827 carried out. This does not itself define the specifics of how governance is to be applied, but it does
2828 provide an unambiguous set of procedures that should ensure consistent actions which Participants
2829 agree are fair and account for sufficient input on the subjects to which governance will be applied.

2830 The Participants may be part of the working group that codifies the Governance Framework and
2831 Processes. When complete, the Participants must acknowledge and agree to abide by the products
2832 generated through application of this structure.

2833 The Governance Framework and Processes are often documented in the charter of a body created or
2834 designated to oversee governance. This is discussed further in the next section. Note that the
2835 Governance Processes should also include those necessary to modify the Governance Framework itself.

2836 An important function of Leadership is not only to initiate but also be the consistent champion of
 2837 governance. Those responsible for carrying out governance mandates must have Leadership who
 2838 makes it clear to Participants that expressed Policies are seen as a means to realizing established goals
 2839 and that compliance with governance is required.

2840 **5.1.2.3 Carrying Out Governance**



2841
 2842 *Figure 51 Carrying out governance model*

2843 **Rule**

2844 A Rule is a prescribed guide for carrying out activities and processes leading to desired results,
 2845 e.g. the operational realization of policies.

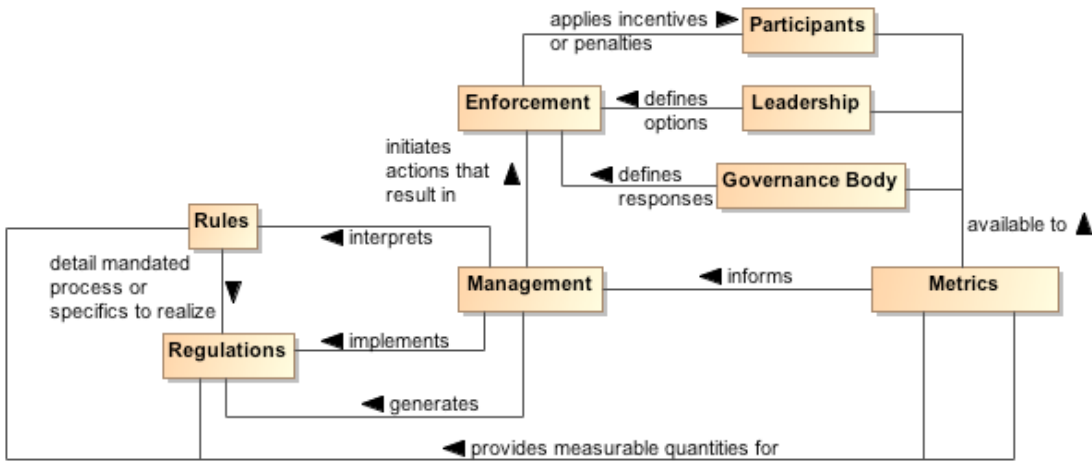
2846 **Regulation**

2847 A Regulation is a mandated process or the specific details that derive from the interpretation of
 2848 Rules and lead to measureable quantities against which compliance can be measured.

2849 To carry out governance, Leadership charts a Governance Body to promulgate the Rules needed to
 2850 make the Policies operational. The Governance Body acts in line with Governance Processes for its rule-
 2851 making process and other functions. Whereas Governance is the setting of Policies and defining the
 2852 Rules that provide an operational context for Policies, the operational details of governance are likely
 2853 delegated by the Governance Body to Management. Management generates Regulations that specify
 2854 details for Rules and other procedures to implement both Rules and Regulations. For example,
 2855 Leadership could set a Policy that all authorized parties should have access to data, the Governance
 2856 Body would promulgate a Rule that PKI certificates are required to establish identity of authorized parties,
 2857 and Management can specify a Regulation of who it deems to be a recognized PKI issuing body. In
 2858 summary, Policy is a predicate to be satisfied and Rules prescribe the activities by which that satisfying
 2859 occurs. A number of rules may be required to satisfy a given policy; the carrying out of a rule may
 2860 contribute to several policies being realized.

2861 Whereas the Governance Framework and Processes are fundamental for having Participants
 2862 acknowledge and commit to compliance with governance, the Rules and Regulations provide operational
 2863 constraints which may require resource commitments or other levies on the Participants. It is important
 2864 for Participants to consider the framework and processes to be fair, unambiguous, and capable of being
 2865 carried out in a consistent manner and to have an opportunity to formally accept or ratify this situation.
 2866 Rules and Regulations, however, do not require individual acceptance by any given participant although
 2867 some level of community comment is likely to be part of the Governance Processes. Having agreed to
 2868 governance, the Participants are bound to comply or be subject to prescribed mechanisms for
 2869 enforcement.

2870 **5.1.2.4 Ensuring Governance Compliance**



2871
2872 *Figure 52 Ensuring governance compliance model*

2873 Setting Rules and Regulations does not ensure effective governance unless compliance can be
 2874 measured and Rules and Regulations can be enforced. Metrics are those conditions and quantities that
 2875 can be measured to characterize actions and results. Rules and Regulations MUST be based on
 2876 collected Metrics or there will be no way for Management to assess compliance. The Metrics are
 2877 available to the Participants, the Leadership, and the Governance Body so what is measured and the
 2878 results of measurement are clear to everyone.

2879 The Leadership in its relationship with Participants will have certain options that can be used for
 2880 Enforcement. A common option may be to effect future funding. The Governance Body defines specific
 2881 enforcement responses, such as what degree of compliance is necessary for full funding to be restored.
 2882 It is up to Management to identify compliance shortfalls and to initiate the Enforcement process.

2883 Note, enforcement does not strictly need to be negative consequences. Management can use Metrics to
 2884 identify exemplars of compliance and Leadership can provide options for rewarding the Participants. It is
 2885 likely the Governance Body that defines awards or other incentives.

2886 **5.1.2.5 Considerations for Multiple Governance Chains**

2887 As noted in section 5.1.2, instances of the governance model often occur as a tiered arrangement, with
 2888 governance at some level delegating specific authority and responsibility to accomplish a focused portion
 2889 of the original level's mandate. For example, a corporation may encompass several lines of business and
 2890 each line of business governs its own affairs in a manner that is consistent with and contributes to the
 2891 goals of the parent organization. Within the line of business, an IT group may be given the mandate to
 2892 provide and maintain IT resources, giving rise to IT governance.

2893 In addition to tiered governance, there are likely to be multiple governance chains working in parallel. For
 2894 example, a company making widgets likely has policies intended to ensure they make high quality
 2895 widgets and make an impressive profit for their shareholders. On the other hand, Sarbanes-Oxley is a
 2896 parallel governance chain in the United States that specifies how the management must handle its
 2897 accounting and information that needs to be given to its shareholders. The parallel chains may just be
 2898 additive or may be in conflict and require some harmonization.

2899 Being distributed and representing different ownership domains, a SOA participant is likely under the
 2900 jurisdiction of multiple governance domains simultaneously and may individually need to resolve
 2901 consequent conflicts. The governance domains may specify precedence for governance conformance or
 2902 it may fall to the discretion of the participant to decide on the course of actions they believe appropriate.

2903 **5.1.3 Governance Applied to SOA**

2904 **5.1.3.1 Where SOA Governance is Different**

2905 SOA governance is often discussed in terms of IT governance, but rather than a parent-child relationship,
2906 Figure 53 shows the two as siblings of the general governance described in section 5.1.2. There are
2907 obvious dependencies and a need for coordination between the two, but the idea of aligning IT with
2908 business already demonstrates that resource providers and resource consumers must be working
2909 towards common goals if they are to be productive and efficient. While SOA governance will be shown to
2910 be active in the area of infrastructure, it is a specialized concern for having a dependable platform to
2911 support service interaction; a host of traditional IT issues is considered to be out of scope. A SOA
2912 governance plan for an enterprise will not resolve shortcomings with the enterprise IT governance.

2913 Governance in the context of SOA is that organization of services: that promotes their visibility; that
2914 facilitates interaction among service participants; and that directs that the results of service interactions
2915 are those real world effects as described within the service description and constrained by policies and
2916 contracts as assembled in the execution context.

2917 SOA governance must specifically account for control across different ownership domains, i.e. all the
2918 participants may not be under the jurisdiction of a single governance authority. However, for governance
2919 to be effective, the participants must agree to recognize the authority of the Governance Body and must
2920 operate within the Governance Framework and through the Governance Processes so defined.

2921 SOA governance must account for interactions across ownership boundaries, which likely also implies
2922 across enterprise governance boundaries. For such situations, governance emphasizes the need for
2923 agreement that some Governance Framework and Governance Processes have jurisdiction, and the
2924 governance defined must satisfy the Goals of the Participants for cooperation to continue. A standards
2925 development organization such as OASIS is an example of voluntary agreement to governance over a
2926 limited domain to satisfy common goals.

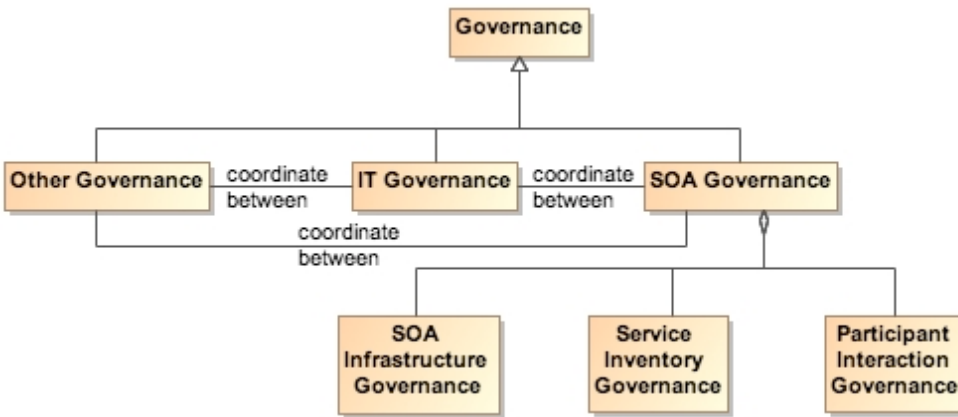
2927 The specifics discussed in the figures in the previous sections are equally applicable to governance
2928 across ownership boundaries as it is within a single boundary. There is a charter agreed to when
2929 Participants become members of the organization, and this charter sets up the structures and processes
2930 that will be followed. Leadership may be shared by the leadership of the overall organization and the
2931 leadership of individual groups themselves chartered per the Governance Processes. There are
2932 Rules/Regulations specific to individual efforts for which Participants agree to local goals, and
2933 Enforcement can be loss of voting rights or under extreme circumstances, expulsion from the group.

2934 Thus, the major difference for SOA governance is an appreciation for the cooperative nature of the
2935 enterprise and its reliance on furthering common goals if productive participation is to continue.

2936 **5.1.3.2 What Must be Governed**

2937 An expected benefit of employing SOA principles is the ability to quickly bring resources to bear to deal
2938 with unexpected and evolving situations. This requires a great deal of confidence in the underlying
2939 capabilities that can be accessed and in the services that enable the access. It also requires
2940 considerable flexibility in the ways these resources can be employed. Thus, SOA governance requires
2941 establishing confidence and trust while instituting a solid framework that enables flexibility, indicating a
2942 combination of strict control over a limited set of foundational aspects but minimum constraints beyond
2943 those bounds.

2944



2945
2946 *Figure 53 Relationship among types of governance*

2947 SOA governance applies to three aspects of service definition and use:

- 2948 • SOA infrastructure – the “plumbing” that provides utility functions that enable and support the use
- 2949 of the service
- 2950 • Service inventory – the requirements on a service to permit it to be accessed within the
- 2951 infrastructure
- 2952 • Participant interaction – the consistent expectations with which all participants are expected to
- 2953 comply

2954 **5.1.3.2.1 Governance of SOA Infrastructure**

2955 The SOA infrastructure is likely composed of several families of SOA services that provide access to
 2956 fundamental computing business services. These include, among many others, services such as
 2957 messaging, security, storage, discovery, and mediation. The provisioning of an infrastructure on which
 2958 these services may be accessed and the general realm of those contributing as utility functions of the
 2959 infrastructure are a traditional IT governance concern. In contrast, the focus of SOA governance is how
 2960 the existence and use of the services enables the SOA ecosystem.

