



Reference Architecture Foundation for Service Oriented Architecture Version 1.0

Comment [PFB1]: Still to be done:
- Check all diagrams for consistency of models and terms;
- consistent **bolding** of formal terms (only first occurrence?) and hyperlinks from *first* mention of a term to its definition in the text;

Starting Point: Draft 04 21 Dec 2010

This Interim Draft PFB, 06-01-2011

Specification URIs:

This Version:

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

Previous Version:

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

Latest Version:

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

Technical Committee:

OASIS Service Oriented Architecture Reference Model TC

Chair(s):

Ken Laskey, MITRE Corporation

Editor(s):

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

Related work:

This specification is related to:

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

Abstract:

This document specifies the OASIS Reference Architecture 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 ANSI¹/IEEE² 1471-2000, (now ISO³/IEC⁴ 42010-2007) Standard. The SOA-RAF is of value to Enterprise Architects, Business and IT Architects as well as CIOs and other senior executives involved in strategic business and IT planning.

The SOA-RAF has three main views: the ~~Ecosystem View~~ Participation in a SOA Ecosystem view which focuses on the way that participants are part of a Service Oriented Architecture ecosystem; the ~~Realizing Services~~ Realization of a SOA Ecosystem view which addresses the requirements for constructing a SOA-based system in a SOA ecosystem ~~Service Oriented Architecture~~; and the ~~Owning Service Oriented Architecture~~ Ownership in a SOA Ecosystem view which focuses on the governance and management of what is meant to own a SOA-based systems.

Status:

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

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

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

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

¹ American National Standards Institute

² Institute of Electrical and Electronics Engineers

³ International Organization for Standardization

⁴ International Electrotechnical Commission

Notices

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

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

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

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

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

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

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

OASIS takes no position regarding the validity or scope of any intellectual property or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; neither does it represent that it has made any effort to identify any such rights. Information on OASIS' procedures with respect to rights in any document or deliverable produced by an OASIS Technical Committee can be found on the OASIS website. Copies of claims of rights made available for publication and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary

rights by implementers or users of this OASIS Committee Specification or OASIS Standard, can be obtained from the OASIS TC Administrator. OASIS makes no representation that any information or list of intellectual property rights will at any time be complete, or that any claims in such list are, in fact, Essential Claims.

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

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

Table of Contents

1	Introduction	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?	12
1.1.3	Relationship to the OASIS Reference Model for SOA.....	12
1.1.4	Relationship to other Reference Architectures	13
1.1.5	Relationship to other SOA Open Standards.....	13
1.1.6	Expectations set by the Reference Architecture Foundation.....	15
1.2	Service Oriented Architecture – An Ecosystems Perspective.....	15
1.3	Viewpoints, Views and Models	16
1.3.1	ANSI/IEEE 1471-2000::ISO/IEC 42010-2007	16
1.3.2	UML Modeling Notation.....	18
1.4	SOA-RAF Viewpoints	18
1.4.1	<i>Participation in a SOA ecosystem</i> viewpoint	19
1.4.2	<i>Realization of a SOA ecosystem</i> viewpoint.....	20
1.4.3	Process Viewpoint.....	20
1.5	Terminology	20
1.5.1	Usage of Terms.....	20
1.6	References	21
1.6.1	Normative References	21
1.6.2	Non-Normative References.....	21
2	Architectural Goals and Principles	23
2.1	Goals and Critical Success Factors of the Reference Architecture Foundation...	23
2.1.1	Goals.....	24
2.1.2	Critical Success Factors.....	25
2.2	Principles of this Reference Architecture Foundation	26
3	<i>Participation in a SOA ecosystem</i> view.....	29
3.1	SOA Ecosystem Model.....	31
3.1.1	Participants, Actors and Delegates	32
3.1.2	Roles in Social Structures	34
3.1.3	Shared State	38
3.1.4	Policies and Contracts	39
3.1.5	Resource and Ownership.....	41
3.2	Action in a SOA Ecosystem Model	43
3.2.1	Action and Joint Action.....	4443
3.2.2	Needs and Capabilities	4645
3.2.3	Services Reflecting Business.....	4746

