



Reference Architecture Foundation for Service Oriented Architecture Version 1.0

Working Draft 07

11 May 2012

Specification URIs:

This version:

[Working Draft – no public URL](#)

Previous working draft version:

<http://www.oasis-open.org/apps/org/workgroup/soa-rm-ra/download.php/45861> (21 April)

Previous published version:

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

<http://docs.oasis-open.org/soa-rm/soa-ra/v1.0/soa-ra-cd-02.html>

<http://docs.oasis-open.org/soa-rm/soa-ra/v1.0/soa-ra-cd-02.doc>

Technical Committee:

OASIS Service Oriented Architecture Reference Model TC

Chair:

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

Editors:

Peter Brown (peter@peterbrown.com), Individual Member

Jeff A. Estefan (jeffrey.a.estefan@jpl.nasa.gov), Jet Propulsion Laboratory

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

Francis G. McCabe (fmccabe@gmail.com), Individual Member

Danny Thornton (danny.thornton@ngc.com), Northrop Grumman

Related work:

This specification is related to:

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

Abstract:

This document specifies the OASIS Reference Architecture Foundation for Service Oriented Architecture (SOA-RAF). It follows from the concepts and relationships defined in the OASIS Reference Model for Service Oriented Architecture. While it remains abstract in nature, the current document describes the foundation upon which specific SOA concrete architectures can be built.

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

The SOA-RAF follows the recommended practice of describing architecture in terms of models, views, and viewpoints, as prescribed in the ANSI/IEEE 1471-2000 (now ISO/IEC 42010-2007) Standard.

It has three main views: the *Participation in a SOA Ecosystem* view which focuses on the way that participants are part of a Service Oriented Architecture ecosystem; the *Realization of a SOA Ecosystem* view which addresses the requirements for constructing a SOA-based system in a

SOA ecosystem; and the *Ownership in a SOA Ecosystem* view which focuses on what is meant to own a SOA-based system.

The SOA-RAF is of value to Enterprise Architects, Business and IT Architects as well as CIOs and other senior executives involved in strategic business and IT planning.

Status:

This [Working Draft](#) (WD) has been produced by one or more TC Members; it has not yet been voted on by the TC or [approved](#) as a Committee Draft (Committee Specification Draft or a Committee Note Draft). The OASIS document [Approval Process](#) begins officially with a TC vote to approve a WD as a Committee Draft. A TC may approve a Working Draft, revise it, and re-approve it any number of times as a Committee Draft.

Citation format:

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

[SOA-RAF]

Reference Architecture Foundation for Service Oriented Architecture Version 1.0. 06 July 2011.
OASIS Committee Specification Draft 03 / Public Review Draft 02. <http://docs.oasis-open.org/soa-raf/soa-raf/v1.0/csprd02/soa-raf-v1.0-csprd02.html>.

Notices

Copyright © OASIS Open 2011. All Rights Reserved.

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

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

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

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

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

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

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

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

Table of Contents

1	Introduction	10
1.1	Context for Reference Architecture for SOA	10
1.1.1	What is a Reference Architecture?	10
1.1.2	What is this Reference Architecture?	10
1.1.3	Relationship to the OASIS Reference Model for SOA	11
1.1.4	Relationship to other Reference Architectures	11
1.1.5	Expectations set by this Reference Architecture Foundation	11
1.2	Service Oriented Architecture – An Ecosystems Perspective	12
1.3	Viewpoints, Views and Models	12
1.3.1	ANSI/IEEE 1471-2000:ISO/IEC 42010-2007	12
1.3.2	UML Modeling Notation	13
1.4	SOA-RAF Viewpoints	14
1.4.1	Participation in a SOA Ecosystem Viewpoint	14
1.4.2	Realization of a SOA Ecosystem Viewpoint	14
1.4.3	Ownership in a SOA Ecosystem Viewpoint	15
1.5	Terminology	15
1.6	References	15
1.6.1	Normative References	15
1.6.2	Non-Normative References	16
2	Architectural Goals and Principles	17
2.1	Goals and Critical Success Factors of the Reference Architecture Foundation	17
2.1.1	Goals	17
2.1.1.1	Effectiveness	17
2.1.1.2	Confidence	17
2.1.1.3	Scalability	17
2.1.2	Critical Success Factors	18
2.1.2.1	Action	18
2.1.2.2	Trust	18
2.1.2.3	Interaction	18
2.1.2.4	Control	18
2.2	Principles of this Reference Architecture Foundation	18
3	Participation in a SOA Ecosystem View	20
3.1	SOA Ecosystem Model	21
3.2	Social Structure in a SOA Ecosystem Model	22
3.2.1	Stakeholders, Participants, Actors and Delegates	24
3.2.2	Social Structures and Roles	25
3.2.2.1	Authority, Rights, and Responsibilities	25
3.2.2.2	Permissions and Obligations	26
3.2.2.3	Service Roles	27
3.2.3	Needs, Requirements and Capabilities	28
3.2.4	Resource and Ownership	30
3.2.4.1	Resource	30
3.2.4.2	Ownership	31
3.2.5	Establishing Execution Context	31

3.2.5.1	Trust and Risk	32
3.2.5.2	Policies and Contracts	33
3.2.5.3	Communication	34
3.2.5.4	Semantics and Semantic Engagement	34
3.3	Action in a SOA Ecosystem Model	35
3.3.1	Services Reflecting Business	36
3.3.2	Activity, Action, and Joint Action	36
3.3.3	State and Shared State	39
3.4	Architectural Implications	39
3.4.1	Social structures	39
3.4.2	Resource and Ownership	39
3.4.3	Policies and Contracts	40
3.4.4	Communications as a Means of Mediating Action	40
3.4.5	Semantics	40
3.4.6	Trust and Risk	40
3.4.7	Needs, Requirements and Capabilities	41
3.4.8	The Importance of Action	41
4	Realization of a SOA Ecosystem view	42
4.1	Service Description Model	42
4.1.1	The Model for Service Description	43
4.1.1.1	Elements Common to General Description	43
4.1.1.2	Assigning Values to Description Instances	45
4.1.1.3	Model Elements Specific to Service Description	47
4.1.2	Use of Service Description	51
4.1.2.1	Service Description in support of Service Interaction	51
4.1.2.2	Description and Invoking Actions Against a Service	53
4.1.2.3	The Question of Multiple Business Functions	54
4.1.2.4	Service Description, Execution Context, and Service Interaction	54
4.1.3	Relationship to Other Description Models	56
4.1.4	Architectural Implications	57
4.2	Service Visibility Model	58
4.2.1	Visibility to Business	58
4.2.2	Visibility	59
4.2.2.1	Awareness	60
4.2.2.2	Willingness	62
4.2.2.3	Reachability	63
4.2.3	Architectural Implications	64
4.3	Interacting with Services Model	64
4.3.1	Interaction Dependencies	65
4.3.2	Actions and Events	65
4.3.3	Message Exchange	66
4.3.3.1	Message Exchange Patterns (MEPs)	67
4.3.3.2	Request/Response MEP	68
4.3.3.3	Event Notification MEP	68
4.3.4	Composition of Services	68
4.3.5	Service Composition of Business Processes and Collaborations	69
4.3.5.1	Service-Oriented Business Processes	70

4.3.5.2 Service-Oriented Business Collaborations	72
4.3.6 Architectural Implications of Interacting with Services	73
4.4 Policies and Contracts Model	74
4.4.1 Policy and Contract Representation	75
4.4.1.1 Policy Framework	75
4.4.2 Policy and Contract Enforcement	76
4.4.2.1 Enforcing Simple Policy Constraints	76
4.4.2.2 Conflict Resolution	76
4.4.3 Architectural Implications.....	77
5 Ownership in a SOA Ecosystem View	78
5.1 Governance Model	78
5.1.1 Understanding Governance	78
5.1.1.1 Terminology	78
5.1.1.2 Relationship to Management	79
5.1.1.3 Why is SOA Governance Important?	79
5.1.1.4 Governance Stakeholders and Concerns	79
5.1.2 A Generic Model for Governance	80
5.1.2.1 Motivating Governance	80
5.1.2.2 Setting Up Governance.....	81
5.1.2.3 Carrying Out Governance	82
5.1.2.4 Ensuring Governance Compliance	83
5.1.2.5 Considerations for Multiple Governance Chains	83
5.1.3 Governance Applied to SOA	84
5.1.3.1 Where SOA Governance is Different	84
5.1.3.2 What Must be Governed	84
5.1.3.3 Overarching Governance Concerns.....	86
5.1.3.4 Considerations for SOA Governance.....	87
5.1.4 Architectural Implications of SOA Governance.....	88
5.2 Security Model	88
5.2.1 Secure Interaction Concepts.....	89
5.2.1.1 Confidentiality	89
5.2.1.2 Integrity	89
5.2.1.3 Authentication	89
5.2.1.4 Authorization	90
5.2.1.5 Non-repudiation	91
5.2.1.6 Availability	91
5.2.2 Where SOA Security is Different.....	92
5.2.3 Security Threats	92
5.2.4 Security Responses	93
5.2.4.1 Privacy Enforcement.....	93
5.2.4.2 Integrity Protection	93
5.2.4.3 Message Replay Protection	94
5.2.4.4 Auditing and Logging	94
5.2.4.5 Graduated engagement	94
5.2.5 Architectural Implications of SOA Security.....	94
5.3 Management Model	95
5.3.1 Management.....	95
5.3.2 Management Means and Relationships	99

5.3.2.1 Management Policy	100
5.3.2.2 Network Management	100
5.3.2.3 Security Management	100
5.3.2.4 Usage Management.....	100
5.3.3 Management and Governance.....	100
5.3.4 Management and Contracts	101
5.3.4.1 Management for Contracts and Policies	101
5.3.4.2 Contracts	101
5.3.4.3 Policies	103
5.3.4.4 Service Description and Management	103
5.3.5 Management for Monitoring and Reporting.....	104
5.3.6 Management for Infrastructure	104
5.3.7 Architectural Implication of the SOA Management.....	105
5.4 SOA Testing Model.....	105
5.4.1 Traditional Software Testing as Basis for SOA Testing	105
5.4.1.1 Types of Testing	106
5.4.1.2 Range of Test Conditions	106
5.4.1.3 Configuration Management of Test Artifacts.....	106
5.4.2 Testing and the SOA Ecosystem	106
5.4.3 Elements of SOA Testing	107
5.4.3.1 What is to be Tested	107
5.4.3.2 How Testing is to be Done	110
5.4.3.3 Who Performs the Testing	110
5.4.3.4 How Testing Results are Reported	111
5.4.4 Testing SOA Services.....	112
5.4.4.1 Progression of SOA Testing.....	112
5.4.4.2 Testing Traditional Dependencies vs. Service Interactions.....	113
5.4.4.3 Use of Service Mocks	113
5.4.4.4 Providers of Service Mocks	114
5.4.4.5 Fundamental Questions for SOA Testing	114
5.4.5 Architectural Implications for SOA Testing.....	115
6 Conformance	116
A. Acknowledgements	117
B. Index of Defined Terms.....	118
C. The Unified Modeling Language, UML	119
D. Critical Factors Analysis.....	120
D.1 Goals.....	120
D.2 Critical Success Factors	120
D.3 Requirements	120
E. Relationship to other SOA Open Standards.....	121
E.1 Navigating the SOA Open Standards Landscape Around Architecture	121
E.2 The Service-Aware Interoperability Framework: Canonical	122
E.3 IEEE Reference Architecture	123
E.4 RM-ODP.....	123

Table of Figures

Figure 1 - Model elements described in the Participation in a SOA Ecosystem view	20
Figure 2 - SOA Ecosystem Model	21
Figure 3 - Social Structure Model	23
Figure 4 – Stakeholders, Actors, Participants and Delegates	24
Figure 5 - Social Structures, Roles and Action	26
Figure 6 - Roles in a Service	28
Figure 7 - Cycle of Needs, Requirements, and Fulfillment	29
Figure 8 - Resources	30
Figure 9 - Willingness and Trust	32
Figure 10 – Policies, Contracts and Constraints	33
Figure 11: An Activity, expressed informally as a graph of Actions	37
Figure 12: Activity involving Actions across an ownership boundary	37
Figure 13 - Model Elements Described in the Realization of a SOA Ecosystem view	42
Figure 14 - General Description	44
Figure 15 - Representation of a Description	45
Figure 16 - Service Description	47
Figure 17 - Service Interface Description	48
Figure 18 - Service Functionality	49
Figure 19 - Model for Policies and Contracts as related to Service Participants	50
Figure 20 - Policies and Contracts, Metrics, and Compliance Records	51
Figure 21 - Relationship between Action and Components of Service Description Model	52
Figure 22 - Execution Context	55
Figure 23 - Interaction Description	56
Figure 24 - Visibility to Business	59
Figure 25 - Mediated Awareness	61
Figure 26 - Awareness in a SOA Ecosystem	62
Figure 27 - Business, Description and Willingness	63
Figure 28 - Service Reachability	63
Figure 29 - Interaction dependencies	65
Figure 30 - A "message" denotes either an action or an event	66
Figure 31 - Fundamental SOA message exchange patterns (MEPs)	67
Figure 32 - Simple model of service composition	69
Figure 33 - Abstract example of orchestration	71
Figure 34 - General Orchestration Pattern	71
Figure 35 - Abstract example of choreography	73
Figure 36 - Policies and Contracts	75
Figure 37 - Model Elements Described in the Ownership in a SOA Ecosystem View	78
Figure 38 - Motivating Governance	80
Figure 39 - Setting Up Governance	81
Figure 40 - Carrying Out Governance	82

Figure 41 - Ensuring Governance Compliance83
Figure 42 - Relationship Among Types of Governance.....85
Figure 43 - Authentication90
Figure 44 - Authorization.....91
Figure 45 - Management model in SOA ecosystem97
Figure 46 - Management Means and Relationships in a SOA ecosystem99
Figure 47 - Management of the service interaction102
Figure 48 - Example UML class diagram—Resources119
Figure 49 - SOA Reference Architecture Positioning122

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 bridges the area between business and IT. It is neither wholly IT nor wholly business, but is of both worlds. Neither business nor IT completely own, govern and manage this SOA ecosystem. Both sets of concerns must be accommodated for the SOA ecosystem to fulfill its purposes.¹

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

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

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

1.1 Context for Reference Architecture for SOA

1.1.1 What is a Reference Architecture?

A reference architecture models the abstract architectural elements in the domain of interest independent of the technologies, protocols, and products that are used to implement a specific solution for the domain. It differs from a reference model in that a reference model describes the important concepts and relationships in the domain focusing on what distinguishes the elements of the domain; a reference architecture elaborates further on the model to show a more complete picture that includes showing what is involved in realizing the modeled entities, while staying independent of any particular solution but instead applies to a class of solutions.

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

1.1.2 What is this Reference Architecture?

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

¹ By *business* we refer to any activity that people are engaged in. We do not restrict the scope of SOA ecosystems to commercial applications.

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

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

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

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

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

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

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

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

69 **1.1.4 Relationship to other Reference Architectures**

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

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

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

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

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

88 1.2 Service Oriented Architecture – An Ecosystems 89 Perspective

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

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

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

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

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

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

113 1.3 Viewpoints, Views and Models

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

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

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

127 The standard defines the following terms²:

128 **Architecture**

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

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

131 **Architectural Description**

132 A collection of products that document the architecture.

133 **System**

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

135 **System Stakeholder**

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

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

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

145 **View**

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

147 **Viewpoint**

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

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

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

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

162 **Model**

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

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

167 **1.3.2 UML Modeling Notation**

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

174 Appendix C introduces the UML notation that is used in this document.

175 **1.4 SOA-RAF Viewpoints**

176 The SOA-RAF specifies three views (described in detail in Sections 3, 4, and 5) that conform to three
 177 viewpoints: *Participation in a SOA Ecosystem*, *Realization of a SOA Ecosystem*, and *Ownership in a SOA*
 178 *Ecosystem*. There is a one-to-one correspondence between viewpoints and views (see Table 1).

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

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

180 **1.4.1 Participation in a SOA Ecosystem Viewpoint**

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

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

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

191 **1.4.2 Realization of a SOA Ecosystem Viewpoint**

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

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

198 1.4.3 Ownership in a SOA Ecosystem Viewpoint

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

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

207 1.5 Terminology

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

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

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

215 Usage of Terms

216 Certain terms **are** used in this document to denote concepts **that are formally defined here and are**
217 **intended to be** used with **the** specific meanings **indicated**. Where **mention** is **first** made of a formally
218 defined concept **and the prescribed meaning is intended**, we use a **bold font**. **When this occurrence**
219 **appears in the text substantially in advance of the formal definition, it is** also **hyperlinked** to the definition
220 in the body of the text. **A list of all such terms is included in the Index of Terms at Appendix B**. Where a
221 more colloquial or informal meaning is intended, these words are used without special emphasis.

Field Code Changed

Comment [PFB1]: Issue 8

To be implemented throughout document (only section 3 done so far)

222 1.6 References

223 1.6.1 Normative References

- 224 [ANSI/IEEE 1471] *IEEE Recommended Practice for Architectural Description of Software-Intensive*
225 *Systems*, American National Standards Institute/Institute for Electrical and
226 *Electronics Engineers*, September 21, 2000.
- 227 [ISO/IEC 10746] International Organization for Standardization and International Electrotechnical
228 Commission, *Information Technology – Open Distributed Processing –*
229 *Reference Model*, September 15, 1996.
- 230 [ISO/IEC 42010] International Organization for Standardization and International Electrotechnical
231 Commission, *System and software engineering – Recommended practice for*
232 *architectural description of software-intensive systems*, July 15, 2007.
- 233 [RFC2119] S. Bradner, *Key words for use in RFCs to Indicate Requirement Levels*,
234 <http://www.ietf.org/rfc/rfc2119.txt>, IETF RFC 2119, March 1997.
- 235 [SOA-RM] OASIS Standard, "Reference Model for Service Oriented Architecture 1.0, 12
236 October 2006. <http://docs.oasis-open.org/soa-rm/v1.0/soa-rm.pdf>
- 237 [UML 2] *Unified Modeling Language: Superstructure*, Ver. 2.1.1, OMG Adopted
238 Specification, OMG document formal/2007-02-05, Object Management Group,
239 Needham, MA, February 5, 2007.
- 240 [WSA] David Booth, et al., "Web Services Architecture", W3C Working Group Note,
241 World Wide Web Consortium (W3C) (Massachusetts Institute of Technology,
242 European Research Consortium for Informatics and Mathematics, Keio
243 University), February, 2004. <http://www.w3.org/TR/2004/NOTE-ws-arch-20040211/>

245 **1.6.2 Non-Normative References**

246 **[BLOOMBERG/SCHMELZER]**

247 Jason Bloomberg and Ronald Schmelzer, *Service Orient or Be Doomed!*, John
248 Wiley & Sons: Hoboken, NJ, 2006.

249 **[DCMI]** Dublin Core Metadata Initiative, <http://dublincore.org>.

250 **[IEEE-829]** *IEEE Standard for Software Test Documentation*, Institute for Electrical and
251 Electronics Engineers, 16 September 1998

252 **[ISO 11179]** ISO/IEC 11179, Information Technology -- Metadata registries (MDR), accessible
253 from <http://metadata-standards.org/11179/>.

254 **[NEWCOMER/LOMOW]**

255 Eric Newcomer and Greg Lomow, *Understanding SOA with Web Services*,
256 Addison-Wesley: Upper Saddle River, NJ, 2005.

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

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

261 **[ISO/IEC 27002]** International Organization for Standardization and International Electrotechnical
262 Commission, *Information technology -- Security techniques -- Code of practice
263 for information security management*, 2007

264 **[SOA NAV]** Heather Kreger and Jeff Estefan (Eds.), "Navigating the SOA Open Standards
265 Landscape Around Architecture," Joint Paper, The Open Group, OASIS, and
266 OMG, July 2009. [http://www.oasis-
267 open.org/committees/download.php/32911/wp_soa_harmonize_d1.pdf](http://www.oasis-open.org/committees/download.php/32911/wp_soa_harmonize_d1.pdf)

268 2 Architectural Goals and Principles

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

271 2.1 Goals and Critical Success Factors of the Reference 272 Architecture Foundation

273 There are three principal goals:

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

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

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

286 These goals and success factors are expanded in the following subsections.

287 2.1.1 Goals

288 2.1.1.1 Effectiveness

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

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

295 2.1.1.2 Confidence

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

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

302 2.1.1.3 Scalability

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

307 2.1.2 Critical Success Factors

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

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

314 2.1.2.1 Action

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

318 2.1.2.2 Trust

319 The viability of a SOA ecosystem depends on participants being able to effectively measure the
320 trustworthiness of the system and of participants. Trust is a private assessment of a participant's belief in
321 the integrity and reliability of the SOA ecosystem (see Section 3.2.5.1).

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

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

330 2.1.2.3 Interaction

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

334 2.1.2.4 Control

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

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

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

346 2.2 Principles of this Reference Architecture Foundation

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

348 Technology Neutrality

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

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

357 Implications: The Reference Architecture Foundation must be technology neutral, meaning that we
358 assume that technology will continue to evolve, and that over the lifetime of this
359 architecture that multiple, potentially competing technologies will co-exist. Another
360 immediate implication of technology independence is that greater effort is needed on the
361 part of architects and other decision makers to construct systems based on this
362 architecture.

363 Parsimony

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

366 Rationale: The hallmark of good design is parsimony, or "less is better." It promotes better
367 understandability or comprehension of a domain of discourse by avoiding gratuitous
368 complexity, while being sufficiently rich to meet requirements.

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

370 Distinction of Concerns

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

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

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

390 Applicability

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

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

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

401 **3 Participation in a SOA Ecosystem View**

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

Comment [PFB2]: Do we want these quotations in a standards document?

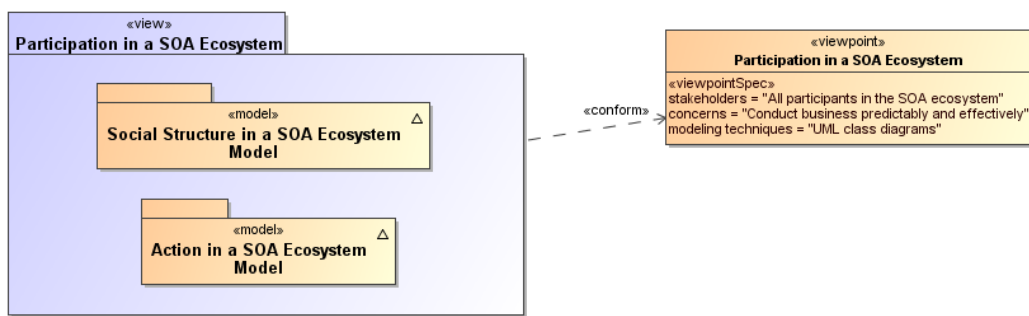
413 The *Participation in a SOA Ecosystem* view in the SOA-RAF focuses on the constraints and context in
414 which people conduct business using a SOA-based system. By business we mean any shared activity
415 whose objective is to satisfy particular **needs** of each participant. To effectively employ the SOA
416 paradigm, the architecture must take into account the fact and implications of different **ownership**
417 domains, and how best to organize and utilize capabilities that are distributed across those different
418 ownership domains. These are the main architectural issues that the Participation in a SOA Ecosystem
419 view tries to address.

420 The subsections below expand on the abstract Reference Model by identifying more fully and with more
421 specificity what challenges need to be addressed in order to successfully apply the SOA paradigm.
422 Although this view does not provide a specific recipe, it does identify the important things that need to be
423 considered and resolved within an ecosystem context.

424 The main models in this view are:

- 425 • The **SOA Ecosystem Model** introduces the main relationships between the social structure and
426 the SOA-based System, as well as the key role played by the hybrid concept of participant in
427 both.
- 428 • the **Social Structure in a SOA Ecosystem Model** introduces the key elements that underlie the
429 relationships between participants and that must be considered as pre-conditions in order to
430 effectively bring needs and capabilities together across **ownership boundaries**;
- 431 • the **Action in a SOA Ecosystem Model** introduces the key concepts involved in service **actions**,
432 and shows how **joint action** and **real-world effect** are the target outcomes that motivate
433 interacting in a SOA ecosystem.

Comment [PFB3]: Issue 32, part



Comment [PFB4]: 2012-03-02: Need to add "SOA Ecosystem Model" as new model in <<view>> package

434 *Figure 1 - Model elements described in the Participation in a SOA Ecosystem view*

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

439 The dominant mode of **communication** within a SOA ecosystem is electronic, supported by IT resources
 440 and artifacts. The **stakeholders** (see next section) are nonetheless people: since there is inherent
 441 indirection involved when people and systems interact using electronic means, we lay the foundations for
 442 how *communication* can be used to represent and enable action. However, it is important to understand
 443 that these communications are usually a means to an end and not the primary interest of the participants
 444 of the ecosystem.

Comment [PFB5]: Issue 32, part


445 3.1 SOA Ecosystem Model

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

450 Together, these ideas describe an environment in which business functions (realized in the form of
 451 services) address business needs. Service implementations utilize capabilities to produce specific (real
 452 world) effects that fulfill those business needs. Both those using the services, and the capabilities
 453 themselves, may be distributed across ownership domains, with different **policies** and conditions of use
 454 in force – this environment is referred to as a **SOA Ecosystem** and is modeled in Figure 2.

455 The role of a service in a SOA Ecosystem is to enable effective **business solutions** in this environment.
 456 Any technology system created to deliver a service in such an environment is referred to as a **SOA-**
 457 **based system**. SOA is thus a paradigm that guides the identification, design, implementation (i.e.,
 458 organization), and utilization of such services. SOA-based systems act as technology-based proxies for
 459 activity that would otherwise be carried out within and between social structures.

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



Comment [PFB6]: Issue 32, part

466 Figure 2 - SOA Ecosystem Model

467 SOA Ecosystem

468 An environment encompassing one or more **social structure(s)** and **SOA-based system(s)** that
 469 interact together to enable effective **business solutions**

470 SOA-based System

471 A technology system created to deliver a service within a **SOA Ecosystem**

472 Social Structures are defined and described in more detail in the next model, shown in Figure 3.

473 **Stakeholders, Actors, and Participants** are formally defined in Section 3.2.1.

474 Participants (as stakeholders and as actors), SOA-based systems, and the environment (or context)
 475 within which they all operate, taken together forms the SOA ecosystem. Participants (or their **delegates**)
 476 interact with a SOA-based system - in the role of actors - and are also members of a social structure - in
 477

478 the role of stakeholders. Here we explicitly note that stakeholders and, thus, participants are people³
479 because machines alone cannot truly have a stake in the outcomes of a social structure. Delegates may
480 be human and nonhuman but are not directly stakeholders. Stakeholders, both Participants and **Non-**
481 **participants**, may potentially benefit from the services delivered by the SOA-based system. Again, this is
482 discussed more fully in Section 3.2.1.

483 The SOA ecosystem may reflect the SOA-based activities within a particular enterprise or of a wider
484 network of one or more enterprises and individuals; these are modeled in and discussed with respect to
485 Figure 3. Although a SOA-based system is essentially an IT concern, it is nonetheless a system
486 engineered deliberately to be able to function in a SOA ecosystem. In this context, a service is the
487 mechanism that brings a SOA-based system capability together with stakeholder needs in the wider
488 ecosystem.

489 Several interdependent concerns are important in our view of a SOA ecosystem. The ecosystem includes
490 stakeholders who are participants in the development, deployment and governance and use of a system
491 and its services; or who may not participate in certain activities but are nonetheless affected by the
492 system. Actors – whether stakeholder **participants** or delegates who act only on behalf of participants
493 (without themselves having any stake in the actions that they have been tasked to perform) – are
494 engaged in **actions** which have an impact on the real world and whose meaning and intent are
495 determined by implied or agreed-to semantics. This is discussed further in relation to the model in Figure
496 4 and elaborated more fully in Section 3.3.

497 **3.2 Social Structure in a SOA Ecosystem Model**

498 The Social Structure Model explains the relationships between stakeholders and the social context in
499 which they operate, within and between distinct boundaries. It is also the foundation for understanding
500 security, governance and management in the SOA ecosystem.

501 Actions undertaken by people (whether natural or legal persons) are performed in a *social context* that
502 defines the relationships between *them*. That context is **provided by social structures existing in society**
503 and the roles played by each person ~~is~~ as **a-stakeholders in those structures.**

504 Whether informal peer groups, associations, enterprises, corporations, government agencies, or entire
505 nations, these structures interact with each other in the world, using treaties, contracts, market rules,
506 handshakes, negotiations and – when necessary – have recourse to arbitration and legislation. They
507 interact because there is a mutual benefit in doing so: one has something that the other can provide.
508 They interact across defined or implicit **ownership boundaries** that define the limits of one structure (and
509 the limits of its **authority, responsibilities, capabilities, etc.**) and the beginning of another.

510 Social structures, together with their **constitution**, their stakeholders, their mission and goals, need
511 therefore to be understood when examining the role that technology plays. Technology systems play an
512 increasing role in carrying out many of the functions performed by such structures and therefore model
513 real-world procedures. The technology systems serve as proxies in digital space for these real-world
514 structures and procedures. The SOA paradigm is particularly concerned with designing, configuring and
515 managing such systems across ownership boundaries precisely because this mirrors the real-world
516 interactions between discrete structures and across their ownership boundaries.

517 A stakeholder in a social structure will be involved in many “actions” that do not involve a SOA-based
518 system. Although such actions and the roles relating to them are outside the scope of this Reference
519 Architecture Foundation, they may nonetheless result in constraining or otherwise impacting a given SOA
520 ecosystem – for example, a new item of legislation that regulates service interactions. The terms ‘actor’
521 and ‘action’ used throughout the document refer thus only to SOA-based systems.

³ ‘People’ and ‘person’ must be understood as both humans and ‘legal persons’, such as companies, who have rights and responsibilities similar to ‘natural persons’ (humans)

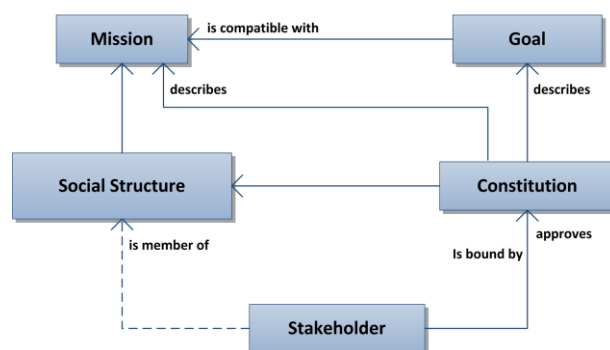


Figure 3 - Social Structure Model

Social Structure

A nexus of relationships amongst people brought together for a specific purpose, the structure's mission.

Comment [PFB7]: Issue 287, part

The social structure is established with an implied or explicitly defined mission, usually reflected in the goals laid down in the social structure's constitution or other 'charter'. Although goals are often expressed in terms of general ambitions for the social structure's work or of desired end states, objectives are expressed more formally in terms of specific, measurable, and achievable action required to realize those states. Action in the context of a social structure is discussed in Section 3.3.

Comment [PFB8]: Issues 28, 53

A social structure may involve any number of persons as stakeholders and a large number of different relationships may exist among them. The organizing principle for these relationships is the social structure's mission. Any given person can be a stakeholder in multiple social structures and a social structure itself can be a stakeholder in its own right as part of a larger one or in another social structure entirely. These multiple roles can result in disagreements, particularly when the mission or goals of different social structures do not align.

A social structure can take different forms. An enterprise is a common kind of social structure with its distinct legal personality; an online community group might represent a social structure of peers that is very loose, albeit with a shared mission. A market represents a social structure of buyers and sellers. Legislation in different geo-political areas (from local and regional to national or global) provides a framework in which social structures can operate.

A social structure will further its goals in one of two ways:

- by acting alone, using its own resources;
- interacting with other structures and using their resources.

Many interactions take place within social structures. Some interactions may or may not cross ownership boundaries depending on the scale and internal organization of the structure (an enterprise, for example, can itself be composed of sub-enterprises). Our focus is on interactions between social structures, particularly as they determine the way that technology systems need to interact. Systems that are designed to do this are SOA-based systems.

The nature and extent of the interactions that take place will reflect, often implicitly, degrees of trust between people and the very specific circumstances of each person at the time, and over the course, of their interactions. It is in the nature of a SOA ecosystem that these relationships are rendered more explicit and are formalized as a central part of what the [SOA-RM] refers to as Execution Context.

Comment [PFB9]: Issue 44, part

The validity of the interactions between social structures is not always clear and is often determined ultimately by relevant legislation. For example, when a customer buys a book over the Internet, the validity of the transaction may be determined by the place of incorporation of the book vendor, the residence of the buyer, or a combination of both. Such legal jurisdiction qualification is typically buried in the fine print of the service description.

560 **Constitution**

561 | A set of rules, written or unwritten, that formalize the mission, goals, scope, and functioning of a
562 **social structure**.

563 Every social structure functions according to rules by which people interact with each other within the
564 structure. In some cases, this is based on an explicit agreement; in other cases, participants behave as
565 though they agree to the constitution without a formal agreement. In still other cases, participants abide
566 by the rules with some degree of reluctance. In all cases, the constitution may change over time; in those
567 cases of implicit agreement, the change can occur quickly. [Section 5.1 contains a detailed discussion of](#)
568 [governance and SOA.](#)

Comment [KJL10]: Issue 31 for edits in this paragraph

569 **3.2.1 Stakeholders, Participants, Actors and Delegates**

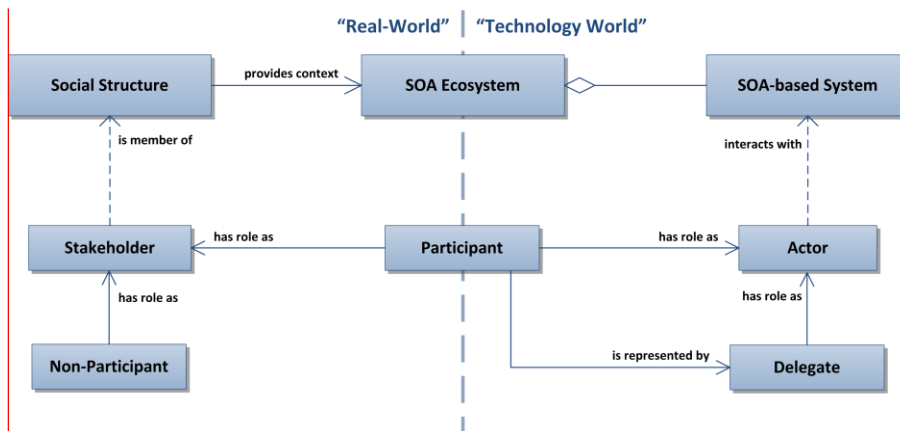
570 A social structure represents the interests of a collection of people who have **rights** and **responsibilities**
571 within the structure. People have a “stake” in such a social structure, and when that social structure is
572 part of a SOA Ecosystem, the people continue to interact through their roles as stakeholders. In addition,
573 people – either directly or through their delegates - interact with SOA-based (technology) systems. Here,
574 the people interact through their roles as actors interacting with specific system-level activity.

575 A person who participates in a social structure as a stakeholder *and* interacts with a SOA-based system
576 as an actor is defined as an ecosystem **Participant**. The concept of participant is particularly important as
577 it reflects a hybrid role of a Stakeholder concerned with expressing needs and seeing those needs fulfilled
578 *and* an Actor directly involved with system-level activity that result in necessary effects.

579 The hybrid role of Participant provides a bridge between social structures within the wider (real-world)
580 ecosystem – in particular the world of the stakeholder – and the more specific (usually technology-
581 focused) system – the world of the actor.

582 The concept of the ecosystem therefore embraces all aspects of the “real world”, human-centered, social
583 structures that are concerned with business interactions together with the technology-centered SOA-
584 based system that deliver services:

585
586



Comment [PFB11]: Issue 32, part; Issue 280, part;

587
588 *Figure 4 – Stakeholders, Actors, Participants and Delegates*

589 **Stakeholder**

590 | A person with an interest (a ‘stake’) in a social structure.

591 Not all stakeholders necessarily participate in all activities in the SOA ecosystem; indeed, the interest of
592 non-participant stakeholders may be to realize the benefits of a well-functioning ecosystem and not suffer
593 unwanted consequences. Non-participant stakeholders cannot all or always be identified in advance but
594 due account is often taken of such stakeholder types, including potential customers, beneficiaries, and

595 other affected third parties. A stakeholder may be a participant with respect to some activities and a non-
596 participant with respect to others.

597 **Actor**

598 A role played either by a Participant or its Delegate and that interacts with a SOA-based
599 system.

600 **Participant**

601 A person who plays a role both in the SOA ecosystem as a stakeholder and with the SOA-
602 based system as an actor either

- 603 • directly, in the case of a human participant; or
- 604 • indirectly, via a delegate.

605 Not all participants are necessarily benign to the social structure: such “negative stakeholders” might
606 deliberately seek a negative impact on the ecosystem (such as hackers or criminals) and social structures
607 will work to ensure that they are not able to operate as welcome participants.

608 **Non-Participant**

609 A person who plays no role as a participant in a social structure’s activities but nonetheless
610 has an interest in, or is affected by, such activities.

611 **Delegate**

612 A role played by a human or an automated or semi-automated agent and acting on behalf of a
613 participant but not directly sharing the participant’s stake in the outcome.

614 Many actors interact with a SOA-based system, including software agents that permit people to offer, and
615 interact with, services; delegates that represent the interests of other participants; or security agents
616 charged with managing the security of the ecosystem. Note that automated agents are *always* delegates,
617 in that they act on behalf of a participant.

618 In the different models of the SOA-RAF, the term actor is used when action is being considered at the
619 level of the SOA-based system and when it is not relevant who is carrying out the action. However, if the
620 actor is acting explicitly *on behalf of* a participant, then we use the term delegate. This underlines the
621 importance of delegation in SOA-based systems, whether the delegation is of work procedures carried
622 out by human agents who have no stake in the actions with which they are tasked but act on behalf of a
623 participant who does; or whether the delegation is performed by technology (automation). On the other
624 hand, if it is important to emphasize that when the actor is also a stakeholder in the ecosystem, then we
625 use the term participant. This also underlines the pivotal role played by a participant, in a unique position
626 between the social structure and the SOA-based system, in the broader ecosystem.

627 The difference between a participant and a delegate is that a delegate acts on behalf of a participant and
628 must have the authority to do so. Because of this, every social structure needs to clearly define the roles
629 assigned to actors (whether participants or delegates) in carrying out activity within its domain.

630 **3.2.2 Social Structures and Roles**

631 Social structures are abstractions: they cannot directly perform actions with SOA-based systems – only
632 actors can, whether they be participants acting under their own volition or delegates (human or not)
633 simply following the instructions of participants. An actor advances the objectives of a social structure
634 through its interaction with SOA-based systems, influencing actions that deliver results. The specifics of
635 the interaction depend on the roles defined by the social structure that the actor may assume or have
636 conferred and the nature of the relationships between the stakeholders concerned. These relationships
637 can introduce constraints on an actor when engaged in an action. These points are illustrated in Figure 5.

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

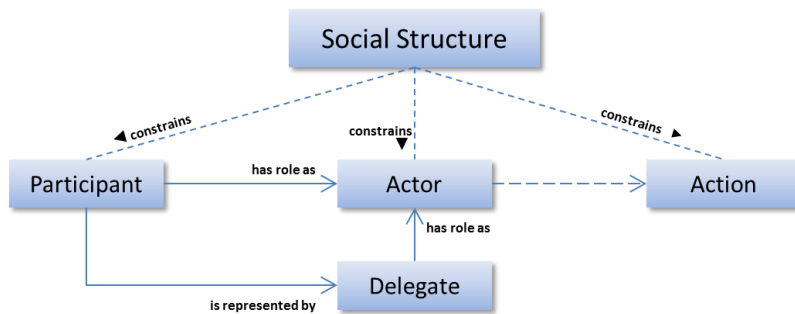
640 **3.2.2.1 Authority, Rights, and Responsibilities**

641 One participant with appropriate authority in the social structure may formally designate a role for a
642 delegate or another participant, with associated rights and responsibilities, and that authority may even

643 qualify a period during which the designated role may be valid. In addition, while many roles are clearly
 644 identified, with appropriate names and definitions of responsibilities, it is also possible to separately
 645 bestow rights, bestow or assume responsibilities and so on, often in a temporary fashion. For example,
 646 when a company president delegates certain responsibilities on another person, this does not imply that
 647 the other person has become company president. Likewise, a company president may bestow on
 648 someone else her role during a period of time that she is on vacation or otherwise unreachable with the
 649 understanding that she will re-assume the role when she returns from vacation.

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

654 The social structure is responsible for establishing the authority by which actors carry out actions in line
 655 with defined constraints:



656
 657 *Figure 5 - Social Structures, Roles and Action*

658 **Authority**

659 A **right** conferred on a participant to ensure that actions are carried out consistent with the
 660 objectives of a social structure.

661 Actions are carried out by actors, either participants themselves or delegates acting on their behalf, by
 662 interacting with the SOA-based system.

663 **Right**

664 A predetermined **permission** conferred upon an **actor** to perform some **action** or assume a role
 665 in relation to the **social structure**.

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

668 **Responsibility**

669 A predetermined **obligation** on a **participant** to ensure that some **action is performed** or assume
 670 a role in relation to other **participants**.

671 Responsibility implies human agency and thus aligns with participants and potentially human delegates
 672 but not with nonhuman delegates. This applies even if the consequences of such responsibility can
 673 impact other (human and non-human) actors. Having authority often implies having responsibility.