2961 By characterizing the environment as containing families of SOA services, the assumption is that there
 2962 may be multiple approaches to providing the business services or variations in the actual business
 2963 services provided. For example, discovery could be based on text search, on metadata search, on
 2964 approximate matches when exact matches are not available, and numerous other variations. The
 2965 underlying implementation of search algorithms are not the purview of SOA governance, but the access
 2966 to the resulting service infrastructure enabling discovery must be stable, reliable, and extremely robust to
 2967 all operating conditions. Such access enables other specialized SOA services to use the infrastructure in
 2968 dependable and predictable ways, and is where governance is important.

2969 **5.1.3.2.2 Governance of the Service Inventory**

2970 Given an infrastructure in which other SOA services can operate, a key governance issue is which SOA
 2971 services to allow in the ecosystem. The major concern SHOULD be a definition of well-behaved services,
 2972 where the required behavior will likely inherit their characteristics from experiences with distributed
 2973 computing but will also evolve with SOA experience. A major requirement for ensuring well-behaved
 2974 services is collecting sufficient metrics to know how the service affects the SOA infrastructure and
 2975 whether it complies with established infrastructure policies.

2976 Another common concern of service approval is whether there will be duplication of function by multiple
 2977 services. Some governance models talk to a tightly controlled environment where a primary concern is to
 2978 avoid any service duplication. Other governance models talk to a market of services where the
 2979 consumers have wide choices. For the latter, it is anticipated that the better services will emerge from
 2980 market consensus and the availability of alternatives will drive innovation.

2981 It is likely that some combination of control and openness will emerge, possibly with a different
2982 appropriate balance for different categories of use. For SOA governance, the issue is less which services
2983 are approved but rather ensuring that sufficient description is available to support informed decisions for
2984 appropriate use. Thus, SOA governance SHOULD concentrate on identifying the required attributes to
2985 adequately describe a service, the required target values of the attributes, and the standards for defining
2986 the meaning of the attributes and their target values. Governance may also specify the processes by
2987 which the attribute values are measured and the corresponding certification that some realized attribute
2988 set may imply.

2989 For example, unlimited access for using a service may require a degree of life cycle maturity that has
2990 demonstrated sufficient testing over a certain size community. Alternately, the policy may specify that a
2991 service in an earlier phase of its life cycle may be made available to a smaller, more technically
2992 sophisticated group in order to collect the metrics that would eventually allow the service to advance its
2993 life cycle status.

2994 This aspect of governance is tightly connected to description because, given a well-behaved set of
2995 services, it is the responsibility of the consumer (or policies promulgated by the consumer's organization)
2996 to decide whether a service is sufficient for that consumer's intended use. The goal is to avoid global
2997 governance specifying criteria that are too restrictive or too lax for the local needs of which global
2998 governance has little insight.

2999 Such an approach to specifying governance allows independent domains to describe services in local
3000 terms while still having the services available for informed use across domains. In addition, changes to
3001 the attribute sets within a domain can be similarly described, thus supporting the use of newly described
3002 resources with the existing ones without having to update the description of all the legacy content.

3003 **5.1.3.2.3 Governance of Participant Interaction**

3004 Finally, given a reliable services infrastructure and a predictable set of services, the third aspect of
3005 governance is prescribing what is required during a service interaction.

3006 Governance would specify adherence to service interface and service reachability parameters and would
3007 require that the result of an interaction MUST correspond to the real world effects as contained in the
3008 service description. Governance would ensure preconditions for service use are satisfied, in particular
3009 those related to security aspects such as user authentication, authorization, and non-repudiation. If
3010 conflicts arise, governance would specify resolution processes to ensure appropriate agreements,
3011 policies, and conditions are met.

3012 It would also rely on sufficient monitoring by the SOA infrastructure to ensure services remain well-
3013 behaved during interactions, e.g. do not use excessive resources or exhibit other prohibited behavior.
3014 Governance would also require that policy agreements as documented in the execution context for the
3015 interaction are observed and that the results and any after effects are consistent with the agreed policies.
3016 It is likely that in this area the governance will focus on more contractual and legal aspects rather than the
3017 precursor descriptive aspects. SOA governance may prescribe the processes by which SOA-specific
3018 policies are allowed to change, but there are likely more business-specific policies that will be governed
3019 by processes outside SOA governance.

3020 **5.1.3.3 Overarching Governance Concerns**

3021 There are numerous governance related concerns whose effects span the three areas just discussed.
3022 One is the area of standards, how these are mandated, and how the mandates may change. The Web
3023 Services standards stack is an example of relevant standards where a significant number are still under
3024 development. In addition, while there are notional scenarios that guide what standards are being
3025 developed, the fact that many of these standards do not yet exist precludes operational testing of their
3026 adequacy or effectiveness as a necessary and sufficient set.

3027 That said, standards are critical to creating a SOA ecosystem where SOA services can be introduced,
3028 used singularly, and combined with other services to deliver complex business functionality. As with
3029 other aspects of SOA governance, the Governance Body should identify the minimum set felt to be
3030 needed and rigorously enforce that that set be used where appropriate. The Governance Body must take
3031 care to expand and evolve the mandated standards in a predictable manner and with sufficient technical

3032 guidance that new services will be able to coexist as much as possible with the old, and changes to
3033 standards do not cause major disruptions.

3034 Another area that may see increasing activity as SOA expands will be additional regulation by
3035 governments and associated legal institutions. New laws are likely that will deal with transactions which
3036 are service based, possibly including taxes on the transactions. Disclosures laws are likely to mandate
3037 certain elements of description so both the consumer and provider act in a predictable environment and
3038 are protected from ambiguity in intent or action. Such laws are likely to spawn rules and regulations that
3039 will influence the metrics collected for evaluation of compliance.

3040 **5.1.3.4 Considerations for SOA Governance**

3041 The Reference Architecture definition of a loosely coupled system is one in which the constraints on the
3042 interactions between components is minimal: sufficient to permit interoperation without additional
3043 constraints that may be an artifact of implementation technology. While governance experience for
3044 standalone systems provides useful guides, we must be careful not to apply constraints that would
3045 preclude the flexibility, agility, and adaptability we expect to realize from a SOA ecosystem.

3046 One of the strengths of SOA is it can make effective use of diversity rather than requiring monolithic
3047 solutions. Heterogeneous organizations can interact without requiring each conforms to uniform tools,
3048 representation, and processes. However, with this diversity comes the need to adequately define those
3049 elements necessary for consistent interaction among systems and participants, such as which
3050 communication protocol, what level of security, which vocabulary for payload content of messages. The
3051 solution is not always to lock down these choices but to standardize alternatives and standardize the
3052 representations through which an unambiguous identification of the alternative chosen can be conveyed.
3053 For example, the URI standard specifies the URI string, including what protocol is being used, what is the
3054 target of the message, and how may parameters be attached. It does not limit the available protocols, the
3055 semantics of the target address, or the parameters that can be transferred. Thus, as with our definition of
3056 loose coupling, it provides absolute constraints but minimizes which constraints it imposes.

3057 There is not a one-size-fits-all governance but a need to understand the types of things governance will
3058 be called on to do in the context of the goals of SOA. It is likely that some communities will initially desire
3059 and require very stringent governance policies and procedures while other will see need for very little.
3060 Over time, best practices will evolve, likely resulting in some consensus on a sensible minimum and,
3061 except in extreme cases where it is demonstrated to be necessary, a loosening of strict governance
3062 toward the best practice mean.

3063 A question of how much governance may center on how much time governance activities require versus
3064 how quickly is the system being governed expected to respond to changing conditions. For large single
3065 systems that take years to develop, the governance process could move slowly without having a serious
3066 negative impact. For example, if something takes two years to develop and the steps involved in
3067 governance take two months to navigate, then the governance can go along in parallel and may not have
3068 a significant impact on system response to changes. Situations where it takes as long to navigate
3069 governance requirements as it does to develop a response are examples where governance may need to
3070 be reevaluated as to whether it facilitates or inhibits the desired results. Thus, the speed at which
3071 services are expected to appear and evolve needs to be considered when deciding the processes for
3072 control. The added weight of governance should be appropriate for overall goals of the application
3073 domain and the service environment.

3074 Governance, as with other aspects of any SOA implementation, should start small and be conceptualized
3075 in a way that keeps it flexible, scalable, and realistic. A set of useful guidelines would include:

- 3076 • Do not hardwire things that will inevitably change. For example, develop a system that uses the
3077 representation of policies rather than code the policies into the implementations.
- 3078 • Avoid setting up processes that demo well for three services without considering how it will work
3079 for 300. Similarly, consider whether the display of status and activity for a small number of
3080 services will also be effective for an operator in a crisis situation looking at dozens of services,
3081 each with numerous, sometimes overlapping and sometimes differing activities.
- 3082 • Maintain consistency and realism. A service solution responding to a natural disaster cannot be
3083 expected to complete a 6-week review cycle but be effective in a matter of hours.

3084 **5.1.4 Architectural Implications of SOA Governance**

3085 The description of SOA governance indicates numerous architectural requirements on the SOA
3086 ecosystem:

- 3087 • Governance is expressed through policies and assumes multiple use of focused policy modules
3088 that can be employed across many common circumstances. This requires the existence of:
 - 3089 ○ descriptions to enable the policy modules to be visible, where the description includes a
3090 unique identifier for the policy and a sufficient, and preferably a machine process-able,
3091 representation of the meaning of terms used to describe the policy, its functions, and its
3092 effects;
 - 3093 ○ one or more discovery mechanisms that enable searching for policies that best meet the
3094 search criteria specified by the service participant; where the discovery mechanism will
3095 have access to the individual policy descriptions, possibly through some repository
3096 mechanism;
 - 3097 ○ accessible storage of policies and policy descriptions, so service participants can access,
3098 examine, and use the policies as defined.
- 3099 • Governance requires that the participants understand the intent of governance, the structures
3100 created to define and implement governance, and the processes to be followed to make
3101 governance operational. This requires the existence of:
 - 3102 ○ an information collection site, such as a Web page or portal, where governance
3103 information is stored and from which the information is always available for access;
 - 3104 ○ a mechanism to inform participants of significant governance events, such as changes in
3105 policies, rules, or regulations;
 - 3106 ○ accessible storage of the specifics of Governance Processes;
 - 3107 ○ SOA services to access automated implementations of the Governance Processes
- 3108 • Governance policies are made operational through rules and regulations. This requires the
3109 existence of:
 - 3110 ○ descriptions to enable the rules and regulations to be visible, where the description
3111 includes a unique identifier and a sufficient, and preferably a machine process-able,
3112 representation of the meaning of terms used to describe the rules and regulations;
 - 3113 ○ one or more discovery mechanisms that enable searching for rules and regulations that
3114 may apply to situations corresponding to the search criteria specified by the service
3115 participant; where the discovery mechanism will have access to the individual
3116 descriptions of rules and regulations, possibly through some repository mechanism;
 - 3117 ○ accessible storage of rules and regulations and their respective descriptions, so service
3118 participants can understand and prepare for compliance, as defined.
 - 3119 ○ SOA services to access automated implementations of the Governance Processes.
- 3120 • Governance implies management to define and enforce rules and regulations. Management is
3121 discussed more specifically in section **Error! Reference source not found.**, but in a parallel to
3122 governance, management requires the existence of:
 - 3123 ○ an information collection site, such as a Web page or portal, where management
3124 information is stored and from which the information is always available for access;
 - 3125 ○ a mechanism to inform participants of significant management events, such as changes
3126 in rules or regulations;
 - 3127 ○ accessible storage of the specifics of processes followed by management.
- 3128 • Governance relies on metrics to define and measure compliance. This requires the existence of:
 - 3129 ○ the infrastructure monitoring and reporting information on SOA resources;
 - 3130 ○ possible interface requirements to make accessible metrics information generated or
3131 most easily accessed by the service itself.

