



Reference Architecture Foundation for Service Oriented Architecture Version 1.0

Working Draft 08

06 July 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/46409> (3 July 2012)

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	11
1.1	Context for Reference Architecture for SOA	11
1.1.1	What is a Reference Architecture?	11
1.1.2	What is this Reference Architecture?	11
1.1.3	Relationship to the OASIS Reference Model for SOA	12
1.1.4	Relationship to other Reference Architectures	12
1.1.5	Expectations set by this Reference Architecture Foundation	12
1.2	Service Oriented Architecture – An Ecosystems Perspective	13
1.3	Viewpoints, Views and Models	13
1.3.1	ANSI/IEEE 1471-2000:ISO/IEC 42010-2007	13
1.3.2	UML Modeling Notation	14
1.4	SOA-RAF Viewpoints	15
1.4.1	Participation in a SOA Ecosystem Viewpoint	15
1.4.2	Realization of a SOA Ecosystem Viewpoint	15
1.4.3	Ownership in a SOA Ecosystem Viewpoint	16
1.5	Terminology	16
1.6	References	16
1.6.1	Normative References	16
1.6.2	Non-Normative References	17
2	Architectural Goals and Principles	18
2.1	Goals and Critical Success Factors of the Reference Architecture Foundation	18
2.1.1	Goals	18
2.1.1.1	Effectiveness	18
2.1.1.2	Confidence	18
2.1.1.3	Scalability	18
2.1.2	Critical Success Factors	19
2.1.2.1	Action	19
2.1.2.2	Trust	19
2.1.2.3	Interaction	19
2.1.2.4	Control	19
2.2	Principles of this Reference Architecture Foundation	19
3	Participation in a SOA Ecosystem View	21
3.1	SOA Ecosystem Model	22
3.2	Social Structure in a SOA Ecosystem Model	23
3.2.1	Stakeholders, Participants, Actors and Delegates	25
3.2.2	Social Structures and Roles	26
3.2.2.1	Authority, Rights, and Responsibilities	27
3.2.2.2	Permissions and Obligations	28
3.2.2.3	Service Roles	28
3.2.3	Needs, Requirements and Capabilities	29
3.2.4	Resource and Ownership	31
3.2.4.1	Resource	31
3.2.4.2	Ownership	32

3.2.5	Establishing Execution Context	32
3.2.5.1	Trust and Risk.....	33
3.2.5.2	Policies and Contracts	34
3.2.5.3	Communication.....	35
3.2.5.4	Semantics and Semantic Engagement.....	35
3.3	Action in a SOA Ecosystem Model.....	36
3.3.1	Services Reflecting Business.....	37
3.3.2	Activity, Action, and Joint Action	37
3.3.3	State and Shared State	39
3.4	Architectural Implications	40
3.4.1	Social structures.....	40
3.4.2	Resource and Ownership	40
3.4.3	Policies and Contracts	40
3.4.4	Communications as a Means of Mediating Action.....	40
3.4.5	Semantics	41
3.4.6	Trust and Risk	41
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	59
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.....	63
4.3	Interacting with Services Model	64
4.3.1	Interaction Dependencies.....	64
4.3.2	Actions and Events	65
4.3.3	Message Exchange	65
4.3.3.1	Message Exchange Patterns (MEPs).....	66
4.3.3.2	Request/Response MEP	67
4.3.3.3	Event Notification MEP	68

4.3.4	Composition of Services	68
4.3.5	Implementing Service Composition	69
4.3.5.1	Service-Oriented Business Processes	69
4.3.5.2	Service-Oriented Business Collaborations	70
4.3.6	Architectural Implications of Interacting with Services	72
4.4	Policies and Contracts Model	73
4.4.1	Policy and Contract Representation	73
4.4.2	Policy and Contract Enforcement	74
4.4.2.1	Enforcing Simple Policy Constraints	75
4.4.2.2	Conflict Resolution	75
4.4.3	Architectural Implications.....	75
5	Ownership in a SOA Ecosystem View	77
5.1	Governance Model	77
5.1.1	Understanding Governance	77
5.1.1.1	Terminology	77
5.1.1.2	Relationship to Management.....	78
5.1.1.3	Why is SOA Governance Important?.....	78
5.1.1.4	Governance Stakeholders and Concerns.....	78
5.1.2	A Generic Model for Governance	79
5.1.2.1	Motivating Governance	79
5.1.2.2	Setting Up Governance.....	80
5.1.2.3	Carrying Out Governance	81
5.1.2.4	Ensuring Governance Compliance	82
5.1.2.5	Considerations for Multiple Governance Chains.....	82
5.1.3	Governance Applied to SOA	83
5.1.3.1	Where SOA Governance is Different.....	83
5.1.3.2	What Must be Governed	83
5.1.3.3	Overarching Governance Concerns	85
5.1.3.4	Considerations for SOA Governance.....	86
5.1.4	Architectural Implications of SOA Governance.....	87
5.2	Security Model	87
5.2.1	Secure Interaction Concepts.....	88
5.2.1.1	Confidentiality	88
5.2.1.2	Integrity.....	88
5.2.1.3	Authentication	88
5.2.1.4	Authorization	89
5.2.1.5	Non-repudiation	90
5.2.1.6	Availability	90
5.2.2	Where SOA Security is Different.....	90
5.2.3	Security Threats	90
5.2.4	Security Responses	91
5.2.4.1	Privacy Enforcement	91
5.2.4.2	Integrity Protection	92
5.2.4.3	Message Replay Protection.....	92
5.2.4.4	Auditing and Logging	92

5.2.4.5 Graduated engagement	93
5.2.5 Identity and Access Control.....	93
5.2.5.1 Identity Propagation	93
5.2.5.2 Access Control Approaches	94
5.2.6 Architectural Implications of SOA Security.....	96
5.3 Management Model.....	97
5.3.1 Management.....	97
5.3.2 Management Means and Relationships	100
5.3.2.1 Management Policy.....	100
5.3.2.2 Network Management.....	101
5.3.2.3 Security Management	101
5.3.2.4 Usage Management.....	101
5.3.3 Management and Governance.....	101
5.3.4 Management and Contracts.....	102
5.3.4.1 Management for Contracts and Policies.....	102
5.3.4.2 Contracts.....	102
5.3.4.3 Policies	104
5.3.4.4 Service Description and Management.....	104
5.3.5 Management for Monitoring and Reporting.....	105
5.3.6 Management for Infrastructure.....	105
5.3.7 Architectural Implication of the SOA Management.....	106
5.4 SOA Testing Model.....	106
5.4.1 Traditional Software Testing as Basis for SOA Testing	106
5.4.1.1 Types of Testing	107
5.4.1.2 Range of Test Conditions	107
5.4.1.3 Configuration Management of Test Artifacts	107
5.4.2 Testing and the SOA Ecosystem	107
5.4.2.1 Testing and the Consumer Communities	107
5.4.2.2 Testing and the Evolving SOA Ecosystem	108
5.4.3 Elements of SOA Testing	108
5.4.3.1 What is to be Tested	108
5.4.3.2 How Testing is to be Done.....	111
5.4.3.3 Who Performs the Testing	111
5.4.3.4 How Testing Results are Reported	112
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	114
5.4.4.4 Providers of Service Mocks	114
5.4.4.5 Fundamental Questions for SOA Testing.....	115
5.4.5 Architectural Implications for SOA Testing.....	115
6 Conformance	117
A. Acknowledgements	118
B. Index of Defined Terms.....	119
C. Critical Factors Analysis.....	120

C.1 Goals.....	120
C.2 Critical Success Factors	120
C.3 Requirements	120
D. Relationship to other SOA Open Standards.....	121
D.1 Navigating the SOA Open Standards Landscape Around Architecture.....	121
D.2 The Service-Aware Interoperability Framework: Canonical	122
D.3 IEEE Reference Architecture	123
D.4 RM-ODP	123

Table of Figures

Figure 1 - Model elements described in the Participation in a SOA Ecosystem view	21
Figure 2 - SOA Ecosystem Model	22
Figure 3 - Social Structure Model	24
Figure 4 – Stakeholders, Actors, Participants and Delegates	25
Figure 5 - Social Structures, Roles and Action	27
Figure 6 - Roles in a Service	29
Figure 7 - Cycle of Needs, Requirements, and Fulfillment	30
Figure 8 - Resources	31
Figure 9 - Willingness and Trust	33
Figure 10 – Policies, Contracts and Constraints	34
Figure 11: An Activity, expressed informally as a graph of Actions	38
Figure 12: Activity involving Actions across an ownership boundary	38
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 - Service Reachability	63
Figure 28 - Interaction dependencies	65
Figure 29 - A 'message' denotes either an action or an event	65
Figure 30 - Fundamental SOA message exchange patterns (MEPs)	67
Figure 31 - Simple model of service composition	68
Figure 32 - Abstract example of a simple business process exposed as a service	70
Figure 33 - Abstract example of a more complex composition that relies on collaboration	71
Figure 34 - Policies and Contracts	74
Figure 35 - Model Elements Described in the Ownership in a SOA Ecosystem View	77
Figure 36 - Motivating Governance	79
Figure 37 - Setting Up Governance	80
Figure 38 - Carrying Out Governance	81
Figure 39 - Ensuring Governance Compliance	82
Figure 40 - Relationship Among Types of Governance	84

Figure 41 - Authorization.....89
Figure 42 - Management model in SOA ecosystem98
Figure 43 - Management Means and Relationships in a SOA ecosystem100
Figure 44 - Management of the service interaction.....103
Figure 45 - 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 The UML presented should not be treated blindly or automatically; the models are intended to formalize
175 the concepts and relationships defined and described in the text but the nature of the RAF means that it
176 still concerns an abstract layer rather than an implementable layer.

Comment [PFB1]: Issue 40

177 **1.4 SOA-RAF Viewpoints**

178 The SOA-RAF specifies three views (described in detail in Sections 3, 4, and 5) that conform to three
 179 viewpoints: *Participation in a SOA Ecosystem*, *Realization of a SOA Ecosystem*, and *Ownership in a SOA*
 180 *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

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

183 **1.4.1 Participation in a SOA Ecosystem Viewpoint**

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

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

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

194 **1.4.2 Realization of a SOA Ecosystem Viewpoint**

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

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

201 1.4.3 Ownership in a SOA Ecosystem Viewpoint

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

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

210 1.5 Terminology

211 The keywords “MUST”, “MUST NOT”, “REQUIRED” (and by extension, “REQUIRES”), “SHALL”, “SHALL
212 NOT”, “SHOULD”, “SHOULD NOT”, “RECOMMENDED”, “MAY”, and “OPTIONAL” in this document are
213 to be interpreted as described in [RFC2119].

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

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

218 Usage of Terms

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

225 1.6 References

226 1.6.1 Normative References

- 227 [ANSI/IEEE 1471] *IEEE Recommended Practice for Architectural Description of Software-Intensive*
228 *Systems*, American National Standards Institute/Institute for Electrical and
229 *Electronics Engineers*, September 21, 2000.
- 230 [ISO/IEC 10746] International Organization for Standardization and International Electrotechnical
231 Commission, *Information Technology – Open Distributed Processing –*
232 *Reference Model*, September 15, 1996.
- 233 [ISO/IEC IS 19793] [International Organization for Standardization and International](#)
234 [Electromechanical Commission, Information Technology – Open Distributed](#)
235 [Processing – Use of UML for ODP System Specification, 2008 \(Also published](#)
236 [as ITU-T recommendation X.906\).](#)
- 237 [ISO/IEC 42010] International Organization for Standardization and International Electrotechnical
238 Commission, *System and software engineering – Recommended practice for*
239 *architectural description of software-intensive systems*, July 15, 2007.
- 240 [RFC2119] S. Bradner, *Key words for use in RFCs to Indicate Requirement Levels*,
241 <http://www.ietf.org/rfc/rfc2119.txt>, IETF RFC 2119, March 1997.
- 242 [SOA-RM] OASIS Standard, "Reference Model for Service Oriented Architecture 1.0, 12
243 October 2006. <http://docs.oasis-open.org/soa-rm/v1.0/soa-rm.pdf>
- 244 [UML 2] *Unified Modeling Language: Superstructure*, Ver. 2.1.1, OMG Adopted
245 Specification, OMG document formal/2007-02-05, Object Management Group,
246 Needham, MA, February 5, 2007.
- 247 [WSA] David Booth, et al., "Web Services Architecture", W3C Working Group Note,
248 World Wide Web Consortium (W3C) (Massachusetts Institute of Technology,

249 European Research Consortium for Informatics and Mathematics, Keio
250 University), February, 2004. [http://www.w3.org/TR/2004/NOTE-ws-arch-](http://www.w3.org/TR/2004/NOTE-ws-arch-20040211/)
251 [20040211/](http://www.w3.org/TR/2004/NOTE-ws-arch-20040211/)

252 1.6.2 Non-Normative References

253 [BLOOMBERG/SCHMELZER]

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

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

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

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

261 [LININGTON] Peter Linington, Zoran Milosevic, Akira Tanaka, Antonio Vallecillo, *Building*
262 *Enterprise Systems with ODP*, Chapman & Hall / CRC, 2012.

263 [NEWCOMER/LOMOW]

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

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

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

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

273 [SOA NAV] Heather Kreger and Jeff Estefan (Eds.), "Navigating the SOA Open Standards
274 Landscape Around Architecture," Joint Paper, The Open Group, OASIS, and
275 OMG, July 2009. [http://www.oasis-](http://www.oasis-open.org/committees/download.php/32911/wp_soa_harmonize_d1.pdf)
276 [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)

277 2 Architectural Goals and Principles

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

280 2.1 Goals and Critical Success Factors of the Reference 281 Architecture Foundation

282 There are three principal goals:

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

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

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

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

296 2.1.1 Goals

297 2.1.1.1 Effectiveness

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

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

304 2.1.1.2 Confidence

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

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

311 2.1.1.3 Scalability

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

316 2.1.2 Critical Success Factors

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

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

323 2.1.2.1 Action

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

327 2.1.2.2 Trust

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

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

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

339 2.1.2.3 Interaction

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

343 2.1.2.4 Control

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

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

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

355 2.2 Principles of this Reference Architecture Foundation

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

357 Technology Neutrality

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

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

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

372 Parsimony

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

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

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

379 Distinction of Concerns

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

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

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

399 Applicability

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

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

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

410 **3 Participation in a SOA Ecosystem View**

411
412
413
414
415
416
417
418
419
420
421

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

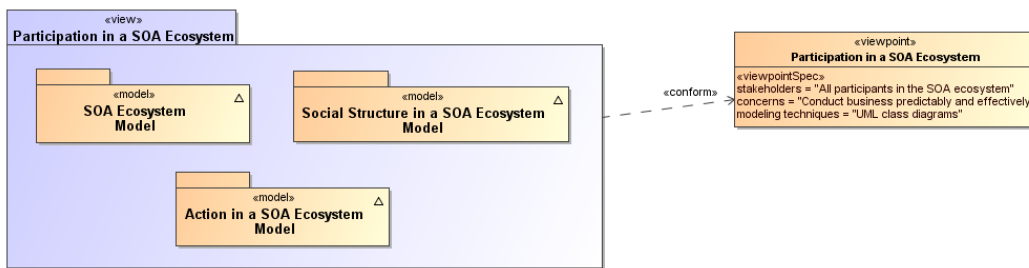
422 The *Participation in a SOA Ecosystem* view in the SOA-RAF focuses on the constraints and context in
423 which people conduct business using a SOA-based system. By business we mean any shared activity
424 whose objective is to satisfy particular **needs** of each participant. To effectively employ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

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

433 The main models in this view are:

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

Comment [PFB2]: Issue 32, part



Comment [PFB3]: Issue 309

443
444 *Figure 1 - Model elements described in the Participation in a SOA Ecosystem view*

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

448 The dominant mode of **communication** within a SOA ecosystem is electronic, supported by IT resources
449 and artifacts. The **stakeholders** (see next section) are nonetheless people: since there is inherent
450 indirection involved when people and systems interact using electronic means, we lay the foundations for

451 how *communication* can be used to represent and enable action. However, it is important to understand
452 that these communications are usually a means to an end and not the primary interest of the participants
453 of the ecosystem.

Comment [PFB4]: Issue 32, part


454 3.1 SOA Ecosystem Model

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

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

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

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



Comment [PFB5]: Issue 32, part

475
476 Figure 2 - SOA Ecosystem Model

477 SOA Ecosystem

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

480 SOA-based System

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

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

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

484 Participants (as stakeholders and as actors), SOA-based systems, and the environment (or context)
485 within which they all operate, taken together forms the SOA ecosystem. Participants (or their **delegates**)
486 interact with a SOA-based system - in the role of actors - and are also members of a social structure - in
487 the role of stakeholders. Here we explicitly note that stakeholders and, thus, participants are people³
488 because machines alone cannot truly have a stake in the outcomes of a social structure. Delegates may
489 be human and nonhuman but are not directly stakeholders. Stakeholders, both Participants and **Non-**

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)

490 **participants**, may potentially benefit from the services delivered by the SOA-based system. Again, this is
491 discussed more fully in Section 3.2.1.

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

498 Several interdependent concerns are important in our view of a SOA ecosystem. The ecosystem includes
499 stakeholders who are participants in the development, deployment and **governance** and use of a system
500 and its services; or who may not participate in certain activities but are nonetheless affected by the
501 system. Actors – whether stakeholder **participants** or delegates who act only on behalf of participants
502 (without themselves having any stake in the actions that they have been tasked to perform) – are
503 engaged in **actions** which have an impact on the real world and whose meaning and intent are
504 determined by implied or agreed-to semantics. This is discussed further in relation to the model in Figure
505 4 and elaborated more fully in Section 3.3.

506 **3.2 Social Structure in a SOA Ecosystem Model**

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

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

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

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

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

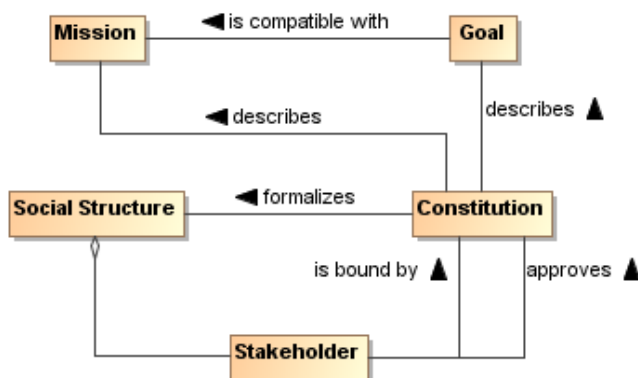


Figure 3 - Social Structure Model

532
533

534 **Social Structure**

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

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

Comment [PFB6]: Reworded Issues 28, 53

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

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

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

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

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

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

Comment [PFB7]: Issue 44, part

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

570 **Constitution**

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

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

Comment [KJL8]: Issue 31 for edits in this paragraph

579 **3.2.1 Stakeholders, Participants, Actors and Delegates**

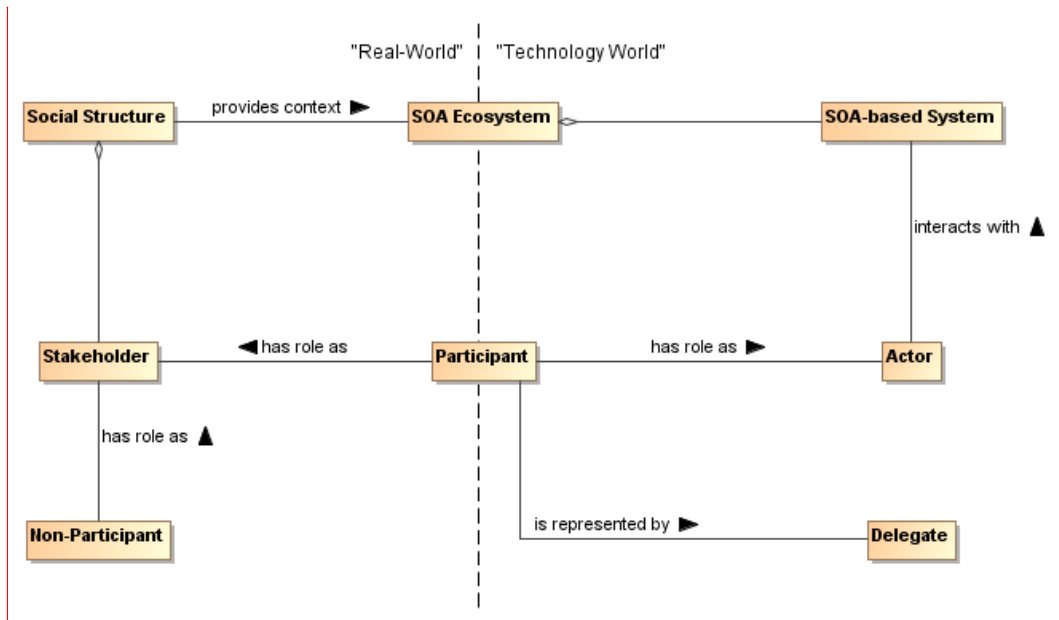
580 A social structure represents the interests of a collection of people who have **rights** and **responsibilities**
581 within the structure. People have a 'stake' in such a social structure, and when that social structure is part
582 of a SOA Ecosystem, the people continue to interact through their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

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

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

592 The concept of the ecosystem therefore embraces all aspects of the 'real world', human-centered, social
593 structures that are concerned with business interactions together with the technology-centered SOA-
594 based system that deliver services:

Comment [PFB9]: Issue 32, part; Issue 280, part



595
596 *Figure 4 – Stakeholders, Actors, Participants and Delegates*

597 **Stakeholder**

598 | A person with an interest (a 'stake') **in a social structure**.

599 Not all stakeholders necessarily participate in all activities in the SOA ecosystem; indeed, the interest of
600 non-participant stakeholders may be to realize the benefits of a well-functioning ecosystem and not suffer
601 unwanted consequences. Non-participant stakeholders cannot all or always be identified in advance but
602 due account is often taken of such stakeholder types, including potential customers, beneficiaries, and
603 other affected third parties. A stakeholder may be a participant with respect to some activities and a non-
604 participant with respect to others.

605 Actor

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

608 Participant

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

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

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

616 Non-Participant

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

619 Delegate

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

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

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

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

638 3.2.2 Social Structures and Roles

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


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

648 **3.2.2.1 Authority, Rights, and Responsibilities**

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

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

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



664
665 *Figure 5 - Social Structures, Roles and Action*

666 **Authority**

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

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

671 **Right**

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

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

676 **Responsibility**

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

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

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

684 3.2.2.2 Permissions and Obligations

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

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

695 **Permission**

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

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

701 **Obligation**

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

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

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

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

713 3.2.2.3 Service Roles

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

718 **Provider**

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

720 **Consumer**

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

722 **Mediator**

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

725 **Owner**

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

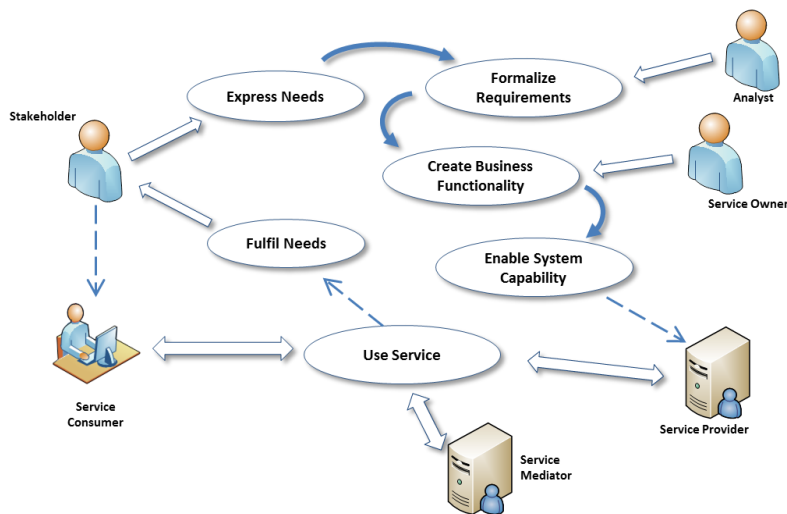


Figure 6 - Roles in a Service

727
728

729 Service consumers typically initiate interactions, but this is not necessarily true in all situations.
730 Additionally, several stakeholders may be involved in a service interaction supporting a given consumer.

731 The roles of service provider and service consumer are often seen as symmetrical, which is also not
732 entirely correct. A stakeholder tends to express a **Need** in non-formal terms: “I want to buy that book”.
733 The type of need that a service is intended to fulfill has to be formalized and encapsulated by designers
734 and developers as a **Requirement**. This Requirement should then be reflected in the target service, as a
735 **Capability** that, when accessed via a service, delivers a **Real World Effect** to an arbitrary consumer:
736 “The chosen book is ordered for the consumer.” It thus fulfills the need that has been defined for an
737 archetypal consumer.

738 Specific and particular customers may not experience a need exactly as captured by the service: “I don’t
739 want to pay that much for the book”, “I wanted an eBook version”, etc. There can therefore be a process
740 of implicit and explicit negotiation between the consumer and the service, aimed at finding a ‘best fit’
741 between the consumer’s specific need and the capabilities of the service that are available and consistent
742 with the service provider’s offering. This process may continue up until the point that the consumer is able
743 to accept what is on offer as being the best fit and finally ‘invokes’ the service. ‘Execution context’ has
744 thus been established. Conditions and agreements that contribute to the execution context are discussed
745 throughout this Reference Architecture.

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

750 3.2.3 Needs, Requirements and Capabilities

Comment [PFB10]: Moved from Action Model section

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

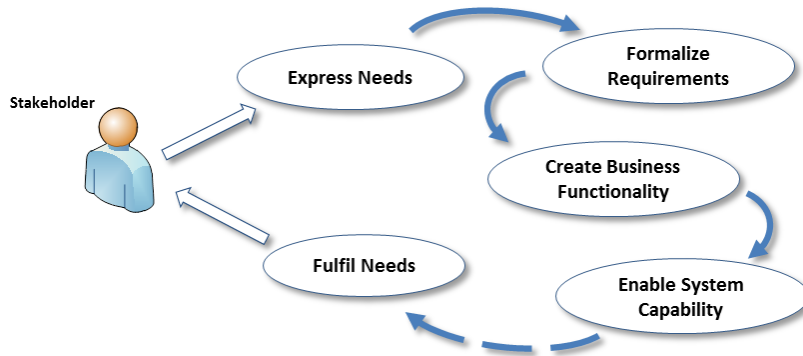
754 There is a reason that participants are engaged: they have different **needs** and have or apply different
755 capabilities for satisfying them. These are core to the concept of a service. The SOA-RM defines a
756 service as “the mechanism by which needs and capabilities are brought together”. This idea of services
757 being a mechanism ‘between’ needs and capabilities was introduced in order to emphasize capability as
758 the notional or existing **business functionality** that would address a well-defined need. Service is
759 therefore the *implementation* of such business functionality *such that it is accessible* through a well-

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

762 Business Functionality

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

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



770
771 Figure 7 - Cycle of Needs, Requirements, and Fulfillment

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

780 Need

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

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

784 Requirement

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

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

789 Capability

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

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

794 Real World Effect

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

Comment [PFB11]: Issues 56 and 57
– text moved from Action section

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

Comment [KL12]: Moved to complete Issues 56 & 57

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

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

808 3.2.4 Resource and Ownership

809 3.2.4.1 Resource

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

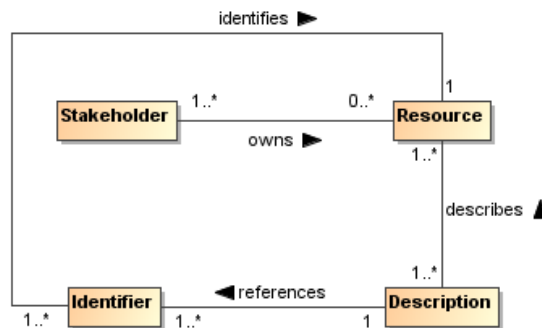


Figure 8 - Resources

812
813

814 Resource

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

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

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

822 Identifier

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

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

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

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

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

841 [Resources frequently have descriptions and the descriptions themselves may be considered resources.](#)
842 [This is discussed in Section 4.1.1. Resource description may link to other resources and their](#)
843 [descriptions: for example, a service description may link to a policy that constrains the conditions of use](#)
844 [of the service.](#)

Comment [KJL13]: Issue 307

845 3.2.4.2 Ownership

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

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

850 Ownership

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

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

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

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

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

872 Ownership Boundary

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


876 3.2.5 Establishing Execution Context

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

Comment [PFB14]: 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

882 **3.2.5.1 Trust and Risk**

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



884
885 *Figure 9 - Willingness and Trust*

886 **Willingness**

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

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

893 **Trust**

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

896 **Risk**

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

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

Comment [PFB15]: Issue 44, part

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

908

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

Comment [PFB16]: Issue 44, part

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

Comment [PFB17]: Issue 44, part

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


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

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

Comment [PFB18]: Issue 44, part

930 3.2.5.2 Policies and Contracts

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



936
937 Figure 10 – Policies, Contracts and Constraints

Comment [PFB19]: New diagram

938 Policy

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

Comment [PFB20]: Issue 282

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

946 **Contract**

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

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

Comment [PFB21]: Issue 46

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

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

963 **3.2.5.3 Communication**

964 **Communication**

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

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

Comment [PFB22]: Reworded
Issue 48

Comment [PFB23]: Issue 247

971 A communication may involve any number of recipients. In some situations, the sender may not be aware
972 of the recipient. However, without both a sender and a recipient, there is no communication. A given
973 communication can be a simple one-way transmission and not require a response by the recipient.
974 However, interaction does, necessarily, involve communication.

Comment [KJL24]: Reworded
Issue 50

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

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

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

Comment [PFB25]: Issue 48

985 **3.2.5.4 Semantics and Semantic Engagement**

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

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

995 Understanding and interpreting those assertions in a SOA-based system allows all the actors in any
996 particular joint action to 'know' what may be expected of them. An actor can potentially 'understand' an
997 assertion in a number of ways, but it is specifically the process of arriving at a *shared* understanding that
998 is important in the ecosystem. This process is semantic engagement and it takes place in different forms
999 throughout the SOA ecosystem. It can be instantaneous or progressively achieved. Participants – who
1000 play the role both as actors in the SOA-based system and as stakeholders in social structures and the
1001 wider ecosystem – can be pivotal in resolving problems of understanding and determining when there is a
1002 level of engagement appropriate and sufficient to the particular context.

1003 **Semantic Engagement**

1004 The process by which an **actor** engages with a set of assertions based on that actor's
1005 interpretation and understanding of those assertions.

1006 Different actors have differing capabilities and requirements for understanding assertions. This is true for
1007 both human and non-human actors. For example, a purchase order process does not require that a
1008 message forwarding agent 'understand' the purchase order, but a processing agent does need to
1009 'understand' the purchase order in order to know what to do with the order once received.

1010 The impact of any assertion can only be fully understood in terms of specific social contexts that
1011 necessarily include the actors that are involved. For example, a policy statement that governs the actions
1012 relating to a particular resource may have a different impact or purpose for the participant that owns the
1013 resource than for the actor that is trying to access it: the former understands the purpose of the policy as
1014 a statement of enforcement - the latter understands it as a statement of constraint.

1015 **3.3 Action in a SOA Ecosystem Model**

1016 Participants cannot always achieve desired results by leveraging resources in their own ownership
1017 domain. This unfulfilled need leads them to seek and leverage services provided by other participants and
1018 using resources beyond their ownership and control. The participants identify service providers with which
1019 they think they can interact to achieve their objective and engage in joint action with those other actors
1020 (service providers) in order to bring about the desired outcome. The SOA ecosystem provides the
1021 environment in which this happens.

1022 An action model is put forth a-priori by the service provider, and is effectively an undertaking by the
1023 service provider that the actions – identified in the action model and invoked consistent with the process
1024 model – will result in the described real world effect. The action model describes the actions leading to a
1025 real-world effect. A potential service consumer – who is interested in a particular outcome to satisfy their
1026 need – must understand those actions as capable of achieving that desired outcome.

1027 When the consumer 'invokes' a service, a joint action is started as identified in the action model,
1028 consistent with the temporal sequence as defined by the process model, and where the consumer and
1029 the provider are the two parties of the joint action. Additionally, the consumer can be assured that the
1030 identified real-world effects will be accomplished through evidence provided via the service description.

1031 Since the service provider does not know about all potential service consumers, the service provider may
1032 also describe what additional constraints are necessary in order for the service consumer to invoke
1033 particular actions, and thus participate in the joint action. These additional constraints, along with others
1034 that might not be listed, are preconditions for the joint action to occur and/or continue (as per the process
1035 model), and are referred to in the SOA-RM as execution context. Execution context goes all the way from
1036 human beings involved in aligning policies, semantics, network connectivity and communication
1037 protocols, to the automated negotiation of security protocols and end-points as the individual actions
1038 proceed through the process model.

1039 Also, it is important to note that both actions and real world effect are recursive in nature, in the sense
1040 that they can often be broken down into more and more granularity depending on how they are examined
1041 and what level of detail is important.

1042 All of these things are important to getting to the core of participants' concern in a SOA ecosystem: the
1043 ability to leverage resources or capabilities to achieve a desired outcome, and in particular where those
1044 resources or capabilities do not belong to them or are beyond their direct control. i.e., that are outside of
1045 their ownership boundary.

1046 | In order to use such resources, participants must be able to identify their own needs; **state those needs** in
1047 | the form of requirements; compose **or identify a suitable** business solution **using** resources or capabilities
1048 | that will meet their needs; and engage in joint action – the coordinated set of actions that participants
1049 | pursue in order to achieve measurable results in furtherance of their goals.

Comment [PFB26]: Issue 53, part

1050 | In order to act in a way that is appropriate and consistent, participants must communicate with each other
1051 | about their own goals, objectives and policies, and those of others. This is the main concern of Semantic
1052 | Engagement.

1053 | A key aspect of joint action revolves around the trust that both parties must exhibit in order to participate
1054 | in the joint action. The willingness to act and a mutual understanding of both the information exchanged
1055 | and the expected results is the particular focus of Sections 3.2.5.1 and 3.2.5.4.

1056 | **3.3.1 Services Reflecting Business**

1057 | The SOA paradigm often emphasizes the interface through which service interaction is accomplished.
1058 | While this enables predictable integration in the sense of traditional software development, the prescribed
1059 | interface alone does not guarantee that services will be composable into business solutions.

1060 | **Business Solution**

1061 | A set of defined interactions that combine implemented or notional **business functionality** in
1062 | order to address a set of business needs.

1063 | **Composability**

1064 | The ability to combine individual services, each providing defined **business functionality**, so as
1065 | to provide more complex **business solutions**.

1066 | To achieve composability, capabilities must be identified that serve as building blocks for business
1067 | solutions. In a SOA ecosystem, these building blocks are captured as services representing well-defined
1068 | business functions, operating under well-defined policies and other constraints, and generating well-
1069 | defined real world effects. These service building blocks should be relatively stable so as not to force
1070 | repeated changes in the compositions that utilize them, but should also embody SOA attributes that
1071 | readily support creating compositions that can be varied to reflect changing circumstances.

1072 | The SOA paradigm emphasizes both composition of services and opacity of how a given service is
1073 | implemented. With respect to opacity, the SOA-RM states that the service could carry out its described
1074 | functionality through one or more automated and/or manual processes that in turn could invoke other
1075 | available services.

1076 | Any composition can itself be made available as a service and the details of the business functionality,
1077 | conditions of use, and effects are among the information documented in its service description.

1078 | Composability is important because many of the benefits of a SOA approach assume multiple uses for
1079 | services, and multiple use requires that the service deliver a business function that is reusable in multiple
1080 | business solutions. Simply providing a Web Service interface for an existing IT artifact does not, in
1081 | general, create opportunities for sharing business functions. Furthermore, the use of tools to auto-
1082 | generate service software interfaces will not guarantee services that can effectively be used within
1083 | compositions if the underlying code represents programming constructs rather than business functions. In
1084 | such cases, services that directly expose the software details will be as brittle to change as the underlying
1085 | code and will not exhibit the characteristic of loose coupling.

1086 | **3.3.2 Activity, Action, and Joint Action**

1087 | In general terms, entities act in order to **fulfill particular objectives**. More precisely, they generate activity.
1088 | An activity is made up of specific Actions (or other Activities) and is formally defined in **[ISO/IEC 10746]**
1089 | as “a single-headed directed acyclic graph of actions...”⁴ It is most clearly understood diagrammatically:

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




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

1090
1091
1092
1093
1094
1095
1096
1097
1098
1099

What constitutes an Action or an Activity will be a matter of context. For the SOA-RAF, an Action represents the smallest and most discrete activity that needs to be modeled for a given Viewpoint. The form of Activity that is of most interest within a SOA ecosystem is that involving Actions as defined below and their interaction across ownership boundaries (and thus involving interaction between more than one actor) – we call this **joint action**. In Figure 12 below, one line of activity (on the left) can be completed thru Action₃ without crossing any ownership boundary but the alternative path, starting at Action₄, can only be completed as a result of joint action across an ownership boundary:




Figure 12: Activity involving Actions across an ownership boundary

1100
1101
1102
1103
1104
1105
1106
1107
1108

Action

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

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

1109 In order for multiple actors to participate in a joint action, they must each act according to their role within
1110 the joint action. This is achieved through communication and messaging.

1111 Communication – the formulation, transmission, receipt and interpretation of messages – is the
1112 foundation of all joint actions within the SOA ecosystem, given the inherent separation – often across
1113 ownership boundaries – of actors in the system.

1114 Communication between actors requires that they play the roles of ‘sender’ or ‘receiver’ of messages as
1115 appropriate to a particular action – although it is not necessarily required that they both be active
1116 simultaneously.

1117 An actor sends a message in order to communicate with other actors. The communication itself is often
1118 not intended as part of the desired real world effect but rather includes messages that seek to establish,
1119 manage, monitor report on, and guide the joint action throughout its execution.

1120 Like communication, joint action usually involves different actors. However, joint action – resulting from
1121 the deliberate actions undertaken by different actors – *intentionally* impacts shared state within the
1122 system leading to real world effects.

1123 **Joint Action**

1124 The coordinated set of **actions** involving the efforts of two or more **actors** to achieve an effect.

1125 Note that the effect of a joint action is *not* always equivalent to one or more effects of the individual
1126 actions of the actors involved, i.e., it may be more than the sum of the parts.

1127 Different perspectives lead to either communication or joint action as being considered most important.
1128 For example, from the [viewpoint-perspective](#) of ecosystem security, the integrity of the communications
1129 may be dominant; from the [perspective viewpoint](#) of ecosystem governance, the integrity of the joint
1130 action may be dominant.