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

676 **3.2.2.2 Permissions and Obligations**

677 People will assume and perform roles according to their actual or perceived rights and responsibilities,
 678 with or without explicit authority. In the context of a SOA ecosystem, human abilities and skills are
 679 relevant as they equip individuals with knowledge, information and tools that may be necessary to have
 680 meaningful and productive interactions with a view to achieving a desired outcome. For example, a
 681 person who needs a particular book, and has both the right and responsibility of purchasing the book from

682 a given bookseller, will not have that need met from the online delegate of that bookstore if he does not
683 know how to use a web browser. Equally, just because someone does have the requisite knowledge or
684 skills does not entitle them *per se* to interact with a specific system.

685 Assuming or accepting rights and responsibilities depend on two important types of constraints that are
686 relevant to a SOA ecosystem: Permission and Obligation.

687 **Permission**

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

690 Note that permissions are distinct from ability, which refers to whether an actor has the capacity to
691 perform the action. Permission does not always involve acting on behalf of anyone, nor does it imply or
692 require the capacity to perform the action.

693 **Obligation**

694 A constraint that prescribes the **actions** that an **actor** must (or must not) perform and/or the
695 **states** the actor must (or must not) attain or maintain.

696 An example of obligations is the case where the service **consumer** and **provider** (see below) have
697 entered into an agreement to provide and consume a service such that the consumer is obligated to pay
698 for the service and the provider is obligated to provide the service – based on the terms of the contract.

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

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

705 **3.2.2.3 Service Roles**

706 As in roles generally, a participant can play one or more in the SOA ecosystem, depending on the
707 context. A participant may be playing a role of a service provider in one relationship while simultaneously
708 playing the role of a consumer in another. Roles inherent to the SOA paradigm include **Consumer**,
709 **Provider**, **Owner**, and **Mediator**.

710 **Provider**

711 A role assumed by a **participant** who is offering a service.

712 **Consumer**

713 A role assumed by a **participant** who is interacting with a service in order to fulfill a **need**.

714 **Mediator**

715 A role assumed by a **participant** to facilitate interaction and connectivity in the offering and use of
716 services.

717 **Owner**

718 A role assumed by a **participant** who is claiming and exercising **ownership** over a service.




Figure 6 - Roles in a Service

719
 720
 721 Service consumers typically initiate interactions, but this is not necessarily true in all situations.
 722 Additionally, several stakeholders may be involved in a service interaction supporting a given consumer.
 723 The roles of service provider and service consumer are often seen as symmetrical, which is also not
 724 entirely correct. A stakeholder tends to express a **Need** in non-formal terms: "I want to buy that book".
 725 The type of need that a service is intended to fulfill has to be formalized and encapsulated by designers
 726 and developers as a **Requirement**. This Requirement should then be reflected in the target service, as a
 727 **Capability** that, when accessed via a service, delivers a **Real World Effect** to an arbitrary consumer:
 728 "The chosen book is ordered for the consumer." It thus fulfills the need that has been defined for an
 729 archetypal consumer.

730 Specific and particular customers may not experience a need exactly as captured by the service: "I don't
 731 want to pay that much for the book", "I wanted an eBook version", etc. There can therefore be a process
 732 of implicit and explicit negotiation between the consumer and the service, aimed at finding a 'best fit'
 733 between the consumer's specific need and the capabilities of the service that are available and consistent
 734 with the service provider's offering. This process may continue up until the point that the consumer is able
 735 to accept what is on offer as being the best fit and finally 'invokes' the service. 'Execution context' has
 736 thus been established. Conditions and agreements that contribute to the execution context are discussed
 737 throughout this Reference Architecture.

738 Service mediation by a participant can take many forms and may invoke and use other services in order
 739 to fulfill such mediation. For example, it might use a service registry in order to identify possible service
 740 partners; or, in our book-buying example, it might provide a price comparison service, suggest alternative
 741 suppliers, different language editions or delivery options.

3.2.3 Needs, Requirements and Capabilities

Comment [PFB12]: Moved from Action Model section

742
 743 Participants in a SOA ecosystem often need other participants to *do* something, leveraging a **capability**
 744 that they do not themselves possess. For example, a customer requiring a book may call upon a service
 745 provider to deliver the book. Likewise, the service provider needs the customer to pay for it.


746 There is a reason that participants are engaged: they have different **needs** and have or apply different
 747 capabilities for satisfying them. These are core to the concept of a service. The SOA-RM defines a
 748 service as "the mechanism by which needs and capabilities are brought together". This idea of services
 749 being a mechanism "between" needs and capabilities was introduced in order to emphasize capability as
 750 the notional or existing business functionality that would address a well-defined need. Service is therefore
 751 the *implementation* of such business functionality *such that it is accessible* through a well-defined

752 interface. A capability that is isolated (i.e., it is inaccessible to potential consumers) is emphatically not a
753 service.

754 **Business functionality**

755 A defined set of business-aligned tasks that provide recognizable business value to consumer
756 **stakeholders** and possibly others in the **SOA ecosystem**.

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



762
763 *Figure 7 - Cycle of Needs, Requirements, and Fulfillment*

764 In a SOA context, the stakeholder expresses a need (for example, the consumer who states that “I want
765 to buy a book”) and looks to an appropriate service to fulfill that need and assesses issues such as the
766 trustworthiness, intent and **willingness** of a particular provider. This ecosystem communication continues
767 up to the point when the stakeholder is ready to act. The stakeholder will then interact with a provider by
768 invoking a service (for example, by ordering the book using an online bookseller) and engaging in
769 relevant actions with the system (at this point, in a role as an actor, interacting with the system through a
770 browser or mobile device, validating the purchase, submitting billing and delivery details) with a view to
771 achieving the desired real world effect (having the book delivered).

772 **Need**

773 A general statement expressed by a **stakeholder** of something deemed necessary.

774 A need may be formalized as one or more requirements that must be fulfilled in order to achieve a stated
775 goal.

776 **Requirement**

777 A formal statement of a desired result (a **real world effect**) that, if achieved, will satisfy a **need**.

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

781 **Capability**

782 An ability to deliver a **real world effect**.

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

786 **Real World Effect**

787 A measurable change to the shared state of pertinent entities, relevant to and experienced by
788 specific stakeholders of an ecosystem.

Comment [PFB13]: Issues 56 and 57 – text moved from Action section

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

Comment [KJL14]: Moved to complete comments 56 & 57


792 [Objectives refer to real world effects that participants believe are achievable by a specific action or set of](#)
793 [actions that deliver appropriate changes in shared state, as distinct from a more generally stated 'goal'.](#)
794 [For example, someone may wish to have enough light to read a book. In order to satisfy that goal, the](#)
795 [reader walks over to flip a light switch. The *objective* is to change the state of the light bulb, by turning on](#)
796 [the lamp, whereas the *goal* is to be able to read. The *real world effect* is more light being available to](#)
797 [enable the person to read.](#)

798 [While an effect is any measurable change resulting from an action, a SOA ecosystem is concerned more](#)
799 [specifically with real world effects.](#)

800 3.2.4 Resource and Ownership

801 3.2.4.1 Resource

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



804
805 *Figure 8 - Resources*

806 Resource

807 An identifiable entity that has value to a **stakeholder**.

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

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

814 Identifier

815 [A sequence of characters that unambiguously indicates a particular resource.](#)

816 [Identifiers are assigned by social structures according to context, policies and procedures considered](#)
817 [sufficient for that structure's purposes.](#)

818 [For example, a group of otherwise unrelated humans are all, in a given context, employees of a particular](#)
819 [company and managed there as human resources. That company's policy is to assign each employee a](#)
820 [unique identifier number and has processes in place to do this, including verifying documentary evidence](#)
821 [\(such as a birth certificate or ID\). Each set of policies and procedures will reflect the needs of the social](#)
822 [structure for its particular context. Resources are typically used or managed by different stakeholder](#)
823 [groups, each of which may need to identify those resources in some particular way. As such, a given](#)
824 [resource may have multiple identifiers, each valid for a different context. In a SOA ecosystem, it is good](#)

825 | [practice to use globally unique identifiers \(for example, Internationalized Resource Identifiers, or IRIs\)](#)
826 | [irrespective of any other resource identifier that might be in use for a particular context.](#)

827 | The ability to identify a resource is important in interactions to determine such things as rights and
828 | authorizations, to understand what functions are being performed and what the results mean, and to
829 | ensure repeatability or characterize differences with future interactions. Many interactions within a SOA
830 | ecosystem take place across ownership boundaries. Identifiers provide the means for all resources
831 | important to a given SOA-based system to be *unambiguously* identifiable at any moment and in any
832 | interaction.

833 | 3.2.4.2 Ownership

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

836 | Typically, the ownership relationship is one of control: the owner of a resource can control some aspect of
837 | the resource.

838 | Ownership

839 | A set of claims, expressed as **rights** and **responsibilities** that a **stakeholder** has in relation to a
840 | **resource**; it may include the right to transfer that ownership, or some subset of rights and
841 | responsibilities, to another entity.

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

848 | A stakeholder who owns a resource may delegate some or all of these rights and responsibilities to
849 | others, but typically retains the responsibility to see that the delegated rights and responsibilities are
850 | exercised as intended

851 | A crucial property that distinguishes ownership from a more limited right to use is the right to transfer
852 | rights and responsibilities totally and irrevocably to another. When [participants](#) use but do not own a
853 | resource, [they](#) may not [be allowed to](#) transfer the right to use the resource to a third [participant](#). The
854 | owner of the resource maintains the rights and responsibilities of being able to authorize others to use the
855 | owned resource.

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

860 | Ownership Boundary

861 | The extent of **ownership** asserted by a **stakeholder** [or a social structure](#) over a set of
862 | **resources** and for which **rights** and **responsibilities** are claimed and (usually) recognized by
863 | other stakeholders.

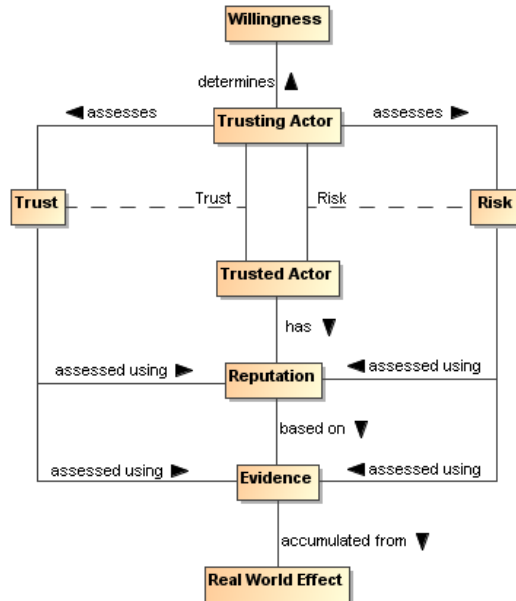
864 | 3.2.5 Establishing Execution Context

865 | In a SOA ecosystem, providers and consumers of services may be, or may be acting on behalf of,
866 | different owners, and thus the interaction between the provider and the consumer of a given service [may](#)
867 | necessarily cross an ownership boundary. It is important to identify these ownership boundaries in a SOA
868 | ecosystem [and](#) successfully crossing them [in a key aspect of establishing execution context](#). [This is turn](#)
869 | requires [that](#) the elements identified in the following sections be addressed.

Comment [PFB15]: Issue 245
Former sections 3.2.5 (“Trust and
Risk”) thru 3.2.8 (“Semantics and
Semantic Engagement”) are now
sub-sections under this heading

870 **3.2.5.1 Trust and Risk**

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



872
873 *Figure 9 - Willingness and Trust*

874 **Willingness**

875 The internal commitment of a human **actor** (or of an automated non-human agent acting on a
876 participant's behalf) to carry out its part of an interaction.

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

881 **Trust**

882 The private assessment or internal perception of one **actor** that another actor will perform
883 **actions** in accordance with an assertion regarding a desired **real world effect**.

884 **Risk**

885 The private assessment or internal perception of the likelihood that certain undesirable **real world**
886 **effects** will result from **actions** taken and the consequences or implications of such.

887 Trust is involved in all interactions and each actor will play a role as either (or alternately) a "trusting" actor
888 and a "trusted" actor. These roles are needed in order that all actors can trust all others in any given
889 interaction, at least to the extent required for continuance of the interaction. The degree and nature of that
890 trust is likely to be different for each actor, most especially when those actors are in different ownership
891 boundaries.

Comment [PFB16]: Issue 44, part

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

896

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

Comment [PFB17]: Issue 44, part

901 Trust is based on the confidence that the trusting actor has accurately and sufficiently gathered and
902 assessed evidence to the degree appropriate for the situation being assessed.

Comment [PFB18]: Issue 44, part

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

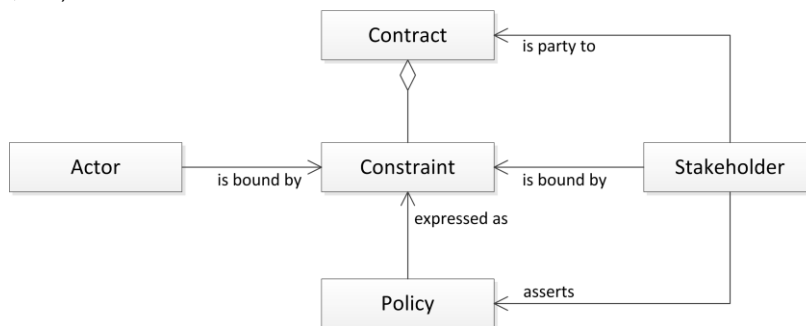
906 The relevance of trust to interaction depends on the assessment of risk. If there is little or no perceived
907 risk, or the risk can be covered by another party who accepts responsibility for it, then the degree of trust
908 may be less or not relevant in assessing possible actions. For example,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

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

Comment [PFB19]: Issue 44, part

918 3.2.5.2 Policies and Contracts

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



924
925 Figure 10 – Policies, Contracts and Constraints

Comment [PFB20]: New diagram

926 Policy

927 An expression of constraints made by a stakeholder that the stakeholder commits to uphold and,
928 if desired or necessary, enforce. The constraints are usually stated as permissions and
929 obligations that affect the behavior of stakeholders or of any actor acting on their behalf.

Comment [PFB21]: Issue 282

930 Policies have an owner – the stakeholder who asserts and takes responsibility for the policy. This owner
931 may or may not be the owner of the object of the policy. These constraints may affect the stakeholder
932 asserting the policy or any other stakeholder involved. The constraints themselves represent some
933 measurable limitation on the state or behavior of the object of the policy, or of those who interact with it.

934 **Contract**

935 An agreement made by two or more **participants** (the contracting parties) on a set of conditions
936 (or contractual terms) together with a set of constraints that govern their behavior and/or **state** in
937 fulfilling those conditions.

938 A service provider's policy may become a service provider/consumer contract when a service consumer
939 agrees to the provider's policy. That agreement may be formal, or may be informal. If a consumer's policy
940 and a provider's policy are mutually exclusive, then some form of negotiation (involving human
941 interactions) or mediation must resolve the mutual exclusion before the service consumer/provider
942 interaction can occur. Note, this also applies if the consumer instead of the provider introduces the policy.

Comment [PFB22]: Issue 46

943 Both policies and contracts imply a desire to see constraints respected and enforced. Stakeholders are
944 responsible for ensuring that any constraints in the policy or contract are enforced, although the actual
945 enforcement may be delegated to a different mechanism. A contract does not necessarily oblige the
946 contracting parties to act (for example to use a service) but it does constrain how they act if and when the
947 condition covered by the contract occurs (for example, when a service is invoked and used).

948 The realization of policies and contracts is discussed in Section 4.4 and contracts in the context of
949 management are discussed in Section 5.3.4.

950 **3.2.5.3 Communication**

951 **Communication**

952 A process involving the exchange of information between a sender and one or more recipients
953 and that ideally culminates in mutual understanding between them.

954 A communication involves a message, a sender of the message and at least one intended recipient, who
955 must be able to correctly interpret the message – or at least those parts of the message relevant to
956 sender and recipient in the particular context. Each must perform its respective role in order for the
957 communication to be successful; failing which communication is not effective.

Comment [PFB23]: Issue 48

Comment [PFB24]: Issue 247

958 A communication may involve any number of recipients. In some situations, the sender may not be aware
959 of the recipient. However, without both a sender and a recipient, there is no communication. A given
960 communication can be a simple one-way transmission and does not necessarily involve interaction
961 between the parties require a response; it can be a simple one-way transmission requiring no further
962 action by the recipient. However, interaction does, necessarily, involve communication.

Comment [KJL25]: Comment 50

963 Message interpretation can itself be characterized in terms of **semantic engagement**: the proper
964 understanding of a message in a given context.

965 We can characterize the necessary modes of interpretation in terms of a shared understanding of a
966 common vocabulary (or mediation among vocabularies) and of the purpose of the communication. More
967 formally, we can say that a communication has a combination of message and purpose.

968 In a SOA ecosystem, senders and recipients can be stakeholders, participants or actors, depending on
969 whether execution context is being established or a specific interaction with the SOA-based system is in
970 progress. Communications need not resemble human speech: indeed system-level machine-to-machine
971 communication is typically highly stylized in form. It may take a particular form and involve terms not
972 found in everyday human communication.

Comment [PFB26]: Issue 48

973 **3.2.5.4 Semantics and Semantic Engagement**

974 Shared understanding is vital to a trusted and effective ecosystem and is a prerequisite to joint action
975 being carried out as intended. Semantics are therefore pervasive throughout SOA ecosystems and
976 important in communications as described above, as well as a driver for policies and other aspects of the
977 ecosystem.

978 In order to arrive at a shared understanding wherever this is necessary within the ecosystem, a
979 message's recipient must effectively understand and process assertions ~~statements~~, made in the
980 sender's message, in a manner appropriate and sufficient to the particular context. Within a SOA-based
981 system, non-human actors must at least be able to parse a message correctly (syntax) and act on the
982 message's statements in a manner consistent with the sender's intent.

983 Understanding and interpreting those assertions in a SOA-based system allows all the actors in any
984 particular joint action to “know” what may be expected of them. An actor can potentially “understand” an
985 assertion in a number of ways, but it is specifically the process of arriving at a *shared* understanding that
986 is important in the ecosystem. This process is semantic engagement and it takes place in different forms
987 throughout the SOA ecosystem. It can be instantaneous or progressively achieved. Participants – who
988 play the role both as actors in the SOA-based system and as stakeholders in social structures and the
989 wider ecosystem – can be pivotal in resolving problems of understanding and determining when there is a
990 level of engagement appropriate and sufficient to the particular context.

991 **Semantic Engagement**

992 The process by which an **actor** engages with a set of assertions based on that actor’s
993 interpretation and understanding of those assertions.

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

998 The impact of any assertion can only be fully understood in terms of specific social contexts that
999 necessarily include the actors that are involved. For example, a policy statement that governs the actions
1000 relating to a particular resource may have a different impact or purpose for the participant that owns the
1001 resource than for the actor that is trying to access it: the former understands the purpose of the policy as
1002 a statement of enforcement - the latter understands it as a statement of constraint.

1003 **3.3 Action in a SOA Ecosystem Model**

1004 Participants cannot always achieve desired results by leveraging resources in their own ownership
1005 domain. This unfulfilled need leads them to seek and leverage services provided by other participants and
1006 using resources beyond their ownership and control. The participants identify service providers with which
1007 they think they can interact to achieve their objective and engage in joint action with those other actors
1008 (service providers) in order to bring about the desired outcome. The SOA ecosystem provides the
1009 environment in which this happens.

1010 An action model is put forth a-priori by the service provider, and is effectively an undertaking by the
1011 service provider that the actions – identified in the action model and invoked consistent with the process
1012 model – will result in the described real world effect. The action model describes the actions leading to a
1013 real-world effect. A potential service consumer – who is interested in a particular outcome to satisfy their
1014 need – must understand those actions as capable of achieving that desired outcome.

1015 When the consumer “invokes” a service, a joint action is started as identified in the action model,
1016 consistent with the temporal sequence as defined by the process model, and where the consumer and
1017 the provider are the two parties of the joint action. Additionally, the consumer can be assured that the
1018 identified real-world effects will be accomplished through evidence provided via the service description.

1019 Since the service provider does not know about all potential service consumers, the service provider may
1020 also describe what additional constraints are necessary in order for the service consumer to invoke
1021 particular actions, and thus participate in the joint action. These additional constraints, along with others
1022 that might not be listed, are preconditions for the joint action to occur and/or continue (as per the process
1023 model), and are referred to in the SOA-RM as execution context. Execution context goes all the way from
1024 human beings involved in aligning policies, semantics, network connectivity and communication
1025 protocols, to the automated negotiation of security protocols and end-points as the individual actions
1026 proceed through the process model.

1027 Also, it is important to note that both actions and real world effect are fractalrecursive in nature, in the
1028 sense that they can often be broken down into more and more granularity depending on how they are
1029 examined and what level of detail is important.

1030 All of these things are important to getting to the core of participants’ concern in a SOA ecosystem: the
1031 ability to leverage resources or capabilities to achieve a desired outcome, and in particular where those
1032 resources or capabilities do not belong to them or are beyond their direct control. i.e., that are outside of
1033 their ownership boundary.

1034 | In order to use such resources, participants must be able to identify their own needs; **state those needs** in
1035 | the form of requirements; compose **or identify a suitable** business solution **using** resources or capabilities
1036 | that will meet their needs; and engage in joint action – the coordinated set of actions that participants
1037 | pursue in order to achieve measurable results in furtherance of their goals.

Comment [PFB27]: Issue 53, part

1038 | In order to act in a way that is appropriate and consistent, participants must communicate with each other
1039 | about their own goals, objectives and policies, and those of others. This is the main concern of Semantic
1040 | Engagement.

1041 | A key aspect of joint action revolves around the trust that both parties must exhibit in order to participate
1042 | in the joint action. The willingness to act and a mutual understanding of both the information exchanged
1043 | and the expected results is the particular focus of Sections 3.2.5.1 and 3.2.5.4.

1044 | **3.3.1 Services Reflecting Business**

1045 | The SOA paradigm often emphasizes the interface through which service interaction is accomplished.
1046 | While this enables predictable integration in the sense of traditional software development, the prescribed
1047 | interface alone does not guarantee that services will be composable into business solutions.

1048 | **Business Solution**

1049 | A set of defined interactions that combine implemented or notional **business functionality** in
1050 | order to address a set of business needs.

1051 | **Composability**

1052 | The ability to combine individual services, each providing defined **business functionality**, so as
1053 | to provide more complex **business solutions**.

1054 | To achieve composability, capabilities must be identified that serve as building blocks for business
1055 | solutions. In a SOA ecosystem, these building blocks are captured as services representing well-defined
1056 | business functions, operating under well-defined policies and other constraints, and generating well-
1057 | defined real world effects. These service building blocks should be relatively stable so as not to force
1058 | repeated changes in the compositions that utilize them, but should also embody SOA attributes that
1059 | readily support creating compositions that can be varied to reflect changing circumstances.

1060 | The SOA paradigm emphasizes both composition of services and opacity of how a given service is
1061 | implemented. With respect to opacity, the SOA-RM states that the service could carry out its described
1062 | functionality through one or more automated and/or manual processes that in turn could invoke other
1063 | available services.


1064 | Any composition can itself be made available as a service and the details of the business functionality,
1065 | conditions of use, and effects are among the information documented in its service description.

1066 | Composability is important because many of the benefits of a SOA approach assume multiple uses for
1067 | services, and multiple use requires that the service deliver a business function that is reusable in multiple
1068 | business solutions. Simply providing a Web Service interface for an existing IT artifact does not, in
1069 | general, create opportunities for sharing business functions. Furthermore, the use of tools to auto-
1070 | generate service software interfaces will not guarantee services that can effectively be used within
1071 | compositions if the underlying code represents programming constructs rather than business functions. In
1072 | such cases, services that directly expose the software details will be as brittle to change as the underlying
1073 | code and will not exhibit the characteristic of loose coupling.

1074 | **3.3.2 Activity, Action, and Joint Action**

1075 | In general terms, entities act in order to **fulfill particular objectives**. More precisely, they generate activity.
1076 | An activity is made up of specific Actions (or other Activities) and is formally defined in **[ISO/IEC 10746]**
1077 | as “a single-headed directed acyclic graph of actions...”⁴ It is most clearly understood diagrammatically:

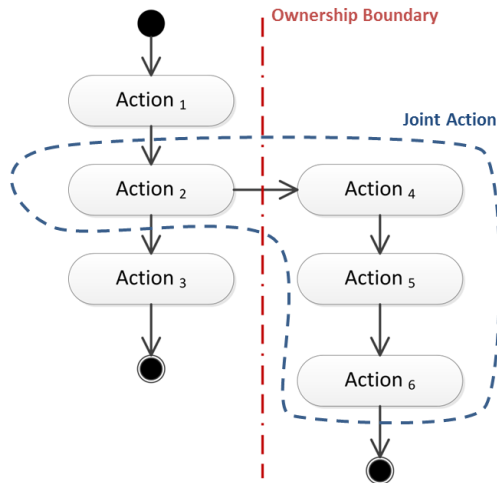
⁴ See **[ISO/IEC 10746] Part 2: Foundations**



1078
1079 *Figure 11: An Activity, expressed informally as a graph of Actions, with a single Start point and alternative End points*

1080 What constitutes an Action or an Activity will be a matter of context. For the SOA-RAF, an Action
1081 represents the smallest and most discrete activity that needs to be modeled for a given Viewpoint.

1082 The form of Activity that is of most interest within a SOA ecosystem is that involving Actions as defined
1083 below and their interaction across ownership boundaries (and thus involving interaction between more
1084 than one actor) – we call this **joint action**. In Figure 12 below, one line of activity (on the left) can be
1085 completed thru Action₃ without crossing any ownership boundary but the alternative path, starting at
1086 Action₄, can only be completed as a result of joint action across an ownership boundary:



1087
1088 *Figure 12: Activity involving Actions across an ownership boundary*

1089 **Action**

1090 The application of intent by an actor to cause an effect.

1091 The aspect of action that distinguishes it from mere force or accident is that someone *intends* that the
1092 action achieves a desired objective or effect. This definition of action is very general. In the case of SOA,
1093 we are mostly concerned with actions that take place within a system and have specific effects on the
1094 SOA ecosystem – defined in section 3.2.3 as what we call real world effects. The actual real world effect
1095 of an action, however, may go beyond the intended effect.

1096 In order for multiple actors to participate in a joint action, they must each act according to their role within
1097 the joint action. This is achieved through communication and messaging.

1098 Communication – the formulation, transmission, receipt and interpretation of messages – is the
1099 foundation of all joint actions within the SOA ecosystem, given the inherent separation – often across
1100 ownership boundaries – of actors in the system.

1101 Communication between actors requires that they play the roles of ‘sender’ or ‘receiver’ of messages as
1102 appropriate to a particular action – although it is not necessarily required that they both be active
1103 simultaneously.

1104 An actor sends a message in order to communicate with other actors. The communication itself is often
1105 not intended as part of the desired real world effect but rather includes messages that seek to establish,
1106 manage, monitor report on, and guide the joint action throughout its execution.

1107 Like communication, joint action usually involves different actors. However, joint action – resulting from
1108 the deliberate actions undertaken by different actors – *intentionally* impacts shared state within the
1109 system leading to real world effects.

1110 **Joint Action**

1111 The coordinated set of **actions** involving the efforts of two or more **actors** to achieve an effect.

1112 Note that the effect of a joint action is *not* always equivalent to one or more effects of the individual
1113 actions of the actors involved, i.e., it may be more than the sum of the parts.

1114 | Different [perspectives](#) lead to either communication or joint action as being considered most important.
1115 For example, from the viewpoint of ecosystem security, the integrity of the communications may be
1116 dominant; from the viewpoint of ecosystem governance, the integrity of the joint action may be dominant.

1117

1118 3.3.3 State and Shared State

1119 State

1120 The condition of an entity at a particular time.

1121 State is characterized by a set of facts that is true of the entity. In principle, the total state of an entity (or
1122 the world as a whole) is unbounded. In practice, we are concerned only with a subset of the state of an
1123 entity that is measurable and useful in a given context.

1124 For example, the total state of a light bulb includes the temperature of the filament of the bulb, the
1125 composition of the glass, the dirt that is on the bulb's surface and so on. However, someone needing
1126 more light to read by is only really interested in whether the bulb is 'on' or 'off' and if it is working properly.
1127 That individual's characterization of the state of the bulb reduces to the fact: 'bulb is now on'.

1128 In a SOA ecosystem, there is a distinction between the set of facts about an entity that only that entity can
1129 access and the set of facts that may be accessible to others, s, notably actors in the SOA-based system.

1130 Private State

1131 That part of an entity's **state** that is knowable by, and accessible to, only that entity.

1132 Shared State

1133 That part of an entity's **state** that is knowable by, and may be accessible to, other actors.

1134 Note that shared state does not imply that the state *is* accessible to other actors. It simply refers to that
1135 subset of state that *may* be accessed by other actors. This will principally be the case when actors need
1136 to participate in joint actions.

1137 It is the aggregation of the shared states of pertinent entities that constitutes the desired effect of a joint
1138 action. Thus the change to this shared state is what is experienced in the wider ecosystem as a real world
1139 effect

1140 3.4 Architectural Implications

1141 3.4.1 Social structures

1142 A SOA ecosystem's participants are organized into various forms of social structure. Not all social
1143 structures are hierarchical: a SOA ecosystem should be able to incorporate peer-to-peer forms of
1144 organization as well as hierarchic structures. In addition, it should be possible to identify and manage any
1145 constitutional agreements that define the social structures present in a SOA ecosystem.

- 1146 • Different social structures have different rules of engagement but predictable behavior is one of
1147 the underpinnings of trust. This therefore requires mechanisms to:
 - 1148 ○ express constitutions and other organizing principles of participants;
 - 1149 ○ inherit rules of engagement from parent to child social structures.
- 1150 • Social structures have roles and members and this impacts who may be authorized to act and in
1151 what circumstances. This requires mechanisms to:
 - 1152 ○ identify and manage members of social structures
 - 1153 ○ Identify and manage attributes of the members
 - 1154 ○ describe roles and role adoption
- 1155 • Social structures overlap and interact, giving rise to situations in which rules of engagement may
1156 conflict. In addition, a given actor may be a member of multiple social structures and the social
1157 structures may be associated with different jurisdictions. This requires mechanisms to:
 - 1158 ○ identify the social structures that are active during a series of joint actions;
 - 1159 ○ identify and resolve conflicts and inconsistencies.

1160 3.4.2 Resource and Ownership

1161 Communication about and between, visibility into, and leveraging of resources requires the unambiguous
1162 identification of those resources. Ensuring unambiguous identities implies

- 1163 • Mechanism for assigning and guaranteeing uniqueness of globally unique identifiers
- 1164 • Identifying the extent of the enterprise over which the identifier needs to be understandable and
- 1165 unique
- 1166 • Mechanism and framework for ensuring the longevity of identifiers (i.e., they cannot just change
- 1167 arbitrarily)

1168 3.4.3 Policies and Contracts

- 1169 • Policies are expressed as constraints
 - 1170 ○ Policies MUST be expressed
 - 1171 ○ Constraints MUST be enforceable
 - 1172 ○ Management of potentially large numbers of policies MUST be achievable
- 1173 • Policies have owners
 - 1174 ○ Policies SHOULD be established by social structures.
- 1175 • Policies may not be consistent with one another
 - 1176 ○ Policy conflict resolution techniques MUST exist and be in place
- 1177 • Agreements are accepted constraints ~~agreed to~~
 - 1178 ○ Contracts SHOULD be enforced by mechanisms of the social structure

1179 3.4.4 Communications as a Means of Mediating Action

1180 Using message exchange for mediating action implies

- 1181 • Ensuring correct identification of the structure of messages:
 - 1182 ○ Identifying the syntax of the message;
 - 1183 ○ Identifying the vocabularies used in the communication
 - 1184 ○ Identifying the higher-level structure of the communication, such as policy assertion,
 - 1185 contract enforcement, etc.
- 1186 • A principal objective of communication is to mediate action
 - 1187 ○ Messages convey actions and events
 - 1188 ○ Receiving a message is an action, but is not the same action as the action conveyed by
 - 1189 the message
 - 1190 ○ Actions are associated with objectives of the actors involved
 - 1191 ▪ Explicit representation of objectives may facilitate automated processing of
 - 1192 messages
 - 1193 ○ An actor agreeing to adopt an objective becomes responsible for that objective

1194 3.4.5 Semantics

1195 Semantics is pervasive in a SOA ecosystem. There are many forms of utterance that are relevant to the

1196 ecosystem: apart from communicated content there are mission and policy statements, goals, objectives,

1197 descriptions, and agreements which are all forms of utterance.

1198 The operation of the SOA ecosystem is significantly enhanced if

- 1199 • A careful distinction is made between public semantics and private semantics. In particular, it
- 1200 MUST be possible for actors to process content such as communications, descriptions and
- 1201 policies solely on the basis of the public semantics of those utterances.
- 1202 • A well founded semantics ensures that any assertions that are essential to the operator of the
- 1203 ecosystem (such as policy statements, and descriptions) have carefully chosen written
- 1204 expressions and associated decision procedures.
- 1205 • The role of vocabularies as a focal point for multiple actors to be able to understand each other is
- 1206 critical. While no two actors can fully share their interpretation of elements of vocabularies,
- 1207 ensuring that they do understand the public meaning of vocabularies' elements is essential.

1208 3.4.6 Trust and Risk

1209 In traditional systems, the balance between trust and risk is achieved by severely restricting interactions

1210 and by controlling the participants of a system.

1211 It is important that actors are able to explicitly reason about both trust and risk in order to effectively
1212 participate in a SOA ecosystem. The more open and public the SOA ecosystem is, the more important it
1213 is for actors to be able to reason about their participation.

1214 **3.4.7 Needs, Requirements and Capabilities**

1215 In the process of capturing needs as requirements, and the subsequent requirements decomposition and
1216 allocation processes need to be informed by capabilities that already exist.

- 1217 • Architecture needs to
- 1218 ○ Take into account existing capabilities available as services

1219 **3.4.8 The Importance of Action**

1220 Participants participate in a SOA ecosystem in order to get their needs met. This involves action; both
1221 individual actions and joint actions.

1222 Any architectural realization of a SOA ecosystem should address:

- 1223 • How actions are modeled:
- 1224 ○ Identifying the performer or agent of the action;
- 1225 ○ the target of the action; and the
- 1226 ○ verb of the action.

1227 Any explicit models of joint action should take into account

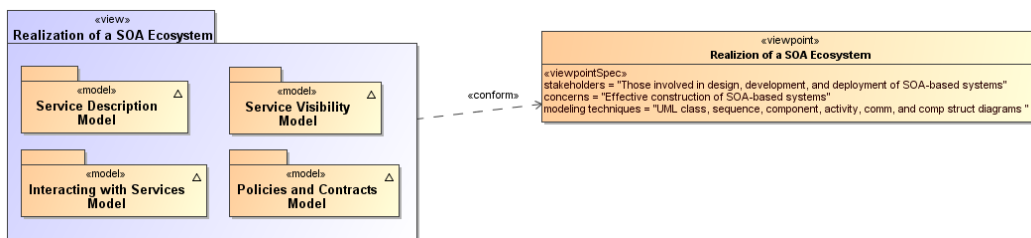
- 1228 • The choreography or orchestration that defines the joint action.
- 1229 • The potential for multiple joint actions to be layered on top of each other

1230 **4 Realization of a SOA Ecosystem view**

1231 *Make everything as simple as possible but no simpler.*
1232 Albert Einstein

1233
1234 The *Realization of a SOA Ecosystem* view focuses on elements that are needed to support the discovery
1235 of and interaction with services. The key questions asked are "What are services, what support is needed
1236 and how are they realized?"

1237 The models in this view include the Service Description Model, the Service Visibility Model, the Interacting
1238 with Services Model, and the Policies and Contracts Model.



1239
1240 *Figure 13 - Model Elements Described in the Realization of a SOA Ecosystem view*

1241 The Service Description Model informs the participants of what services exist and the conditions under
1242 which they can be used. [The Policies and Contracts Model elaborates on the conditions under which](#)
1243 [service use is prescribed and agreements among participants in the SOA ecosystem. Some of these](#)
1244 [conditions follow from policies and agreements on policy that flow from the Policies and Contracts Model.](#)
1245 The information in the service description as augmented by details of policy provides the basis for
1246 visibility as defined in the SOA Reference Model and captured in the Service Visibility Model. Finally, the
1247 process by which services are used under the defined conditions and agreements is described in the
1248 Interacting with Services Model.

Comment [KJL28]: Issue 168

1249 **4.1 Service Description Model**

1250 A service description is an artifact, often document-based, that defines or references the information
1251 needed to use, deploy, manage and otherwise control a service. This includes not only the information
1252 and behavior models associated with a service that define [interaction via](#) the service interface but also
1253 includes information needed to decide whether the service is appropriate for the current needs of the
1254 service consumer. Thus, the service description should also include information such as service
1255 reachability, service functionality, and the policies associated with a service.

Comment [KJL29]: Issue 170
(see also Issue 176)

1256 A service description artifact may be a single document or it may be an interlinked set of documents. For
1257 the purposes of this model, differences in representation are to be ignored, but the implications of a "web
1258 of documents" are discussed later in this section.

1259 There are several points to note regarding service description:

- 1260 • The Reference Model states that one of the hallmarks of SOA is the large amount of associated
1261 description. The model presented below focuses on the description of services but it is equally
1262 important to consider the descriptions of the consumer, other participants, and needed resources
1263 other than services.
- 1264 • Descriptions are inherently incomplete but may be determined as *sufficient* when it is possible for
1265 the participants to access and use the described services based only on the descriptions
1266 provided. This means that, at one end of the spectrum, a description along the lines of "That
1267 service on that machine" may be sufficient for the intended audience. On the other extreme, a
1268 service description with a machine-process-able description of the semantics of its operations

- 1269 and real world effects may be required for services accessed via automated service discovery
 1270 and planning systems.
- 1271 • Descriptions come with context, i.e. a given description comprises information needed to
 1272 adequately support the context. For example, a list of items can define a version of a service, but
 1273 for many contexts an indicated version number is sufficient without the detailed list. The current
 1274 model focuses on the description needed by a service consumer to understand what the service
 1275 does, under what conditions the service will do it, how well the service does it, and what steps are
 1276 needed by the consumer to initiate and complete a service interaction. Such information also
 1277 enables the service provider to clearly specify what is being provided and the intended conditions
 1278 of use.
 - 1279 • Descriptions change over time as, for example, the ingredients and nutrition information for food
 1280 labeling continues to evolve. A requirement for transparency of transactions may require
 1281 additional description for those associated contexts.
 - 1282 • Description always proceeds from a basis of what is considered "common knowledge". This may
 1283 be social conventions that are commonly expected or possibly codified in law. It is impossible to
 1284 describe everything and it can be expected that a mechanism as far reaching as SOA will also
 1285 connect entities where there is inconsistent "common" knowledge.
 - 1286 • Descriptions become the collection point of information related to a service or any other resource,
 1287 but it is not necessarily the originating point or the motivation for generating this information. In
 1288 particular, given a SOA service as the access to an underlying capability, the service may point to
 1289 some of the capability's previously generated description, e.g. a service providing access to a
 1290 data store may also have access to information indicating the freshness of the data.

1291 These points emphasize that there is no one "right" description for all contexts and for all time. Several
 1292 descriptions for the same subject may exist at the same time, and this emphasizes the importance of the
 1293 description referencing source material maintained by that material's owner rather than having multiple
 1294 copies that become out of synch and inconsistent.

1295 It may also prove useful for a description assembled for one context to cross-reference description
 1296 assembled for another context as a way of referencing ancillary information without overburdening any
 1297 single description. Rather than a single artifact, description can be thought of as a web of documents that
 1298 enhance the total available description.

1299 This Reference Architecture Foundation uses the term service description for consistency with the
 1300 concept defined in the Reference Model. Some SOA literature treats the idea of a "service contract" as
 1301 equivalent to service description. In the SOA-RAF, the term service description is preferred. Replacing the
 1302 term "service description" with the term "service contract" implies that just one side of the interaction is
 1303 governing and misses the point that a single set of policies identified by a service description may lead to
 1304 numerous contracts, i.e. service level agreements, leveraging the same description.

1305 4.1.1 The Model for Service Description

1306 Figure 14 shows Service Description as a subclass of the general Description class, where Description is
 1307 a subclass of the resource class as defined in Section 3.2.4.1.⁵ In addition, each resource is assumed to
 1308 have a description. The following section discusses the relationships among elements of general
 1309 description and the subsequent sections focus on service description. Other descriptions, such as those
 1310 of participants, are important to SOA but are not individually elaborated in this document.