3132 5.2 Security Model

3133 Security is one aspect of confidence – the confidence in the integrity, reliability, and confidentiality of the
3134 system. In particular, security focuses on those aspects of assurance that involve the accidental or malign
3135 intent of other people to damage or compromise trust in the system and on the availability of SOA-based
3136 systems to perform desired capability.

3137 Security

3138 Security concerns the set of mechanisms for ensuring and enhancing trust and confidence in the
3139 SOA ecosystem.

3140 Providing for security for Service Oriented Architecture is somewhat different than for other contexts;
3141 although many of the same principles apply equally to SOA and to other systems. The fact that SOA
3142 embraces crossing ownership boundaries makes the issues involved with moving data more visible.

3143 As well as securing the movement of data within and across ownership boundaries, security often
3144 revolves around *resources*: the need to guard certain resources against inappropriate access – whether
3145 reading, writing or otherwise manipulating those resources. The basic resource model that informs our
3146 discussion is outlined in Section 3.3.3.

3147 Any comprehensive security solution must take into account the people that are using, maintaining and
3148 managing the SOA. Furthermore, the relationships between them must also be incorporated: any security
3149 assertions that may be associated with particular interactions originate in the people that are behind the
3150 interaction.

3151 We analyze security in terms of the social structures that define the legitimate permissions, obligations
3152 and roles of people in relation to the system, and mechanisms that must be put into place to realize a
3153 secure system. The former are typically captured in a series of security policy statements; the latter in
3154 terms of security *guards* that ensure that policies are enforced.

3155 How and when to apply these derived security policy mechanisms is directly associated with the
3156 assessment of the *threat model* and a *security response model*. The threat model identifies the kinds of
3157 threats that directly impact the message and/or application of constraints, and the response model is the
3158 proposed mitigation to those threats. Properly implemented, the result can be an acceptable level of risk
3159 to the safety and integrity of the system.

3160 5.2.1 Secure Interaction Concepts

3161 We can characterize secure interactions in terms of key security concepts [ISO/IEC 27002]:
3162 confidentiality, integrity, authentication, authorization, non-repudiation, and availability. The concepts for
3163 secure interactions are well defined in other standards and publications. The security concepts here are
3164 not defined but rather related to the SOA ecosystem perspective of this reference architecture foundation.

3165 5.2.1.1 Confidentiality

3166 Confidentiality concerns the protection of privacy of participants in their interactions. Confidentiality refers
3167 to the assurance that unauthorized entities are not able to read messages or parts of messages that are
3168 transmitted.

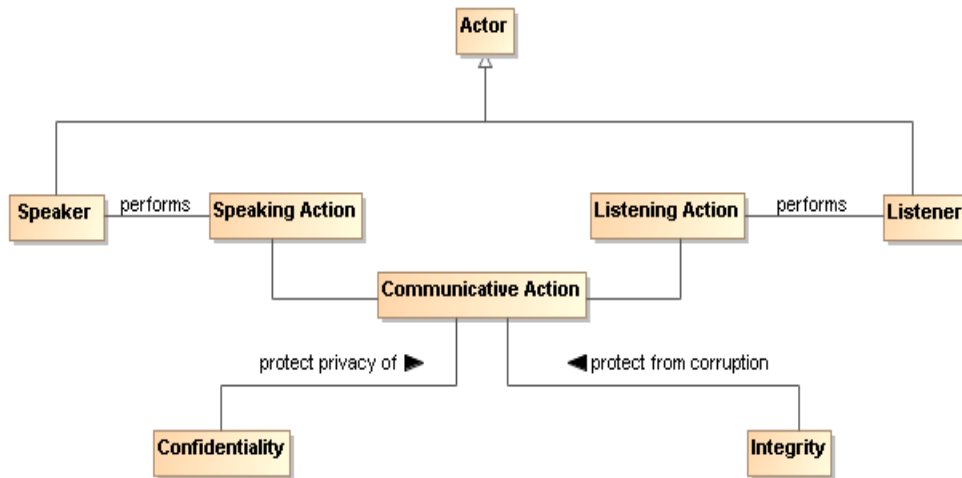
3169 Note that confidentiality has degrees: in a completely confidential exchange, third parties would not even
3170 be aware that a confidential exchange has occurred. In a partially confidential exchange, the identities of
3171 the participants may be known but the content of the exchange obscured.

3172 5.2.1.2 Integrity

3173 Integrity concerns the protection of information that is exchanged – either from unauthorized writing or
3174 inadvertent corruption. Integrity refers to the assurance that information that has been exchanged has not
3175 been altered.

3176 Integrity is different from confidentiality in that messages that are sent from one participant to another
3177 may be obscured to a third party, but the third party may still be able to introduce his own content into the
3178 exchange without the knowledge of the participants.

3179 Figure 54 applies confidentiality and integrity to communicative action, see Section 3.1.3 for a description
3180 of communicative action.




3181
3182 *Figure 54 Confidentiality and Integrity*

3183 The communicative action is the joint action involved in the exchange of messages. Section 5.2.4
3184 describes common computing techniques for providing confidentiality and integrity during message
3185 exchanges.

3186 5.2.1.3 Authentication

3187 Authentication concerns the identity of the participants in an exchange. Authentication refers to the
3188 means by which one participant can be assured of the identity of other participants.

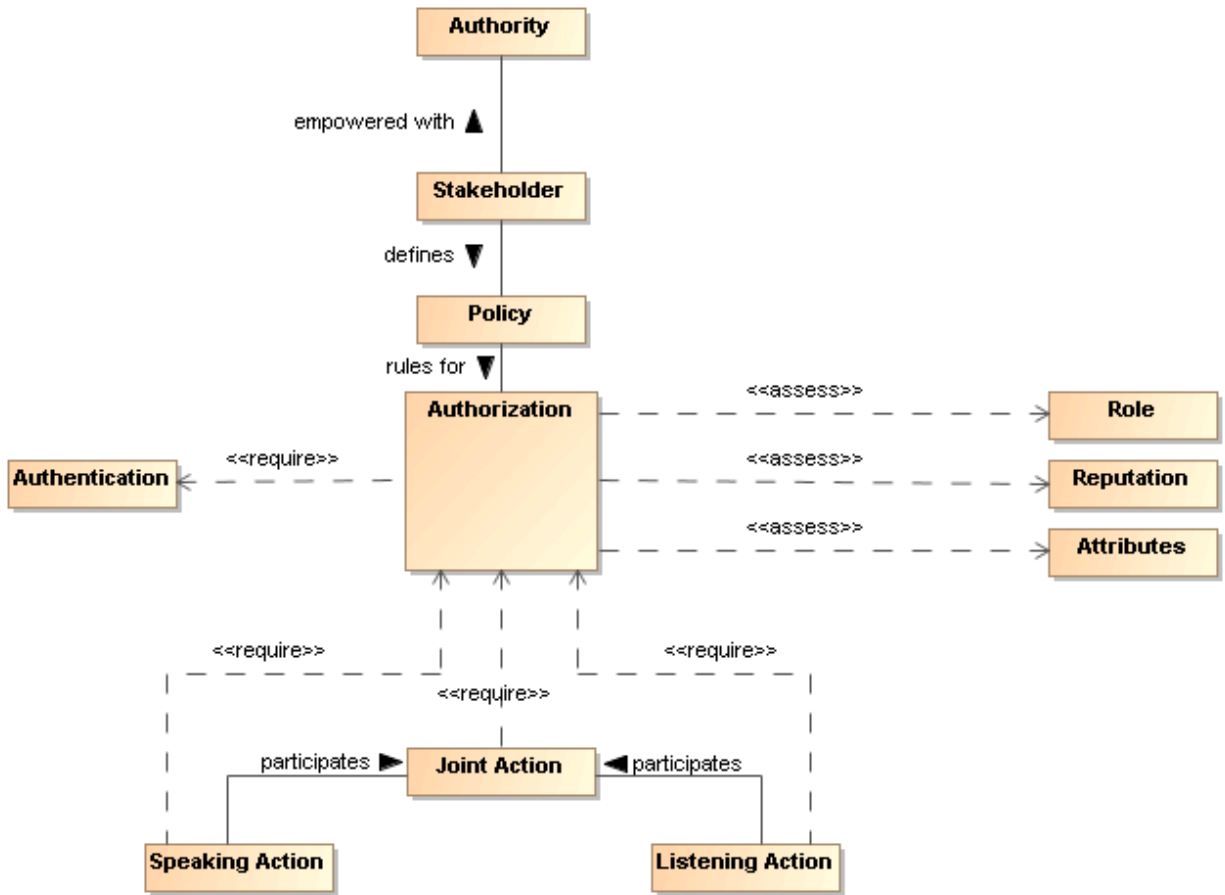
3189 Figure 55 applies authentication to the identity of participants.



3191
3192 *Figure 55 Authentication*

3193 5.2.1.4 Authorization

3194 Authorization concerns the legitimacy of the interaction. Authorization refers to the means by which a
3195 stakeholder may be assured that the information and actions that are exchanged are either explicitly or
3196 implicitly approved.



3197
3198 *Figure 56 Authorization*

3199 The roles and attributes which provide a participant’s credentials are expanded to include reputation.
3200 Reputation often helps determine willingness to interact, for example, reviews of a service provider are
3201 likely to influence the decision to interact with the service provider. The roles, reputation, and attributes
3202 are represented as assertions measured by authorization decision points.

3203 The role of policy for security is to permit stakeholders to express their choices. In Figure 56, a policy is a
3204 written constraint and the role, reputation, and attribute assertions are evaluated according to the
3205 constraints in the authorization policy. A combination of security mechanisms and their control via
3206 explicit policies can form the basis of an authorization solution.

3207 **5.2.1.5 Non-repudiation**

3208 Non-repudiation concerns the accountability of participants. To foster trust in the performance of a system
3209 used to conduct shared activities it is important that the participants are not able to later deny their
3210 actions: to repudiate them. Non-repudiation refers to the means by which a participant may not, at a later
3211 time, successfully deny having participated in the interaction or having performed the actions as reported
3212 by other participants.

3213 **5.2.1.6 Availability**

3214 Availability concerns the ability of systems to use and offer the services for which they were designed.
3215 One of the threats against availability is the so-called denial of service attack in which attackers attempt to
3216 prevent legitimate access to the system.

3217 We differentiate here between general availability – which includes aspects such as systems reliability –
3218 and availability as a security concept where we need to respond to active threats to the system.

3219 **5.2.2 Where SOA Security is Different**

3220 The core security concepts are fundamental to all social interactions. The evolution of sharing
3221 information using a SOA requires the flexibility to dynamically secure computing interactions in a
3222 computing ecosystem where the owning social groups, roles, and authority are constantly changing as
3223 described in section 5.1.3.1.

3224 SOA policy-based security can be more adaptive for a computing ecosystem than previous computing
3225 technologies allow for, and typically involves a greater degree of distributed mechanisms.

3226 Standards for security, as is the case with all aspects of SOA, play a large role in flexible security on a
3227 global scale. SOA security may also involve greater auditing and reporting to adhere to regulatory
3228 compliance established by governance structures.

3229 **5.2.3 Security Threats**

3230 There are a number of ways in which an attacker may attempt to compromise the security of a system.
3231 The two primary sources of attack are third parties attempting to subvert interactions between legitimate
3232 participants and an entity that is participating but attempting to subvert its partner(s). The latter is
3233 particularly important in a SOA where there may be multiple ownership boundaries and trust boundaries.

3234 The threat model lists some common threats that relate to the core security concepts listed in Section
3235 5.2.1. Each technology choice in the realization of a SOA can potentially have many threats to consider.