1131 **3.3.3 State and Shared State**

1132 **State**

1133 The condition of an entity at a particular time.

1134 State is characterized by a set of facts that is true of the entity. In principle, the total state of an entity (or
1135 the world as a whole) is unbounded. In practice, we are concerned only with a subset of the state of an
1136 entity that is measurable and useful in a given context.

1137 For example, the total state of a light bulb includes the temperature of the filament of the bulb, the
1138 composition of the glass, the dirt that is on the bulb’s surface and so on. However, [someone needing](#)
1139 [more light to read by-is only really](#) interested in whether the bulb is ‘on’ or ‘off’ [and if it is working properly](#).
1140 That [individual](#)’s characterization of the state of the bulb reduces to the fact: “bulb is now on”.

1141 In a SOA ecosystem, there is a distinction between the set of facts about an entity that only that entity can
1142 access and the set of facts that may be accessible to others, notably actors in the SOA-based system.

1143 **Private State**

1144 That part of an entity’s **state** that is knowable by, and accessible to, only that entity.

1145 **Shared State**

1146 That part of an entity’s **state** that is knowable by, and may be accessible to, other actors.

1147 Note that shared state does not imply that the state *is* accessible to other actors. It simply refers to that
1148 subset of state that *may* be accessed by other actors. This will principally be the case when actors need
1149 to participate in joint actions.

1150 It is the aggregation of the shared states of pertinent entities that constitutes the desired effect of a joint
1151 action. Thus the change to this shared state is what is experienced in the wider ecosystem as a real world
1152 effect

1153 3.4 Architectural Implications

1154 3.4.1 Social structures

1155 A SOA ecosystem's participants are organized into various forms of social structure. Not all social
1156 structures are hierarchical: a SOA ecosystem **SHOULD** be able to incorporate peer-to-peer forms of
1157 organization as well as hierarchic structures. In addition, it **SHOULD** be possible to identify and manage
1158 any constitutional agreements that define the social structures present in a SOA ecosystem.

- 1159 • Different social structures have different rules of engagement but predictable behavior is one of
1160 the underpinnings of trust. ~~This therefore requires mechanisms to~~ **Mechanisms MUST therefore**
1161 **be available to:**
 - 1162 ○ express constitutions and other organizing principles of participants;
 - 1163 ○ inherit rules of engagement from parent to child social structures.
- 1164 • Social structures have roles and members and this impacts who may be authorized to act and in
1165 what circumstances. ~~This requires mechanisms to~~ **Mechanisms MUST be available to:**
 - 1166 ○ identify and manage members of social structures
 - 1167 ○ Identify and manage attributes of the members
 - 1168 ○ describe roles and role adoption
- 1169 • Social structures overlap and interact, giving rise to situations in which rules of engagement may
1170 conflict. In addition, a given actor may be a member of multiple social structures and the social
1171 structures may be associated with different jurisdictions. ~~This requires mechanisms~~
1172 **to** **Mechanisms MUST be available to:**
 - 1173 ○ identify the social structures that are active during a series of joint actions;
 - 1174 ○ identify and resolve conflicts and inconsistencies.

1175 3.4.2 Resource and Ownership

1176 Communication about and between, visibility into, and leveraging of resources requires the unambiguous
1177 identification of those resources. ~~Ensuring unambiguous identities implies~~ **Mechanisms MUST be**
1178 **available for:**

- 1179 • ~~Mechanism for a~~ Assigning and guaranteeing uniqueness of globally unique identifiers
- 1180 • Identifying the extent of the enterprise over which the identifier needs to be understandable and
1181 unique
- 1182 • ~~Mechanism and framework for e~~ Ensuring the longevity of identifiers (i.e., they cannot just change
1183 arbitrarily)

1184 3.4.3 Policies and Contracts

- 1185 • Policies are expressed as constraints:
 - 1186 ○ Policies **MUST** be expressed
 - 1187 ○ Constraints **MUST** be enforceable
 - 1188 ○ Management of potentially large numbers of policies **MUST** be achievable
- 1189 • Policies have owners:
 - 1190 ○ Policies **SHOULD** be established by social structures.
- 1191 • Policies may not be consistent with one another:
 - 1192 ○ Policy conflict resolution techniques **MUST** exist and be in place
- 1193 • Agreements are accepted constraints:
 - 1194 ○ Contracts **SHOULD** be enforced by mechanisms of the social structure

1195 3.4.4 Communications as a Means of Mediating Action

1196 Using message exchange for mediating action implies

- 1197 • ~~Ensuring correct identification of t~~ The structure of messages **MUST be validated by:**
 - 1198 ○ Identifying the syntax of the message;
 - 1199 ○ Identifying the vocabularies used in the communication

- 1200 ○ Identifying the higher-level structure of the communication, such as policy assertion,
1201 contract enforcement, etc.
- 1202 • A principal objective of communication is to mediate action, therefore:
- 1203 ○ Messages **SHOULD** convey actions and events
- 1204 ○ Receiving a message is an action, but is not the same action as the action conveyed by
1205 the message
- 1206 ○ Actions are associated with objectives of the actors involved
- 1207 ▪ Explicit representation of objectives may facilitate automated processing of
1208 messages
- 1209 ○ An actor agreeing to adopt an objective becomes responsible for that objective

Comment [PFB27]: Is this section still valid?

1210 3.4.5 Semantics

1211 Semantics is pervasive in a SOA ecosystem. There are many forms of utterance that are relevant to the
1212 ecosystem: apart from communicated content there are mission and policy statements, goals, objectives,
1213 descriptions, and agreements which are all forms of utterance.

1214 The operation of the SOA ecosystem is significantly enhanced if

- 1215 • A careful distinction is made between public semantics and private semantics. In particular, it
1216 **MUST** be possible for actors to process content such as communications, descriptions and
1217 policies solely on the basis of the public semantics of those utterances.
- 1218 • A well founded semantics **MUST** ensure that any assertions ~~that are~~ essential to the operator of
1219 the ecosystem (such as policy statements, and descriptions) have carefully chosen written
1220 expressions and associated decision procedures.
- 1221 • The role of vocabularies as a focal point for multiple actors to be able to understand each other is
1222 critical. While no two actors can fully share their interpretation of elements of vocabularies,
1223 ~~ensuring that they do~~ **SHOULD** be able to understand the intended public meaning of
1224 vocabularies' elements ~~is essential~~.

1225 3.4.6 Trust and Risk

1226 In traditional systems, the balance between trust and risk is achieved by severely restricting interactions
1227 and by controlling the participants of a system.

1228 ~~It is important that actors~~ Actors **MUST** be able to explicitly reason about both trust and risk in order to
1229 effectively participate in a SOA ecosystem. The more open and public the SOA ecosystem is, the more
1230 important it is for actors to be able to reason about their participation.

1231 3.4.7 Needs, Requirements and Capabilities

1232 In the process of capturing needs as requirements, and the subsequent requirements decomposition and
1233 allocation processes need to be informed by capabilities that already exist.

- 1234 • Architecture ~~needs to~~ **MUST** take into account existing capabilities available as services

1235 3.4.8 The Importance of Action

1236 Participants participate in a SOA ecosystem in order to get their needs met. This involves action; both
1237 individual actions and joint actions.

1238 Any architectural realization of a SOA ecosystem ~~should~~ **SHOULD** address:

- 1239 • How actions are modeled:
 - 1240 ○ Identifying the performer or agent of the action;
 - 1241 ○ the target of the action; and the
 - 1242 ○ verb of the action.

1243 Any explicit models of joint action ~~should~~ **SHOULD** take into account

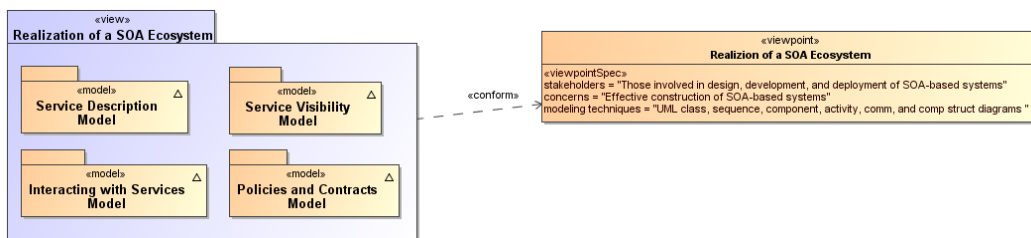
- 1244 • The ~~choreography or orchestration possible compositions~~ that defines the joint action.
- 1245 • The potential for multiple joint actions to be layered on top of each other

1246 **4 Realization of a SOA Ecosystem view**

1247 *Make everything as simple as possible but no simpler.*
1248 Albert Einstein

1249
1250 The *Realization of a SOA Ecosystem* view focuses on elements that are needed to support the discovery
1251 of and interaction with services. The key questions asked are "What are services, what support is needed
1252 and how are they realized?"

1253 The models in this view include the Service Description Model, the Service Visibility Model, the Interacting
1254 with Services Model, and the Policies and Contracts Model.



1255
1256 *Figure 13 - Model Elements Described in the Realization of a SOA Ecosystem view*

1257 The Service Description Model informs the participants of what services exist and the conditions under
1258 which they can be used. [The Policies and Contracts Model elaborates on the conditions under which](#)
1259 [service use is prescribed and agreements among participants in the SOA ecosystem.](#) The information in
1260 the service description as augmented by details of policy provides the basis for visibility as defined in the
1261 SOA Reference Model and captured in the Service Visibility Model. Finally, the process by which services
1262 are used under the defined conditions and agreements is described in the Interacting with Services
1263 Model.

Comment [KJL28]: Issue 168

1264 **4.1 Service Description Model**

1265 A service description is an artifact, often document-based, that defines or references the information
1266 needed to use, deploy, manage and otherwise control a service. This includes not only the information
1267 and behavior models associated with a service that define [interaction via](#) the service interface but also
1268 includes information needed to decide whether the service is appropriate for the current needs of the
1269 service consumer. Thus, the service description should also include information such as service
1270 reachability, service functionality, and the policies associated with a service.

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

1271 A service description artifact may be a single document or it may be an interlinked set of documents. For
1272 the purposes of this model, differences in representation are to be ignored, but the implications of a 'web
1273 of documents' are discussed later in this section.

1274 There are several points to note regarding service description:

- 1275 • The Reference Model states that one of the hallmarks of SOA is the large amount of associated
1276 description. The model presented below focuses on the description of services but it is equally
1277 important to consider the descriptions of the consumer, other participants, and needed resources
1278 other than services.
- 1279 • Descriptions are inherently incomplete but may be determined as *sufficient* when it is possible for
1280 the participants to access and use the described services based only on the descriptions
1281 provided. This means that, at one end of the spectrum, a description along the lines of "That
1282 service on that machine" may be sufficient for the intended audience. On the other extreme, a
1283 service description with a machine-process-able description of the semantics of its **operations**
1284 and real world effects may be required for services accessed via automated service discovery
1285 and planning systems.

- 1286
- 1287
- 1288
- 1289
- 1290
- 1291
- 1292
- 1293
- 1294
- 1295
- 1296
- 1297
- 1298
- 1299
- 1300
- 1301
- 1302
- 1303
- 1304
- 1305
- Descriptions come with context, i.e. a given description comprises information needed to adequately support the context. For example, a list of items can define a version of a service, but for many contexts an indicated version number is sufficient without the detailed list. The current model focuses on the description needed by a service consumer to understand what the service does, under what conditions the service will do it, how well the service does it, and what steps are needed by the consumer to initiate and complete a service interaction. Such information also enables the service provider to clearly specify what is being provided and the intended conditions of use.
 - Descriptions change over time as, for example, the ingredients and nutrition information for food labeling continues to evolve. A requirement for transparency of transactions may require additional description for those associated contexts.
 - Description always proceeds from a basis of what is considered 'common knowledge'. This may be social conventions that are commonly expected or possibly codified in law. It is impossible to describe everything and it can be expected that a mechanism as far reaching as SOA will also connect entities where there is inconsistent 'common' knowledge.
 - Descriptions become the collection point of information related to a service or any other resource, but it is not necessarily the originating point or the motivation for generating this information. In particular, given a SOA service as the access to an underlying capability, the service may point to some of the capability's previously generated description, e.g. a service providing access to a data store may also have access to information indicating the freshness of the data.

1306 These points emphasize that there is no one 'right' description for all contexts and for all time. Several
1307 descriptions for the same subject may exist at the same time, and this emphasizes the importance of the
1308 description referencing source material maintained by that material's owner rather than having multiple
1309 copies that become out of synch and inconsistent.

1310 It may also prove useful for a description assembled for one context to cross-reference description
1311 assembled for another context as a way of referencing ancillary information without overburdening any
1312 single description. Rather than a single artifact, description can be thought of as a web of documents that
1313 enhance the total available description.

1314 This Reference Architecture Foundation uses the term service description for consistency with the
1315 concept defined in the Reference Model. Some SOA literature treats the idea of a 'service contract' as
1316 equivalent to service description. In the SOA-RAF, the term service description is preferred. Replacing the
1317 term 'service description' with the term 'service contract' implies that just one side of the interaction is
1318 governing and misses the point that a single set of policies identified by a service description may lead to
1319 numerous contracts, i.e. service level agreements, leveraging the same description.

1320 4.1.1 The Model for Service Description

1321 Figure 14 shows Service Description as a subclass of the general Description class. *As well as describing*
1322 *a Resource (as we saw in Section 3.2.4.1), a Description is also a subclass of the Resource class.* In
1323 addition, each resource is assumed to *have* a description⁵. The following section discusses the
1324 relationships among elements of general description and the subsequent sections focus on service
1325 description. Other descriptions, such as those of participants, are important to SOA but are not
1326 individually elaborated in this document.

Comment [PFB30]: Issue 307

Comment [PFB31]: Issue 290, part

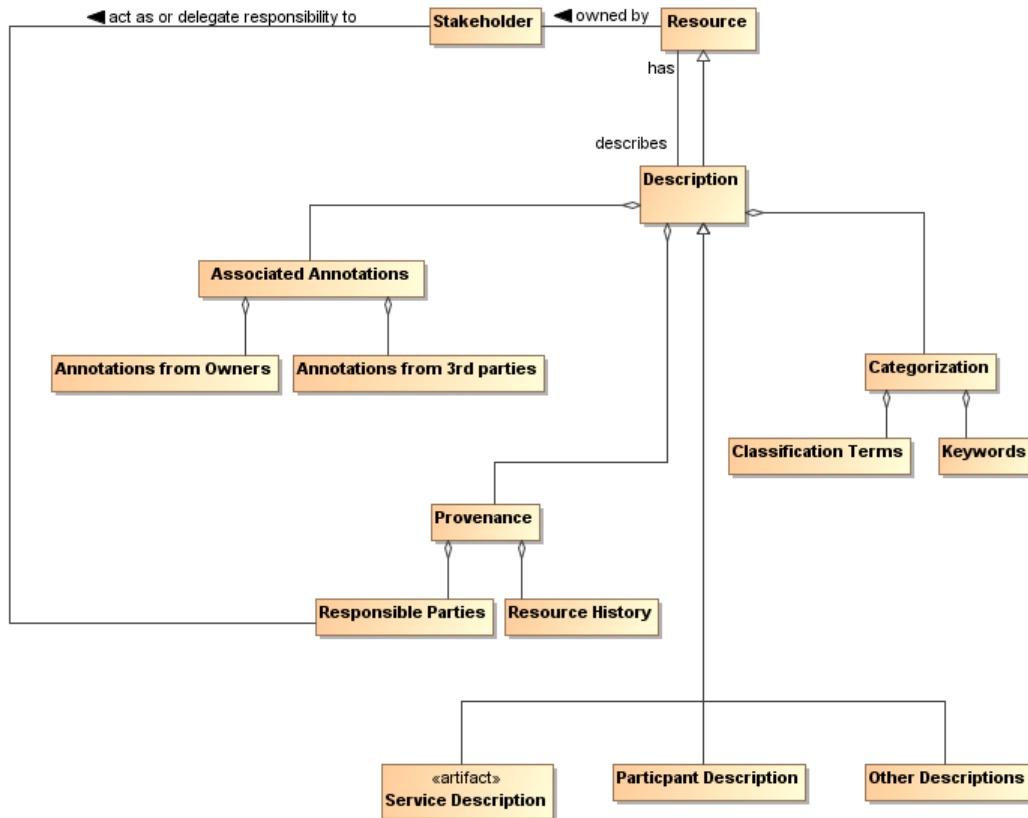
1327 4.1.1.1 Elements Common to General Description

1328 The general Description class is composed of a number of elements that are expected to be common
1329 among all descriptions supporting a service-oriented architecture. A registry/*repository* often contains a
1330 subset of the description instance, where the chosen subset is identified as that which facilitates
1331 discovery. Additional information contained in a more complete description may be needed to initiate and
1332 continue interaction.

Comment [KJL32]: Issue 173

⁵ The description itself can have further descriptive data such as its version or last revision. The model emphasizes this point but should not be interpreted too rigorously as allowing endless recursion.

1333



Comment [PFB33]: Issue 290, part

1334
1335

Figure 14 - General Description

1336 4.1.1.1.1 Provenance

1337 While the resource Identifier provides the means to know which subject and subject description are being
 1338 considered, Provenance as related to the Description class provides information that reflects on the
 1339 quality or usability of the subject. Provenance specifically identifies the stakeholder (human, defined role,
 1340 organization, etc.) that assumes responsibility for the resource being described and tracks historic
 1341 information that establishes a context for understanding what the resource provides and how it has
 1342 changed over time. Responsibilities may be directly assumed by the stakeholder who owns a resource
 1343 (see Section 3.2.4.2) or the Owner may designate Responsible Parties for the various aspects of
 1344 maintaining the resource and provisioning it for use by others. There may be more than one stakeholder
 1345 identified under Responsible Parties; for example, one stakeholder may be responsible for code
 1346 maintenance while another is responsible for provisioning of the executable code.

1347 4.1.1.1.2 Keywords and Classification Terms

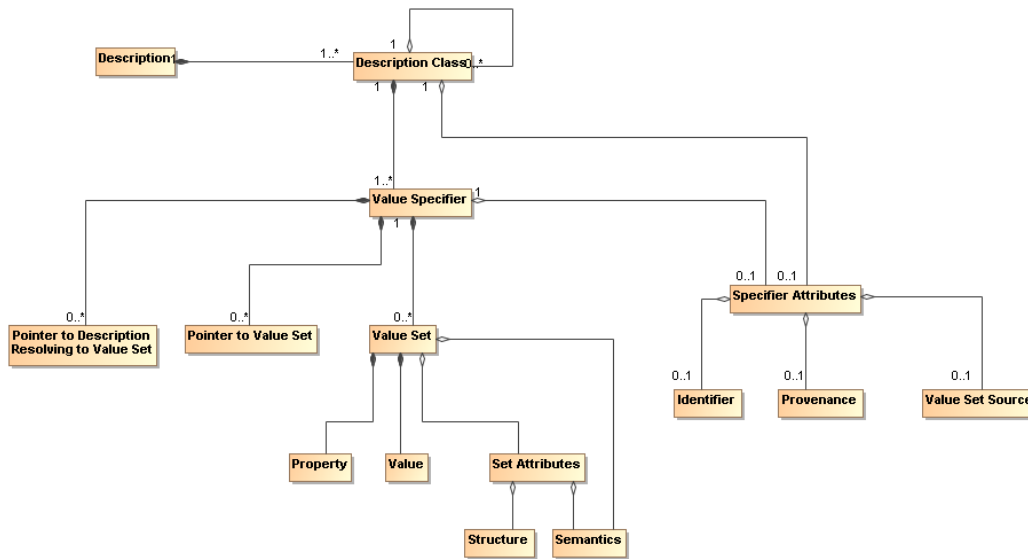
1348 A traditional element of description has been to associate the resource being described with predefined
 1349 keywords or classification taxonomies that derive from referenceable formal definitions and vocabularies.
 1350 This Reference Architecture Foundation does not prescribe which vocabularies or taxonomies may be
 1351 referenced, nor does it limit the number of keywords or classifications that may be associated with the
 1352 resource. It does, however, state that a normative definition of any terms or keywords SHOULD be
 1353 referenced, whether that be a representation in a formal ontology language, a pointer to an online

1354 dictionary, or any other accessible source. See Section 4.1.1.2 for further discussion on associating
1355 semantics with assigned values.

1356 4.1.1.1.3 Associated Annotations

1357 The general description instance may also reference associated documentation that is in addition to that
1358 considered necessary in this model. For example, the owner of a service may have documentation on
1359 best practices for using the service. Alternately, a third party may certify a service based on their own
1360 criteria and certification process; this may be vital information to other prospective consumers if they were
1361 willing to accept the certification in lieu of having to perform another certification themselves. Note, while
1362 the examples of Associated Documentation presented here are related to services, the concept applies
1363 equally to description of other entities.

1364 4.1.1.2 Assigning Values to Description Instances



1365
1366

Figure 15 - Representation of a Description

1367 Figure 14 shows the template for a general description, but individual description instances depend on
1368 the ability to associate meaningful values with the identified elements. Figure 15 shows a model for a
1369 collection of information that provides for value assignment and traceability for both the meaning and the
1370 source of a value. The model is not meant to replace existing or future schema or other structures that
1371 have or will be defined for specific implementations, but it is meant as guidance for the information such
1372 structures need to capture to generate sufficient description. It is expected that tools will be developed to
1373 assist the user in populating description and auto-filling many of these fields, and in that context, this
1374 model provides guidance to the tool developers.

1375 In Figure 15, each class has an associated value specifier or is made up of components that eventually
1376 resolve to a value specifier. For example, Description has several components, one of which is
1377 Categorization, which would have an associated value specifier.

1378 A value specifier consists of

- 1379 • a collection of value sets with associated property-value pairs, pointers to such value sets, or
- 1380 pointers to descriptions that eventually resolve to value sets that describe the component; and
- 1381 • attributes that qualify the value specifier and the value sets it contains.

1382 The qualifying attributes for the value specifier include

- 1383
- 1384
- 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 [PFB34]: Issue 291

1389 If the value specifier is contained within a higher-level component (such as Service Description containing
1390 Service Functionality), the component may assume values from the attributes of its container.

1391 Note, provenance as a qualifying attribute of a value specifier is different from provenance as part of an
1392 instance of Description. Provenance for a service identifies those who own and are responsible for the
1393 service, as described in Section 3.2.4. Provenance for a value specifier identifies who is responsible for
1394 choosing and assigning values to the value sets that comprise the value specifier. It is assumed that
1395 granularity at the value specifier level is sufficient and provenance is not required for each value set.

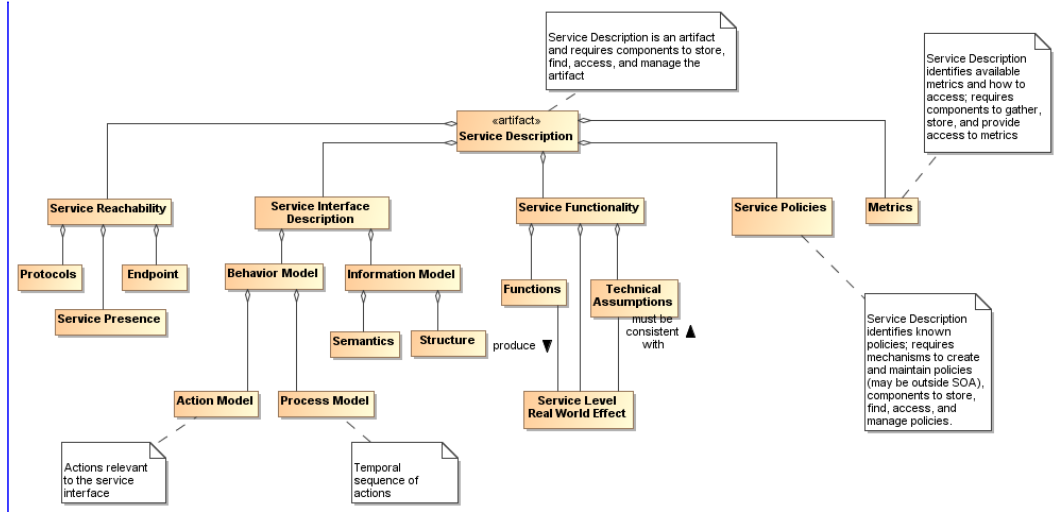
1396 The value set also has attributes that define its structure and semantics.

- 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.

1407 The value specifier may indicate a default semantic model for its component value sets and the individual
1408 value sets may provide an override.

1409 The property-value pair construct is introduced for the value set to emphasize the need to identify
1410 unambiguously both what is being specified and what is a consistent associated value. The further
1411 qualifying of Structure and Semantics in the Set Attributes allows for flexibility in defining the form of the
1412 associated values.

1413 **4.1.1.3 Model Elements Specific to Service Description**



1414 *Figure 16 - Service Description*

Comment [KL35]: Issues 66, part, and 176

1415

1416 The major elements for the Service Description subclass follow directly from the areas discussed in the
 1417 Reference Model. Here, we discuss the detail shown in Figure 16 and the purpose served by each
 1418 element of service description. For example, Service Policies as included in Figure 16 indicate those
 1419 policies that affect conditions of use of the service; however, while the description may link to detailed
 1420 policy documents, it is not the purpose of description to justify or elaborate on the rationale for the
 1421 policies. Similarly, Service Interface Description as included in Figure 16 captures information about what
 1422 interactions are supported by the service via its Behavior Model and the information exchange needed to
 1423 carry out those interactions in accordance with the service's Information Model; it is not the coded
 1424 interface.

1425 Note, the intent in the subsections that follow is to describe how a particular element, such as the service
 1426 interface description, is reflected in the service description, not to elaborate on the details of that element.

1427 **4.1.1.3.1 Service Interface Description**

1428 As noted in the Reference Model, the service interface is the means for interacting with a service. For the
 1429 SOA-RAF and as shown in Section 4.3 the service interface supports an exchange of messages, where

- 1430 • the message conforms to a referenceable message exchange pattern (MEP, covered below in
- 1431 Section 4.3.3.1),
- 1432 • the message payload conforms to the structure and semantics of the indicated information model,
- 1433 • the messages are used to denote events related to or actions against the service, where the
- 1434 actions are specified in the action model and any required sequencing of actions is specified in
- 1435 the process model.

1436 The Service Interface Description element as shown in Figure 17 includes the information needed to carry
 1437 out this message exchange in order to realize the service behavior described. In addition to the
 1438 Information Model that conveys the Semantics and Structure of the message, the Service Interface
 1439 Description indicates what behavior can be expected through interactions conveyed in the Action and
 1440 Process Models.

Comment [KL36]: Issue 176, part

Comment [KL37]: Issues 176, part; and 292

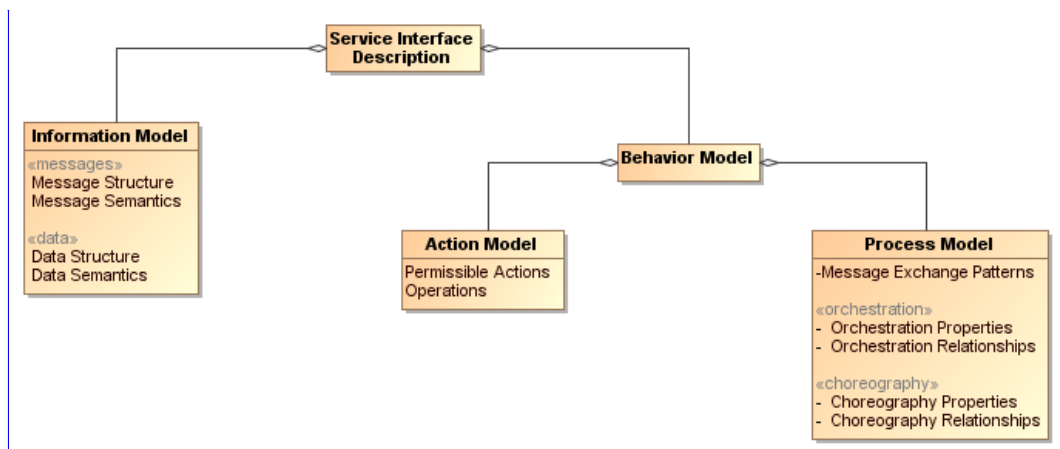


Figure 17 - Service Interface Description

1441
1442

1443 Note we distinguish the structure and semantics of the message from that of the underlying **protocol** that
1444 conveys the message. The message structure may include nested structures that are independently
1445 defined, such as an enclosing envelope structure and an enclosed data structure.

1446 These aspects of messages are discussed in more detail in Section 4.3.2.

1447 **4.1.1.3.2 Service Reachability**

1448 Service reachability, as modeled in Section 4.2.2.3 enables service participants to locate and interact with
1449 one another. To support service reachability, the service description should indicate the **endpoints** (also
1450 modeled and defined in that section) to which a service consumer can direct messages to invoke actions
1451 and the protocol to be used for message exchange using that endpoint.

1452 As generally applied to an action, the endpoint is the conceptual location where one applies an action;
1453 with respect to service description, it is the actual address where a message is sent.

1454 **4.1.1.3.3 Service Functionality**

1455 While the service interface and service reachability are concerned with the mechanics of using a service,
1456 service functionality and performance metrics (discussed in Section 4.1.1.3.4) describe what can be
1457 expected as a result of interacting with a service. Service Functionality, shown in Figure 16 as part of the
1458 overall Service Description model and extended in Figure 18, is a clear expression of service function(s)
1459 and the real world effects of invoking the function. The Functions represent business activities in some
1460 domain that produce the desired real world effects.

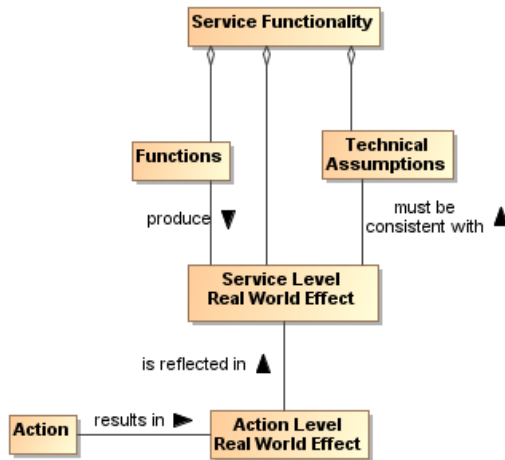


Figure 18 - Service Functionality

1461
1462

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

In Figure 16 and Figure 18, we specifically refer to [the descriptions of Service Level and Action Level Real World Effects](#).

Comment [KL38]: Issue 176

1471 Service Level Real World Effect

1472 A specific change in the **state** or the information returned as a result of interacting with a service.

1473 Action Level Real World Effect

1474 A specific change in the **state** or the information returned as a result of interacting through a
1475 specific action.

1476 Service description describes the service as a whole while the component aspects should contribute to
1477 that whole. Thus, while individual Actions may contribute to the real world effects to be realized from
1478 interaction with the service, there would be a serious disconnect for Actions to contribute real world
1479 effects that could not consistently be reflected in the Service Level Real World Effects and thus the
1480 Service Functionality. The relationship to Action Level Real World Effects and the implications on defining
1481 the scope of a service are discussed in Section 4.1.2.1.

1482 Elements of Service Functionality may be expressed as natural language text, reference an existing
1483 taxonomy of functions or other formal model.

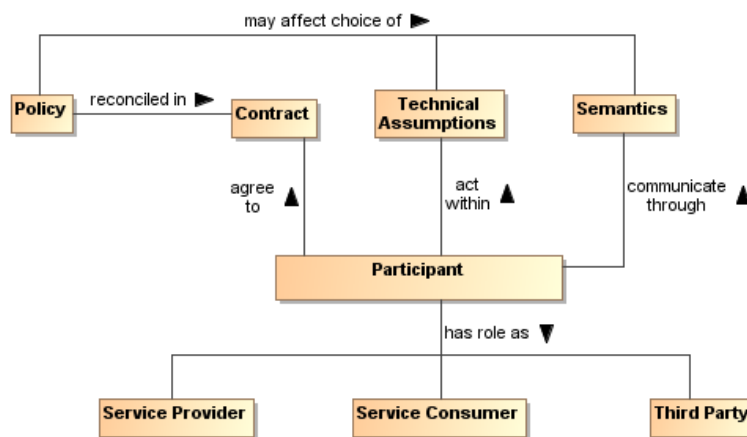
1484 4.1.1.3.4 Service Policies, Metrics, and Compliance Records

1485 Policies prescribe the conditions and constraints for interacting with a service and impact the willingness
1486 to continue visibility with the other participants. Whereas technical constraints are statements of 'physical'
1487 fact, policies are subjective assertions made by the service provider (sometimes as passed on from
1488 higher authorities).

1489 The service description provides a central location for identifying what policies have been asserted by the
1490 service provider. The specific representation of the policy, e.g. in some formal policy language, is outside
1491 of the service description. The service description would reference the normative definition of the policy.

1492 Policies may also be asserted by other [service](#) participants, as illustrated by the model shown in Figure
1493 19. Policies that are generally applicable to any interaction with the service are asserted by the service
1494 provider and included in the Service Policies section of the service description.

Comment [KL39]: Issue 179, part



Comment [PFB40]: Issue 179, part

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

1495
1496

1497 In Figure 19, we specifically refer to policies at the service level. In a similar manner to that discussed for
1498 Service Level vs. Action Level Real World Effects in Section 4.1.1.3.3, individual Actions may have
1499 associated policies stating conditions for performing the action, but these must be reflected in and be
1500 consistent with the policies made visible at the service level and thus the description of the service as a
1501 whole. The relationship to Action Level Policies and the implications on defining the scope of a service
1502 are discussed in Section 4.1.2.1.

1503 As noted in Figure 19, the policies asserted may be reflected as Technical Assumptions/Constraints that
1504 available services or their underlying capabilities must be capable of meeting; it may similarly affect the
1505 semantics that can be used. For example of the former, there may be a policy that specifies the surge
1506 capacity to be accommodated by a server, but a service that is not designed to make use of the larger
1507 server capacity would not satisfy the intent of the policy and would not be appropriate to use. For the
1508 latter, a policy may require that only services that support interaction via a community-sponsored
1509 vocabulary can be used.

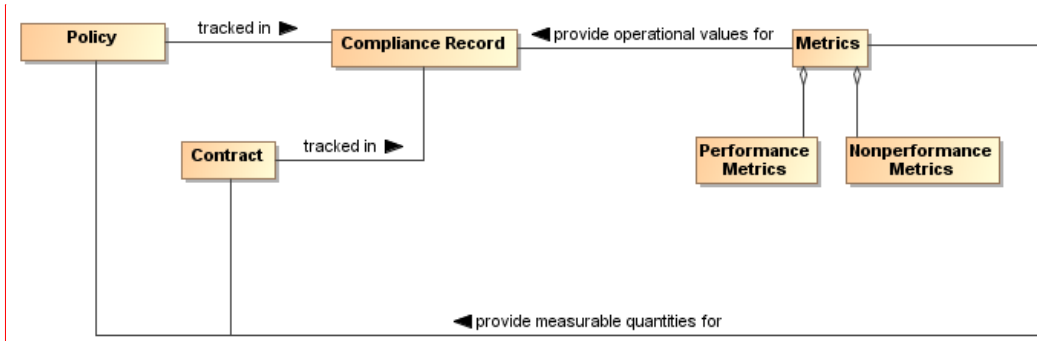
1510 Contracts are agreements among the service participants. The contract may reconcile inconsistent
1511 policies asserted by the participants or may specify details of the interaction. Service level agreements
1512 (SLAs) are one of the commonly used categories of contracts.

Comment [KJL41]: Issue 179, part

1513 The definition and later enforcement of policies and contracts are predicated on the potential for
1514 measurement; the relationships among the relevant concepts are shown in the model in Figure 20.
1515 Performance Metrics identify quantities that characterize the speed and quality of realizing the real world
1516 effects produced using the SOA service; in addition, policies and contracts may depend on
1517 nonperformance metrics, such as whether a license is in place to use the service. Some of these metrics
1518 may reflect the underlying capability, some metrics may reflect processing of the SOA service, and some
1519 metrics may include expected network overhead. The metrics should be carefully defined to avoid
1520 confusion in exactly what is being reported, for example, a case where the service processing time is
1521 reported as if it were the total time including the capability and network processing but is only measuring
1522 the service processing. Some of these metrics reflect the underlying capability, e.g. a SOA service cannot
1523 respond in two seconds if the underlying capability is expected to take five seconds to do its processing;
1524 some metrics reflect the SOA service, e.g. the additional overhead introduced when making data access
1525 requests across the network.

Comment [KJL42]: Issue 254

1526



Comment [PFB43]: Issue 66, part

1527
1528

Figure 20 - Policies and Contracts, Metrics, and Compliance Records

1529 As with many quantities, the metrics associated with a service are not themselves defined by this Service
1530 Description Model because it is not known *a priori* which metrics are being collected or otherwise checked
1531 by the services, the SOA infrastructure, or other resources that participate in the SOA interactions.
1532 However, the service description SHOULD provide a placeholder (possibly through a link to an externally
1533 compiled list) for identifying which metrics are available and how these can be accessed.

1534 The use of metrics to evaluate compliance and the results of compliance evaluation SHOULD be
1535 maintained in compliance records and the means to access the compliance records MAY be included in
1536 the Service Policies portion of the service description. For example, the description may be in the form of
1537 static information (e.g. over the first year of operation, this service had a 91% availability), a link to a
1538 dynamically generated metric (e.g. over the past 30 days, the service has had a 93.3% availability), or
1539 access to a dynamic means to check the service for current availability (e.g., a ping). The relationship
1540 between service **presence** and the presence of the individual actions that can be invoked is discussed
1541 under Reachability in Section 4.2.2.3.