Comment [PFB30]: This contradicts 3.1.3.1, which states only that "Description describes Resource"

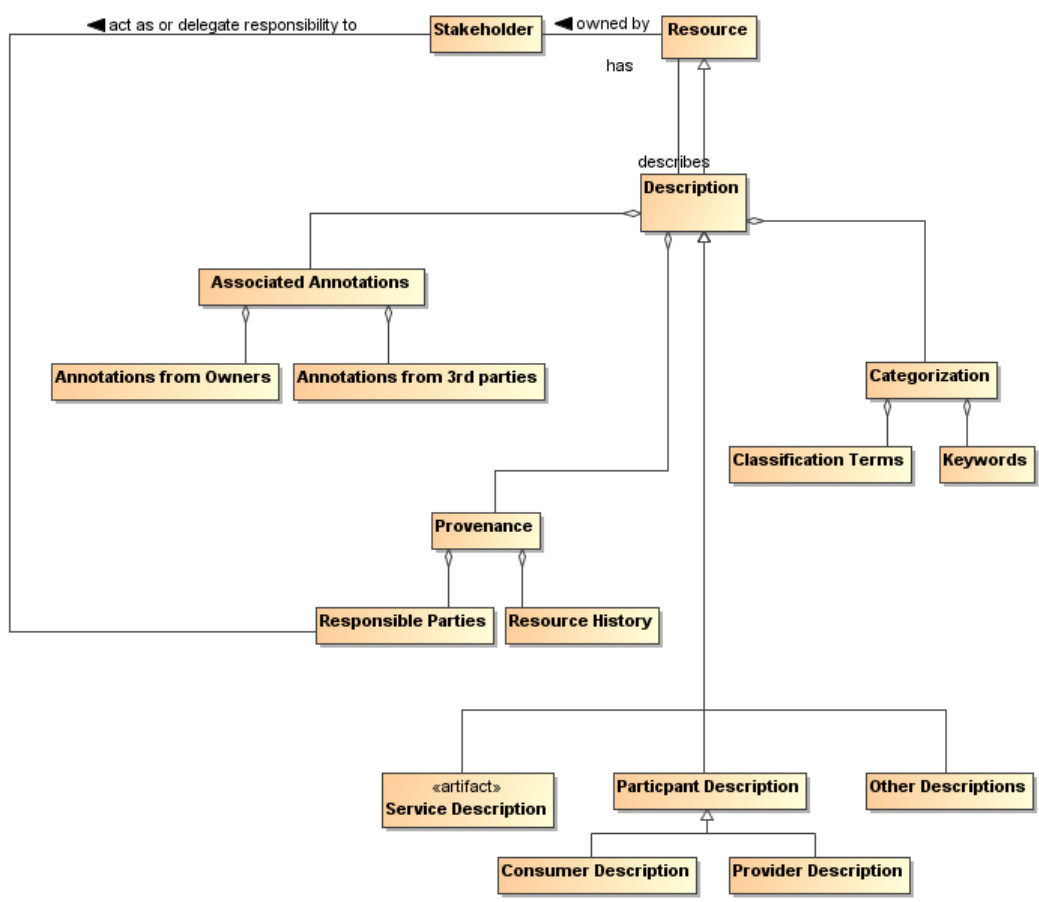
1311 4.1.1.1 Elements Common to General Description

1312 The general Description class is composed of a number of elements that are expected to be common
 1313 among all descriptions supporting a service-oriented architecture. A registry/[repository](#) often contains a
 1314 subset of the description instance, where the chosen subset is identified as that which facilitates

Comment [KJL31]: Issue 173

⁵ [Resources have descriptions and a description can be considered as a type of resource \(and the description itself has further data such as its version or last revision\) The model emphasizes this point but should not be interpreted too rigorously as allowing endless recursion.](#)

1315 discovery. Additional information contained in a more complete description may be needed to initiate and
 1316 continue interaction.
 1317



Comment [PFB32]: Recursion loop – every description is a resource, that requires a description. Modifications needed for Issue 290 remove Consumer and Provider Description classes.

1318
 1319 *Figure 14 - General Description*

1320 **4.1.1.1.1 Provenance**

1321 While the resource Identifier provides the means to know which subject and subject description are being
 1322 considered, Provenance as related to the Description class provides information that reflects on the
 1323 quality or usability of the subject. Provenance specifically identifies the stakeholder (human, defined role,
 1324 organization, etc.) that assumes responsibility for the resource being described and tracks historic
 1325 information that establishes a context for understanding what the resource provides and how it has
 1326 changed over time. Responsibilities may be directly assumed by the stakeholder who owns a resource
 1327 (see Section 3.2.4.2) or the Owner may designate Responsible Parties for the various aspects of
 1328 maintaining the resource and provisioning it for use by others. There may be more than one stakeholder
 1329 identified under Responsible Parties; for example, one stakeholder may be responsible for code
 1330 maintenance while another is responsible for provisioning of the executable code.

1331 **4.1.1.1.2 Keywords and Classification Terms**

1332 A traditional element of description has been to associate the resource being described with predefined
 1333 keywords or classification taxonomies that derive from referenceable formal definitions and vocabularies.

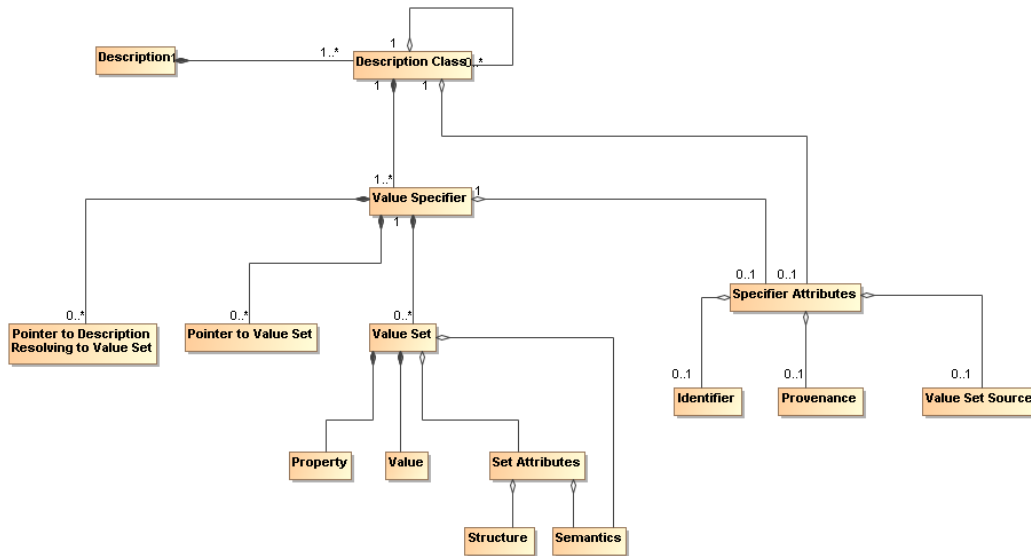
1334 This Reference Architecture Foundation does not prescribe which vocabularies or taxonomies may be
 1335 referenced, nor does it limit the number of keywords or classifications that may be associated with the
 1336 resource. It does, however, state that a normative definition of any terms or keywords SHOULD be
 1337 referenced, whether that be a representation in a formal ontology language, a pointer to an online
 1338 dictionary, or any other accessible source. See Section 4.1.1.2 for further discussion on associating
 1339 semantics with assigned values.

1340 **4.1.1.1.3 Associated Annotations**

1341 The general description instance may also reference associated documentation that is in addition to that
 1342 considered necessary in this model. For example, the owner of a service may have documentation on
 1343 best practices for using the service. Alternately, a third party may certify a service based on their own
 1344 criteria and certification process; this may be vital information to other prospective consumers if they were
 1345 willing to accept the certification in lieu of having to perform another certification themselves. Note, while
 1346 the examples of Associated Documentation presented here are related to services, the concept applies
 1347 equally to description of other entities.

1348 **4.1.1.2 Assigning Values to Description Instances**

1349



1350

1351 *Figure 15 - Representation of a Description*

1352 Figure 14 shows the template for a general description, but individual description instances depend on
 1353 the ability to associate meaningful values with the identified elements. Figure 15 shows a model for a
 1354 collection of information that provides for value assignment and traceability for both the meaning and the
 1355 source of a value. The model is not meant to replace existing or future schema or other structures that
 1356 have or will be defined for specific implementations, but it is meant as guidance for the information such
 1357 structures need to capture to generate sufficient description. It is expected that tools will be developed to
 1358 assist the user in populating description and auto-filling many of these fields, and in that context, this
 1359 model provides guidance to the tool developers.

1360 In Figure 15, each class has an associated value specifier or is made up of components that eventually
 1361 resolve to a value specifier. For example, Description has several components, one of which is
 1362 Categorization, which would have an associated value specifier.

1363 A value specifier consists of

- 1364
- 1365
- 1366
- a collection of value sets with associated property-value pairs, pointers to such value sets, or pointers to descriptions that eventually resolve to value sets that describe the component; and
 - attributes that qualify the value specifier and the value sets it contains.

1367 The qualifying attributes for the value specifier include

- 1368
- 1369
- 1370
- 1371
- 1372
- 1373
- an optional identifier that would allow the value set to be defined, accessed, and reused elsewhere;
 - provenance information that identifies the **party-person** (individual, ~~role~~, or organization) who has responsibility for assigning the value sets to any description component;
 - an optional source of the value set, if appropriate and meaningful, e.g. if a particular data source is mandated.

Comment [PFB33]: Issue 291

1374 If the value specifier is contained within a higher-level component (such as Service Description containing
1375 Service Functionality), the component may assume values from the attributes of its container.

1376 Note, provenance as a qualifying attribute of a value specifier is different from provenance as part of an
1377 instance of Description. Provenance for a service identifies those who own and are responsible for the
1378 service, as described in Section 3.2.4. Provenance for a value specifier identifies who is responsible for
1379 choosing and assigning values to the value sets that comprise the value specifier. It is assumed that
1380 granularity at the value specifier level is sufficient and provenance is not required for each value set.

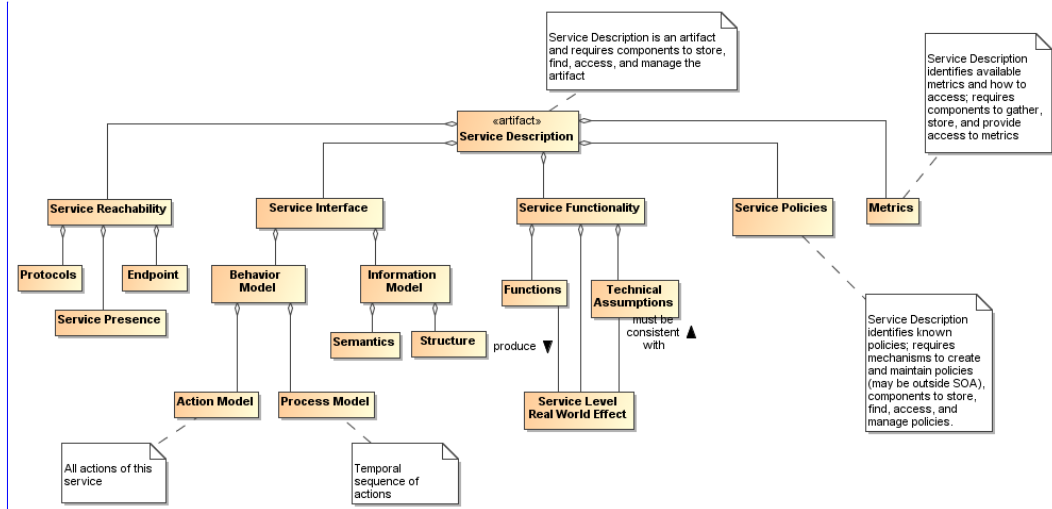
1381 The value set also has attributes that define its structure and semantics.

- 1382
- 1383
- 1384
- 1385
- 1386
- 1387
- 1388
- 1389
- 1390
- 1391
- The semantics of the value set property should be associated with a semantic context conveying the meaning of the property within the execution context, where the semantic context could vary from a free text definition to a formal ontology.
 - For numeric values, the structure would provide the numeric format of the value and the “semantics” would be conveyed by a dimensional unit with an identifier to an authoritative source defining the dimensional unit and preferred mechanisms for its conversion to other dimensional units of like type.
 - For nonnumeric values, the structure would provide the data structure for the value representation and the semantics would be an associated semantic model.
 - For pointers, architectural guidelines would define the preferred addressing scheme.

1392 The value specifier may indicate a default semantic model for its component value sets and the individual
1393 value sets may provide an override.

1394 The property-value pair construct is introduced for the value set to emphasize the need to identify
1395 unambiguously both what is being specified and what is a consistent associated value. The further
1396 qualifying of Structure and Semantics in the Set Attributes allows for flexibility in defining the form of the
1397 associated values.

1398 **4.1.1.3 Model Elements Specific to Service Description**



Comment [KJL34]: To be changed as per 3/29/2012 email and attachment (will resolve Issues 66, 176) Add event model or remove action and process models? (Issue 292)

1399
1400 *Figure 16 - Service Description*

1401 The major elements for the Service Description subclass follow directly from the areas discussed in the
1402 Reference Model. Here, we discuss the detail shown in Figure 16 and the purpose served by each
1403 element of service description. For example, Service Policies as included in Figure 16 indicate those
1404 policies that affect conditions of use of the service; however, while the description may link to detailed
1405 policy documents, it is not the purpose of description to justify or elaborate on the rationale for the
1406 policies. Similarly, Service Interface Description as included in Figure 16 captures information about what
1407 interactions are supported by the service via its Behavior Model and the information exchange needed to
1408 carry out those interactions in accordance to the service's Information Model; it is not the coded interface.

1409 Note, the intent in the subsections that follow is to describe how a particular element, such as the service
1410 interface description, is reflected in the service description, not to elaborate on the details of that element.

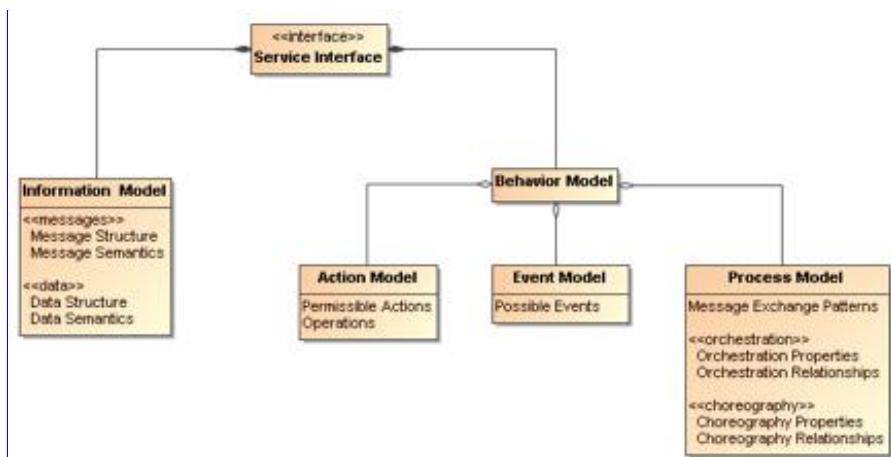
1411 **4.1.1.3.1 Service Interface Description**

1412 As noted in the Reference Model, the service interface is the means for interacting with a service. For the
1413 SOA-RAF and as shown in Section 4.3 the service interface supports an exchange of messages, where

- 1414 • the message conforms to a referenceable message exchange pattern (MEP),
- 1415 • the message payload conforms to the structure and semantics of the indicated information model,
- 1416 • the messages are used to denote events or actions against the service, where the actions are
- 1417 specified in the action model and any required sequencing of actions is specified in the process
- 1418 model.

1419 The Service Interface Description element as shown in Figure 17 includes the information needed to carry
1420 out this message exchange in order to realize the service behavior described. In addition to the
1421 Information Model that conveys the Semantics and Structure of the message, the Service Interface
1422 Description indicates what behavior can be expected through interactions conveyed in the Action and
1423 Process Models.

Comment [KJL35]: Issue 176, part



Comment [KJL36]: Change per 3/29/2012 email and attachment

Comment [PFB37]: TO BE DONE: Needs slight modifications in light of changes to previous fig.

1424
1425 | Figure 17 - Service Interface [Description](#)

1426 Note we distinguish the structure and semantics of the message from that of the underlying protocol that
1427 conveys the message. The message structure may include nested structures that are independently
1428 defined, such as an enclosing envelope structure and an enclosed data structure.
1429 These aspects of messages are discussed in more detail in Section 4.3.2.

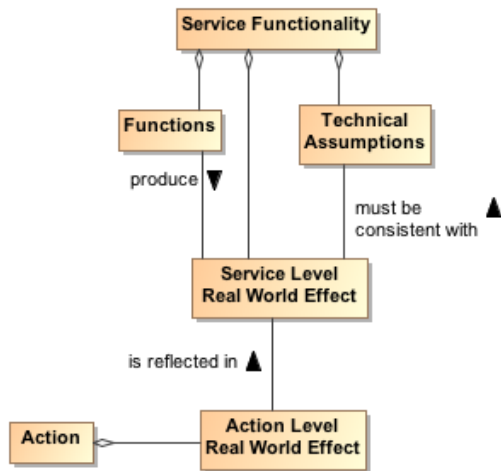
1430 4.1.1.3.2 Service Reachability

1431 Service reachability, as modeled in Section 4.2.2.3 enables service participants to locate and interact with
1432 one another. To support service reachability, the service description should indicate the endpoints to
1433 which a service consumer can direct messages to invoke actions and the protocol to be used for
1434 message exchange using that endpoint.

1435 As generally applied to an action, the endpoint is the conceptual location where one applies an action;
1436 with respect to service description, it is the actual address where a message is sent.

1437 4.1.1.3.3 Service Functionality

1438 While the service interface and service reachability are concerned with the mechanics of using a service,
1439 service functionality and performance metrics (discussed in Section 4.1.1.3.4) describe what can be
1440 expected as a result of interacting with a service. Service Functionality, shown in Figure 16 as part of the
1441 overall Service Description model and extended in Figure 18, is a clear expression of service function(s)
1442 and the real world effects of invoking the function. The Functions represent business activities in some
1443 domain that produce the desired real world effects.



1444

1445 *Figure 18 - Service Functionality*

1446 The Service Functionality may also be limited by technical assumptions/constraints that underlie the
 1447 effects that can result. Technical constraints are defined as domain specific restrictions and may express
 1448 underlying physical limitations, such as flow speeds must be below sonic velocity or disk access that
 1449 cannot be faster than the maximum for its host drive. Technical constraints are related to the underlying
 1450 capability accessed by the service. In any case, the real world effects must be consistent with the
 1451 technical assumptions/constraints.

1452 In Figure 16 and Figure 18, we specifically refer to [the descriptions of](#) Service Level and Action Level real
 1453 world effects.

Comment [KJL38]: Issue 176

1454 **Service Level Real World Effect**

1455 A service level real world effect is a specific change in the state or the information returned as a
 1456 result of interacting with a service.

1457 **Action Level Real World Effect**

1458 An action level real world effect is a specific change in the state or the information returned as a
 1459 result of interacting through a specific action.

1460 Service description describes the service as a whole while the component aspects should contribute to
 1461 that whole. Thus, while individual Actions may contribute to the real world effects to be realized from
 1462 interaction with the service, there would be a serious disconnect for Actions to contribute real world
 1463 effects that could not consistently be reflected in the Service Level Real World Effects and thus the
 1464 Service Functionality. The relationship to Action Level Real World Effects and the implications on defining
 1465 the scope of a service are discussed in Section 4.1.2.1.

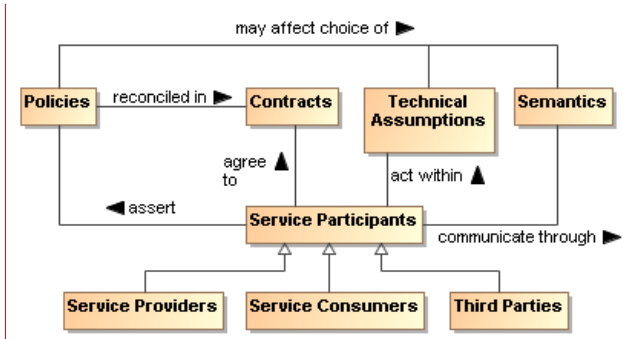
1466 Elements of Service Functionality may be expressed as natural language text, reference an existing
 1467 taxonomy of functions or other formal model.

1468 **4.1.1.3.4 Service Policies, Metrics, and Compliance Records**

1469 Policies prescribe the conditions and constraints for interacting with a service and impact the willingness
 1470 to continue visibility with the other participants. Whereas technical constraints are statements of "physical"
 1471 fact, policies are subjective assertions made by the service provider (sometimes as passed on from
 1472 higher authorities).

1473 The service description provides a central location for identifying what policies have been asserted by the
 1474 service provider. The specific representation of the policy, e.g. in some formal policy language, is outside
 1475 of the service description. The service description would reference the normative definition of the policy.

1476 Policies may also be asserted by other service participants, as illustrated by the model shown in Figure
 1477 19. Policies that are generally applicable to any interaction with the service are asserted by the service
 1478 provider and included in the Service Policies section of the service description.



Comment [PFB39]: Issue 179 - modifications needed

1479
 1480 Figure 19 - Model for Policies and Contracts as related to Service Participants

1481 In Figure 19, we specifically refer to policies at the service level. In a similar manner to that discussed for
 1482 Service Level vs. Action Level Real World Effects in Section 4.1.1.3.3, individual Actions may have
 1483 associated policies stating conditions for performing the action, but these must be reflected in and be
 1484 consistent with the policies made visible at the service level and thus the description of the service as a
 1485 whole. The relationship to Action Level Policies and the implications on defining the scope of a service
 1486 are discussed in Section 4.1.2.1.

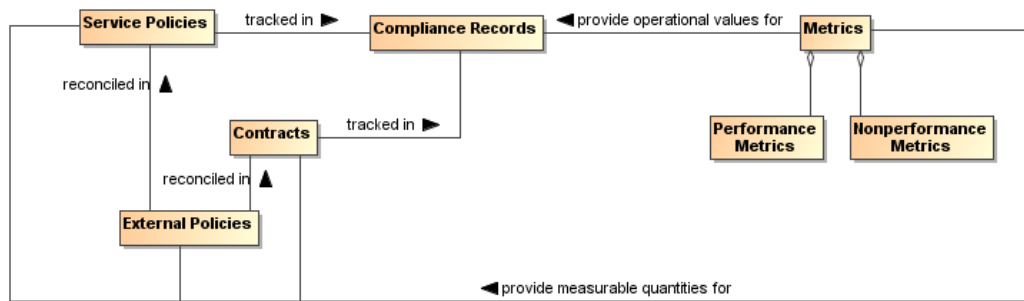
1487 As noted in Figure 19, the policies asserted may be reflected as Technical Assumptions/Constraints that
 1488 available services or their underlying capabilities must be capable of meeting; it may similarly affect the
 1489 semantics that can be used. For example of the former, there may be a policy that specifies the surge
 1490 capacity to be accommodated by a server, but a service that is not designed to make use of the larger
 1491 server capacity would not satisfy the intent of the policy and would not be appropriate to use. For the
 1492 latter, a policy may require that only services that support interaction via a community-sponsored
 1493 vocabulary can be used.

1494 Contracts are agreements among the service participants. The contract may reconcile inconsistent
 1495 policies asserted by the participants or may specify details of the interaction. Service level agreements
 1496 (SLAs) are one of the commonly used categories of contracts.

1497 The definition and later enforcement of policies and contracts are predicated on the potential for
 1498 measurement; the relationships among the relevant concepts are shown in the model in Figure 20.
 1499 Performance Metrics identify quantities that characterize the speed and quality of realizing the real world
 1500 effects produced using the SOA service; in addition, policies and contracts may depend on
 1501 nonperformance metrics, such as whether a license is in place to use the service. Some of these metrics
 1502 may reflect the underlying capability, some metrics may reflect processing of the SOA service, and some
 1503 metrics may include expected network overhead. The metrics should be carefully defined to avoid
 1504 confusion in exactly what is being reported, for example, a case where the service processing time is
 1505 reported as if it were the total time including the capability and network processing but is only measuring
 1506 the service processing. Some of these metrics reflect the underlying capability, e.g. a SOA service cannot
 1507 respond in two seconds if the underlying capability is expected to take five seconds to do its processing;
 1508 some metrics reflect the SOA service, e.g. the additional overhead introduced when making data access
 1509 requests across the network.

Comment [KJL40]: Issue 254

1510



1511
1512 *Figure 20 - Policies and Contracts, Metrics, and Compliance Records*

1513 As with many quantities, the metrics associated with a service are not themselves defined by this Service
1514 Description Model because it is not known *a priori* which metrics are being collected or otherwise checked
1515 by the services, the SOA infrastructure, or other resources that participate in the SOA interactions.
1516 However, the service description SHOULD provide a placeholder (possibly through a link to an externally
1517 compiled list) for identifying which metrics are available and how these can be accessed.

1518 The use of metrics to evaluate compliance and the results of compliance evaluation SHOULD be
1519 maintained in compliance records and the means to access the compliance records MAY be included in
1520 the Service Policies portion of the service description. For example, the description may be in the form of
1521 static information (e.g. over the first year of operation, this service had a 91% availability), a link to a
1522 dynamically generated metric (e.g. over the past 30 days, the service has had a 93.3% availability), or
1523 access to a dynamic means to check the service for current availability (e.g., a ping). The relationship
1524 between service presence and the presence of the individual actions that can be invoked is discussed
1525 under Reachability in Section 4.2.2.3.

1526 Note, even when policies relate to the perspective of a single participant, policy compliance can be
1527 measured and policies may be enforceable without contractual agreement with other participants. While
1528 certain elements of contracts and contract compliance are likely private, public aspects of compliance
1529 should be reflected in the compliance record information referenced in the service description. [This](#)
1530 [provides input to evidence that supports determining willingness as described in Section 3.2.5.1.](#)

Comment [PFB41]: Issue 70

1531 4.1.2 Use of Service Description

1532 4.1.2.1 Service Description in support of Service Interaction

1533 If we assume we have awareness, the service participants must still establish willingness and presence to
1534 ensure full visibility (See Section 4.2) and to interact with the service. Service description provides
1535 necessary information for many aspects of preparing for and carrying through with interaction. Recall the
1536 fundamental definition of a SOA service is a mechanism to access an underlying capability; the service
1537 description describes this mechanism and its use. It lays the groundwork for what can occur, whereas
1538 service interaction comprises the specifics through which real-world effects are realized.

1569 The conditions under which an action can be invoked may depend on policies associated with the action.
1570 The Action Level Policies MUST be reflected in (or subsumed by) the Service Policies because such
1571 policies may be critical to determining whether the conditions for use of the service are consistent with the
1572 policies asserted by the service consumer. [For example, if an action requires interaction with another
1573 service and that other service has licensing requirements, then the service with such an action also has
1574 the same requirement.](#) The Service Policies are included in the service description.

Comment [KJL46]: Issue 77

1575 Similarly, the result of invoking an action is one or more real world effects, and any Action Level Real
1576 World Effects MUST be reflected in the Service Level Real World Effect included in the service
1577 description. The unambiguous expression of action level policies and real world effects as service
1578 counterparts is necessary to adequately describe what constitutes the service interaction. [For example, if
1579 an action allows for the tracking of user preferences, then the service with such an action results in the
1580 same real world effect.](#)

Comment [KJL47]: Similar to Issue 77 but never explicitly entered

1581 An adequate service description MUST provide a consumer with information needed to determine if the
1582 service policies, the (business) functions, and service-level real world effects are of interest, and there is
1583 nothing in the technical constraints that preclude use of the service.

1584 Note at the service level, the business functions are not concerned with the action or process models.
1585 These models are detailed separately.

1586 The service description is not intended to be isolated documentation but rather an integral part of service
1587 use. Changes in service description SHOULD immediately be made known to consumers and potential
1588 consumers.

1589 4.1.2.2 Description and Invoking Actions Against a Service

Comment [PFB48]: Invocation?

1590 At this point, let us assume the descriptions were sufficient to establish willingness; see Section 4.2.2.2.
1591 Figure 21 indicates the service endpoint establishes where to actually carry out the interaction. This is
1592 where we start considering the action and process models.

1593 The action model identifies the multiple actions a user can perform against a service and the user would
1594 perform these in the context of the process model as specified or referenced under the Service Interface
1595 [Description](#) portion of Service Description. For a given business function, there is a corresponding
1596 process model, where any process model may involve multiple actions. From the above discussion of
1597 model elements of description we may conclude (1) actions have reachability information, including
1598 endpoint and presence, (2) presence of service is some aggregation of presence of its actions, (3) action
1599 preconditions and service dependencies do not affect presence although these may affect successful
1600 completion.

Comment [PFB49]: Issue 183

1601 Having established visibility, the interaction can proceed. Given a business function, the consumer knows
1602 what will be accomplished (the service functionality), the conditions under which interaction will proceed
1603 (service policies), and the process that must be followed (the process model). The remaining question is
1604 how the description information for structure and semantics enable interaction.

1605 We have established the importance of the process model in identifying relevant actions and their
1606 sequence. Interaction proceeds through messages and thus it is the syntax and semantics of the
1607 messages with which we are here concerned. A common approach is to define the structure and
1608 semantics that can appear as part of a message; then assemble the pieces into messages; and,
1609 associate messages with actions. Actions make use of structure and semantics as defined in the
1610 information model to describe its legal messages.

1611 The process model identifies actions to be performed against a service and the sequence for performing
1612 the actions. For a given action, the Reachability portion of description indicates the protocol bindings that
1613 are available, the endpoint corresponding to a binding, and whether there is presence at that endpoint. An
1614 interaction is through the exchange of messages that conform to the structure and semantics defined in
1615 the information model and the message sequence conforming to the action's identified MEP. The result is
1616 some portion of the real world effect that must be assessed and/or processed (e.g. if an error exists, that
1617 part that covers the error processing would be invoked).

1618 **4.1.2.3 The Question of Multiple Business Functions**

1619 Action level effects and policies MUST be reflected at the service level for service description to support
1620 visibility.

1621 It is assumed that a SOA service represents an identifiable business function to which policies can be
1622 applied and from which desired business effects can be obtained. While contemporary discussions of
1623 SOA services and supporting standards do not constrain what actions or combinations of actions can or
1624 should be defined for a service, the SOA-RAF considers the implications of service description in defining
1625 the range of actions appropriate for an individual SOA service.

1626 Consider the situation if a given SOA service is the mechanism for access to multiple independent (but
1627 loosely related) business functions. These are not multiple effects from a single function but multiple
1628 functions with potentially different sets of effects for each function. A service can have multiple actions a
1629 user may perform against it, and this does not change with multiple business functions. As an individual
1630 business function corresponds to a process model, so multiple business functions imply multiple process
1631 models. The same action may be used in multiple process models but the aggregated service presence
1632 would be specific to each business function because the components being aggregated may be different
1633 between process models. In summary, for a service with multiple business functions, each function has
1634 (1) its own process model and dependencies, (2) its own aggregated presence, and (3) possibly its own
1635 list of policies and real world effects.

1636 A common variation on this theme is for a single service to have multiple endpoints for different levels of
1637 quality of service (QoS), e.g. Gold, Silver, and Bronze. Different QoS imply separate statements of policy,
1638 separate endpoints, possibly separate dependencies, and so on. One could say the QoS variation does
1639 not require this because there can be a single QoS policy that encompasses the variations, and all other
1640 aspects of the service would be the same except for the endpoint used for each QoS. However, the
1641 different aspects of policy at the service level would need to be mapped to endpoints, and this introduces
1642 an undesirable level of coupling across the elements of description. In addition, it is obvious that
1643 description at the service level can become very complicated if the number of combinations is allowed to
1644 grow.

Comment [KJL50]: Issue 257

1645 One could imagine a service description that is basically a container for action descriptions, where each
1646 action description is self-contained; however, this would lead to duplication of description components
1647 across actions. If common description components are factored, this either is limited to components
1648 common across all actions or requires complicated tagging to capture the components that often but do
1649 not universally apply.

1650 If a provider cannot describe a service as a whole but must describe every action, this leads to the
1651 situation where it may be extremely difficult to construct a clear and concise service description that can
1652 effectively support discovery and use without tedious logic to process the description and assemble the
1653 available permutations. In effect, if adequate description of an action begins to look like description of a
1654 service, it may be best to have it as a separate service.

1655 Recall, more than one service can access the same underlying capability, and this is appropriate if a
1656 different real world effect is to be exposed. Along these lines, one can argue that different QoS are
1657 different services because getting a response in one minute rather than one hour is more than a QoS
1658 difference; it is a fundamental difference in the business function being provided.

1659 As a best practice, the criteria for whether a service is appropriately scoped may be the ease or difficulty
1660 in creating an unambiguous service description. A consequence of having tightly-scoped services is there
1661 will likely be a greater reliance on combining services, i.e. more fundamental business functions, to create
1662 more advanced business functions. This is consistent with the principles of service oriented architecture
1663 and is the basic position of this Reference Architecture Foundation, although not an absolute
1664 requirement. Combining services increases the reliance on understanding and implementing the concepts
1665 of orchestration, choreography, and other approaches yet to be developed; these are discussed in more
1666 detail in section 4.4 Interacting with Services.

1667 **4.1.2.4 Service Description, Execution Context, and Service Interaction**


1668 The service description MUST provide sufficient information to support service visibility, including the
1669 willingness of service participants to interact. However, the corresponding descriptions for providers and

1670 consumers may both contain policies, technical assumptions, constraints on semantics, and other
1671 technical and procedural conditions that must be aligned to define the terms of willingness. The
1672 agreements that encapsulate the necessary alignment form the basis upon which interactions may
1673 proceed – in the Reference Model, this collection of agreements and the necessary environmental
1674 support establish the execution context.

1675 To illustrate ~~the concept of the~~ execution context of a service interaction, consider a Web-based system
1676 for timecard entry. For an employee onsite at an employer facility, the execution context requires a
1677 computer connected to the local network and the employee must enter their network ID and password.
1678 Relevant policies include that the employee must maintain the most recent anti-virus software and virus
1679 definitions for any computer connected to the network.

Comment [PFB51]: Issue 185

1680 For the same employee connecting from offsite, the execution context specifies the need for a computer
1681 with installed VPN software and a security token to negotiate the VPN connection. The execution context
1682 also includes proxy settings as needed to connect to the offsite network. The employee must still comply
1683 with the requirements for onsite computers and access, but the offsite execution context includes
1684 additional items before the employee can access the same underlying capability and realize the same
1685 real world effects, i.e. the timecard entries.

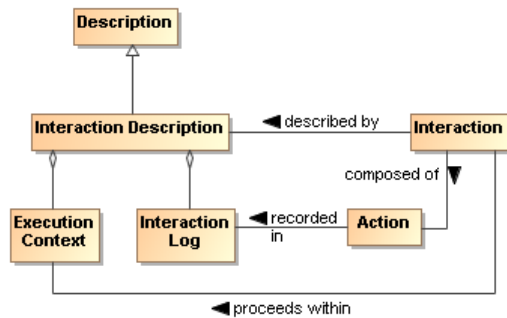


1686
1687 *Figure 22 - Execution Context*

1688 Figure 22 shows a few broad categories found in execution context. These are not meant to be
1689 comprehensive. Other items may need to be included to provide a sufficient description of the interaction
1690 conditions. Any other items not explicitly noted in the model but needed to set the environment SHOULD
1691 be included in the execution context.

1692 While the execution context captures the conditions under which interaction can occur, it does not capture
1693 the specific service invocations that do occur in a specific interaction. A service interaction as modeled in
1694 Figure 23 introduces the concept of an Interaction Description that is composed of both the Execution
1695 Context and an Interaction Log. The execution context specifies the set of conditions under which the
1696 interaction occurs and the interaction log captures the sequence of service interactions that occur within
1697 the execution context. This sequence should follow the Process Model but can include details beyond
1698 those specified there. For example, the Process Model may specify an action that results in identifying a
1699 data source, and the identified source is used in a subsequent action. The Interaction Log would record
1700 the specific data source used.

1701 The execution context can be thought of as a container in which the interaction occurs and the interaction
1702 log captures what happens inside the container. This combination is needed to support auditability and
1703 repeatability of the interactions.



1704

1705 *Figure 23 - Interaction Description*

1706 SOA allows flexibility to accomplish both repeatability and reusability. In facilitating reusability, a service
 1707 can be updated without disrupting the user experience of the service. So, Google can improve their
 1708 ranking algorithm without notifying the user about the details of the update.

1709 However, it may also be vital for the consumer to be able to recreate past results or to generate
 1710 consistent results in the future, and information such as what conditions, which services, and which
 1711 versions of those services were used is indispensable in retracing one's path. The interaction log is a
 1712 critical part of the resulting real world effects because it defines how the effects were generated and
 1713 possibly the meaning of observed effects. This increases in importance as dynamic composability
 1714 becomes more feasible. In essence, a result has limited value if one does not know how it was generated.

1715 The interaction log SHOULD be a detailed trace for a specific interaction, and its reuse is limited to
 1716 duplicating that interaction. An execution context can act as a template for identical or similar interactions.
 1717 Any given execution context MAY define the conditions of future interactions.

1718 Such uses of execution context imply (1) a standardized format for capturing execution context and (2) a
 1719 subclass of general description could be defined to support visibility of saved execution contexts. The
 1720 specifics of the relevant formats and descriptions are beyond the scope of this document.

1721 A service description is unlikely to track interaction descriptions or the constituent execution contexts or
 1722 interaction logs that include mention of the service. However, as appropriate, linking to specific instances
 1723 of either of these could be done through associated annotations.

1724 4.1.3 Relationship to Other Description Models

1725 While the representation shown in Figure 15 is derived from considerations related to service description,
 1726 it is acknowledged that other metadata standards are relevant and should, as possible, be incorporated
 1727 into this work. Two standards of particular relevance are the Dublin Core Metadata Initiative (DCMI)
 1728 [DCMI] and ISO 11179 [ISO 11179], especially Part 5.

1729 When the service description (or even the general description class) is considered as the DCMI
 1730 "resource", Figure 15 aligns nicely with the DCMI resource model. While some differences exist, these are
 1731 mostly in areas where DCMI goes into detail that is considered beyond the scope of the current
 1732 Reference Architecture Foundation. For example, DCMI defines classes of "shared semantics" whereas
 1733 this Reference Architecture Foundation considers that an identification of relevant semantic models is
 1734 sufficient. Likewise, the DCMI "description model" goes into the details of possible syntax encodings
 1735 whereas for the Reference Architecture Framework it is sufficient to identify the relevant formats.

1736 With respect to ISO 11179 Part 5, the metadata fields defined in that reference may be used without
 1737 prejudice as the properties in Figure 15. Additionally, other defined metadata sets may be used by the
 1738 service provider if the other sets are considered more appropriate, i.e. it is fundamental to this reference
 1739 architecture to identify the need and the means to make vocabulary declarations explicit but it is beyond
 1740 the scope to specify which vocabularies are to be used. In addition, the identification of domain of the
 1741 properties and range of the values has not been included in the current Reference Architecture
 1742 discussion, but the text of ISO 11179 Part 5 can be used consistently with the model prescribed in this
 1743 document.

1744 Description as defined here considers a wide range of applicability and support of the principles of service
1745 oriented architecture. Other metadata models can be used in concert with the model presented here
1746 because most of these focus on a finer level of detail that is outside the present scope, and so provide a
1747 level of implementation guidance that can be applied as appropriate.

1748 **4.1.4 Architectural Implications**

1749 The definition of service description indicates numerous architectural implications on the SOA ecosystem:

- 1750 • It changes over time and its contents will reflect changing needs and context. This requires the
1751 existence of:
 - 1752 ○ mechanisms to support the storage, referencing, and access to normative definitions of
1753 one or more versioning schemes that may be applied to identify different aggregations of
1754 descriptive information, where the different schemes may be versions of a versioning
1755 scheme itself;
 - 1756 ○ configuration management mechanisms to capture the contents of each aggregation and
1757 apply a unique identifier in a manner consistent with an identified versioning scheme;
 - 1758 ○ one or more mechanisms to support the storage, referencing, and access to conversion
1759 relationships between versioning schemes, and the mechanisms to carry out such
1760 conversions.
- 1761 • Description makes use of defined semantics, where the semantics may be used for
1762 categorization or providing other property and value information for description classes. This
1763 requires the existence of:
 - 1764 ○ semantic models that provide normative descriptions of the utilized terms, where the
1765 models may range from a simple dictionary of terms to an ontology showing complex
1766 relationships and capable of supporting enhanced reasoning;
 - 1767 ○ mechanisms to support the storage, referencing, and access to these semantic models;
 - 1768 ○ configuration management mechanisms to capture the normative description of each
1769 semantic model and to apply a unique identifier in a manner consistent with an identified
1770 versioning scheme;
 - 1771 ○ one or more mechanisms to support the storage, referencing, and access to conversion
1772 relationships between semantic models, and the mechanisms to carry out such
1773 conversions.
- 1774 • Descriptions include reference to policies defining conditions of use. In this sense, policies are
1775 also resources that need to be visible, discoverable, and accessible. This requires the existence
1776 of (as also enumerated under governance):
 - 1777 ○ description of policies, including a unique identifier for the policy and a sufficient, and
1778 preferably a machine processable, representation of the meaning of terms used to
1779 describe the policy, its functions, and its effects;
 - 1780 ○ one or more discovery mechanisms that enable searching for policies that best meet the
1781 search criteria specified by the service participant; where the discovery mechanism has
1782 access to the individual policy descriptions, possibly through some repository
1783 mechanism;
 - 1784 ○ accessible storage of policies and policy descriptions, so service participants can access,
1785 examine, and use the policies as defined.
- 1786 • Descriptions include references to metrics that describe the operational characteristics of the
1787 subjects being described. This requires the existence of (as partially enumerated under
1788 governance):
 - 1789 ○ the infrastructure monitoring and reporting information on SOA resources;
 - 1790 ○ possible interface requirements to make accessible metrics information generated;
 - 1791 ○ mechanisms to catalog and enable discovery of which metrics are available for a
1792 described resources and information on how these metrics can be accessed;
 - 1793 ○ mechanisms to catalog and enable discovery of compliance records associated with
1794 policies and contracts that are based on these metrics.
- 1795 • Descriptions of the interactions are important for enabling auditability and repeatability, thereby
1796 establishing a context for results and support for understanding observed change in performance
1797 or results. This requires the existence of:

- 1798 o one or more mechanisms to capture, describe, store, discover, and retrieve interaction
1799 logs, execution contexts, and the combined interaction descriptions;
1800 o one or more mechanisms for attaching to any results the means to identify and retrieve
1801 the interaction description under which the results were generated.
- 1802 • Descriptions may capture very focused information subsets or can be an aggregate of numerous
1803 component descriptions. Service description is an example of an aggregate for which manual
1804 maintenance of the whole would not be feasible. This requires the existence of:
 - 1805 o tools to facilitate identifying description elements that are to be aggregated to assemble
1806 the composite description;
 - 1807 o tools to facilitate identifying the sources of information to associate with the description
1808 elements;
 - 1809 o tools to collect the identified description elements and their associated sources into a
1810 standard, referenceable format that can support general access and understanding;
 - 1811 o tools to automatically update the composite description as the component sources
1812 change, and to consistently apply versioning schemes to identify the new description
1813 contents and the type and significance of change that occurred.
 - 1814 • The description is the source of vital information in establishing willingness to interact with a
1815 resource, reachability to make interaction possible, and compliance with relevant conditions of
1816 use. This requires the existence of:
 - 1817 o one or more discovery mechanisms that enable searching for described resources that
1818 best meet the criteria specified by a service participant;
 - 1819 o tools to appropriately track users of the descriptions and notify them when a new version
1820 of the description is available.