3236 **Message alteration**

3237 If an attacker is able to modify the content (or even the order) of messages that are exchanged
3238 without the legitimate participants being aware of it then the attacker has successfully
3239 compromised the security of the system. In effect, the participants may unwittingly serve the
3240 needs of the attacker rather than their own.

3241 An attacker may not need to completely replace a message with his own to achieve his objective:
3242 replacing the identity of the beneficiary of a transaction may be enough.

3243 **Message interception**

3244 If an attacker is able to intercept and understand messages exchanged between participants,
3245 then the attacker may be able to gain advantage. This is probably the most commonly understood
3246 security threat.

3247 **Man in the middle**

3248 In a man-in-the-middle attack, the legitimate participants believe that they are interacting with
3249 each other; but are in fact interacting with the attacker. The attacker attempts to convince each
3250 participant that he is their correspondent; whereas in fact he is not.

3251 In a successful man-in-the-middle attack, legitimate participants will often not have a true
3252 understanding of the state of the other participants. The attacker can use this to subvert the
3253 intentions of the participants.

3254 **Spoofing**

3255 In a spoofing attack, the attacker convinces a participant that he is really someone else –
3256 someone that the participant would normally trust.

3257 **Denial of service attack**

3258 In a denial of service (DoS) attack, the attacker attempts to prevent legitimate users from making
3259 use of the service. A DoS attack is easy to mount and can cause considerable harm: by
3260 preventing legitimate interactions, or by slowing them down enough, the attacker may be able to
3261 simultaneously prevent legitimate access to a service and to attack the service by another
3262 means.

3263 A variation of the DoS attack is the Distributed Denial of Service attack. In a DDoS attack the
3264 attacker uses multiple agents to the attack the target. In some circumstances this can be
3265 extremely difficult to counteract effectively.

3266 One of the features of a DoS attack is that it does not require valid interactions to be effective:
3267 responding to invalid messages also takes resources and that may be sufficient to cripple the
3268 target.

3269 **Replay attack**

3270 In a replay attack, the attacker captures the message traffic during a legitimate interaction and
3271 then replays part of it to the target. The target is persuaded that a similar transaction to the
3272 previous one is being repeated and it will respond as though it were a legitimate interaction.

3273 A replay attack may not require that the attacker understand any of the individual
3274 communications; the attacker may have different objectives (for example attempting to predict
3275 how the target would react to a particular request).

3276 **False repudiation**

3277 In false repudiation, a user completes a normal transaction and then later attempts to deny that
3278 the transaction occurred. For example, a customer may use a service to buy a book using a credit
3279 card; then, when the book is delivered, refuse to pay the credit card bill claiming that *someone*
3280 *else* must have ordered the book.

3281 **5.2.4 Security Responses**

3282 Security goals are never absolute: it is not possible to guarantee 100% confidentiality, non-repudiation,
3283 etc. However, a well designed and implemented security response model can ensure acceptable levels of
3284 security risk. For example, using a well-designed cipher to encrypt messages may make the cost of
3285 breaking communications so great and so lengthy that the information obtained is valueless.

3286 Performing threat assessments, devising mitigation strategies, and determining acceptable levels of risk
3287 are the foundation for an effective process to mitigating threats in a cost-effective way.²⁰ The choice in
3288 hardware and software to realize a SOA will be the basis for threat assessments and mitigation
3289 strategies. The stakeholders of a specific SOA implementation should determine acceptable levels of risk
3290 based on threat assessments and the cost of mitigating those threats.

3291 **5.2.4.1 Privacy Enforcement**

3292 The most efficient mechanism to assure confidentiality is the encryption of information. Encryption is
3293 particularly important when messages must cross trust boundaries; especially over the Internet. Note that
3294 encryption need not be limited to the content of messages: it is possible to obscure even the existence of
3295 messages themselves through encryption and 'white noise' generation in the communications channel.

3296 The specifics of encryption are beyond the scope of this architecture. However, we are concerned about
3297 how the connection between privacy-related policies and their enforcement is made.

3298 A policy enforcement point for enforcing privacy may take the form of an automatic function to encrypt
3299 messages as they leave a trust boundary; or perhaps simply ensuring that such messages are suitably
3300 encrypted.

3301 Any policies relating to the level of encryption being used would then apply to these centralized
3302 messaging functions.

3303 **5.2.4.2 Integrity Protection**

3304 To protect against message tampering or inadvertent message alteration, and to allow the receiver of a
3305 message to authenticate the sender, messages may be accompanied by a digital signature. Digital

²⁰ 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.

3306 signatures provide a means to detect if signed data has been altered. This protection can also extend to
3307 authentication and non-repudiation of a sender.

3308 A common way a digital signature is generated is with the use of a private key that is associated with a
3309 public key and a digital certificate. The private key of some entity in the system is used to create a digital
3310 signature for some set of data. Other entities in the system can check the integrity of the signed data set
3311 via signature verification algorithms. Any changes to the data that was signed will cause signature
3312 verification to fail, which indicates that integrity of the data set has been compromised.

3313 A party verifying a digital signature must have access to the public key that corresponds to the private key
3314 used to generate the signature. A digital certificate contains the public key of the owner, and is itself
3315 protected by a digital signature created using the private key of the issuing Certificate Authority (CA).

3316 **5.2.4.3 Message Replay Protection**

3317 To protect against replay attacks, messages may contain information that can be used to detect replayed
3318 messages. The simplest requirement to prevent replay attacks is that each message that is ever sent is
3319 unique. For example, a message may contain a message ID, a timestamp, and the intended destination.

3320 By storing message IDs, and comparing each new message with the store, it becomes possible to verify
3321 whether a given message has been received before (and therefore should be discarded).

3322 The timestamp may be included in the message to help check for message freshness. Messages that
3323 arrive after their message ID could have been cleared (after receiving the same message some time
3324 previously) may also have been replayed. A common means for representing timestamps is a useful part
3325 of an interoperable replay detection mechanism.

3326 The destination information is used to determine if the message was misdirected or replayed. If the
3327 replayed message is sent to a different endpoint than the destination of the original message, the replay
3328 could go undetected if the message does not contain information about the intended destination.

3329 In the case of messages that are replies to prior messages, it is also possible to include seed information
3330 in the prior messages that is randomly and uniquely generated for each message that is sent out. A
3331 replay attack can then be detected if the reply does not embed the random number that corresponds to
3332 the original message.

3333 **5.2.4.4 Auditing and Logging**

3334 False repudiation involves a participant denying that it authorized a previous interaction. An effective
3335 strategy for responding to such a denial is to maintain careful and complete logs of interactions which can
3336 be used for auditing purposes. The more detailed and comprehensive an audit trail is, the less likely it is
3337 that a false repudiation would be successful.

3338 The countermeasures assume that the non-repudiation tactic (e.g. digital signatures) is not undermined
3339 itself. For example, if private key is stolen and used by an adversary, even extensive logging cannot
3340 assist in rejecting a false repudiation.

3341 Unlike many of the security responses discussed here, it is likely that the scope for automation in
3342 rejecting a repudiation attempt is limited to careful logging.

3343 **5.2.4.5 Graduated engagement**

3344 The key to managing and responding to DoS attacks is to be careful in the use of resources when
3345 responding to interaction. Put simply, a system has a choice to respond to a communication or to ignore
3346 it. In order to avoid vulnerability to DoS attacks a service provider should be careful not to commit
3347 resources beyond those implied by the current state of interactions; this permits a graduation in
3348 commitment by the service provider that mirrors any commitment on the part of service consumers and
3349 attackers alike.

3350 **5.2.5 Architectural Implications of SOA Security**

3351 Providing SOA security in an ecosystem of governed services has the following implications on the policy
3352 support and the distributed nature of mechanisms used to assure SOA security:

- 3353 • Security expressed through policies have the same architectural implications as described in
3354 Section 4.4.4 for policies and contracts architectural implications.
- 3355 • Security policies require mechanisms to support security description administration, storage, and
3356 distribution.
- 3357 • Service descriptions supporting security policies should:
 - 3358 ○ have a meta-structure sufficiently rich to support security policies;
 - 3359 ○ be able to reference one or more security policy artifacts;
 - 3360 ○ have a framework for resolving conflicts between security policies.
- 3361 • The mechanisms that make-up the execution context in secure SOA-based systems should:
 - 3362 ○ provide protection of the confidentiality and integrity of message exchanges;
 - 3363 ○ be distributed so as to provide centralized or decentralized policy-based identification,
3364 authentication, and authorization;
 - 3365 ○ ensure service availability to consumers;
 - 3366 ○ be able to scale to support security for a growing ecosystem of services;
 - 3367 ○ be able to support security between different communication technologies;
- 3368 • Common security services include:
 - 3369 ○ services that abstract encryption techniques;
 - 3370 ○ services for auditing and logging interactions and security violations;
 - 3371 ○ services for identification;
 - 3372 ○ services for authentication;
 - 3373 ○ services for authorization;
 - 3374 ○ services for intrusion detection and prevention;
 - 3375 ○ services for availability including support for quality of service specifications and metrics.

3376 5.3 Management Model

3377 Management

3378 Management is the control of the use, configuration, and availability of resources in accordance
3379 with the policies of the stakeholders involved.

3380 There are three separate but linked domains of interest within the management of SOA-based systems.
3381 The first and most obvious is the management and support of the resources that are involved in any
3382 complex system – of which SOA-based systems are excellent examples. The second is the promulgation
3383 and enforcement of the policies and contracts agreed to by the stakeholders in SOA-based systems. The
3384 third domain is the management of the relationships of the participants in SOA-based systems – both to
3385 each other and to the services that they use and offer.

3386 There are many artifacts in a large system that may need management. As soon as there is the possibility
3387 of more than one instance of a thing, the issue of managing those things becomes relevant. Historically,
3388 systems management capabilities have been organized by the following functional groups known as
3389 “FCAPS” functions (based on ITU-T Rec. M.3400 (02/2000), “TMN Management Functions”): Fault
3390 management, configuration management, account management, performance and security management.

3391 In the context of SOA we see many possible resources that may require management: services, service
3392 descriptions, service capabilities, policies, contracts, roles, relationships, security, and infrastructure
3393 elements. In addition, given the ecosystem nature of SOA, it is also potentially necessary to manage the
3394 business relationships between participants in the SOA.


3395 Managing systems that may be used across ownership boundaries raises issues that are not normally
3396 present when managing a system within a single ownership domain. For example, care is required
3397 managing a service when the owner of the service, the provider of the service, the host of the service and
3398 access mediators to the service may all belong to different stakeholders. In addition, it may be important

3399 to allow service consumers to communicate their requirements to the service provider so that they are
 3400 satisfied in a timely manner.

3401 A given service may be provided and consumed in more than one version. Version control of services is
 3402 important both for service providers and service consumers (who may need to ensure certainty in the
 3403 version of the service they are interacting with).

3404 In fact, managing a service has quite a few similarities to using a service: suggesting that we can use the
 3405 service oriented model to manage SOA-based systems as well as provide them. A management service
 3406 would be distinguished from a non-management service more by the nature of the capabilities involved
 3407 (i.e., capabilities that relate to managing services) than by any intrinsic difference.

3408 In this model, we show how the SOA framework may apply to managing services as well as using and
 3409 offering them. There are, of course, some special considerations that apply to service management which
 3410 we bring out: namely that we will be managing the life-cycle of services, managing any service level
 3411 attributes, managing dependencies between services and so on.



3412
 3413 *Figure 57 Managing resources in a SOA*

3414 The core concept in management is that of a manageability capability:

3415 **Manageability Capability**

3416 The manageability capability of a resource is the capability that allows it to be managed with
 3417 respect to some property. Note that manageability capabilities are not necessarily part of the
 3418 managed entities themselves.

3419 Manageability capabilities are the core resources that management systems use to manage:
 3420 each resource that may be managed in some way has a number of aspects that may be
 3421 managed. For example, a service’s life-cycle may be manageable, as may its Quality of Service
 3422 parameter; a policy may also be managed for life-cycle but Quality of Service would not normally
 3423 apply.