3.2.4 Trust and Risk	4847
3.2.5 Using Joint Action for Communication	4948
3.2.6 Semantics and Semantic Engagement	50
3.2.7 Using Joint Action for Real World Effect	51
3.3 Architectural Implications.....	5351
3.3.1 Social structures.....	5351
3.3.2 The Importance of Action	5351
3.3.3 Communications as a Means of Mediating Action.....	5452
3.3.4 Semantics	5452
3.3.5 Trust and Risk	5452
3.3.6 Policies and Contracts	5553
4 <i>Realization of a SOA ecosystem view</i>	5654
4.1 Service Description Model	5654
4.1.1 The Model for Service Description	5856
4.1.2 Use Of Service Description	6866
4.1.3 Relationship to Other Description Models	7572
4.1.4 Architectural Implications	7573
4.2 Service Visibility Model	7775
4.2.1 Visibility to Business.....	7775
4.2.2 Visibility	7977
4.2.3 Architectural Implications	8482
4.3 Interacting with Services Model	8583
4.3.1 Interaction Dependencies.....	8583
4.3.2 Actions and Events	8684
4.3.3 Message Exchange.....	8785
4.3.4 Composition of Services	9088
4.3.5 Architectural Implications of Interacting with Services.....	9492
4.4 Policies and Contracts Model	9694
4.4.1 Policy and Contract Representation.....	9694
4.4.2 Policy and Contract Enforcement.....	9896
4.4.3 Architectural Implications	10097
5 <i>Ownership in a SOA ecosystem View</i>	10199
5.1 Governance Model	10199
5.1.1 Understanding Governance	102100
5.1.2 A Generic Model for Governance.....	103101
5.1.3 Governance Applied to SOA	108106
5.1.4 Architectural Implications of SOA Governance	114112
5.2 Security Model	115113
5.2.1 Secure Interaction Concepts	116114

5.2.2 Where SOA Security is Different	119117
5.2.3 Security Threats	119117
5.2.4 Security Responses	120118
5.2.5 Architectural Implications of SOA Security	123121
5.3 Management Model	123121
5.3.1 Management and Governance.....	127125
5.3.2 Management Contracts and Policies.....	127125
5.3.3 Management Infrastructure	127125
5.3.4 Service Life-cycle	128126
5.4 SOA Testing Model.....	128126
5.4.1 Traditional Software Testing as Basis for SOA Testing.....	128126
5.4.2 Testing and the SOA Ecosystem	130128
5.4.3 Elements of SOA Testing	131129
5.4.4 Testing SOA Services	138135
5.4.5 Architectural Implications for SOA Testing	142140
6 Conformance.....	143141
A. Acknowledgements	144142
B. Glossary of Terms.....	145143
C. The Unified Modeling Language, UML.....	148151
D. Critical Factors Analysis.....	149152
D.1 Goals.....	149152
Critical Success Factors.....	149152
Requirements.....	149152
CFA Diagrams.....	149152