1821 4.2 Service Visibility Model

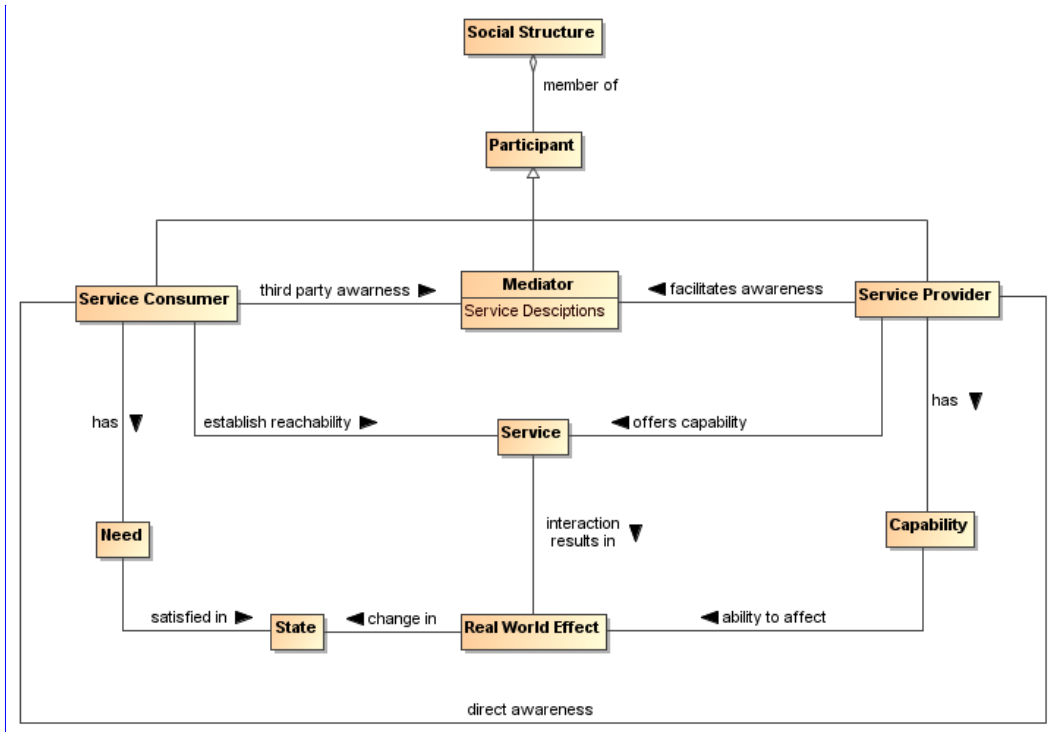
1822 One of the key requirements for participants interacting with each other in the context of a SOA
1823 ecosystem is achieving visibility: before services can interoperate, the participants have to be visible to
1824 each other using whatever means are appropriate. The Reference Model analyzes visibility in terms of
1825 awareness, willingness, and reachability. In this section, we explore how visibility may be achieved.

1826 4.2.1 Visibility to Business

1827 The relationship of visibility to the SOA ecosystem encompasses both human social structures and
1828 automated IT mechanisms. Figure 24 depicts a business setting that is a basis for visibility as related to
1829 the Social Structure Model (Figure 3) in the Participation in a SOA Ecosystem view (see Section 3.1). The
1830 participants acting in the various roles of service consumers, mediators, and service providers may have
1831 direct awareness or mediated awareness where mediated awareness is achieved through some third
1832 party. A consumer's willingness to use a service is reflected by the consumer's presumption of satisfying
1833 goals and needs as these compare with information provided in the service description. Service providers
1834 offer capabilities that have real world effects that result in a change in state. Reachability of the service by
1835 the consumer may lead to interactions that change the state of the SOA ecosystem. The consumer can
1836 measure the change of state to determine if the claims made by description and the real world effects of
1837 consuming the service meet the consumer's needs.

Comment [PFB52]: Issue 294 –
wording available in issue sheet

1838



Comment [KJL53]: To be modified per Issue 85 and possibly Issue 294

Resolve 294 by removing Social Structure and Participant from model.

1840
1841

Figure 24 - Visibility to Business

1842 Visibility and interoperability in a SOA ecosystem requires more than location and interface information. A
1843 meta-model for this broader view of visibility is depicted in Section 4.1. In addition to providing improved
1844 awareness of service capabilities through description of information such as reachability, behavior
1845 models, information models, functionality, and metrics, the service description may identify policies
1846 valuable for determination of willingness to interact.

1847 A mediator using service descriptions may provide event notifications to both consumers and providers
1848 about information relating to the descriptions. One example of this capability is a publish/subscribe model
1849 where the mediator allows consumers to subscribe to service description version changes made by the
1850 provider. Likewise, the mediator may provide notifications to the provider of consumers that have
1851 subscribed to service description updates.

1852 Another important capability in a SOA ecosystem is the ability to narrow visibility to trusted members
1853 within a social structure. Mediators for awareness may provide policy based access to service
1854 descriptions allowing for the dynamic formation of awareness between trusted members.

1855 4.2.2 Visibility

1856 Attaining visibility is described in terms of steps that lead to visibility. Different participant communities can
1857 bring different contexts for visibility within a single social structure, and the same general steps can be
1858 applied to each of the contexts to accomplish visibility.

1859 Attaining SOA visibility requires

- 1860 • service description creation and maintenance,
- 1861 • processes and mechanisms for achieving awareness of and accessing descriptions,
- 1862 • processes and mechanisms for establishing willingness of participants,
- 1863 • processes and mechanisms to determine reachability.

1864 Visibility may occur in stages, i.e. a participant can become aware enough to look or ask for further
1865 description, and with this description, the participant can decide on willingness, possibly requiring
1866 additional description. For example, if a potential consumer has a need for a tree cutting (business)
1867 service, the consumer can use a web search engine to find web sites of providers. The web search
1868 engine (a mediator) gives the consumer links to relevant web pages and the consumer can access those
1869 descriptions. For those prospective providers that satisfy the consumer's criteria, the consumer's
1870 willingness to interact increases. The consumer may contact several tree services to get detailed cost
1871 information (or arrange for an estimate) and may ask for references (further description). The consumer is
1872 likely to establish full visibility and proceed with interaction with the tree service that mutually establishes
1873 visibility.

1874 4.2.2.1 Awareness

1875 An important means for a service participant to be aware of another participant is to have access to a
1876 description of that participant and for the description to have sufficient completeness to establish the other
1877 requirements of visibility.

1878 Awareness is inherently a function of a participant; awareness can be established without any action on
1879 the part of the target participant other than the target providing appropriate descriptions. Awareness is
1880 often discussed in terms of consumer awareness of providers but the concepts are equally valid for
1881 provider awareness of consumers.

1882 Awareness can be decomposed into: creating the descriptions, making them available, and discovering
1883 the descriptions. Discovery can be initiated or it can be by notification. Initiated discovery for business
1884 may require formalization of the required capabilities and resources to achieve business goals.

1885 Achieving awareness in a SOA can range from word of mouth to formal service descriptions in a
1886 standards-based registry-repository. Some other examples of achieving awareness in a SOA are the use
1887 of a web page containing description information, email notifications of descriptions, and document based
1888 descriptions.

1889 A mediator for awareness is a third party participant whose use provides awareness to one or more
1890 consumers of one or more services. Direct awareness is awareness between a consumer and provider
1891 without the use of a third party. The use of a registry/repository can provide awareness as can a Web
1892 page displaying similar information.

1893 Direct awareness may be the result of having previously established an execution context, or direct
1894 awareness may include determining the presence of services and then querying the service directly for
1895 description. As an example, a priori visibility of some sensor device may provide the means for interaction
1896 or a query for standardized sensor device metadata may be broadcast to multiple locations. If
1897 acknowledged, the service interface for the device may directly provide description to a consumer so the
1898 consumer can determine willingness to interact.

1899 The same medium for awareness may be direct in one context and may be mediated in another context.
1900 For example, a service provider may maintain a web site with links to the provider's descriptions of
1901 services giving the consumers direct awareness to the provider's services. Alternatively, a community
1902 may maintain a ~~mediated~~ web site with [a search interface that makes use of an index of these \(and](#)
1903 [possibly other\) links to various provider](#) descriptions of services, ~~for and the web site could be used by~~
1904 any number of consumers. More than one ~~mediator approach to mediation~~ may be involved, as different
1905 ~~mediators sources of description~~ may specialize in different ~~mediation~~ functions [whose use provides](#)
1906 [mediation](#).

1907 Descriptions may be formal or informal. Section 4.1, provides a comprehensive model for service
1908 description that can be used to mediate visibility. Using consistent description taxonomies and standards
1909 based mediated awareness helps provide more effective awareness.

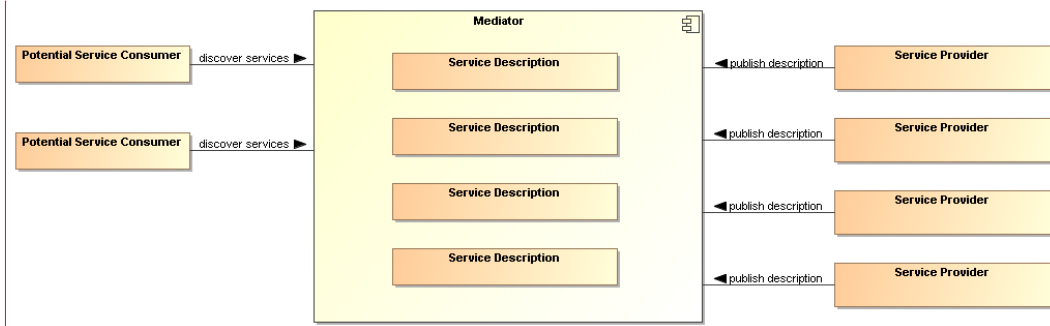
Comment [KJL54]: Issue 187,
188

1910 4.2.2.1.1 Mediated Awareness

1911 Mediated awareness promotes ~~simplification of the overall services infrastructure loose coupling by~~
1912 ~~keeping the consumers and services from explicitly referring to each other. Mediation lets interaction vary~~
1913 ~~independently~~. Rather than all potential service consumers being informed on a continual basis about all

Comment [PFB55]: Issue 189

1914 services, there is a known or agreed upon facility or location that stores and supports discovery and/or
1915 notification related to the service description.



Comment [PFB56]: "Mediator" needs to be replaced with "Mediated Awareness" as component label Issue 190

1916
1917 | Figure 25 - Mediated ~~Service~~ Awareness

Comment [PFB57]: Issue 190

1918 In Figure 25, the potential service consumers perform queries or are notified in order to locate those
1919 services that satisfy their needs. As an example, the telephone book is a mediating registry where
1920 individuals perform manual searches to locate services (i.e. the yellow pages). The telephone book is
1921 also a mediated registry for solicitors to find and notify potential customers (i.e. the white pages).

1922 In mediated service awareness for large and dynamic numbers of service consumers and service
1923 providers, the benefits of utilizing the ~~awareness~~ mediator typically far outweigh the management issues
1924 associated with it. Some of the benefits of mediated service awareness are

Comment [PFB58]: Issue 192

- 1925 • Potential service consumers have a known location for searching thereby eliminating needless
1926 and random searches
- 1927 • Typically a consortium of interested parties (or a sufficiently large corporation) signs up to host
1928 the mediation facility
- 1929 • Standardized tools and methods can be developed and promulgated to promote interoperability
1930 and ease of use.

1931 However, mediated awareness can have some risks associated with it:

- 1932 • A single point of failure. If the ~~awareness mediator/mediation service~~ fails then a large number of
1933 service providers and consumers are potentially adversely affected.
- 1934 • A single point of control. If the ~~central mediation service/awareness mediator~~ is owned by, or
1935 controlled by, someone other than the service consumers and/or providers then the latter may be
1936 put at a competitive disadvantage based on policies of the discovery provider.

Comment [PFB59]: Issue 193

Comment [PFB60]: Issue 194

1937 A common mechanism for mediated awareness is a registry/repository. The registry stores links or
1938 pointers to service description artifacts. The repository in this example is the storage location for the
1939 service description artifacts. Service descriptions can be pushed (publish/subscribe for example) or pulled
1940 from the registry/repository mediator.

1941 Registries/repositories may be referred to as federated when supported functions, such as responding to
1942 discovery requests, are distributed across multiple registry/repository instances.


1943 4.2.2.1.2 Awareness in Complex Social Structures

1944 Awareness applies to one or more communities within one or more social structures where a community
1945 consists of at least one description provider and one description consumer. These communities may be
1946 part of the same social structure or be part of different ones.

1947 In Figure 26, awareness can be between consumers and providers within a single community, multiple
1948 communities, or all communities in the social structure. ~~The Within a social structure, awareness can be~~
1949 ~~encouraged or restricted awareness through its policies,~~ and these policies can affect participant
1950 willingness. The information about policies should be incorporated in the relevant descriptions. ~~The social~~
1951 ~~structure also governs the conditions for establishing contracts, the results of which are reflected in the~~
1952 ~~execution context if interaction is to proceed~~ Additionally, the conditions for establishing contracts are
1953 ~~governed within a social structure.~~

Comment [PFB61]: Issue 195

Comment [PFB62]: Issue 196



1954
1955 *Figure 26 - Awareness in a SOA Ecosystem*

1956 IT policy/contract mechanisms can be used by visibility mechanisms to provide awareness between
1957 communities. The IT mechanisms for awareness may incorporate trust mechanisms to enable awareness
1958 between trusted communities. For example, government organizations may want to limit awareness of an
1959 organization's services to specific communities of interest.

1960 Another common business model for awareness is maximizing awareness to communities within the
1961 social structure, the traditional market place business model. A centralized awareness-mediator often
1962 arises as a provider for this global visibility, a gatekeeper of visibility so to speak. For example, Google is
1963 a centralized awareness-mediator for accessing information on the web. As another example, television
1964 networks have centralized entities providing a level of awareness to communities that otherwise could not
1965 be achieved without going through the television network.

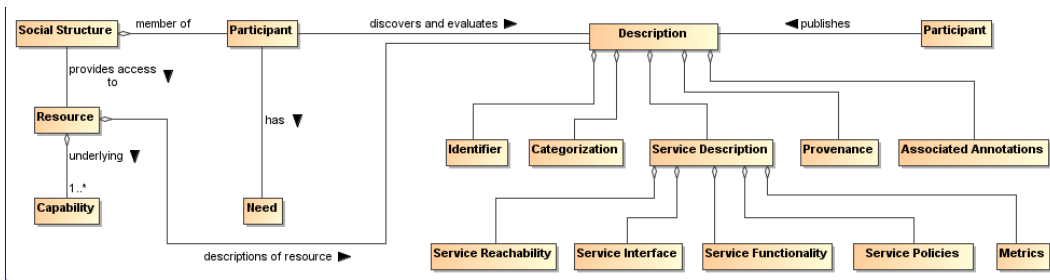
1966 However, mediators have motivations, and they may be selective in which information they choose to
1967 make available to potential consumers. For example, in a secure environment, the mediator may enforce
1968 security policies and make information selectively available depending on the security clearance of the
1969 consumers.

1970 **4.2.2.2 Willingness**

1971 Having achieved awareness, participants use descriptions to help determine their willingness to interact
1972 with another participant. Both awareness and willingness are determined prior to consumer/provider
1973 interaction.

1974
1975

1976



Comment [KJL63]: Modify for issue 261. Also, Change per 3/29/2012 email and attachment even though this figure wasn't explicitly noted.

1977
1978

Figure 27 - Business, Description and Willingness

1979 Figure 27 relates elements of the *Participation in a SOA Ecosystem* view, and elements from the Service
 1980 Description Model to willingness. By having a willingness to interact within a particular social structure, the
 1981 social structure provides the participant access to capabilities based on conditions the social structure
 1982 finds appropriate for its context. The participant can use these capabilities to satisfy goals and objectives
 1983 as specified by the participant's needs.

1984 In Figure 27, information used to determine willingness is defined by Description. Information referenced
 1985 by Description may come from many sources. For example, a mediator for descriptions may provide 3rd
 1986 party annotations for reputation. Another source for reputation may be a participant's own history of
 1987 interactions with another participant. [The contribution of real world effects to providing evidence and](#)
 1988 [establishing the reputation of a participant is discussed with relation to Figure 9.](#)

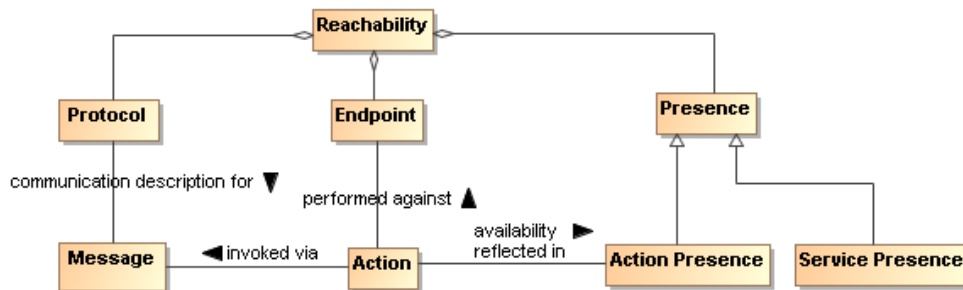
Comment [KJL64]: Issue 262 (part)

1989 A participant inspects functionality for potential satisfaction of needs. Identity is associated with any
 1990 participant, however, identity may or may not be verified. If available, participant reputation may be a
 1991 deciding factor for willingness to interact. Policies and contracts referenced by the description may be
 1992 particularly important to determine the agreements and commitments required for business interactions.
 1993 Provenance may be used for verification of authenticity of a resource.

1994 Mechanisms that aid in determining willingness make use of the artifacts referenced by descriptions of
 1995 services. Mechanisms for establishing willingness could be as simple as rendering service description
 1996 information for human consumption to automated evaluation of functionality, policies, and contracts by a
 1997 rules engine. The rules engine for determining willingness could operate as a policy decision procedure
 1998 as defined in Section 4.4.

1999 4.2.2.3 Reachability

2000 Reachability involves knowing the endpoint, protocol, and presence of a service. At a minimum,
 2001 reachability requires information about the location of the service and the protocol describing the means
 2002 of communication.



2003

2004 Figure 28 - Service Reachability

2005 Endpoint

2006 An endpoint is a reference-able entity, processor or resource against which an action can be
 2007 performed.

2008 **Protocol**

2009 A protocol is a structured means by which details of a service interaction mechanism are defined.

2010 **Presence**

2011 Presence is the measurement of reachability of a service at a particular point in time.

2012 A protocol defines a structured method of communication. Presence is determined by interaction through
2013 a communication protocol. Presence may not be known in many cases until the interaction begins. To
2014 overcome this problem, IT mechanisms may make use of presence protocols to provide the current
2015 up/down status of a service.

2016 Service reachability enables service participants to locate and interact with one another. Each action may
2017 have its own endpoint and also its own protocols associated with the endpoint and whether there is
2018 presence for the action through that endpoint. Presence of a service is an aggregation of the presence of
2019 the service's actions, and the service level may aggregate to some degraded or restricted presence if
2020 some action presence is not confirmed. For example, if error processing actions are not available, the
2021 service can still provide required functionality if no error processing is needed. This implies reachability
2022 relates to each action as well as applying to the service/business as a whole.

2023 **4.2.3 Architectural Implications**

2024 Visibility in a SOA ecosystem has the following architectural implications on mechanisms providing
2025 support for awareness, willingness, and reachability:

- 2026 • Mechanisms providing support for awareness have the following minimum capabilities:
 - 2027 ○ creation of Description, preferably conforming to a standard Description format and
 - 2028 structure;
 - 2029 ○ publishing of Description directly to a consumer or through a third party mediator;
 - 2030 ○ discovery of Description, preferably conforming to a standard for Description discovery;
 - 2031 ○ notification of Description updates or notification of the addition of new and relevant
 - 2032 Descriptions;
 - 2033 ○ classification of Description elements according to standardized classification schemes.
- 2034 • In a SOA ecosystem with complex social structures, awareness may be provided for specific
2035 communities of interest. The architectural mechanisms for providing awareness to communities
2036 of interest require support for:
 - 2037 ○ policies that allow dynamic formation of communities of interest;
 - 2038 ○ trust that awareness can be provided for and only for specific communities of interest, the
 - 2039 bases of which is typically built on encryption technologies.
- 2040 • The architectural mechanisms for determining willingness to interact require support for:
 - 2041 ○ verification of identity and credentials of the provider and/or consumer;
 - 2042 ○ access to and understanding of description;
 - 2043 ○ inspection of functionality and capabilities;
 - 2044 ○ inspection of policies and/or contracts.
- 2045 • The architectural mechanisms for establishing reachability require support for:
 - 2046 ○ the location or address of an endpoint;
 - 2047 ○ verification and use of a service interface by means of a communication protocol;
 - 2048 ○ determination of presence with an endpoint which may only be determined at the point of
 - 2049 interaction but may be further aided by the use of a presence protocol for which the
 - 2050 endpoints actively participate.

2051 **4.3 Interacting with Services Model**

2052 Interaction is the activity involved in using a service to access capability in order to achieve a particular
2053 desired real world effect, where real world effect is the actual result of using a service. An interaction can
2054 be characterized by a sequence of communicative actions. Consequently, interacting with a service, i.e.
2055 participating in joint action with the service—usually accomplished mediated by a series of message

Comment [KJL65]: Issue 202

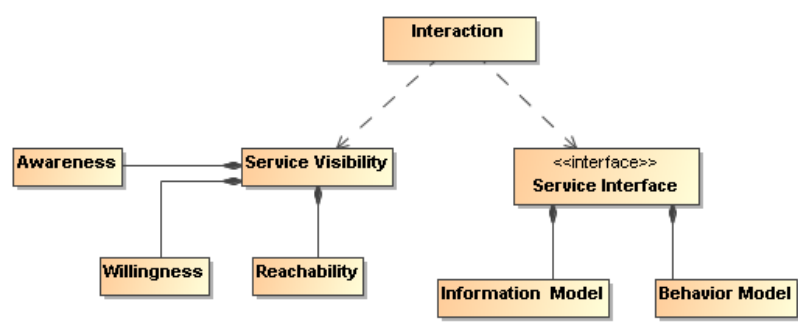
2056 exchanges—involves individual actions performed by both the service and the consumer.⁷ Note that a
2057 participant (or delegate acting on behalf of the participant) can be the sender of a message, the receiver
2058 of a message, or both.

2059 4.3.1 Interaction Dependencies

2060 Recall from the Reference Model that service visibility is the capacity for those with needs and those with
2061 capabilities to be able to interact with each other, and that the service interface is the means by which the
2062 underlying capabilities of a service are accessed. Ideally, the details of the underlying service
2063 implementation are abstracted away by the service interface. [Service] interaction therefore has a direct
2064 dependency on the visibility of the service as well as its implementation-neutral interface (see Figure 29).
2065 Service visibility is composed of awareness, willingness, and reachability, and these are discussed in
2066 Section 4.2. The information related to the service interface description is discussed in Section 4.1.1.3.1,
2067 and the specifics of interaction are detailed in the remainder of Section 4.3. Service visibility is modeled in
2068 Section 4.2.2.

Comment [KJL66]: Comment 203

Comment [KJL67]: Change per
3/29/2012 email and attachment

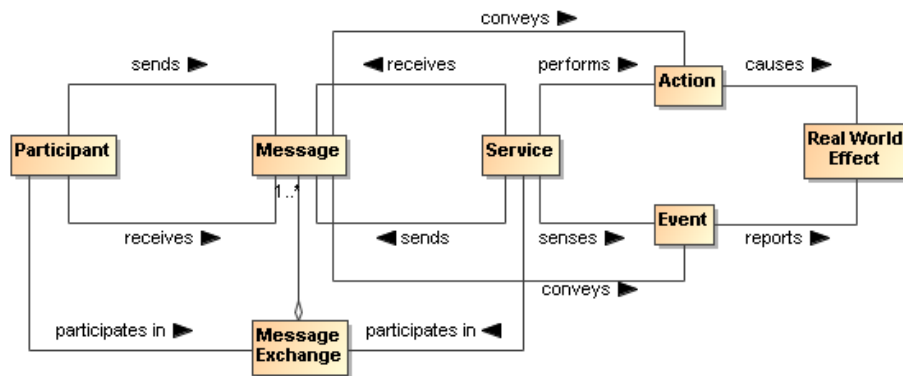


2069
2070 Figure 29 - Interaction dependencies

2071 4.3.2 Actions and Events

2072 The SOA-RAF uses message exchange between service participants to denote actions performed
2073 against and by the service, and to denote events that report on real world effects that are caused by the
2074 service actions. A visual model of the relationship between these concepts is shown in Figure 30.

⁷ In order for multiple actors to participate in a joint action, they must each act according to their role within the joint action. For SOA-based systems, this is achieved through a message exchange style of communication. The concept of “joint action” is further described in Section 3.3.2.



2075
2076 *Figure 30 - A "message" denotes either an action or an event*

2077 Both actions and events, realized by the SOA services, are denoted by the messages. The Reference
2078 Model states that the action model characterizes the "permissible set of actions that may be invoked
2079 against a service." We extend that notion here to include events ~~as part of the event model~~ and that
2080 messages are intended for invoking actions or for notification of events.
2081 In Section 3.3.2 we saw that participants interact with each other in order to participate in joint actions. A
2082 joint action is not itself the same thing as the result of the joint action. When a joint action is participated in
2083 with a service, the real world effect that results may be reported in the form of an event notification.

Comment [PFB68]: Issue 292

2084 4.3.3 Message Exchange

2085 *Message exchange* is the means by which service participants (or their delegates) interact with each
2086 other. There are two primary modes of interaction: joint actions that cause real world effects and
2087 notification of events that report real world effects⁸.

2088 A message exchange is used to affect an action when the messages contain the appropriately formatted
2089 content, are directed towards a particular action in accordance with the action model, and the delegates
2090 involved interpret the message appropriately.

2091 A message exchange is also used to communicate event notifications. An event is an occurrence that is
2092 of interest to some participant; in our case when some real world effect has occurred. Just as action
2093 messages have formatting requirements, so do event notification messages. In this way, the Information
2094 Model of a service must specify the syntax (structure), and semantics (meaning) of the action messages
2095 and event notification messages as part of a service interface. It must also specify the syntax and
2096 semantics of any data that is carried as part of a payload of the action or event notification message. The
2097 Information Model is described in greater detail in the Service Description Model (see Section 4.1).

2098 In addition to the Information Model that describes the syntax and semantics of the messages and data
2099 payloads, exception conditions and error handling in the event of faults (e.g., network outages, improper
2100 message formats, etc.) must be specified or referenced as part of the Service Description.

2101 When a message is used to invoke an action, the correct interpretation typically requires the receiver to
2102 perform an operation, which itself invokes a set of private, internal actions. These *operations* represent
2103 the sequence of (private) actions a service must perform in order to validly participate in a given joint
2104 action.

2105 Similarly, the correct consequence of realizing a real world effect may be to initiate the reporting of that
2106 real world effect via an event notification.

⁸ The notion of "joint" in joint action implies that you have to have a speaker *and* a listener in order to interact.

2107 **Message Exchange**

2108 The means by which joint action and event notifications are coordinated by service participants
 2109 (or delegates).

2110 **Operations**

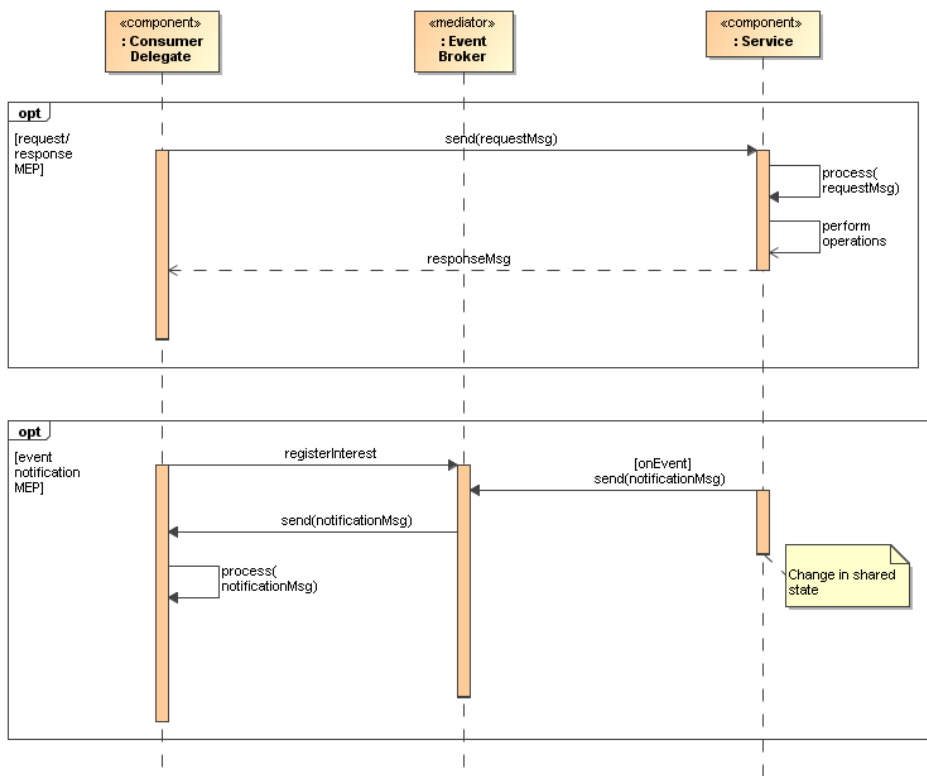
2111 The sequence of actions a service must perform in order to validly participate in a given joint
 2112 action.

2113 **4.3.3.1 Message Exchange Patterns (MEPs)**

2114 The basic temporal aspect of service interaction can be characterized by two fundamental message
 2115 exchange patterns (MEPs):

- 2116 • Request/response to represent how actions cause a real world effect
- 2117 • Event notification to represent how events report a real world effect

2118 This is by no means a complete list of all possible MEPs used for inter- or intra-enterprise messaging but
 2119 it does represent those that are most commonly used in exchange of information and reporting changes
 2120 in state both within organizations and across organizational boundaries.



2121
 2122 *Figure 31 - Fundamental SOA message exchange patterns (MEPs)*

2123 Recall from the Reference Model that the Process Model characterizes “the temporal relationships
 2124 between and temporal properties of actions and events associated with interacting with the service.”
 2125 Thus, MEPs are a key element of the Process Model. The meta-level aspects of the Process Model (just
 2126 as with the Action Model) are provided as part of the Service Description Model (see Section 4.1).

2127 In the UML sequence diagram shown in Figure 31 it is assumed that the service participants (consumer
 2128 and provider) have delegated message handling to hardware or software delegates acting on their behalf.
 2129 In the case of the service consumer, this is represented by the *Consumer Delegate* component. In the

2130 case of the service provider, the delegate is represented by the *Service* component. The message
2131 interchange model illustrated represents a logical view of the MEPs and not a physical view. In other
2132 words, specific hosts, network protocols, and underlying messaging system are not shown, as these tend
2133 to be implementation specific. Although such implementation-specific elements are considered outside
2134 the scope of this document, they are important considerations in modeling the SOA execution context.
2135 Recall from the Reference Model that the *execution context* of a service interaction is “the set of
2136 infrastructure elements, process entities, policy assertions and agreements that are identified as part of
2137 an instantiated service interaction, and thus forms a path between those with needs and those with
2138 capabilities.”

2139 **4.3.3.2 Request/Response MEP**

2140 In a request/response MEP, the Consumer Delegate component sends a request message to the Service
2141 component. The Service component then processes the request message. Based on the content of the
2142 message, the Service component performs the service operation and the associated private actions.
2143 Following the completion of these operations, a response message is returned to the Consumer Delegate
2144 component. The response could be that a step in a process is complete, the initiation of a follow-on
2145 operation, or the return of requested information.⁹

2146 Although the sequence diagram shows a *synchronous* interaction (because the sender of the request
2147 message, i.e., Consumer Delegate, is blocked from continued processing until a response is returned
2148 from the Service) other variations of request/response are valid, including *asynchronous* (non-blocking)
2149 interaction through use of queues, channels, or other messaging techniques.

2150 What is important to convey here is that the request/response MEP represents action, which causes a
2151 real world effect, irrespective of the underlying messaging techniques and messaging infrastructure used
2152 to implement the request/response MEP.

2153 **4.3.3.3 Event Notification MEP**

2154 An event is made visible to interested consumers by means of an event notification message exchange
2155 that reports a real world effect; specifically, a change in shared state between service participants. The
2156 basic event notification MEP takes the form of a one-way message sent by a notifier component (in this
2157 case, the Service component) and received by components with an interest in the event (here, the
2158 Consumer Delegate component).


2159 Often the sending component may not be fully aware of all the components that wish to receive the
2160 notification; particularly in so-called publish/subscribe (“pub/sub”) situations. In event notification message
2161 exchanges, it is rare to have a tightly-coupled link between the sending and the receiving component(s)
2162 for a number of practical reasons. One of the most common needs for pub/sub messaging is the potential
2163 for network outages or communication interrupts that can result in loss of notification of events. Therefore,
2164 a third-party mediator component is often used to decouple the sending and receiving components.

2165 Although this is typically an implementation issue, because this type of third-party decoupling is so
2166 common in event-driven systems, it is warranted for use in modeling this type of message exchange in
2167 the SOA-RAF. This third-party intermediary is shown in Figure 31 as an Event Broker mediator. As with
2168 the request/response MEP, no distinction is made between synchronous versus asynchronous
2169 communication, although asynchronous message exchange is illustrated in the UML sequence diagram
2170 depicted in Figure 31.

2171 **4.3.4 Composition of Services**

2172 Composition of services is the act of aggregating or “composing” a single service from one or more other
2173 services. A simple model of service composition is illustrated in Figure 32.

⁹ 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.



2174
2175 *Figure 32 - Simple model of service composition*

2176 Here, Service A is a service that has an exposed interface IServiceA, which is available to the Consumer Delegate and relies on two other services in its implementation. The Consumer Delegate does not know
2177 that Services B and C are used by Service A, or whether they are used in serial or parallel, or if their
2178 operations succeed or fail. The Consumer Delegate only cares about the success or failure of Service A.
2179 The exposed interfaces of Services B and C (IService B and IServiceC) are not necessarily hidden from
2180 the Consumer Delegate; only the fact that these services are used as part of the composition of Service B
2181 A. In this example, there is no practical reason the Consumer Delegate could not interact with Service B
2182 or Service C in some other interaction scenario.
2183

2184 [While the service composition is opaque from the Consumer Delegate's perspective, it is transparent to](#)
2185 [the service owner. This transparency is necessary for service management. It is possible for a service](#)
2186 [composition to be opaque from one perspective and transparent from another. For example, a service](#)
2187 [may appear to be a single service from the Consumer's Delegate's perspective, but is transparently](#)
2188 [composed of one or more services from a service management perspective. A Service Management](#)
2189 [capability needs to be able to have visibility into the composition in order](#) to properly manage the
2190 dependencies between the services used in constructing the composite service—including managing the
2191 service's lifecycle. The subject of services as management entities is described and modeled in the
2192 *Ownership in a SOA Ecosystem* View of the SOA-RAF and is not further elaborated in this section. The
2193 point to be made here is that there can be different levels of opaqueness or transparency when it comes
2194 to visibility of service composition.

Comment [KJL69]: Issue 93

2195 Services can be composed in a variety of ways, including direct consumer-to-service interaction, by using
2196 programming techniques, or [using an intermediary, such as an orchestration engine leveraging higher](#)
2197 [level orchestration languages, they can be aggregated by means of an aggregation engine approach that](#)
2198 [leverages a service composition scripting language.](#) Such approaches are further elaborated in the
2199 following sub-sections [on service-oriented business processes and collaborations.](#)

Comment [KJL70]: Issue 94, 204

2200 4.3.5 Service Composition of Business Processes and Collaborations

2201 The concepts of business processes and collaborations in the context of [transactions and](#) exchanges
2202 across organizational boundaries are described and modeled as part of the *Participation in a SOA*
2203 *Ecosystem* view of this reference architecture (see Section 3). Here, we focus on the belief that the
2204 [principles involved in the of composition of services \(including but not limited to loose coupling, selective](#)
2205 [transparency and opacity, dynamic interactions\)](#) can be applied to business processes and
2206 collaborations. Of course, business processes and collaborations traditionally represent complex, multi-
2207 step business functions that may involve multiple participants, including internal users, external
2208 customers, and trading partners. Therefore, such complexities cannot simply be ignored when
2209 transforming traditional business processes and collaborations to their service-oriented variants.

Comment [KJL71]: Issue 264

2210 Business Process

2211 A set of [one or more linked steps](#) (activities) that are performed [in accordance with predefined](#)
2212 [logic in order to achieve a required to achieve a certain](#) business outcome.

Comment [PFB72]: Issue 206

2213 Business Collaboration

2214 A [set of interactions among business participants where each participant agrees to perform](#)
2215 [activities that in aggregate will produce a required business outcome.](#)

2216 [Realizing the required business outcomes often involve a combination of business processes and](#)
2217 [business collaborations. Collaborations may be among actors executing formal business practices;](#)
2218 [business processes may call upon other actors who accomplish their activities through collaborative](#)
2219 [efforts. The techniques discussed in the following can be applied to any combination of services that](#)
2220 [instantiate service-oriented business processes or are used as part of service-oriented business](#)
2221 [collaborations.](#)

2222 **4.3.5.1 Service-Oriented Business Processes**

2223 Service orientation as applied to business processes (i.e., “service-oriented business processes”)
2224 [includes both \(1\) abstracting as services the participating activities and rules governing business](#)
2225 [processes and \(2\) using the resulting service to realize the effects of the abstracted process](#)
2226 [means that the aggregation or composition of all of the abstracted activities, flows, and rules that govern a business](#)
2227 [process can themselves be abstracted as a service \[BLOOMBERG/SCHMELZER\].](#)

Comment [PFB73]: Issue 208

2228 When business processes are [implemented as abstracted in this manner and accessed through](#) SOA
2229 services, all of the concepts used to describe and model composition of services that were articulated in
2230 Section 4.3.4 apply. However, there are some important differences between a [composite](#) service that
2231 represents an abstraction of a business process and a [composite](#) service that represents a single-step
2232 business interaction. Business processes have temporal properties and can range from short-lived
2233 processes that execute on the order of minutes or hours to long-lived processes that can execute for
2234 weeks, months, or even years. Further, these processes may involve many participants. [These and are](#)
2235 [may be](#) important considerations for the consumer of a service-oriented business process, [and For](#)
2236 [example, a consumer may need to know details of the business process in order to have confidence in](#)
2237 [the resulting real world effects. In such cases,](#) these temporal properties [along with the meta-level](#)
2238 [aspects of any sub-processes must may need to](#) be articulated as part of the meta-level aspects of the
2239 service-oriented business process in its Service Description, [along with the meta-level aspects of any](#)
2240 [sub-processes that may be of use or need to be visible to the service consumer.](#)

Comment [KJL74]: Issues 95, 209

Comment [kjl75]: Statement relevant with or without service being a composite.

2241 [In addition, a workflow activity represents a unit of work that some actor acting in a described role \(i.e.,](#)
2242 [role player\) is asked to perform. Activities can be broken down into steps with each step representing a](#)
2243 [task for the role player to perform.](#) A technique that is used to compose service-oriented business
2244 processes that are hierarchical (top-down) and self-contained in nature is known as *orchestration*.