3424 **Life-cycle manageability**

3425 A manageability capability associated with a resource that permits the life cycle of the resource to
 3426 be managed. As noted above, the life-cycle manageability capability of a resource is unlikely to
 3427 reside within the resource itself (you cannot tell a system that is not running to start itself).

3428 The life-cycle management of a resource typically refers to how the resource is created, how it is
 3429 destroyed and what dependencies there might exist that must be simultaneously managed.

3430 **Configuration manageability**

3431 A capability that permits the configuration of resources to be managed. Service configuration, in
 3432 particular, may be complex in cases where there are dependencies between services and other
 3433 resources.

3434 **Event monitoring manageability**

3435 Managing the reporting of events and faults is one of the key lower-level manageability
 3436 capabilities.

3437 **Accounting manageability**

3438 A capability associated with resources that allows for the use of those resources to be measured
 3439 and accounted for. This implies that not only can the *use* of resources be properly measured, but
 3440 also that those *using* those resources also be properly identified.

3441 Accounting for the use of resources by participants in the SOA supports the proper budgeting and
3442 allocation of funding by participants.

3443 **Quality of service manageability**

3444 A manageability capability associated with a resource that permits any quality of service
3445 associated with the resource to be managed. Classic examples of this include bandwidth
3446 requirements and offerings associated with a service.

3447 **Business performance manageability**

3448 A manageability capability that is associated with services that permits the service's business
3449 performance to be monitored and managed. In particular, if there are business-level service level
3450 agreements that apply to a service, being able to monitor and manage those SLAs is an
3451 important role for management systems.

3452 Building support for arbitrary business monitoring is likely to be challenging. However, given a
3453 *measure* for determining a service's compliance to business service level agreements,
3454 management systems can monitor that performance in a way that is entirely similar to other
3455 management tasks.

3456 **Policy manageability**

3457 Where the policies associated with a resource may be complex and dynamic, so those policies
3458 themselves may require management. The ability to manage those policies (such as
3459 promulgating policies, retiring policies and ensuring that policy decision points and enforcement
3460 points are current) is a management function.

3461 In the particular case of policies, there is a special relationship between management and
3462 policies. Just like other artifacts, policies require management in a SOA. However, much of
3463 management is about *applying* policies also: where governance is often about what the policies
3464 regarding artifacts and services should be, a key management role is to ensure that those
3465 policies are consistently applied.

3466 **Management service**

3467 A management service is a service that manages other services and resources.

3468 **Management Policy**

3469 A management policy is a policy whose topic is a management topic. Just as with other aspects
3470 of a SOA, the management of resources within the SOA may be governed by management
3471 policies, contracts (such as SLAs).

3472 In a deployed system, it may well be that different aspects of the management of a given service are
3473 managed by different management services. For example, the life-cycle management of services often
3474 involves managing dependencies between services and resource requirements. Managing quality of
3475 service is often very specific to the service itself; for example, quality of service attributes for a video
3476 streaming service are quite different to those for a banking system.

3477 There are additional concepts of management that often also apply to IT management:

3478 **Systems management**

3479 Systems management refers to enterprise-wide maintenance and administration of distributed
3480 computer systems.

3481 **Network management**

3482 Network management refers to the maintenance and administration of large-scale networks such
3483 as computer networks and telecommunication networks. Systems and network management
3484 execute a set of functions required for controlling, planning, deploying, coordinating, and
3485 monitoring the distributed computer systems and the resources of a network.

3486 However, for the purposes of this Reference Architecture, while recognizing their importance, we do not
3487 focus on systems management or network management.

3488 - the specific identifier is not prescribed by this Reference Architecture but the structure and semantics of
3489 the identifier must be indicated for the identifier value to be properly used. For example, part of identity
3490 may include version identification.

3491 For this, the configuration management plan or similar document from which the version number is
3492 derived must be identified.

3493

3494 **5.3.1 Management and Governance**

3495 The primary role of governance in the context of SOA is to allow the stakeholders in the SOA to be able
3496 to negotiate and set the key policies that govern the running of the system. Recall that in an ecosystems
3497 perspective, the goal is less to have complete fine-grained control but more to enable the individual
3498 participants to work together. Policies that are set at the governance of a SOA will tend to focus on the
3499 rules of engagement between participants – what kind of interacts are permissible, how to resolve
3500 disputes, and so on.

3501 While governance may be primarily focused on setting policies, management is more focused on
3502 realization and enforcement of policies.

3503 **5.3.2 Management Contracts and Policies**

3504 As we noted above, management can often be viewed as the application of contracts and policies to
3505 ensure the smooth running of the SOA. Policies play an important part in managing systems both as
3506 artifacts that need to be managed and as the guiding constraints to determine how the SOA should be
3507 managed.

3508 **5.3.2.1 Policies**

3509 "Although provision of management capabilities enables a service to become manageable, the extent and
3510 degree of permissible management are defined in management policies that are associated with the
3511 services. Management policies are used to define the obligations for, and permissions to, managing the
3512 service." **[WSA]**

3513 On the other hand, a policy without any means of enforcing it is vacuous. In the case of management
3514 policy, we rely on a management infrastructure to realize and enforce management policy.

3515 **5.3.3 Management Infrastructure**

3516 In order for a service or other resource to be manageable there must be a corresponding manageability
3517 capability that can effect that management. The particulars of this capability will vary somewhat
3518 depending on the nature of the capability. For example, a service life-cycle manageability capability
3519 requires the ability to start a service, to stop the service, and potentially to pause the service. Conversely,
3520 in order to manage document-like artifacts, such as service descriptions, the capability of storing the
3521 artifacts, controlling access to those artifacts, allowing updates of the artifacts to be deployed are all
3522 important capabilities for managing them.

3523

3524 Elements of a basic service management infrastructure should include the following characteristics:

3525

- 3526 • Integrate with existing security services
- 3527 • Monitoring
- 3528 • Heartbeat and Ping
- 3529 • Alerting
- 3530 • Pause/Restore/Restart Service Access
- 3531 • Logging, Auditing, Non-Repudiation

- 3532 • Runtime Version Management
- 3533 • Complement other infrastructure services (discovery, messaging, mediation)
- 3534
- 3535 * Message Routing and Redirection
- 3536 * Failover
- 3537 * Load-balancing
- 3538
- 3539 * QoS, Management of Service Level Objects and Agreements
- 3540 * Availability
- 3541 * Response Time
- 3542 * Throughput
- 3543
- 3544 • Fault and Exception Management
- 3545

3546 **5.3.4 Service Life-cycle**

3547 Managing a service's life cycle involves managing the establishment of the service, managing its steady-
3548 state performance, and managing its termination. The most obvious feature of this is that a service cannot
3549 manage its own life cycle (imagine asking a non-functioning service to start). Another important
3550 consideration is that services may have resource requirements that must be established at various points
3551 in the services' life cycles. These dependencies may take the form of other services being established;
3552 possibly even services that are not exposed by the service's own interface.

3553
3554

3555 **5.4 SOA Testing Model**

3556
3557
3558

*Program testing can be used to show the presence of bugs,
but never to show their absence!*
Edsger Dijkstra

3559 Testing for SOA combines the typical challenges of software testing and certification with the additional
3560 needs of accommodating the distributed nature of the resources, the greater access of a more
3561 unbounded consumer population, and the desired flexibility to create new solutions from existing
3562 components over which the solution developer has little if any control. The purpose of testing is to
3563 demonstrate a required level of reliability, correctness, and effectiveness that enable prospective
3564 consumers to have adequate confidence in using a service. Adequacy is defined by the consumer based
3565 on the consumer's needs and context of use. As the Dijkstra quote points out, absolute correctness and
3566 completeness cannot be proven by testing; however, for SOA, it is critical for the prospective consumer to
3567 know what testing has been performed, how it has been performed, and what were the results.

3568 **5.4.1 Traditional Software Testing as Basis for SOA Testing**

3569 SOA services are largely software artifacts and can leverage the body of experience that has evolved
3570 around software testing. IEEE-829 specifies the basic set of software test documents while allowing
3571 flexibility for tailored use. As such, the document structure can also provide guidance to SOA testing.

3572 IEEE-829 covers test specification and test reporting through use of the following document types:

- 3573 • *Test plan* documenting the scope (what will be tested, both which entity and what features of the
3574 entity), the approach (how it will be tested), and the needed resources (who will do the testing, for
3575 how long), with details contained in the:

- 3576 • *Test design specification*: features to be tested, test conditions (e.g. test cases, test procedures
3577 needed) and expected results (criteria for passing test); entrance and exit criteria
- 3578 • *Test case specification*: test data used for input and expected output
- 3579 • *Test procedure specification*: steps required to run the test, including any set-up preconditions
- 3580 • *Test item transmittal* to identify the test items being transmitted for testing
- 3581 • *Test log* to record what occurred during test, i.e. which tests run, who ran, what order, what happened
- 3582 • *Test incident report* to capture any event that happened during test which requires further
3583 investigation
- 3584 • *Test summary* as a management report summarizing test run and results, conclusions
- 3585 In summary, IEEE-829 captures (1) what was tested, (2) how it was tested, e.g. the test procedure used,
3586 and (3) the results of the test.

3587 **5.4.1.1 Types of Testing**

3588 There are numerous aspects of testing that, in total, work to establish that an entity is (1) built as required
3589 per policies and related specifications prescribed by the entity's owner, and (2) delivers the functionality
3590 required by its intended users. This is often referred to as verification and validation.

3591 Policies, as described in Section 4.4, that are related to testing may prescribe but are not limited to the
3592 business processes to be followed, the standards with which an implementation must comply, and the
3593 qualifications of and restrictions on the users. In addition to the functional requirements prescribing what
3594 an entity does, there may also be non-functional performance and/or quality metrics that state how well
3595 the entity does it. The relation of these policies to SOA testing is discussed further below.

3596 The identification of policies is the purview of governance (section 5.1) and the assuring of compliance
3597 (including response to noncompliance) with policies is a matter for management (section **Error!**
3598 **Reference source not found.**).

3599 **5.4.1.2 Range of Test Conditions**

3600 Test conditions and expected responses are detailed in the test case specification. The test conditions
3601 should be designed to cover the areas for which the entity's response must be documented and may
3602 include:

- 3603 • nominal conditions
- 3604 • boundaries and extremes of expected conditions
- 3605 • breaking point where the entity has degraded below a certain level or has otherwise ceased
3606 effective functioning
- 3607 • random conditions to investigate unidentified dependencies among combinations of conditions
- 3608 • errors conditions to test error handling

3609 The specification of how each of these conditions should be tested for SOA resources, including the
3610 infrastructure elements of the SOA ecosystem, is beyond the scope of this Reference Architecture but is
3611 an area that will evolve along with operational SOA experience.

3612 **5.4.1.3 Configuration Management of Test Artifacts**

3613 The test item transmittal provides an unambiguous identification of the entity being tested, thus
3614 REQUIRING that the configuration of the entity is appropriately tracked and documented. In addition, the
3615 test documents (such as those specified by IEEE-829) MUST also be under a documented and
3616 appropriately audited configuration management process, as should other resources used for testing.
3617 The description of each artifact would follow the general description model as discussed in section
3618 4.1.1.1; in particular, it would include a version number for the artifact and reference to the documentation
3619 describing the versioning scheme from which the version number is derived.

3620

3621 [EDITOR'S NOTE: TO WHAT EXTENT SHOULD CM BE EXPLICITLY INCLUDED IN THE MANAGEMENT
3622 SECTION?]