Table of Figures

Figure 1- SOA Open Standards Work	15
Figure 2 Critical Factors Analysis of the Reference Architecture	24
Figure 3 Model elements described in the Participation in a SOA ecosystem view	30
Figure 4 Social Structure	31
Figure 5 Actors, Participants and Delegates	33
Figure 7 Service Roles	37
Figure 8 Policies and Contracts	40
Figure 9 Resources	41
Figure 10 Resource Ownership	4342
Figure 11 Actions, Real World Effect and Events	44
Figure 12 Trusting Actor and Willingness	4847
Figure 13 Communication as Joint Action	49
Figure 14 Model Elements Described in the Realization of a SOA ecosystem view ..	5654
Figure 15 General Description	5957
Figure 16 Representation of a Description	6159
Figure 17 Service Description	6364
Figure 18 Service Interface	6462
Figure 19 Service Functionality	6563
Figure 20 Model for Policies and Contracts as related to Service Participants	6664
Figure 21 Action-Level and Service-Level Policies	6765
Figure 22 Policies and Contracts, Metrics, and Compliance Records	6866
Figure 23 Relationship Between Action and Service Description Components	6967
Figure 24 Execution Context	7374
Figure 25 Interaction Description	7472
Figure 26 Visibility to Business	7876
Figure 27 Mediated Service Awareness	8078
Figure 28 Awareness In a SOA Ecosystem	8280
Figure 29 Business, Description and Willingness	8384
Figure 30 Service Reachability	8482
Figure 18 Interaction dependencies	8684
Figure 19 A "message" conveys either an action or an event	8684
Figure 20 Fundamental SOA message exchange patterns (MEPs)	8886
Figure 21 Simple model of service composition	9088
Figure 22 Abstract example of orchestration of service-oriented business process ..	9290

Figure 23 Abstract example of choreography of service-oriented business collaboration.	9492
Figure 37 Policies and Contracts	9795
Figure 38 Model Elements Described in the Ownership in a SOA ecosystem View	10199
Figure 39 Motivating governance model	104102
Figure 40 Setting up governance model	105103
Figure 41 Carrying out governance model	106104
Figure 42 Ensuring governance compliance model	107105
Figure 43 Relationship among types of governance	110108
Figure 44 Confidentiality and Integrity	117115
Figure 45 Authentication	117115
Figure 46 Authorization	118116
Figure 47 Managing resources in a SOA	124122
Figure 48 Example UML class diagram—Resources.	148151

1 Introduction

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

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

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

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

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

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

1.1.2 What is this Reference Architecture?

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

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

- [Use of](#) resources that are distributed across ownership boundaries;
- people and systems interacting with each other, also across ownership boundaries;
- security, management and governance that are similarly distributed across ownership boundaries; and
- interaction between people and systems that is primarily through the exchange of messages with reliability that is appropriate for the intended uses and purposes.

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

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

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

1.1.3 Relationship to the OASIS Reference Model for SOA

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

The SOA-RAF goes further in that it shows how SOA-based systems can be realized – albeit in an abstract way. As noted above, SOA-based systems are better thought of as dynamic systems rather than stand-alone software products. Consequently, how they

79 are used and managed is at least as important architecturally as how they are
80 constructed.

81 **1.1.4 Relationship to other Reference Architectures**

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

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

91 In addition, this Reference Architecture Foundation, as the title illustrates, is intended to
92 provide foundational models on which to build other reference architectures and
93 eventual concrete architectures. The relationship to other industry reference
94 architectures for SOA and related SOA open standards is described below.

95 **1.1.5 Relationship to other SOA Open Standards**

96 The white paper “Navigating the SOA Open Standards Landscape Around Architecture”
97 issued jointly by OASIS, OMG, and The Open Group **[SOA-NAV]** was written to help
98 the SOA community at large navigate the myriad of overlapping technical products
99 produced by these organizations with specific emphasis on the “A” in SOA, i.e.,
100 Architecture.

101 The white paper explains and positions standards for SOA reference models,
102 ontologies, reference architectures, maturity models, modeling languages, and
103 standards work on SOA governance. It outlines where the works are similar, highlights
104 the strengths of each body of work, and touches on how the work can be used together
105 in complementary ways. It is also meant as a guide to users for selecting those
106 specifications most appropriate for their needs.

107 While the understanding of SOA and SOA Governance concepts provided by these
108 works is similar, the evolving standards are written from different perspectives. Each
109 specification supports a similar range of opportunity, but has provided different depths
110 of detail for the perspectives on which they focus. ~~Therefore, a~~ Although the definitions
111 and expressions may differ ~~somewhat~~, there is agreement on the fundamental concepts
112 of SOA and SOA Governance.