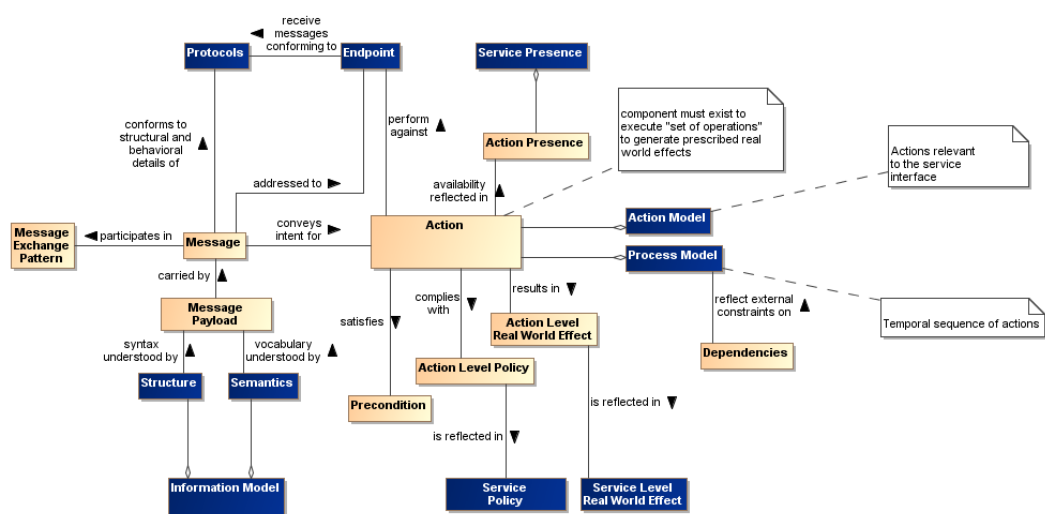
1542 Note, even when policies relate to the perspective of a single participant, policy compliance can be
1543 measured and policies may be enforceable without contractual agreement with other participants. While
1544 certain elements of contracts and contract compliance are likely private, public aspects of compliance
1545 should be reflected in the compliance record information referenced in the service description. [This](#)
1546 [provides input to evidence that supports determining willingness as described in Section 3.2.5.1.](#)

Comment [PFB44]: Issue 70

1547 4.1.2 Use of Service Description

1548 4.1.2.1 Service Description in support of Service Interaction

1549 If we assume we have awareness, the service participants must still establish willingness and presence to
1550 ensure full visibility (See Section 4.2) and to interact with the service. Service description provides
1551 necessary information for many aspects of preparing for and carrying through with interaction. Recall the
1552 fundamental definition of a SOA service is a mechanism to access an underlying capability; the service
1553 description describes this mechanism and its use. It lays the groundwork for what can occur, whereas
1554 service interaction comprises the specifics through which real-world effects are realized.



Comment [KJL45]: Issues 66, part; and 256

1555
1556

Figure 21 - Relationship between Action and Components of Service Description Model

1557 Figure 21 combines the models in the subsections of Section 4.1.1 to concisely relate action and the
1558 relevant components of the Service Description model. The purpose of Figure 21 is to demonstrate that
1559 the components of service description go beyond arbitrary documentation and form the critical set of
1560 information needed to define the what and how of action. In Figure 21, the leaf nodes from Figure 16 are
1561 shown in blue.

1562 Action is typically invoked via a Message where the structure and behavioral processing details of the
1563 message conform to an identified Protocol and is directed to the address of the identified endpoint, and
1564 the message payload conforms to the service Information Model.

Comment [PFB46]: Issue 182

1565 The availability of an action is reflected in the Action Presence and each Action Presence contributes to
1566 the overall Service Presence; this is discussed further in Section 4.2.2.3. Each action has its own
1567 endpoint and protocols are associated with the endpoint⁶. The endpoint and service presence are also
1568 part of the service description.

1569 An action may have preconditions where a Precondition is something that needs to be in place before an
1570 action can occur, e.g. confirmation of a precursor action. Whether preconditions are satisfied is evaluated
1571 when an actor tries to perform the action and not before. Presence for an action means an actor can
1572 initiate it and is independent of whether the preconditions are satisfied. However, the successful
1573 completion of the action may depend on whether its preconditions were satisfied. [The service as a whole
1574 may assume responsibility for providing fallback if a precondition is not met, and the service description
1575 may indicate functionality without explicitly containing details of how preconditions are satisfied or
1576 otherwise mitigated.](#)

Comment [KJL47]: Issue 75

1577 Analogous to the relationship between actions and preconditions, the Process Model may imply
1578 Dependencies for succeeding steps in a process, e.g. that a previous step has successfully completed, or
1579 may be isolated to a given step. An example of the latter would be a dependency that the host server has
1580 scheduled maintenance and access attempts at these times would fail. Dependencies related to the
1581 process model do not affect the presence of a service although these may affect whether the business
1582 function successfully completes. [The service as a whole may assume responsibility for providing fallback
1583 if a dependency is not met, and the service description may indicate functionality without explicitly
1584 containing details of how dependencies are satisfied or otherwise mitigated.](#)

Comment [KJL48]: Issue 76

⁶ This is analogous to a WSDL 2.0 interface operation (WSDL 1.1 portType) having one or more defined bindings and the service identifies the endpoints (WSDL 1.1 ports) corresponding to the bindings.

1585 The conditions under which an action can be invoked may depend on policies associated with the action.
1586 The Action Level Policies MUST be reflected in (or subsumed by) the Service Policies because such
1587 policies may be critical to determining whether the conditions for use of the service are consistent with the
1588 policies asserted by the service consumer. [For example, if an action requires interaction with another
1589 service and that other service has licensing requirements, then the service with such an action also has
1590 the same requirement.](#) The Service Policies are included in the service description.

Comment [KJL49]: Issue 77

1591 Similarly, the result of invoking an action is one or more real world effects, and any Action Level Real
1592 World Effects MUST be reflected in the Service Level Real World Effect included in the service
1593 description. The unambiguous expression of action level policies and real world effects as service
1594 counterparts is necessary to adequately describe what constitutes the service interaction. [For example, if
1595 an action allows for the tracking of user preferences, then the service with such an action results in the
1596 same real world effect.](#)

Comment [KJL50]: Similar to Issue 77
but never explicitly entered

1597 An adequate service description MUST provide a consumer with information needed to determine if the
1598 service policies, the (business) functions, and service-level real world effects are of interest, and there is
1599 nothing in the technical constraints that preclude use of the service.

1600 Note at the service level, the business functions are not concerned with the action or process models.
1601 These models are detailed separately.

1602 The service description is not intended to be isolated documentation but rather an integral part of service
1603 use. Changes in service description SHOULD immediately be made known to consumers and potential
1604 consumers.

1605 4.1.2.2 Description and Invoking Actions Against a Service

Comment [PFB51]: Issue 308

1606 At this point, let us assume the descriptions were sufficient to establish willingness; see Section 4.2.2.2.
1607 Figure 21 indicates the service endpoint establishes where to actually carry out the interaction. This is
1608 where we start considering the action and process models.

1609 The action model identifies the multiple actions a user can perform against a service and the user would
1610 perform these in the context of the process model as specified or referenced under the Service Interface
1611 [Description](#) portion of Service Description. For a given business function, there is a corresponding
1612 process model, where any process model may involve multiple actions. From the above discussion of
1613 model elements of description we may conclude (1) actions have reachability information, including
1614 endpoint and presence, (2) presence of service is some aggregation of presence of its actions, (3) action
1615 preconditions and service dependencies do not affect presence although these may affect successful
1616 completion.

Comment [PFB52]: Issue 183

1617 Having established visibility, the interaction can proceed. Given a business function, the consumer knows
1618 what will be accomplished (the service functionality), the conditions under which interaction will proceed
1619 (service policies), and the process that must be followed (the process model). The remaining question is
1620 how the description information for structure and semantics enable interaction.

1621 We have established the importance of the process model in identifying relevant actions and their
1622 sequence. Interaction proceeds through messages and thus it is the syntax and semantics of the
1623 messages with which we are here concerned. A common approach is to define the structure and
1624 semantics that can appear as part of a message; then assemble the pieces into messages; and,
1625 associate messages with actions. Actions make use of structure and semantics as defined in the
1626 information model to describe its legal messages.

1627 The process model identifies actions to be performed against a service and the sequence for performing
1628 the actions. For a given action, the Reachability portion of description indicates the protocol bindings that
1629 are available, the endpoint corresponding to a binding, and whether there is presence at that endpoint. An
1630 interaction is through the exchange of messages that conform to the structure and semantics defined in
1631 the information model and the message sequence conforming to the action's identified MEP. The result is
1632 some portion of the real world effect that must be assessed and/or processed (e.g. if an error exists, that
1633 part that covers the error processing would be invoked).

1634 **4.1.2.3 The Question of Multiple Business Functions**

1635 Action level effects and policies MUST be reflected at the service level for service description to support
1636 visibility.

1637 It is assumed that a SOA service represents an identifiable business function to which policies can be
1638 applied and from which desired business effects can be obtained. While contemporary discussions of
1639 SOA services and supporting standards do not constrain what actions or combinations of actions can or
1640 should be defined for a service, the SOA-RAF considers the implications of service description in defining
1641 the range of actions appropriate for an individual SOA service.

1642 Consider the situation if a given SOA service is the mechanism for access to multiple independent (but
1643 loosely related) business functions. These are not multiple effects from a single function but multiple
1644 functions with potentially different sets of effects for each function. A service can have multiple actions a
1645 user may perform against it, and this does not change with multiple business functions. As an individual
1646 business function corresponds to a process model, so multiple business functions imply multiple process
1647 models. The same action may be used in multiple process models but the aggregated service presence
1648 would be specific to each business function because the components being aggregated may be different
1649 between process models. In summary, for a service with multiple business functions, each function has
1650 (1) its own process model and dependencies, (2) its own aggregated presence, and (3) possibly its own
1651 list of policies and real world effects.

1652 A common variation on this theme is for a single service to have multiple endpoints for different levels of
1653 quality of service (QoS), e.g. Gold, Silver, and Bronze. Different QoS imply separate statements of policy,
1654 separate endpoints, possibly separate dependencies, and so on. One could say the QoS variation does
1655 not require this because there can be a single QoS policy that encompasses the variations, and all other
1656 aspects of the service would be the same except for the endpoint used for each QoS. However, the
1657 different aspects of policy at the service level would need to be mapped to endpoints, and this introduces
1658 an undesirable level of coupling across the elements of description. In addition, it is obvious that
1659 description at the service level can become very complicated if the number of combinations is allowed to
1660 grow.

Comment [KL53]: Issue 257

1661 One could imagine a service description that is basically a container for action descriptions, where each
1662 action description is self-contained; however, this would lead to duplication of description components
1663 across actions. If common description components are factored, this either is limited to components
1664 common across all actions or requires complicated tagging to capture the components that often but do
1665 not universally apply.

1666 If a provider cannot describe a service as a whole but must describe every action, this leads to the
1667 situation where it may be extremely difficult to construct a clear and concise service description that can
1668 effectively support discovery and use without tedious logic to process the description and assemble the
1669 available permutations. In effect, if adequate description of an action begins to look like description of a
1670 service, it may be best to have it as a separate service.

1671 Recall, more than one service can access the same underlying capability, and this is appropriate if a
1672 different real world effect is to be exposed. Along these lines, one can argue that different QoS are
1673 different services because getting a response in one minute rather than one hour is more than a QoS
1674 difference; it is a fundamental difference in the business function being provided.

1675 As a best practice, the criteria for whether a service is appropriately scoped may be the ease or difficulty
1676 in creating an unambiguous service description. A consequence of having tightly-scoped services is there
1677 will likely be a greater reliance on combining services, i.e. more fundamental business functions, to create
1678 more advanced business functions. This is consistent with the principles of service oriented architecture
1679 and is the basic position of this Reference Architecture Foundation, although not an absolute
1680 requirement. Combining services increases the reliance on understanding and implementing the concepts
1681 of orchestration, choreography, and other approaches yet to be developed; these are discussed in more
1682 detail in section 4.4 Interacting with Services.

1683 **4.1.2.4 Service Description, Execution Context, and Service Interaction**

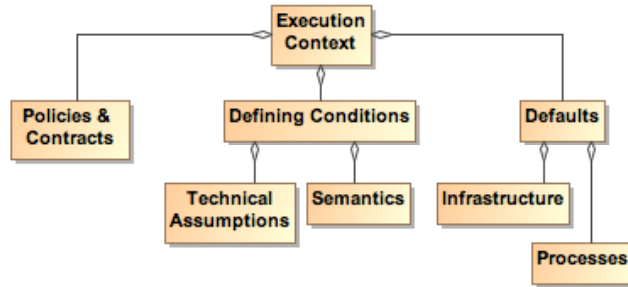
1684 The service description MUST provide sufficient information to support service visibility, including the
1685 willingness of service participants to interact. However, the corresponding descriptions for providers and

1686 consumers may both contain policies, technical assumptions, constraints on semantics, and other
1687 technical and procedural conditions that must be aligned to define the terms of willingness. The
1688 agreements that encapsulate the necessary alignment form the basis upon which interactions may
1689 proceed – in the Reference Model, this collection of agreements and the necessary environmental
1690 support establish the execution context.

1691 | To illustrate ~~the concept of the~~ execution context of a service interaction, consider a Web-based system
1692 for timecard entry. For an employee onsite at an employer facility, the execution context requires a
1693 computer connected to the local network and the employee must enter their network ID and password.
1694 Relevant policies include that the employee must maintain the most recent anti-virus software and virus
1695 definitions for any computer connected to the network.

Comment [PFB54]: Issue 185

1696 For the same employee connecting from offsite, the execution context specifies the need for a computer
1697 with installed VPN software and a security token to negotiate the VPN connection. The execution context
1698 also includes proxy settings as needed to connect to the offsite network. The employee must still comply
1699 with the requirements for onsite computers and access, but the offsite execution context includes
1700 additional items before the employee can access the same underlying capability and realize the same
1701 real world effects, i.e. the timecard entries.



1702
1703 *Figure 22 - Execution Context*

1704 Figure 22 shows a few broad categories found in execution context. These are not meant to be
1705 comprehensive. Other items may need to be included to provide a sufficient description of the interaction
1706 conditions. Any other items not explicitly noted in the model but needed to set the environment SHOULD
1707 be included in the execution context.

1708 While the execution context captures the conditions under which interaction can occur, it does not capture
1709 the specific service invocations that do occur in a specific interaction. A service interaction as modeled in
1710 Figure 23 introduces the concept of an Interaction Description that is composed of both the Execution
1711 Context and an Interaction Log. The execution context specifies the set of conditions under which the
1712 interaction occurs and the interaction log captures the sequence of service interactions that occur within
1713 the execution context. This sequence should follow the Process Model but can include details beyond
1714 those specified there. For example, the Process Model may specify an action that results in identifying a
1715 data source, and the identified source is used in a subsequent action. The Interaction Log would record
1716 the specific data source used.

1717 The execution context can be thought of as a container in which the interaction occurs and the interaction
1718 log captures what happens inside the container. This combination is needed to support auditability and
1719 repeatability of the interactions.

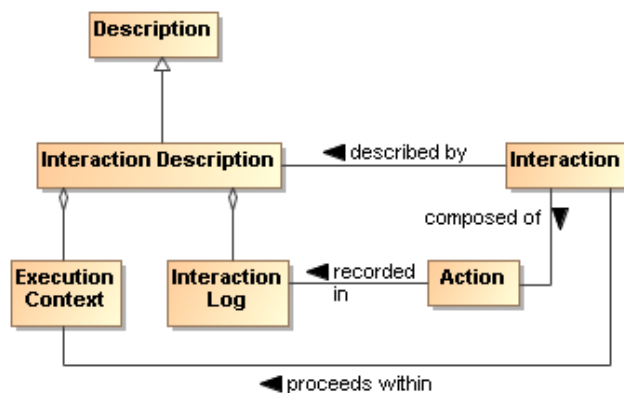


Figure 23 - Interaction Description

1720
1721
1722
1723
1724
1725
1726
1727
1728
1729
1730
1731
1732
1733
1734
1735
1736
1737
1738
1739

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

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

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

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

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

4.1.3 Relationship to Other Description Models

1740
1741
1742
1743
1744
1745
1746
1747
1748
1749
1750
1751
1752
1753
1754
1755
1756

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

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

With respect to ISO 11179 Part 5, the metadata fields defined in that reference may be used without prejudice as the properties in Figure 15. Additionally, other defined metadata sets may be used by the service provider if the other sets are considered more appropriate, i.e. it is fundamental to this reference architecture to identify the need and the means to make vocabulary declarations explicit but it is beyond the scope to specify which vocabularies are to be used. In addition, the identification of domain of

1757 properties and range of the values has not been included in the current Reference Architecture
1758 discussion, but the text of ISO 11179 Part 5 can be used consistently with the model prescribed in this
1759 document.

1760 Description as defined here considers a wide range of applicability and support of the principles of service
1761 oriented architecture. Other metadata models can be used in concert with the model presented here
1762 because most of these focus on a finer level of detail that is outside the present scope, and so provide a
1763 level of implementation guidance that can be applied as appropriate.

1764 4.1.4 Architectural Implications

1765 The definition of service description ~~indicates~~ has numerous architectural implications ~~on~~ for the SOA
1766 ecosystem:

- 1767 • The real world effects that the service description definition support must be consistent with the
1768 technical assumptions/constraints
- 1769 • ~~#-~~The service description definition changes over time and its contents will reflect changing needs
1770 and context. The service description definition MUST therefore have:
 - 1771 ○ mechanisms to support the storage, referencing, and access to normative definitions of
1772 one or more versioning schemes that may be applied to identify different aggregations of
1773 descriptive information, where the different schemes may be versions of a versioning
1774 scheme itself;
 - 1775 ○ configuration management mechanisms to capture the contents of each aggregation and
1776 apply a unique identifier in a manner consistent with an identified versioning scheme;
 - 1777 ○ one or more mechanisms to support the storage, referencing, and access to conversion
1778 relationships between versioning schemes, and the mechanisms to carry out such
1779 conversions.
- 1780 • Description makes use of defined semantics, where the semantics **MAY** be used for
1781 categorization or providing other property and value information for description classes. In such
1782 cases, the service description MUST have:
 - 1783 ○ semantic models that provide normative descriptions of the utilized terms, where the
1784 models may range from a simple dictionary of terms to an ontology showing complex
1785 relationships and capable of supporting enhanced reasoning;
 - 1786 ○ mechanisms to support the storage, referencing, and access to these semantic models;
 - 1787 ○ configuration management mechanisms to capture the normative description of each
1788 semantic model and to apply a unique identifier in a manner consistent with an identified
1789 versioning scheme;
 - 1790 ○ one or more mechanisms to support the storage, referencing, and access to conversion
1791 relationships between semantic models, and the mechanisms to carry out such
1792 conversions.
- 1793 • Once awareness exists, the service participants MUST still establish willingness and presence to
1794 ensure full visibility (See Section 4.2).
- 1795 • Given a business function, the consumer knows what will be accomplished (the service
1796 functionality), the conditions under which interaction will proceed (service policies), and the
1797 process that MUST be followed (the process model).
- 1798 • Actions MAY have associated policies stating conditions for performing the action, but these
1799 MUST be reflected in and be consistent with the policies made visible at the service level and
1800 thus the description of the service as a whole.
- 1801 • Policies asserted MAY be reflected as Technical Assumptions/Constraints that available services
1802 or their underlying capabilities MUST be capable of meeting.
- 1803 • Descriptions include reference to policies defining conditions of use. In this sense, policies are
1804 also resources that need to be visible, discoverable, and accessible. The service description (as
1805 also enumerated under governance) MUST have:
 - 1806 ○ description of policies, including a unique identifier for the policy and a sufficient,
1807 preferably machine processable, representation of the meaning of terms used to describe
1808 the policy, its functions, and its effects;

- 1809 | o [a method to](#) enable searching for policies that best meet the search criteria specified by
1810 | the service participant; where the discovery mechanism has access to the individual
1811 | policy descriptions, possibly through some repository mechanism;
1812 | o accessible storage of policies and policy descriptions, so service participants can access,
1813 | examine, and use the policies as defined.
- 1814 | • Descriptions include references to metrics that describe the operational characteristics of the
1815 | subjects being described. [The service description definition](#) (as [also](#) partially enumerated under
1816 | governance) **MUST have**:
- 1817 | o infrastructure monitoring and reporting information on SOA resources;
1818 | o possible interface requirements to make accessible metrics information generated;
1819 | o mechanisms to catalog and enable discovery of which metrics are available for a
1820 | described resources and information on how these metrics can be accessed;
1821 | o mechanisms to catalog and enable discovery of compliance records associated with
1822 | policies and contracts that are based on these metrics.
- 1823 | • Descriptions of the interactions are important for enabling auditability and repeatability, thereby
1824 | establishing a context for results and support for understanding observed change in performance
1825 | or results. [Thus, the service description definition MUST have](#):
- 1826 | o one or more mechanisms to capture, describe, store, discover, and retrieve interaction
1827 | logs, execution contexts, and the combined interaction descriptions;
1828 | o one or more mechanisms for attaching to any results the means to identify and retrieve
1829 | the interaction description under which the results were generated.
- 1830 | • Descriptions may capture very focused information subsets or can be an aggregate of numerous
1831 | component descriptions. Service description is an example of an aggregate for which manual
1832 | maintenance of the whole would not be feasible. [Thus, the service description definition MUST](#)
1833 | [have](#):
- 1834 | o tools to facilitate identifying description elements that are to be aggregated to assemble
1835 | the composite description;
1836 | o tools to facilitate identifying the sources of information to associate with the description
1837 | elements;
1838 | o tools to collect the identified description elements and their associated sources into a
1839 | standard, referenceable format that can support general access and understanding;
1840 | o tools to automatically update the composite description as the component sources
1841 | change, and to consistently apply versioning schemes to identify the new description
1842 | contents and the type and significance of change that occurred.
- 1843 | • The description is the source of vital information in establishing willingness to interact with a
1844 | resource, reachability to make interaction possible, and compliance with relevant conditions of
1845 | use. [Thus, the service description definition MUST have](#):
- 1846 | o one or more discovery mechanisms that enable searching for described resources that
1847 | best meet the criteria specified by a service participant;
1848 | o tools to appropriately track users of the descriptions and notify them when a new version
1849 | of the description is available.
- 1850 | • [The service description MUST provide sufficient information to support service visibility, including](#)
1851 | [the willingness of service participants to interact. However, the corresponding descriptions for](#)
1852 | [providers and consumers may both contain policies, technical assumptions, constraints on](#)
1853 | [semantics, and other technical and procedural conditions that must be aligned to define the terms](#)
1854 | [of willingness](#)

1855 | 4.2 Service Visibility Model

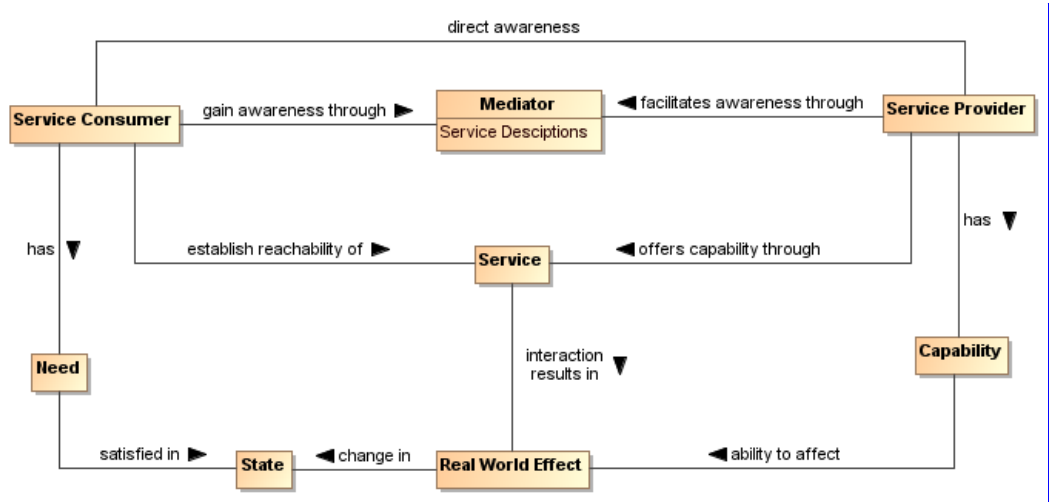
1856 | One of the key requirements for participants interacting with each other in the context of a SOA
1857 | ecosystem is achieving visibility: before services can interoperate, the participants have to be visible to
1858 | each other using whatever means are appropriate. The Reference Model analyzes visibility in terms of
1859 | awareness, willingness, and reachability. In this section, we explore how visibility may be achieved.

1860 **4.2.1 Visibility to Business**

1861 The relationship of visibility to the SOA ecosystem encompasses both human social structures and
 1862 automated IT mechanisms. Figure 24 depicts a business setting that is a basis for visibility as related to
 1863 the Social Structure Model (Figure 3) in the Participation in a SOA Ecosystem view (see Section 3.1). The
 1864 participants acting in the various roles of service consumers, mediators, and service providers may have
 1865 direct awareness or mediated awareness where mediated awareness is achieved through some third
 1866 party. A consumer's willingness to use a service is reflected by the consumer's presumption of satisfying
 1867 goals and needs as these compare with information provided in the service description. Service providers
 1868 offer capabilities that have real world effects that result in a change in state. Reachability of the service by
 1869 the consumer may lead to interactions that change the state of the SOA ecosystem. The consumer can
 1870 measure the change of state to determine if the claims made by description and the real world effects of
 1871 consuming the service meet the consumer's needs.

Comment [PFB55]: Reworded Issue 294

1872
1873



Comment [KL56]: Issues 85, part; and 301

1874
1875

Figure 24 - Visibility to Business

1876 Visibility and interoperability in a SOA ecosystem requires more than location and interface information. A
 1877 meta-model for this broader view of visibility is depicted in Section 4.1. In addition to providing improved
 1878 awareness of service capabilities through description of information such as reachability, behavior
 1879 models, information models, functionality, and metrics, the service description may identify policies
 1880 valuable for determination of willingness to interact.

1881 A mediator using service descriptions may provide event notifications to both consumers and providers
 1882 about information relating to the descriptions. One example of this capability is a publish/subscribe model
 1883 where the mediator allows consumers to subscribe to service description version changes made by the
 1884 provider. Likewise, the mediator may provide notifications to the provider of consumers that have
 1885 subscribed to service description updates.

1886 Another important capability in a SOA ecosystem is the ability to narrow visibility to trusted members
 1887 within a social structure. Mediators for awareness may provide policy based access to service
 1888 descriptions allowing for the dynamic formation of awareness between trusted members.

1889 **4.2.2 Visibility**

1890 Attaining visibility is described in terms of steps that lead to visibility. Different participant communities can
 1891 bring different contexts for visibility within a single social structure, and the same general steps can be
 1892 applied to each of the contexts to accomplish visibility.

1893 Attaining SOA visibility requires

- 1894 • service description creation and maintenance,
- 1895 • processes and mechanisms for achieving awareness of and accessing descriptions,
- 1896 • processes and mechanisms for establishing willingness of participants,
- 1897 • processes and mechanisms to determine reachability.

1898 Visibility may occur in stages, i.e. a participant can become aware enough to look or ask for further
 1899 description, and with this description, the participant can decide on willingness, possibly requiring
 1900 additional description. For example, if a potential consumer has a need for a tree cutting (business)
 1901 service, the consumer can use a web search engine to find web sites of providers. The web search
 1902 engine (a mediator) gives the consumer links to relevant web pages and the consumer can access those
 1903 descriptions. For those prospective providers that satisfy the consumer's criteria, the consumer's
 1904 willingness to interact increases. The consumer may contact several tree services to get detailed cost
 1905 information (or arrange for an estimate) and may ask for references (further description). The consumer is
 1906 likely to establish full visibility and proceed with interaction with the tree service that mutually establishes
 1907 visibility.

1908 4.2.2.1 Awareness

Comment [KL57]: changes in this section per Issue 302

1909 An important means for a serviceone participant to be aware of another participant is to have access to a
 1910 description of that participant and for the description to have be sufficiently completeness to establish
 1911 support the other requirements of visibility.

1912 Awareness ~~is inherently a function of a participant; awareness~~ can be established without any action on
 1913 the part of the target participant other than the target providing appropriate descriptions. Awareness is
 1914 often discussed in terms of consumer awareness of providers but the concepts are equally valid for
 1915 provider awareness of consumers.

1916 Awareness can be decomposed into: creating the descriptions, making them available, and discovering
 1917 the descriptions. Discovery can be initiated or it can be by notification. ~~Initiated discovery for business~~
 1918 ~~may require formalization of the required capabilities and resources to achieve business goals.~~

1919 Achieving awareness in a SOA ecosystem can range from word of mouth to formal service descriptions in
 1920 a standards-based registry-/repository. Some other examples of achieving awareness in a SOA
 1921 ecosystem are the use of a web page containing description information, email notifications of
 1922 descriptions, and document based descriptions.

1923 A mediator for awareness is a third party participant whose use provides awareness to one or more
 1924 consumers of one or more services. Direct awareness is awareness between a consumer and provider
 1925 without the use of a third party. The use of a registry/repository can provide awareness as can a Web
 1926 page displaying similar information.

1927 Direct awareness may be the result of having previously established an execution context, or direct
 1928 awareness may include determining the presence of services and then querying the service directly for
 1929 description. As an example, a priori visibility of some sensor device may provide the means for interaction
 1930 or a query for standardized sensor device metadata may be broadcast to multiple locations. If
 1931 acknowledged, the service interface for the device may directly provide description to a consumer so the
 1932 consumer can determine willingness to interact.

1933 The same medium for awareness may be direct in one context and may be mediated in another context.
 1934 For example, a service provider may maintain a web site with links to the provider's descriptions of
 1935 services giving the consumers direct awareness to the provider's services. Alternatively, a community
 1936 may maintain a web site with a search interface that makes use of an index of these (and possibly other)
 1937 descriptions of services, and the web site could be used by any number of consumers. More than one
 1938 approach to mediation may be involved, as different sources of description may specialize in different
 1939 functions whose use provides mediation.

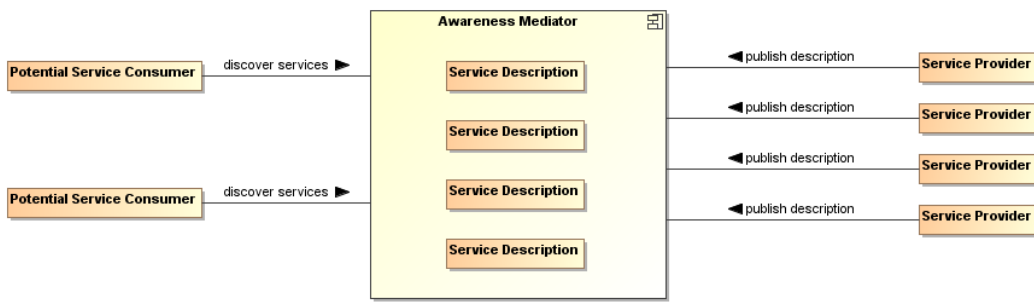
Comment [KL58]: Issue 187, 188

1940 Descriptions may be formal or informal. Section 4.1, provides a comprehensive model for service
 1941 description that can be used to mediate visibility. Using consistent description taxonomies and standards
 1942 based mediated awareness helps provide more effective awareness.

1943 **4.2.2.1.1 Mediated Awareness**

1944 Mediated awareness promotes **simplification of the overall services infrastructure**. Rather than all
1945 potential service consumers being informed on a continual basis about all services, there is a known or
1946 agreed upon facility or location that stores and supports discovery and/or notification related to the
1947 service description.

Comment [PFB59]: Issue 189



Comment [PFB60]: Issue 190

1948
1949 *Figure 25 - Mediated -Awareness*

1950 In Figure 25, the potential service consumers perform queries or are notified in order to locate those
1951 services that satisfy their needs. As an example, the telephone book is a mediating registry where
1952 individuals perform manual searches to locate services (i.e. the yellow pages). The telephone book is
1953 also a mediated registry for solicitors to find and notify potential customers (i.e. the white pages).

1954 In mediated service awareness for large and dynamic numbers of service consumers and service
1955 providers, the benefits of utilizing the **awareness** mediator typically far outweigh the management issues
1956 associated with it. Some of the benefits of mediated service awareness are

Comment [PFB61]: Issue 192

- 1957 • Potential service consumers have a known location for searching thereby eliminating needless
1958 and random searches
- 1959 • Typically a consortium of interested parties (or a sufficiently large corporation) **signs up to serves**
1960 **as the host of** the mediation facility
- 1961 • Standardized tools and methods can be developed and promulgated to promote interoperability
1962 and ease of use.

1963 However, mediated awareness can have some risks associated with it:

- 1964 • A single point of failure. If the **awareness mediator/mediation service** fails then a large number of
1965 service providers and consumers are potentially adversely affected.
- 1966 • A single point of control. If the **central mediation service awareness mediator** is owned by, or
1967 controlled by, someone other than the service consumers and/or providers then the latter may be
1968 put at a competitive disadvantage based on policies of the discovery provider.

Comment [PFB62]: Issue 193

Comment [PFB63]: Issue 194

1969 A common mechanism for mediated awareness is a registry/repository. The registry stores links or
1970 pointers to service description artifacts. The repository in this example is the storage location for the
1971 service description artifacts. Service descriptions can be pushed (publish/subscribe for example) or pulled
1972 from the registry/repository mediator.

1973 Registries/repositories may be referred to as federated when supported functions, such as responding to
1974 discovery requests, are distributed across multiple registry/repository instances.

1975 **4.2.2.1.2 Awareness in Complex Social Structures**


1976 Awareness applies to **one or more communities within** one or more social structures where **a community**
1977 **consists of there is** at least one description provider and one description consumer. **These communities**
1978 **may be part of Awareness may occur within** the same social structure or **be part of different ones across**
1979 **social structures**.

1980 In Figure 26, awareness can be between **a limited set of** consumers and providers within a single
1981 **community, multiple communities, or all communities in the** social structure. Within a social structure,
1982 awareness can be encouraged or restricted through policies and these policies can affect participant

Comment [PFB64]: Issue 302, part

Comment [PFB65]: Issue 195

1983 willingness. The information about policies should be incorporated in the relevant descriptions.
 1984 Additionally, the conditions for establishing contracts are governed within a social structure.



Comment [PFB66]: Issue 196
 Comment [PFB67]: Issue 302, part
 Formatted: Centered

Figure 26 - Awareness in a SOA Ecosystem

1985
 1986
 1987 IT policy/contract mechanisms can be used by visibility mechanisms to provide awareness between [social](#)
 1988 [structures, including communities](#). ~~The IT mechanisms for awareness may incorporate~~ trust mechanisms
 1989 to enable awareness between trusted [social structures](#)~~communities~~. For example, government
 1990 organizations may want to limit awareness of an organization's services to specific communities of
 1991 interest.

1992 Another common business model for awareness is maximizing awareness to [communities](#)~~those~~ within
 1993 the social structure, the traditional market place business model. A centralized awareness-mediator often
 1994 arises as a provider for this global visibility, a gatekeeper of visibility so to speak. For example, Google is
 1995 a centralized awareness-mediator for accessing information on the web. As another example, television
 1996 networks have centralized entities providing a level of awareness to communities that otherwise could not
 1997 be achieved without going through the television network.

1998 However, mediators have motivations, and they may be selective in which information they choose to
 1999 make available to potential consumers. For example, in a secure environment, the mediator may enforce
 2000 security policies and make information selectively available depending on the security clearance of the
 2001 consumers.

2002 4.2.2.2 Willingness

2003 Having achieved awareness, participants use descriptions to help determine their willingness to interact
 2004 with another participant. Both awareness and willingness are determined prior to consumer/provider
 2005 interaction.

2006 **Error! Reference source not found.**By [having establishing](#) a willingness to interact within a particular
 2007 social structure ([see Section 3.2.5.1](#)), the social structure provides the participant access to capabilities
 2008 based on conditions the social structure finds appropriate for its context. The participant can use these
 2009 capabilities to satisfy goals and objectives as specified by the participant's needs.

2010 ~~Information~~ information used to determine willingness is [defined provided](#) by Description ([see Section 4.1.1](#)).
 2011 Information referenced by Description may come from many sources. For example, a mediator for
 2012 descriptions may provide 3rd party annotations for reputation. Another source for reputation may be a
 2013 participant's own history of interactions with another participant. [The contribution of real world effects to](#)
 2014 [providing evidence and establishing the reputation of a participant is discussed with relation to](#) Figure 9.

2015 A participant inspects functionality for potential satisfaction of needs. Identity is associated with any
 2016 participant, however, identity may or may not be verified. If available, participant reputation may be a
 2017 deciding factor for willingness to interact. Policies and contracts referenced by the description may be
 2018 particularly important to determine the agreements and commitments required for business interactions.
 2019 Provenance may be used for verification of authenticity of a resource.

2020 Mechanisms that aid in determining willingness make use of the artifacts referenced by descriptions of
 2021 services. Mechanisms for establishing willingness could be as simple as rendering service description
 2022 information for human consumption to automated evaluation of functionality, policies, and contracts by a
 2023 rules engine. The rules engine for determining willingness could operate as a policy decision procedure
 2024 as defined in Section 4.4.

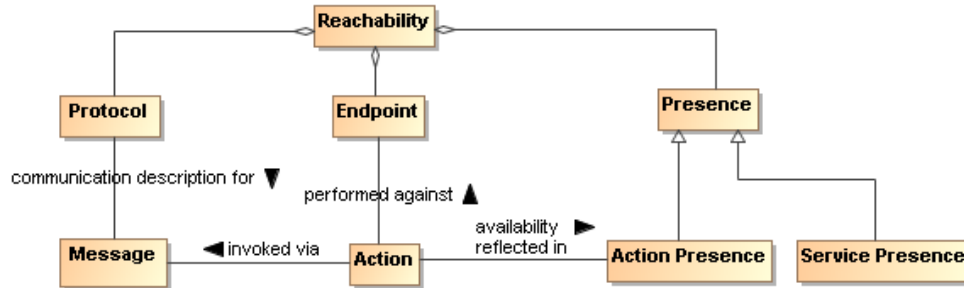
Comment [KJL68]: Figure deleted
 Issue 262 (part)

Comment [KJL69]: Issue 262 (part)

Comment [KJL70]: Issue 262 (part)

2025 **4.2.2.3 Reachability**

2026 Reachability involves knowing the endpoint, protocol, and presence of a service. At a minimum,
2027 reachability requires information about the location of the service and the protocol describing the means
2028 of communication.



2029
2030 *Figure 27 - Service Reachability*

2031 **Endpoint**

2032 A reference-able entity, processor or **resource** against which an **action** can be performed.

2033 **Protocol**

2034 A structured means by which details of a service interaction mechanism are defined.

2035 **Presence**

2036 The measurement of reachability of a service at a particular point in time.

2037 A protocol defines a structured method of communication. Presence is determined by interaction through
2038 a communication protocol. Presence may not be known in many cases until the interaction begins. To
2039 overcome this problem, IT mechanisms may make use of presence protocols to provide the current
2040 up/down status of a service.