3623 **5.4.2 Testing and the SOA Ecosystem**

3624 [EDITOR'S NOTE: THE EMPHASIS THOUGH MUCH OF THE RA IS THE LARGER ECOSYSTEM BUT WE NEED
3625 WORDS IN SECTION 3 TO ACKNOWLEDGE THE EXISTENCE OF THE ENTERPRISE AND THAT AN
3626 ENTERPRISE (AS COMMONLY INTERPRETED) IS LIKELY MORE CONSTRAINED AND MORE PRECISELY
3627 DESCRIBED FOR THE CONTEXT OF THE ENTERPRISE. THE ECOSYSTEM PERSPECTIVE, THOUGH, IS
3628 STILL APPLICABLE FOR THE FOLLOWING REASONS:

- 3629
- 3630 1. A GIVEN ENTERPRISE MAY COMPRISE NUMEROUS CONSTITUENT ENTERPRISES THAT
3631 RESEMBLE THE INDEPENDENT ENTITIES DESCRIBED FOR THE ECOSYSTEM. AN ENTERPRISE
3632 MAY ATTEMPT TO REDUCE VARIATIONS AMONG THE CONSTITUENTS BUT THE ECOSYSTEM VIEW
3633 ENABLES SOA TO BENEFIT THE ENTERPRISE WITHOUT REQUIRING THE ENTERPRISE ISSUES TO
3634 BE FULLY RESOLVED.
 - 3635 2. RESOURCES SPECIFICALLY MOTIVATED BY THE CONTEXT OF THE ENTERPRISE CAN BE MORE
3636 READILY USED IN A DIFFERENT CONTEXT IF ECOSYSTEM CONSIDERATIONS ARE INCLUDED AT
3637 AN EARLY STAGE. THE CHANGE IN A CONTEXT MAY BE A FUNDAMENTAL CHANGE IN THE
3638 ENTERPRISE OR THE NEWLY DISCOVERED APPLICABILITY OF ENTERPRISE RESOURCES TO USE
3639 OUTSIDE THE ENTERPRISE.

3640

3641 IN THIS REFERENCE ARCHITECTURE, REFERENCE TO THE SOA ECOSYSTEM APPLIES BUT WITH
3642 POSSIBLY LESS GENERALITY TO AN ENTERPRISE USE OF SOA.]

3643 Testing of SOA artifacts for use in the SOA ecosystem differs from traditional software testing for several
3644 reasons. First, a highly touted benefit of SOA is to enable unanticipated consumers to make use of
3645 services for unanticipated purposes. Examples of this could include the consumer using a service for a
3646 result that was not considered the primary one by the provider, or the service may be used in combination
3647 with other services in a scenario that is different from the one considered when designing for the initial
3648 target consumer community. It is unlikely that a new consumer will push the services back to anything
3649 resembling the initial test phase to test the new use, and thus additional paradigms for testing are
3650 necessary. Some testing may depend on the availability of test resources made available as a service
3651 outside the initial test community, while some testing is likely to be done as part of limited use in the
3652 operational setting. The potential responsibilities related to such "consumer testing" is discussed further
3653 below.

3654 Secondly, in addition to consumers who interact with a service to realize the described real world effects,
3655 the developer community is also intended to be a consumer. In the SOA vision of reuse, the developer
3656 will compose new solutions using existing services, where the existing services provides access to some
3657 desired real world effects that are needed by the new solution. The new solution is a consumer of the
3658 existing services, enabling repeated interactions with the existing services playing the role of reusable
3659 components. Note, those components are used at the locations where they individually reside and are not
3660 typically duplicated for the new solution. The new solution may itself be offered as a SOA service, and a
3661 consumer of the service composition representing the new solution may be totally unaware of the
3662 component services being used. (See section 4.3.4 for further discussion on service compositions.)

3663 Another difference from traditional testing is that the distributed, unbounded nature of the SOA ecosystem
3664 makes it unlikely to have an isolated test environment that duplicates the operational environment. A
3665 traditional testing approach often makes use of a test system that is identical to the eventual operational
3666 system but isolated for testing. After testing is successfully completed, the tested entity would be
3667 migrated to the operational environment, or the test environment may be delivered as part of the system
3668 to become operational. This is not feasible for the SOA ecosystem as a whole.

3669 SOA services must be testable in the environment and under the conditions that can be encountered in
3670 the operational SOA ecosystem. As the ecosystem is in a state of constant change, so some level of
3671 testing is continuous through the lifetime of the service, leveraging utility services used by the ecosystem
3672 infrastructure to monitor its own health and respond to situations that could lead to degraded
3673 performance. This implies the test resources must incorporate aspects of the SOA paradigm, and a

3674 category of services may be created to specifically support and enable effective monitoring and
3675 continuous testing for resources participating in the SOA ecosystem.

3676 While SOA within an enterprise may represent a more constrained and predictable operational
3677 environment, the composability and unanticipated use aspects are highly touted within the enterprise.
3678 The expanded perspective on testing may not be as demanding within an enterprise but fuller
3679 consideration of the ecosystem enables the enterprise to be more responsive should conditions change.

3680 **5.4.3 Elements of SOA Testing**

3681 IEEE-829 identifies fundamental aspects of testing, and many of these should carry over to SOA testing:
3682 in particular, the identification of what is to be tested, how it is to be tested, and by whom the testing is to
3683 be done. While IEEE-829 identifies a suggested document tree, the availability of these documents in the
3684 SOA ecosystem is an additional matter of concern that will be discussed below.

3685 **5.4.3.1 What is to be Tested**

3686 The focus of this discussion is the SOA service. It is recognized that the infrastructure components of
3687 any SOA environment are likely to also be SOA services and, as such, will fall under the same testing
3688 guidance. Other resources that contribute to a SOA environment may not be SOA services, but will be
3689 expected to satisfy the intent if not the letter of guidance presented here. Specific differences for such
3690 resources are as yet largely undefined and further elaboration is beyond the scope of this Reference
3691 Architecture.

3692 The following discussion often focuses on a singular SOA service but it is implicit that any service may be
3693 a composite of other services. As such, testing the functionality of a composite service may effectively be
3694 testing an end-to-end business process that is being provided by the composite service. If new versions
3695 are available for the component services, appropriate end-to-end testing of the composite may be
3696 required in order to verify that the composite functionality is still adequately provided. The level of
3697 required testing of an updated composite will depend on policies of those providing the service, policies of
3698 those using the service, and mission criticality of those depending on the service results.

3699 The SOA service to be tested MUST be unambiguously identified as specified by its applicable
3700 configuration management scheme. Specifying such a scheme is beyond the scope of this Reference
3701 Architecture other than to say the scheme should be documented and itself under configuration
3702 management.

3703 **5.4.3.1.1 Origin of Test Requirements**

3704 In the Service Description model (Figure 21), the aspects of a service that need to be described are:

- 3705 • the service functionality and technical assumptions that underlie the functionality;
- 3706 • the policies that describe conditions of use;
- 3707 • the service interface that defines information exchange with the service;
- 3708 • service reachability that identifies how and where message exchange is to occur; and
- 3709 • metrics access for any participant to have information on how a service is performing.

3710 Service testing must provide adequate assurance that each of these aspects is operational as defined.

3711 The information in the service description comes from different sources. The functionality is defined
3712 through whatever process identifies needs and the community for which these needs will be addressed.
3713 The process may be ad hoc as serves the prospective service owner or strictly governed, but defining the
3714 functionality is an essential first step in development. It is also an early and ongoing focus of testing to
3715 ensure the service accurately reflects the described functionality and the described functionality
3716 accurately addresses the consumer needs.

3717 Policies define the conditions of development and conditions of use for a service and are typically
3718 specified as part of the governance process. Policies constraining service development, such as coding
3719 standards and best practices, require appropriate testing and auditing during development to ensure
3720 compliance. While the governance process will identify development policies, these are likely to originate
3721 from the technical community responsible for development activities. Policies that define conditions of

3722 use often define business practices that service owners and providers or those responsible for the SOA
3723 infrastructure want followed. These policies are initially tested during service development and are
3724 continuously monitored during the operational lifetime of the service.

3725 The testing of the service interface and service reachability are often related but essentially reflect
3726 different motivations and needs. The service interface is specified as a joint product of the service
3727 owners and providers who define service functionality, the prospective consumer community, the service
3728 developer, and the governance process. The semantics of the information model must align with the
3729 semantics of those who consume the service in order for there to be meaningful exchange of information.
3730 The structure of the information is influenced by the consumer semantics and the requirements and
3731 constraints of the representation as interpreted by the service developer. The service process model that
3732 defines actions which can be performed against a service and any temporal dependencies derive from
3733 the defined functionality and may be influenced by the development process. Any of these constraints
3734 may be identified and expressed as policy through the governance process.

3735 Service reachability conditions are the purview of the service provider who identifies the service endpoint
3736 and the protocols recognized at the endpoint. These may be constrained by governance decisions on
3737 how endpoint addresses may be allocated and what protocols should be used.

3738 While the considerations for defining the service interface derive from several sources, testing of the
3739 service interface is more straightforward and isolated in the testing process. At any point where the
3740 interface is modified or exposes a new resource, the message exchange should be monitored both to
3741 ensure the message reaches its intended destination and it is parsed correctly once received. Once an
3742 interface has been shown to function properly, it is unlikely it will fail later unless something fundamental
3743 to the service changes.

3744 The service interface is also tested when the service endpoint changes. Testing of the endpoint ensures
3745 message exchange can occur at the time of testing and the initial testing shows the interface is being
3746 processed properly at the new endpoint. Functioning of a service endpoint at one time does not
3747 guarantee it is functioning at another time, e.g. the server with the endpoint address may be down,
3748 making testing of service reachability a continual monitoring function through the life of the service's use
3749 of the endpoint. Also, while testing of the service endpoint is a necessary and most commonly noted part
3750 of the test regiment, it is not in itself sufficient to ensure the other aspects of testing discussed in this
3751 section.

3752 Finally, governance is impossible without the collection of metrics against which service behavior can be
3753 assessed. Metrics are also a key indicator for consumers to decide if a service is adequate for their
3754 needs. For instance, the average response time or the recent availability can be determining factors even
3755 if there are no rules or regulations promulgated through the governance process against which these
3756 metrics are assessed. The available metrics are a combination of those expected by the consumer
3757 community and those mandated through the governance process. The total set of metrics will evolve
3758 over time with SOA experience. Testing of the services that gather and provide access to the metrics will
3759 follow testing as described in this section, but for an individual service, testing will ensure that the metrics
3760 access indicated in the service description is accurate.

3761 The individual test requirements highlight aspects of the service that testing must consider but testing
3762 must establish more than isolated behavior. The emphasis is the holistic results of interacting with the
3763 service in the SOA environment. Recall that the execution context is the set of agreements between a
3764 consumer and a provider that define the conditions under which service interaction occurs. The
3765 agreements are expected to be predominantly the acceptance of the standard conditions as enumerated
3766 by the service provider, but it may include the identification of alternate conditions that will govern the
3767 interaction.

3768 For example, the provider may prefer a policy where it can sell the contact information of its consumers
3769 but will honor the request of a consumer to keep such information private. The identification of the
3770 alternate privacy policy is part of the execution context, and it is the application of and compliance with
3771 this policy that operational monitoring will attempt to measure. The collection of metrics showing this
3772 condition is indeed met when chosen is considered part of the ongoing testing of the service.

3773 Other variations in the execution context also require monitoring to ensure that different combinations of
3774 conditions perform together as desired. For example, if a new privacy policy takes additional resources to

3775 apply, this may affect quality of service and propagate other effects. These could not be tested during the
3776 original testing if the alternate policy did not exist at that time.

3777 **5.4.3.1.2 Testing Against Non-Functional Requirements**

3778 Testing against non-functional requirements constitutes testing of business usability of the service. In a
3779 marketplace of services, non-functional characteristics may be the primary differentiator between services
3780 that produce essentially the same real world effects.

3781 As noted in the previous section, non-functional characteristics are often associated with policies or other
3782 terms of use and may be collected in service level contracts offered by the service providers. Non-
3783 functional requirements may also reflect the network and hardware infrastructure that support
3784 communication with the service, and changes may impact quality of service. The service consumer and
3785 even the service provider may not be aware of all such infrastructure changes but the changes may
3786 manifest in shared states that impact the usability of the service.