Comment [PFB76]: Issue 265

Comment [PFB77]: Issue 210

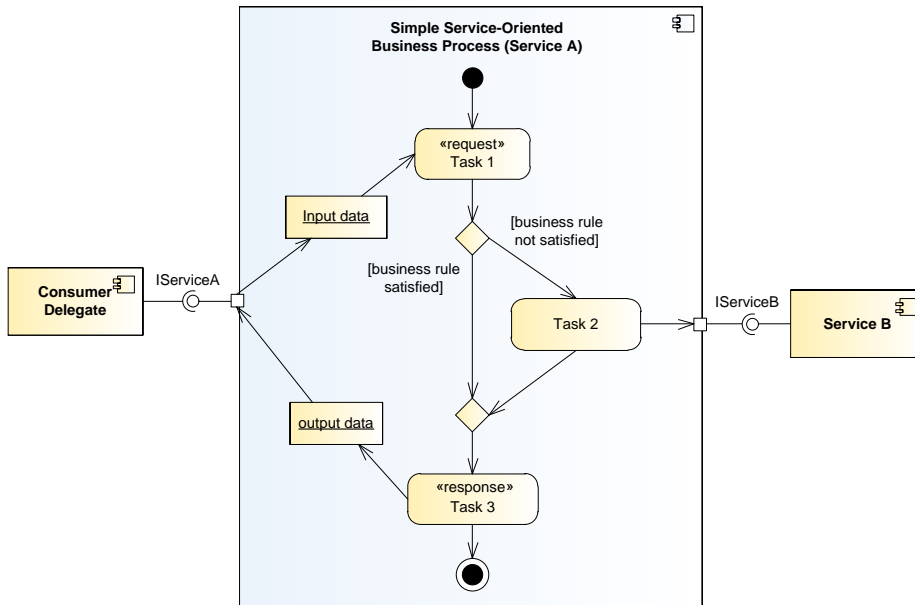
2245 **Orchestration**

2246 A technique used to compose service-oriented business processes that are executed and
2247 coordinated by an actor acting as “conductor.”

2248 [In orchestration, the conductor organizes, controls, and is accountable for the final expected outcome.](#)
2249 [Among the many ways of implementing business processes, a prevalent implementation is using the](#)
2250 [orchestration engine and orchestration language \(domain-specific language designed specifically to](#)
2251 [simplify programming\). An orchestration is typically implemented using a scripting approach to compose](#)
2252 [service-oriented business processes. This typically involves use of a standards-based orchestration](#)
2253 [scripting language. In terms of automation, an orchestration can be mechanized using a business](#)
2254 [process orchestration engine, which is a hardware or software component \(delegate\) responsible for](#)
2255 [acting in the role of central conductor/coordinator responsible for executing the flows that comprise the](#)
2256 [orchestration.](#)

Comment [KJL78]: Issues 96 & 216 plus agreed rewrite

2257 A simple generic example of such an orchestration is illustrated in Figure 33. [Here, Service A is the](#)
2258 [orchestrating service that controls interaction with the orchestrated service, Service B.](#)




2259
2260 Figure 33 - Abstract example of orchestration

Comment [PFB79]: Issue 213

2261 Here, we use a UML activity diagram to model the simple service-oriented business process as it allows
2262 us to capture the major elements of business processes such as the set of related tasks to be performed,
2263 linking between tasks in a logical flow, data that is passed between tasks, and any relevant business
2264 rules that govern the transitions between tasks. A task is a unit of work that an individual, system, or
2265 organization performs and can be accomplished in one or more steps or subtasks. While subtasks can be
2266 readily modeled, they are not illustrated in the orchestration model in Figure 33.

2267 This particular example is based on a request/response MEP and captures how one particular task (Task
2268 2) actually utilizes an externally-provided service, Service B. The entire service-oriented business process
2269 is exposed as Service A that is accessible via its externally visible interface, IServiceA.

2270 Although not explicitly shown in the orchestration model above, it is assumed that there exists a software
2271 or hardware component, i.e., orchestration engine that executes the process flow. Recall that a central
2272 concept to orchestration is that process flow is coordinated and executed by a single conductor delegate;
2273 hence the name "orchestration." [This is illustrated more generally in Figure 34.](#)



Comment [PFB80]: Figure needs formalizing

2274
2275 Figure 34 - General Orchestration Pattern

2276 **4.3-5.14.3.5.2 Service-Oriented Business Collaborations**

2277 Whereas orchestration requires a central controller to execute a predefined business process, service
2278 composition can also be accomplished as a simultaneous cooperation between actors without the
2279 presence of a central control. For such a collaboration, the actors, often considered to be acting as peers,
2280 proceed according to prior agreements for information flow and actions. Business collaborations typically
2281 represent the interaction involved in executing business transactions.

2282 It is important to note that business collaborations represent "peer"-style interactions; in other words,
2283 peers in a business collaboration act as equals. This means that unlike the orchestration of business
2284 processes, there is no single or central entity that coordinates or "conducts" a business collaboration.
2285 These peer styles of interactions typically occur between trading partners that span organizational
2286 boundaries.

2287 Business collaborations can also be service-enabled. For purposes of this Reference Architecture
2288 Foundation, we refer to these such interactions as "service-oriented business collaborations." Service-
2289 oriented business collaborations do not necessarily imply exposing the entire peer-style business
2290 collaboration as a service itself but rather the collaboration uses service-based interchanges.

2291 The technique that is used to compose service-oriented business collaborations in which multiple parties
2292 collaborate in a peer-style as part of some larger business transaction by exchanging messages with
2293 trading partners and external organizations (e.g., suppliers) is known as choreography
2294 [NEWCOMER/LOMOW].

2295 **Choreography**

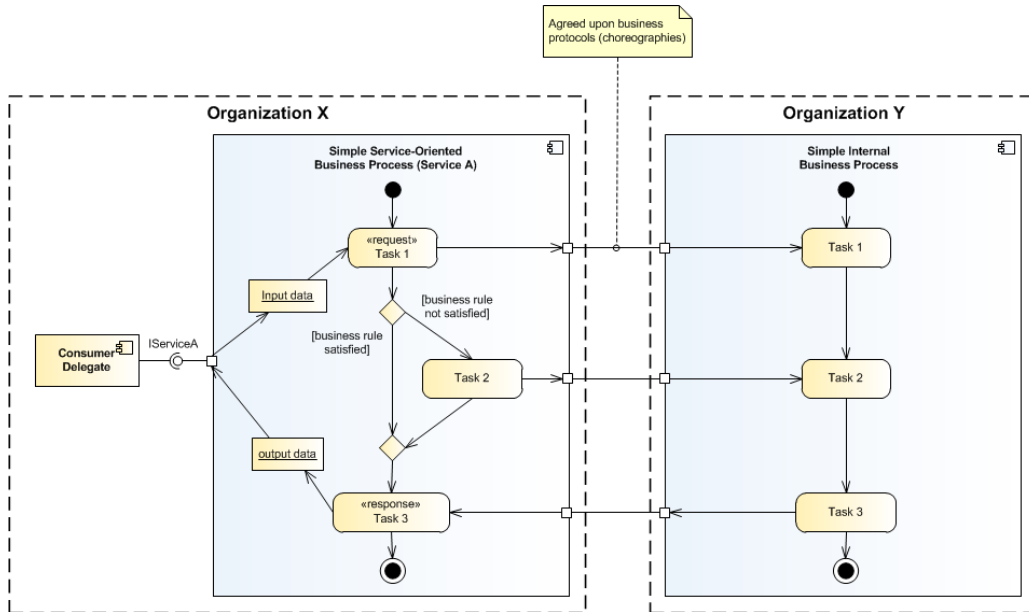
2296 A technique used to engage independent business services into collaborative efforts in order to
2297 achieve a common business outcome based on collective agreements between participants and
2298 with no one in charge over the entire collaboration

2299 For choreography, multiple parties collaborate in a peer-style communication as part of some larger
2300 business transaction by exchanging messages with trading partners and external organizations (e.g.,
2301 suppliers). [NEWCOMER/LOMOW] Choreography It differs from orchestration primarily in that each party
2302 in a business collaboration describes its part in the service interaction. Service-oriented business
2303 collaborations do not necessarily imply exposing the entire peer-style business collaboration as a service
2304 itself but rather the collaboration uses service-based interchanges. Note that choreography as we have
2305 defined it here should not be confused with the term process choreography, which is defined in the
2306 Participation in a SOA Ecosystem view as "the description of the possible interactions that may take place
2307 between two or more participants to fulfill an objective." This is an example of domain-specific
2308 nomenclature that often leads to confusion and why we are making note of it here.

2309 A simple generic example of a choreography is illustrated in Figure 35.

Comment [kjl81]:
Issues 218-221.
Agreed to replace current wording in lines 2185-90 with the essence of proposal from Boris (col. L) but need to incorporate wording that conveys sense of business collaboration being 'simultaneous cooperation between actors that do not require any centralized orchestration'.

Comment [KJL82]: Issue 98



2310
2311 | Figure 35 - Abstract example of choreography

2312 This example, which is a variant of the orchestration example illustrated earlier (in Figure 33), adds trust
 2313 boundaries between two organizations; namely, Organization X and Organization Y. It is assumed that
 2314 these two organizations are peer entities that have an interest in a business collaboration, for example,
 2315 Organization X and Organization Y could be trading partners. Organization X retains the service-oriented
 2316 business process Service A, which is exposed to internal consumers via its provided service interface,
 2317 IServiceA. Organization Y also has a business process that is involved in the business collaboration;
 2318 however, for this example, it is an internal business process that is not exposed to potential consumers
 2319 either within or outside its organizational boundary.

2320 In Figure 35, the communications between Organization X and Organization Y are shown through ports
 2321 where there are “agreed-upon business protocols (choreographies)”. These ports do not explicitly show
 2322 service interfaces in order to emphasize that in the example these are not intended to be generally
 2323 available to any actor in the SOA ecosystem; however, the interfaces should adhere to the principles
 2324 involved in the composition of services.

2325 The scripting language that is message exchanges used for the choreography need to define how and
 2326 when to pass control from one trading partner to another, i.e., between Organization X and Organization
 2327 Y. Defining the business protocols used in the business collaboration involves precisely specifying the
 2328 visible message exchange behavior of each of the parties involved in the protocol, without revealing
 2329 internal implementation details [NEWCOMER/LOMOW].

Comment [PFB83]: Issue 225:
can a message exchange “define”
anything? “indicate”?

2330 In a peer-style service-oriented business collaboration, a choreography scripting language must be
 2331 capable of describing the coordination of those service-oriented processes that cross organizational
 2332 boundaries.

Comment [PFB84]: Issue 226

2333 4.3.6 Architectural Implications of Interacting with Services

2334 Interacting with Services has the following architectural implications on mechanisms that facilitate service
 2335 interaction:

- 2336 • A well-defined service Information Model that:
 - 2337 ○ describes the syntax and semantics of the messages used to denote actions and events;
 - 2338 ○ describes the syntax and semantics of the data payload(s) contained within messages;
 - 2339 ○ documents exception conditions in the event of faults due to network outages, improper
 2340 message/data formats, etc.;

- 2341 ○ is both human readable and machine processable;
- 2342 ○ is referenceable from the Service Description artifact.
- 2343 • A well-defined service Behavior Model (as defined in the SOA-RM) that:
 - 2344 ○ characterizes the knowledge of the actions invoked against the service and events that
 - 2345 report real world effects as a result of those actions;
 - 2346 ○ characterizes the temporal relationships and temporal properties of actions and events
 - 2347 associated in a service interaction;
 - 2348 ○ describe activities involved in a workflow activity that represents a unit of work;
 - 2349 ○ describes the role (s) performed in a service-oriented business process or service-
 - 2350 oriented business collaboration;
 - 2351 ○ is both human readable and machine processable;
 - 2352 ○ is referenceable from the Service Description artifact.
- 2353 • Service composition mechanisms to support orchestration of service-oriented business processes and
- 2354 choreography of service-oriented business collaborations such as:
 - 2355 ○ Declarative and programmatic compositional languages;
 - 2356 ○ Orchestration and/or choreography engines that support multi-step processes as part of a
 - 2357 short-lived or long-lived business transaction;
 - 2358 ○ Orchestration and/or choreography engines that support compensating transactions in
 - 2359 the presences of exception and fault conditions.
- 2360 • Infrastructure services that provides mechanisms to support service interaction, including but not
- 2361 limited to:
 - 2362 ○ mediation services within service interactions based on shared such as message and
 - 2363 event brokers, providers, and/or buses that provide message translation/transformation,
 - 2364 gateway capability, message persistence, reliable message delivery, and/or intelligent
 - 2365 routing semantics;
 - 2366 ○ binding services that support translation and transformation of multiple application-level
 - 2367 protocols to standard network transport protocols;
 - 2368 ○ auditing and logging services that provide a data store and mechanism to record
 - 2369 information related to service interaction activity such as message traffic patterns,
 - 2370 security violations, and service contract and policy violations
 - 2371 ○ security services that provides centralized authorization and authentication support, etc.,
 - 2372 which provide protection against common security threats in a SOA ecosystem;
 - 2373 ○ monitoring services such as hardware and software mechanisms that both monitor the
 - 2374 performance of systems that host services and network traffic during service interaction,
 - 2375 and are capable of generating regular monitoring reports.
- 2376 • A layered and tiered service component architecture that supports multiple message exchange
- 2377 patterns (MEPs) in order to:
 - 2378 ○ promote the industry best practice of separation of concerns that facilitates flexibility in
 - 2379 the presence of changing business requirements;
 - 2380 ○ promote the industry best practice of separation of roles in a service development
 - 2381 lifecycle such that subject matter experts and teams are structured along areas of
 - 2382 expertise;
 - 2383 ○ support numerous standard interaction patterns, peer-to-peer interaction patterns,
 - 2384 enterprise integration patterns, and business-to-business integration patterns.

Comment [KJL85]: Issue 227

Comment [KJL86]: Issue 300

Comment [PFB87]: Issue 101

Comment [PFB88]: Issue 229

Comment [PFB89]: Issue 102, 230

2385 4.4 Policies and Contracts Model

2386 A common phenomenon of many machines and systems is that the scope of potential behavior is much
 2387 broader than is actually needed for a particular circumstance. This is especially true of a system as
 2388 powerful as a SOA ecosystem. As a result, the behavior and performance of the system tend to be under-
 2389 constrained by the implementation; instead, the actual behavior is expressed by means of policies of
 2390 some form. Policies define the choices that stakeholders make; these choices are used to guide the
 2391 actual behavior of the system to the desired behavior and performance.

2392 As noted in Section 3.2.5.2, a policy is an expression of constraints of some form that is promulgated by a
 2393 stakeholder who has the responsibility of ensuring that the constraint is enforceable. In contrast,

Comment [PFB90]: Issue 232

2394 | contracts are **agreements** between participants. ~~However, like policies, it is a necessary part of contracts~~
2395 | ~~that they are enforceable.~~

Comment [PFB91]: Issue 232

2396 | While responsibility for enforcement may differ, both contracts and policies share a common characteristic
2397 | – there is a constraint that must be enforced. In both cases, the mechanisms needed to enforce
2398 | constraints are likely to be identical; in this model, we focus on the issues involved in representing
2399 | policies and contracts and on some of the principles behind their enforcement.

2400 4.4.1 Policy and Contract Representation

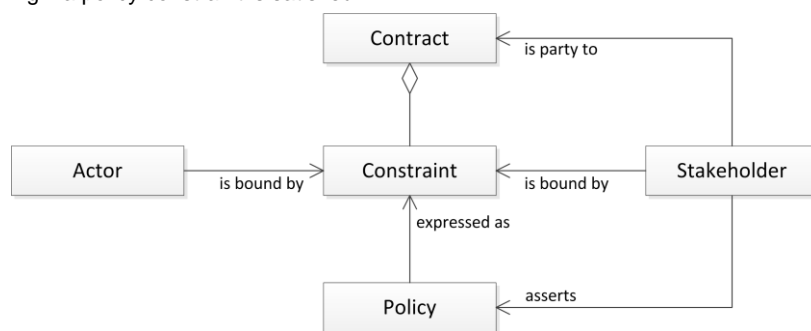
2401 | A **policy constraint** is a specific kind of constraint: the ontology of policies and contracts includes the
2402 | core concepts of permission, obligation, owner, and subject. In addition, it may be necessary to be able
2403 | combine policy constraints and to be able to resolve policy conflicts.

2404 4.4.1.1 Policy Framework

2405 Policy Framework

2406 | A policy framework is a language in which policy constraints may be expressed.

2407 | A policy framework combines syntax for expressing policy constraints together with a decision procedure
2408 | for determining if a policy constraint is satisfied.



2409 |
2410 | *Figure 36 - Policies and Contracts*

2411 | We can characterize a policy framework in terms of a logical framework and an ontology of policies. The
2412 | policy ontology details specific kinds of policy constraints that can be expressed; and the logical
2413 | framework is a 'glue' that allows us to express combinations of policies.

2414 Logical Framework

2415 | A logical framework is a linguistic framework consisting of a syntax – a way of writing expressions
2416 | – and a semantics – a way of interpreting the expressions.

2417 Policy Ontology

2418 | A policy ontology is a formalization of a set of concepts that are relevant to forming policy
2419 | expressions.

2420 | For example, a policy ontology that allows identification of simple constraints – such as the existence of a
2421 | property, or that a value of a property should be compared to a fixed value – is often enough to express
2422 | many basic constraints.

2423 | Included in many policy ontologies are the basic signals of permissions and obligations. Some policy
2424 | frameworks are sufficiently constrained that there is no possibility of representing an obligation; in which
2425 | case there is often no need to 'call out' the distinction between permissions and obligations.

2426 | The logical framework is also a strong determiner of the expressivity of the policy framework: the richer
2427 | the logical framework, the richer the set of policy constraints that can be expressed. However, there is a
2428 | strong inverse correlation ~~between such that increasing~~ expressivity ~~and yields less~~ ease and ~~greater~~
2429 | ~~inefficiency of implementation.~~

Comment [PFB92]: Issue 272

2430 | In the discussion that follows we assume the following basic policy ontology:

2431 **Policy Owner**

2432 A policy owner is a stakeholder that asserts and enforces the policy.

2433 **Policy Subject**

2434 A policy subject is an actor who is subject to the constraints of a policy or contract.

2435 **Policy Constraint**

2436 A policy constraint is a measurable and enforceable proposition that characterizes the constraint
2437 that the policy is about.

2438 **Policy Object**

2439 A policy object is an identifiable state, action or resource that is potentially constrained by the
2440 policy.

2441 **4.4.2 Policy and Contract Enforcement**

2442 The enforcement of policy constraints has to address two core problems: how to enforce the atomic policy
2443 constraints, and how to enforce combinations of policy constraints. In addition, it is necessary to address
2444 the resolution of policy conflicts. [Contracts are the documented agreement between two or more parties](#)
2445 [but otherwise have the same enforcement requirements as policies.](#)

Comment [KJL93]: Issue 237

2446 **4.4.2.1 Enforcing Simple Policy Constraints**

2447 The two primary kinds of policy constraint – permission and obligation – naturally lead to different styles
2448 of enforcement. A permission constraint must typically be enforced prior to the policy subject invoking the
2449 policy object. On the other hand, an obligation constraint must typically be enforced after the fact through
2450 some form of auditing process and remedial action.

2451 For example, if a communications policy required that all communication be encrypted, this is enforceable
2452 at the point of communication: any attempt to communicate a message that is not encrypted can be
2453 blocked.

2454 Similarly, an obligation to pay for services rendered is enforced by ensuring that payment arrives within a
2455 reasonable period of time. Invoices are monitored for prompt (or lack of) payment.

2456 The key concepts in enforcing both forms of policy constraint are the policy decision and the policy
2457 enforcement.

2458 **Policy Decision**

2459 A policy decision is a determination as to whether a given policy constraint is satisfied.

2460 A policy decision is effectively a measurement of some state – typically a portion of the SOA ecosystem's
2461 **shared state**. This implies a certain *timeliness* in the measuring: a measurement that is too early or is too
2462 late does not actually help in determining if the policy constraint is satisfied appropriately.

2463 **Policy Enforcement**

2464 A policy enforcement is the use of a mechanism which limits the behavior and/or state of policy
2465 subjects to comply with a policy decision.

2466 A policy enforcement implies the use of some mechanism to ensure compliance with a policy decision.
2467 The range of mechanisms is completely dependent on the kinds of atomic policy constraints that the
2468 policy framework may support. As noted above, the two primary styles of constraint – permission and
2469 **obligation** –lead to different styles of enforcement.

2470 **4.4.2.2 Conflict Resolution**

2471 Whenever it is possible that more than one policy constraint applies in a given situation, there is the
2472 potential that the policy constraints themselves are not mutually consistent. For example, a policy
2473 constraint that requires communication to be encrypted and a policy constraint that requires an
2474 administrator to read every communication conflict with each other – the two policy constraints cannot
2475 both be satisfied concurrently.

2476 In general, with sufficiently rich policy frameworks, it is not possible to always resolve policy conflicts
2477 automatically. However, a reasonable approach is to augment the policy decision process with simple
2478 policy conflict resolution rules; with the potential for *escalating* a policy conflict to human adjudication.

2479 **Policy Conflict**

2480 A policy conflict exists between two or more policy constraints in a policy decision process if the
2481 satisfaction of one or more policy constraints leads directly to the violation of one or more other
2482 policy constraints.

2483 **Policy Conflict Resolution**

2484 A policy conflict resolution rule is a way of determining which policy constraints should prevail if a
2485 policy conflict occurs.

2486 The inevitable consequence of policy conflicts is that it is not possible to guarantee that all policy
2487 constraints are satisfied at all times. This, in turn, implies a certain *flexibility* in the application of policy
2488 constraints: each individual constraint may not always be honored.

2489 **4.4.3 Architectural Implications**

2490 The key choices that must be made in a system of policies center on the policy framework, policy
2491 enforcement, and conflict resolution

- 2492 • There SHOULD be a standard policy framework that is adopted across ownership domains within the
2493 SOA ecosystem:
 - 2494 ○ This framework MUST permit the expression of simple policy constraints
 - 2495 ○ The framework MAY allow (to a varying extent) the combination of policy constraints,
2496 including
 - 2497 • Both positive and negative constraints
 - 2498 • Conjunctions and disjunctions of constraints
 - 2499 • The quantification of constraints
 - 2500 ○ The framework MUST at least allow the policy subject and the policy object to be identified as
2501 well as the policy constraint.
 - 2502 ○ The framework MAY allow further structuring of policies into modules, inheritance between
2503 policies and so on.
- 2504 • There SHOULD be mechanisms that facilitate the application of policies:
 - 2505 ○ There SHOULD be mechanisms that allow policy decisions to be made, consistent with the
2506 policy frameworks.
 - 2507 ○ There SHOULD be mechanisms to enforce policy decisions
 - 2508 • There SHOULD be mechanisms to support the measurement of whether certain
2509 policy constraints are satisfied, or to what degree they are satisfied.
 - 2510 • Such enforcement mechanisms MAY include support for both permission-style
2511 constraints and obligation-style constraints.
 - 2512 • Enforcement mechanisms MAY support the simultaneous enforcement of multiple
2513 policy constraints across multiple points in the SOA ecosystem.
 - 2514 ○ There SHOULD be mechanisms to resolve policy conflicts
 - 2515 • This MAY involve escalating policy conflicts to human adjudication.
 - 2516 ○ There SHOULD be mechanisms that support the management and promulgation of policies.

2517 **5 Ownership in a SOA Ecosystem View**

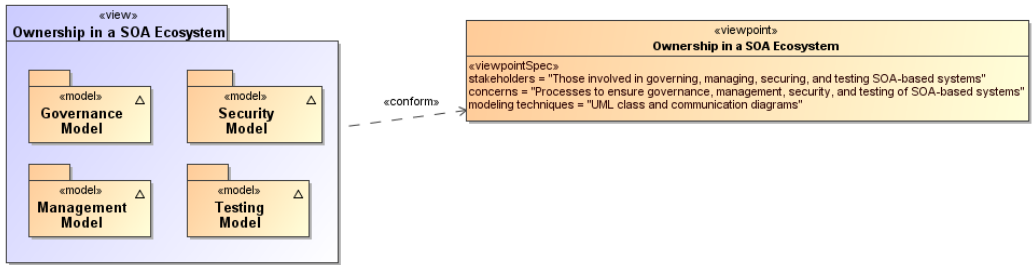
2518 *Governments are instituted among Men,*
2519 *deriving their just power from the consent of the governed*
2520 *American Declaration of Independence*

2521
2522 The *Ownership in a SOA Ecosystem View* focuses on the issues, requirements and responsibilities
2523 involved in owning a SOA-based system.

2524 Ownership of a SOA-based system in a SOA ecosystem raises significantly different challenges to
2525 owning other complex systems – such as Enterprise suites – because there are strong limits on the
2526 control and authority of any one party when a system spans multiple ownership domains.

2527 Even when a SOA-based system is deployed internally within an organization, there are multiple internal
2528 stakeholders involved and there may not be a simple hierarchy of control and management. Thus, an
2529 early consideration of how multiple boundaries affect SOA-based systems provides a firm foundation for
2530 dealing with them in whatever form they are found rather than debating whether the boundaries should
2531 exist.

2532 This view focuses on the governance and management of SOA-based systems, on the security
2533 challenges involved in running a SOA-based system, and testing challenges.



2534
2535 *Figure 37 - Model Elements Described in the Ownership in a SOA Ecosystem View*

2536 The following subsections present models of these functions.

2537 **5.1 Governance Model**

2538 The Reference Model defines Service Oriented Architecture as an architectural paradigm for organizing
2539 and utilizing distributed capabilities that may be under the control of different ownership domains [SOA-
2540 RM]. Consequently, it is important that organizations that plan to engage in service interactions adopt
2541 governance policies and procedures sufficient to ensure that there is standardization across both internal
2542 and external organizational boundaries to promote the effective creation and use of SOA-based services.

2543 **5.1.1 Understanding Governance**

2544 **5.1.1.1 Terminology**

2545 Governance is about making decisions that are aligned with the overall organizational strategy and
2546 culture of the enterprise. **[Gartner]** It specifies the decision rights and accountability framework to
2547 encourage desirable behaviors **[Weill/Ross-MIT Sloan School]** towards realizing the strategy and
2548 defines incentives (positive or negative) towards that end. It is less about overt control and strict
2549 adherence to rules, and more about guidance and effective and equitable usage of resources to ensure
2550 sustainability of an organization’s strategic objectives. **[TOGAF v8.1]**

2551 To accomplish this, governance requires organizational structure and processes and must identify who
2552 has authority to define and carry out its mandates. It must address the following questions:

- 2553 1. what decisions must be made to ensure effective management and use?,
- 2554 2. who should make these decisions?,
- 2555 3. how will these decisions be made and monitored? , and
- 2556 4. how will these decisions be communicated?

2557 The intent is to achieve goals, add value, and reduce risk.

2558 Within a single ownership domain such as an enterprise, generally there is a hierarchy of governance
2559 structures. Some of the more common enterprise governance structures include corporate governance,
2560 technology governance, IT governance, and architecture governance [TOGAF v8.1]. These governance
2561 structures can exist at multiple levels (global, regional, and local) within the overall enterprise.

2562 It is often asserted that SOA governance is a specialization of IT governance as there is a natural
2563 hierarchy of these types of governance structures; however, the focus of SOA governance is less on
2564 decisions to ensure effective management and use of IT as it is to ensure effective management and use
2565 of SOA-based systems. Certainly, SOA governance must still answer the basic questions also associated
2566 with IT governance, i.e., who should make the decisions, and how these decisions will be made and
2567 monitored.

2568 **5.1.1.2 Relationship to Management**

2569 There is often confusion centered on the relationship between governance and management. As
2570 described earlier, governance is concerned with decision making. Management, on the other hand, is
2571 concerned with execution. Put another way, governance describes the world as leadership wants it to be;
2572 management executes activities that intend to make the leadership's desired world a reality. Where
2573 governance determines who has the authority and responsibility for making decisions and the
2574 establishment of guidelines for how those decisions should be made, management is the actual process
2575 of making, implementing, and measuring the impact of those decisions [Loeb]. Consequently,
2576 governance and management work in concert to ensure a well-balanced and functioning organization as
2577 well as an ecosystem of inter-related organizations. In the sections that follow, we elaborate further on the
2578 relationship between governance and management in terms of setting and enforcing service policies,
2579 contracts, and standards as well as addressing issues surrounding regulatory compliance.

2580 **5.1.1.3 Why is SOA Governance Important?**

2581 One of the hallmarks of SOA that distinguishes it from other architectural paradigms for distributed
2582 computing is the ability to provide a uniform means to offer, discover, interact with and use capabilities
2583 (as well the ability to compose new capabilities from existing ones) all in an environment that transcends
2584 domains of ownership. Consequently, ownership, and issues surrounding it, such as obtaining acceptable
2585 terms and conditions (T&Cs) in a contract, is one of the primary topics for SOA governance. Generally, IT
2586 governance does not include T&Cs, for example, as a condition of use as its primary concern.

2587 Just as other architectural paradigms, technologies, and approaches to IT are subject to change and
2588 evolution, so too is SOA. Setting policies that allow change management and evolution, establishing
2589 strategies for change, resolving disputes that arise, and ensuring that SOA-based systems continue to
2590 fulfill the goals of the business are all reasons why governance is important to SOA.

2591 **5.1.1.4 Governance Stakeholders and Concerns**

2592 As noted in Section 3.2.1 the participants in a service interaction include the service provider, the service
2593 consumer, and other interested or unintentional third parties. Depending on the circumstances, it may
2594 also include the owners of the underlying capabilities that the SOA services access. Governance must
2595 establish the policies and rules under which duties and responsibilities are defined and the expectations
2596 of participants are grounded. The expectations include transparency in aspects where transparency is
2597 mandated; trust in the impartial and consistent application of governance; and assurance of reliable and
2598 robust behavior throughout the SOA ecosystem.

2599 **5.1.2 A Generic Model for Governance**

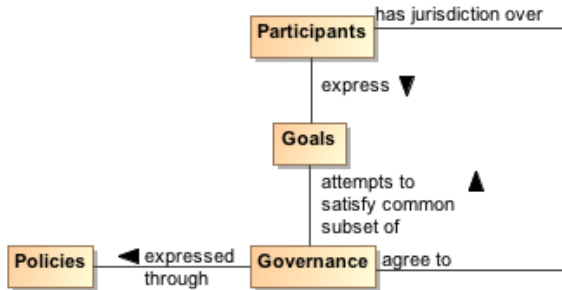
2600 **Governance**

2601 Governance is the prescribing of conditions and constraints consistent with satisfying common
2602 goals and the structures and processes needed to define and respond to actions taken towards
2603 realizing those goals.

2604 The following is a generic model of governance represented by segmented models that begin with
2605 motivation and proceed through measuring compliance. It is not all-encompassing but a focused subset
2606 that captures the aspects necessary to describe governance for SOA. It does not imply that practical
2607 application of governance is a single, isolated instance of these models; in reality, there may be
2608 hierarchical and parallel chains of governance that deal with different aspects or focus on different goals.
2609 This is discussed further in section 5.1.2.5. The defined models are simultaneously applicable to each of
2610 the overlapping instances.

2611 A given enterprise may already have portions of these models in place. To a large extent, the models
2612 shown here are not specific to SOA; discussions on direct applicability begin in section 5.1.3.

2613 **5.1.2.1 Motivating Governance**



2614
2615 *Figure 38 - Motivating Governance*

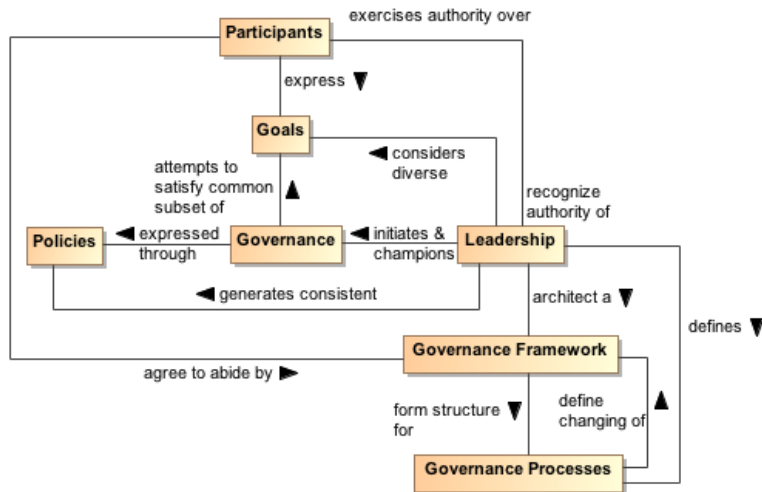
2616 An organizational domain such as an enterprise is made up of participants who may be individuals or
2617 groups of individuals forming smaller organizational units within the enterprise. The overall business
2618 strategy should be consistent with the goals of the participants; otherwise, the business strategy would
2619 not provide value to the participants and governance towards those ends becomes difficult if not
2620 impossible. This is not to say that an instance of governance simultaneously satisfies all the goals of all
2621 the participants; rather, the goals of any governance instance must sufficiently satisfy a useful subset of
2622 each participant's goals so as to provide value and ensure the cooperation of all the participants.

2623 A policy is the formal characterization of the conditions and constraints that governance deems as
2624 necessary to realize the goals which it is attempting to satisfy. Policy may identify required conditions or
2625 actions or may prescribe limitations or other constraints on permitted conditions or actions. For example,
2626 a policy may prescribe that safeguards must be in place to prevent unauthorized access to sensitive
2627 material. It may also prohibit use of computers for activities unrelated to the specified work assignment.
2628 Policy is made operational through the promulgation and implementation of Rules and Regulations (as
2629 defined in section 5.1.2.3).

2630 As noted in section 4.4.2, policy may be asserted by any participant or on behalf of the participant by its
2631 organization. Part of the purpose of governance is to arbitrate among diverse goals of participants and
2632 the diverse policies articulated to realize those goals. The intent is to form a consistent whole that allows
2633 governance to minimize ambiguity about its purpose. While resolving all ambiguity would be an ideal, it is
2634 unlikely that all inconsistencies will be identified and resolved before governance becomes operational.

2635 For governance to have effective jurisdiction over participants, there must be some degree of agreement
2636 by all participants that they will abide by the governance mandates. A minimal degree of agreement often
2637 presages participants who "slow-roll" if not actively rejecting compliance with Policies that express the
2638 specifics of governance.

2639 **5.1.2.2 Setting Up Governance**



2640
2641 *Figure 39 - Setting Up Governance*

2642 **Leadership**

2643 Leadership is the entity who has the responsibility and authority to generate consistent policies
2644 through which the goals of governance can be expressed and to define and champion the
2645 structures and processes through which governance is realized.

2646 **Governance Framework**

2647 The Governance Framework is a set of organizational structures that enable governance to be
2648 consistently defined, clarified, and as needed, modified to respond to changes in its domain of
2649 concern.

2650 **Governance Processes**

2651 Governance Processes are the defined set of activities that are performed within the Governance
2652 Framework to enable the consistent definition, application, and as needed, modification of Rules
2653 that organize and regulate the activities of participants for the fulfillment of expressed policies.
2654 (See section 5.1.2.3 for elaboration on the relationship of Governance Processes and Rules.)

2655 As noted earlier, governance requires an appropriate organizational structure and identification of who
2656 has authority to make governance decisions. In Figure 39, the entity with governance authority is
2657 designated the Leadership. This is someone, possibly one or more of the participants, which participants
2658 recognize as having authority for a given purpose or over a given set of issues or concerns.

2659 The Leadership is responsible for prescribing or delegating a working group to prescribe the Governance
2660 Framework that forms the structure for Governance Processes that define how governance is to be
2661 carried out. This does not itself define the specifics of how governance is to be applied, but it does
2662 provide an unambiguous set of procedures that should ensure consistent actions which participants agree
2663 are fair and account for sufficient input on the subjects to which governance is applied.

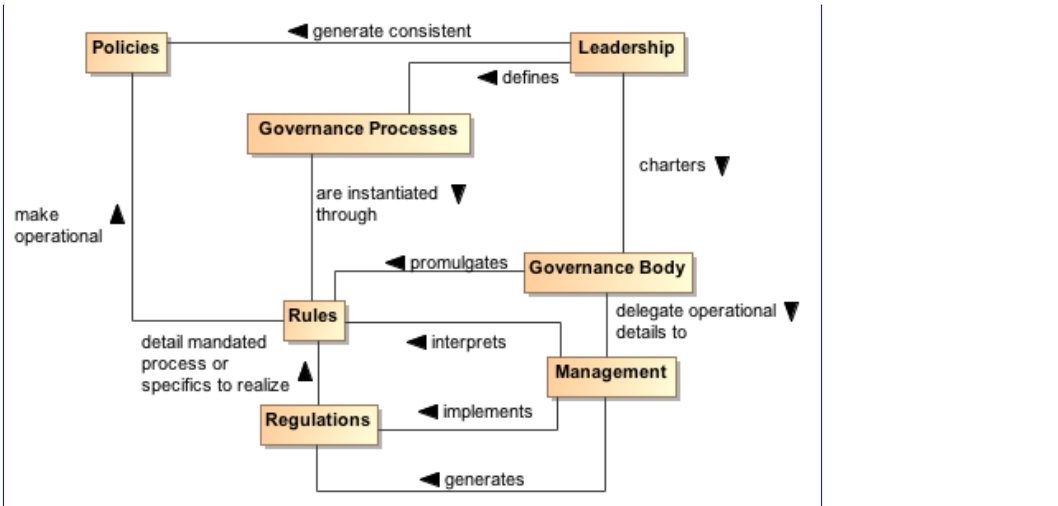
2664 The participants may be part of the working group that codifies the Governance Framework and
2665 Processes. When complete, the participants must acknowledge and agree to abide by the products
2666 generated through application of this structure.

2667 The Governance Framework and Processes are often documented in the constitution or charter of a body
2668 created or designated to oversee governance. This is discussed further in the next section. Note that the
2669 Governance Processes should also include those necessary to modify the Governance Framework itself.

2670 An important function of Leadership is not only to initiate but also to be the consistent champion of
2671 governance. Those responsible for carrying out governance mandates must have Leadership who make

2672 it clear to participants that expressed Policies are seen as a means to realizing established goals and that
2673 compliance with governance is required.

2674 5.1.2.3 Carrying Out Governance



Comment [KJL94]: Comment 115: Regulations derived from Rules

2675
2676 *Figure 40 - Carrying Out Governance*

2677 **Rule**

2678 A Rule is a prescribed guide for carrying out activities and processes leading to desired results,
2679 e.g. the operational realization of policies.

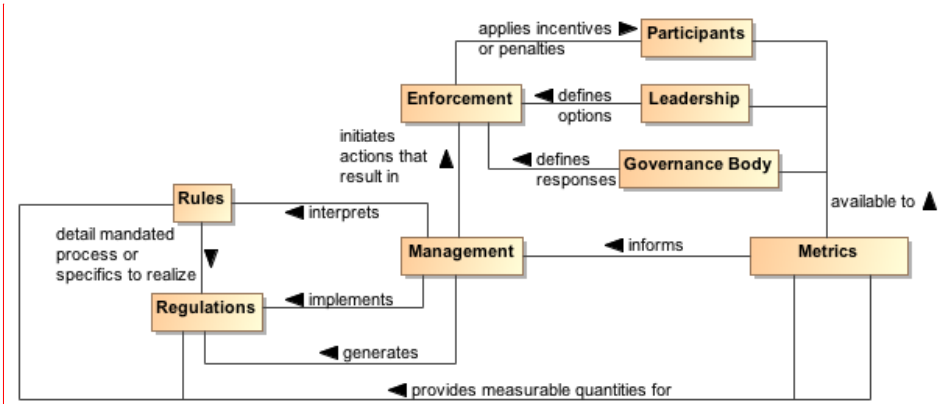
2680 **Regulation**

2681 A Regulation is a mandated process or the specific details that derive from the interpretation of
2682 Rules and lead to measureable quantities against which compliance can be measured.

2683 To carry out governance, Leadership charts a Governance Body to promulgate the Rules needed to
2684 make the Policies operational. The Governance Body acts in line with Governance Processes for its rule-
2685 making process and other functions. Whereas Governance is the setting of Policies and defining the
2686 Rules that provide an operational context for Policies, Governance Body may delegate the operational
2687 details of governance to Management. Management generates Regulations that specify details for Rules
2688 and other procedures to implement both Rules and Regulations. For example, Leadership could set a
2689 Policy that all authorized parties should have access to data, the Governance Body would promulgate a
2690 Rule that PKI certificates are required to establish identity of authorized parties, and Management can
2691 specify a Regulation of who it deems to be a recognized PKI issuing body. In summary, Policy is a
2692 predicate to be satisfied and Rules prescribe the activities by which that satisfying occurs. A number of
2693 rules may be required to satisfy a given policy; the carrying out of a rule may contribute to several policies
2694 being realized.

2695 Whereas the Governance Framework and Processes are fundamental for having participants
2696 acknowledge and commit to compliance with governance, the Rules and Regulations provide operational
2697 constraints that may require resource commitments or other levies on the participants. It is important for
2698 participants to consider the framework and processes to be fair, unambiguous, and capable of being
2699 carried out in a consistent manner and to have an opportunity to formally accept or ratify this situation.
2700 Rules and Regulations, however, do not require individual acceptance by any given participant although
2701 some level of community comment may be part of the Governance Processes. Having agreed to
2702 governance, the participants are bound to comply or be subject to prescribed mechanisms for
2703 enforcement.

2704 **5.1.2.4 Ensuring Governance Compliance**



Comment [PFB95]: Issue 115: Regulations derived from Rules (reverse arrow)

2705
2706 *Figure 41 - Ensuring Governance Compliance*

2707 Setting Rules and Regulations does not ensure effective governance unless compliance can be
2708 measured and Rules and Regulations can be enforced. Metrics are those conditions and quantities that
2709 can be measured to characterize actions and results. Rules and Regulations MUST be based on
2710 collected Metrics or there is no means for Management to assess compliance. The Metrics are available
2711 to the participants, the Leadership, and the Governance Body so what is measured and the results of
2712 measurement are clear to everyone.

2713 The Leadership in its relationship with participants has certain options that can be used for Enforcement.
2714 A common option may be to affect future funding. The Governance Body defines specific enforcement
2715 responses, such as what degree of compliance is necessary for full funding to be restored. It is up to
2716 Management to identify compliance shortfalls and to initiate the Enforcement process.