113 The following is a summary taken from **[SOA-NAV]** of the positioning and guidance on
114 the specifications:

- 115 • The OASIS Reference Model for SOA (SOA RM) is the most abstract of the
116 specifications positioned. It is used for understanding core SOA concepts
- 117 • The Open Group SOA Ontology extends, refines, and formalizes some of the
118 core concepts of the SOA RM. It is used for understanding core SOA concepts
119 and facilitates a model-driven approach to SOA development.

- 120 • The OASIS Reference Architecture Foundation for SOA (this document) is an
121 abstract, foundational reference architecture addressing [the a broader](#)
122 [ecosystem viewpoint](#) for building and interacting within the SOA paradigm. It is
123 used for understanding different elements of SOA, the completeness of SOA
124 architectures and implementations, and considerations for reaching across
125 ownership boundaries where there is no single authoritative entity for SOA and
126 SOA governance.
- 127 • The Open Group SOA Reference Architecture is a layered architecture from
128 consumer and provider perspective with cross cutting concerns describing these
129 architectural building blocks and principles that support the realizations of SOA. It
130 is used for understanding the different elements of SOA, deployment of SOA in
131 enterprise, basis for an industry or organizational reference architecture,
132 implication of architectural decisions, and positioning of vendor products in a
133 SOA context.
- 134 • The Open Group SOA Governance Framework is a governance domain
135 reference model and method. It is for understanding SOA governance in
136 organizations. The OASIS Reference Architecture for SOA Foundation contains
137 an abstract discussion of governance principles as applied to SOA across
138 boundaries
- 139 • The Open Group SOA Integration Maturity Model (OSIMM) is a means to assess
140 an organization's maturity within a broad SOA spectrum and define a roadmap
141 for incremental adoption. It is used for understanding the level of SOA maturity in
142 an organization
- 143 • The Object Management Group SoaML Specification supports services modeling
144 UML extensions. It can be seen as an instantiation of a subset of the Open
145 Group RA used for representing SOA artifacts in UML.

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

Comment [PFB2]: Nuanced to distinguish from the first of the three 'formal' viewpoints

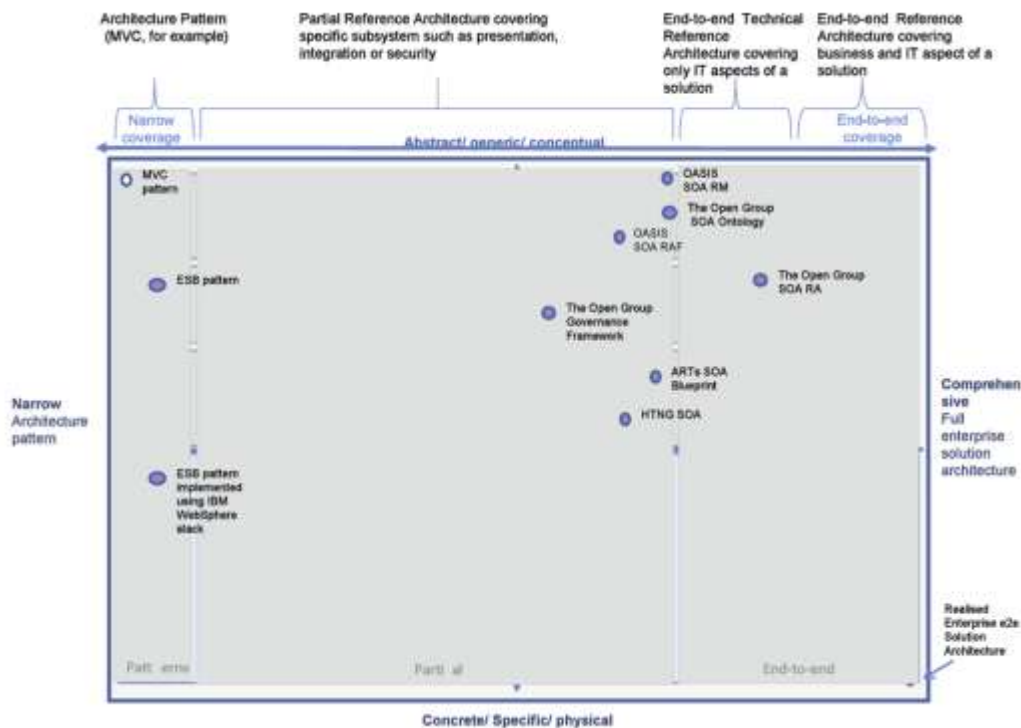


Figure 1- SOA Reference Architecture Positioning (from "Navigating the SOA Open Standards Landscape Around Architecture," © OASIS, OMG, The Open Group).