2041 Service reachability enables service participants to locate and interact with one another. Each action may
2042 have its own endpoint and also its own protocols associated with the endpoint and whether there is
2043 presence for the action through that endpoint. Presence of a service is an aggregation of the presence of
2044 the service's actions, and the service level may aggregate to some degraded or restricted presence if
2045 some action presence is not confirmed. For example, if error processing actions are not available, the
2046 service can still provide required functionality if no error processing is needed. This implies reachability
2047 relates to each action as well as applying to the service/business as a whole.

2048 **4.2.3 Architectural Implications**

2049 Visibility in a SOA ecosystem has the following architectural implications on mechanisms providing
2050 support for awareness, willingness, and reachability:

- 2051 • Mechanisms providing support for awareness **MUST** have the following minimum capabilities:
 - 2052 ○ creation of Description, preferably conforming to a standard Description format and
 - 2053 structure;
 - 2054 ○ publishing of Description directly to a consumer or through a third party mediator;
 - 2055 ○ discovery of Description, preferably conforming to a standard for Description discovery;
 - 2056 ○ notification of Description updates or notification of the addition of new and relevant
 - 2057 Descriptions;
 - 2058 ○ classification of Description elements according to standardized classification schemes.
- 2059 • In a SOA ecosystem with complex social structures, awareness **MAY** be provided for specific
2060 communities of interest. The architectural mechanisms for providing awareness to communities
2061 of interest **MUST** support:
 - 2062 ○ policies that allow dynamic formation of communities of interest;

- 2063 ○ trust that awareness can be provided for and only for specific communities of interest, the
2064 bases of which are typically built on encryption technologies.
- 2065 • The architectural mechanisms for determining willingness to interact **MUST** support:
2066 ○ verification of identity and credentials of the provider and/or consumer;
2067 ○ access to and understanding of description;
2068 ○ inspection of functionality and capabilities;
2069 ○ inspection of policies and/or contracts.
 - 2070 • The architectural mechanisms for establishing reachability **MUST** support:
2071 ○ the location or address of an endpoint;
2072 ○ verification and use of a service interface by means of a communication protocol;
2073 ○ determination of presence with an endpoint which **MAY** only be determined at the point of
2074 interaction but **MAY** be further aided by the use of a presence protocol for which the
2075 endpoints actively participate.

2076 4.3 Interacting with Services Model

2077 Interaction is the activity involved in using a service to access capability in order to achieve a particular
2078 desired real world effect, where real world effect is the actual result of using a service. An interaction can
2079 be characterized by a sequence of communicative actions. Consequently, interacting with a service, i.e.
2080 participating in joint action with the service—usually accomplished mediated by a series of message
2081 exchanges—involves individual actions performed by both the service and the consumer.⁷ Note that a
2082 participant (or delegate acting on behalf of the participant) can be the sender of a message, the receiver
2083 of a message, or both.

Comment [KJL71]: Issue 202

2084 4.3.1 Interaction Dependencies

2085 Recall from the Reference Model that service visibility is the capacity for those with needs and those with
2086 capabilities to be able to interact with each other, and that the service interface is the means by which the
2087 underlying capabilities of a service are accessed. Ideally, the details of the underlying service
2088 implementation are abstracted away by the service interface. [Service] interaction therefore has a direct
2089 dependency on the visibility of the service as well as its implementation-neutral interface (see Figure 28).
2090 Service visibility is composed of awareness, willingness, and reachability, and these are discussed in
2091 Section 4.2. The information related to the service interface description is discussed in Section 4.1.1.3.1,
2092 and the specifics of interaction are detailed in the remainder of Section 4.3. Service visibility is modeled in
2093 Section 4.2.2.

Comment [KJL72]: Issue 203

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.

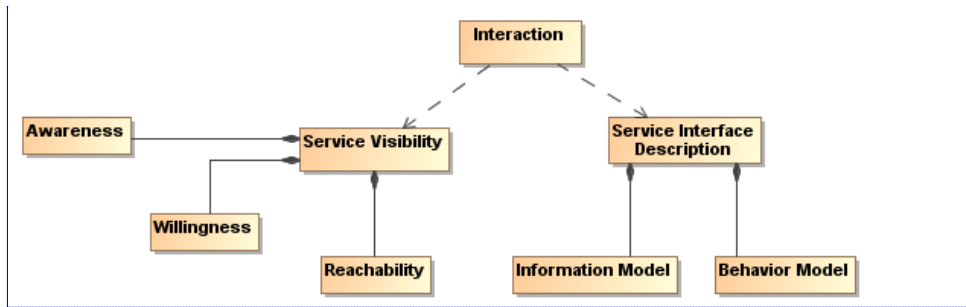


Figure 28 - Interaction dependencies

2094
2095

2096 4.3.2 Actions and Events

2097 The SOA-RAF uses message exchange between service participants to denote actions performed
2098 against and by the service, and to denote events that report on real world effects that are caused by the
2099 service actions. A visual model of the relationship between these concepts is shown in Figure 29.

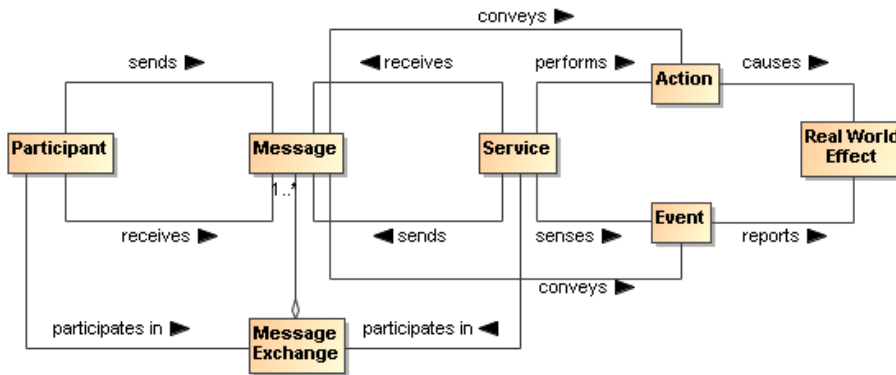


Figure 29 - A 'message' denotes either an action or an event

2100
2101

2102 Both actions and events, realized by the SOA services, are denoted by the messages. The Reference
2103 Model states that the action model characterizes the "permissible set of actions that may be invoked
2104 against a service." We extend that notion here to include events ~~as part of the event model~~ and that
2105 messages are intended for invoking actions or for notification of events.

2106 In Section 3.3.2 we saw that participants interact with each other in order to participate in joint actions. A
2107 joint action is not itself the same thing as the result of the joint action. When a joint action is participated in
2108 with a service, the real world effect that results may be reported in the form of an event notification.

2109 4.3.3 Message Exchange

2110 *Message exchange* is the means by which service participants (or their delegates) interact with each
2111 other. There are two primary modes of interaction: joint actions that cause real world effects and
2112 notification of events that report real world effects⁸.

2113 A message exchange is used to affect an action when the messages contain the appropriately formatted
2114 content, are directed towards a particular action in accordance with the action model, and the delegates
2115 involved interpret the message appropriately.

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

2116 A message exchange is also used to communicate event notifications. An event is an occurrence that is
2117 of interest to some participant; in our case when some real world effect has occurred. Just as action
2118 messages have formatting requirements, so do event notification messages. In this way, the Information
2119 Model of a service must specify the syntax (structure), and semantics (meaning) of the action messages
2120 and event notification messages as part of a service interface. It must also specify the syntax and
2121 semantics of any data that is carried as part of a payload of the action or event notification message. The
2122 Information Model is described in greater detail in the Service Description Model (see Section 4.1).

2123 In addition to the Information Model that describes the syntax and semantics of the messages and data
2124 payloads, exception conditions and error handling in the event of faults (e.g., network outages, improper
2125 message formats, etc.) must be specified or referenced as part of the Service Description.

2126 When a message is used to invoke an action, the correct interpretation typically requires the receiver to
2127 perform an operation, which itself invokes a set of private, internal actions. These **operations** represent
2128 the sequence of (private) actions a service must perform in order to validly participate in a given joint
2129 action.

2130 Similarly, the correct consequence of realizing a real world effect may be to initiate the reporting of that
2131 real world effect via an event notification.

2132 **Message Exchange**

2133 The means by which joint action and event notifications are coordinated by service participants
2134 (or delegates).

2135 **Operations**

2136 The sequence of actions a service must perform in order to validly participate in a given joint
2137 action.

2138 **4.3.3.1 Message Exchange Patterns (MEPs)**

2139 The basic temporal aspect of service interaction can be characterized by two fundamental message
2140 exchange patterns (MEPs):

- 2141 • Request/response to represent how actions cause a real world effect
- 2142 • Event notification to represent how events report a real world effect

2143 This is by no means a complete list of all possible MEPs used for inter- or intra-enterprise messaging but
2144 it does represent those that are most commonly used in exchange of information and reporting changes
2145 in state both within organizations and across organizational boundaries.

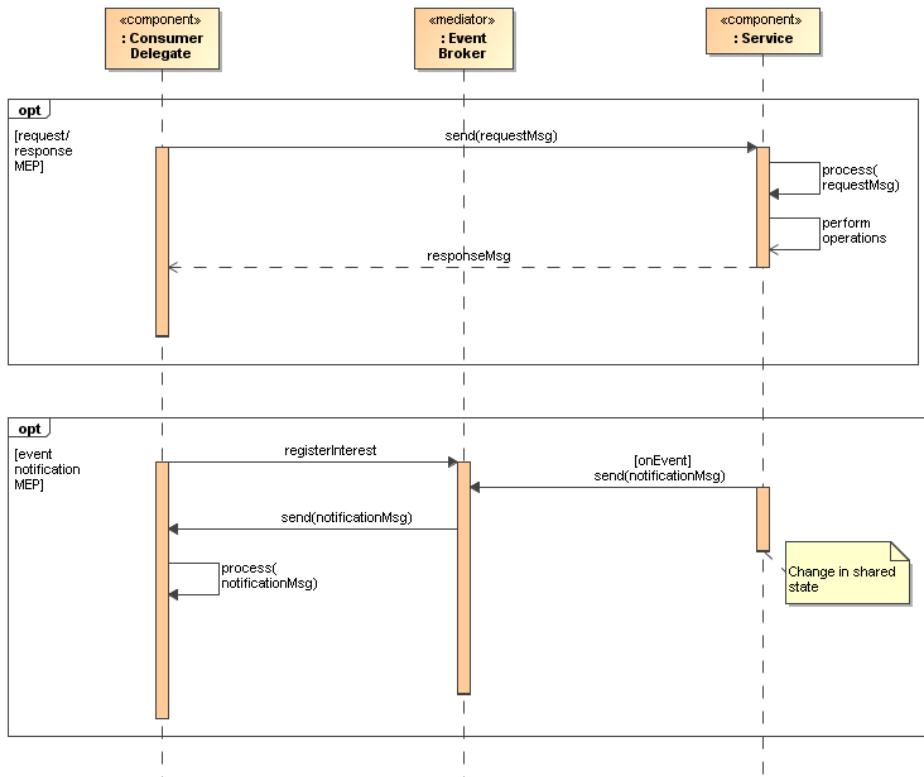


Figure 30 - Fundamental SOA message exchange patterns (MEPs)

2146
2147

2148 Recall from the Reference Model that the Process Model characterizes “the temporal relationships
2149 between and temporal properties of actions and events associated with interacting with the service.”
2150 Thus, MEPs are a key element of the Process Model. The meta-level aspects of the Process Model (just
2151 as with the Action Model) are provided as part of the Service Description Model (see Section 4.1).

2152 In the UML sequence diagram shown in Figure 30 it is assumed that the service participants (consumer
2153 and provider) have delegated message handling to hardware or software delegates acting on their behalf.
2154 In the case of the service consumer, this is represented by the *Consumer Delegate* component. In the
2155 case of the service provider, the delegate is represented by the *Service* component. The message
2156 interchange model illustrated represents a logical view of the MEPs and not a physical view. In other
2157 words, specific hosts, network protocols, and underlying messaging system are not shown, as these tend
2158 to be implementation specific. Although such implementation-specific elements are considered outside
2159 the scope of this document, they are important considerations in modeling the SOA execution context.
2160 Recall from the Reference Model that the *execution context* of a service interaction is “the set of
2161 infrastructure elements, process entities, policy assertions and agreements that are identified as part of
2162 an instantiated service interaction, and thus forms a path between those with needs and those with
2163 capabilities.”

2164 4.3.3.2 Request/Response MEP

2165 In a request/response MEP, the Consumer Delegate component sends a request message to the Service
2166 component. The Service component then processes the request message. Based on the content of the
2167 message, the Service component performs the service operation and the associated private actions.

2168 Following the completion of these operations, a response message is returned to the Consumer Delegate
2169 component. The response could be that a step in a process is complete, the initiation of a follow-on
2170 operation, or the return of requested information.⁹

2171 Although the sequence diagram shows a *synchronous* interaction (because the sender of the request
2172 message, i.e., Consumer Delegate, is blocked from continued processing until a response is returned
2173 from the Service) other variations of request/response are valid, including *asynchronous* (non-blocking)
2174 interaction through use of queues, channels, or other messaging techniques.

2175 What is important to convey here is that the request/response MEP represents action, which causes a
2176 real world effect, irrespective of the underlying messaging techniques and messaging infrastructure used
2177 to implement the request/response MEP.

2178 4.3.3.3 Event Notification MEP


2179 An event is made visible to interested consumers by means of an event notification message exchange
2180 that reports a real world effect; specifically, a change in shared state between service participants. The
2181 basic event notification MEP takes the form of a one-way message sent by a notifier component (in this
2182 case, the Service component) and received by components with an interest in the event (here, the
2183 Consumer Delegate component).

2184 Often the sending component may not be fully aware of all the components that wish to receive the
2185 notification; particularly in so-called publish/subscribe ('pub/sub') situations. In event notification message
2186 exchanges, it is rare to have a tightly-coupled link between the sending and the receiving component(s)
2187 for a number of practical reasons. One of the most common needs for pub/sub messaging is the potential
2188 for network outages or communication interrupts that can result in loss of notification of events. Therefore,
2189 a third-party mediator component is often used to decouple the sending and receiving components.

2190 Although this is typically an implementation issue, because this type of third-party decoupling is so
2191 common in event-driven systems, it is warranted for use in modeling this type of message exchange in
2192 the SOA-RAF. This third-party intermediary is shown in Figure 30 as an Event Broker mediator. As with
2193 the request/response MEP, no distinction is made between synchronous versus asynchronous
2194 communication, although asynchronous message exchange is illustrated in the UML sequence diagram
2195 depicted in Figure 30.

2196 4.3.4 Composition of Services

2197 Composition of services is the act of aggregating or 'composing' a single service from one or more other
2198 services. A simple model of service composition is illustrated in Figure 31.



2199 Figure 31 - Simple model of service composition
2200

2201 Here, Service A is a service that has an exposed interface IServiceA, which is available to the Consumer
2202 Delegate and relies on two other services in its implementation. The Consumer Delegate does not know
2203 that Services B and C are used by Service A, or whether they are used in serial or parallel, or if their

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.

2204 operations succeed or fail. The Consumer Delegate only cares about the success or failure of Service A.
2205 The exposed interfaces of Services B and C (IService B and IServiceC) are not necessarily hidden from
2206 the Consumer Delegate; only the fact that these services are used as part of the composition of Service
2207 A. In this example, there is no practical reason the Consumer Delegate could not interact with Service B
2208 or Service C in some other interaction scenario.

2209 ~~While the service composition is opaque from the Consumer Delegate's perspective, it is transparent to~~
2210 ~~the service owner. This transparency is necessary for service management. It is possible for a service~~
2211 ~~composition to be opaque from one perspective and transparent from another. For example, a service~~
2212 ~~may appear to be a single service from the Consumer's Delegate's perspective, but is transparently~~
2213 ~~composed of one or more services from a service management perspective. A Service Management~~
2214 ~~capability needs to be able to have visibility into the composition in order~~ to properly manage the
2215 dependencies between the services used in constructing the composite service—including managing the
2216 service's lifecycle. The subject of services as management entities is described and modeled in the
2217 *Ownership in a SOA Ecosystem* View of the SOA-RAF and is not further elaborated in this section. The
2218 point to be made here is that there can be different levels of opaqueness or transparency when it comes
2219 to visibility of service composition.

Comment [KJL75]: Issue 93

2220 Services can be composed in a variety of ways, including direct consumer-to-service interaction, by using
2221 programming techniques, or ~~using an intermediary, such as an orchestration engine leveraging higher~~
2222 ~~level orchestration languages, they can be aggregated by means of an aggregation engine approach that~~
2223 ~~leverages a service composition scripting language.~~ Such approaches are further elaborated in the
2224 following sub-sections ~~on service-oriented business processes and collaborations.~~

Comment [KJL76]: Issue 94, 204

2225 4.3.5 Implementing Service Composition

2226 Services are implemented through a combination of processes and collaboration. The concepts involved
2227 and that would be used in the context of exchanges both within and across organizational boundaries are
2228 described and modeled as part of the *Participation in a SOA Ecosystem* view of this reference
2229 architecture (see Section 3).

Comment [PFB77]:
Whole of Section 4.3.5 re-written

2230 The principles involved in the composition of services (including but not limited to loose coupling,
2231 selective transparency and opacity, dynamic interactions) are equally applicable to services which
2232 implement business processes and collaborations. Business processes and collaborations represent
2233 complex, multi-step business functions that may involve multiple participants, including internal users,
2234 external customers, and trading partners. Therefore, such complexities cannot simply be ignored when
2235 transforming traditional business processes and collaborations to their service-oriented variants.

Comment [KJL78]: Issue 264

2236 While business processes are primarily concerned with describing how services are invoked and
2237 executed, business collaborations are more concerned with how actors (usually from different
2238 organizations) interact to realize a desired effect.

2239 Collaborations can include processes (for example, when one actor executes a particular activity
2240 according to the predefined steps of a process) as much as processes can include collaborations (a
2241 predefined step of a particular process may include agreed-upon activities provided by other participants).

2242 The techniques discussed below can be applied to any combination of services that instantiate service-
2243 oriented business processes or service-oriented business collaborations.

2244 4.3.5.1 Service-Oriented Business Processes

2245 Service orientation as applied to a business process includes

- 2246 • abstracting the set of activities and rules governing a business process; and
- 2247 • composing and exposing the resultant logic as a reusable service.

Comment [PFB79]: Issue 208

2248 When business processes are implemented as SOA services, all of the concepts used to describe and
2249 model composition of services that were articulated in Section 4.3.4 apply.

Comment [KJL80]: Issues 95, 209

2250 Business processes have temporal properties and can be short-lived or long-lived. Further, these
2251 processes may involve many participants and may be important considerations for the consumer of a
2252 service-oriented business process. For example, a consumer may need to know certain details of the
2253 business process in order to have confidence in the resulting real world effects. For business processes

Comment [PFB81]: Issues 96, 210,
216, and 265 (obsolete with re-
write)

2254 implemented as SOA-based services, ensuring that the meta-level aspects of the service-oriented
 2255 business process are included in its Service Description can provide needed insight for the consumer.

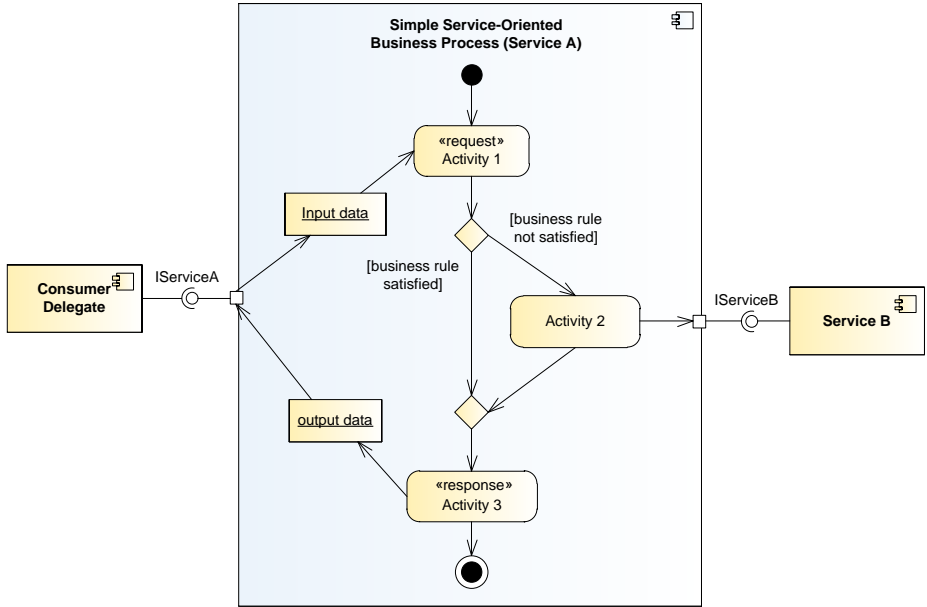


Figure 32 - Abstract example of a simple business process exposed as a service

Comment [PFB82]: Issue 213

2256
2257

2258 In Figure 32, we use a UML activity diagram to model the simple service-oriented business process. This
 2259 allows us to capture the major elements, such as the set of related activities to be performed (an activity
 2260 being made up of one or more related actions, as explained in Section 3.3.2); the links between these
 2261 activities in a logical flow; data that is passed between activities, and any relevant business rules that
 2262 govern the transitions between the activities. While specific actions and activities can be readily modeled
 2263 in more detail, they are not illustrated in the model in Figure 32.

2264 This example is based on a request/response MEP and captures how one process can leverage
 2265 fulfillment of a particular activity (Activity 2) leverages by calling upon an externally-provided service,
 2266 Service B. The entire service-oriented business process is exposed as Service A that is accessible via its
 2267 externally visible interface, IServiceA. It is the availability of this external interface, and the description of
 2268 what the service intends, that distinguishes this from a simple business process.

2269 Although not explicitly shown in the model above, it is assumed that there exists a software or hardware
 2270 component that executes the process flow (Functionality of Service A). However, human actors may also
 2271 take part. This may be particularly important in cases where the automation fails and human intervention
 2272 becomes necessary.

Comment [KJL83]: Issue 235

2273 4.3.5.2 Service-Oriented Business Collaborations

2274 Whereas a service can execute according to a predefined business process determined by one
 2275 organization, service composition can also be accomplished as a cooperation, or business collaboration,
 2276 between actors in different organizations and systems.

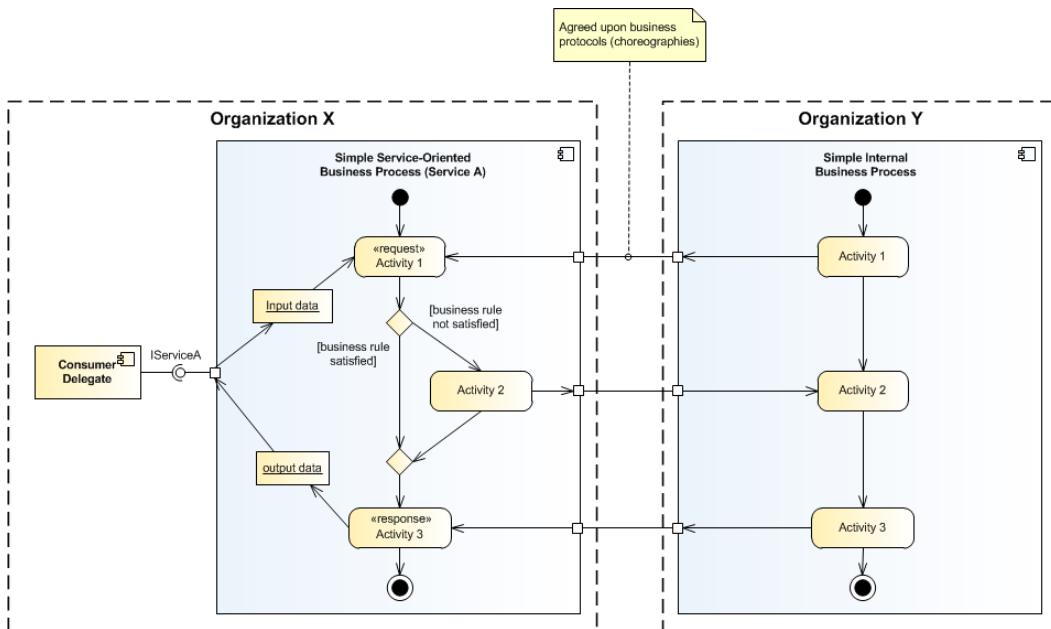
2277 In a service-oriented business collaboration, multiple participants interact in a peer-style communication
 2278 as part of some larger business transaction by exchanging messages with trading partners and external
 2279 organizations (e.g., suppliers) [NEWCOMER/LOMOW]. Participants do not necessarily expose the
 2280 entirety of their respective capabilities but rather use service-based interactions to access those
 2281 capabilities needed to fulfill the collaboration.

Comment [KJL84]:
Issues 98, 218-221.

2282 Service-orientation as applied to a business collaboration includes:

- 2283 - ability of participants to individually provide and commit to what is required during an interaction for a
- 2284 collaboration to be successfully realized, including acceptance of preconditions and expected
- 2285 outcome;
- 2286 - availability of service functionality sufficient to realize the effects expected from the business
- 2287 collaboration;
- 2288 - willingness of participants to engage in interactions that are required as part of the collaboration;
- 2289 - availability of shared state and notifications to all participants who require them, such that they can
- 2290 fulfill their respective parts of the collaboration.

2291 Any service contributing to such a service-oriented business collaboration participates “as is”, without
 2292 modification, and consistent with its own service description. Each contributing service is only an
 2293 instrument in the collaboration and is not typically “aware” of its own contribution except as would be
 2294 conveyed through inputs, access to shared states, or event notifications that are generally available to the
 2295 service.



2296
 2297 *Figure 33 - Abstract example of a more complex composition that relies on collaboration*

2298 Figure 33, which is a variant of the example illustrated earlier (in Figure 32), includes trust boundaries
 2299 between two organizations; namely, Organization X and Organization Y. It is assumed that these two
 2300 organizations are peer entities that have an interest in a business collaboration, for example,
 2301 Organization X and Organization Y could be trading partners. Organization X retains the service-oriented
 2302 business process Service A, which is exposed to internal consumers via its provided service interface,
 2303 IServiceA. Organization Y also has a business process that is involved in the business collaboration; in
 2304 this example, it is an internal business process but it could also be exposed to potential consumers either
 2305 within or outside its organizational boundary if it is designed as a reusable service in accordance with
 2306 SOA design principles.

2307 In Figure 33, the communications between Organization X and Organization Y are shown through ports
 2308 where there are “agreed-upon business protocols” that also cover the order in which activities are carried
 2309 out. These ports do not explicitly show service interfaces in order to emphasize that in the example these
 2310 are not intended to be generally available to any actor in the SOA ecosystem; however, the interfaces
 2311 should adhere to the principles involved in the composition of services.

2312 The message exchanges that are used need to specify how and when to initiate activity by the other
 2313 trading partner, i.e., communication between Organization X and Organization Y. Defining the business
 2314 protocols used in the business collaboration involves precisely specifying the visible message exchange

Comment [PFB85]: Issue 225
 (Obsolete)

2315 behavior and order of each of the parties involved in the protocol, without revealing internal
2316 implementation details [NEWCOMER/LOMOW]. This is consistent with the Information and Behavior
2317 Models discussed in the Reference Model and as part of service description in section 4.1.

2318 | Business processes and collaboration are thus both facets of SOA service composition. The degree to
2319 which one predominates over the other (and the mix of the two that emerges) will be a reflection of many
2320 factors including the relative autonomy of participants and actors, the desired flexibility of a system, the
2321 extent of trust involved and the assessment of risk, among others.

Comment [PFB86]: Issue 226
(Obselete)

Comment [KJL87]: Issue 236
(Obselete)

2322 4.3.6 Architectural Implications of Interacting with Services

2323 Interacting with Services has the following architectural implications on mechanisms that facilitate service
2324 interaction:

- 2325 | • A well-defined service Information Model **MUST be provided** that:
 - 2326 ○ describes the syntax and semantics of the messages used to denote actions and events;
 - 2327 ○ describes the syntax and semantics of the data payload(s) contained within messages;
 - 2328 ○ documents exception conditions in the event of faults due to network outages, improper
 - 2329 message/data formats, etc.;
 - 2330 ○ is both human readable and machine processable;
 - 2331 ○ is referenceable from the Service Description artifact.
- 2332 | • A well-defined service Behavior Model (as defined in the SOA-RM) **MUST be provided** that:
 - 2333 ○ characterizes the knowledge of the actions invoked against the service and events that
 - 2334 report real world effects as a result of those actions;
 - 2335 ○ characterizes the temporal relationships and temporal properties of actions and events
 - 2336 associated in a service interaction;
 - 2337 ○ describe activities involved in a workflow activity that represents a unit of work;
 - 2338 ○ describes the role (s) performed in a service-oriented business process or service-
 - 2339 oriented business collaboration;
 - 2340 ○ is both human readable and machine processable;
 - 2341 ○ is referenceable from the Service Description artifact.
- 2342 | • Mechanisms **MUST be included** to support ~~orchestration of composition of~~ service-oriented business
2343 processes and ~~choreography of~~ service-oriented business collaborations such as:
 - 2344 ○ Declarative and programmatic compositional languages;
 - 2345 ○ Orchestration and/or choreography engines that support multi-step processes as part of a
 - 2346 short-lived or long-lived business transaction;
 - 2347 ○ Orchestration and/or choreography engines that support compensating transactions in
 - 2348 the presences of exception and fault conditions.
- 2349 | • Infrastructure services **MUST be specified** that provides mechanisms to support service interaction,
2350 including but not limited to:
 - 2351 ○ mediation ~~services within service interactions based on shared such as message and~~
2352 ~~event brokers, providers, and/or buses that provide message translation/transformation,~~
2353 ~~gateway capability, message persistence, reliable message delivery, and/or intelligent~~
2354 ~~routing~~ semantics;
 - 2355 ○ ~~binding services that support~~ translation and transformation of multiple application-level
2356 protocols to standard network transport protocols;
 - 2357 ○ auditing and logging ~~services~~ that provide a data store and mechanism to record
2358 information related to service interaction activity such as message traffic patterns,
2359 security violations, and service contract and policy violations
 - 2360 ○ security ~~services~~ that provides ~~centralized~~ authorization and authentication support, etc.,
2361 which provide protection against common security threats in a SOA ecosystem;
 - 2362 ○ monitoring ~~services~~ such as hardware and software mechanisms that both monitor the
2363 performance of systems that host services and network traffic during service interaction,
2364 and are capable of generating regular monitoring reports.
- 2365 | • In a service-oriented business collaboration, any language used **MUST be capable of describing the**
2366 coordination required of those service-oriented business processes that cross organizational
2367 boundaries. This **SHOULD** provide for contingencies, in case of an upset or when automation fails,
2368 including any necessary human intervention.

Comment [KJL88]: Issue 227

Comment [KJL89]: Issue 300

Comment [PFB90]: Issue 101

Comment [PFB91]: Issue 229

- 2369 • A layered and tiered service component architecture that supports multiple message exchange
2370 patterns (MEPs) in order to:
- 2371 ○ promote the industry best practice of separation of concerns that facilitates flexibility in
2372 the presence of changing business requirements;
 - 2373 ○ promote the industry best practice of separation of roles in a service development
2374 lifecycle such that subject matter experts and teams are structured along areas of
2375 expertise;
 - 2376 ○ support numerous standard interaction patterns, peer-to-peer interaction patterns,
2377 enterprise integration patterns, and business-to-business integration patterns.

Comment [PFB92]: Issue 102, 230

2378 4.4 Policies and Contracts Model

2379 A common phenomenon of many machines and systems is that the scope of potential behavior is much
2380 broader than is actually needed for a particular circumstance. This is especially true of a system as
2381 powerful as a SOA ecosystem. As a result, the behavior and performance of the system tend to be under-
2382 constrained by the implementation; instead, the actual behavior is expressed by means of policies of
2383 some form. Policies define the choices that stakeholders make; these choices are used to guide the
2384 actual behavior of the system to the desired behavior and performance.

2385 As noted in Section 3.2.5.2, a policy is an expression of constraints that is promulgated by a stakeholder
2386 who has the responsibility of ensuring that the constraint is enforceable. In contrast, contracts are
2387 agreements between participants. However, like policies, it is a necessary part of contracts that they are
2388 enforceable.

Comment [PFB93]: Issue 232

Comment [PFB94]: Issue 232

2389 While responsibility for enforcement may differ, both contracts and policies share a common characteristic
2390 – there is a constraint that must be enforced. In both cases, the mechanisms needed to enforce
2391 constraints are likely to be identical; in this model, we focus on the issues involved in representing
2392 policies and contracts and on some of the principles behind their enforcement.

2393 4.4.1 Policy and Contract Representation

2394 A policy constraint is a specific kind of constraint: the ontology of policies and contracts includes the core
2395 concepts of permission, obligation, owner, and subject. In addition, it may be necessary to be able to
2396 combine policy constraints and to be able to resolve policy conflicts.

2397 4.4.1.1 Policy Framework

2398 Policy Framework

2399 A policy framework is a language in which policy constraints may be expressed.

2400 A policy framework combines syntax for expressing policy constraints together with a decision procedure
2401 for determining if a policy constraint is satisfied.




Figure 34 - Policies and Contracts

2402
2403

2404 We can characterize a policy framework in terms of a **logical framework** and an ontology of policies. The
2405 **policy ontology** details specific kinds of policy constraints that can be expressed; and the logical
2406 framework is a 'glue' that allows us to express combinations of policies.

2407 **Logical Framework**

2408 A linguistic framework consisting of a syntax – a way of writing expressions – and a semantics –
2409 a way of interpreting the expressions.

2410 **Policy Ontology**

2411 A formalization of a set of concepts that are relevant to forming policy expressions.

2412 For example, a policy ontology that allows identification of simple constraints – such as the existence of a
2413 property, or that a value of a property should be compared to a fixed value – is often enough to express
2414 many basic constraints.

2415 Included in many policy ontologies are the basic signals of permissions and obligations. Some policy
2416 frameworks are sufficiently constrained that there is no possibility of representing an obligation; in which
2417 case there is often no need to 'call out' the distinction between permissions and obligations.

2418 The logical framework is also a strong determiner of the expressivity of the policy framework: the richer
2419 the logical framework, the richer the set of policy constraints that can be expressed. However, there is a
2420 strong inverse correlation such that increasing expressivity yields less ease and greater inefficiency of
2421 implementation.

Comment [PFB95]: Issue 272

2422 In the discussion that follows we assume the following basic policy ontology:

2423 **Policy Owner**

2424 A **stakeholder** that asserts and enforces the **policy**.

2425 **Policy Subject**

2426 An **actor** whose action, or a resource whose maintenance or use, is constrained by a **policy-or-**
2427 **contract**.

Comment [KJL96]: Issue 310

2428 **Policy Constraint**

2429 A measurable and enforceable proposition-assertion that characterizes the constraint that the
2430 found within a policy-is-about.

Comment [KJL97]: Issue 311

2431 **Policy Object**

2432 An identifiable **state, action** or **resource** that is potentially constrained by the **policy**.

2433 **4.4.2 Policy and Contract Enforcement**

2434 The enforcement of policy constraints has to address two core problems: how to enforce the atomic policy
2435 constraints, and how to enforce combinations of policy constraints. In addition, it is necessary to address
2436 the resolution of policy conflicts. Contracts are the documented agreement between two or more parties
2437 but otherwise have the same enforcement requirements as policies.

Comment [KJL98]: Issue 237

2438 4.4.2.1 Enforcing Simple Policy Constraints

2439 The two primary kinds of policy constraint – permission and obligation – naturally lead to different styles
2440 of enforcement. A permission constraint must typically be enforced prior to the policy subject invoking the
2441 policy object. On the other hand, an obligation constraint must typically be enforced after the fact through
2442 some form of auditing process and remedial action.

2443 For example, if a communications policy required that all communication be encrypted, this is enforceable
2444 at the point of communication: any attempt to communicate a message that is not encrypted can be
2445 blocked.

2446 Similarly, an obligation to pay for services rendered is enforced by ensuring that payment arrives within a
2447 reasonable period of time. Invoices are monitored for prompt (or lack of) payment.

2448 The key concepts in enforcing both forms of policy constraint are the policy decision and the policy
2449 enforcement.

2450 **Policy Decision**

2451 A determination as to whether a given **policy constraint** is satisfied.

2452 A policy decision is effectively a measurement of some state – typically a portion of the SOA ecosystem's
2453 **shared state**. This implies a certain *timeliness* in the measuring: a measurement that is too early or is too
2454 late does not actually help in determining if the policy constraint is satisfied appropriately.

2455 **Policy Enforcement**

2456 A mechanism that limits the behavior and/or **state of policy subjects** to comply with a **policy**
2457 **decision**.

2458 A policy enforcement implies the use of some mechanism to ensure compliance with a policy decision.

2459 The range of mechanisms is completely dependent on the kinds of atomic policy constraints that the
2460 policy framework may support. As noted above, the two primary styles of constraint – permission and
2461 **obligation** –lead to different styles of enforcement.

2462 4.4.2.2 Conflict Resolution

2463 Whenever it is possible that more than one policy constraint applies in a given situation, there is the
2464 potential that the policy constraints themselves are not mutually consistent. For example, a policy
2465 constraint that requires communication to be encrypted and a policy constraint that requires an
2466 administrator to read every communication conflict with each other – the two policy constraints cannot
2467 both be satisfied concurrently.

2468 In general, with sufficiently rich policy frameworks, it is not possible to always resolve policy conflicts
2469 automatically. However, a reasonable approach is to augment the policy decision process with simple
2470 policy conflict resolution rules; with the potential for *escalating* a policy conflict to human adjudication.

2471 **Policy Conflict**

2472 A state in a **policy decision** process in which the satisfaction of one or more **policy constraints**
2473 leads directly to the violation of one or more other policy constraints.

2474 **Policy Conflict Resolution**

2475 A rule determining which policy constraint(s) should prevail if a **policy conflict** occurs.

2476 The inevitable consequence of policy conflicts is that it is not possible to guarantee that all policy
2477 constraints are satisfied at all times. This, in turn, implies certain *flexibility* in the application of policy
2478 constraints: each individual constraint may not always be honored.