2717 Note, enforcement does not strictly need to be negative consequences. Management can use Metrics to
2718 identify exemplars of compliance and Leadership can provide options for rewarding the participants. The
2719 Governance Body defines awards or other incentives.

2720 **5.1.2.5 Considerations for Multiple Governance Chains**

2721 As noted in section 5.1.2, instances of the governance model often occur as a tiered arrangement, with
2722 governance at some level delegating specific authority and responsibility to accomplish a focused portion
2723 of the original level's mandate. For example, a corporation may encompass several lines of business and
2724 each line of business governs its own affairs in a manner that is consistent with and contributes to the
2725 goals of the parent organization. Within the line of business, an IT group may be given the mandate to
2726 provide and maintain IT resources, giving rise to IT governance.

2727 In addition to tiered governance, there may be multiple governance chains working in parallel. For
2728 example, a company making widgets has policies intended to ensure they make high quality widgets and
2729 make an impressive profit for their shareholders. On the other hand, Sarbanes-Oxley is a parallel
2730 governance chain in the United States that specifies how the management must handle its accounting
2731 and information that needs to be given to its shareholders. The parallel chains may just be additive or
2732 may be in conflict and require some harmonization.

2733 Being distributed and representing different ownership domains, a SOA participant falls under the
2734 jurisdiction of multiple governance domains simultaneously and may individually need to resolve
2735 consequent conflicts. The governance domains may specify precedence for governance conformance or
2736 it may fall to the discretion of the participant to decide on the course of actions they believe appropriate.

2737 5.1.3 Governance Applied to SOA

2738 5.1.3.1 Where SOA Governance is Different

2739 SOA governance is often discussed in terms of IT governance, but rather than a parent-child relationship,
2740 Figure 42 shows the two as siblings within the general governance described in section 5.1.2. There are
2741 obvious dependencies and a need for coordination between the two, but the idea of aligning IT with
2742 business already demonstrates that resource providers and resource consumers must be working
2743 towards common goals if they are to be productive and efficient. While SOA governance is shown to be
2744 active in the area of infrastructure, it is a specialized concern for having a dependable platform to support
2745 service interaction; a range of traditional IT issues is therefore out of scope of this document. A SOA
2746 governance plan for an enterprise will not of itself resolve shortcomings with the enterprise's IT
2747 governance.

2748 Governance in the context of SOA is that organization of services: that promotes their visibility; that
2749 facilitates interaction among service participants; and that directs that the results of service interactions
2750 are those real world effects as described within the service description and constrained by policies and
2751 contracts as assembled in the execution context.

2752 SOA governance must specifically account for control across different ownership domains, i.e. all the
2753 participants may not be under the jurisdiction of a single governance authority. However, for governance
2754 to be effective, the participants must agree to recognize the authority of the Governance Body and must
2755 operate within the Governance Framework and through the Governance Processes so defined.

2756 SOA governance must account for interactions across ownership boundaries, which may also imply
2757 across enterprise governance boundaries. For such situations, governance emphasizes the need for
2758 agreement that some Governance Framework and Governance Processes have jurisdiction, and the
2759 governance defined must satisfy the goals of the participants for cooperation to continue. A standards
2760 development organization such as OASIS is an example of voluntary agreement to governance over a
2761 limited domain to satisfy common goals.

2762 The specifics discussed in the figures in the previous sections are equally applicable to governance
2763 across ownership boundaries as it is within a single boundary. There is a charter agreed to when
2764 participants become members of the organization, and this charter sets up the structures and processes
2765 to be followed. Leadership may be shared by the leadership of the overall organization and the leadership
2766 of individual groups themselves chartered per the Governance Processes. There are Rules/Regulations
2767 specific to individual efforts for which participants agree to local goals, and Enforcement can be loss of
2768 voting rights or under extreme circumstances, expulsion from the group.

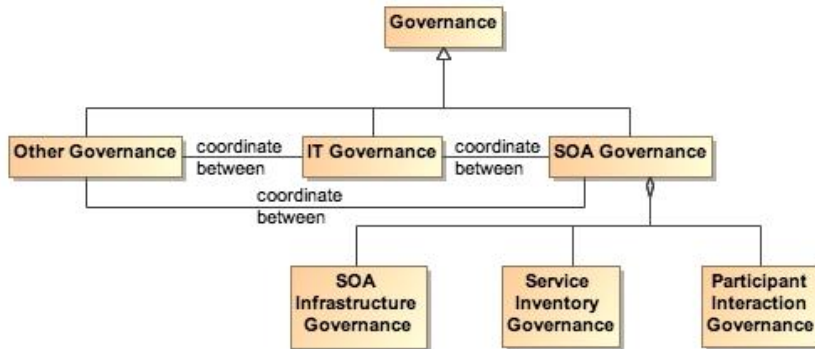
2769 Thus, the major difference for SOA governance is an appreciation for the cooperative nature of the
2770 enterprise and its reliance on furthering common goals if productive participation is to continue.

2771 5.1.3.2 What Must be Governed

2772 An expected benefit of employing SOA principles is the ability to quickly bring resources to bear to deal
2773 with unexpected and evolving situations. This requires a great deal of confidence in the underlying
2774 capabilities that can be accessed and in the services that enable the access. It also requires considerable
2775 flexibility in the ways these resources can be employed. Thus, SOA governance requires establishing
2776 confidence and trust (see Section 3.2.5.1) while instituting a solid framework that enables flexibility,
2777 indicating a combination of strict control over a limited set of foundational aspects but minimum
2778 constraints beyond those bounds.

2779

Comment [PFB96]: Issue 117



2780
2781 *Figure 42 - Relationship Among Types of Governance*

2782 SOA governance applies to three aspects of service definition and use:

- 2783 • SOA infrastructure – the “plumbing” that provides utility functions that enable and support the use
- 2784 of the service
- 2785 • Service inventory – the requirements on a service to permit it to be accessed within the
- 2786 infrastructure
- 2787 • Participant interaction – the consistent expectations with which all participants are expected to
- 2788 comply

2789 **5.1.3.2.1 Governance of SOA Infrastructure**

2790 The SOA infrastructure is likely composed of several families of SOA services that provide access to
2791 fundamental computing business services. These include, among many others, services such as
2792 messaging, security, storage, discovery, and mediation. The provisioning of an infrastructure on which
2793 these services may be accessed and the general realm of those contributing as utility functions of the
2794 infrastructure are a traditional IT governance concern. In contrast, the focus of SOA governance is how
2795 the existence and use of the services enables the SOA ecosystem.

2796 By characterizing the environment as containing families of SOA services, the assumption is that there
2797 may be multiple approaches to providing the business services or variations in the actual business
2798 services provided. For example, discovery could be based on text search, on metadata search, on
2799 approximate matches when exact matches are not available, and numerous other variations. The
2800 underlying implementation of search algorithms are not the purview of SOA governance, but the access
2801 to the resulting service infrastructure enabling discovery must be stable, reliable, and extremely robust to
2802 all operating conditions. Such access enables other specialized SOA services to use the infrastructure in
2803 dependable and predictable ways, and is where governance is important.

2804 **5.1.3.2.2 Governance of the Service Inventory**

2805 Given an infrastructure in which other SOA services can operate, a key governance issue is which SOA
2806 services to allow in the ecosystem. The major concern SHOULD be a definition of well-behaved services,
2807 where the required behavior will inherit their characteristics from experiences with distributed computing
2808 but also evolve with SOA experience. A major requirement for ensuring well-behaved services is
2809 collecting sufficient metrics to know how the service affects the SOA infrastructure and whether it
2810 complies with established infrastructure policies.

2811 Another common concern of service approval is whether there is a possibility of duplication of function by
2812 multiple services. Some governance models talk to a tightly controlled environment where a primary
2813 concern is to avoid any service duplication. Other governance models talk to a market of services where
2814 the consumers have wide choices. For the latter, it is anticipated that the better services will emerge from
2815 market consensus and the availability of alternatives will drive innovation.

2816 Some combination of control and openness will emerge, possibly with a different appropriate balance for
2817 different categories of use. For SOA governance, the issue is less which services are approved but rather
2818 ensuring that sufficient description is available to support informed decisions for appropriate use. Thus,
2819 SOA governance SHOULD concentrate on identifying the required attributes to adequately describe a
2820 service, the required target values of the attributes, and the standards for defining the meaning of the
2821 attributes and their target values. Governance may also specify the processes by which the attribute
2822 values are measured and the corresponding certification that some realized attribute set may imply.

2823 For example, unlimited access for using a service may require a degree of life cycle maturity that has
2824 demonstrated sufficient testing over a certain size community. Alternately, the policy may specify that a
2825 service in an earlier phase of its life cycle may be made available to a smaller, more technically
2826 sophisticated group in order to collect the metrics that would eventually allow the service to advance its
2827 life cycle status.

2828 This aspect of governance is tightly connected to description because, given a well-behaved set of
2829 services, it is the responsibility of the consumer (or policies promulgated by the consumer's organization)
2830 to decide whether a service is sufficient for that consumer's intended use. The goal is to avoid global
2831 governance specifying criteria that are too restrictive or too lax for the local needs of which global
2832 governance has little insight.

2833 Such an approach to specifying governance allows independent domains to describe services in local
2834 terms while still having the services available for informed use across domains. In addition, changes to
2835 the attribute sets within a domain can be similarly described, thus supporting the use of newly described
2836 resources with the existing ones without having to update the description of the entire legacy content.

2837 **5.1.3.2.3 Governance of Participant Interaction**

2838 Finally, given a reliable services infrastructure and a predictable set of services, the third aspect of
2839 governance is prescribing what is required during a service interaction.

2840 Governance would specify adherence to service interface and service reachability parameters and would
2841 require that the result of an interaction MUST correspond to the real world effects as contained in the
2842 service description. Governance would ensure preconditions for service use are satisfied, in particular
2843 those related to security aspects such as user authentication, authorization, and non-repudiation. If
2844 conflicts arise, governance would specify resolution processes to ensure appropriate agreements,
2845 policies, and conditions are met.

2846 It would also rely on sufficient monitoring by the SOA infrastructure to ensure services remain well-
2847 behaved during interactions, e.g. do not use excessive resources or exhibit other prohibited behavior.
2848 Governance would also require that policy agreements as documented in the execution context for the
2849 interaction are observed and that the results and any after effects are consistent with the agreed policies.
2850 Here, governance focuses more on contractual and legal aspects rather than the precursor descriptive
2851 aspects. SOA governance may prescribe the processes by which SOA-specific policies are allowed to
2852 change, but there are probably more business-specific policies that will be governed by processes
2853 outside SOA governance.

2854 **5.1.3.3 Overarching Governance Concerns**

2855 There are numerous governance related concerns whose effects span the three areas just discussed.
2856 One is the area of standards, how these are mandated, and how the mandates may change. The Web
2857 Services standards stack is an example of relevant standards where a significant number are still under
2858 development. In addition, while there are notional scenarios that guide what standards are being
2859 developed, the fact that many of these standards do not yet exist precludes operational testing of their
2860 adequacy or effectiveness as a necessary and sufficient set.

2861 That said, standards are critical to creating a SOA ecosystem where SOA services can be introduced,
2862 used singularly, and combined with other services to deliver complex business functionality. As with other
2863 aspects of SOA governance, the Governance Body should identify the minimum set felt to be needed and
2864 rigorously enforce that that set be used where appropriate. The Governance Body takes care to expand
2865 and evolve the mandated standards in a predictable manner and with sufficient technical guidance that
2866 new services are able to coexist as much as possible with the old, and changes to standards do not
2867 cause major disruptions.

2868 Another area that may see increasing activity as SOA expands is additional regulation by governments
2869 and associated legal institutions. New laws may deal with transactions that are service based, possibly
2870 including taxes on the transactions. Disclosure laws may mandate certain elements of description so both
2871 the consumer and provider act in a predictable environment and are protected from ambiguity in intent or
2872 action. Such laws spawn rules and regulations that will influence the metrics collected for evaluation of
2873 compliance.

2874 **5.1.3.4 Considerations for SOA Governance**

2875 The Reference Architecture definition of a loosely coupled system is one in which the constraints on the
2876 interactions between components are minimal: sufficient to permit interoperability without additional
2877 constraints that may be an artifact of implementation technology. While governance experience for
2878 standalone systems provides useful guides, we must be careful not to apply constraints that would
2879 preclude the flexibility, agility, and adaptability we expect to realize from a SOA ecosystem.

2880 One of the strengths of the SOA paradigm is it can make effective use of diversity rather than requiring
2881 monolithic solutions. Heterogeneous organizations can interact without requiring each conforms to
2882 uniform tools, representation, and processes. However, with this diversity comes the need to adequately
2883 define those elements necessary for consistent interaction among systems and participants, such as
2884 which communication protocol, what level of security, which vocabulary for payload content of messages.
2885 The solution is not always to lock down these choices but to standardize alternatives and standardize the
2886 representations through which an unambiguous identification of the alternative chosen can be conveyed.
2887 For example, the URI standard specifies the URI string, including what protocol is being used, what is the
2888 target of the message, and how parameters may be attached. It does not limit the available protocols, the
2889 semantics of the target address, or the parameters that can be transferred. Thus, as with our definition of
2890 loose coupling, it provides absolute constraints but minimizes which constraints it imposes.

2891 There is not a one-size-fits-all governance but a need to understand the types of things governance is
2892 called upon to do in the context of the goals of the SOA paradigm. Some communities may initially desire
2893 and require very stringent governance policies and procedures while others see need for very little. Over
2894 time, best practices will evolve, resulting in some consensus on a sensible minimum and, except in
2895 extreme cases where it is demonstrated to be necessary, a loosening of strict governance toward the
2896 best practice mean.

2897 A question of how much governance may center on how much time governance activities require versus
2898 how quickly is the system being governed expected to respond to changing conditions. For large single
2899 systems that take years to develop, the governance process could move slowly without having a serious
2900 negative impact. For example, if something takes two years to develop and the steps involved in
2901 governance take two months to navigate, then the governance can go along in parallel and may not have
2902 a significant impact on system response to changes. Situations where it takes as long to navigate
2903 governance requirements as it does to develop a response are examples where governance may need to
2904 be reevaluated as to whether it facilitates or inhibits the desired results. Thus, the speed at which services
2905 are expected to appear and evolve needs to be considered when deciding the processes for control. The
2906 added weight of governance should be appropriate for overall goals of the application domain and the
2907 service environment.

2908 Governance, as with other aspects of any SOA implementation, should start small and be conceptualized
2909 in a way that keeps it flexible, scalable, and realistic. A set of useful guidelines would include:

- 2910 • Do not hardwire things that will inevitably change. For example, develop a system that uses the
2911 representation of policies rather than code the policies into the implementations.
- 2912 • Avoid setting up processes that demo well for three services without considering how they may
2913 work for 300. Similarly, consider whether the display of status and activity for a small number of
2914 services will also be effective for an operator in a crisis situation looking at dozens of services,
2915 each with numerous, sometimes overlapping and sometimes differing activities.
- 2916 • Maintain consistency and realism. A service solution responding to a natural disaster cannot be
2917 expected to complete a 6-week review cycle but be effective in a matter of hours.

2918 5.1.4 Architectural Implications of SOA Governance

2919 The description of SOA governance indicates numerous architectural requirements on the SOA
2920 ecosystem:

- 2921 • Governance is expressed through policies and assumes multiple use of focused policy modules
2922 that can be employed across many common circumstances. This requires the existence of:
 - 2923 ○ descriptions to enable the policy modules to be visible, where the description includes a
2924 unique identifier for the policy and a sufficient, and preferably a machine process-able,
2925 representation of the meaning of terms used to describe the policy, its functions, and its
2926 effects;
 - 2927 ○ one or more discovery mechanisms that enable searching for policies that best meet the
2928 search criteria specified by the service participant; where the discovery mechanism will
2929 have access to the individual policy descriptions, possibly through some repository
2930 mechanism;
 - 2931 ○ accessible storage of policies and policy descriptions, so service participants can access,
2932 examine, and use the policies as defined.
- 2933 • Governance requires that the participants understand the intent of governance, the structures
2934 created to define and implement governance, and the processes to be followed to make
2935 governance operational. This requires the existence of:
 - 2936 ○ an information collection site, such as a Web page or portal, where governance
2937 information is stored and from which the information is always available for access;
 - 2938 ○ a mechanism to inform participants of significant governance events, such as changes in
2939 policies, rules, or regulations;
 - 2940 ○ accessible storage of the specifics of Governance Processes;
 - 2941 ○ SOA services to access automated implementations of the Governance Processes
- 2942 • Governance policies are made operational through rules and regulations. This requires the
2943 existence of:
 - 2944 ○ descriptions to enable the rules and regulations to be visible, where the description
2945 includes a unique identifier and a sufficient, and preferably a machine process-able,
2946 representation of the meaning of terms used to describe the rules and regulations;
 - 2947 ○ one or more discovery mechanisms that enable searching for rules and regulations that
2948 may apply to situations corresponding to the search criteria specified by the service
2949 participant; where the discovery mechanism will have access to the individual
2950 descriptions of rules and regulations, possibly through some repository mechanism;
 - 2951 ○ accessible storage of rules and regulations and their respective descriptions, so service
2952 participants can understand and prepare for compliance, as defined.
 - 2953 ○ SOA services to access automated implementations of the Governance Processes.
 - 2954 ○ Governance implies management to define and enforce rules and regulations.
2955 Management is discussed more specifically in section 5.3, but in a parallel to
2956 governance, management requires the existence of:
 - 2957 ○ an information collection site, such as a Web page or portal, where management
2958 information is stored and from which the information is always available for access;
 - 2959 ○ a mechanism to inform participants of significant management events, such as changes
2960 in rules or regulations;
 - 2961 ○ accessible storage of the specifics of processes followed by management.
- 2962 • Governance relies on metrics to define and measure compliance. This requires the existence of:
 - 2963 ○ the infrastructure monitoring and reporting information on SOA resources;
 - 2964 ○ possible interface requirements to make accessible metrics information generated or
2965 most easily accessed by the service itself.

2966 5.2 Security Model

2967 Security is one aspect of confidence – the confidence in the integrity, reliability, and confidentiality of the
2968 system. In particular, security focuses on those aspects of assurance that involve the accidental or malign
2969 intent of other people to damage or compromise trust in the system and on the availability of SOA-based
2970 systems to perform desired capability.

2971 Security

2972 The set of mechanisms for ensuring and enhancing trust and confidence in the SOA ecosystem.

2973 Providing for security for Service Oriented Architecture is somewhat different than for other contexts;
2974 although many of the same principles apply equally to SOA and to other systems. The fact that SOA
2975 embraces crossing ownership boundaries makes the issues involved with moving data more visible.

2976 As well as securing the movement of data within and across ownership boundaries, security often
2977 revolves around resources: the need to guard certain resources against inappropriate access – whether
2978 reading, writing or otherwise manipulating those resources.

2979 Any comprehensive security solution must take into account the people that are using, maintaining and
2980 managing SOA-based systems. Furthermore, the relationships between them must also be incorporated:
2981 any security assertions that may be associated with particular interactions originate in the people that are
2982 behind the interaction.

2983 We analyze security in terms of the social structures that define the legitimate permissions, obligations
2984 and roles of people in relation to the system, and mechanisms that must be put into place to realize a
2985 secure system. The former are typically captured in a series of security policy statements; the latter in
2986 terms of security guards that ensure that policies are enforced.

2987 How and when to apply these derived security policy mechanisms is directly associated with the
2988 assessment of the *threat model* and a *security response model*. The threat model identifies the kinds of
2989 threats that directly impact the message and/or application of constraints and the response model is the
2990 proposed mitigation to those threats. Properly implemented, the result can be an acceptable level of risk
2991 to the safety and integrity within the SOA ecosystem.

2992 5.2.1 Secure Interaction Concepts

2993 We can characterize secure interactions in terms of key security concepts [ISO/IEC 27002]:
2994 confidentiality, integrity, authentication, authorization, non-repudiation, and availability. The concepts for
2995 secure interactions are well defined in other standards and publications. The security concepts here are
2996 not defined but rather related to the SOA ecosystem perspective of the SOA-RAF.

2997 5.2.1.1 Confidentiality

2998 Confidentiality concerns the protection of privacy of participants in their interactions. Confidentiality refers
2999 to the assurance that unauthorized entities are not able to read messages or parts of messages that are
3000 transmitted.

3001 Note that confidentiality has degrees: in a completely confidential exchange, third parties would not even
3002 be aware that a confidential exchange has occurred. In a partially confidential exchange, the identities of
3003 the participants may be known but the content of the exchange obscured.

3004 5.2.1.2 Integrity

3005 Integrity concerns the protection of information that is exchanged – either from unauthorized writing or
3006 inadvertent corruption. Integrity refers to the assurance that information that has been exchanged has not
3007 been altered.


3008 Integrity is different from confidentiality in that messages that are sent from one participant to another
3009 may be obscured to a third party, but the third party may still be able to introduce his own content into the
3010 exchange without the knowledge of the participants.

3011 Section 5.2.4 describes common computing techniques for providing confidentiality and integrity during
3012 message exchanges.

3013 5.2.1.3 Authentication

3014 ~~Authentication concerns the identity of the participants in an exchange and refers to the means by which~~
3015 ~~one participant can be assured of the identity of other participants. Figure 43 applies authentication to the~~
3016 ~~identity of participants. Authentication is concerned with addressing a need to adequately identify actors~~
3017 ~~in a potential interaction or joint action.~~

3018 A combination of **identifiers** (as discussed in section 3.2.4.1) and other attributes of an actor is typically
 3019 used to achieve this.
 3020 The set of attribute values that claim to identify a specific actor are matched against the set of reference
 3021 values expected for that actor and that are maintained by some trusted authority.
 3022 If the comparison results in a sufficient match, authentication has been achieved.
 3023 Which specific set of attributes is considered adequate will be context-dependent and specifying such
 3024 sets is not within the scope of the SOA-RAF.
 3025 Authentication merely provides an assertion that an actor is the person or agent that it claims to be. Of
 3026 itself, it does not provide a “green light” to proceed with the interaction – this is rather the concern of
 3027 authorization, covered below.



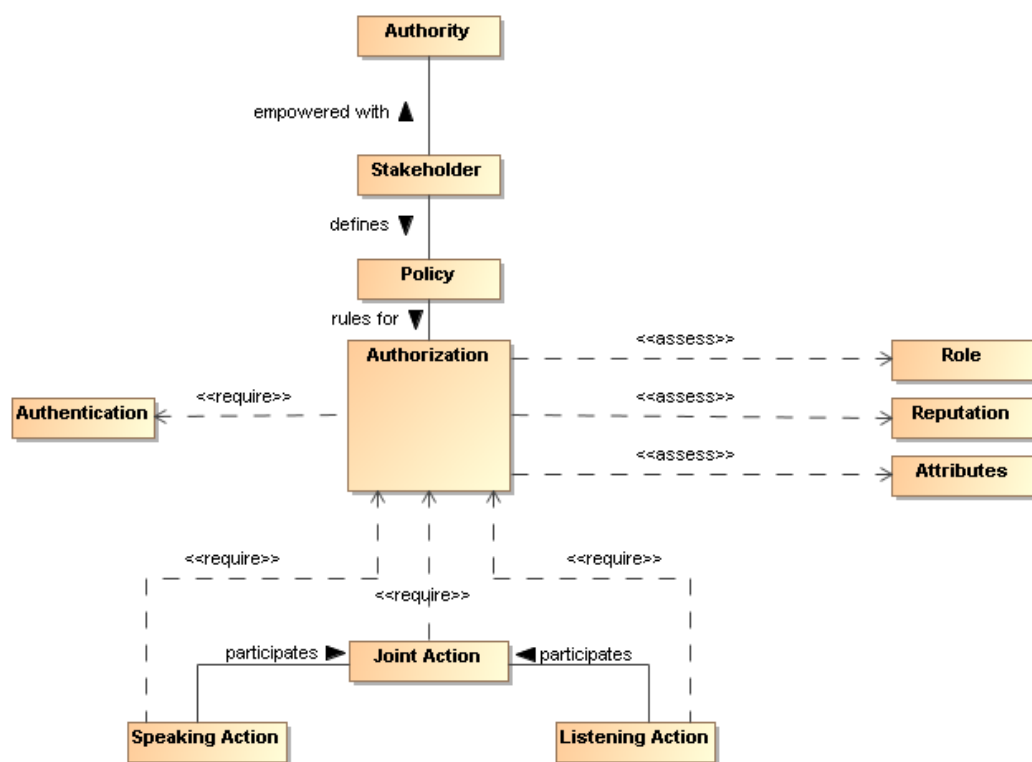
3028
3029 *Figure 43 - Authentication*

Comment [PFB97]:
 Issue 40
 New text reflects deletion of
 'identity' early in section 3.
 Figure must be updated
 Or, delete?

3030 **5.2.1.4 Authorization**

3031 Authorization concerns the legitimacy of the interaction. Authorization refers to the means by which a
 3032 stakeholder may be assured that the information and actions that are exchanged are either explicitly or
 3033 implicitly approved.

Comment [PFB98]: Issue 128 – speaking and listening actions need clarifying; remove both



3034
3035 *Figure 44 - Authorization*

3036 The roles and attributes which provide a participant's credentials are expanded to include reputation.
3037 Reputation often helps determine willingness to interact, for example, reviews of a service provider will
3038 influence the decision to interact with the service provider. The roles, reputation, and attributes are
3039 represented as assertions measured by authorization decision points.

3040 The role of policy for security is to permit stakeholders to express their choices. In Figure 44, a policy is a
3041 written constraint and the role, reputation, and attribute assertions are evaluated according to the
3042 constraints in the authorization policy. A combination of security mechanisms and their control via explicit
3043 policies can form the basis of an authorization solution.

3044 5.2.1.5 Non-repudiation

3045 Non-repudiation concerns the accountability of participants. To foster trust in the performance of a system
3046 used to conduct shared activities, it is important that the participants are not able to later deny their
3047 actions: to repudiate them. Non-repudiation refers to the means by which a participant may not, at a later
3048 time, successfully deny having participated in the interaction or having performed the actions as reported
3049 by other participants.

3050 5.2.1.6 Availability

3051 Availability concerns the ability of systems to use and offer the services for which they were designed.
3052 One of the threats against availability is the so-called denial of service attack in which attackers attempt to
3053 prevent legitimate access to the system.

3054 We differentiate here between general availability – which includes aspects such as systems reliability –
3055 and availability as a security concept where we need to respond to active threats to the system.

3056 5.2.2 Where SOA Security is Different

3057 The core security concepts are fundamental to all social interactions. The evolution of sharing information
3058 within a SOA ecosystem requires the flexibility to dynamically secure computing interactions where the
3059 owning social groups, roles, and authority are constantly changing as described in section 5.1.3.1.

3060 SOA policy-based security can be more adaptive than previous computing technologies, and typically
3061 involves a greater degree of distributed mechanisms.

3062 Standards for security, as is the case with all aspects of SOA implementation and use, play a large role in
3063 flexible security on a global scale. SOA security may also involve greater auditing and reporting to adhere
3064 to regulatory compliance established by governance structures.

3065 5.2.3 Security Threats

3066 There are a number of ways in which an attacker may attempt to compromise the security of a system.
3067 The two primary sources of attack are third parties attempting to subvert interactions between legitimate
3068 participants and an entity that is participating but attempting to subvert other participants. The latter is
3069 particularly important in a SOA **ecosystem** where there may be multiple ownership boundaries and trust
3070 boundaries.

Comment [PFB99]: Issue 135

3071 The threat model lists some common threats that relate to the core security concepts listed in Section
3072 5.2.1. Each technology choice in the realization of a SOA **based system** can potentially have many
3073 threats to consider.

Comment [PFB100]: Issue 135

3074 Message alteration

3075 If an attacker is able to modify the content (or even the order) of messages that are exchanged
3076 without the legitimate participants being aware of it then the attacker has successfully
3077 compromised the security of the system. In effect, the participants may unwittingly serve the
3078 needs of the attacker rather than their own.

3079 An attacker may not need to completely replace a message with his own to achieve his objective:
3080 replacing the identity of the **initially intended recipient** of a transaction may be enough.

Comment [PFB101]: Issue 136

3081 Message interception

3082 If an attacker is able to intercept and understand messages exchanged between participants,
3083 then the attacker may be able to gain advantage. This is probably the most commonly understood
3084 security threat.

3085 Man in the middle

3086 In a man-in-the-middle attack, the legitimate participants believe that they are interacting with
3087 each other; but are in fact interacting with the attacker. The attacker attempts to convince each
3088 participant that he is their correspondent; whereas in fact he is not.

3089 In a successful man-in-the-middle attack, legitimate participants do not have an accurate
3090 understanding of the state of the other participants. The attacker can use this to subvert the
3091 intentions of the participants.

3092 Spoofing

3093 In a spoofing attack, the attacker convinces a participant that he is really someone else –
3094 someone that the participant would normally trust.

3095 Denial of service attack

3096 In a denial of service (DoS) attack, the attacker attempts to prevent legitimate users from making
3097 use of the service. A DoS attack is easy to mount and can cause considerable harm: by
3098 preventing legitimate interactions, or by slowing them down enough, the attacker may be able to
3099 simultaneously prevent legitimate access to a service and to attack the service by another
3100 means.

3101 A variation of the DoS attack is the Distributed Denial of Service (DDoS) attack. In a DDoS attack
3102 the attacker uses multiple agents to the attack the target. In some circumstances this can be
3103 extremely difficult to counteract effectively.

3104 One of the features of a DoS attack is that it does not require valid interactions to be effective:
3105 responding to invalid messages also takes resources and that may be sufficient to cripple the
3106 target.

3107 **Replay attack**

3108 In a replay attack, the attacker captures the message traffic during a legitimate interaction and
3109 then replays part of it to the target. The target is persuaded that a similar transaction to the
3110 previous one is being repeated and it responds as though it were a legitimate interaction.

3111 A replay attack may not require that the attacker understand any ~~of the~~ individual
3112 ~~message communications~~; the attacker may have different objectives (for example attempting to
3113 predict how the target would react to a particular request).

3114 **False repudiation**

3115 In false repudiation, a user completes a normal transaction and then later attempts to deny that
3116 the transaction occurred. For example, a customer may use a service to buy a book using a credit
3117 card; then, when the book is delivered, refuse to pay the credit card bill claiming that *someone*
3118 *else* must have ordered the book.

3119 **5.2.4 Security Responses**

3120 Security goals are never absolute: it is not possible to guarantee 100% confidentiality, non-repudiation,
3121 etc. However, a well designed and implemented security response model can ensure acceptable levels of
3122 security risk. For example, using a well-designed cipher to encrypt messages may make the cost of
3123 breaking communications so great and so lengthy that the information obtained is valueless.

3124 Performing threat assessments, devising mitigation strategies, and determining acceptable levels of risk
3125 are the foundation for an effective process to mitigating threats in a cost-effective way.¹⁰ The choice in
3126 hardware and software to realize a SOA ~~implementation~~ will be a basis for threat assessments and
3127 mitigation strategies. The stakeholders of a specific SOA implementation should determine acceptable
3128 levels of risk based on threat assessments and the cost of mitigating those threats.

Comment [PFB102]: Issue 141

3129 **5.2.4.1 Privacy Enforcement**

3130 The most efficient mechanism to assure confidentiality is the encryption of information. Encryption is
3131 particularly important when messages must cross trust boundaries; especially over the Internet. Note that
3132 encryption need not be limited to the content of messages: it is possible to obscure even the existence of
3133 messages themselves through encryption and 'white noise' generation in the communications channel.

3134 The specifics of encryption are beyond the scope of this architecture. However, we are concerned about
3135 how the connection between privacy-related policies and their enforcement is made.

3136 A policy enforcement point for enforcing privacy may take the form of an automatic function to encrypt
3137 messages as they leave a trust boundary; or perhaps simply ensuring that such messages are suitably
3138 encrypted.

3139 Any policies relating to the level of encryption being used would then apply to these centralized
3140 messaging functions.

3141 **5.2.4.2 Integrity Protection**

3142 To protect against message tampering or inadvertent message alteration, and to allow the receiver of a
3143 message to authenticate the sender, messages may be accompanied by a digital signature. Digital

¹⁰ 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 ecosystem for this stakeholder.

3144 signatures provide a means to detect if signed data has been altered. This protection can also extend to
3145 authentication and non-repudiation of a sender.

3146 A common way a digital signature is generated is with the use of a private key that is associated with a
3147 public key and a digital certificate. The private key of some entity in the system is used to create a digital
3148 signature for some set of data. Other entities in the system can check the integrity of the signed data set
3149 via signature verification algorithms. Any changes to the data that was signed will cause signature
3150 verification to fail, which indicates that integrity of the data set has been compromised.

3151 A party verifying a digital signature must have access to the public key that corresponds to the private key
3152 used to generate the signature. A digital certificate contains the public key of the owner, and is itself
3153 protected by a digital signature created using the private key of the issuing Certificate Authority (CA).

3154 5.2.4.3 Message Replay Protection

3155 To protect against replay attacks, messages may contain information that can be used to detect replayed
3156 messages. The simplest requirement to prevent replay attacks is that each message that is ever sent is
3157 unique. For example, a message may contain a message ID, a timestamp, and the intended destination.

3158 By storing message IDs, and comparing each new message with the store, it becomes possible to verify
3159 whether a given message has been received before (and therefore should be discarded).

3160 The timestamp may be included in the message to help check for message freshness. Messages that
3161 arrive after their message ID could have been cleared (after receiving the same message some time
3162 previously) may also have been replayed. A common means for representing timestamps is a useful part
3163 of an interoperable replay detection mechanism.

3164 The destination information is used to determine if the message was misdirected or replayed. If the
3165 replayed message is sent to a different endpoint than the destination of the original message, the replay
3166 could go undetected if the message does not contain information about the intended destination.

3167 In the case of messages that are replies to prior messages, it is also possible to include seed information
3168 in the prior messages that is randomly and uniquely generated for each message that is sent out. A
3169 replay attack can then be detected if the reply does not embed the random number that corresponds to
3170 the original message.

3171 5.2.4.4 Auditing and Logging

3172 False repudiation involves a participant denying that it authorized a previous interaction. An effective
3173 strategy for responding to such a denial is to maintain careful and complete logs of interactions that can
3174 be used for auditing purposes. The more detailed and comprehensive an audit trail is, the less likely it is
3175 that a false repudiation would be successful.

3176 The countermeasures assume that the non-repudiation tactic (e.g. digital signatures) is not undermined
3177 itself. For example, if private key is stolen and used by an adversary, even extensive logging cannot
3178 assist in rejecting a false repudiation.

3179 Unlike many of the security responses discussed here, it is likely that the scope for automation in
3180 rejecting a repudiation attempt is limited [in the immediate future](#) to careful logging.

Comment [KJL103]: Issue 275

3181 5.2.4.5 Graduated engagement

3182 The key to managing and responding to DoS attacks is to be careful in the use of resources when
3183 responding to interaction. Put simply, a system has a choice to respond to a communication or to ignore
3184 it. In order to avoid vulnerability to DoS attacks a service provider should be careful not to commit
3185 resources beyond those implied by the current state of interactions; this permits a graduation in
3186 commitment by the service provider that mirrors any commitment on the part of service consumers and
3187 attackers alike.

3188 5.2.5 Architectural Implications of SOA Security

3189 Providing SOA security in an ecosystem of governed services has the following implications on the policy
3190 support and the distributed nature of mechanisms used to assure SOA security:

- 3191
- 3192
- 3193
- 3194
- 3195
- 3196
- 3197
- 3198
- 3199
- 3200
- 3201
- 3202
- 3203
- 3204
- 3205
- 3206
- 3207
- 3208
- 3209
- 3210
- 3211
- 3212
- 3213
- Security expressed through policies has the same architectural implications as described in Section 4.4.3 for policies and contracts architectural implications.
 - Security policies require mechanisms to support security description administration, storage, and distribution.
 - Service descriptions supporting security policies should:
 - have a meta-structure sufficiently rich to support security policies;
 - be able to reference one or more security policy artifacts;
 - have a framework for resolving conflicts between security policies.
 - The mechanisms that make-up the execution context in secure SOA-based systems should:
 - provide protection of the confidentiality and integrity of message exchanges;
 - be distributed so as to provide centralized or decentralized policy-based identification, authentication, and authorization;
 - ensure service availability to consumers;
 - be able to scale to support security for a growing ecosystem of services;
 - be able to support security between different communication technologies;
 - Common security services include:
 - services that abstract encryption techniques;
 - services for auditing and logging interactions and security violations;
 - services for identification;
 - services for authentication;
 - services for authorization;
 - services for intrusion detection and prevention;
 - services for availability including support for quality of service specifications and metrics.

3214 5.3 Management Model

3215 5.3.1 Management

3216 Management is a process of controlling resources in accordance with the policies and principles defined
3217 by Governance.

3218 There are three separate but linked domains of interest within the management of a SOA ecosystem:

- 3219
- 3220
- 3221
- 3222
- 3223
- 3224
1. the management and support of the resources that are involved in any complex structures – of which SOA ecosystems are excellent examples;
 2. the promulgation and enforcement of the policies and service contracts agreed to by the stakeholders in the SOA ecosystem;
 3. the management of the relationships of the participants – both to each other and to the services that they use and offer.

3225 There are many artifacts related to management. Historically, systems management capabilities have
3226 been organized by the “FCAPS” functions (based on ITU-T Rec. M.3400 (02/2000), “TMN Management
3227 Functions”):

- 3228
- 3229
- 3230
- 3231
- fault management,
 - configuration management,
 - account management,
 - performance and security management.

3232 The primary task of the functional groups is to concentrate on maintaining systems in a trusted, active,
3233 and accessible state.

3234 In the context of the SOA ecosystem, we see many possible resources that may require management
3235 such as services, service descriptions, service contracts, policies, roles, relationships, security, people
3236 and systems that implement services and infrastructure elements. In addition, given the ecosystem
3237 nature, it is also potentially necessary to manage the business relationships between participants.

3238 Successful operation of a SOA ecosystem requires trust among the stakeholders and between them and
3239 the SOA-based system elements. In contrast, regular systems in technology are not necessarily operated
3240 or used in an environment requiring trust before the stakeholders make use of the system. Indeed, many
3241 of these systems exist in hierarchical management structures, within which use may be mandated by

3242 legal requirement, executive decision, or good business practice in furthering the business' strategy. The
3243 pre-condition of trust in the SOA ecosystem is rooted both in the principles of service orientation and in
3244 the distributed, authoritative ownership of independent services. Even for hierarchical management
3245 structures applied to a SOA ecosystem, the service in use should have a contractual basis rather than
3246 solely being mandated.

3247 Trust may be established through agreements/contracts, policies, or implicitly through observation of
3248 repeated interactions with others. Explicit trust is usually accompanied by formalized documents suitable
3249 for management. Implicit trust adds fragility to the management of a SOA ecosystem because failure to
3250 maintain consistent and predictable interactions will undermine the trust between participants and within
3251 the ecosystem as a whole.

3252 Management in a SOA ecosystem is thus concerned with management taking actions that will establish
3253 the condition of trust that must be present before engaging in service interactions. These concerns should
3254 largely be handled within the governance of the ecosystem. The policies, agreements, and practices
3255 defined through governance provide the boundaries within which management operates and for which
3256 management must provide enforcement and feedback. However, governance alone cannot foresee all
3257 circumstances but must offer sufficient guidance where agreement between all stakeholders cannot be
3258 reached. Management in these cases must be flexible and adaptable to handle unanticipated conditions
3259 without unnecessarily breaking trust relationships.

3260 Service management is the process – manual, automated, or a combination – of proactively monitoring
3261 and controlling the behavior of a service or a set of services. Service management operates under
3262 constraints attributed to the business and social context. Specific policies may be used to govern cross-
3263 boundary relationships. Managing solutions based on such policies (and that may be used across
3264 ownership boundaries) raises issues that are not typically present when managing a service within a
3265 single ownership domain. Care is therefore required in managing a service when the owner of the
3266 service, the provider of the service, the host of the service and mediators to the service may all belong to
3267 different stakeholders.

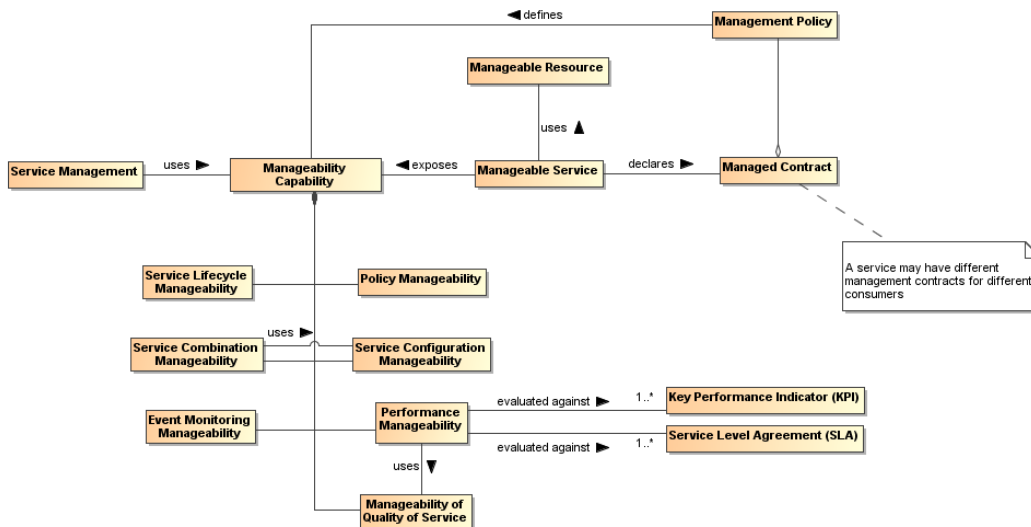
3268 Cross-boundary service management takes place in, at least, the following situations:

- 3269 • using combinations of services that belong to different ownership domains
- 3270 • using of services that mediate between ownership domains
- 3271 • sharing monitoring and reporting means and results.