3787 In general, a change in the non-functional requirements results in a change to the execution context, but
3788 as with any collection of information that constitutes a description, the execution context is unable to
3789 explicitly capture all non-functional requirements that may apply. A change in non-functional
3790 requirements, whether explicitly part of the execution context or an implicit contributor, may require
3791 retesting of the service even if its functionality and the implementation of the functionality has not
3792 changed. Depending on the circumstances, retesting may require a formal recertifying of end-to-end
3793 behavior or more likely will be part of the continuous monitoring that applies throughout the service
3794 lifetime.

3795 **5.4.3.1.3 Testing Content and the Interests of Consumers**

3796 As noted in section 5.4.1.1, testing may involve verification of conformance with respect to policies and
3797 technical specifications and validation with respect to sufficiency of functionality to meet some prescribed
3798 use. It may also include demonstration of performance and quality aspects. For some of these items,
3799 such as demonstrating the business processes followed in developing the service or the use of standards
3800 in implementing the service, the testing or relevant auditing is done internal to the service development
3801 process and follows traditional software testing and quality assurance. If it is believed of value to
3802 potential consumers, information about such testing could be included in the service description.
3803 However, it is not required that all test or compliance artifacts be available to consumers, as many of the
3804 details tested may be part of the opacity of the service implementation.

3805 Some aspects of the service being tested will reflect directly on the real world effects realized through
3806 interaction with the service. In these cases, it is more likely that testing results will be directly relevant to
3807 potential consumers. For example, if the service was designed to correspond to certain elements of a
3808 business process or that a certain workflow is followed, testing should verify that the real world effects
3809 reflect that the business process or workflow were satisfactorily captured.

3810 The testing may also need to demonstrate that specified conditions of use are satisfied. For example,
3811 policies may be asserted that require certain qualifications of or impose restrictions on the consumers
3812 who may interact with the service. The service testing must demonstrate that the service independently
3813 enforces the policies or it provides the required information exchanges with the SOA ecosystem so other
3814 resources can ensure the specified conditions.

3815 The completeness of the testing, both in terms of the features tested and the range of parameters for
3816 which response is tested, depends on the context of expected use: the more critical the use, the more
3817 complete the testing. There are always limits on the resources available for testing, if nothing else than
3818 the service must be available for use in a finite amount of time.

3819 This again emphasizes the need for adequate documentation to be available. If the original testing is
3820 very thorough, it may be adequate for less demanding uses in the future. If the original testing was more
3821 constrained, then well-documented test results establish the foundation on which further testing can be
3822 defined and executed.

3823 **5.4.3.2 How Testing is to be Done**

3824 Testing should follow well-defined methodologies and, if possible, should reuse test artifacts that have
3825 proven generally useful for past testing. For example, IEEE-829 notes that test cases are separated from
3826 test designs to allow for use in more than one design and to allow for reuse in other situations. In the
3827 SOA ecosystem, description of such artifacts, as with description of a service, enables awareness of the
3828 item and describes how the artifact may be accessed or used.

3829 As with traditional testing, the specific test procedures and test case inputs are important so the tests are
3830 unambiguously defined and entities can be retested in the future. Automated testing and regression
3831 testing may be more important in the SOA ecosystem in order to re-verify a service is still acceptable
3832 when incorporated in a new use. For example, if a new use requires the services to deal with input
3833 parameters outside the range of initial testing, the tests could be rerun with the new parameters. If the
3834 testing resources are available to consumers within the SOA ecosystem, the testing as designed by test
3835 professionals could be consumed through a service accessed by consumers, and their results could
3836 augment those already in place. This is discussed further in the next section.

3837 **5.4.3.3 Who Performs the Testing**

3838 As with any software, the first line of testing is unit testing done by software developers. It is likely that
3839 initial testing will be done by those developing the software but may also be done independently by other
3840 developers. For SOA development, unit testing is likely confined to a development sandbox isolated from
3841 the SOA ecosystem.

3842 SOA testing will differ from traditional software testing in that testing beyond the development sandbox
3843 must incorporate aspects of the SOA ecosystem, and those doing the testing must be familiar with both
3844 the characteristics and responses of the ecosystem and the tools, especially those available as services,
3845 to facilitate and standardize testing. Test professionals will know what level of assurance must be
3846 established as the exposure of the service to the ecosystem and ecosystem to the service increases
3847 towards operational status. These test professionals may be internal resources to an organization or may
3848 evolve as a separate discipline provided through external contracting.

3849 As noted above, it is unlikely that a complete duplicate of the SOA ecosystem will be available for isolated
3850 testing, and thus use of ecosystem resources will manifest as a transition process rather than a step
3851 change from a test environment to an operational one. This is especially true for new composite services
3852 that incorporate existing operational services to achieve the new functionality. The test professionals will
3853 need to understand the available resources and the ramifications of this transition.

3854 As with current software development, a stage beyond work by test professionals will make use of a
3855 select group of typical users, commonly referred to as beta testers, to report on service response during
3856 typical intended use. This establishes fitness by the consumers, providing final validation of previously
3857 verified processes, requirements, and final implementation.

3858 In traditional software development, beta testing is the end of testing for a given version of the software.
3859 However, although the initial test phase can establish an appropriate level of confidence consistent with
3860 the designed use for the initial target consumer community, the operational service will exist in an
3861 evolving ecosystem, and later conditions of use may differ from those thought to be sufficient during the
3862 initial testing. Thus, operational monitoring becomes an extension of testing through the service lifetime.
3863 This continuous testing will attempt to ensure that a service does not consume an inordinate amount of
3864 ecosystem resources or display other behavior that degrades the ecosystem, but it will not undercover
3865 functional errors that may surface over time.

3866 As with any software, it is the responsibility of the consumers to consider the reasonableness of solutions
3867 in order to spot errors in either the software or the way the software is being used. This is especially
3868 important for consumers with unanticipated uses that may go beyond the original test conditions. It is
3869 unlikely the consumers will initiate a new round of formal testing unless the new use requires a
3870 significantly higher level of confidence in the service. Rather the consumer becomes a new extension to
3871 the testing regiment. Obvious testing would include a sanity check of results during the new use.
3872 However, if the details of legacy testing are associated with the service through the service description
3873 and if testing resources are available through automated testing services, then the new consumers can
3874 rerun and extend previous testing to include the extended test conditions. If the test results are

3875 acceptable, these can be added to the documentation of previous results and become the extended basis
3876 for future decisions by prospective consumers on the appropriateness of the service. If the results are not
3877 acceptable or in some way questionable, the responsible party for the service or testing professionals can
3878 be brought in to decide if remedial action is necessary.

3879 **5.4.3.4 How Testing Results are Reported**

3880 For any SOA service, an accurate reporting of the testing a service has undergone and the results of the
3881 testing is vital to consumers deciding whether a service is appropriate for intended use. Appropriateness
3882 may be defined by a consumer organization and require specific test regiments culminating in a
3883 certification; appropriateness could be established by accepting testing and certifications that have been
3884 conferred by others.

3885 The testing and certification information should be identified in the service description. Referring to the
3886 general description model of *Figure 25*, tests conducted by or under a request from the service owner (see
3887 Ownership in section 3.3.4) would be captured under Annotations from Owners. Testing done by others,
3888 such as consumers with unanticipated uses, could be associated through Annotations from 3rd Parties.
3889 The annotations should clearly indicate what was tested, how the testing was done, who did the testing,
3890 and the testing results. The clear description of each of these artifacts and of standardized testing
3891 protocols for various levels of sophistication and completeness of testing would enable a common
3892 understanding and comparison of test coverage. It will also make it more straightforward to conduct and
3893 report on future testing, facilitating the maintenance of the service description.

3894 Consumer testing and the reporting of results raises additional issues. While stating who did the testing
3895 is mandatory, there may be formal requirements for authentication of the tester to ensure traceability of
3896 the testing claims. In some circumstances, persons or organizations would not be allowed to state testing
3897 claims unless the tester was an approved entity. In other cases, ensuring the tester had a valid email
3898 may be sufficient. In either case, it would be at the discretion of the potential consumer to decide what
3899 level of authentication was acceptable and which testers are considered authoritative in the context of
3900 their anticipated use.

3901 Finally, in a world of openly shared information, we would see an ever-expanding set of testing
3902 information as new uses and new consumers interact with a service. In reality, these new uses may
3903 represent proprietary processes or classified use that should only be available to authorized parties.
3904 Testing information, as with other elements of description, may require special access controls to ensure
3905 appropriate access and use.

3906 **5.4.4 Testing SOA Services**

3907 Testing of SOA services should be consistent with the SOA paradigm. In particular, testing resources
3908 and artifacts should be visible in support of service interaction between providers and consumers, where
3909 here the interaction is between the testing resource and the tester. In addition, the idea of opacity of the
3910 implementation should limit the details that need to be available for effective use of the test resources.
3911 Testing that requires knowledge of the internal structure of the service or its underlying capability should
3912 be performed as part of unit testing in the development sandbox, and should represent a minimum level
3913 of confidence before the service begins its transition to further testing and eventual operation in the SOA
3914 ecosystem.

3915 **5.4.4.1 Progression of SOA Testing**

3916 Software testing is a gradual exercise going from micro inspection to testing macro effects. The first step
3917 in testing is likely the traditional code reviews. SOA considerations would account for the distributed
3918 nature of SOA, including issues of distributed security and best practices to ensure secure resources. It
3919 would also set the groundwork for opacity of implementation, hiding programming details and simplifying
3920 the use of the service.

3921 Code review is likely followed by unit testing in a development sandbox isolated from the operational
3922 environment. The unit testing is done with full knowledge of the service internal structure and knowledge
3923 of resources representing underlying capabilities. It tests the interface to ensure exchanged messages
3924 are as specified in the service description and the messages can be parsed and interpreted as intended.

3925 Unit testing also verifies intended functionality and that the software has dealt correctly with internal
3926 dependencies, such as structure of a file system or access to other dedicated resources.

3927 Some aspects of unit testing require external dependencies be satisfied, and this is often done using
3928 mock objects to substitute for the external resources. In particular, it will likely be necessary to include
3929 mocks of existing operational services, both those provided as part of the SOA infrastructure and services
3930 from other providers.

3931 **Service Mock**

3932 A service mock is an entity that mimics some aspect of the performance of an operational service
3933 without committing to the real world effects that the operational service would produce.

3934 Mocks are discussed in detail in sections 5.4.4.3 and 5.4.4.4.

3935 After unit testing has demonstrated an adequate level of confidence in the service, the testing must
3936 transition from the tightly controlled environment of the development sandbox to an environment that
3937 more clearly resembles the operational SOA ecosystem or, at a minimum, the intended enterprise. While
3938 sandbox testing will use simple mocks of some aspects of the SOA environment, such as an interface to
3939 a security service without the security service functionality, the dynamic nature of SOA makes a full
3940 simulation infeasible to create or maintain. This is especially true when a new composite service makes
3941 use of operational services provided by others. Thus, at some point before testing is complete, the
3942 service will need to demonstrate its functionality by using resources and dealing with conditions that only
3943 exist in the full ecosystem or the intended enterprise. Some of these resources may still provide test
3944 interfaces -- more on this below -- but the interfaces will be accessible via the SOA environment and not
3945 just implemented for the sandbox.

3946 At this stage, the opacity of the service becomes important as the details of interacting with the service
3947 now rely on correct use of the service interface and not knowledge of the service internals. The workings
3948 of the service will only be observable through the real world effects realized through service interactions
3949 and external indications that conditions of use, such as user authentication, are satisfied. Monitoring the
3950 behavior of the service will depend on service interfaces that expose internal monitoring or provide
3951 required information to the SOA infrastructure monitoring function. The monitoring required to test a new
3952 service is likely to have significant overlap with the monitoring the SOA infrastructure includes to monitor
3953 its own health and to identify and isolate behavior outside of acceptable bounds. This is exactly what is
3954 needed as part of service testing, and it is reasonable to assume that the ecosystem transition includes
3955 use of operational monitoring rather than solely dedicated monitoring for each service being tested.