2479 4.4.3 Architectural Implications

2480 The key choices that must be made in a system of policies center on the policy framework, policy
2481 enforcement, and conflict resolution

- 2482 • There **SHOULD** be a standard policy framework that is adopted across ownership domains within the
2483 SOA ecosystem:

- 2484 ○ This framework **MUST** permit the expression of simple policy constraints
- 2485 ○ The framework **MAY** allow (to a varying extent) the combination of policy constraints,
- 2486 including
 - 2487 • Both positive and negative constraints
 - 2488 • Conjunctions and disjunctions of constraints
 - 2489 • The quantification of constraints
- 2490 ○ The framework **MUST** at least allow the policy subject and the policy object to be identified as
- 2491 well as the policy constraint.
- 2492 ○ The framework **MAY** allow further structuring of policies into modules, inheritance between
- 2493 policies and so on.
- 2494 • There **SHOULD** be mechanisms that facilitate the application of policies:
 - 2495 ○ There **SHOULD** be mechanisms that allow policy decisions to be made, consistent with the
 - 2496 policy frameworks.
 - 2497 ○ There **SHOULD** be mechanisms to enforce policy decisions
 - 2498 • There **SHOULD** be mechanisms to support the measurement of whether certain
 - 2499 policy constraints are satisfied, or to what degree they are satisfied.
 - 2500 • Such enforcement mechanisms **MAY** include support for both permission-style
 - 2501 constraints and obligation-style constraints.
 - 2502 • Enforcement mechanisms **MAY** support the simultaneous enforcement of multiple
 - 2503 policy constraints across multiple points in the SOA ecosystem.
 - 2504 ○ There **SHOULD** be mechanisms to resolve policy conflicts
 - 2505 • This **MAY** involve escalating policy conflicts to human adjudication.
 - 2506 ○ There **SHOULD** be mechanisms that support the management and promulgation of policies.

2507 **5 Ownership in a SOA Ecosystem View**

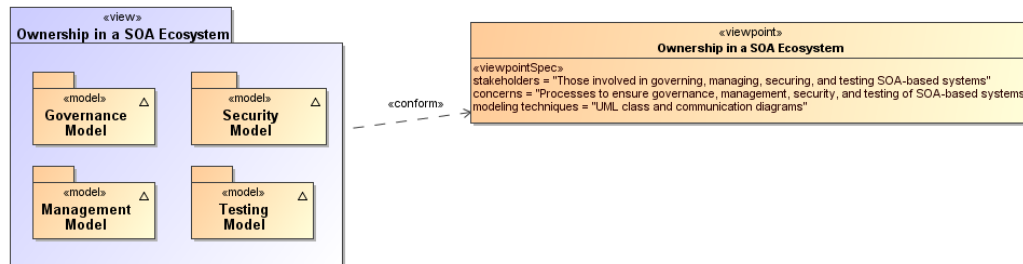
2508 *Governments are instituted among Men,*
2509 *deriving their just power from the consent of the governed*
2510 *American Declaration of Independence*

2511
2512 The *Ownership in a SOA Ecosystem View* focuses on the issues, requirements and responsibilities
2513 involved in owning a SOA-based system.

2514 Ownership of a SOA-based system in a SOA ecosystem raises significantly different challenges to
2515 owning other complex systems – such as Enterprise suites – because there are strong limits on the
2516 control and authority of any one party when a system spans multiple ownership domains.

2517 Even when a SOA-based system is deployed internally within an organization, there are multiple internal
2518 stakeholders involved and there may not be a simple hierarchy of control and management. Thus, an
2519 early consideration of how multiple boundaries affect SOA-based systems provides a firm foundation for
2520 dealing with them in whatever form they are found rather than debating whether the boundaries should
2521 exist.

2522 This view focuses on the governance and management of SOA-based systems, on the security
2523 challenges involved in running a SOA-based system, and testing challenges.



2524
2525 *Figure 35 - Model Elements Described in the Ownership in a SOA Ecosystem View*

2526 The following subsections present models of these functions.

2527 **5.1 Governance Model**

2528 The Reference Model defines Service Oriented Architecture as an architectural paradigm for organizing
2529 and utilizing distributed capabilities that may be under the control of different ownership domains [SOA-
2530 RM]. Consequently, it is important that organizations that plan to engage in service interactions adopt
2531 governance policies and procedures sufficient to ensure that there is standardization across both internal
2532 and external organizational boundaries to promote the effective creation and use of SOA-based services.

2533 **5.1.1 Understanding Governance**

2534 **5.1.1.1 Terminology**

2535 Governance is about making decisions that are aligned with the overall organizational strategy and
2536 culture of the enterprise. [Gartner] It specifies the decision rights and accountability framework to
2537 encourage desirable behaviors [Weill/Ross-MIT Sloan School] towards realizing the strategy and
2538 defines incentives (positive or negative) towards that end. It is less about overt control and strict
2539 adherence to rules, and more about guidance and effective and equitable usage of resources to ensure
2540 sustainability of an organization's strategic objectives. [TOGAF v8.1]

2541 To accomplish this, governance requires organizational structure and processes and must identify who
2542 has authority to define and carry out its mandates. It must address the following questions:

- 2543 1. what decisions must be made to ensure effective management and use?,
- 2544 2. who should make these decisions?,
- 2545 3. how will these decisions be made and monitored? , and
- 2546 4. how will these decisions be communicated?

2547 The intent is to achieve goals, add value, and reduce risk.

2548 Within a single ownership domain such as an enterprise, generally there is a hierarchy of governance
2549 structures. Some of the more common enterprise governance structures include corporate governance,
2550 technology governance, IT governance, and architecture governance **[TOGAF v8.1]**. These governance
2551 structures can exist at multiple levels (global, regional, and local) within the overall enterprise.

2552 It is often asserted that SOA governance is a specialization of IT governance as there is a natural
2553 hierarchy of these types of governance structures; however, the focus of SOA governance is less on
2554 decisions to ensure effective management and use of IT as it is to ensure effective management and use
2555 of SOA-based systems. Certainly, SOA governance must still answer the basic questions also associated
2556 with IT governance, i.e., who should make the decisions, and how these decisions will be made and
2557 monitored.

2558 **5.1.1.2 Relationship to Management**

2559 There is often confusion centered on the relationship between governance and management. As
2560 described earlier, governance is concerned with decision making. Management, on the other hand, is
2561 concerned with execution. Put another way, governance describes the world as **leadership** wants it to
2562 be; management executes activities that intend to make the leadership's desired world a reality. Where
2563 governance determines who has the authority and responsibility for making decisions and the
2564 establishment of guidelines for how those decisions should be made, management is the actual process
2565 of making, implementing, and measuring the impact of those decisions **[Loeb]**. Consequently,
2566 governance and management work in concert to ensure a well-balanced and functioning organization as
2567 well as an ecosystem of inter-related organizations. In the sections that follow, we elaborate further on the
2568 relationship between governance and management in terms of setting and enforcing service policies,
2569 contracts, and standards as well as addressing issues surrounding regulatory compliance.

2570 **5.1.1.3 Why is SOA Governance Important?**

2571 One of the hallmarks of SOA that distinguishes it from other architectural paradigms for distributed
2572 computing is the ability to provide a uniform means to offer, discover, interact with and use capabilities
2573 (as well the ability to compose new capabilities from existing ones) all in an environment that transcends
2574 domains of ownership. Consequently, ownership, and issues surrounding it, such as obtaining acceptable
2575 terms and conditions (T&Cs) in a contract, is one of the primary topics for SOA governance. Generally, IT
2576 governance does not include T&Cs, for example, as a condition of use as its primary concern.

2577 Just as other architectural paradigms, technologies, and approaches to IT are subject to change and
2578 evolution, so too is SOA. Setting policies that allow change management and evolution, establishing
2579 strategies for change, resolving disputes that arise, and ensuring that SOA-based systems continue to
2580 fulfill the goals of the business are all reasons why governance is important to SOA.

2581 **5.1.1.4 Governance Stakeholders and Concerns**

2582 As noted in Section 3.2.1 the participants in a service interaction include the service provider, the service
2583 consumer, and other interested or unintentional third parties. Depending on the circumstances, it may
2584 also include the owners of the underlying capabilities that the SOA services access. Governance must
2585 establish the policies and rules under which duties and responsibilities are defined and the expectations
2586 of participants are grounded. The expectations include transparency in aspects where transparency is
2587 mandated; trust in the impartial and consistent application of governance; and assurance of reliable and
2588 robust behavior throughout the SOA ecosystem.

2589 **5.1.2 A Generic Model for Governance**

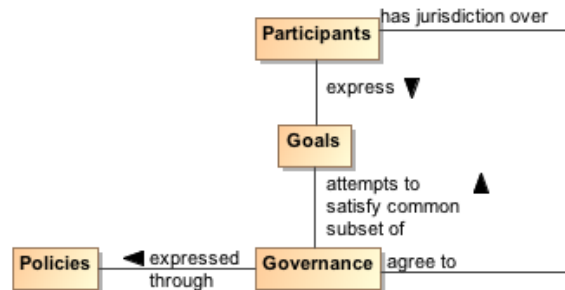
2590 **Governance**

2591 The prescription of conditions and constraints consistent with satisfying common goals and the
2592 structures and processes needed to define and respond to actions taken towards realizing those
2593 goals.

2594 The following is a generic model of governance represented by segmented models that begin with
2595 motivation and proceed through measuring compliance. It is not all-encompassing but a focused subset
2596 that captures the aspects necessary to describe governance for SOA. It does not imply that practical
2597 application of governance is a single, isolated instance of these models; in reality, there may be
2598 hierarchical and parallel chains of governance that deal with different aspects or focus on different goals.
2599 This is discussed further in section 5.1.2.5. The defined models are simultaneously applicable to each of
2600 the overlapping instances.

2601 A given enterprise may already have portions of these models in place. To a large extent, the models
2602 shown here are not specific to SOA; discussions on direct applicability begin in section 5.1.3.

2603 **5.1.2.1 Motivating Governance**



2604 *Figure 36 - Motivating Governance*

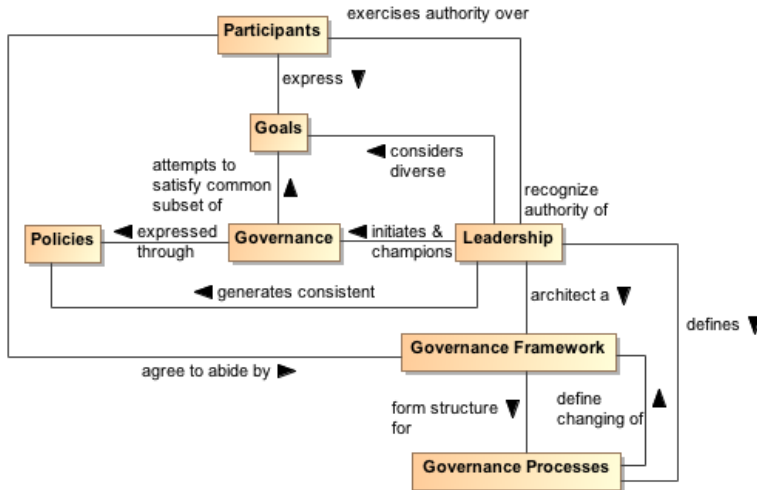
2605 An organizational domain such as an enterprise is made up of participants who may be individuals or
2606 groups of individuals forming smaller organizational units within the enterprise. The overall business
2607 strategy should be consistent with the goals of the participants; otherwise, the business strategy would
2608 not provide value to the participants and governance towards those ends becomes difficult if not
2609 impossible. This is not to say that an instance of governance simultaneously satisfies all the goals of all
2610 the participants; rather, the goals of any governance instance must sufficiently satisfy a useful subset of
2611 each participant's goals so as to provide value and ensure the cooperation of all the participants.
2612

2613 A policy is the formal characterization of the conditions and constraints that governance deems as
2614 necessary to realize the goals which it is attempting to satisfy. Policy may identify required conditions or
2615 actions or may prescribe limitations or other constraints on permitted conditions or actions. For example,
2616 a policy may prescribe that safeguards must be in place to prevent unauthorized access to sensitive
2617 material. It may also prohibit use of computers for activities unrelated to the specified work assignment.
2618 Policy is made operational through the promulgation and implementation of Rules and Regulations (as
2619 defined in section 5.1.2.3).

2620 As noted in section 4.4.2, policy may be asserted by any participant or on behalf of the participant by its
2621 organization. Part of the purpose of governance is to arbitrate among diverse goals of participants and
2622 the diverse policies articulated to realize those goals. The intent is to form a consistent whole that allows
2623 governance to minimize ambiguity about its purpose. While resolving all ambiguity would be an ideal, it is
2624 unlikely that all inconsistencies will be identified and resolved before governance becomes operational.

2625 For governance to have effective jurisdiction over participants, there must be some degree of agreement
2626 by all participants that they will abide by the governance mandates. A minimal degree of agreement often
2627 presages participants who 'slow-roll' if not actively rejecting compliance with policies that express the
2628 specifics of governance.

2629 **5.1.2.2 Setting Up Governance**



2630
2631 *Figure 37 - Setting Up Governance*

2632 **Leadership**

2633 The entity having the **responsibility** and **authority** to generate consistent **policies** through which
2634 the goals of **governance** can be expressed and to define and champion the structures and
2635 processes through which governance is realized.

2636 **Governance Framework**

2637 The set of organizational structures that enable **governance** to be consistently defined, clarified,
2638 and as needed, modified to respond to changes in its domain of concern.

2639 **Governance Process**

2640 The defined set of activities performed within the **Governance Framework** to enable the
2641 consistent definition, application, and as needed, modification of **rules** that organize and regulate
2642 the activities of participants for the fulfillment of expressed policies. (See section 5.1.2.3 for
2643 elaboration on the relationship of Governance Processes and Rules.)

2644 As noted earlier, governance requires an appropriate organizational structure and identification of who
2645 has authority to make governance decisions. In Figure 37, the entity with governance authority is
2646 designated the Leadership. This is someone, possibly one or more of the participants, which participants
2647 recognize as having authority for a given purpose or over a given set of issues or concerns.

2648 The leadership is responsible for prescribing or delegating a working group to prescribe the governance
2649 framework that forms the structure for governance processes that define how governance is to be carried
2650 out. This does not itself define the specifics of how governance is to be applied, but it does provide an
2651 unambiguous set of procedures that should ensure consistent actions which participants agree are fair
2652 and account for sufficient input on the subjects to which governance is applied.

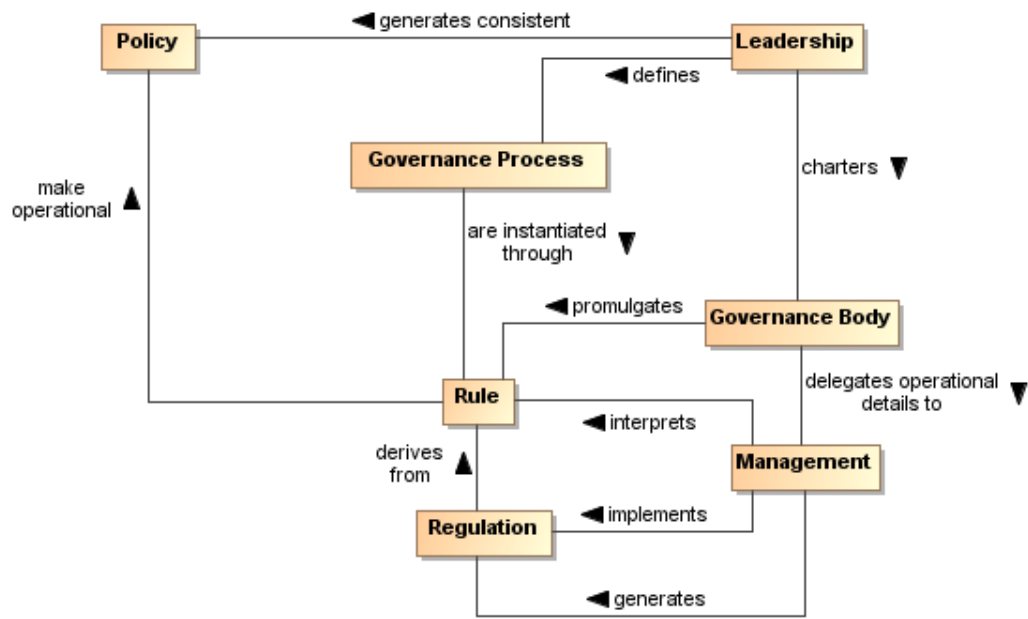
2653 The participants may be part of the working group that codifies the governance framework and
2654 processes. When complete, the participants must acknowledge and agree to abide by the products
2655 generated through application of this structure.

2656 The governance framework and processes are often documented in the constitution or charter of a body
2657 created or designated to oversee governance. This is discussed further in the next section. Note that the
2658 governance processes should also include those necessary to modify the governance framework itself.

2659 An important function of leadership is not only to initiate but also be the consistent champion of
2660 governance. Those responsible for carrying out governance mandates must have leadership who make it

2661 clear to participants that expressed policies are seen as a means to realizing established goals and that
2662 compliance with governance is required.

2663 5.1.2.3 Carrying Out Governance



Comment [KL99]: Issue 115, part

Figure 38 - Carrying Out Governance

2664
2665
2666 **Rule**

2667 A prescribed guide for carrying out activities and processes leading to desired results, e.g. the
2668 operational realization of policies.

2669 **Regulation**

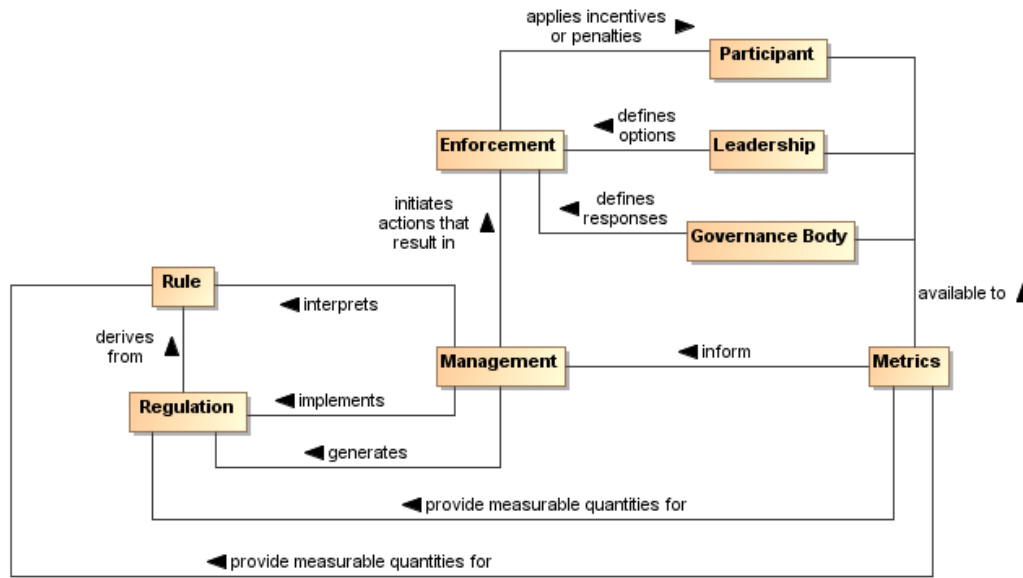
2670 A mandated process or the specific details that derive from the interpretation of **rules** and lead to
2671 measureable quantities against which compliance can be measured.

2672 To carry out governance, leadership charters a governance body to promulgate the rules needed to make
2673 the policies operational. The governance body acts in line with governance processes for its rule-making
2674 process and other functions. Whereas governance is the setting of policies and defining the rules that
2675 provide an operational context for policies, governance body may delegate the operational details of
2676 governance to management. Management generates regulations that specify details for rules and other
2677 procedures to implement both rules and regulations. For example, leadership could set a policy that all
2678 authorized parties should have access to data, the governance body would promulgate a rule that PKI
2679 certificates are required to establish identity of authorized parties, and management can specify a
2680 regulation of who it deems to be a recognized PKI issuing body. In summary, policy is a predicate to be
2681 satisfied and rules prescribe the activities by which that satisfying occurs. A number of rules may be
2682 required to satisfy a given policy; the carrying out of a rule may contribute to several policies being
2683 realized.

2684 Whereas the governance framework and processes are fundamental for having participants acknowledge
2685 and commit to compliance with governance, the rules and regulations provide operational constraints that
2686 may require resource commitments or other levies on the participants. It is important for participants to
2687 consider the framework and processes to be fair, unambiguous, and capable of being carried out in a
2688 consistent manner and to have an opportunity to formally accept or ratify this situation. rules and
2689 regulations, however, do not require individual acceptance by any given participant although some level

2690 of community comment may be part of the governance processes. Having agreed to governance, the
 2691 participants are bound to comply or be subject to prescribed mechanisms for enforcement.

2692 **5.1.2.4 Ensuring Governance Compliance**



Comment [PFB100]: Issue 115, part

2693
 2694 *Figure 39 - Ensuring Governance Compliance*

2695 Setting rules and regulations does not ensure effective governance unless compliance can be measured
 2696 and rules and regulations can be enforced. Metrics are those conditions and quantities that can be
 2697 measured to characterize actions and results. Rules and regulations must be based on collected metrics,
 2698 or there is no means for management to assess compliance. The metrics are available to the participants,
 2699 the leadership, and the governance body so what is measured and the results of measurement are clear
 2700 to everyone.

2701 The leadership in its relationship with participants has certain options that can be used for enforcement. A
 2702 common option may be to affect future funding. The governance body defines specific enforcement
 2703 responses, such as what degree of compliance is necessary for full funding to be restored. It is up to
 2704 management to identify compliance shortfalls and to initiate the enforcement process.

2705 Note, enforcement does not strictly need to be negative consequences. Management can use metrics to
 2706 identify exemplars of compliance and leadership can provide options for rewarding the participants. The
 2707 governance body defines awards or other incentives.

2708 **5.1.2.5 Considerations for Multiple Governance Chains**

2709 As noted in section 5.1.2, instances of the governance model often occur as a tiered arrangement, with
 2710 governance at some level delegating specific authority and responsibility to accomplish a focused portion
 2711 of the original level's mandate. For example, a corporation may encompass several lines of business and
 2712 each line of business governs its own affairs in a manner that is consistent with and contributes to the
 2713 goals of the parent organization. Within the line of business, an IT group may be given the mandate to
 2714 provide and maintain IT resources, giving rise to IT governance.

2715 In addition to tiered governance, there may be multiple governance chains working in parallel. For
 2716 example, a company making widgets has policies intended to ensure they make high quality widgets and
 2717 make an impressive profit for their shareholders. On the other hand, Sarbanes-Oxley is a parallel
 2718 governance chain in the United States that specifies how the management must handle its accounting

2719 and information that needs to be given to its shareholders. The parallel chains may just be additive or
2720 may be in conflict and require some harmonization.
2721 Being distributed and representing different ownership domains, a SOA participant falls under the
2722 jurisdiction of multiple governance domains simultaneously and may individually need to resolve
2723 consequent conflicts. The governance domains may specify precedence for governance conformance or
2724 it may fall to the discretion of the participant to decide on the course of actions they believe appropriate.

2725 **5.1.3 Governance Applied to SOA**

2726 **5.1.3.1 Where SOA Governance is Different**

2727 SOA governance is often discussed in terms of IT governance, but rather than a parent-child relationship,
2728 Figure 40 shows the two as siblings within the general governance described in section 5.1.2. There are
2729 obvious dependencies and a need for coordination between the two, but the idea of aligning IT with
2730 business already demonstrates that resource providers and resource consumers must be working
2731 towards common goals if they are to be productive and efficient. While SOA governance is shown to be
2732 active in the area of infrastructure, it is a specialized concern for having a dependable platform to support
2733 service interaction; a range of traditional IT issues is therefore out of scope of this document. A SOA
2734 governance plan for an enterprise will not of itself resolve shortcomings with the enterprise's IT
2735 governance.

2736 Governance in the context of SOA is that organization of services: that promotes their visibility; that
2737 facilitates interaction among service participants; and that directs that the results of service interactions
2738 are those real world effects as described within the service description and constrained by policies and
2739 contracts as assembled in the execution context.

2740 SOA governance must specifically account for control across different ownership domains, i.e. all the
2741 participants may not be under the jurisdiction of a single governance authority. However, for governance
2742 to be effective, the participants must agree to recognize the authority of the governance body and must
2743 operate within the Governance Framework and through the Governance Processes so defined.

2744 SOA governance must account for interactions across ownership boundaries, which may also imply
2745 across enterprise governance boundaries. For such situations, governance emphasizes the need for
2746 agreement that some governance framework and governance processes have jurisdiction, and the
2747 governance defined must satisfy the goals of the participants for cooperation to continue. A standards
2748 development organization such as OASIS is an example of voluntary agreement to governance over a
2749 limited domain to satisfy common goals.

2750 The specifics discussed in the figures in the previous sections are equally applicable to governance
2751 across ownership boundaries as it is within a single boundary. There is a charter agreed to when
2752 participants become members of the organization, and this charter sets up the structures and processes
2753 to be followed. Leadership may be shared by the leadership of the overall organization and the leadership
2754 of individual groups themselves chartered per the governance processes. There are rules and regulations
2755 specific to individual efforts for which participants agree to local goals, and enforcement can be loss of
2756 voting rights or under extreme circumstances, expulsion from the group.

2757 Thus, the major difference for SOA governance is an appreciation for the cooperative nature of the
2758 enterprise and its reliance on furthering common goals if productive participation is to continue.

2759 **5.1.3.2 What Must be Governed**

2760 An expected benefit of employing SOA principles is the ability to quickly bring resources to bear to deal
2761 with unexpected and evolving situations. This requires a great deal of confidence in the underlying
2762 capabilities that can be accessed and in the services that enable the access. It also requires considerable
2763 flexibility in the ways these resources can be employed. Thus, SOA governance requires establishing
2764 confidence and trust (see Section 3.2.5.1) while instituting a solid framework that enables flexibility,
2765 indicating a combination of strict control over a limited set of foundational aspects but minimum
2766 constraints beyond those bounds.

2767

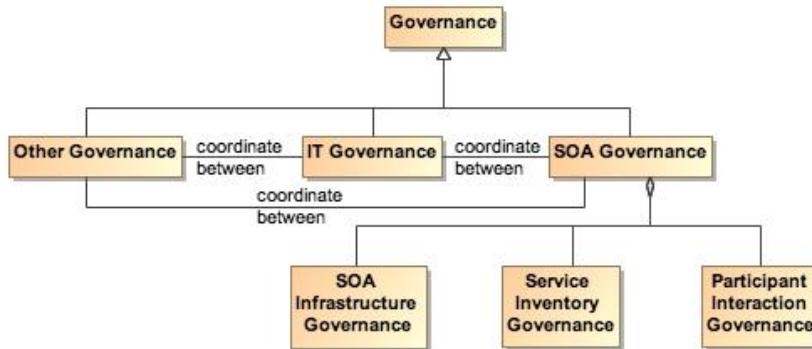


Figure 40 - Relationship Among Types of Governance

2768
2769
2770
2771
2772
2773
2774
2775
2776

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

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

2777 **5.1.3.2.1 Governance of SOA Infrastructure**

2778 The SOA infrastructure is likely composed of several families of SOA services that provide access to
2779 fundamental computing business services. These include, among many others, services such as
2780 messaging, security, storage, discovery, and mediation. The provisioning of an infrastructure on which
2781 these services may be accessed and the general realm of those contributing as utility functions of the
2782 infrastructure are a traditional IT governance concern. In contrast, the focus of SOA governance is how
2783 the existence and use of the services enables the SOA ecosystem.

2784 By characterizing the environment as containing families of SOA services, the assumption is that there
2785 may be multiple approaches to providing the business services or variations in the actual business
2786 services provided. For example, discovery could be based on text search, on metadata search, on
2787 approximate matches when exact matches are not available, and numerous other variations. The
2788 underlying implementation of search algorithms are not the purview of SOA governance, but the access
2789 to the resulting service infrastructure enabling discovery must be stable, reliable, and extremely robust to
2790 all operating conditions. Such access enables other specialized SOA services to use the infrastructure in
2791 dependable and predictable ways, and is where governance is important.

2792 **5.1.3.2.2 Governance of the Service Inventory**

2793 Given an infrastructure in which other SOA services can operate, a key governance issue is which SOA
2794 services to allow in the ecosystem. The major concern should be a definition of well-behaved services,
2795 where the required behavior will inherit their characteristics from experiences with distributed computing
2796 but also evolve with SOA experience. A major requirement for ensuring well-behaved services is
2797 collecting sufficient metrics to know how the service affects the SOA infrastructure and whether it
2798 complies with established infrastructure policies.

2799 Another common concern of service approval is whether there is a possibility of duplication of function by
2800 multiple services. Some governance models talk to a tightly controlled environment where a primary
2801 concern is to avoid any service duplication. Other governance models talk to a market of services where
2802 the consumers have wide choices. For the latter, it is anticipated that the better services will emerge from
2803 market consensus and the availability of alternatives will drive innovation.

2804 Some combination of control and openness will emerge, possibly with a different appropriate balance for
2805 different categories of use. For SOA governance, the issue is less which services are approved but rather
2806 ensuring that sufficient description is available to support informed decisions for appropriate use. Thus,
2807 SOA governance should concentrate on identifying the required attributes to adequately describe a
2808 service, the required target values of the attributes, and the standards for defining the meaning of the
2809 attributes and their target values. Governance may also specify the processes by which the attribute
2810 values are measured and the corresponding certification that some realized attribute set may imply.

2811 For example, unlimited access for using a service may require a degree of life cycle maturity that has
2812 demonstrated sufficient testing over a certain size community. Alternately, the policy may specify that a
2813 service in an earlier phase of its life cycle may be made available to a smaller, more technically
2814 sophisticated group in order to collect the metrics that would eventually allow the service to advance its
2815 life cycle status.

2816 This aspect of governance is tightly connected to description because, given a well-behaved set of
2817 services, it is the responsibility of the consumer (or policies promulgated by the consumer's organization)
2818 to decide whether a service is sufficient for that consumer's intended use. The goal is to avoid global
2819 governance specifying criteria that are too restrictive or too lax for the local needs of which global
2820 governance has little insight.

2821 Such an approach to specifying governance allows independent domains to describe services in local
2822 terms while still having the services available for informed use across domains. In addition, changes to
2823 the attribute sets within a domain can be similarly described, thus supporting the use of newly described
2824 resources with the existing ones without having to update the description of the entire legacy content.

2825 **5.1.3.2.3 Governance of Participant Interaction**

2826 Finally, given a reliable services infrastructure and a predictable set of services, the third aspect of
2827 governance is prescribing what is required during a service interaction.

2828 Governance would specify adherence to service interface and service reachability parameters and would
2829 require that the result of an interaction correspond to the real world effects as contained in the service
2830 description. Governance would ensure preconditions for service use are satisfied, in particular those
2831 related to security aspects such as user authentication, authorization, and non-repudiation. If conflicts
2832 arise, governance would specify resolution processes to ensure appropriate agreements, policies, and
2833 conditions are met.

2834 It would also rely on sufficient monitoring by the SOA infrastructure to ensure services remain well-
2835 behaved during interactions, e.g. do not use excessive resources or exhibit other prohibited behavior.
2836 Governance would also require that policy agreements as documented in the execution context for the
2837 interaction are observed and that the results and any after effects are consistent with the agreed policies.
2838 Here, governance focuses more on contractual and legal aspects rather than the precursor descriptive
2839 aspects. SOA governance may prescribe the processes by which SOA-specific policies are allowed to
2840 change, but there are probably more business-specific policies that will be governed by processes
2841 outside SOA governance.

2842 **5.1.3.3 Overarching Governance Concerns**

2843 There are numerous governance related concerns whose effects span the three areas just discussed.
2844 One is the area of standards, how these are mandated, and how the mandates may change. The Web
2845 Services standards stack is an example of relevant standards where a significant number are still under
2846 development. In addition, while there are notional scenarios that guide what standards are being
2847 developed, the fact that many of these standards do not yet exist precludes operational testing of their
2848 adequacy or effectiveness as a necessary and sufficient set.

2849 That said, standards are critical to creating a SOA ecosystem where SOA services can be introduced,
2850 used singularly, and combined with other services to deliver complex business functionality. As with other
2851 aspects of SOA governance, the governance body should identify the minimum set felt to be needed and
2852 rigorously enforce that that set be used where appropriate. The governance body takes care to expand
2853 and evolve the mandated standards in a predictable manner and with sufficient technical guidance that
2854 new services are able to coexist as much as possible with the old, and changes to standards do not
2855 cause major disruptions.

2856 Another area that may see increasing activity as SOA expands is additional regulation by governments
2857 and associated legal institutions. New laws may deal with transactions that are service based, possibly
2858 including taxes on the transactions. Disclosure laws may mandate certain elements of description so both
2859 the consumer and provider act in a predictable environment and are protected from ambiguity in intent or
2860 action. Such laws spawn rules and regulations that will influence the metrics collected for evaluation of
2861 compliance.

2862 **5.1.3.4 Considerations for SOA Governance**

2863 The Reference Architecture definition of a loosely coupled system is one in which the constraints on the
2864 interactions between components are minimal: sufficient to permit interoperation without additional
2865 constraints that may be an artifact of implementation technology. While governance experience for
2866 standalone systems provides useful guides, we must be careful not to apply constraints that would
2867 preclude the flexibility, agility, and adaptability we expect to realize from a SOA ecosystem.

2868 One of the strengths of the SOA paradigm is it can make effective use of diversity rather than requiring
2869 monolithic solutions. Heterogeneous organizations can interact without requiring each conforms to
2870 uniform tools, representation, and processes. However, with this diversity comes the need to adequately
2871 define those elements necessary for consistent interaction among systems and participants, such as
2872 which communication protocol, what level of security, which vocabulary for payload content of messages.
2873 The solution is not always to lock down these choices but to standardize alternatives and standardize the
2874 representations through which an unambiguous identification of the alternative chosen can be conveyed.
2875 For example, the URI standard specifies the URI string, including what protocol is being used, what is the
2876 target of the message, and how parameters may be attached. It does not limit the available protocols, the
2877 semantics of the target address, or the parameters that can be transferred. Thus, as with our definition of
2878 loose coupling, it provides absolute constraints but minimizes which constraints it imposes.

2879 There is not a one-size-fits-all governance but a need to understand the types of things governance is
2880 called upon to do in the context of the goals of the SOA paradigm. Some communities may initially desire
2881 and require very stringent governance policies and procedures while others see need for very little. Over
2882 time, best practices will evolve, resulting in some consensus on a sensible minimum and, except in
2883 extreme cases where it is demonstrated to be necessary, a loosening of strict governance toward the
2884 best practice mean.

2885 A question of how much governance may center on how much time governance activities require versus
2886 how quickly is the system being governed expected to respond to changing conditions. For large single
2887 systems that take years to develop, the governance process could move slowly without having a serious
2888 negative impact. For example, if something takes two years to develop and the steps involved in
2889 governance take two months to navigate, then the governance can go along in parallel and may not have
2890 a significant impact on system response to changes. Situations where it takes as long to navigate
2891 governance requirements as it does to develop a response are examples where governance may need to
2892 be reevaluated as to whether it facilitates or inhibits the desired results. Thus, the speed at which services
2893 are expected to appear and evolve needs to be considered when deciding the processes for control. The
2894 added weight of governance should be appropriate for overall goals of the application domain and the
2895 service environment.

2896 Governance, as with other aspects of any SOA implementation, should start small and be conceptualized
2897 in a way that keeps it flexible, scalable, and realistic. A set of useful guidelines would include:

- 2898 • Do not hardwire things that will inevitably change. For example, develop a system that uses the
2899 representation of policies rather than code the policies into the implementations.
- 2900 • Avoid setting up processes that demo well for three services without considering how they may
2901 work for 300. Similarly, consider whether the display of status and activity for a small number of
2902 services will also be effective for an operator in a crisis situation looking at dozens of services,
2903 each with numerous, sometimes overlapping and sometimes differing activities.
- 2904 • Maintain consistency and realism. A service solution responding to a natural disaster cannot be
2905 expected to complete a 6-week review cycle but be effective in a matter of hours.

2906 5.1.4 Architectural Implications of SOA Governance

2907 The description of SOA governance indicates numerous architectural requirements on the SOA
2908 ecosystem:

- 2909 • Governance is expressed through policies and assumes multiple use of focused policy modules
2910 that can be employed across many common circumstances. The **following are thus REQUIRED:**
 - 2911 ○ descriptions to enable the policy modules to be visible, where the description **SHOULD**
2912 includes a unique identifier for the policy ~~and as well as~~ a sufficient, and preferably a
2913 machine process-able, representation of the meaning of terms used to describe the
2914 policy, its functions, and its effects;
 - 2915 ○ one or more discovery mechanisms that enable searching for policies that best meet the
2916 search criteria specified by ~~the service-a~~ participant; where the discovery mechanism will
2917 have access to the individual policy descriptions, possibly through some repository
2918 mechanism;
 - 2919 ○ accessible storage of policies and policy descriptions, so ~~service~~-participants can access,
2920 examine, and use the policies as defined.
- 2921 • Governance requires that the participants understand the intent of governance, the structures
2922 created to define and implement governance, and the processes to be followed to make
2923 governance operational. This **REQUIRES:**
 - 2924 ○ an information collection site, such as a Web page or portal, where governance
2925 information is stored and from which the information is always available for access;
 - 2926 ○ a mechanism to inform participants of significant governance events, such as changes in
2927 policies, rules, or regulations;
 - 2928 ○ accessible storage of the specifics of Governance Processes;
 - 2929 ○ SOA services to access automated implementations of the Governance Processes
- 2930 • Governance policies are made operational through rules and regulations. This **REQUIRES:**
 - 2931 ○ descriptions to enable the rules and regulations to be visible, where the description
2932 **SHOULD** includes a unique identifier and a sufficient, and preferably a machine process-
2933 able, representation of the meaning of terms used to describe the rules and regulations;
 - 2934 ○ one or more discovery mechanisms that enable searching for rules and regulations that
2935 may apply to situations corresponding to the search criteria specified by ~~the service-a~~
2936 participant; where the discovery mechanism will have access to the individual
2937 descriptions of rules and regulations, possibly through some repository mechanism;
 - 2938 ○ accessible storage of rules and regulations and their respective descriptions, so ~~service~~
2939 participants can understand and prepare for compliance, as defined.
 - 2940 ○ SOA services to access automated implementations of the Governance Processes.
- 2941 • Governance implies management to define and enforce rules and regulations. Management
2942 is discussed more specifically in section 5.3, but in a parallel to governance, management
2943 **REQUIRES:**
 - 2944 ○ an information collection site, such as a Web page or portal, where management
2945 information is stored and from which the information is always available for access;
 - 2946 ○ a mechanism to inform participants of significant management events, such as changes
2947 in rules or regulations;
 - 2948 ○ accessible storage of the specifics of processes followed by management.
- 2949 • Governance relies on metrics to define and measure compliance. This **REQUIRES:**
 - 2950 ○ the infrastructure monitoring and reporting information on SOA resources;
 - 2951 ○ possible interface requirements to make accessible metrics information generated or
2952 most easily accessed by the service itself.