3272 These situations are particularly important in ecosystems that are highly decentralized, in which the
3273 participants interact as peers as well as in the “master-servant” mode.

3274 The management model shown in Figure 45 conveys how the SOA paradigm applies to managing
3275 services. Services management operates via service metadata, such as properties associated with
3276 service lifecycles and with service use, which are typically collected in or accessed through the service
3277 description.

3278



A service may have different management contracts for different consumers

3279
3280 *Figure 45 - Management model in SOA ecosystem*

3281 The service metadata of interest is that set of service properties that is manageable. These manageability
3282 properties are generally identifiable for any service consumed or supplied within the ecosystem. The
3283 necessary existence of these properties within the SOA ecosystem motivates the following definitions:

3284 **Manageability**

3285 Manageability of a resource is the capability that allows it to be controlled, monitored, and
3286 reported on with respect to some properties.

3287 **Manageability property**

3288 Manageability property is the property used in the manageability of a resource. Manageability
3289 property is the fundamental unit of management in systems management.

3290 Note that manageability is not necessarily a part of the managed entities themselves and are generally
3291 considered to be external to the managed entities.

3292 Each resource may be managed through a number of aspects of management, and the resources may
3293 be grouped based on similar aspects. For example, resources may be grouped according to the aspect
3294 referred to as “Configuration Manageability” for the collection of services. Some resources may not be
3295 managed under a particular capability if there are no manageability aspects with a clear meaning or use.
3296 As an example, all resources within a SOA ecosystem have a lifecycle that is meaningful within the
3297 ecosystem. Thus, all resources are manageable under Lifecycle Manageability. In contrast, not all
3298 resources report or handle events. Thus, Event Manageability is only concerned with those resources for
3299 which events are meaningful.

3300 **Life-cycle Manageability** of a service typically refers to how the service is created, how it is **destroyed**
3301 **retired** and how service versions must be managed. This manageability is a feature of the SOA
3302 ecosystem because the service cannot manage its own life cycle. **Related properties may include the**
3303 **necessary state of the ecosystem for the creation and retirement of the service and the state of the**
3304 **ecosystem following the retirement of the service. The SOA ecosystem distinguishes between service**
3305 **composition and service aggregation: retiring of service composition leads to retiring of all services**
3306 **comprising the composition while retiring of service aggregation assumes that comprising services have**
3307 **their own life-cycle and can be used in another aggregation.**

Comment [KJL104]: Issue 146

Comment [PFB105]: Issue 145

3308 Another important consideration is that services may have resource requirements, **such as concurrent**
3309 **connectivity to a data source, which** must be established at various points in the services' life cycles.
3310 However, actual providers of these resources may not be known at the time of the service creation and,
3311 thus, have to be managed at service run-time.

Comment [KJL106]: Issue 147

3312 **Combination Manageability** of a service addresses management of service characteristics that allow for
3313 creating and changing combinations in which the service participates or that the service combines itself.
3314 Known models of such combinations are aggregations and compositions. Examples of patterns of
3315 combinations are choreography and orchestration. In cases of business collaboration, combination of
3316 services appears as cooperation of services. Combination Manageability drives implementation of the
3317 Service Composability Principle of service orientation.

Comment [PFB107]: Issue 145

3318 Service combination manageability resonates with the methodology of process management.
3319 Combination Manageability may be applied at different phases of service creation and execution and, in
3320 some cases, can utilize Configuration Manageability.

3321 Service combinations typically contribute the most in delivering business values to the stakeholders.
3322 Managing service combinations is the one of the most important tasks and features of the SOA
3323 ecosystem.

3324 **Configuration Manageability** of a service allows managing the identity of and the interactions among
3325 internal elements of the service. for example, a use of data encryption for internal inter-component
3326 communication in particular deployment conditions. Also, Configuration Manageability correlates with the
3327 management of service versions and configuration of the deployment of new services into the ecosystem.
3328 Configuration Management differs from the Combination Manageability in the scope and scale of
3329 manageability, and addresses lower level concerns than the architectural combination of services.

Comment [PFB108]: Issue 145

3330 **Event Monitoring Manageability** allows managing the categories of events of interest related to services
3331 and reporting recognized events to the interested stakeholders. Such events may be the ones that trigger
3332 service invocations as well as execution of particular functionality provided by the service. For example,
3333 an execution of a set of financial market risk services, which implements choreography pattern, may be
3334 started if certain financial event occurs in a stock exchange.

Comment [PFB109]: Issue 145

3335 Event Monitoring Manageability is a key lower-level manageability aspect, in which the service provider
3336 and associated stakeholders are interested. Monitored events may be internal or external to the SOA
3337 ecosystem. For example, a disaster in the oil industry, which is outside the SOA ecosystem of the Insurer,
3338 can trigger the service's functionality that is responsible for immediate or constant monitoring of oil prices
3339 in the oil trading exchanges and, respectively, modify the premium paid by the insured oil companies.

3340 **Performance Manageability** of a service allows controlling the service results, shared and sharable real
3341 world effects against the business goals and objectives of the service. This manageability assumes
3342 monitoring of the business performance as well as the management of this monitoring itself. Performance
3343 Manageability includes business and technical performance manageability through a performance criteria
3344 set, such as business key performance indicators (KPI) and service-level agreements (SLA).

3345 The performance business- and technical-level characteristics of the service should be known from the
3346 service contract. The service provider and consumer must be able to monitor and measure these
3347 characteristics or be informed about the results measured by a third party. An example of such monitoring
3348 would be when the comparison of service performance results against an SLA is not satisfactory to the
3349 consumer, and as a consequence, the consumer# may replace the service by a service from a
3350 competitor.

Comment [PFB110]: Issue 145

3351 Performance Manageability is the instrument for providing compliance of the service with its service
3352 contracts. Performance Manageability utilizes Manageability of Quality of Service.

3353 **Manageability of Quality of Service** deals with management of service non-functional characteristics
3354 that may be of significant value to the service consumers and other stakeholders in the SOA ecosystem.
3355 A classic example- of this is managing bandwidth offerings associated with a service.

3356 Manageability of quality of service assumes that the properties associated with service qualities are
3357 monitored during the service execution. Results of monitoring may be challenged-compared against an
3358 SLA or a KPI, which results in the continuous validation of how the service contract is preserved by the
3359 service provider.

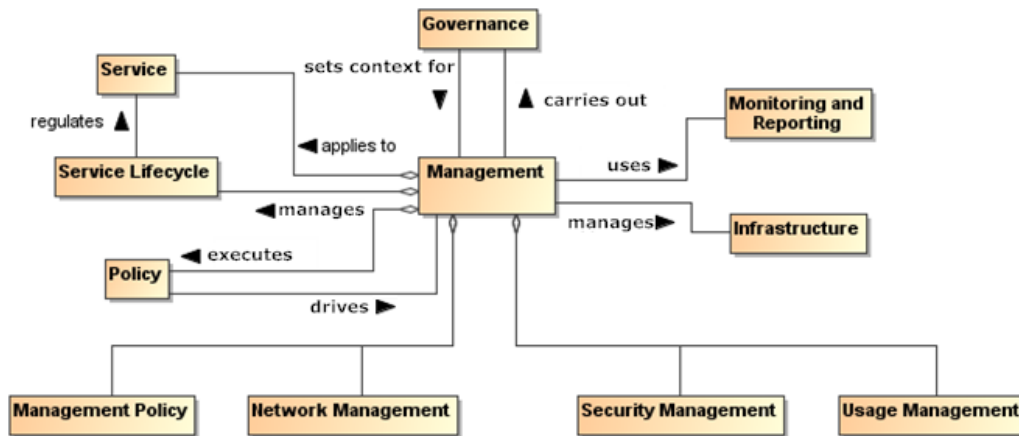
3360 **Policy Manageability** allows additions, changes and replacements of the policies associated with a
3361 resource in the SOA ecosystem. The ability to manage those policies (such as promulgating policies,
3362 retiring policies and ensuring that policy decision points and enforcement points are current) enables the
3363 ecosystem to apply policies and *evaluate* the results.

3364 The capability to manage, i.e. use a particular manageability, requires policies from governance to be
 3365 translated into detailed rules and regulations which are measured and monitored providing corresponding
 3366 feedback for enforcement. At the same time, the execution of a management capability MUST adhere to
 3367 certain policies governing the management itself. For example, a management has to enforce and control
 3368 policies of compliance with particular industry regulation while the management is obliged by another
 3369 policy to report on the compliance status periodically.

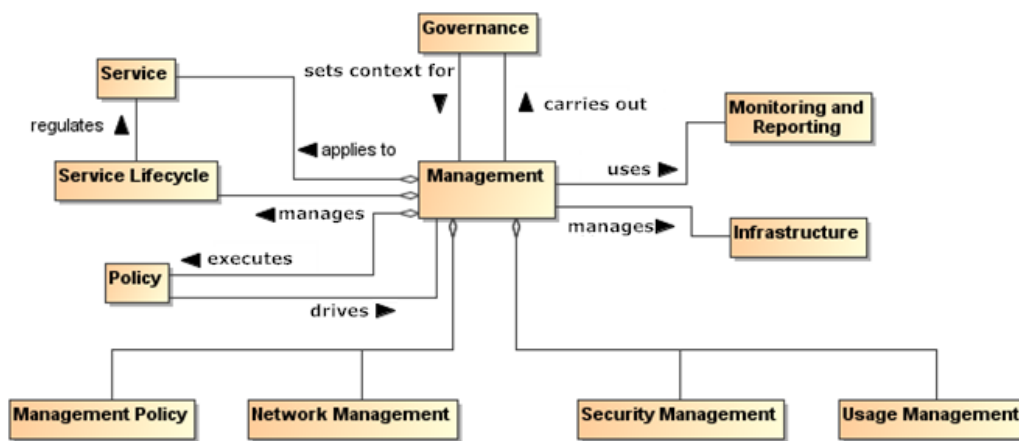
3370 Management of SOA ecosystem recognizes the manageability challenge and requires manageability
 3371 properties to be considered for all aforementioned manageability cases. In the following sub-sections, we
 3372 describe how these properties are used in the management as well as some relationships between
 3373 management and other components of SOA ecosystem.

3374 5.3.2 Management Means and Relationships

3375 A minimal set of management issues for the SOA ecosystem is shown in Figure
 3376 46



3377 Figure 46 - Management Means and Relationships in a SOA ecosystem
 3378 and elaborated in the following sections.
 3379



Comment [PFB111]: Issue 152 – diagram updated

3380 Figure 46 - Management Means and Relationships in a SOA ecosystem
 3381

3382 5.3.2.1 Management Policy

3383 The management of resources within the SOA ecosystem may be governed by management policies. In
3384 a deployed SOA-based solution, it may well be that different aspects of the management of a given
3385 service are managed by different management services. For example, the life-cycle management of
3386 services often involves managing service versions. Managing quality of service is often very specific to
3387 the service itself; for example, quality of service attributes for a video streaming service are quite different
3388 to those for a banking system.

3389 ~~Additional concepts of management also apply to IT management.~~

3390 5.3.2.2 Network Management

3391 Network management deals with the maintenance and administration of large scale physical networks
3392 such as computer networks and telecommunication networks. Specifics of the networks may affect
3393 service interactions from performance and operational perspectives.

3394 Network and related system management execute a set of functions required for controlling, planning,
3395 deploying, coordinating, and monitoring the distributed services in the SOA ecosystem. However, while
3396 recognizing their importance, the specifics of systems management or network management are out of
3397 scope for this Reference Architecture Foundation.

Comment [KJL112]: Issue 153

3398 5.3.2.3 Security Management

3399 Security Management includes identification of roles, permissions, access rights, and policy attributes
3400 defining security boundaries and events that may trigger a security response.

3401 Security management within a SOA ecosystem is essential to maintaining the trust relationships between
3402 participants residing in different ownership domains. Security management must consider not just the
3403 internal properties related to interactions between participants but ecosystem properties that preserve the
3404 integrity of the ecosystem from external threats.

3405 5.3.2.4 Usage Management

3406 Usage Management is concerned with how resources are used, including:

- 3407 • how the resource is accessed, who is using the resource, and the state of the resource (access
3408 properties);
 - 3409 • controlling or shaping demand for resources to optimize the overall operation of the ecosystem
3410 (demand properties);
 - 3411 • assigning costs to the use of resources and distributing those cost assignments to the
3412 participants in an appropriate manner (financial properties).
- 3413

3414 5.3.3 Management and Governance

3415 The primary role of governance in the context of a SOA ecosystem is to foster an atmosphere of
3416 predictability, trust, and efficiency, and it accomplishes this by allowing the stakeholders to negotiate and
3417 set the key policies that govern the running of the SOA-based solution. Recall that in an ecosystem
3418 perspective, the goal of governance is less to have complete fine-grained control but more to enable the
3419 individual participants to work together.

3420 Policies for a SOA ecosystem will tend to focus on the rules of engagement between participants; for
3421 example, what kinds of interactions are permissible, how disputes are resolved, etc. While governance
3422 may primarily focus on setting policies, management will focus on the realization and enforcement of
3423 policies. Effective management in the SOA ecosystem requires an ability for governance to understand
3424 the consequences of its policies, guidelines, and principles, and to adjust those as needed when
3425 inconsistencies or ambiguity become evident from the operation of the management functions. This
3426 understanding and adjustment must be facilitated by the results of management and so the mechanisms
3427 for providing feedback from management into governance must exist.

3428 Governance operates via specialized activities and, thus, should be managed itself. Governance policies
3429 are included in the Governance Framework and Processes, and driven by the enterprise business model,
3430 business objectives and strategies. Where corporate management policies exist, these are usually guided
3431 and directed by the corporate executives. In peer relationships, governance policies are set by either an
3432 external entity and accepted by the peers or by the peers themselves. This creates the appropriate
3433 authoritative level for the policies used for the management of the Governance Framework and
3434 Processes. Management to operationalize governance controls the life-cycle of the governing policies,
3435 including procedures and processes, for modifying the Governance Framework and Processes.

3436 **5.3.4 Management and Contracts**

3437 **5.3.4.1 Management for Contracts and Policies**

3438 As we noted above, management can often be viewed as the application of contracts and individual
3439 policies to ensure the smooth running of the SOA ecosystem. Policies and service contracts specify the
3440 service characteristics that have to be monitored, analyzed and managed. These also play an important
3441 role as the guiding constraints for management, as well as being artifacts (e.g., policy and contractual
3442 documents) that also need to be managed.

3443 **5.3.4.2 Contracts**

3444 As described in sections “Participation in a SOA Ecosystem view” and “Realization of a SOA Ecosystem
3445 view”, there are several types of contractual information in the SOA ecosystem. From the management
3446 perspective, three basic types of the contractual information relate to:

- 3447 • relationship between service provider and consumer;
- 3448 • communication with the service;
- 3449 • control of the quality of the service execution.

3450 When a consumer prepares to interact with a service, the consumer and the service provider must come
3451 to an agreement on the service features and characteristics that will be provided by the service and made
3452 available to the consumer. This agreement is known as a service contract.

3453 **Service Contract**

3454 An implicit or explicit documented agreement between the service consumer and service provider
3455 about the use of the service based on

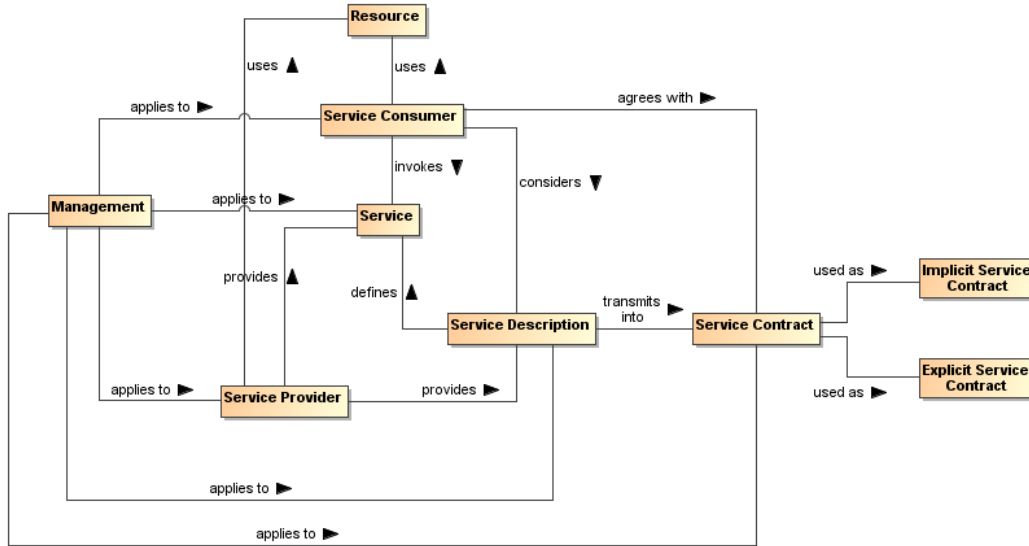
- 3456 • the commitment by a service provider to provide service functionality and results consistent
3457 with identified real world effects and
- 3458 • the commitment by a service consumer to interact with the service per specific means and
3459 per specified policies,

3460 where both consumer and provider actions are in the manner described in the service description.

3461 The service description provides the basis for the service contract and, in some situations, may be used
3462 as an implicit default service contract. In addition, the service description may set mandatory aspects of a
3463 service contract, e.g. for security services, or may specify acceptable alternatives. As an example of
3464 alternatives, the service description may identify which versions of a vocabulary will be recognized, and
3465 the specifics of the contract are satisfied when the consumer uses one of the alternatives. Another
3466 alternative could have a consumer identify a policy they require be satisfied, e.g. a standard privacy policy
3467 on handling personal information, and a provider that is prepared to accept a policy request would
3468 indicate acceptance as part of the service contract by continuing with the interaction. In each of these
3469 cases, the actions of the participants are consistent with an implicit service contract without the existence
3470 of a formal agreement between the participants.

3471 In the case of business services, it is anticipated that the service contract may take an explicit form and
3472 the agreement between business consumer and business service provider is formalized. Formalization
3473 requires up-front interactions between service consumer and service provider. In many business
3474 interactions, especially between business organizations within or across corporate boundaries, a
3475 consumer needs a contractual assurance from the provider or wants to explicitly indicate choices among
3476 alternatives, e.g., only use a subset of the business functionality offered by the service and pay a prorated
3477 cost.

3478



3479
3480 *Figure 47 - Management of the service interaction*

3481 Consequently, an implicit service contract is an agreement (1) on the consumer side with the terms,
3482 conditions, features and interaction means specified in the service description "as is" or (2) a selection
3483 from alternatives that are made available through mechanisms included in the service description, and
3484 neither of these require any a priori interactions between the service consumer and the service provider.
3485 For example, a browser interface may display a checked box indicating the consumer agrees to accept
3486 future advertisement; the consumer can uncheck the box to indicate advertisements should not be sent.

Comment [KJL113]: Issue 158

3487 An explicit service contract always requires a form of interaction between the service consumer and the
3488 service provider prior to the service invocation. This interaction may regard the choice or selection of the
3489 subset of the elements of the service description or other alternatives introduced through the formal
3490 agreement process that would be applicable to the interaction with the service and affect related joint
3491 action.

3492 Any form of explicit contract couples the service consumer and provider. While explicit contracts may be
3493 necessary or desirable in some cases, such as in supply chain management, commerce often uses a mix
3494 of implicit and explicit contracts, and a service provider may offer (via service description) a conditional
3495 shift from implicit to explicit contract. For example, Twitter offers an implicit contract on the use of its APIs
3496 to any application with the limit on the amount of service invocations; if the application needs to use more
3497 invocations, one has to enter into the explicit fee-based contract with the provider. A case where an
3498 implicit contract transforms into an explicit contract may be illustrated when one buys a new computer and
3499 it does not work. The buyer returns the computer for repair under the manufacturer's warranty as stated
3500 by an implicit purchase contract. However, if the repair does not fix the problem and the seller offers an
3501 upgraded model in replacement, the buyer may agree to an explicit contract that limits the rights of the
3502 buyer to make the explicit agreement public.

3503 Control of the quality of the service execution, often represented as a service level agreement (SLA), is
3504 performed by service monitoring systems and includes both technical and operational business controls.
3505 SLA is a part of the service contract and, because of the individual nature of such contracts, may vary
3506 from one service contract to another, even for the same consumer. Typically, a particular SLA in the
3507 service contract is a concrete instance of the SLA declared in the service description.

3508 Management of the service contracts is based on management policies that may be mentioned in the
3509 service description and in the service contracts. Management of the service contracts is mandatory for
3510 consumer relationship management. In the case of explicit service contracts, the contracts have to be
3511 created, stored, maintained, reviewed/controlled and archived/destroyed as needed. All the activities are

3512 management concerns. Explicit service contracts may be stored in specialized repositories that provide
3513 appropriate level of security.

3514 Management of the service interfaces is based on several management policies that regulate

- 3515 • availability of interfaces specified in the service contracts,
- 3516 • accessibility of interfaces,
- 3517 • procedures for interface changes,
- 3518 • interface versions as well as the versions of all parts of the interfaces,
- 3519 • traceability of the interfaces and their versions back to the service description document.

3520 Management of the SLA is integral to the management of service monitoring and operational service
3521 behavior at run-time. An SLA usually enumerates service characteristics and expected performances of
3522 the service. Since an SLA carries the connotation of a "promise", monitoring is needed to know if the
3523 promise is being kept. Existence of an SLA itself does not guarantee that the consumer will be provided
3524 with the service level specified in the service contract.

3525 The use of an SLA in a SOA ecosystem can be wider than just an agreement on technical performances.
3526 An SLA may contain remedies for situations where the promised service cannot be maintained, or the
3527 real world effect cannot be achieved due to developments subsequent to the agreement. A service
3528 consumer that acts accordingly to realize the real world effect may be compensated for the breach of the
3529 SLA if the effect is not realized.

3530 Management of the SLA includes, among others, policies to change, update, and replace the SLA. This
3531 aspect concerns service Execution Context because the business logic associated with a defined
3532 interface may differ in different Execution Contexts and affect the overall performance of the service.

3533 5.3.4.3 Policies

3534 "Although provision of management capabilities enables a service to become manageable, the extent and
3535 degree of permissible management are defined in management policies that are associated with the
3536 services. Management policies are used to define the obligations for, and permissions to, managing the
3537 service" [WSA]. Management policies, in essence, are the realization of governing rules and regulations.
3538 As such, some management policies may target services while other policies may target the management
3539 of the services.

3540 In practice, a policy without any means of enforcing it is vacuous. In the case of management policy, we
3541 rely on a management infrastructure to realize and enforce management policy.

3542 5.3.4.4 Service Description and Management

3543 The service description identifies several management objects such as a set of service interfaces and
3544 related set of SLAs. Service behavioral characteristics and performances specified in the SLA depend on
3545 the interface type and its Execution Context. In the service description, a service consumer can find
3546 references to management policies, SLA metrics, and the means of accessing measured values that
3547 together increase assurance in the service quality. At the same time, service description is an artifact that
3548 needs to be managed.

3549 In the SOA ecosystem, the service description is the assembled information that describes the service but
3550 it may be reported or displayed in different presentations. While each separate version of the service has
3551 one and only one service description, different categories of service consumers may focus their interests
3552 on different aspects of the service description. Thus, the same service description may be displayed not
3553 only in different languages but also with different cultural and professional accents in the content.

3554 New service description may be issued to reflect changes and update in the service. If the change in the
3555 service does not affect its service description, the new service version may have the same service
3556 description as the previous version except for the updated version identifier. For example, a service
3557 description may stay the same if bugs were fixed in the service. However, if a change in the service
3558 influences any aspects of the service quality that can affect the real world effect resulting from
3559 interactions with the service, the service description MUST reflect this change even if there are no
3560 changes to the service interface.

3561 Management of the service description as well as of the explicit service contracts is essential for delivery
3562 of the service to the consumer satisfaction. This management can also prevent business problems rooted
3563 in poor communication between the service consumers and the service providers.

3564 Thus, management of service description contains, among others, management of the service description
3565 presentations, the life-cycles of the service descriptions, service description distribution practices and
3566 storage of the service descriptions and related service contracts. Collections of service descriptions in the
3567 enterprise may manifest a need for specialized registries and/or repositories. Depending on the enterprise
3568 policies, an allocation of purposes and duties of registries and repositories may vary but this topic is
3569 beyond the current scope.

3570 **5.3.5 Management for Monitoring and Reporting**

3571 The successful application of management relies on the monitoring and reporting aspects of management
3572 to enable the control aspect. Monitoring in the context of management consists of measuring values of
3573 managed aspects and evaluating that measurement in relationship to some expectation. Monitoring in a
3574 SOA ecosystem is enabled through the use of mechanisms by resources for exposing managed aspects.
3575 In the SOA framework, this mechanism may be a service for obtaining the measurement. Alternatively,
3576 the measurement may be monitored by means of event generation containing updated values of the
3577 managed aspect.

3578 Approaches to monitoring may use a polling strategy in which the measurements are requested from
3579 resources in periodic intervals, in a pull strategy in which the measurements are requested from
3580 resources at random times, or in a push strategy in which the measurements are supplied by the resource
3581 without request. The push strategy can be used in a periodic update approach or in an "update on
3582 change" approach. Management services must be capable of handling these different approaches to
3583 monitoring.

3584 Reporting is the complement to monitoring. Where monitoring is responsible for obtaining measurements,
3585 reporting is responsible for distributing those measurements to interested stakeholders. The separation
3586 between monitoring and reporting is made to include the possibility that data obtained through monitoring
3587 might not be used until an event impacting the ecosystem occurs or the measurement requires further
3588 processing to be useful. In the SOA framework, reporting is provided using services for requesting
3589 measurement reports. These reports may consist of raw measurement data, formatted collections of data,
3590 or the results of analysis performed on measurement data from collections of different managed aspects.
3591 Reporting is also used to support logging and auditing capabilities, where the reporting mechanisms
3592 create log or audit entries.

3593 **5.3.6 Management for Infrastructure**

3594 All of the properties, policies, interactions, resources, and management are only possible if a SOA
3595 ecosystem infrastructure provides support for managed capabilities. Each managed capability imposes
3596 different requirements on the capabilities supplied by the infrastructure in SOA ecosystem and requires
3597 that those capabilities be usable as services or at the very least be interoperable.

3598 While not providing a full list of infrastructural elements of a SOA ecosystem, we list some examples here:

- 3599 1. Registries and repositories for services, policies, and related descriptions and contracts
- 3600 2. Synchronous and asynchronous communication channels for service interactions (e.g., network,
3601 e-mail, message routing with ability of mediating transport protocols, etc.)
- 3602 3. Recovery capabilities
- 3603 4. Security controls

3604 A SOA ecosystem infrastructure, enabling service management, should also support:

- 3605 1. Management enforcement and control means
- 3606 2. Monitoring and SLA validation controls
- 3607 3. Testing and Reporting capabilities

3608 | The combination of manageability [properties, related](#) capabilities and infrastructure elements constitutes
3609 a certain level of SOA management maturity. While several maturity models exist, this topic is out of the
3610 scope of the current document.

3611 **5.3.7 Architectural Implication of the SOA Management**

3612 SOA Management is one of the fundamental elements of the SOA ecosystem; it impacts all aspects of a
3613 service life-cycle, service activities and actions, and a service usage. The key choices that must be made
3614 centre in management means, methods and manageability properties:

- 3615 • Every resource of the SOA ecosystem and, particularly, services MUST provide manageability
3616 properties
 - 3617 ○ The set of manageability properties SHOULD include as minimum such properties as life-
3618 cycle, combination, configuration, event monitoring, performance, quality of services, and
3619 policy manageability
 - 3620 ○ Combinations of manageability properties MAY be used in different management
3621 methods and tools
- 3622 • Manageability properties and applicable policies SHOULD be appropriately described in the
3623 services description and contracts
- 3624 • Management processes SHOULD operate (control, enforce and provide a feedback to the
3625 governance) via policies, agreements/contracts, and practices defined through governance
- 3626 • Management functions and information MAY be realized as services and, thus, MUST be
3627 managed itself
- 3628 • Management in the cases, where sufficient guidance is unavailable or for which agreement
3629 between all stakeholders cannot be reached, MUST be flexible and adaptable to handle
3630 unanticipated conditions without unnecessarily breaking trust relationships
- 3631 • Management SHOULD engage a monitoring mechanism to enable manageability. Monitoring
3632 HAS to include
 - 3633 ○ Access mechanisms to collected SLA metrics
 - 3634 ○ Assessment mechanisms to compare metrics against policies and contracts
- 3635 • Results of monitoring and reporting MUST be made accessible to participants in different
3636 ownership domains.

Comment [PFB114]: Issue 160

3637 **5.4 SOA Testing Model**

3638 Testing for SOA combines the typical challenges of software testing and certification with the additional
3639 needs of accommodating the distributed nature of the resources, the greater access of a more
3640 unbounded consumer population, and the desired flexibility to create new solutions from existing
3641 components over which the solution developer has little if any control. The purpose of testing is to
3642 demonstrate a required level of reliability, correctness, and effectiveness that enable prospective
3643 consumers to have adequate confidence in using a service. Adequacy is defined by the consumer based
3644 on the consumer's needs and context of use. Absolute correctness and completeness cannot be proven
3645 by testing; however, for SOA, it is critical for the prospective consumer to know what testing has been
3646 performed, how it has been performed, and what were the results.

3647 **5.4.1 Traditional Software Testing as Basis for SOA Testing**

3648 SOA services are largely software artifacts and can leverage the body of experience that has evolved
3649 around software testing. [IEEE-829] specifies the basic set of software test documents while allowing
3650 flexibility for tailored use. Many testing frameworks are available but the SOA-RAF does not prescribe the
3651 use of any one in particular and choice will be driven by a framework that offers the right amount and
3652 level of testing. As such, the document structure can also provide guidance to SOA testing.

Comment [PFB115]: Issue 299

3653 IEEE-829 covers test specification and test reporting through use of the following document types:

- 3654 • *Test plan* documenting the scope (what is to be tested, both which entity and what features of the
3655 entity), the approach (how it is tested), and the needed resources (who does the testing, for how
3656 long), with details contained in the:
- 3657 • *Test design specification*: features to be tested, test conditions (e.g. test cases, test procedures
3658 needed) and expected results (criteria for passing test); entrance and exit criteria
- 3659 • *Test case specification*: test data used for input and expected output

- 3660 • *Test procedure specification*: steps required to run the test, including any set-up preconditions
- 3661 • *Test item transmittal* to identify the test items being transmitted for testing
- 3662 • *Test log* to record what occurred during test, i.e. which tests run, who ran, what order, what
- 3663 happened
- 3664 • *Test incident report* to capture any event that happened during test which requires further
- 3665 investigation
- 3666 • *Test summary* as a management report summarizing test run and results, conclusions

3667 In summary, IEEE-829 captures (1) what was tested, (2) how it was tested, e.g. the test procedure used,
3668 and (3) the results of the test.

3669 **5.4.1.1 Types of Testing**

3670 There are numerous aspects of testing that, in total, work to establish that an entity is (1) built as required
3671 per policies and related specifications prescribed by the entity's owner, and (2) delivers the functionality
3672 required by its intended users. This is often referred to as verification and validation.

3673 Policies, as described in Section 4.4, that are related to testing may prescribe but are not limited to the
3674 business processes to be followed, the standards with which an implementation must comply, and the
3675 qualifications of and restrictions on the users. In addition to the functional requirements prescribing what
3676 an entity does, there may also be non-functional performance and/or quality metrics that state how well
3677 the entity does it. The relation of these policies to SOA testing is discussed further below.

3678 The identification of policies is the purview of governance (section 5.1) and the assuring of compliance
3679 (including response to noncompliance) with policies is a matter for management (section 5.3).

3680 **5.4.1.2 Range of Test Conditions**

3681 Test conditions and expected responses are detailed in the test case specification. The test conditions
3682 should be designed to cover the areas for which the entity's response must be documented and may
3683 include:

- 3684 • nominal conditions
- 3685 • boundaries and extremes of expected conditions
- 3686 • breaking point where the entity has degraded below a certain level or has otherwise ceased
- 3687 effective functioning
- 3688 • random conditions to investigate unidentified dependencies among combinations of conditions
- 3689 • errors conditions to test error handling

3690 The specification of how each of these conditions should be tested for SOA resources, including the
3691 infrastructure elements of the SOA ecosystem, is beyond the scope of this document but is an area that
3692 evolves along with operational SOA experience.

3693 **5.4.1.3 Configuration Management of Test Artifacts**

3694 The test item transmittal provides an unambiguous identification of the entity being tested, thus
3695 REQUIRING that the configuration of the entity is appropriately tracked and documented. In addition, the
3696 test documents (such as those specified by IEEE-829) MUST also be under a documented and
3697 appropriately audited configuration management process, as should other resources used for testing. The
3698 description of each artifact would follow the general description model as discussed in section 4.1.1.1; in
3699 particular, it would include a version number for the artifact and reference to the documentation
3700 describing the versioning scheme from which the version number is derived.

3701 **5.4.2 Testing and the SOA Ecosystem**

3702 Testing of SOA artifacts for use in the SOA ecosystem differs from traditional software testing for several
3703 reasons. First, a highly touted benefit of SOA is to enable unanticipated consumers to make use of
3704 services for unanticipated purposes. Examples of this could include the consumer using a service for a
3705 result that was not considered the primary one by the provider, or the service may be used in combination
3706 with other services in a scenario that is different from the one considered when designing for the initial
3707 target consumer community. It is unlikely that a new consumer will push the services back to anything

3708 resembling the initial test phase to test the new use, and thus additional paradigms for testing are
3709 necessary. Some testing may depend on the availability of test resources made available as a service
3710 outside the initial test community, while some testing is likely to be done as part of limited use in the
3711 operational setting. The potential responsibilities related to such "consumer testing" are discussed further
3712 below.

3713 Secondly, in addition to consumers who interact with a service to realize the described real world effects,
3714 the developer community is also intended to be a consumer. In the SOA vision of reuse, the developer
3715 composes new solutions using existing services, where the existing services provides access to some
3716 desired real world effects that are needed by the new solution. The new solution is a consumer of the
3717 existing services, enabling repeated interactions with the existing services playing the role of reusable
3718 components. Note, those components are used at the locations where they individually reside and are not
3719 typically duplicated for the new solution. The new solution may itself be offered as a SOA service, and a
3720 consumer of the service composition representing the new solution may be totally unaware of the
3721 component services being used. (See section 4.3.4 for further discussion on service compositions.)

3722 Another difference from traditional testing is that the distributed, unbounded nature of the SOA ecosystem
3723 makes it unlikely to have an isolated test environment that duplicates the operational environment. A
3724 traditional testing approach often makes use of a test system that is identical to the eventual operational
3725 system but isolated for testing. After testing is successfully completed, the tested entity would be migrated
3726 to the operational environment, or the test environment may be delivered as part of the system to become
3727 operational. This is not feasible for the SOA ecosystem as a whole.

3728 SOA services must be testable in the environment and under the conditions that can be encountered in
3729 the operational SOA ecosystem. As the ecosystem is in a state of constant change, so some level of
3730 testing is continuous through the lifetime of the service, leveraging utility services used by the ecosystem
3731 infrastructure to monitor its own health and respond to situations that could lead to degraded
3732 performance. This implies the test resources must incorporate aspects of the SOA paradigm, and a
3733 category of services may be created to specifically support and enable effective monitoring and
3734 continuous testing for resources participating in the SOA ecosystem.

3735 While SOA within an enterprise may represent a more constrained and predictable operational
3736 environment, the composability and unanticipated use aspects are highly touted within the enterprise. The
3737 expanded perspective on testing may not be as demanding within an enterprise but fuller consideration of
3738 the ecosystem enables the enterprise to be more responsive should conditions change.

3739 **5.4.3 Elements of SOA Testing**

3740 IEEE-829 identifies fundamental aspects of testing, and many of these should carry over to SOA testing:
3741 in particular, the identification of what is to be tested, how it is to be tested, and by whom the testing is to
3742 be done. While IEEE-829 identifies a suggested document tree, the availability of these documents in the
3743 SOA ecosystem is discussed below.

3744 **5.4.3.1 What is to be Tested**

3745 The focus of this discussion is the SOA service. It is recognized that the infrastructure components of any
3746 SOA environment are likely to also be SOA services and, as such, falls under the same testing guidance.
3747 Other resources that contribute to a SOA environment may not be SOA services, but are expected to
3748 satisfy the intent if not the letter of guidance presented here. Specific differences for such resources are
3749 as yet largely undefined and further elaboration is beyond the scope of the SOA-RAF.

3750 The following discussion often focuses on a singular SOA service but it is implicit that any service may be
3751 a composite of other services. As such, testing the functionality of a composite service may effectively be
3752 testing an end-to-end business process that is being provided by the composite service. If new versions
3753 are available for the component services, appropriate end-to-end testing of the composite may be
3754 required in order to verify that the composite functionality is still adequately provided. The level of
3755 required testing of an updated composite depends on policies of those providing the service, policies of
3756 those using the service, and mission criticality of those depending on the service results.

3757 The SOA service to be tested MUST be unambiguously identified as specified by its applicable
3758 configuration management scheme. Specifying such a scheme is beyond the scope of the SOA-RAF
3759 other than to say the scheme should be documented and itself under configuration management.

3760 **5.4.3.1.1 Origin of Test Requirements**

3761 In the Service Description model (Figure 16), the aspects of a service that need to be described are:

- 3762 • the service functionality and technical assumptions that underlie the functionality;
- 3763 • the policies that describe conditions of use;
- 3764 • the service interface that defines information exchange with the service;
- 3765 • service reachability that identifies how and where message exchange is to occur; and
- 3766 • metrics access for any participant to have information on how a service is performing.

3767 Service testing must provide adequate assurance that each of these aspects is operational as defined.

3768 The information in the service description comes from different sources. The functionality is defined
3769 through whatever process identifies needs and the community for which these needs are addressed. The
3770 process may be ad hoc as serves the prospective service owner or strictly governed, but defining the
3771 functionality is an essential first step in development. It is also an early and ongoing focus of testing to
3772 ensure the service accurately reflects the described functionality and the described functionality
3773 accurately addresses the consumer needs.

3774 Policies define the conditions of development and conditions of use for a service and are typically
3775 specified as part of the governance process. Policies constraining service development, such as coding
3776 standards and best practices, require appropriate testing and auditing during development to ensure
3777 compliance. While the governance process identifies development policies, these are likely to originate
3778 from the technical community responsible for development activities. Policies that define conditions of use
3779 often define business practices that service owners and providers or those responsible for the SOA
3780 infrastructure want followed. These policies are initially tested during service development and are
3781 continuously monitored during the operational lifetime of the service.

3782 The testing of the service interface and service reachability are often related but essentially reflect
3783 different motivations and needs. The service interface is specified as a joint product of the service owners
3784 and providers who define service functionality, the prospective consumer community, the service
3785 developer, and the governance process. The semantics of the information model must align with the
3786 semantics of those who consume the service in order for there to be meaningful exchange of information.
3787 The structure of the information is influenced by the consumer semantics and the requirements and
3788 constraints of the representation as interpreted by the service developer. The service process model that
3789 defines actions which can be performed against a service and any temporal dependencies derive from
3790 the defined functionality and may be influenced by the development process. Any of these constraints
3791 may be identified and expressed as policy through the governance process.

3792 Service reachability conditions are the purview of the service provider who identifies the service endpoint
3793 and the protocols recognized at the endpoint. These may be constrained by governance decisions on
3794 how endpoint addresses may be allocated and what protocols should be used.

3795 While the considerations for defining the service interface derive from several sources, testing of the
3796 service interface is more straightforward and isolated in the testing process. At any point where the
3797 interface is modified or exposes a new resource, the message exchange should be monitored both to
3798 ensure the message reaches its intended destination and it is parsed correctly once received. Once an
3799 interface has been shown to function properly, it is unlikely to fail later unless something fundamental to
3800 the service changes.

3801 The service interface is also tested when the service endpoint changes. Testing of the endpoint ensures
3802 message exchange can occur at the time of testing and the initial testing shows the interface is being
3803 processed properly at the new endpoint. Functioning of a service endpoint at one time does not
3804 guarantee it is functioning at another time, e.g. the server with the endpoint address may be down,
3805 making testing of service reachability a continual monitoring function through the life of the service's use
3806 of the endpoint. Also, while testing of the service endpoint is a necessary and most commonly noted part
3807 of the test regiment, it is not in itself sufficient to ensure the other aspects of testing discussed in this
3808 section.

3809 Finally, governance is impossible without the collection of metrics against which service behavior can be
3810 assessed. Metrics are also a key indicator for consumers to decide if a service is adequate for their
3811 needs. For instance, the average response time or the recent availability can be determining factors even
3812 if there are no rules or regulations promulgated through the governance process against which these
3813 metrics are assessed. The available metrics are a combination of those expected by the consumer
3814 community and those mandated through the governance process. The total set of metrics will evolve over
3815 time with SOA experience. Testing of the services that gather and provide access to the metrics will follow
3816 testing as described in this section, but for an individual service, testing will ensure that the metrics
3817 access indicated in the service description is accurate.