3956 Use of SOA monitoring resources during the explicit testing phase sets the stage for monitoring and a
3957 level of continual testing throughout the service lifetime.

3958 **5.4.4.2 Testing Traditional Dependencies vs. Service Interactions**

3959 A SOA service is not required to make use of other operational services beyond what may be required for
3960 monitoring by the ecosystem infrastructure. The service can implement hardcoded dependencies which
3961 have been tested in the development sandbox through the use of dedicated mocks. While coordination
3962 may be required with real data sources during integration testing, the dependencies can be constrained to
3963 things that can be tested in a more traditional manner. Policies can also be set to restrict access to pre-
3964 approved users, and thus the question of unanticipated users and unanticipated uses can be eliminated.
3965 Operational readiness can be defined in terms of what can be proven in isolated testing. While all this
3966 may provide more confidence in the service for its designed purpose, such a service will not fully
3967 participate in the benefits or challenges of the ecosystem. This is akin to filling a swimming pool with sea
3968 water and having someone in the pool say they are swimming in the ocean.

3969 In considering the testing needed for a fully participating service, consider the example of a new
3970 composite service that combines the real world effects and complies with the conditions of use of five
3971 existing operational services. The developer of the composite service does not own any of the
3972 component services and has limited, if any, ability to get the distributed owners to do any customization.
3973 The developer also is limited by the principle of opacity to information comprising the service description,
3974 and does not know internal details of the component services. The developer of the composite service
3975 must use the component services as they exist as part of the SOA environment, including what is

3976 provided to support testing by new users. This introduces requirements for what is needed in the way of
3977 service mocks.

3978 **5.4.4.3 Use of Service Mocks**

3979 Service mocks enables the tested service to respond to specific features of an operational service that is
3980 being used as a component. It allows service testing to proceed without needing access to or with only
3981 limited engagement with the component service. Mocks can also mimic difficult to create situations for
3982 which it is desired to test the new service response. For composite services using multiple component
3983 services, mocks may be used in combination to function for any number of the components. Note, when
3984 using service mocks, it is important to remember that it is not the component service that is being tested
3985 (although anomalous behavior may be uncovered during testing) but the use of the component in the new
3986 composite.

3987 Individual service mocks can emphasize different features of the component service they represent but
3988 any given mock does not have to mimic all features. For example, a mock of the service interface can
3989 echo a sent message and demonstrate the message is reaching its intended destination. A mock could
3990 go further and parse the sent message to demonstrate the message not only reached its destination but
3991 was understood. As a final step, the mock could report back what actions would have been taken by the
3992 component service and what real world effects would result. If the response mimicked the operational
3993 response, functional testing could proceed as if the real world effect actually occurred.

3994 There are numerous ways to provide mock functionality. The service mock could be a simulation of the
3995 operational service and return simulated results in a realistic response message or event notification. It is
3996 also possible for the operational service to act as its own mock and simply not execute the commit stage
3997 of its functionality. The service mock could use a combination of simulation and service action without
3998 commit to generate a report of what would have occurred during the defined interaction with the
3999 operational service.

4000 As the service proceeds through testing, mocks should be systematically replaced by the component
4001 resources accessed through their operational interfaces. Before beta testing begins, end-to-end testing,
4002 i.e. proceeding from the beginning of the service interaction to the resulting real world results, should be
4003 accomplished using component resources via their operational interfaces.

4004 **5.4.4.4 Providers of Service Mocks**

4005 In traditional testing, it is often the test professionals who design and develop the mocks, but in the
4006 distributed world of SOA, this may not be efficient or desirable.

4007 In the development sandbox, it is likely the new service developer or test professionals working with the
4008 developer will create mocks adequate for unit testing. Given that most of this testing is to verify the new
4009 service is performing as designed, it is not necessary to have high fidelity models of other resources
4010 being accessed. In addition, given opacity of SOA implementation, the developer of the new service may
4011 not have sufficient detailed knowledge of a component service to build a detailed mock of the component
4012 service functionality. Sharing existing mocks at this stage may be possible but the mocks would need to
4013 be implemented in the sandbox, and for simple models it is likely easier to build the mock from scratch.

4014 As testing begins its transition to the wider SOA environment, mocks may be available as services. For
4015 existing resources, it is possible that an Open Source model could evolve where service mocks of
4016 available functions can be catalogued and used during initial interaction of the tested service and the
4017 operational environment. Widely used functions may have numerous service mocks, some mimicking
4018 detailed conditions within the SOA infrastructure. However, the Open Source model is less likely to be
4019 sufficient for specialty services that are not widely used by a large consumer community.

4020 The service developer is probably best qualified for also developing more detailed service mocks or for
4021 mock modes of operational services. This implies that in addition to their operational interfaces, services
4022 will routinely provide test interfaces to enable service mocks to be used as services. As noted above, a
4023 new service developer wanting to build a mock of component services is limited to the description
4024 provided by the component service developer or owner. The description typically will detail real world
4025 effects and conditions of use but will not provide implementation details, some of which may be
4026 proprietary. Just as important in the SOA ecosystem, if it becomes standard protocol for developers to

4027 create service mocks of their own services, a new service developer is only responsible for building his
4028 own mocks and can expect other mocks to be available from other developers. This reduces duplication
4029 of effort where multiple developers would be trying to build the same mocks from the same insufficient
4030 information. Finally, a service developer is probably best qualified to know when and how a service mock
4031 should be updated to reflect modified functionality or message exchange.

4032 It is also possible that testing organizations will evolve to provide high-fidelity test harnesses for new
4033 services. The harnesses would allow new services to plug into a test environment and would facilitate
4034 accessing mocks of component services. However, it will remain a constant challenge for such
4035 organizations to capture evolving uses and characteristics of service interactions in the real SOA
4036 environment and maintain the fidelity and accuracy of the test systems.

4037 **5.4.4.5 Fundamental Questions for SOA Testing**

4038 In order for the transition to the SOA operational environment to proceed, it is necessary to answer two
4039 fundamental questions:

- 4040 • Who provides what testing resources for the SOA operational environment, e.g. mocks of
4041 interfaces, mocks of functionality, monitoring tools?
- 4042 • What testing needs to be accomplished before operational environment resources can be
4043 accessed for further testing?

4044 The discussion in section 5.4.4.4 notes various levels of sophistication of service mocks and different
4045 communities are likely to be responsible for different levels. Section 5.4.4.4 advocates a significant role
4046 for service developers, but there needs to be community consensus that such mocks are needed and that
4047 service developers will agree to fulfilling this role. There is also a need for consensus as to what tools
4048 should be available as services from the SOA infrastructure.

4049 As for use of the service mocks and SOA environment monitoring services, practical experience is
4050 needed upon which guidelines can be established for when a new service has been adequately tested to
4051 proceed with a greater level of exposure with the SOA environment. Malfunctioning services could cause
4052 serious problems if they cannot be identified and isolated. On the other hand, without adequate testing
4053 under SOA operational conditions, it is unlikely that problems can be uncovered and corrected before
4054 they reach an operational stage.

4055 As noted in section 5.4.4.2, some of these questions can be avoided by restricting services to more
4056 traditional use scenarios. However, such restriction will limit the effectiveness of SOA use and the result
4057 will resemble the constraints of traditional integration activities we are trying to move beyond.

4058 **5.4.5 Architectural Implications for SOA Testing**

4059 The discussion of SOA Testing indicates numerous architectural implications on the SOA ecosystem:

- 4060 • The distributed, boundary-less nature of the SOA ecosystem makes it infeasible to create
4061 and maintain a single mock of the entire ecosystem to support testing activities.
- 4062 • A standard suite of monitoring services needs to be defined, developed, and maintained.
4063 This should be done in a manner consistent with the evolving nature of the ecosystem.
- 4064 • Services should provide interfaces that support access in a test mode.
- 4065 • Testing resources must be described and their descriptions must be catalogued in a
4066 manner that enables their discovery and access.
- 4067 • Guidelines for testing and ecosystem access need to be established and the ecosystem
4068 must be able to enforce those guidelines asserted as policies.
- 4069 • Services should be available to support automated testing and regression testing.
- 4070 • Services should be available to facilitate updating service description by anyone who has
4071 performed testing of a service.

4072

6 Conformance

4073

This Reference Architecture Foundation is an abstract architecture, which means that it is especially difficult to construct automated tests for conformance to the architecture. However, in order to be conformant to this architecture, it should be possible to identify in a concrete implementation the key concepts and components of this architecture, albeit in abstracted form.

4074

4075

4076

[EDITOR'S NOTE: The conformance section is not complete]

4077

4078

4079

A. Acknowledgements

4080 The following individuals have participated in the creation of this specification and are gratefully
4081 acknowledged:

4082 **Participants:**

4083 Rex Brooks, Individual Member
4084 Peter Brown, Pensive.eu
4085 Scott Came, Search Group Inc.
4086 Joseph Chiusano, Booz Allen Hamilton
4087 Robert Ellinger, Northrop Grumman Corporation
4088 David Ellis, Sandia National Laboratories
4089 Jeff A. Estefan, Jet Propulsion Laboratory
4090 Don Flinn, Individual Member
4091 Anil John, Johns Hopkins University
4092 Ken Laskey, MITRE Corporation
4093 Boris Lublinsky, Nokia Corporation
4094 Francis G. McCabe, Individual Member
4095 Christopher McDaniels, USSTRATCOM
4096 Tom Merkle, Lockheed Martin Corporation
4097 Jyoti Namjoshi, Patni Computer Systems Ltd.
4098 Duane Nickull, Adobe Inc.
4099 James Odell, Associate
4100 Michael Poulin, Fidelity Investments
4101 Michael Stiefel, Associate
4102 Danny Thornton, Northrop Grumman
4103 Timothy Vibbert, Lockheed Martin Corporation
4104 Robert Vitello, New York Dept. of Labor

4105 B. Critical Factors Analysis

4106 A critical factors analysis (CFA) is an analysis of the key properties of a project. A CFA is analyzed in
4107 terms of the goals of the project, the critical factors that will lead to its success and the measurable
4108 requirements of the project implementation that support the goals of the project. CFA is particularly
4109 suitable for capturing quality attributes of a project, often referred to as “non-functional” or “other-than-
4110 functional” requirements: for example, security, scalability, wide-spread adoption, and so on. As such,
4111 CFA complements rather than attempts to replace other requirements capture techniques.

4112 B.1 Goals

4113 A goal is an overall target that you are trying to reach with the project. Typically, goals are hard to
4114 measure by themselves. Goals are often directed at the potential consumer of the product rather than the
4115 technology developer.

4116 Critical Success Factors

4117 A critical success factor (CSF) is a property, sub-goal that directly supports a goal and there is strong
4118 belief that without it the goal is unattainable. CSFs themselves are not necessarily measurable in
4119 themselves.

4120 Requirements

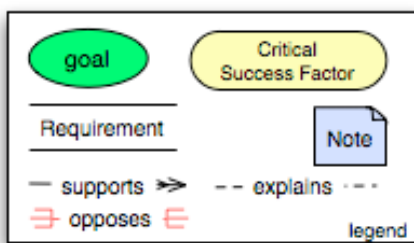
4121 A requirement is a specific measurable property that directly supports a CSF. The key here is
4122 measurability: it should be possible to unambiguously determine if a requirement has been met. While
4123 goals are typically directed at consumers of the specification, requirements are focused on technical
4124 aspects of the specification.

4125 CFA Diagrams

4126 It can often be helpful to illustrate graphically the key concepts and relationships between them. Such
4127 diagrams can act as effective indices into the written descriptions of goals etc., but is not intended to
4128 replace the text.

4129 The legend:

4130



4131 illustrates the key elements of the graphical notation. Goals are written in round ovals, critical success
4132 factors are written in round-ended rectangles and requirements are written using open-ended rectangles.
4133 The arrows show whether a CSF/goal/requirement is supported by another element or opposed by it. This
4134 highlights the potential for conflict in requirements.
4135