2953 5.2 Security Model

2954 Security is one aspect of confidence – the confidence in the integrity, reliability, and confidentiality of the
2955 system. In particular, security focuses on those aspects of assurance that involve the accidental or malign
2956 intent of other people to damage or compromise trust in the system and on the availability of SOA-based
2957 systems to perform desired capability.

2958 Security

2959 The set of mechanisms for ensuring and enhancing **trust** and confidence in the **SOA ecosystem**.

2960 Providing for security for Service Oriented Architecture is somewhat different than for other contexts;
2961 although many of the same principles apply equally to SOA and to other systems. The fact that SOA
2962 embraces crossing ownership boundaries makes the issues involved with moving data more visible.

2963 As well as securing the movement of data within and across ownership boundaries, security often
2964 revolves around resources: the need to guard certain resources against inappropriate access – whether
2965 reading, writing or otherwise manipulating those resources.

2966 Any comprehensive security solution must take into account the people that are using, maintaining and
2967 managing SOA-based systems. Furthermore, the relationships between them must also be incorporated:
2968 any security assertions that may be associated with particular interactions originate in the people that are
2969 behind the interaction.

2970 We analyze security in terms of the social structures that define the legitimate permissions, obligations
2971 and roles of people in relation to the system, and mechanisms that must be put into place to realize a
2972 secure system. The former are typically captured in a series of security policy statements; the latter in
2973 terms of security guards that ensure that policies are enforced.

2974 How and when to apply these derived security policy mechanisms is directly associated with the
2975 assessment of the *threat model* and a *security response model*. The threat model identifies the kinds of
2976 threats that directly impact the message and/or application of constraints and the response model is the
2977 proposed mitigation to those threats. Properly implemented, the result can be an acceptable level of risk
2978 to the safety and integrity within the SOA ecosystem.

2979 5.2.1 Secure Interaction Concepts

2980 We can characterize secure interactions in terms of key security concepts **[ISO/IEC 27002]**:
2981 confidentiality, integrity, authentication, authorization, non-repudiation, and availability. The concepts for
2982 secure interactions are well defined in other standards and publications. The security concepts here are
2983 not defined but rather related to the SOA ecosystem perspective of the SOA-RAF.

2984 5.2.1.1 Confidentiality

2985 Confidentiality is concerned with the protection of privacy of participants in their interactions.
2986 Confidentiality refers to the assurance that unauthorized entities are not able to read messages or parts
2987 of messages that are transmitted.

2988 Note that confidentiality has degrees: in a completely confidential exchange, third parties would not even
2989 be aware that a confidential exchange has occurred. In a partially confidential exchange, the identities of
2990 the participants may be known but the content of the exchange obscured.

2991 5.2.1.2 Integrity

2992 Integrity is concerned with the protection of information that is exchanged – either from unauthorized
2993 writing or inadvertent corruption. Integrity refers to the assurance that information that has been
2994 exchanged has not been altered.

2995 Integrity is different from confidentiality in that messages that are sent from one participant to another
2996 may be obscured to a third party, but the third party may still be able to introduce his own content into the
2997 exchange without the knowledge of the participants.

2998 Section 5.2.4 describes common computing techniques for providing confidentiality and integrity during
2999 message exchanges.

3000 5.2.1.3 Authentication

3001 Authentication is concerned with addressing a need to adequately identify actors in a potential interaction
3002 or joint action.

3003 A combination of identifiers (as discussed in section 3.2.4.1) and other attributes of an actor is typically
 3004 used to achieve this.
 3005 The set of attribute values that claim to identify a specific actor are matched against the set of reference
 3006 values expected for that actor and that are maintained by some trusted authority.
 3007 If the comparison results in a sufficient match, authentication has been achieved.
 3008 Which specific set of attributes is considered an adequate basis for comparison will be context-dependent
 3009 and specifying such sets is not within the scope of the SOA-RAF.
 3010 Authentication merely provides an assertion that an actor is the person or agent that it claims to be. Of
 3011 itself, it does not provide a 'green light' to proceed with the interaction – this is rather the concern of
 3012 authorization, covered below.

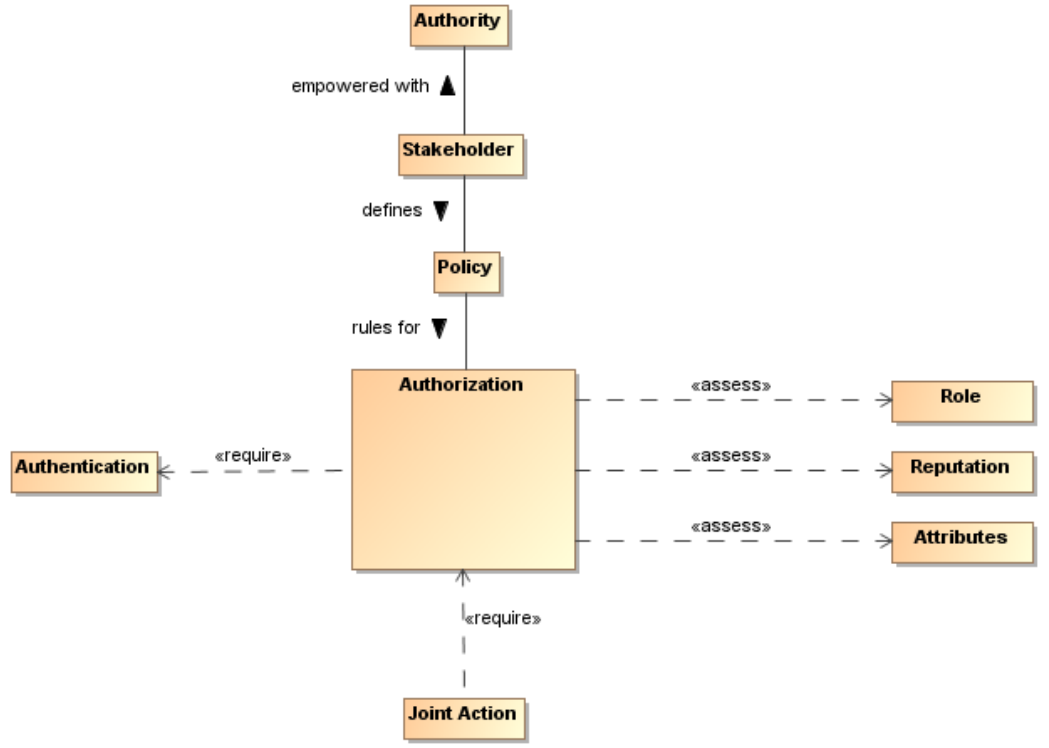
Comment [PFB101]: Issue 40, part
 New text reflects deletion of
 'identity' early in section 3.

Comment [KJL102]: Issue 40, part
 (Figure deleted)

3013 **5.2.1.4 Authorization**

3014 Authorization concerns the legitimacy of the interaction. Authorization refers to the means by which a
 3015 stakeholder may be assured that the information and actions that are exchanged are either explicitly or
 3016 implicitly approved.

Comment [PFB103]: Issue 128



3017 *Figure 41 - Authorization*

3018
 3019 The roles and attributes which provide a participant's credentials are expanded to include reputation.
 3020 Reputation often helps determine willingness to interact, for example, reviews of a service provider will
 3021 influence the decision to interact with the service provider. The roles, reputation, and attributes are
 3022 represented as assertions measured by authorization decision points.

3023 The role of policy for security is to permit stakeholders to express their choices. In Figure 41, a policy is a
 3024 written constraint and the role, reputation, and attribute assertions are evaluated according to the
 3025 constraints in the authorization policy. A combination of security mechanisms and their control via explicit
 3026 policies can form the basis of an authorization solution.

3027 **5.2.1.5 Non-repudiation**

3028 Non-repudiation concerns the accountability of participants. To foster trust in the performance of a system
3029 used to conduct shared activities, it is important that the participants are not able to later deny their
3030 actions: to repudiate them. Non-repudiation refers to the means by which a participant may not, at a later
3031 time, successfully deny having participated in the interaction or having performed the actions as reported
3032 by other participants.

3033 **5.2.1.6 Availability**

3034 Availability concerns the ability of systems to use and offer the services for which they were designed.
3035 One of the threats against availability is the so-called denial of service attack in which attackers attempt to
3036 prevent legitimate access to the system.
3037 We differentiate here between general availability – which includes aspects such as systems reliability –
3038 and availability as a security concept where we need to respond to active threats to the system.

3039 **5.2.2 Where SOA Security is Different**

3040 The core security concepts are fundamental to all social interactions. The evolution of sharing information
3041 within a SOA ecosystem requires the flexibility to dynamically secure computing interactions where the
3042 owning social groups, roles, and authority are constantly changing as described in section 5.1.3.1.

3043 SOA policy-based security can be more adaptive than previous computing technologies, and typically
3044 involves a greater degree of distributed mechanisms.

3045 Standards for security, as is the case with all aspects of SOA implementation and use, play a large role in
3046 flexible security on a global scale. SOA security may also involve greater auditing and reporting to adhere
3047 to regulatory compliance established by governance structures.

3048 **5.2.3 Security Threats**

3049 There are a number of ways in which an attacker may attempt to compromise the security of a system.
3050 The two primary sources of attack are third parties attempting to subvert interactions between legitimate
3051 participants and an entity that is participating but attempting to subvert other participants. The latter is
3052 particularly important in a SOA ecosystem where there may be multiple ownership boundaries and trust
3053 boundaries.

3054 The threat model lists some common threats that relate to the core security concepts listed in Section
3055 5.2.1. Each technology choice in the realization of a SOA-based system can potentially have many
3056 threats to consider.

3057 **Message alteration**

3058 If an attacker is able to modify the content (or even the order) of messages that are exchanged without
3059 the legitimate participants being aware of it then the attacker has successfully compromised the security
3060 of the system. In effect, the participants may unwittingly serve the needs of the attacker rather than their
3061 own.

3062 An attacker may not need to completely replace a message with his own to achieve his objective:
3063 replacing the identity of the initially intended recipient of a transaction may be enough.

Comment [PFB104]: Issue 136

3064 **Message interception**

3065 If an attacker is able to intercept and understand messages exchanged between participants, then the
3066 attacker may be able to gain advantage. This is probably the most commonly understood security threat.

3067 **Man in the middle**

3068 In a man-in-the-middle attack, the legitimate participants believe that they are interacting with each other;
3069 but are in fact interacting with the attacker. The attacker attempts to convince each participant that he is
3070 their correspondent; whereas in fact he is not.

3071 In a successful man-in-the-middle attack, legitimate participants do not have an accurate understanding
3072 of the state of the other participants. The attacker can use this to subvert the intentions of the participants.

3073 **Spoofing**

3074 In a spoofing attack, the attacker convinces a participant that he is really someone else – someone that
3075 the participant would normally trust.

3076 **Denial of service attack**

3077 In a denial of service (DoS) attack, the attacker attempts to prevent legitimate users from making use of
3078 the service. A DoS attack is easy to mount and can cause considerable harm: by preventing legitimate
3079 interactions, or by slowing them down enough, the attacker may be able to simultaneously prevent
3080 legitimate access to a service and to attack the service by another means.

3081 A variation of the DoS attack is the Distributed Denial of Service (DDoS) attack. In a DDoS attack the
3082 attacker uses multiple agents to attack the target. In some circumstances this can be extremely
3083 difficult to counteract effectively.

3084 One of the features of a DoS attack is that it does not require valid interactions to be effective: responding
3085 to invalid messages also takes resources and that may be sufficient to cripple the target.

3086 **Replay attack**

3087 In a replay attack, the attacker captures the message traffic during a legitimate interaction and then
3088 replays part of it to the target. The target is persuaded that a similar transaction to the previous one is
3089 being repeated and it responds as though it were a legitimate interaction.

3090 A replay attack may not require that the attacker understand any ~~of the~~ individual
3091 ~~message communications~~; the attacker may have different objectives (for example attempting to predict
3092 how the target would react to a particular request).

3093 **False repudiation**

3094 In false repudiation, a user completes a normal transaction and then later attempts to deny that the
3095 transaction occurred. For example, a customer may use a service to buy a book using a credit card; then,
3096 when the book is delivered, refuse to pay the credit card bill claiming that someone else must have
3097 ordered the book.

3098 **5.2.4 Security Responses**

3099 Security goals are never absolute: it is not possible to guarantee 100% confidentiality, non-repudiation,
3100 etc. However, a well designed and implemented security response model can ensure acceptable levels of
3101 security risk. For example, using a well-designed cipher to encrypt messages may make the cost of
3102 breaking communications so great and so lengthy that the information obtained is valueless.

3103 Performing threat assessments, devising mitigation strategies, and determining acceptable levels of risk
3104 are the foundation for an effective process to mitigating threats in a cost-effective way.¹⁰ The choice in
3105 hardware and software to realize a SOA implementation will be a basis for threat assessments and
3106 mitigation strategies. The stakeholders of a specific SOA implementation should determine acceptable
3107 levels of risk based on threat assessments and the cost of mitigating those threats.

Comment [PFB105]: Issue 141

3108 **5.2.4.1 Privacy Enforcement**

3109 The most efficient mechanism to assure confidentiality is the encryption of information. Encryption is
3110 particularly important when messages must cross trust boundaries; especially over the Internet. Note that
3111 encryption need not be limited to the content of messages: it is possible to obscure even the existence of
3112 messages themselves through encryption and 'white noise' generation in the communications channel.

3113 The specifics of encryption are beyond the scope of this architecture. However, we are concerned about
3114 how the connection between privacy-related policies and their enforcement is made.

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.

3115 A policy enforcement point for enforcing privacy may take the form of an automatic function to encrypt
3116 messages as they leave a trust boundary; or perhaps simply ensuring that such messages are suitably
3117 encrypted.

3118 Any policies relating to the level of encryption being used would then apply to these centralized
3119 messaging functions.

3120 **5.2.4.2 Integrity Protection**

3121 To protect against message tampering or inadvertent message alteration, and to allow the receiver of a
3122 message to authenticate the sender, messages may be accompanied by a digital signature. Digital
3123 signatures provide a means to detect if signed data has been altered. This protection can also extend to
3124 authentication and non-repudiation of a sender.

3125 A common way a digital signature is generated is with the use of a private key that is associated with a
3126 public key and a digital certificate. The private key of some entity in the system is used to create a digital
3127 signature for some set of data. Other entities in the system can check the integrity of the signed data set
3128 via signature verification algorithms. Any changes to the data that was signed will cause signature
3129 verification to fail, which indicates that integrity of the data set has been compromised.

3130 A party verifying a digital signature must have access to the public key that corresponds to the private key
3131 used to generate the signature. A digital certificate contains the public key of the owner, and is itself
3132 protected by a digital signature created using the private key of the issuing Certificate Authority (CA).

3133 **5.2.4.3 Message Replay Protection**

3134 To protect against replay attacks, messages may contain information that can be used to detect replayed
3135 messages. The simplest requirement to prevent replay attacks is that each message that is ever sent is
3136 unique. For example, a message may contain a message ID, a timestamp, and the intended destination.

3137 By storing message IDs, and comparing each new message with the store, it becomes possible to verify
3138 whether a given message has been received before (and therefore should be discarded).

3139 The timestamp may be included in the message to help check for message freshness. Messages that
3140 arrive after their message ID could have been cleared (after receiving the same message some time
3141 previously) may also have been replayed. A common means for representing timestamps is a useful part
3142 of an interoperable replay detection mechanism.

3143 The destination information is used to determine if the message was misdirected or replayed. If the
3144 replayed message is sent to a different endpoint than the destination of the original message, the replay
3145 could go undetected if the message does not contain information about the intended destination.

3146 In the case of messages that are replies to prior messages, it is also possible to include seed information
3147 in the prior messages that is randomly and uniquely generated for each message that is sent out. A
3148 replay attack can then be detected if the reply does not embed the random number that corresponds to
3149 the original message.

3150 **5.2.4.4 Auditing and Logging**

3151 False repudiation involves a participant denying that it authorized a previous interaction. An effective
3152 strategy for responding to such a denial is to maintain careful and complete logs of interactions that can
3153 be used for auditing purposes. The more detailed and comprehensive an audit trail is, the less likely it is
3154 that a false repudiation would be successful.

3155 The countermeasures assume that the non-repudiation tactic (e.g. digital signatures) is not undermined
3156 itself. For example, if private key is stolen and used by an adversary, even extensive logging cannot
3157 assist in rejecting a false repudiation.

3158 Unlike many of the security responses discussed here, it is likely that the scope for automation in
3159 rejecting a repudiation attempt is limited [in the immediate future](#) to careful logging.

Comment [KL106]: Issue 275

3160 5.2.4.5 Graduated engagement

3161 The key to managing and responding to DoS attacks is to be careful in the use of resources when
3162 responding to interaction. Put simply, a system has a choice to respond to a communication or to ignore
3163 it. In order to avoid vulnerability to DoS attacks a service provider should be careful not to commit
3164 resources beyond those implied by the current state of interactions; this permits a graduation in
3165 commitment by the service provider that mirrors any commitment on the part of service consumers and
3166 attackers alike.

3167 5.2.5 Identity and Access Control

3168 5.2.5.1 Identity Propagation

3169 As service-oriented solutions have risen in popularity, standards bodies have developed various
3170 specifications used to convey identity. In solutions utilizing these standards, an assertion of a subject's
3171 identity is issued by an asserting party, and is sent to a relying party. The recipient's assurance of the
3172 subject's identity is based on its trust of the asserting party. When web service orchestration solutions
3173 combine multiple distributed components and services throughout a network, each component of the
3174 solution may need to understand the identity of the user that originated a request in order to provide
3175 proper access control to data.

3176 The following sections cover two approaches for identity propagation, revolving around the choice of the
3177 asserting party. The first approach, where the message sender (or service consumer) is the asserting
3178 party, is a "sender vouches" approach. The second approach, where a token service is trusted to issue
3179 assertions about subjects, is a "token service"-based approach. Each has implications.

3180 5.2.5.1.1 Sender-Vouches Approaches

3181 In a "sender vouches" approach, the message sender vouches for the identity of an end-user. Here, an
3182 assertion of a user's identity is created and sent (or propagated) from the service consumer to a service
3183 provider. A challenge with this approach is that the recipient service may then propagate the claim to
3184 further service participants. As the number of service participants grows and the relative "distance"
3185 between the originator of the request and the final recipient service provider increases in a service
3186 transaction, it can become increasingly more difficult to positively prove the identity of every actor in the
3187 service chain.

3188 As (K. Smith, "Mitigating Risks Associated with Transitive Trust in Service Based Identity Propagation",
3189 Information Security Journal: A Global Perspective, 21:2, 71-78, April 2012) discusses, some of the
3190 challenges that such approaches bring involve:

- 3191 1. Trust of Message Senders. In such end-to-end scenarios, the trust of the claim of the identity of
3192 the user is always based on the trust of the message sender(s) passing the claim.
- 3193 2. Risk of Vulnerabilities in Intermediaries. Because the called services are basing the assumption
3194 of the identity of the propagated end-user on the claim passed to the service by the message
3195 sender (which the service does not control), a risk is that intermediary services within the
3196 transaction may become compromised and may inaccurately send false identity claims.
3197 Depending on the exact messaging syntax, an intermediary service could potentially manipulate
3198 the claim about the originating user or substitute another claim about another user not intended
3199 for use in the transaction. There could also be impersonation of the intermediary services,
3200 affecting the reliability of the transaction.
- 3201 3. Degrading Trust. Because the trust of the claim of the identity of the user is based on the trust of
3202 the message senders, the more intermediaries there are, trust degrades as the distance between
3203 the end-user and the service being called becomes greater. Trust of the identity of the originating
3204 user is therefore dependent on the trust of every sender in the chain to properly pass the claim.
- 3205 4. Loss of Context. Could an assertion of identity that was created for one purpose be maliciously
3206 used or unintentionally used for another purpose? Could someone, for example, use a signed
3207 claim about a user's identity to empty the user's bank account, when the issuer of the claim only
3208 intended it to be used for another purpose?

3209 Approaches for mitigating risks in sender-vouches approaches involve a careful combination of SOA
3210 Security governance, limiting the amount of message senders trusted to propagate identity, requiring
3211 services to keep explicit trust lists of consumers trusted to propagate identity, limiting the re-use of
3212 assertions beyond a certain number of points, establishing conditions of use for propagated assertions,
3213 keeping track of the history of the assertion in the transaction, and use of digital signatures by the original
3214 asserting party. It should be mentioned that the risks identified in this section can be reduced, but not
3215 completely mitigated.

3216 **5.2.5.1.2 Token Service-based Approaches**

3217 This approach revolves around the asserting party being a token service trusted to vouch for a user's
3218 identity. Typically, this approach involves the combination of an Identity Provider (IdP) and Security Token
3219 Service (STS), where an application redirects user authentication to the IdP/STS, and where the token
3220 service authenticates the user and returns a token (typically signed) to the application.

3221 By using this approach, the application, or service consumer, propagates the token to a service provider.
3222 The service provider's trust of the token is therefore based on its trust of the token service (and not the
3223 service consumer calling the service). Because trust of the token revolves around the service trusted to
3224 issue tokens, such an approach reduces many of the risks detailed in the last section. Beyond the initial
3225 consumer-service request, the recipient service may then propagate the claim to further service
3226 participants.

3227 There are risks associated with the re-use of tokens beyond the initial consumer-provider interaction.
3228 Most token service protocols and specifications, therefore, provide the capability for "refreshing" tokens
3229 for reuse in such situations. In this case, each service retrieving a token may request that the token
3230 service issue a "refresh token" that can be propagated for a subsequent service call. By utilizing refresh
3231 tokens, it removes the risks associated with reuse.

3232 Unlike sender-vouches approaches for identity propagation, such an approach revolves around a
3233 centralized token service. There may, there are be architectural implications related to performance and
3234 availability. If all services in the enterprise need interact with the token service, then the token service
3235 must always be available, and there will most likely be latency related to interacting with a token service,
3236 resulting in some performance overhead. It is therefore advised that solutions that provide elastic
3237 scalability be used to ensure that token services are readily available to respond to requests.

3238 **5.2.5.2 Access Control Approaches**

3239 Access control revolves around security policy. If access control policy can be discovered and processed,
3240 and if authorization credentials of users can be retrieved, access control can be successfully enforced.
3241 Architectural flexibility for authorization in a Service-Oriented Architecture is achieved by logically
3242 separating duties into Policy Decision Points (PDPs) and Policy Enforcement Points (PEPs). A PDP is the
3243 point at which access control decisions are made, based on an expressed access control policy and a
3244 subject's authorization credentials. The enforcement of the decision is delegated to a PEP. Some
3245 standards, such as XACML (the eXtensible Access Control Markup Language), decompose the policy
3246 model further into Policy Administration Points (PAPs) that create policy and the Policy Information Points
3247 (PIPs) that query attributes for subjects requesting access to resources. There are many strategies of
3248 how PDPs and PEPs can work together, each with architectural implications that have an impact on
3249 security, performance and scalability, using any and all of the approaches discussed in this paper. For
3250 this reason, we have found the following to be the best approach:

3251 The use of certain types of enterprise security services may dictate the architectural approaches used, so
3252 implementations in certain environments may need to understand implications. This section provides an
3253 overview of such approaches.

3254 **5.2.5.2.1 Centralized Access Control Approaches**

3255 The first approach to policy enforcement is a purely centralized approach. A centralized policy server acts
3256 as a PDP, and is called to make a yes or no decision for a subject request for resource. In this model, all
3257 policy is globally managed, and services request authorization decisions from the policy server. The
3258 policy server may retrieve security attributes about subjects requesting resources, and makes decisions

3259 based on a subject's credentials and enterprise policy for the requested resource. In this model, the policy
3260 server PDP returns a "yes" or "no" authorization decision and enforces the decision.

3261 A positive aspect of this approach is information hiding; all services that request authorization decisions
3262 do not know exactly why decisions are made, and at no point in the lifecycle of the message are user
3263 authorization credentials revealed. This can be a benefit from an information security standpoint. Since
3264 applications and services may not have a need to know about authorization credentials of the user or the
3265 specific policy of a resource, this model protects that information in cases where this information may be
3266 sensitive or confidential. The second positive aspect of this model involves centralized access to the
3267 latest policy. If enterprise policies are employed, an enterprise is always assured of using the most up-to-
3268 date version of that policy.

3269 There can be, however, negatives about this model regarding performance and availability. If all services
3270 in the enterprise need to connect to a central authorization server for every request, then the server must
3271 always be available. If the authorization server ever goes down, you have two grim choices: allow all
3272 access for all subjects (turning off security), or deny all access to all subjects (denial of service). It is
3273 unlikely that either of these is an acceptable course of action. Complicating the issue is that calls to such
3274 a policy service are usually cryptographically protected in order to have high assurance of the integrity
3275 and identity of the policy server. That cryptography, combined with network latency of each request,
3276 slows down the response time of all services and applications that are forced to call the policy server.
3277 Centralization can be a "virtual" concept, meaning that various mechanisms for load-balancing and
3278 failover for many authorization servers can be used when using the centralized policy management and
3279 enforcement model. Because of the performance, availability, and scalability concerns, any centralized
3280 solution should be coupled with alternative approaches for greater flexibility.

3281 **5.2.5.2 Decentralized Access Control Approaches**

3282 In a decentralized approach, the authenticating application retrieves or creates an assertion about the
3283 identity and/or attributes of the end-user and propagates that assertion to a web service, which has a
3284 "local" PDP and PEP. The web service PDP refers to locally expressed policy, and therefore, its PDP can
3285 inspect the policy, combined with the assertion information propagated in, in order to make an access
3286 control decision. If only identity is propagated in to the web service, the web service may retrieve
3287 authorization credentials from an Attribute Service lookup based on the identity. If authorization
3288 credentials are passed in the assertion, the web service PDP may provide access control without calling
3289 any external service, based on locally expressed policy and the authorization credential assertion.

3290 The decentralized model alleviates the performance concerns of the purely central model. Because the
3291 policy is locally expressed, the web service does not have to cryptographically call a PDP over the
3292 network. Instead, all policy is local, and all policy decisions are local, based on global attribute
3293 credentials. Since all policy is declared locally, there is no longer concern of the availability of a central
3294 policy server.

3295 There are two potential concerns with this model. One concern is that there is no information hiding. If an
3296 assertion about the user is propagated into the web service, the web service will have full read access to
3297 the security credentials of the user, which may be sensitive. The second concern revolves around policy
3298 management. In situations where an organization may want to have full control over policy, this purely
3299 decentralized model does not allow it.

3300 **5.2.5.2.3 Hybrid Access Control Approaches**

3301 Because of the significant weaknesses related to performance, availability, and scalability in purely
3302 centralized approaches, and because of the desire to have centralized control of access control policy,
3303 hybrid approaches have emerged in access control. There are two common methods for providing a
3304 "happy medium" between local control of policy (where services express all policy) and central control of
3305 policy (where a central policy server expresses all policy). In these models, each web service expresses
3306 local policy, but leverages global organizational policy in order to make decisions. This alleviates policy
3307 control concerns related to the purely decentralized model.

3308 In these approaches, a web service's PDP must get access to global policy, and combines it with its
3309 locally expressed policy in order to make a decision. These hybrid solutions also require a conflict-
3310 resolution mechanism to be in place (for example, where global policy trumps local policy). Two

3311 implementation approaches for this hybrid strategy are where the web service's PDP retrieves global
3312 policy from a central policy server (a policy retrieval method) and caches it for a period of time for
3313 access control decisions, and where the policy server syndicates global policy to the web service PDP (a
3314 policy push method).

3315 Similar to the purely centralized model, there may be performance and availability issues related to all
3316 services using the policy retrieval method. However, the difference is that global policy is only
3317 downloaded on a periodic basis, compared to the strategy of every request. A potential issue arises
3318 related to the frequency of global policy changes in that there may be a lag time between when a global
3319 policy changes and when those policies are retrieved by the web services. This consideration is
3320 exacerbated when global policy changes are frequent.

3321 The policy push method alleviates performance and availability concerns, as policy is pushed immediately
3322 to the web service PDPs when changes occur. This alleviates the time lag concern of the policy pull
3323 method, because as policy is changed, it is immediately sent to the web service PDPs. Security of
3324 syndication in this model does need to be addressed; the security model needs to protect against entities
3325 "spoofing" the syndication server, replaying syndication policy update messages, tampering with
3326 messages in transit, and denial of service attacks that would prevent syndicated policy messages from
3327 being received by all PDPs. Both methods discussed in this section are effective for access control in a
3328 SOA, but depend on each organization's security and performance requirements.

3329 5.2.6 Architectural Implications of SOA Security

3330 Providing SOA security in an ecosystem of governed services has the following implications on the policy
3331 support and the distributed nature of mechanisms used to assure SOA security:

- 3332 • Security expressed through security messaging policies **SHOULD follow** the same architectural
3333 implications as described in Section 4.4.3 for policies and contracts architectural implications.
- 3334 • Security policies **MUST have** mechanisms to support security description administration, storage,
3335 and distribution.
- 3336 • Service descriptions supporting security policies **SHOULD**:
 - 3337 ○ have a meta-structure sufficiently rich to support security policies;
 - 3338 ○ be able to reference one or more security policy artifacts;
 - 3339 ○ have a framework for resolving conflicts between security policies.
- 3340 • The mechanisms that make-up the execution context in secure SOA-based systems **SHOULD**:
 - 3341 ○ provide protection of the confidentiality and integrity of message exchanges;
 - 3342 ○ be distributed so as to provide centralized or decentralized policy-based identification,
3343 authentication, and authorization;
 - 3344 ○ ensure service availability to consumers;
 - 3345 ○ be able to scale to support security for a growing ecosystem of services;
 - 3346 ○ be able to support security between different communication technologies;
- 3347 • Common security services **SHOULD** include:
 - 3348 ○ services that abstract encryption techniques;
 - 3349 ○ services for auditing and logging interactions and security violations;
 - 3350 ○ ~~services for identification;~~
 - 3351 ○ services for authentication;
 - 3352 ○ services for the retrieval of authorization credentials (attribute services);
 - 3353 ○ services for making authorization decisions (policy decision point services);
 - 3354 ○ services for enforcing access control policies;
 - 3355 ○ services for providing tokens (security token services) representing identity, authorizing
3356 credentials, or authorization decisions that can be used in service interactions;
 - 3357 ○ services for intrusion detection and prevention;
 - 3358 ○ services for availability including support for quality of service specifications and metrics.

3359 5.3 Management Model

3360 5.3.1 Management

3361 Management is a process of controlling resources in accordance with the policies and principles defined
3362 by Governance.

3363 There are three separate but linked domains of interest within the management of a SOA ecosystem:

- 3364 1. the management and support of the resources that are involved in any complex structures – of
3365 which SOA ecosystems are excellent examples;
- 3366 2. the promulgation and enforcement of the policies and service contracts agreed to by the
3367 stakeholders in the SOA ecosystem;
- 3368 3. the management of the relationships of the participants – both to each other and to the services
3369 that they use and offer.

3370 There are many artifacts related to management. Historically, systems management capabilities have
3371 been organized by the FCAPS functions (based on ITU-T Rec. M.3400 (02/2000), *TMN Management*
3372 *Functions*):

- 3373 • fault management,
- 3374 • configuration management,
- 3375 • account management,
- 3376 • performance and security management.

3377 The primary task of the functional groups is to concentrate on maintaining systems in a trusted, active,
3378 and accessible state.

3379 In the context of the SOA ecosystem, we see many possible resources that may require management
3380 such as services, service descriptions, service contracts, policies, roles, relationships, security, people
3381 and systems that implement services and infrastructure elements. In addition, given the ecosystem
3382 nature, it is also potentially necessary to manage the business relationships between participants.

3383 Successful operation of a SOA ecosystem requires trust among the stakeholders and between them and
3384 the SOA-based system elements. In contrast, regular systems in technology are not necessarily operated
3385 or used in an environment requiring trust before the stakeholders make use of the system. Indeed, many
3386 of these systems exist in hierarchical management structures, within which use may be mandated by
3387 legal requirement, executive decision, or good business practice in furthering the business' strategy. The
3388 pre-condition of trust in the SOA ecosystem is rooted both in the principles of service orientation and in
3389 the distributed, authoritative ownership of independent services. Even for hierarchical management
3390 structures applied to a SOA ecosystem, the service in use should have a contractual basis rather than
3391 solely being mandated.

3392 Trust may be established through agreements/contracts, policies, or implicitly through observation of
3393 repeated interactions with others. Explicit trust is usually accompanied by formalized documents suitable
3394 for management. Implicit trust adds fragility to the management of a SOA ecosystem because failure to
3395 maintain consistent and predictable interactions will undermine the trust between participants and within
3396 the ecosystem as a whole.

3397 Management in a SOA ecosystem is thus concerned with management taking actions that will establish
3398 the condition of trust that must be present before engaging in service interactions. These concerns should
3399 largely be handled within the governance of the ecosystem. The policies, agreements, and practices
3400 defined through governance provide the boundaries within which management operates and for which
3401 management must provide enforcement and feedback. However, governance alone cannot foresee all
3402 circumstances but must offer sufficient guidance where agreement between all stakeholders cannot be
3403 reached. Management in these cases must be flexible and adaptable to handle unanticipated conditions
3404 without unnecessarily breaking trust relationships.

3405 Service management is the process – manual, automated, or a combination – of proactively monitoring
3406 and controlling the behavior of a service or a set of services. Service management operates under
3407 constraints attributed to the business and social context. Specific policies may be used to govern cross-
3408 boundary relationships. Managing solutions based on such policies (and that may be used across
3409 ownership boundaries) raises issues that are not typically present when managing a service within a

3410 single ownership domain. Care is therefore required in managing a service when the owner of the
 3411 service, the provider of the service, the host of the service and mediators to the service may all belong to
 3412 different stakeholders.

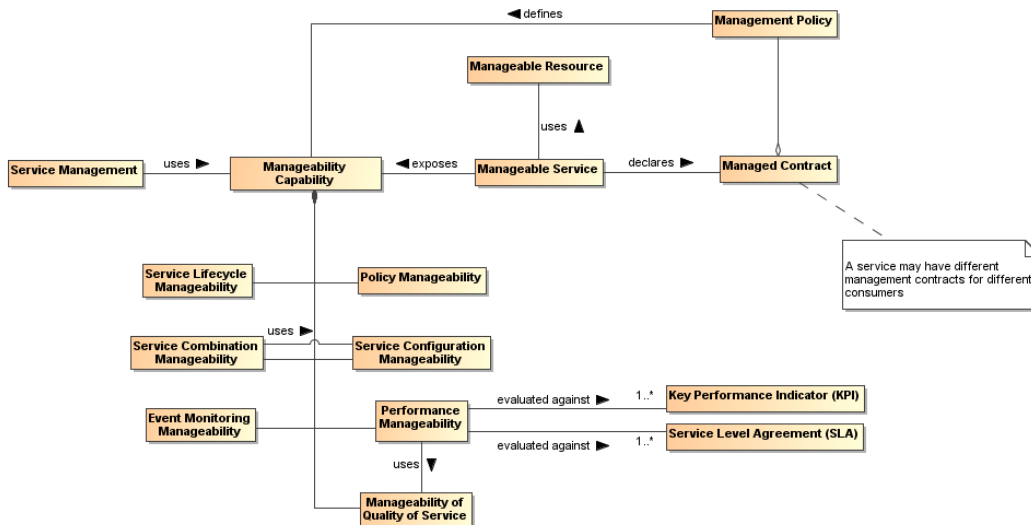
3413 Cross-boundary service management takes place in, at least, the following situations:

- 3414 • using combinations of services that belong to different ownership domains
- 3415 • using of services that mediate between ownership domains
- 3416 • sharing monitoring and reporting means and results.

3417 These situations are particularly important in ecosystems that are highly decentralized, in which the
 3418 participants interact as peers as well as in the 'master-servant' mode.

3419 The management model shown in Figure 42 conveys how the SOA paradigm applies to managing
 3420 services. Services management operates via service metadata, such as properties associated with
 3421 service lifecycles and with service use, which are typically collected in or accessed through the service
 3422 description.

3423



3424

3425 *Figure 42 - Management model in SOA ecosystem*

3426 The service metadata of interest is that set of service properties that is manageable. These manageability
 3427 properties are generally identifiable for any service consumed or supplied within the ecosystem. The
 3428 necessary existence of these properties within the SOA ecosystem motivates the following definitions:

3429 **Manageability**

3430 A capability that allows a **resource** to be controlled, monitored, and reported on with respect to
 3431 some properties.

3432 **Manageability property**

3433 A property used in the **manageability** of a **resource**. The fundamental unit of management in
 3434 systems management.

3435 Note that manageability is not necessarily a part of the managed entities themselves and are generally
 3436 considered to be external to the managed entities.

3437 Each resource may be managed through a number of aspects of management, and the resources may
 3438 be grouped based on similar aspects. For example, resources may be grouped according to the aspect
 3439 referred to as 'Configuration Manageability' for the collection of services. Some resources may not be
 3440 managed under a particular capability if there are no manageability aspects with a clear meaning or use.
 3441 As an example, all resources within a SOA ecosystem have a lifecycle that is meaningful within the

3442 ecosystem. Thus, all resources are manageable under Lifecycle Manageability. In contrast, not all
3443 resources report or handle events. Thus, Event Manageability is only concerned with those resources for
3444 which events are meaningful.

3445 **Life-cycle Manageability** of a service typically refers to how the service is created, how it is **destroyed**
3446 **retired** and how service versions must be managed. This manageability is a feature of the SOA
3447 ecosystem because the service cannot manage its own life cycle. **Related properties may include the**
3448 **necessary state of the ecosystem for the creation and retirement of the service and the state of the**
3449 **ecosystem following the retirement of the service. The SOA ecosystem distinguishes between service**
3450 **composition and service aggregation: retiring of service composition leads to retiring of all services**
3451 **comprising the composition while retiring of service aggregation assumes that comprising services have**
3452 **their own life-cycle and can be used in another aggregation.**

Comment [KJL107]: Issue 146

3453 Another important consideration is that services may have resource requirements, **such as concurrent**
3454 **connectivity to a data source**, which must be established at various points in the services' life cycles.
3455 However, actual providers of these resources may not be known at the time of the service creation and,
3456 thus, have to be managed at service run-time.