3818 The individual test requirements highlight aspects of the service that testing must consider but testing
3819 must establish more than isolated behavior. The emphasis is the holistic results of interacting with the
3820 service in the SOA environment. Recall that the execution context is the set of agreements between a
3821 consumer and a provider that define the conditions under which service interaction occurs. The
3822 agreements are expected to be predominantly the acceptance of the standard conditions as enumerated
3823 by the service provider, but it may include the identification of alternate conditions that will govern the
3824 interaction.

3825 For example, the provider may prefer a policy where it can sell the contact information of its consumers
3826 but will honor the request of a consumer to keep such information private. The identification of the
3827 alternate privacy policy is part of the execution context, and it is the application of and compliance with
3828 this policy that operational monitoring will attempt to measure. The collection of metrics showing this
3829 condition is indeed met when chosen is considered part of the ongoing testing of the service.

3830 Other variations in the execution context also require monitoring to ensure that different combinations of
3831 conditions perform together as desired. For example, if a new privacy policy takes additional resources to
3832 apply, this may affect quality of service and propagate other effects. These could not be tested during the
3833 original testing if the alternate policy did not exist at that time.

3834 **5.4.3.1.2 Testing Against Non-Functional Requirements**

3835 Testing against non-functional requirements constitutes testing of business usability of the service. In a
3836 marketplace of services, non-functional characteristics may be the primary differentiator between services
3837 that produce essentially the same real world effects.

3838 As noted in the previous section, non-functional characteristics are often associated with policies or other
3839 terms of use and may be collected in service level contracts offered by the service providers. Non-
3840 functional requirements may also reflect the network and hardware infrastructure that support
3841 communication with the service, and changes may impact quality of service. The service consumer and
3842 even the service provider may not be aware of all such infrastructure changes but the changes may
3843 manifest in shared states that impact the usability of the service.

3844 In general, a change in the non-functional requirements results in a change to the execution context, but
3845 as with any collection of information that constitutes a description, the execution context is unable to
3846 explicitly capture all non-functional requirements that may apply. A change in non-functional
3847 requirements, whether explicitly part of the execution context or an implicit contributor, may require
3848 retesting of the service even if its functionality and the implementation of the functionality has not
3849 changed. Depending on the circumstances, retesting may require a formal recertifying of end-to-end
3850 behavior or more likely will be part of the continuous monitoring that applies throughout the service
3851 lifetime.

3852 **5.4.3.1.3 Testing Content and the Interests of Consumers**

3853 As noted in section 5.4.1.1, testing may involve verification of conformance with respect to policies and
3854 technical specifications and validation with respect to sufficiency of functionality to meet some prescribed
3855 use. It may also include demonstration of performance and quality aspects. For some of these items,
3856 such as demonstrating the business processes followed in developing the service or the use of standards
3857 in implementing the service, the testing or relevant auditing is done internal to the service development
3858 process and follows traditional software testing and quality assurance. If it is believed of value to potential
3859 consumers, information about such testing could be included in the service description. However, it is not

3860 required that all test or compliance artifacts be available to consumers, as many of the details tested may
3861 be part of the opacity of the service implementation.

3862 Some aspects of the service being tested will reflect directly on the real world effects realized through
3863 interaction with the service. In these cases, it is more likely that testing results will be directly relevant to
3864 potential consumers. For example, if the service was designed to correspond to certain elements of a
3865 business process or that a certain workflow is followed, testing should verify that the real world effects
3866 reflect that the business process or workflow were satisfactorily captured.

3867 The testing may also need to demonstrate that specified conditions of use are satisfied. For example,
3868 policies may be asserted that require certain qualifications of or impose restrictions on the consumers
3869 who may interact with the service. The service testing must demonstrate that the service independently
3870 enforces the policies or it provides the required information exchanges with the SOA ecosystem so other
3871 resources can ensure the specified conditions.

3872 The completeness of the testing, both in terms of the features tested and the range of parameters for
3873 which response is tested, depends on the context of expected use: the more critical the use, the more
3874 complete the testing. There are always limits on the resources available for testing, if nothing else than
3875 the service must be available for use in a finite amount of time.

3876 This again emphasizes the need for adequate documentation to be available. If the original testing is very
3877 thorough, it may be adequate for less demanding uses in the future. If the original testing was more
3878 constrained, then well-documented test results establish the foundation on which further testing can be
3879 defined and executed.

3880 **5.4.3.2 How Testing is to be Done**

3881 Testing should follow well-defined methodologies and, if possible, should reuse test artifacts that have
3882 proven generally useful for past testing. For example, IEEE-829 notes that test cases are separated from
3883 test designs to allow for use in more than one design and to allow for reuse in other situations. In the SOA
3884 ecosystem, description of such artifacts (as with description of a service) enables awareness of the item
3885 and describes how the artifact may be accessed or used.

3886 As with traditional testing, the specific test procedures and test case inputs are important so the tests are
3887 unambiguously defined and entities can be retested in the future. Automated testing and regression
3888 testing may be more important in the SOA ecosystem in order to re-verify a service is still acceptable
3889 when incorporated in a new use. For example, if a new use requires the services to deal with input
3890 parameters outside the range of initial testing, the tests could be rerun with the new parameters. If the
3891 testing resources are available to consumers within the SOA ecosystem, the testing as designed by test
3892 professionals could be consumed through a service accessed by consumers, and their results could
3893 augment those already in place. This is discussed further in the next section.

3894 **5.4.3.3 Who Performs the Testing**

3895 As with any software, the first line of testing is unit testing done by software developers. It is likely that
3896 initial testing will be done by those developing the software but may also be done independently by other
3897 developers. For SOA development, unit testing is likely confined to a development sandbox isolated from
3898 the SOA ecosystem.

3899 SOA testing will differ from traditional software testing in that testing beyond the development sandbox
3900 must incorporate aspects of the SOA ecosystem, and those doing the testing must be familiar with both
3901 the characteristics and responses of the ecosystem and the tools, especially those available as services,
3902 to facilitate and standardize testing. Test professionals will know what level of assurance must be
3903 established as the exposure of the service to the ecosystem and ecosystem to the service increases
3904 towards operational status. These test professionals may be internal resources to an organization or may
3905 evolve as a separate discipline provided through external contracting.

3906 As noted above, it is unlikely that a complete duplicate of the SOA ecosystem will be available for isolated
3907 testing, and thus use of ecosystem resources will manifest as a transition process rather than a step
3908 change from a test environment to an operational one. This is especially true for new composite services
3909 that incorporate existing operational services to achieve the new functionality. The test professionals will
3910 need to understand the available resources and the ramifications of this transition.

3911 As with current software development, a stage beyond work by test professionals will make use of a
3912 select group of typical users, commonly referred to as beta testers, to report on service response during
3913 typical intended use. This establishes fitness by the consumers, providing final validation of previously
3914 verified processes, requirements, and final implementation.

3915 In traditional software development, beta testing is the end of testing for a given version of the software.
3916 However, although the initial test phase can establish an appropriate level of confidence consistent with
3917 the designed use for the initial target consumer community, the operational service will exist in an
3918 evolving ecosystem, and later conditions of use may differ from those thought to be sufficient during the
3919 initial testing. Thus, operational monitoring becomes an extension of testing through the service lifetime.
3920 This continuous testing will attempt to ensure that a service does not consume an inordinate amount of
3921 ecosystem resources or display other behavior that degrades the ecosystem, but it will not undercover
3922 functional errors that may surface over time.

3923 As with any software, it is the responsibility of the consumers to consider the reasonableness of solutions
3924 in order to spot errors in either the software or the way the software is being used. This is especially
3925 important for consumers with unanticipated uses that may go beyond the original test conditions. It is
3926 unlikely the consumers will initiate a new round of formal testing unless the new use requires a
3927 significantly higher level of confidence in the service. Rather the consumer becomes a new extension to
3928 the testing regiment. Obvious testing would include a sanity check of results during the new use.
3929 However, if the details of legacy testing are associated with the service through the service description
3930 and if testing resources are available through automated testing services, then the new consumers can
3931 rerun and extend previous testing to include the extended test conditions. If the test results are
3932 acceptable, these can be added to the documentation of previous results and become the extended basis
3933 for future decisions by prospective consumers on the appropriateness of the service. If the results are not
3934 acceptable or in some way questionable, the responsible party for the service or testing professionals can
3935 be brought in to decide if remedial action is necessary.

3936 **5.4.3.4 How Testing Results are Reported**

3937 For any SOA service, an accurate reporting of the testing a service has undergone and the results of the
3938 testing is vital to consumers deciding whether a service is appropriate for intended use. Appropriateness
3939 may be defined by a consumer organization and require specific test regiments culminating in a
3940 certification; appropriateness could be established by accepting testing and certifications that have been
3941 conferred by others.

3942 The testing and certification information should be identified in the service description. Referring to the
3943 general description model of Figure 14, tests conducted by or under a request from the service owner
3944 (see ownership in section 3.2.4) would be captured under Annotations from Owners. Testing done by
3945 others, such as consumers with unanticipated uses, could be associated through Annotations from 3rd
3946 Parties. The annotations should clearly indicate what was tested, how the testing was done, who did the
3947 testing, and the testing results. The clear description of each of these artifacts and of standardized testing
3948 protocols for various levels of sophistication and completeness of testing would enable a common
3949 understanding and comparison of test coverage. It will also make it more straightforward to conduct and
3950 report on future testing, facilitating the maintenance of the service description.

3951 Consumer testing and the reporting of results raises additional issues. While stating who did the testing is
3952 mandatory, there may be formal requirements for authentication of the tester to ensure traceability of the
3953 testing claims. In some circumstances, persons or organizations would not be allowed to state testing
3954 claims unless the tester was an approved entity. In other cases, ensuring the tester had a valid email may
3955 be sufficient. In either case, it would be at the discretion of the potential consumer to decide what level of
3956 authentication was acceptable and which testers are considered authoritative in the context of their
3957 anticipated use.

3958 Finally, in a world of openly shared information, we would see an ever-expanding set of testing
3959 information as new uses and new consumers interact with a service. In reality, these new uses may
3960 represent proprietary processes or classified use that should only be available to authorized parties.
3961 Testing information, as with other elements of description, may require special access controls to ensure
3962 appropriate access and use.

3963 5.4.4 Testing SOA Services

3964 Testing of SOA services should be consistent with the SOA paradigm. In particular, testing resources and
3965 artifacts should be visible in support of service interaction between providers and consumers, where here
3966 the interaction is between the testing resource and the tester. In addition, the idea of opacity of the
3967 implementation should limit the details that need to be available for effective use of the test resources.
3968 Testing that requires knowledge of the internal structure of the service or its underlying capability should
3969 be performed as part of unit testing in the development sandbox, and should represent a minimum level
3970 of confidence before the service begins its transition to further testing and eventual operation in the SOA
3971 ecosystem.

3972 5.4.4.1 Progression of SOA Testing

3973 Software testing is a gradual exercise going from micro inspection to testing macro effects. The first step
3974 in testing is likely the traditional code reviews. SOA considerations would account for the distributed
3975 nature of SOA, including issues of distributed security and best practices to ensure secure resources. It
3976 would also set the groundwork for opacity of implementation, hiding programming details and simplifying
3977 the use of the service.

3978 Code review is likely followed by unit testing in a development sandbox isolated from the operational
3979 environment. The unit testing is done with full knowledge of the service internal structure and knowledge
3980 of resources representing underlying capabilities. It tests the interface to ensure exchanged messages
3981 are as specified in the service description and the messages can be parsed and interpreted as intended.
3982 Unit testing also verifies intended functionality and that the software has dealt correctly with internal
3983 dependencies, such as structure of a file system or access to other dedicated resources.

3984 Some aspects of unit testing require external dependencies be satisfied, and this is often done using
3985 mock objects to substitute for the external resources. In particular, it will likely be necessary to include
3986 mocks of existing operational services, both those provided as part of the SOA infrastructure and services
3987 from other providers.

3988 Service Mock

3989 A service mock is an entity that mimics some aspect of the performance of an operational service
3990 without committing to the real world effects that the operational service would produce.

3991 Mocks are discussed in detail in sections 5.4.4.3 and 5.4.4.4.

3992 After unit testing has demonstrated an adequate level of confidence in the service, the testing must
3993 transition from the tightly controlled environment of the development sandbox to an environment that
3994 more clearly resembles the operational SOA ecosystem or, at a minimum, the intended enterprise. While
3995 sandbox testing will use simple mocks of some aspects of the SOA environment, such as an interface to
3996 a security service without the security service functionality, the dynamic nature of SOA makes a full
3997 simulation infeasible to create or maintain. This is especially true when a new composite service makes
3998 use of operational services provided by others. Thus, at some point before testing is complete, the
3999 service will need to demonstrate its functionality by using resources and dealing with conditions that only
4000 exist in the full ecosystem or the intended enterprise. Some of these resources may still provide test
4001 interfaces -- more on this below -- but the interfaces will be accessible using the SOA environment and
4002 not just implemented for the sandbox.

4003 At this stage, the opacity of the service becomes important as the details of interacting with the service
4004 now rely on correct use of the service interface and not knowledge of the service internals. The workings
4005 of the service will only be observable through the real world effects realized through service interactions
4006 and external indications that conditions of use, such as user authentication, are satisfied. Monitoring the
4007 behavior of the service will depend on service interfaces that expose internal monitoring or provide
4008 required information to the SOA infrastructure monitoring function. The monitoring required to test a new
4009 service is likely to have significant overlap with the monitoring the SOA infrastructure includes to monitor
4010 its own health and to identify and isolate behavior outside of acceptable bounds. This is exactly what is
4011 needed as part of service testing, and it is reasonable to assume that the ecosystem transition includes
4012 use of operational monitoring rather than solely dedicated monitoring for each service being tested.

4013 Use of SOA monitoring resources during the explicit testing phase sets the stage for monitoring and a
4014 level of continual testing throughout the service lifetime.

4015 **5.4.4.2 Testing Traditional Dependencies vs. Service Interactions**

4016 A SOA service is not required to make use of other operational services beyond what may be required for
4017 monitoring by the ecosystem infrastructure. The service can implement hardcoded dependencies which
4018 have been tested in the development sandbox through the use of dedicated mocks. While coordination
4019 may be required with real data sources during integration testing, the dependencies can be constrained to
4020 things that can be tested in a more traditional manner. Policies can also be set to restrict access to pre-
4021 approved users, and thus the question of unanticipated users and unanticipated uses can be eliminated.
4022 Operational readiness can be defined in terms of what can be proven in isolated testing. While all this
4023 may provide more confidence in the service for its designed purpose, such a service will not fully
4024 participate in the benefits or challenges of the ecosystem. This is akin to filling a swimming pool with sea
4025 water and having someone in the pool say they are swimming in the ocean.

4026 In considering the testing needed for a fully participating service, consider the example of a new
4027 composite service that combines the real world effects and complies with the conditions of use of five
4028 existing operational services. The developer of the composite service does not own any of the component
4029 services and has limited, if any, ability to get the distributed owners to do any customization. The
4030 developer also is limited by the principle of opacity to information comprising the service description, and
4031 does not know internal details of the component services. The developer of the composite service must
4032 use the component services as they exist as part of the SOA environment, including what is provided to
4033 support testing by new users. This introduces requirements for what is needed in the way of service
4034 mocks.

4035 **5.4.4.3 Use of Service Mocks**

4036 Service mocks enable the tested service to respond to specific features of an operational service that is
4037 being used as a component. It allows service testing to proceed without needing access to or with only
4038 limited engagement with the component service. Mocks can also mimic difficult to create situations for
4039 which it is desired to test the new service response. For composite services using multiple component
4040 services, mocks may be used in combination to function for any number of the components. Note, when
4041 using service mocks, it is important to remember that it is not the component service that is being tested
4042 (although anomalous behavior may be uncovered during testing) but the use of the component in the new
4043 composite.

4044 Individual service mocks can emphasize different features of the component service they represent but
4045 any given mock does not have to mimic all features. For example, a mock of the service interface can
4046 echo a sent message and demonstrate the message is reaching its intended destination. A mock could
4047 go further and parse the sent message to demonstrate the message not only reached its destination but
4048 was understood. As a final step, the mock could report back what actions would have been taken by the
4049 component service and what real world effects would result. If the response mimicked the operational
4050 response, functional testing could proceed as if the real world effect actually occurred.

4051 There are numerous ways to provide mock functionality. The service mock could be a simulation of the
4052 operational service and return simulated results in a realistic response message or event notification. It is
4053 also possible for the operational service to act as its own mock and simply not execute the commit stage
4054 of its functionality. The service mock could use a combination of simulation and service action without
4055 commit to generate a report of what would have occurred during the defined interaction with the
4056 operational service.

4057 As the service proceeds through testing, mocks should be systematically replaced by the component
4058 resources accessed through their operational interfaces. Before beta testing begins, end-to-end testing,
4059 i.e. proceeding from the beginning of the service interaction to the resulting real world results, should be
4060 accomplished using component resources via their operational interfaces.

4061 **5.4.4.4 Providers of Service Mocks**

4062 In traditional testing, it is often the test professionals who design and develop the mocks, but in the
4063 distributed world of SOA, this may not be efficient or desirable.

4064 In the development sandbox, it is likely the new service developer or test professionals working with the
4065 developer will create mocks adequate for unit testing. Given that most of this testing is to verify the new
4066 service is performing as designed, it is not necessary to have high fidelity models of other resources
4067 being accessed. In addition, given opacity of SOA implementation, the developer of the new service may
4068 not have sufficient detailed knowledge of a component service to build a detailed mock of the component
4069 service functionality. Sharing existing mocks at this stage may be possible but the mocks would need to
4070 be implemented in the sandbox, and for simple models it is likely easier to build the mock from scratch.

4071 As testing begins its transition to the wider SOA environment, mocks may be available as services. For
4072 existing resources, it is possible that an Open Source model could evolve where service mocks of
4073 available functions can be catalogued and used during initial interaction of the tested service and the
4074 operational environment. Widely used functions may have numerous service mocks, some mimicking
4075 detailed conditions within the SOA infrastructure. However, the Open Source model is less likely to be
4076 sufficient for specialty services that are not widely used by a large consumer community.

4077 The service developer is probably best qualified for also developing more detailed service mocks or for
4078 mock modes of operational services. This implies that in addition to their operational interfaces, services
4079 will routinely provide test interfaces to enable service mocks to be used as services. As noted above, a
4080 new service developer wanting to build a mock of component services is limited to the description
4081 provided by the component service developer or owner. The description typically will detail real world
4082 effects and conditions of use but will not provide implementation details, some of which may be
4083 proprietary. Just as important in the SOA ecosystem, if it becomes standard protocol for developers to
4084 create service mocks of their own services, a new service developer is only responsible for building his
4085 own mocks and can expect other mocks to be available from other developers. This reduces duplication
4086 of effort where multiple developers would be trying to build the same mocks from the same insufficient
4087 information. Finally, a service developer is probably best qualified to know when and how a service mock
4088 should be updated to reflect modified functionality or message exchange.

4089 It is also possible that testing organizations will evolve to provide high-fidelity test harnesses for new
4090 services. The harnesses would allow new services to plug into a test environment and would facilitate
4091 accessing mocks of component services. However, it will remain a constant challenge for such
4092 organizations to capture evolving uses and characteristics of service interactions in the real SOA
4093 environment and maintain the fidelity and accuracy of the test systems.

4094 **5.4.4.5 Fundamental Questions for SOA Testing**

4095 In order for the transition to the SOA operational environment to proceed, it is necessary to answer two
4096 fundamental questions:

- 4097 • Who provides what testing resources for the SOA operational environment, e.g. mocks of
4098 interfaces, mocks of functionality, monitoring tools?
- 4099 • What testing needs to be accomplished before operational environment resources can be
4100 accessed for further testing?

4101 The discussion in section 5.4.4.4 notes various levels of sophistication of service mocks and different
4102 communities are likely to be responsible for different levels. Section 5.4.4.4 advocates a significant role
4103 for service developers, but there needs to be community consensus that such mocks are needed and that
4104 service developers will agree to fulfill this role. There is also a need for consensus as to what tools should
4105 be available as services from the SOA infrastructure.

4106 As for use of the service mocks and SOA environment monitoring services, practical experience is
4107 needed upon which guidelines can be established for when a new service has been adequately tested to
4108 proceed with a greater level of exposure with the SOA environment. Malfunctioning services could cause
4109 serious problems if they cannot be identified and isolated. On the other hand, without adequate testing
4110 under SOA operational conditions, it is unlikely that problems can be uncovered and corrected before
4111 they reach an operational stage.

4112 As noted in section 5.4.4.2, some of these questions can be avoided by restricting services to more
4113 traditional use scenarios. However, such restriction will limit the effectiveness of SOA use and the result
4114 will resemble the constraints of traditional integration activities we are trying to move beyond.

4115 **5.4.5 Architectural Implications for SOA Testing**

4116 The discussion of SOA Testing indicates numerous architectural implications on the SOA ecosystem:

- 4117 • The distributed, boundary-less nature of the SOA ecosystem makes it infeasible to create and
4118 maintain a single mock of the entire ecosystem to support testing activities.
- 4119 • A standard suite of monitoring services needs to be defined, developed, and maintained. This
4120 should be done in a manner consistent with the evolving nature of the ecosystem.
- 4121 • Services should provide interfaces that support access in a test mode.
- 4122 • Testing resources must be described and their descriptions must be catalogued in a manner that
4123 enables their discovery and access.
- 4124 • Guidelines for testing and ecosystem access need to be established and the ecosystem must be
4125 able to enforce those guidelines asserted as policies.
- 4126 • Services should be available to support automated testing and regression testing.
- 4127 • Services should be available to facilitate updating service description by anyone who has
4128 performed testing of a service.

4129 6 Conformance

4130 This Reference Architecture Framework is an abstract architectural description of Service Oriented
4131 Architecture, which means that it is especially difficult to construct tests for conformance to the
4132 architecture. In addition, conformance to an architectural specification does not, by itself, guarantee any
4133 form of interoperability between multiple implementations.

4134 However, it *is* possible to decide whether or not a given architecture is conformant to an architectural
4135 description such as this one. In discussions of conformance we use the term **target architecture** to
4136 identify the (typically concrete) architecture that may be viewable as conforming to the abstract principles
4137 outlined in this document.

4138 Target Architecture

4139 A target architecture is an architectural description of a system that is intended to be viewed as
4140 conforming to the SOA-RAF.

4141 While we cannot guarantee interoperability between target architectures (or more specifically between
4142 applications and systems residing within the ecosystems of those target architectures), interoperability
4143 between target architectures is promoted by conformance to this Reference Architecture Framework as it
4144 reduces the semantic impedance mismatch between the different ecosystems.

4145 The primary measure of conformance is whether given concepts as described in document have
4146 corresponding concepts in the target architecture. Such a correspondence **MUST** honor the relationships
4147 identified within this document for the target architecture to be considered conforming.

4148 For example, in Section 3.2.4.1 we identify resource as a key concept. A resource is associated with an
4149 owner and a number of identifiers. For a target architecture to conform to the SOA-RAF, it must be
4150 possible to find corresponding concepts of resource, identifier and owner within the target architecture:
4151 say *entity*, *token* and *user*. Furthermore, the relationships between *entity*, *token* and *user* **MUST** mirror
4152 the relationships between resource, identifier and owner appropriately.

4153 Clearly, such correspondence is simpler if the terminology within the target architecture is identical to that
4154 in the SOA-RAF. But so long as the 'graph' of concepts and relationships is consistent, that is all that is
4155 required for the target architecture to conform to this Reference Architecture Framework.

4156 ~~[EDITOR'S NOTE: The conformance section is not complete]~~

4157 **A. Acknowledgements**

4158 The following individuals have participated in the work of the technical committee responsible for creation
4159 of this specification and are gratefully acknowledged:

4160 **Participants:**

- 4161 Chris Bashioum, MITRE Corporation
- 4162 Rex Brooks, Individual Member
- 4163 Peter F Brown, Individual Member
- 4164 Scott Came, Search Group Inc.
- 4165 Joseph Chiusano, Booz Allen Hamilton
- 4166 Robert Ellinger, Northrop Grumman Corporation
- 4167 David Ellis, Sandia National Laboratories
- 4168 Jeff A. Estefan, Jet Propulsion Laboratory
- 4169 Don Flinn, Individual Member
- 4170 Anil John, Johns Hopkins University
- 4171 Ken Laskey, MITRE Corporation
- 4172 Boris Lublinsky, Nokia Corporation
- 4173 Francis G. McCabe, Individual Member
- 4174 Christopher McDaniels, USSTRATCOM
- 4175 Tom Merkle, Lockheed Martin Corporation
- 4176 Jyoti Namjoshi, Patni Computer Systems Ltd.
- 4177 Duane Nickull, Adobe Inc.
- 4178 James Odell, Associate
- 4179 Michael Poulin, Fidelity Investments
- 4180 Michael Stiefel, Associate
- 4181 Danny Thornton, Northrop Grumman
- 4182 Timothy Vibbert, Lockheed Martin Corporation
- 4183 Robert Vitello, New York Dept. of Labor

4184 The committee would particularly like to underline the significant writing and conceptualization
4185 contributions made by Chris Bashioum, Rex Brooks, Peter Brown, Dave Ellis, Jeff Estefan, Ken Laskey,
4186 Boris Lublinsky, Frank McCabe, Michael Poulin and Danny Thornton

4187

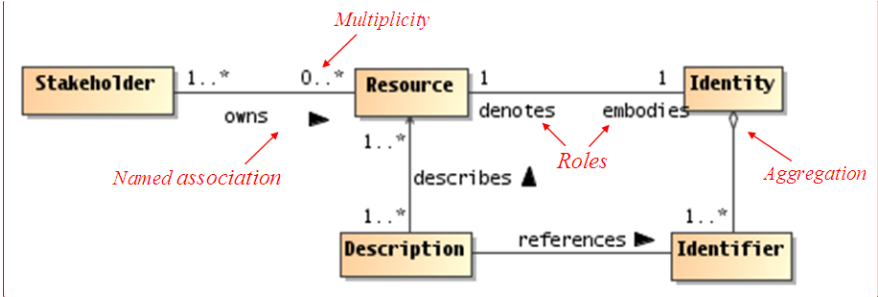
B. Index of Defined Terms

4188	Action.....	37	Policy.....	33
	Action Level Real World Effect	49	Policy Conflict.....	77
	Actor	25	Policy Conflict Resolution	77
	Authority	26	Policy Constraint	76
	Business Collaboration.....	69	Policy Decision	76
	Business functionality.....	29	Policy Enforcement	76
	Business Process	69	Policy Framework.....	75
	Business solution.....	36	Policy Object.....	76
	Capability.....	29	Policy Ontology	75
	Choreography	72	Policy Owner	76
	Communication	34	Policy Subject.....	76
	Composability	36	Presence.....	64
	Constitution	24	Private State	39
	Consumer.....	27	Protocol.....	64
	Contract.....	34	Provider.....	27
	Delegate	25	Real World Effect	29
	Endpoint	63	Regulation.....	82
	Governance.....	80	Requirement	29
	Governance Framework.....	81	Resource.....	30
	Governance Processes.....	81	Responsibility	26
	Identifier.....	30	Right.....	26
	Joint Action.....	38	Risk	32
	Leadership.....	81	Rule.....	82
	Logical Framework	75	Security	89
	Manageability.....	96	Semantic Engagement	35
	Manageability property	96	Service Contract.....	100
	Mediator.....	27	Service Level Real World Effect.....	49
	Message Exchange	67	Service Mock	111
	Need	29	Shared State.....	39
	Non-Participant	25	SOA Ecosystem	21
	Obligation	27	SOA-based System.....	21
	Operations.....	67	Social Structure	23
	Orchestration.....	70	Stakeholder.....	24
	Owner	27	State	39
	Ownership	31	Target Architecture.....	115
	Ownership Boundary	31	Trust	32
	Participant	25	Willingness.....	32
	Permission.....	27		

4189

4190 **C. The Unified Modeling Language, UML**

4191 Figure 48 illustrates an annotated example of a UML class diagram that is used to represent a visual
4192 model depiction of the Resources Model in the *Participation in a SOA Ecosystem* view.



Comment [PFB118]: Revise and update (using only used classes) once all figures are harmonized using same toolset Issue 163

4193
4194 *Figure 48 - Example UML class diagram—Resources*

4195 Lines connecting boxes (classifiers) represent associations between things. An association has two roles
4196 (one in each direction). A role can have cardinality, for example, one or more (“1..*”) stakeholders own
4197 zero or more (“0..*”) resources. The role from classifier A to B is labeled closest to B, and vice versa, for
4198 example, the role between resource to Identity can be read as resource embodies Identity, and Identity
4199 denotes a resource.

4200 Mostly, we use named associations, which are denoted with a verb or verb phrase associated with an
4201 arrowhead. A named association reads from classifier A to B, for example, one or more stakeholders
4202 owns zero or more resources. Named associations are a very effective way to model relationships
4203 between concepts.

4204 An open diamond (at the end of an association line) denotes an aggregation, which is a part-of
4205 relationship, for example, Identifiers are part of Identity (or conversely, Identity is made up of Identifiers).

4206 A stronger form of aggregation is known as composition, which involves using a filled-in diamond at the
4207 end of an association line (not shown in above diagram). For example, if the association between Identity
4208 and Identifier were a composition rather than an aggregation as shown, deleting Identity would also
4209 delete any owned Identifiers. There is also an element of exclusive ownership in a composition
4210 relationship between classifiers, but this usually refers to specific instances of the owned classes
4211 (objects).

4212 This is by no means a complete description of the semantics of all diagram elements that comprise a
4213 UML class diagram, but rather is intended to serve as an illustrative example for the reader. It should be
4214 noted that the SOA-RAF utilizes additional class diagram elements as well as other UML diagram types
4215 such as sequence diagrams and component diagrams. The reader who is unfamiliar with the UML is
4216 encouraged to review one or more of the many useful online resources and book publications available
4217 describing UML (see, for example, www.uml.org).

4218 **D. Critical Factors Analysis**

4219 A critical factors analysis (CFA) is an analysis of the key properties of a project. A CFA is analyzed in
4220 terms of the goals of the project, the critical factors that will lead to its success and the measurable
4221 requirements of the project implementation that support the goals of the project. CFA is particularly
4222 suitable for capturing quality attributes of a project, often referred to as “non-functional” or “other-than-
4223 functional” requirements: for example, security, scalability, wide-spread adoption, and so on. As such,
4224 CFA complements rather than attempts to replace other requirements capture techniques.

4225 **D.1 Goals**

4226 A goal is an overall target that you are trying to reach with the project. Typically, goals are hard to
4227 measure by themselves. Goals are often directed at the potential consumer of the product rather than the
4228 technology developer.

4229 **D.2 Critical Success Factors**

4230 A critical success factor (CSF) is a property, sub-goal that directly supports a goal and there is strong
4231 belief that without it the goal is unattainable. CSFs themselves are not necessarily measurable in
4232 themselves.

4233 **D.3 Requirements**

4234 A requirement is a specific measurable property that directly supports a CSF. The key here is
4235 measurability: it should be possible to unambiguously determine if a requirement has been met. While
4236 goals are typically directed at consumers of the specification, requirements are focused on technical
4237 aspects of the specification.

Comment [KJL119]: Appendix expanded to include HL7 material per comment 298 and setting for appropriate reference to IEEE RA

4238

E. Relationship to other SOA Open Standards

4239 [Numerous efforts have been working in the space of defining standards for SOA and its applications. The](#)
4240 [OASIS SOA-RM Technical Committee and its SOA-RA Technical Subcommittee has established](#)
4241 [communications with several of these efforts in an attempt to coordinate and facilitate among the efforts.](#)
4242 [This appendix notes some of these efforts.](#)

E.1 Navigating the SOA Open Standards Landscape Around Architecture

4244 The white paper “Navigating the SOA Open Standards Landscape Around Architecture” issued jointly by
4245 OASIS, OMG, and The Open Group [**SOA-NAV**] was written to help the SOA community at large
4246 navigate the myriad of overlapping technical products produced by these organizations with specific
4247 emphasis on the “A” in SOA, i.e., Architecture.

4248 The white paper explains and positions standards for SOA reference models, ontologies, reference
4249 architectures, maturity models, modeling languages, and standards work on SOA governance. It outlines
4250 where the works are similar, highlights the strengths of each body of work, and touches on how the work
4251 can be used together in complementary ways. It is also meant as a guide to users for selecting those
4252 specifications most appropriate for their needs.

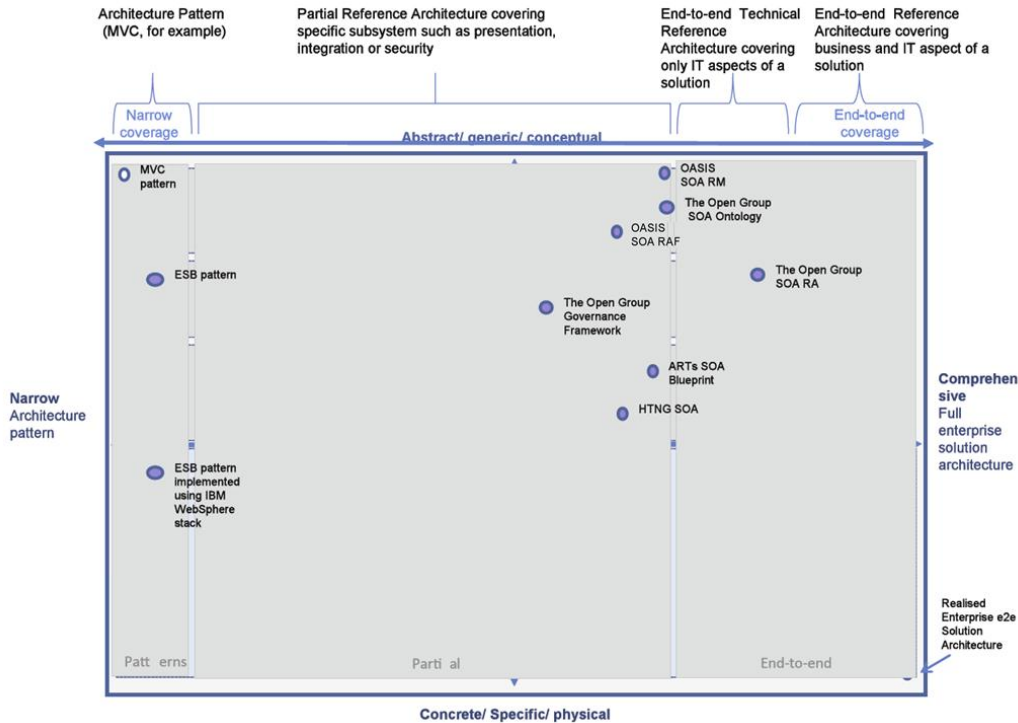
4253 While the understanding of SOA and SOA Governance concepts provided by these works is similar, the
4254 evolving standards are written from different perspectives. Each specification supports a similar range of
4255 opportunity, but has provided different depths of detail for the perspectives on which they focus. Although
4256 the definitions and expressions may differ, there is agreement on the fundamental concepts of SOA and
4257 SOA Governance.

4258 The following is a summary taken from [**SOA-NAV**] of the positioning and guidance on the specifications:

- 4259 • The OASIS Reference Model for SOA (SOA RM) is the most abstract of the specifications
4260 positioned. It is used for understanding core SOA concepts
- 4261 • The Open Group SOA Ontology extends, refines, and formalizes some of the core concepts of
4262 the SOA RM. It is used for understanding core SOA concepts and facilitates a model-driven
4263 approach to SOA development.
- 4264 • The OASIS Reference Architecture Foundation for SOA (this document) is an abstract,
4265 foundational reference architecture addressing a broader ecosystem viewpoint for building and
4266 interacting within the SOA paradigm. It is used for understanding different elements of SOA, the
4267 completeness of SOA architectures and implementations, and considerations for reaching across
4268 ownership boundaries where there is no single authoritative entity for SOA and SOA governance.
- 4269 • The Open Group SOA Reference Architecture is a layered architecture from consumer and
4270 provider perspective with cross cutting concerns describing these architectural building blocks
4271 and principles that support the realizations of SOA. It is used for understanding the different
4272 elements of SOA, deployment of SOA in enterprise, basis for an industry or organizational
4273 reference architecture, implication of architectural decisions, and positioning of vendor products in
4274 a SOA context.
- 4275 • The Open Group SOA Governance Framework is a governance domain reference model and
4276 method. It is for understanding SOA governance in organizations. The OASIS Reference
4277 Architecture for SOA Foundation contains an abstract discussion of governance principles as
4278 applied to SOA across boundaries
- 4279 • The Open Group SOA Integration Maturity Model (OSIMM) is a means to assess an
4280 organization’s maturity within a broad SOA spectrum and define a roadmap for incremental
4281 adoption. It is used for understanding the level of SOA maturity in an organization
- 4282 • The Object Management Group SoaML Specification supports services modeling UML
4283 extensions. It can be seen as an instantiation of a subset of the Open Group RA used for
4284 representing SOA artifacts in UML.

4285 Fortunately, there is a great deal of agreement on the foundational core concepts across the many
4286 independent specifications and standards for SOA. This could be best explained by broad and common
4287 experience of users of SOA and its maturity in the marketplace. It also provides assurance that investing

4288 in SOA-based business and IT transformation initiatives that incorporate and use these specifications and
 4289 standards helps to mitigate risks that might compromise a successful SOA solution.



4290
 4291 Figure 49 - SOA Reference Architecture Positioning (from "Navigating the SOA Open Standards Landscape Around
 4292 Architecture, © OASIS, OMG, The Open Group)

Comment [PFB120]: Issue 298

4293 **E.2 The Service-Aware Interoperability Framework: Canonical**

4294 Readers of the RAF are strongly encouraged to review a document recently published by the Health
 4295 Level Seven (HL7) Architecture Board (ArB) entitled "The Service-Aware Interoperability Framework:
 4296 Canonical." The document was developed over the past four years, and represents a substantive,
 4297 industry-specific (i.e. the large but vertical healthcare industry) effort to surface, define, and discuss in
 4298 detail various aspects of a number of critical success factors involved in implementing large-scale (i.e.
 4299 enterprises-level) architectures with a focus on achieving both intra- and inter-enterprise technical
 4300 interoperability irrespective of the particular exchange mechanism involved, e.g. service interface,
 4301 messages, or structure documents.

4302 In addition to providing an independent validation for the both the general focus as well as some of the
 4303 concrete specifics of the RAF (especially those involving the importance of governance in achieving
 4304 large-scale interoperability), the HL7 document underscores several important aspects of the RAF
 4305 including:

- 4306 1. A validation of one of the RAF's primary claims, i.e. the need to specifically focus on intra- and inter-
 4307 enterprise interoperability as a first-class citizen in any enterprise (or cross-enterprise) architecture
 4308 discussion irrespective of the particular choice of enterprise architecture approach, framework, or
 4309 implementation technology, e.g. TOGAF, Zachman, ODP, SOA, etc. In addition, the HL7 document
 4310 clearly articulates – as the RAF does as well – the difficulties involved in achieving that focus in such
 4311 a manner that it can be manifest in operationally effective and manageable processes and
 4312 deliverables.

4313 2. An agreement as to the critical importance of governance as the root of any successful effort to
4314 implement large-scale, cross-boundary interoperability aimed at achieving a collective shared
4315 purpose-mission or goal. In particular, both documents share the notion that “technical-level”
4316 governance – e.g. service – or message-level technical interchange specifications – must itself be a
4317 manifestation of a higher-level, cross-jurisdictional agreement on desired goals, responsibilities,
4318 accountabilities, and deliverables.

4319 3. A validation of the importance of core SOA constructs as constructs useful in expressing many of the
4320 central aspects of interoperability irrespective of whether a particular interoperability scenario is
4321 actually “realized” using SOA-compatible technologies. (NOTE: Although it might at first appear that
4322 the OASIS document is more “service-focused” than the “service-aware” document from HL7, there
4323 are considerably more similarities than differences in these slightly different foci secondary to the fact
4324 that both documents are intent on describing principles and framework concepts rather than delving
4325 into technical details. There are, however, certain instances where content of the OASIS document
4326 would be likely to find its analogue in SAIF Implementation Guides rather than in the SAIF Canonical
4327 Definition document.)

4328 4. The need for specific, explicit statements of those aspects of a given component that affects its ability
4329 to participate in a reliable, predictable manner in a variety of interoperability scenarios. In particular,
4330 component characteristics must be explicitly expressed in both design-time and run-time contexts as
4331 implicit assumptions are the root of most failures to achieve successfully cross-boundary
4332 interoperability irrespective of the chosen technical details of a particular interoperability instance.

4333 In summary, although the two documents are clearly not identical in their specifics, e.g. there are
4334 differences in the language used to name various concepts, constructs, and relationships; there are some
4335 differences in levels of abstraction regarding certain topics, etc; and although the OASIS RAF is more
4336 directly focused on services as a final implementation architecture than the HL7 SAIF CD, the
4337 commonalities of purpose, content, and approach present in the two documents – documents which were
4338 developed by each organization without any knowledge of the others’ work in what clearly are areas of
4339 common interest and concern – far outweighs their differences. As such, the HL7 ArB and the OASIS
4340 RAF Task Force have agreed to work together going forward to obtain the highest degree of alignment
4341 and harmonization possible between the two documents including the possible development of a joint
4342 document under the auspices of one of the ISO software engineering threads.

4343
4344 The current version of the HL7 document – as well as all future versions – is available at:
4345 <http://www.hl7.org/permalink/?SAIFCDR1PUBLIC>

4346 **E.3 IEEE Reference Architecture**

Comment [PFB121]: Issue 298 – to be completed

4347 TBD

4348 **E.4 RM-ODP**

Comment [PFB122]: Issue 298 – to be completed

4349 TBD