Comment [KJL109]: Issue 147

3457 **Combination Manageability** of a service addresses management of service characteristics that allow for
3458 creating and changing combinations in which the service participates or that the service combines itself.
3459 Known models of such combinations are aggregations and compositions. Examples of patterns of
3460 combinations are choreography and orchestration. **In cases of business collaboration, combination of**
3461 **services appears as cooperation of services.** Combination Manageability drives implementation of the
3462 Service Composability Principle of service orientation.

Comment [PFB110]: Issue 145

3463 Service combination manageability resonates with the methodology of process management.
3464 Combination Manageability may be applied at different phases of service creation and execution and, in
3465 some cases, can utilize Configuration Manageability.

3466 Service combinations typically contribute the most in delivering business values to the stakeholders.
3467 Managing service combinations is the one of the most important tasks and features of the SOA
3468 ecosystem.

3469 **Configuration Manageability** of a service allows managing the identity of and the interactions among
3470 internal elements of the service, **for example, a use of data encryption for internal inter-component**
3471 **communication in particular deployment conditions.** Also, Configuration Manageability correlates with the
3472 management of service versions and configuration of the deployment of new services into the ecosystem.
3473 Configuration Management differs from the Combination Manageability in the scope and scale of
3474 manageability, and addresses lower level concerns than the architectural combination of services.

Comment [PFB111]: Issue 145

3475 **Event Monitoring Manageability** allows managing the categories of events of interest related to services
3476 and reporting recognized events to the interested stakeholders. Such events may be the ones that trigger
3477 service invocations as well as execution of particular functionality provided by the service. **For example,**
3478 **an execution of a set of financial market risk services, which implements choreography pattern, may be**
3479 **started if certain financial event occurs in a stock exchange.**

Comment [PFB112]: Issue 145

3480 Event Monitoring Manageability is a key lower-level manageability aspect, in which the service provider
3481 and associated stakeholders are interested. Monitored events may be internal or external to the SOA
3482 ecosystem. For example, a disaster in the oil industry, which is outside the SOA ecosystem of the Insurer,
3483 can trigger the service's functionality that is responsible for immediate or constant monitoring of oil prices
3484 in the oil trading exchanges and, respectively, modify the premium paid by the insured oil companies.

3485 **Performance Manageability** of a service allows controlling the service results, shared and sharable real
3486 world effects against the business goals and objectives of the service. This manageability assumes
3487 monitoring of the business performance as well as the management of this monitoring itself. Performance
3488 Manageability includes business and technical performance manageability through a performance criteria
3489 set, such as business key performance indicators (KPI) and service-level agreements (SLA).

3490 The performance business- and technical-level characteristics of the service should be known from the
3491 service contract. The service provider and consumer must be able to monitor and measure these
3492 characteristics or be informed about the results measured by a third party. **An example of such monitoring**
3493 **would be when the comparison of service performance results against an SLA is not satisfactory to the**
3494 **consumer, and as a consequence, the consumer# may replace the service by a service from a**
3495 **competitor.**

Comment [PFB113]: Issue 145

3496 Performance Manageability is the instrument for providing compliance of the service with its service
 3497 contracts. Performance Manageability utilizes Manageability of Quality of Service.

3498 **Manageability of Quality of Service** deals with management of service non-functional characteristics
 3499 that may be of significant value to the service consumers and other stakeholders in the SOA ecosystem.
 3500 A classic example- of this is managing bandwidth offerings associated with a service.

3501 Manageability of quality of service assumes that the properties associated with service qualities are
 3502 monitored during the service execution. Results of monitoring may be compared against an SLA or a KPI,
 3503 which results in the continuous validation of how the service contract is preserved by the service provider.

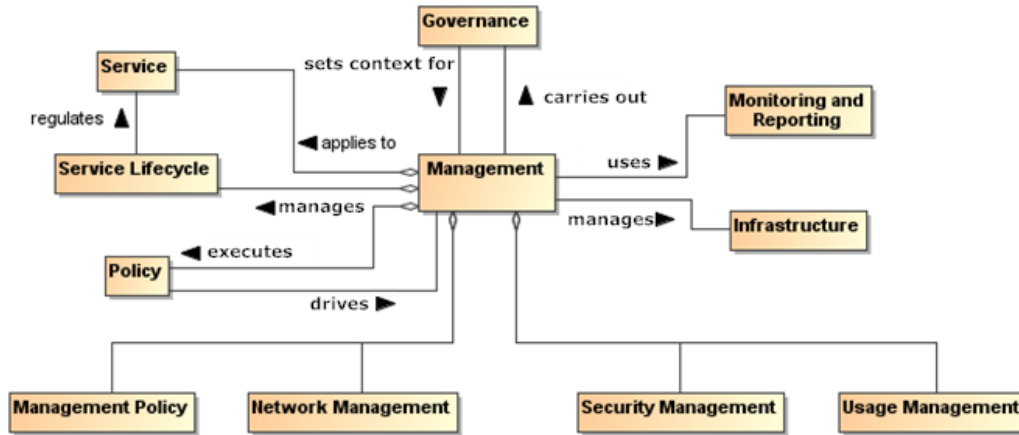
3504 **Policy Manageability** allows additions, changes and replacements of the policies associated with a
 3505 resource in the SOA ecosystem. The ability to manage those policies (such as promulgating policies,
 3506 retiring policies and ensuring that policy decision points and enforcement points are current) enables the
 3507 ecosystem to apply policies and *evaluate* the results.

3508 The capability to manage, i.e. use a particular manageability, requires policies from governance to be
 3509 translated into detailed rules and regulations which are measured and monitored providing corresponding
 3510 feedback for enforcement. At the same time, the execution of a management capability must adhere to
 3511 certain policies governing the management itself. For example, a management has to enforce and control
 3512 policies of compliance with particular industry regulation while the management is obliged by another
 3513 policy to report on the compliance status periodically.

3514 Management of SOA ecosystem recognizes the manageability challenge and requires manageability
 3515 properties to be considered for all aforementioned manageability cases. In the following subsections, we
 3516 describe how these properties are used in the management as well as some relationships between
 3517 management and other components of SOA ecosystem.

3518 5.3.2 Management Means and Relationships

3519 A minimal set of management issues for the SOA ecosystem is shown in Figure 43 and elaborated in the
 3520 following sections.



3521
 3522 *Figure 43 - Management Means and Relationships in a SOA ecosystem*

3523 5.3.2.1 Management Policy

3524 The management of resources within the SOA ecosystem may be governed by management policies. In
 3525 a deployed SOA-based solution, it may well be that different aspects of the management of a given
 3526 service are managed by different management services. For example, the life-cycle management of
 3527 services often involves managing service versions. Managing quality of service is often very specific to
 3528 the service itself; for example, quality of service attributes for a video streaming service are quite different
 3529 to those for a banking system.

3530 5.3.2.2 Network Management

3531 Network management deals with the maintenance and administration of large scale physical networks
3532 such as computer networks and telecommunication networks. Specifics of the networks may affect
3533 service interactions from performance and operational perspectives.

3534 Network and related system management **execute** a set of functions required for controlling, planning,
3535 deploying, coordinating, and monitoring the distributed services in the SOA ecosystem. However, while
3536 recognizing their importance, the specifics of systems management or network management are out of
3537 scope for this Reference Architecture Foundation.

Comment [KJL114]: Issue 153

3538 5.3.2.3 Security Management

3539 Security Management includes identification of roles, permissions, access rights, and policy attributes
3540 defining security boundaries and events that may trigger a security response.

3541 Security management within a SOA ecosystem is essential to maintaining the trust relationships between
3542 participants residing in different ownership domains. Security management must consider not just the
3543 internal properties related to interactions between participants but ecosystem properties that preserve the
3544 integrity of the ecosystem from external threats.

3545 5.3.2.4 Usage Management

3546 Usage Management is concerned with how resources are used, including:

- 3547 • how the resource is accessed, who is using the resource, and the state of the resource (access
3548 properties);
 - 3549 • controlling or shaping demand for resources to optimize the overall operation of the ecosystem
3550 (demand properties);
 - 3551 • assigning costs to the use of resources and distributing those cost assignments to the
3552 participants in an appropriate manner (financial properties).
- 3553

3554 5.3.3 Management and Governance

3555 The primary role of governance in the context of a SOA ecosystem is to foster an atmosphere of
3556 predictability, trust, and efficiency, and it accomplishes this by allowing the stakeholders to negotiate and
3557 set the key policies that govern the running of the SOA-based solution. Recall that in an ecosystem
3558 perspective, the goal of governance is less to have complete fine-grained control but more to enable the
3559 individual participants to work together.

3560 Policies for a SOA ecosystem will tend to focus on the rules of engagement between participants; for
3561 example, what kinds of interactions are permissible, how disputes are resolved, etc. While governance
3562 may primarily focus on setting policies, management will focus on the realization and enforcement of
3563 policies. Effective management in the SOA ecosystem requires an ability for governance to understand
3564 the consequences of its policies, guidelines, and principles, and to adjust those as needed when
3565 inconsistencies or ambiguity become evident from the operation of the management functions. This
3566 understanding and adjustment must be facilitated by the results of management and so the mechanisms
3567 for providing feedback from management into governance must exist.

3568 Governance operates via specialized activities and, thus, should be managed itself. Governance policies
3569 are included in the Governance Framework and Processes, and driven by the enterprise business model,
3570 business objectives and strategies. Where corporate management policies exist, these are usually guided
3571 and directed by the corporate executives. In peer relationships, governance policies are set by either an
3572 external entity and accepted by the peers or by the peers themselves. This creates the appropriate
3573 authoritative level for the policies used for the management of the Governance Framework and
3574 Processes. Management to operationalize governance controls the life-cycle of the governing policies,
3575 including procedures and processes, for modifying the Governance Framework and Processes.

3576 5.3.4 Management and Contracts

3577 5.3.4.1 Management for Contracts and Policies

3578 As we noted above, management can often be viewed as the application of contracts and individual
3579 policies to ensure the smooth running of the SOA ecosystem. Policies and service contracts specify the
3580 service characteristics that have to be monitored, analyzed and managed. These also play an important
3581 role as the guiding constraints for management, as well as being artifacts (e.g., policy and contractual
3582 documents) that also need to be managed.

3583 5.3.4.2 Contracts

3584 As described in sections *Participation in a SOA Ecosystem* view and *Realization of a SOA*
3585 *Ecosystem* view, there are several types of contractual information in the SOA ecosystem. From the
3586 management perspective, three basic types of the contractual information relate to:

- 3587 • relationship between service provider and consumer;
- 3588 • communication with the service;
- 3589 • control of the quality of the service execution.

3590 When a consumer prepares to interact with a service, the consumer and the service provider must come
3591 to an agreement on the service features and characteristics that will be provided by the service and made
3592 available to the consumer. This agreement is known as a service contract.

3593 **Service Contract**

3594 An implicit or explicit documented agreement between the service **consumer** and service
3595 **provider** about the use of the service based on

- 3596 • the commitment by a service provider to provide service functionality and results consistent
3597 with identified **real world effects** and
- 3598 • the commitment by a service **consumer** to interact with the service per specific means and
3599 per specified **policies**,

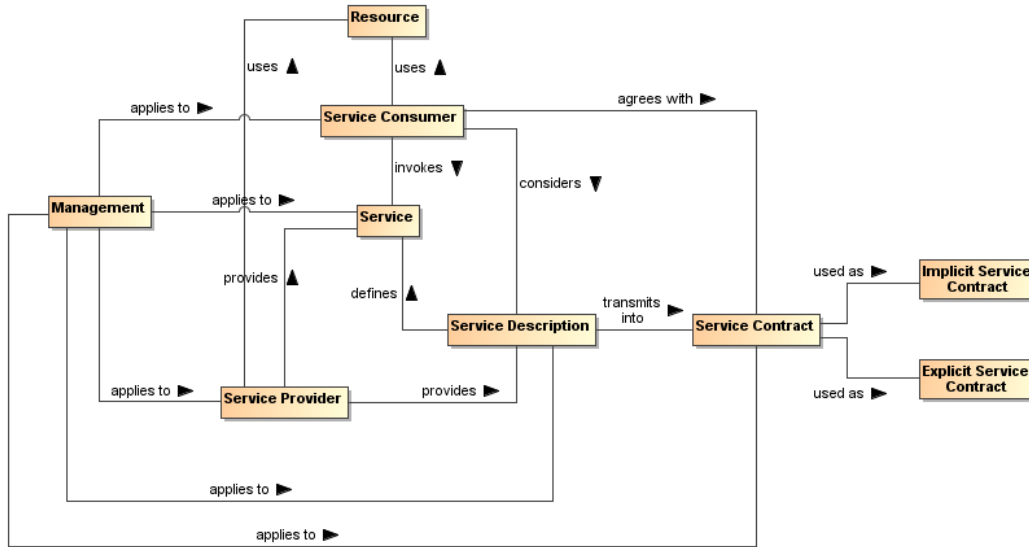
3600 where both consumer and provider actions are in the manner described in the service description.

3601 The service description provides the basis for the service contract and, in some situations, may be used
3602 as an implicit default service contract. In addition, the service description may set mandatory aspects of a
3603 service contract, e.g. for security services, or may specify acceptable alternatives. As an example of
3604 alternatives, the service description may identify which versions of a vocabulary will be recognized, and
3605 the specifics of the contract are satisfied when the consumer uses one of the alternatives. Another
3606 alternative could have a consumer identify a policy they require be satisfied, e.g. a standard privacy policy
3607 on handling personal information, and a provider that is prepared to accept a policy request would
3608 indicate acceptance as part of the service contract by continuing with the interaction. In each of these
3609 cases, the actions of the participants are consistent with an implicit service contract without the existence
3610 of a formal agreement between the participants.

3611 In the case of business services, it is anticipated that the service contract may take an explicit form and
3612 the agreement between business consumer and business service provider is formalized. Formalization
3613 requires up-front interactions between service consumer and service provider. In many business
3614 interactions, especially between business organizations within or across corporate boundaries, a
3615 consumer needs a contractual assurance from the provider or wants to explicitly indicate choices among
3616 alternatives, e.g., only use a subset of the business functionality offered by the service and pay a prorated

3617

COST.



3618

3619

Figure 44 - Management of the service interaction

3620 Consequently, an implicit service contract is an agreement (1) on the consumer side with the terms,
3621 conditions, features and interaction means specified in the service description "as is" or (2) a selection
3622 from alternatives that are made available through mechanisms included in the service description, and
3623 neither of these require any a priori interactions between the service consumer and the service provider.

3624 For example, a browser interface may display a checked box indicating the consumer agrees to accept
3625 future advertisement; the consumer can uncheck the box to indicate advertisements should not be sent.

Comment [KL115]: Issue 158

3626 An explicit service contract always requires a form of interaction between the service consumer and the
3627 service provider prior to the service invocation. This interaction may regard the choice or selection of the
3628 subset of the elements of the service description or other alternatives introduced through the formal
3629 agreement process that would be applicable to the interaction with the service and affect related joint
3630 action.

3631 Any form of explicit contract couples the service consumer and provider. While explicit contracts may be
3632 necessary or desirable in some cases, such as in supply chain management, commerce often uses a mix
3633 of implicit and explicit contracts, and a service provider may offer (via service description) a conditional
3634 shift from implicit to explicit contract. For example, Twitter offers an implicit contract on the use of its APIs
3635 to any application with the limit on the amount of service invocations; if the application needs to use more
3636 invocations, one has to enter into the explicit fee-based contract with the provider. A case where an
3637 implicit contract transforms into an explicit contract may be illustrated when one buys a new computer and
3638 it does not work. The buyer returns the computer for repair under the manufacturer's warranty as stated
3639 by an implicit purchase contract. However, if the repair does not fix the problem and the seller offers an
3640 upgraded model in replacement, the buyer may agree to an explicit contract that limits the rights of the
3641 buyer to make the explicit agreement public.

3642 Control of the quality of the service execution, often represented as a service level agreement (SLA), is
3643 performed by service monitoring systems and includes both technical and operational business controls.
3644 SLA is a part of the service contract and, because of the individual nature of such contracts, may vary
3645 from one service contract to another, even for the same consumer. Typically, a particular SLA in the
3646 service contract is a concrete instance of the SLA declared in the service description.

3647 Management of the service contracts is based on management policies that may be mentioned in the
3648 service description and in the service contracts. Management of the service contracts is mandatory for
3649 consumer relationship management. In the case of explicit service contracts, the contracts have to be
3650 created, stored, maintained, reviewed/controlled and archived/destroyed as needed. All the activities are

3651 management concerns. Explicit service contracts may be stored in specialized repositories that provide
3652 appropriate level of security.

3653 Management of the service interfaces is based on several management policies that regulate

- 3654 • availability of interfaces specified in the service contracts,
- 3655 • accessibility of interfaces,
- 3656 • procedures for interface changes,
- 3657 • interface versions as well as the versions of all parts of the interfaces,
- 3658 • traceability of the interfaces and their versions back to the service description document.

3659 Management of the SLA is integral to the management of service monitoring and operational service
3660 behavior at run-time. An SLA usually enumerates service characteristics and expected performances of
3661 the service. Since an SLA carries the connotation of a 'promise', monitoring is needed to know if the
3662 promise is being kept. Existence of an SLA itself does not guarantee that the consumer will be provided
3663 with the service level specified in the service contract.

3664 The use of an SLA in a SOA ecosystem can be wider than just an agreement on technical performances.
3665 An SLA may contain remedies for situations where the promised service cannot be maintained, or the
3666 real world effect cannot be achieved due to developments subsequent to the agreement. A service
3667 consumer that acts accordingly to realize the real world effect may be compensated for the breach of the
3668 SLA if the effect is not realized.

3669 Management of the SLA includes, among others, policies to change, update, and replace the SLA. This
3670 aspect concerns service Execution Context because the business logic associated with a defined
3671 interface may differ in different Execution Contexts and affect the overall performance of the service.

3672 **5.3.4.3 Policies**

3673 "Although provision of management capabilities enables a service to become manageable, the extent and
3674 degree of permissible management are defined in management policies that are associated with the
3675 services. Management policies are used to define the obligations for, and permissions to, managing the
3676 service" [WSA]. Management policies, in essence, are the realization of governing rules and regulations.
3677 As such, some management policies may target services while other policies may target the management
3678 of the services.

3679 In practice, a policy without any means of enforcing it is vacuous. In the case of management policy, we
3680 rely on a management infrastructure to realize and enforce management policy.

3681 **5.3.4.4 Service Description and Management**

3682 The service description identifies several management objects such as a set of service interfaces and
3683 related set of SLAs. Service behavioral characteristics and performances specified in the SLA depend on
3684 the interface type and its Execution Context. In the service description, a service consumer can find
3685 references to management policies, SLA metrics, and the means of accessing measured values that
3686 together increase assurance in the service quality. At the same time, service description is an artifact that
3687 needs to be managed.

3688 In the SOA ecosystem, the service description is the assembled information that describes the service but
3689 it may be reported or displayed in different presentations. While each separate version of the service has
3690 one and only one service description, different categories of service consumers may focus their interests
3691 on different aspects of the service description. Thus, the same service description may be displayed not
3692 only in different languages but also with different cultural and professional accents in the content.

3693 New service description may be issued to reflect changes and update in the service. If the change in the
3694 service does not affect its service description, the new service version may have the same service
3695 description as the previous version except for the updated version identifier. For example, a service
3696 description may stay the same if bugs were fixed in the service. However, if a change in the service
3697 influences any aspects of the service quality that can affect the real world effect resulting from
3698 interactions with the service, the service description must reflect this change even if there are no changes
3699 to the service interface.

3700 Management of the service description as well as of the explicit service contracts is essential for delivery
3701 of the service to the consumer satisfaction. This management can also prevent business problems rooted
3702 in poor communication between the service consumers and the service providers.

3703 Thus, management of service description contains, among others, management of the service description
3704 presentations, the life-cycles of the service descriptions, service description distribution practices and
3705 storage of the service descriptions and related service contracts. Collections of service descriptions in the
3706 enterprise may manifest a need for specialized registries and/or repositories. Depending on the enterprise
3707 policies, an allocation of purposes and duties of registries and repositories may vary but this topic is
3708 beyond the current scope.

3709 **5.3.5 Management for Monitoring and Reporting**

3710 The successful application of management relies on the monitoring and reporting aspects of management
3711 to enable the control aspect. Monitoring in the context of management consists of measuring values of
3712 managed aspects and evaluating that measurement in relationship to some expectation. Monitoring in a
3713 SOA ecosystem is enabled through the use of mechanisms by resources for exposing managed aspects.
3714 In the SOA framework, this mechanism may be a service for obtaining the measurement. Alternatively,
3715 the measurement may be monitored by means of event generation containing updated values of the
3716 managed aspect.

3717 Approaches to monitoring may use a polling strategy in which the measurements are requested from
3718 resources in periodic intervals, in a pull strategy in which the measurements are requested from
3719 resources at random times, or in a push strategy in which the measurements are supplied by the resource
3720 without request. The push strategy can be used in a periodic update approach or in an 'update on
3721 change' approach. Management services must be capable of handling these different approaches to
3722 monitoring.

3723 Reporting is the complement to monitoring. Where monitoring is responsible for obtaining measurements,
3724 reporting is responsible for distributing those measurements to interested stakeholders. The separation
3725 between monitoring and reporting is made to include the possibility that data obtained through monitoring
3726 might not be used until an event impacting the ecosystem occurs or the measurement requires further
3727 processing to be useful. In the SOA framework, reporting is provided using services for requesting
3728 measurement reports. These reports may consist of raw measurement data, formatted collections of data,
3729 or the results of analysis performed on measurement data from collections of different managed aspects.
3730 Reporting is also used to support logging and auditing capabilities, where the reporting mechanisms
3731 create log or audit entries.

3732 **5.3.6 Management for Infrastructure**

3733 All of the properties, policies, interactions, resources, and management are only possible if a SOA
3734 ecosystem infrastructure provides support for managed capabilities. Each managed capability imposes
3735 different requirements on the capabilities supplied by the infrastructure in SOA ecosystem and requires
3736 that those capabilities be usable as services or at the very least be interoperable.

3737 While not providing a full list of infrastructural elements of a SOA ecosystem, we list some examples here:

- 3738 1. Registries and repositories for services, policies, and related descriptions and contracts
- 3739 2. Synchronous and asynchronous communication channels for service interactions (e.g., network,
3740 e-mail, message routing with ability of mediating transport protocols, etc.)
- 3741 3. Recovery capabilities
- 3742 4. Security controls

3743 A SOA ecosystem infrastructure, enabling service management, should also support:

- 3744 1. Management enforcement and control means
- 3745 2. Monitoring and SLA validation controls
- 3746 3. Testing and Reporting capabilities

3747 The combination of manageability [properties, related](#) capabilities and infrastructure elements constitutes
3748 a certain level of SOA management maturity. While several maturity models exist, this topic is out of the
3749 scope of the current document.

3750

5.3.7 Architectural Implication of the SOA Management

3751
3752
3753

SOA Management is one of the fundamental elements of the SOA ecosystem; it impacts all aspects of a service life-cycle, service activities and actions, and a service usage. The key choices that must be made center on management means, methods and manageability properties:

3754
3755

- Every resource of the SOA ecosystem and, particularly, services **MUST** provide manageability properties

3756
3757
3758

- The set of manageability properties **SHOULD** include as minimum such properties as life-cycle, combination, configuration, event monitoring, performance, quality of services, and policy manageability

3759
3760

- Combinations of manageability properties **MAY** be used in different management methods and tools

3761
3762

- Manageability properties and applicable policies **SHOULD** be appropriately described in the services description and contracts

3763
3764

- Management processes **SHOULD** operate (control, enforce and provide a feedback to the governance) via policies, agreements/contracts, and practices defined through governance

3765
3766

- Management functions and information **MAY** be realized as services and, thus, **MUST** be managed itself

3767
3768

- Management in the cases, where sufficient guidance is unavailable or for which agreement between all stakeholders cannot be reached, **MUST** be flexible and adaptable to handle unanticipated conditions without unnecessarily breaking trust relationships

3769
3770

- Management **SHOULD** engage a monitoring mechanism to enable manageability. Monitoring **HAS to MUST** include

3771
3772

- Access mechanisms to collected SLA metrics
- Assessment mechanisms to compare metrics against policies and contracts

3773
3774
3775

- Results of monitoring and reporting **MUST** be made accessible to participants in different ownership domains.

Comment [PFB116]: Issue 160

3776

5.4 SOA Testing Model

3777
3778
3779
3780
3781
3782
3783
3784
3785

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

3786

5.4.1 Traditional Software Testing as Basis for SOA Testing

3787
3788
3789
3790
3791
3792

SOA services are largely software artifacts and can leverage the body of experience that has evolved around software testing. [IEEE-829] specifies the basic set of software test documents while allowing flexibility for tailored use. Many testing frameworks are available but the SOA-RAF does not prescribe the use of any one in particular and choice will be driven by a framework that offers the right amount and level of testing. As such, IEEE-829 can provide guidance to SOA testing and a point of reference for additional test concerns introduced by a SOA approach.

Comment [PFB117]: Issue 299

3793
3794
3795
3796
3797
3798

IEEE-829 covers test specification and test reporting through use of several document types, including test plans; test design, test case, and test procedure specifications; and documents to identify, log, and report on test occurrences and artifacts. In summary, IEEE-829 captures (1) what was tested, (2) how it was tested, e.g. the test procedure used, and (3) the results of the test. While the SOA-RAF does not require IEEE-829 artifacts, those with responsibilities for testing should consider how aspects of IEEE-829 apply.

3799 5.4.1.1 Types of Testing

3800 There are numerous aspects of testing that, in total, work to establish that an entity is (1) built as required
3801 per policies and related specifications prescribed by the entity's owner, and (2) delivers the functionality
3802 required by its intended users. This is often referred to as verification and validation.

3803 In Section 4.4, Policies are described that can be related to testing. These policies may prescribe but are
3804 not limited to the business processes to be followed. Policies may also prescribe the standards with which
3805 an implementation must comply, as well as the qualifications of and restrictions on the users. In addition
3806 to the functional requirements prescribing what an entity does, there may also be non-functional
3807 performance and/or quality metrics that state how well the entity performs. The relation of these policies
3808 to SOA testing is discussed further below.

Comment [KJL118]: Issue rb40

3809 The identification of policies is the purview of governance (section 5.1) and the assuring of compliance
3810 (including response to noncompliance) with policies is a matter for management (section 5.3).

3811 5.4.1.2 Range of Test Conditions

3812 Test conditions and expected responses are detailed in the test case specification. The test conditions
3813 should be designed to cover the areas for which the entity's response must be documented and may
3814 include:

- 3815 • nominal conditions
- 3816 • boundaries and extremes of expected conditions
- 3817 • breaking point where the entity has degraded below a certain level or has otherwise ceased
3818 effective functioning
- 3819 • random conditions to investigate unidentified dependencies among combinations of conditions
- 3820 • errors conditions to test error handling

3821 The specification of how each of these conditions should be tested for SOA resources, including the
3822 infrastructure elements of the SOA ecosystem, is beyond the scope of this document but is an area that
3823 evolves along with operational SOA experience.

3824 5.4.1.3 Configuration Management of Test Artifacts

3825 The test item transmittal provides an unambiguous identification of the entity being tested, thus
3826 REQUIRING that the configuration of the entity is appropriately tracked and documented. In addition, the
3827 test documents (such as those specified by IEEE-829) must also be under a documented and
3828 appropriately audited configuration management process, as should other resources used for testing. The
3829 description of each artifact would follow the general description model as discussed in section 4.1.1.1; in
3830 particular, it would include a version number for the artifact and reference to the documentation
3831 describing the versioning scheme from which the version number is derived.

3832 5.4.2 Testing and the SOA Ecosystem

3833 Testing of SOA artifacts for use in the SOA ecosystem differs from traditional software testing for several
3834 reasons. These include a difference in what constitutes the consumer community and what constitutes
3835 the evolving environment that comprises the SOA ecosystem. In response, testing must include
3836 considerations for making a service testable throughout its lifetime.

3837 5.4.2.1 Testing and the Consumer Communities

3838 A highly touted benefit of SOA is to enable unanticipated consumers to make use of services for
3839 unanticipated purposes. Examples of this could include the consumer using a service for a result that
3840 was not considered the primary one by the provider or the service may be used in combination with other
3841 services in a scenario that is different from the one considered when designing for the initial target
3842 consumer community. It is unlikely that a new consumer will push the services back to anything
3843 resembling the initial test phase to test the new use, and thus additional paradigms for testing are
3844 necessary. The potential responsibilities related to such "consumer testing" are discussed further below.

Comment [KJL119]: Issues rb42, rb43, rb44 part

3845 In addition to consumers who interact with a service to realize the described **real world effects**, the
3846 developer community is also intended to be a consumer. In the SOA vision of reuse, the developer
3847 composes new solutions using existing services, where the existing services provide desired **real world**
3848 **effects** that are needed by the new solution. The composed solution must be tested for its intended
3849 functionality, and the component service may need particular attention if its use is different from its typical
3850 use as a separate offering. Note, the composition developer is not expected to own a private copy of a
3851 component service, and testing may be dependent on test interfaces provided by the component service.

3852 **5.4.2.2 Testing and the Evolving SOA Ecosystem**

3853 The distributed, unbounded nature of the SOA ecosystem makes it unlikely to have an isolated test
3854 environment that duplicates the operational environment. A traditional testing approach often makes use
3855 of a test system that is identical to the eventual operational system but isolated for testing. After testing is
3856 successfully completed, the tested entity would be migrated to the operational environment, or the test
3857 environment may be delivered as part of the system to become operational. This is not feasible for the
3858 SOA ecosystem as a whole.

Comment [KL120]: Issues rb42, rb43, rb44 part

3859 SOA services must be testable in the environment and under the conditions that can be encountered in
3860 the operational SOA ecosystem. As the ecosystem is in a state of constant change, so some level of
3861 testing is continuous through the lifetime of the service, leveraging utility services used by the ecosystem
3862 infrastructure to monitor its own health and respond to situations that could lead to degraded
3863 performance. This implies the test resources must incorporate aspects of the SOA paradigm, and a
3864 category of services may be created to specifically support and enable effective monitoring and
3865 continuous testing for **resources** participating in the SOA ecosystem.

3866 While SOA within an enterprise may represent a more constrained and predictable operational
3867 environment, the composability and unanticipated use aspects are highly touted within the enterprise.
3868 The expanded perspective on testing may not be as demanding within an enterprise but fuller
3869 consideration of the ecosystem enables the enterprise to be more responsive should conditions change.

3870 **5.4.3 Elements of SOA Testing**

3871 IEEE-829 emphasizes identifying what is to be tested, how it is to be tested, and by whom the testing is to
3872 be done. This is equally applicable to SOA testing.

3873 **5.4.3.1 What is to be Tested**

3874 The focus of this discussion is the SOA service. It is recognized that the infrastructure components of
3875 any SOA environment are likely to also be SOA services and, as such, falls under the same testing
3876 guidance. Other resources that contribute to a SOA environment may not be SOA services, but are
3877 expected to satisfy the intent if not the letter of guidance presented here.

3878 The following discussion often focuses on a singular SOA service but it is implicit that any service may be
3879 a composite of other services. As such, testing the functionality of a composite service may effectively be
3880 testing an end-to-end business process that is being provided by the composite service. If new versions
3881 are available for the component services, appropriate end-to-end testing of the composite may be
3882 required in order to verify that the composite functionality is still adequately provided. The level of
3883 required testing of an updated composite **service** depends on policies of those providing the service,
3884 policies of those using the service, and mission criticality of those depending on the service results.

Comment [KL121]: Issue rb41

3885 The SOA service to be tested **MUST** be unambiguously identified as specified by its applicable
3886 configuration management scheme. Specifying such a scheme is beyond the scope of the SOA-RAF
3887 other than to say the scheme should be documented and itself under configuration management.

3888 **5.4.3.1.1 Origin of Test Requirements**

3889 In the Service Description model (Figure 16), the aspects of a service that need to be described are:

- 3890 • the service functionality and technical assumptions that underlie the functionality;
- 3891 • the policies that describe conditions of use;
- 3892 • the service interface that defines information exchange with the service;

3893 • service reachability that identifies how and where message exchange is to occur; and
3894 • metrics access for any participant to have information on how a service is performing.

3895 Service testing must provide adequate assurance that each of these aspects is operational as defined.

3896 The information in the service description comes from different sources. The functionality is defined
3897 through whatever process identifies needs and the community for which these needs are addressed. The
3898 process may be ad hoc as serves the prospective service owner or strictly governed, but defining the
3899 functionality is an essential first step in development. It is also an early and ongoing focus of testing to
3900 ensure the service accurately reflects the described functionality and the described functionality
3901 accurately addresses the consumer needs.

3902 Policies define the conditions of development and conditions of use for a service and are typically
3903 specified as part of the governance process. Policies constraining service development, such as coding
3904 standards and best practices, require appropriate testing and auditing during development to ensure
3905 compliance. While the governance process identifies development policies, these are likely to originate
3906 from the technical community responsible for development activities. Policies that define conditions of use
3907 often define business practices that service owners and providers or those responsible for the SOA
3908 infrastructure want followed. These policies are initially tested during service development and are
3909 continuously monitored during the operational lifetime of the service.

3910 The testing of the service interface and service reachability are often related but essentially reflect
3911 different motivations and needs. The service interface is specified as a joint product of the service owners
3912 and providers who define service functionality, the prospective consumer community, the service
3913 developer, and the governance process. The semantics of the information model must align with the
3914 semantics of those who consume the service in order for there to be meaningful exchange of information.
3915 The structure of the information is influenced by the consumer semantics and the requirements and
3916 constraints of the representation as interpreted by the service developer. The service process model that
3917 defines actions which can be performed against a service and any temporal dependencies derive from
3918 the defined functionality and may be influenced by the development process. Any of these constraints
3919 may be identified and expressed as policy through the governance process.

3920 Service reachability conditions are the purview of the service provider who identifies the service endpoint
3921 and the protocols recognized at the endpoint. These may be constrained by governance decisions on
3922 how endpoint addresses may be allocated and what protocols should be used.

3923 While the considerations for defining the service interface derive from several sources, testing of the
3924 service interface is more straightforward and isolated in the testing process. At any point where the
3925 interface is modified or exposes a new resource, the message exchange should be monitored both to
3926 ensure the message reaches its intended destination and it is parsed correctly once received. Once an
3927 interface has been shown to function properly, it is unlikely to fail later unless something fundamental to
3928 the service changes.

3929 The service interface is also tested when the service endpoint changes. Testing of the endpoint ensures
3930 message exchange can occur at the time of testing and the initial testing shows the interface is being
3931 processed properly at the new endpoint. Functioning of a service endpoint at one time does not
3932 guarantee it is functioning at another time, e.g. the server with the endpoint address may be down,
3933 making testing of service reachability a continual monitoring function through the life of the service's use
3934 of the endpoint. Also, while testing of the service endpoint is a necessary and most commonly noted part
3935 of the test regiment, it is not in itself sufficient to ensure the other aspects of testing discussed in this
3936 section.

3937 Finally, governance is impossible without the collection of metrics against which service behavior can be
3938 assessed. Metrics are also a key indicator for consumers to decide if a service is adequate for their
3939 needs. For instance, the average response time or the recent availability can be determining factors even
3940 if there are no rules or regulations promulgated through the governance process against which these
3941 metrics are assessed. The available metrics are a combination of those expected by the consumer
3942 community and those mandated through the governance process. The total set of metrics will evolve over
3943 time with SOA experience. Testing of the services that gather and provide access to the metrics will follow
3944 testing as described in this section, but for an individual service, testing will ensure that the metrics
3945 access indicated in the service description is accurate.

3946 The individual test requirements highlight aspects of the service that testing must consider but testing
3947 must establish more than isolated behavior. The emphasis is the holistic results of interacting with the
3948 service in the SOA environment. Recall that the execution context is the set of agreements between a
3949 consumer and a provider that define the conditions under which service interaction occurs. The
3950 agreements are expected to be predominantly the acceptance of the standard conditions as enumerated
3951 by the service provider, but it may include the identification of alternate conditions that will govern the
3952 interaction.

3953 For example, the provider may prefer a policy where it can sell the contact information of its consumers
3954 but will honor the request of a consumer to keep such information private. The identification of the
3955 alternate privacy policy is part of the execution context, and it is the application of and compliance with
3956 this policy that operational monitoring will attempt to measure. The collection of metrics showing this
3957 condition is indeed met when chosen is considered part of the ongoing testing of the service.

3958 Other variations in the execution context also require monitoring to ensure that different combinations of
3959 conditions perform together as desired. For example, if a new privacy policy takes additional resources to
3960 apply, this may affect quality of service and propagate other effects. These could not be tested during the
3961 original testing if the alternate policy did not exist at that time.

3962 **5.4.3.1.2 Testing Against Non-Functional Requirements**

3963 Testing against non-functional requirements constitutes testing of business usability of the service. In a
3964 marketplace of services, non-functional characteristics may be the primary differentiator between services
3965 that produce essentially the same real world effects.

3966 As noted in the previous section, non-functional characteristics are often associated with policies or other
3967 terms of use and may be collected in service level contracts offered by the service providers. Non-
3968 functional requirements may also reflect the network and hardware infrastructure that support
3969 communication with the service, and changes may impact quality of service. The service consumer and
3970 even the service provider may not be aware of all such infrastructure changes but the changes may
3971 manifest in shared states that impact the usability of the service.

3972 In general, a change in the non-functional requirements results in a change to the execution context, but
3973 as with any collection of information that constitutes a description, the execution context is unable to
3974 explicitly capture all non-functional requirements that may apply. A change in non-functional
3975 requirements, whether explicitly part of the execution context or an implicit contributor, may require
3976 retesting of the service even if its functionality and the implementation of the functionality has not
3977 changed. Depending on the circumstances, retesting may require a formal recertifying of end-to-end
3978 behavior or more likely will be part of the continuous monitoring that applies throughout the service
3979 lifetime.

3980 **5.4.3.1.3 Testing Content and the Interests of Consumers**

3981 As noted in section 5.4.1.1, testing may involve verification of conformance with respect to policies and
3982 technical specifications and validation with respect to sufficiency of functionality to meet some prescribed
3983 use. It may also include demonstration of performance and quality aspects. For some of these items,
3984 such as demonstrating the business processes followed in developing the service or the use of standards
3985 in implementing the service, the testing or relevant auditing is done internal to the service development
3986 process and follows traditional software testing and quality assurance. If it is believed of value to potential
3987 consumers, information about such testing could be included in the service description. However, it is not
3988 required that all test or compliance artifacts be available to consumers, as many of the details tested may
3989 be part of the opacity of the service implementation.

3990 Some aspects of the service being tested will reflect directly on the real world effects realized through
3991 interaction with the service. In these cases, it is more likely that testing results will be directly relevant to
3992 potential consumers. For example, if the service was designed to correspond to certain elements of a
3993 business process or that a certain workflow is followed, testing should verify that the real world effects
3994 reflect that the business process or workflow were satisfactorily captured.

3995 The testing may also need to demonstrate that specified conditions of use are satisfied. For example,
3996 policies may be asserted that require certain qualifications of or impose restrictions on the consumers
3997 who may interact with the service. The service testing must demonstrate that the service independently

3998 enforces the policies or it provides the required information exchanges with the SOA ecosystem so other
3999 resources can ensure the specified conditions.

4000 The completeness of the testing, both in terms of the features tested and the range of parameters for
4001 which response is tested, depends on the context of expected use: the more critical the use, the more
4002 complete the testing. There are always limits on the resources available for testing, if nothing else than
4003 the service must be available for use in a finite amount of time.

4004 This again emphasizes the need for adequate documentation to be available. If the original testing is very
4005 thorough, it may be adequate for less demanding uses in the future. If the original testing was more
4006 constrained, then well-documented test results establish the foundation on which further testing can be
4007 defined and executed.

4008 **5.4.3.2 How Testing is to be Done**

4009 Testing should follow well-defined methodologies and, if possible, should reuse test artifacts that have
4010 proven generally useful for past testing. For example, IEEE-829 notes that test cases are separated from
4011 test designs to allow for use in more than one design and to allow for reuse in other situations. As with
4012 description of a service in the SOA ecosystem, description of testing artifacts enables awareness of the
4013 artifact and describes how the artifact may be accessed or used.

4014 As with traditional testing, the specific test procedures and test case inputs are important so the tests are
4015 unambiguously defined and entities can be retested in the future. Automated testing and regression
4016 testing may be more important in the SOA ecosystem in order to re-verify a service is still acceptable
4017 when incorporated in a new use. For example, if a new use requires the services to deal with input
4018 parameters outside the range of initial testing, the tests could be rerun with the new parameters. If the
4019 testing resources are available to consumers within the SOA ecosystem, the testing as designed by test
4020 professionals could be consumed through a service accessed by consumers, and their results could
4021 augment those already in place. This is discussed further in the next section.

4022 **5.4.3.3 Who Performs the Testing**

4023 As with any software, the first line of testing is unit testing done by software developers. It is likely that
4024 initial testing will be done by those developing the software but may also be done independently by other
4025 developers. For SOA development, unit testing is likely confined to a development sandbox isolated from
4026 the SOA ecosystem.

4027 SOA testing will differ from traditional software testing in that testing beyond the development sandbox
4028 must incorporate aspects of the SOA ecosystem, and those doing the testing must be familiar with both
4029 the characteristics and responses of the ecosystem and the tools, especially those available as services,
4030 to facilitate and standardize testing. Test professionals will know what level of assurance must be
4031 established as the exposure of the service to the ecosystem and ecosystem to the service increases
4032 towards operational status. These test professionals may be internal resources to an organization or may
4033 evolve as a separate discipline provided through external contracting.

4034 As noted above, it is unlikely that a complete duplicate of the SOA ecosystem will be available for isolated
4035 testing, and thus use of ecosystem [resources](#) will manifest as a transition process rather than a step
4036 change from a test environment to an operational one. This is especially true for new composite services
4037 that incorporate existing operational services to achieve the new functionality. The test professionals will
4038 need to understand the available resources and the ramifications of this transition.

4039 As with current software development, a stage beyond work by test professionals will make use of a
4040 select group of typical users (commonly referred to as beta testers) to report on service response during
4041 typical intended use. This establishes fitness by the consumers, providing final validation of previously
4042 verified processes, requirements, and final implementation.

4043 In traditional software development, beta testing is the end of testing for a given version of the software.
4044 However, although the initial test phase can establish an appropriate level of confidence consistent with
4045 the designed use for the initial target consumer community, the operational service will exist in an
4046 evolving ecosystem, and later conditions of use may differ from those thought to be sufficient during the
4047 initial testing. Thus, operational monitoring becomes an extension of testing through the service lifetime.
4048 This continuous testing will attempt to ensure that a service does not consume an inordinate amount of

4049 ecosystem resources or display other behavior that degrades the ecosystem, but it will not undercover
4050 functional errors that may surface over time.

4051 As with any software, it is the responsibility of the consumers to consider the reasonableness of solutions
4052 in order to spot errors in either the software or the way the software is being used. This is especially
4053 important for consumers with unanticipated uses that may go beyond the original test conditions. It is
4054 unlikely the consumers will initiate a new round of formal testing unless the new use requires a
4055 significantly higher level of confidence in the service. Rather the consumer becomes a new extension to
4056 the testing regiment. Obvious testing would include a sanity check of results during the new use.
4057 However, if the details of legacy testing are associated with the service through the service description
4058 and if testing resources are available through automated testing services, then the new consumers can
4059 rerun and extend previous testing to include the extended test conditions. If the test results are
4060 acceptable, these can be added to the documentation of previous results and become the extended basis
4061 for future decisions by prospective consumers on the appropriateness of the service. If the results are not
4062 acceptable or in some way questionable, the responsible party for the service or testing professionals can
4063 be brought in to decide if remedial action is necessary.

4064 **5.4.3.4 How Testing Results are Reported**

4065 For any SOA service, an accurate reporting of the testing a service has undergone and the results of the
4066 testing is vital to consumers deciding whether a service is appropriate for intended use. Appropriateness
4067 may be defined by a consumer organization and require specific test regiments culminating in a
4068 certification; appropriateness could be established by accepting testing and certifications that have been
4069 conferred by others.

4070 The testing and certification information should be identified in the service description. Referring to the
4071 general description model of Figure 14, tests conducted by or under a request from the service owner
4072 (see [ownership](#) in section 3.2.4) would be captured under Annotations from Owners. Testing done by
4073 others (such as consumers with unanticipated uses) could be associated through Annotations from 3rd
4074 Parties.

4075 Consumer testing and the reporting of results raise additional issues. While stating who did the testing is
4076 mandatory, there may be formal requirements for authentication of the tester to ensure traceability of the
4077 testing claims. In some circumstances, persons or organizations would not be allowed to state testing
4078 claims unless the tester was an approved entity. In other cases, ensuring the tester had a valid email
4079 may be sufficient. In either case, it would be at the discretion of the potential consumer to decide what
4080 level of authentication was acceptable and which testers are considered authoritative in the context of
4081 their anticipated use.

4082 Finally, in a world of openly shared information, we would see an ever-expanding set of testing
4083 information as new uses and new consumers interact with a service. In reality, these new uses may
4084 represent proprietary processes or classified use that should only be available to authorized parties.
4085 Testing information, as with other elements of description, may require special access controls to ensure
4086 appropriate access and use.

4087 **5.4.4 Testing SOA Services**

4088 Testing of SOA services should be consistent with the SOA paradigm. In particular, testing resources and
4089 artifacts should be visible in support of service interaction between providers and consumers, where here
4090 the interaction is between the testing resource and the tester. In addition, the idea of opacity of the
4091 implementation should limit the details that need to be available for effective use of the test resources.
4092 Testing that requires knowledge of the internal structure of the service or its underlying capability should
4093 be performed as part of unit testing in the development sandbox, and should represent a minimum level
4094 of confidence before the service begins its transition to further testing and eventual operation in the SOA
4095 ecosystem.

4096 **5.4.4.1 Progression of SOA Testing**

4097 Software testing is a gradual exercise going from micro inspection to testing macro effects. The first step
4098 in testing is likely the traditional code reviews. SOA considerations would account for the distributed

4099 nature of SOA, including issues of distributed security and best practices to ensure secure resources. It
4100 would also set the groundwork for opacity of implementation, hiding programming details and simplifying
4101 the use of the service.

4102 Code review is likely followed by unit testing in a development sandbox isolated from the operational
4103 environment. The unit testing is done with full knowledge of the service internal structure and knowledge
4104 of resources representing underlying capabilities. It tests the interface to ensure exchanged messages
4105 are as specified in the service description and the messages can be parsed and interpreted as intended.
4106 Unit testing also verifies intended functionality and that the software has dealt correctly with internal
4107 dependencies, such as structure of a file system or access to other dedicated resources.

4108 Some aspects of unit testing require external dependencies be satisfied, and this is often done using
4109 mock objects to substitute for the external resources. In particular, it will likely be necessary to include
4110 mocks of existing operational services, both those provided as part of the SOA infrastructure and services
4111 from other providers.

4112 **Service Mock**

4113 An entity that mimics some aspect of the performance of an operational service without
4114 committing to the real world effects that the operational service would produce.

4115 Mocks are discussed in detail in sections 5.4.4.3 and 5.4.4.4.

4116 After unit testing has demonstrated an adequate level of confidence in the service, the testing must
4117 transition from the tightly controlled environment of the development sandbox to an environment that
4118 more clearly resembles the operational SOA ecosystem or, at a minimum, the intended enterprise. While
4119 sandbox testing will use simple mocks of some aspects of the SOA environment, such as an interface to
4120 a security service without the security service functionality, the dynamic nature of SOA makes a full
4121 simulation infeasible to create or maintain. This is especially true when a new composite service makes
4122 use of operational services provided by others. Thus, at some point before testing is complete, the
4123 service will need to demonstrate its functionality by using resources and dealing with conditions that only
4124 exist in the full ecosystem or the intended enterprise. Some of these resources may still provide test
4125 interfaces -- more on this below -- but the interfaces will be accessible using the SOA environment and
4126 not just implemented for the sandbox.

4127 At this stage, the opacity of the service becomes important as the details of interacting with the service
4128 now rely on correct use of the service interface and not knowledge of the service internals. The workings
4129 of the service will only be observable through the real world effects realized through service interactions
4130 and external indications that conditions of use, such as user authentication, are satisfied. Monitoring the
4131 behavior of the service will depend on service interfaces that expose internal monitoring or provide
4132 required information to the SOA infrastructure monitoring function. The monitoring required to test a new
4133 service is likely to have significant overlap with the monitoring the SOA infrastructure includes to monitor
4134 its own health and to identify and isolate behavior outside of acceptable bounds. This is exactly what is
4135 needed as part of service testing, and it is reasonable to assume that the ecosystem transition includes
4136 use of operational monitoring rather than solely dedicated monitoring for each service being tested.

4137 Use of SOA monitoring resources during the explicit testing phase sets the stage for monitoring and a
4138 level of continual testing throughout the service lifetime.

4139 **5.4.4.2 Testing Traditional Dependencies vs. Service Interactions**

4140 A SOA service is not required to make use of other operational services beyond what may be required for
4141 monitoring by the ecosystem infrastructure. The service can implement hardcoded dependencies which
4142 have been tested in the development sandbox through the use of dedicated mocks. While coordination
4143 may be required with real data sources during integration testing, the dependencies can be constrained to
4144 things that can be tested in a more traditional manner. Policies can also be set to restrict access to pre-
4145 approved users, and thus the question of unanticipated users and unanticipated uses can be eliminated.
4146 Operational readiness can be defined in terms of what can be proven in isolated testing. While all this
4147 may provide more confidence in the service for its designed purpose, such a service will not fully
4148 participate in the benefits or challenges of the ecosystem. This is akin to filling a swimming pool with sea
4149 water and having someone in the pool say they are swimming in the ocean.

4150 In considering the testing needed for a fully participating service, consider the example of a new
4151 composite service that combines the real world effects and complies with the conditions of use of five
4152 existing operational services. The developer of the composite service does not own any of the component
4153 services and has limited, if any, ability to get the distributed owners to do any customization. The
4154 developer also is limited by the principle of opacity to information comprising the service description, and
4155 does not know internal details of the component services. The developer of the composite service must
4156 use the component services as they exist as part of the SOA environment, including what is provided to
4157 support testing by new users. This introduces requirements for what is needed in the way of service
4158 mocks.

4159 **5.4.4.3 Use of Service Mocks**

4160 Service mocks enable the tested service to respond to specific features of an operational service that is
4161 being used as a component. It allows service testing to proceed without needing access to or with only
4162 limited engagement with the component service. Mocks can also mimic difficult to create situations for
4163 which it is desired to test the new service response. For composite services using multiple component
4164 services, mocks may be used in combination to function for any number of the components. Note, when
4165 using service mocks, it is important to remember that it is not the component service that is being tested
4166 (although anomalous behavior may be uncovered during testing) but the use of the component in the new
4167 composite.

4168 Individual service mocks can emphasize different features of the component service they represent but
4169 any given mock does not have to mimic all features. For example, a mock of the service interface can
4170 echo a sent message and demonstrate the message is reaching its intended destination. A mock could
4171 go further and parse the sent message to demonstrate the message not only reached its destination but
4172 was understood. As a final step, the mock could report back what actions would have been taken by the
4173 component service and what real world effects would result. If the response mimicked the operational
4174 response, functional testing could proceed as if the real world effect actually occurred.

4175 There are numerous ways to provide mock functionality. The service mock could be a simulation of the
4176 operational service and return simulated results in a realistic response message or event notification. It is
4177 also possible for the operational service to act as its own mock and simply not execute the commit stage
4178 of its functionality. The service mock could use a combination of simulation and service action without
4179 commit to generate a report of what would have occurred during the defined interaction with the
4180 operational service.

4181 As the service proceeds through testing, mocks should be systematically replaced by the component
4182 resources accessed through their operational interfaces. Before beta testing begins, end-to-end testing,
4183 i.e. proceeding from the beginning of the service interaction to the resulting real world results, should be
4184 accomplished using component resources via their operational interfaces.

4185 **5.4.4.4 Providers of Service Mocks**

4186 In traditional testing, it is often the test professionals who design and develop the mocks, but in the
4187 distributed world of SOA, this may not be efficient or desirable.

4188 In the development sandbox, it is likely the new service developer or test professionals working with the
4189 developer will create mocks adequate for unit testing. Given that most of this testing is to verify the new
4190 service is performing as designed, it is not necessary to have high fidelity models of other resources
4191 being accessed. In addition, given opacity of SOA implementation, the developer of the new service may
4192 not have sufficient detailed knowledge of a component service to build a detailed mock of the component
4193 service functionality. Sharing existing mocks at this stage may be possible but the mocks would need to
4194 be implemented in the sandbox, and for simple models it is likely easier to build the mock from scratch.

4195 As testing begins its transition to the wider SOA environment, mocks may be available as services. For
4196 existing resources, it is possible that an Open Source model could evolve where service mocks of
4197 available functions can be catalogued and used during initial interaction of the tested service and the
4198 operational environment. Widely used functions may have numerous service mocks, some mimicking
4199 detailed conditions within the SOA infrastructure. However, the Open Source model is less likely to be
4200 sufficient for specialty services that are not widely used by a large consumer community.

4201 The service developer is probably best qualified for also developing more detailed service mocks or for
4202 mock modes of operational services. This implies that in addition to their operational interfaces, services
4203 will routinely provide test interfaces to enable service mocks to be used as services. As noted above, a
4204 new service developer wanting to build a mock of component services is limited to the description
4205 provided by the component service developer or owner. The description typically will detail real world
4206 effects and conditions of use but will not provide implementation details, some of which may be
4207 proprietary. Just as important in the SOA ecosystem, if it becomes standard protocol for developers to
4208 create service mocks of their own services, a new service developer is only responsible for building his
4209 own mocks and can expect other mocks to be available from other developers. This reduces duplication
4210 of effort where multiple developers would be trying to build the same mocks from the same insufficient
4211 information. Finally, a service developer is probably best qualified to know when and how a service mock
4212 should be updated to reflect modified functionality or message exchange.

4213 It is also possible that testing organizations will evolve to provide high-fidelity test harnesses for new
4214 services. The harnesses would allow new services to plug into a test environment and would facilitate
4215 accessing mocks of component services. However, it will remain a constant challenge for such
4216 organizations to capture evolving uses and characteristics of service interactions in the real SOA
4217 environment and maintain the fidelity and accuracy of the test systems.

4218 5.4.4.5 Fundamental Questions for SOA Testing

4219 In order for the transition to the SOA operational environment to proceed, it is necessary to answer two
4220 fundamental questions:

- 4221 • Who provides what testing resources for the SOA operational environment, e.g. mocks of
4222 interfaces, mocks of functionality, monitoring tools?
- 4223 • What testing needs to be accomplished before operational environment resources can be
4224 accessed for further testing?

4225 The discussion in section 5.4.4.4 notes various levels of sophistication of service mocks and different
4226 communities are likely to be responsible for different levels. Section 5.4.4.4 advocates a significant role
4227 for service developers, but there needs to be community consensus that such mocks are needed and that
4228 service developers will agree to fulfill this role. There is also a need for consensus as to what tools should
4229 be available as services from the SOA infrastructure.

4230 As for use of the service mocks and SOA environment monitoring services, practical experience is
4231 needed upon which guidelines can be established for when a new service has been adequately tested to
4232 proceed with a greater level of exposure with the SOA environment. Malfunctioning services could cause
4233 serious problems if they cannot be identified and isolated. On the other hand, without adequate testing
4234 under SOA operational conditions, it is unlikely that problems can be uncovered and corrected before
4235 they reach an operational stage.

4236 As noted in section 5.4.4.2, some of these questions can be avoided by restricting services to more
4237 traditional use scenarios. However, such restriction will limit the effectiveness of SOA use and the result
4238 will resemble the constraints of traditional integration activities we are trying to move beyond.

4239 5.4.5 Architectural Implications for SOA Testing

4240 The discussion of SOA Testing indicates numerous architectural implications that **MUST** be considered
4241 for testing of resources and interactions within the SOA ecosystem:

- 4242 • SOA services **MUST** be testable in the environment and under the conditions that can be
4243 encountered in the operational SOA ecosystem.
- 4244 • The distributed, boundary-less nature of the SOA ecosystem makes it infeasible to create and
4245 maintain a single testing substitute of the entire ecosystem to support testing activities. Test
4246 protocols **MUST** recognize and accommodate for changes to and activities within the ecosystem.
- 4247 • A standard suite of monitoring services **SHOULD** be defined, developed, and maintained. This
4248 **SHOULD** be done in a manner consistent with the evolving nature of the ecosystem.
- 4249 • Services **SHOULD** provide interfaces that support access in a test mode.
- 4250 • Testing resources **MUST** be described and their descriptions **MUST** be catalogued in a manner
4251 that enables their discovery and access.

Comment [KJL122]: Issue rb45

- 4252 • Guidelines for testing and ecosystem access **MUST** be established and the ecosystem **MUST** be
- 4253 able to enforce those guidelines asserted as policies.
- 4254 • Services **SHOULD** be available to support automated testing and regression testing.
- 4255 • Services **SHOULD** be available to facilitate updating service description by authorized
- 4256 participants who has performed testing of a service.

4257 6 Conformance

4258 This Reference Architecture Foundation is an abstract architectural description of Service Oriented
4259 Architecture, which means that it is especially difficult to construct conformance tests ~~for conformance to~~
4260 ~~the architecture~~. In addition, conformance to an abstract architectural specification does not, by itself,
4261 guarantee any form of interoperability between multiple implementations.

4262 However, it is possible to decide whether or not a given architecture is conformant to an architectural
4263 description such as this one. In discussions of conformance we use the term **target architecture** to
4264 identify the (typically concrete) architecture that may be viewable-considered as conforming to the
4265 abstract principles outlined in this document.

4266 Target Architecture

4267 An architectural description of a system that is intended to be viewed as conforming to the SOA-
4268 RAF.

4269 While we cannot guarantee interoperability between target architectures (or more specifically between
4270 applications and systems residing within the ecosystems of those target architectures), the likelihood of
4271 interoperability between target architectures is promoted-increased by conformance to this Reference
4272 Architecture Framework as it ~~reduces the semantic impedance mismatch facilitates semantic engagement~~
4273 between the different ecosystems.

4274 ~~The primary measure of conformance is whether given cC~~ concepts as described in the RAF document
4275 ~~have corresponding concepts SHOULD be expressed and used~~ in the target architecture. If used, sSuch
4276 ~~expression a correspondence MUST honor-reflect~~ the relationships identified within this document ~~for the~~
4277 ~~target architecture to be considered conforming.~~

4278 ~~For example, in Section 3.2.4.1 we identify resource as a key concept. A resource is associated with an~~
4279 ~~owner and a number of identifiers. For a target architecture to conform to the SOA-RAF, it must be~~
4280 ~~possible to find corresponding concepts of resource, identifier and owner within the target architecture:~~
4281 ~~say entity, token and user. Furthermore, the relationships between entity, token and user MUST mirror~~
4282 ~~the relationships between resource, identifier and owner appropriately.~~

4283 Terminology within the target architecture **SHOULD** be identical to that in the RAF and the terms used
4284 refer to the same concepts; and any the 'graph' of concepts and relationships used is MUST be consistent
4285 with the RAF, that is all that is required for the target architecture to conform to this Reference
4286 ~~Architecture Framework.~~

4287 In addition, conformance can be measured in terms of how the various architectural implications
4288 described in the RAF have been addressed in the target architecture. These implications are covered at
4289 the end of the main sections above (sections 3.4, 4.1.4, 4.2.3, 4.3.6, 4.4.3, 5.1.4, 5.2.5, 5.3.7, and 5.4.5).
4290 Many of these sections contain formal conformance requirements ("MAY", "MUST", "SHOULD") in
4291 accordance with [RFC 2119] and the target architecture MUST take account of these implications where
4292 specified in those sections.

4293

A. Acknowledgements

4294 The following individuals have participated in the work of the technical committee responsible for creation
4295 of this specification and are gratefully acknowledged:

4296 **Participants:**

4297 Chris Bashioum, MITRE Corporation
4298 Rex Brooks, Individual Member
4299 Peter F Brown, Individual Member
4300 Scott Came, Search Group Inc.
4301 Joseph Chiusano, Booz Allen Hamilton
4302 Robert Ellinger, Northrop Grumman Corporation
4303 David Ellis, Sandia National Laboratories
4304 Jeff A. Estefan, Jet Propulsion Laboratory
4305 Don Flinn, Individual Member
4306 Anil John, Johns Hopkins University
4307 Ken Laskey, MITRE Corporation
4308 Boris Lublinsky, Nokia Corporation
4309 Francis G. McCabe, Individual Member
4310 Christopher McDaniels, USSTRATCOM
4311 Tom Merkle, Lockheed Martin Corporation
4312 Jyoti Namjoshi, Patni Computer Systems Ltd.
4313 Duane Nickull, Adobe Inc.
4314 James Odell, Associate
4315 Michael Poulin, Fidelity Investments
4316 Michael Stiefel, Associate
4317 Danny Thornton, Northrop Grumman
4318 Timothy Vibbert, Lockheed Martin Corporation
4319 Robert Vitello, New York Dept. of Labor

4320 The committee would particularly like to underline the significant writing and conceptualization
4321 contributions made by Chris Bashioum, Rex Brooks, Peter Brown, Dave Ellis, Jeff Estefan, Ken Laskey,
4322 Boris Lublinsky, Frank McCabe, Michael Poulin and Danny Thornton

4323

B. Index of Defined Terms

4324

Action.....	37	Policy Conflict.....	76
Action Level Real World Effect.....	48	Policy Conflict Resolution.....	76
Actor.....	25	Policy Constraint.....	75
Authority.....	26	Policy Decision.....	75
Business Collaboration.....	68	Policy Enforcement.....	75
Business functionality.....	29	Policy Framework.....	74
Business Process.....	68	Policy Object.....	75
Business solution.....	36	Policy Ontology.....	74
Capability.....	29	Policy Owner.....	75
Communication.....	34	Policy Subject.....	75
Composability.....	36	Presence.....	62
Constitution.....	24	Private State.....	39
Consumer.....	27	Protocol.....	62
Contract.....	34	Provider.....	27
Delegate.....	25	Real World Effect.....	29
Endpoint.....	62	Regulation.....	81
Governance.....	79	Requirement.....	29
Governance Framework.....	80	Resource.....	30
Governance Processes.....	80	Responsibility.....	26
Identifier.....	30	Right.....	26
Joint Action.....	38	Risk.....	32
Leadership.....	80	Rule.....	81
Logical Framework.....	74	Security.....	88
Manageability.....	96	Semantic Engagement.....	35
Manageability property.....	96	Service Contract.....	99
Mediator.....	27	Service Level Real World Effect.....	48
Message Exchange.....	65	Service Mock.....	110
Need.....	29	Shared State.....	39
Non-Participant.....	25	SOA Ecosystem.....	21
Obligation.....	27	SOA-based System.....	21
Operations.....	65	Social Structure.....	23
Owner.....	27	Stakeholder.....	24
Ownership.....	31	State.....	39
Ownership Boundary.....	31	Target Architecture.....	114
Participant.....	25	Trust.....	32
Permission.....	27	Willingness.....	32
Policy.....	33		

4325

4326 C. Critical Factors Analysis

4327 A critical factors analysis (CFA) is an analysis of the key properties of a project. A CFA is analyzed in
4328 terms of the goals of the project, the critical factors that will lead to its success and the measurable
4329 requirements of the project implementation that support the goals of the project. CFA is particularly
4330 suitable for capturing quality attributes of a project, often referred to as 'non-functional' or 'other-than-
4331 functional' requirements: for example, security, scalability, wide-spread adoption, and so on. As such,
4332 CFA complements rather than attempts to replace other requirements capture techniques.

4333 C.1 Goals

4334 A goal is an overall target that you are trying to reach with the project. Typically, goals are hard to
4335 measure by themselves. Goals are often directed at the potential consumer of the product rather than the
4336 technology developer.

4337 C.2 Critical Success Factors

4338 A critical success factor (CSF) is a property, sub-goal that directly supports a goal and there is strong
4339 belief that without it the goal is unattainable. CSFs themselves are not necessarily measurable in
4340 themselves.

4341 C.3 Requirements

4342 A requirement is a specific measurable property that directly supports a CSF. The key here is
4343 measurability: it should be possible to unambiguously determine if a requirement has been met. While
4344 goals are typically directed at consumers of the specification, requirements are focused on technical
4345 aspects of the specification.

4346

D. Relationship to other SOA Open Standards

4347 [Numerous efforts have been working in the space of defining standards for SOA and its applications. The](#)
4348 [OASIS SOA-RM Technical Committee and its SOA-RA Technical Subcommittee has established](#)
4349 [communications with several of these efforts in an attempt to coordinate and facilitate among the efforts.](#)
4350 [This appendix notes some of these efforts.](#)

4351 [D.1 Navigating the SOA Open Standards Landscape Around Architecture](#)

4352 The white paper *Navigating the SOA Open Standards Landscape Around Architecture* issued jointly by
4353 OASIS, OMG, and The Open Group [**SOA-NAV**] was written to help the SOA community at large
4354 navigate the myriad of overlapping technical products produced by these organizations with specific
4355 emphasis on the 'A' in SOA, i.e., Architecture.

4356 The white paper explains and positions standards for SOA reference models, ontologies, reference
4357 architectures, maturity models, modeling languages, and standards work on SOA governance. It outlines
4358 where the works are similar, highlights the strengths of each body of work, and touches on how the work
4359 can be used together in complementary ways. It is also meant as a guide to users for selecting those
4360 specifications most appropriate for their needs.

4361 While the understanding of SOA and SOA Governance concepts provided by these works is similar, the
4362 evolving standards are written from different perspectives. Each specification supports a similar range of
4363 opportunity, but has provided different depths of detail for the perspectives on which they focus. Although
4364 the definitions and expressions may differ, there is agreement on the fundamental concepts of SOA and
4365 SOA Governance.

4366 The following is a summary taken from [**SOA-NAV**] of the positioning and guidance on the specifications:

- 4367 • The OASIS Reference Model for SOA (SOA RM) is the most abstract of the specifications
4368 positioned. It is used for understanding core SOA concepts
- 4369 • The Open Group SOA Ontology extends, refines, and formalizes some of the core concepts of
4370 the SOA RM. It is used for understanding core SOA concepts and facilitates a model-driven
4371 approach to SOA development.
- 4372 • The OASIS Reference Architecture Foundation for SOA (this document) is an abstract,
4373 foundational reference architecture addressing a broader ecosystem viewpoint for building and
4374 interacting within the SOA paradigm. It is used for understanding different elements of SOA, the
4375 completeness of SOA architectures and implementations, and considerations for reaching across
4376 ownership boundaries where there is no single authoritative entity for SOA and SOA governance.
- 4377 • The Open Group SOA Reference Architecture is a layered architecture from consumer and
4378 provider perspective with cross cutting concerns describing these architectural building blocks
4379 and principles that support the realizations of SOA. It is used for understanding the different
4380 elements of SOA, deployment of SOA in enterprise, basis for an industry or organizational
4381 reference architecture, implication of architectural decisions, and positioning of vendor products in
4382 a SOA context.
- 4383 • The Open Group SOA Governance Framework is a governance domain reference model and
4384 method. It is for understanding SOA governance in organizations. The OASIS Reference
4385 Architecture for SOA Foundation contains an abstract discussion of governance principles as
4386 applied to SOA across boundaries
- 4387 • The Open Group SOA Integration Maturity Model (OSIMM) is a means to assess an
4388 organization's maturity within a broad SOA spectrum and define a roadmap for incremental
4389 adoption. It is used for understanding the level of SOA maturity in an organization
- 4390 • The Object Management Group SoaML Specification supports services modeling UML
4391 extensions. It can be seen as an instantiation of a subset of the Open Group RA used for
4392 representing SOA artifacts in UML.

4393 Fortunately, there is a great deal of agreement on the foundational core concepts across the many
4394 independent specifications and standards for SOA. This could be best explained by broad and common
4395 experience of users of SOA and its maturity in the marketplace. It also provides assurance that investing

4396 in SOA-based business and IT transformation initiatives that incorporate and use these specifications and
 4397 standards helps to mitigate risks that might compromise a successful SOA solution.




Figure 45 - SOA Reference Architecture Positioning (from 'Navigating the SOA Open Standards Landscape Around Architecture', © OASIS, OMG, The Open Group)

4398
 4399
 4400

Comment [PFB127]: Issue 298

4401 **D.2 The Service-Aware Interoperability Framework: Canonical**

4402 Readers of the RAF are strongly encouraged to review a document recently published by the Health
 4403 Level Seven (HL7) Architecture Board (ArB) entitled *The Service-Aware Interoperability Framework:*
 4404 *Canonical*. The document was developed over the past four years, and represents a substantive,
 4405 industry-specific (i.e. the large but vertical healthcare industry) effort to surface, define, and discuss in
 4406 detail various aspects of a number of critical success factors involved in implementing large-scale (i.e.
 4407 enterprises-level) architectures with a focus on achieving both intra- and inter-enterprise technical
 4408 interoperability irrespective of the particular exchange mechanism involved, e.g. service interface,
 4409 messages, or structure documents.

4410 In addition to providing an independent validation for the both the general focus as well as some of the
 4411 concrete specifics of the RAF (especially those involving the importance of governance in achieving
 4412 large-scale interoperability), the HL7 document underscores several important aspects of the RAF
 4413 including:

- 4414 1. A validation of one of the RAF's primary claims, i.e. the need to specifically focus on intra- and inter-
 4415 enterprise interoperability as a first-class citizen in any enterprise (or cross-enterprise) architecture
 4416 discussion irrespective of the particular choice of enterprise architecture approach, framework, or
 4417 implementation technology, e.g. TOGAF, Zachman, ODP, SOA, etc. In addition, the HL7 document
 4418 clearly articulates – as the RAF does as well – the difficulties involved in achieving that focus in such
 4419 a manner that it can be manifest in operationally effective and manageable processes and
 4420 deliverables.

- 4421 2. An agreement as to the critical importance of governance as the root of any successful effort to
4422 implement large-scale, cross-boundary interoperability aimed at achieving a collective shared
4423 purpose-mission or goal. In particular, both documents share the notion that 'technical-level'
4424 governance – e.g. service – or message-level technical interchange specifications – must itself be a
4425 manifestation of a higher-level, cross-jurisdictional agreement on desired goals, responsibilities,
4426 accountabilities, and deliverables.
- 4427 3. A validation of the importance of core SOA constructs as constructs useful in expressing many of the
4428 central aspects of interoperability irrespective of whether a particular interoperability scenario is
4429 actually 'realized' using SOA-compatible technologies. (NOTE: Although it might at first appear that
4430 the OASIS document is more 'service-focused' than the 'service-aware' document from HL7, there
4431 are considerably more similarities than differences in these slightly different foci secondary to the fact
4432 that both documents are intent on describing principles and framework concepts rather than delving
4433 into technical details. There are, however, certain instances where content of the OASIS document
4434 would be likely to find its analogue in SAIF Implementation Guides rather than in the SAIF Canonical
4435 Definition document.)
- 4436 4. The need for specific, explicit statements of those aspects of a given component that affects its ability
4437 to participate in a reliable, predictable manner in a variety of interoperability scenarios. In particular,
4438 component characteristics must be explicitly expressed in both design-time and run-time contexts as
4439 implicit assumptions are the root of most failures to achieve successfully cross-boundary
4440 interoperability irrespective of the chosen technical details of a particular interoperability instance.

4441 In summary, although the two documents are clearly not identical in their specifics, e.g. there are
4442 differences in the language used to name various concepts, constructs, and relationships; there are some
4443 differences in levels of abstraction regarding certain topics, etc.; and although the OASIS RAF is more
4444 directly focused on services as a final implementation architecture than the HL7 SAIF CD, the
4445 commonalities of purpose, content, and approach present in the two documents – documents which were
4446 developed by each organization without any knowledge of the others' work in what clearly are areas of
4447 common interest and concern – far outweighs their differences. As such, the HL7 ArB and the OASIS
4448 RAF Task Force have agreed to work together going forward to obtain the highest degree of alignment
4449 and harmonization possible between the two documents including the possible development of a joint
4450 document under the auspices of one of the ISO software engineering threads.

4451 The current version of the HL7 document – as well as all future versions – is available at:

4452 <http://www.hl7.org/permalink/?SAIFCDR1PUBLIC>

4453 **D.3 IEEE Reference Architecture**

Comment [PFB128]: Issue 298, part

4454 As the RAF has been finalized, a new initiative has appeared from the Institute of Electrical and
4455 Electronics Engineers (IEEE) to develop a SOA Reference Architecture. Encouragingly, the working
4456 group established decided not to start from scratch but instead take account of existing work. Its initial
4457 phase of work is currently ongoing (Summer 2012) and is concentrating on assessing both the current
4458 RAF and The Open Group's SOA Reference Architecture. The desire at this stage is to endorse these
4459 two works rather than creating a new one.

4460 **D.4 RM-ODP**

Comment [PFB129]: Issue 298, part

4461 The Reference Model for Open Distributed Processing (the RM-ODP) is an international standard
4462 developed by the ISO and ITU-T standardization organizations [ISO/IEC 10746]. It provides a set of
4463 concepts and structuring rules for describing and building open distributed systems, structured in terms of
4464 five viewpoints, representing concerns of different stakeholders.

4465 The enterprise viewpoint is concerned with understanding and defining organizational context in which
4466 a distributed system is to be built and operated. The main concept here is that of a 'Community', defining
4467 how some set of participants should behave in order to achieve a specific objective. A community is
4468 specified in term of roles played by 'Enterprise Objects' in that community; processes and interactions in
4469 which they are involved; and enterprise policies (obligations, permissions, prohibitions, authorizations)
4470 that constrain the behavior of enterprise objects in these roles. A community also specifies business
4471 services, which are particular descriptions of behavior expressing functionality or capability provided by

4472 one party to others who can use the service to satisfy their own business needs, resulting in an added
4473 value to them.

4474 A special type of enterprise object is 'Party', modeling a natural person or any other entity considered to
4475 have some of the rights, powers and duties of a natural person. Party is one of the accountability
4476 concepts defined in RM-ODP to support the traceability of obligations in the enterprise specification. Note
4477 that a typical enterprise specification will consist of several overlapping and interacting communities that
4478 make up the enterprise. The enterprise language defines a number of kinds of action developed to define
4479 and analyze how responsibilities evolve. These include: 'Commitment', being an action that results in
4480 object performing it to undertake some obligation; 'Delegation', being an action (performed by a principal)
4481 that assigns authority, responsibility or a function to another object (agent); 'Prescription', an action that
4482 establishes a rule; and 'Evaluation', an action that assesses the value of something.

4483 The **information viewpoint** focuses on modeling information within the enterprise of interest. Key
4484 concepts are 'Information Objects' and 'static, dynamic and invariant Schemas'.

4485 The **computational viewpoint** is concerned with describing basic functionality of the processes and
4486 applications supporting the enterprise activities. The key concepts are 'Computational Objects', the
4487 'Interactions' between the computation objects (which can be of several types, namely 'Operations',
4488 'Streams' or 'Signals') which occur at their 'Interfaces'. Further, the 'Environment Contracts' specify non-
4489 functional properties of computational objects and their interactions. The computational specification is
4490 independent of any technology platform and can be implemented using multiple such platforms.

4491 The **engineering viewpoint** is concerned with the mechanisms needed to support distributed interactions
4492 between objects. Many such mechanisms are provided as part of middleware specifications such as Web
4493 Services, .NET, JEE or CORBA.

4494 The **technology viewpoint** is concerned with specifying real-world constrains such as available
4495 hardware components or application platforms with which a new system is to be integrated, specification
4496 of hardware and software components from which the new system is built and selection of implementable
4497 standards.

4498 In order to express relationships between these separately developed specifications, the RM-ODP
4499 standard defines the concept of 'Correspondences' between viewpoints. An ODP specification of a
4500 system thus consists of five viewpoint specifications and correspondences specification that links these
4501 separate specifications together.

4502 The RM-ODP standard also provides a well-defined **conformance framework**, providing links between
4503 specifications and implementations and thus supporting testing.

4504 The ODP viewpoint languages are defined in abstract way and can be supported by several notations.
4505 The use of UML notation in expressing ODP viewpoint languages is defined in a separate ISO standard,
4506 *Use of UML for ODP system specification* ('UML4ODP' for short) [ISO/IEC IS 19793].

4507 From an architectural point of view, there is no significant difference between service-oriented
4508 architectures (SOA) and the architectural framework defined in ODP; some argue that current service-
4509 oriented approaches can be understood as a subset of the more general ODP approach [LININGTON]. In
4510 fact many of the concepts and principles in this SOA-RAF have significant alignment with the RM-ODP
4511 concepts, in particular those covered in the enterprise and computational viewpoints.