1.1.6 Expectations set by the Reference Architecture Foundation

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

1.2 Service Oriented Architecture – An Ecosystems Perspective

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

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

Comment [PFB3]: Tighten up.

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

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

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

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

189 1.3 Viewpoints, Views and Models

190 1.3.1 ANSI/IEEE 1471-2000::ISO/IEC 42010-2007

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

- 195 1. Architectural description identification, version, and overview information
- 196 2. Identification of the system stakeholders and their concerns judged to be relevant
197 to the architecture
- 198 3. Specifications of each viewpoint that has been selected to organize the
199 representation of the architecture and the rationale for those selections
- 200 4. One or more architectural views
- 201 5. A record of all known inconsistencies among the architectural description's
202 required constituents
- 203 6. A rationale for selection of the architecture (in particular, showing how the
204 architecture supports the identified stakeholders' concerns).

205 | The standard defines the following terms⁶:

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

206 **Architecture**

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

210 **Architectural Description**

211 A collection of products that document the architecture.

212 **System**

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

215 **System Stakeholder**

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

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

223 When describing architectures, it is important to identify stakeholder concerns and
224 associate them with viewpoints to insure that those concerns will be addressed in
225 some manner by the models that comprise the views on the architecture. The standard
226 defines views and viewpoints as follows:

227 **View**

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

230 **Viewpoint**

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

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

237 [\[IEEE 1471 Figure 1 of conceptual model of architectural descriptions\]](#)

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

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

249 **Model**

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

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

256 **1.3.2 UML Modeling Notation**

257 An open standard modeling language is used to help visualize structural and behavioral
258 architectural concepts. Although many architecture description languages exist, we
259 have adopted the Unified Modeling Language™ 2 (UML® 2) **[UML 2]** as the main
260 viewpoint modeling language. ~~It should be noted that while UML 2 is used, the~~
261 ~~formalization and recommendation of a UML Profile for SOA is beyond the scope of the~~
262 ~~SOA-RAF. Additionally,~~ Normative UML is used unless otherwise ~~stated but noted~~ it
263 should be noted that ~~UML can it can~~ only partially describe the concepts in each model
264 – it is important to read the text in order to gain a more complete understanding of the
265 concepts being described in each section..

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

267 **1.4 SOA-RAF Viewpoints**

268 The RAF uses three views that conform to three viewpoints: ~~the Ecosystem~~
269 ~~View~~ *Participation in a SOA Ecosystem*; ~~the Realizing Service Oriented Architecture~~
270 ~~Realization of a SOA Ecosystem~~ *View*, and ~~the Ownership in a SOA Ecosystem~~ *Owning*
271 ~~Service Oriented Architectures~~ *View*. There is a one-to-one correspondence between
272 viewpoints and views (see Table 1).

Viewpoint Element	Viewpoint		
	<u>Participation in a SOA Ecosystem</u>	<u>System Realization of a SOA Ecosystem</u>	<u>Process Ownership in a SOA Ecosystem</u>
Main concepts covered	Captures what SOA is meant for people to participate in a <u>service SOA ecosystem</u> .	Deals with the requirements for constructing <u>Captures what is meant to realize a SOA-based system in a SOA ecosystem</u> .	Addresses issues involved in owning and managing a SOA. <u>Captures what is meant to own a SOA-based system in a SOA ecosystem</u>
Stakeholders addressed	All stakeholders participants in the <u>SOA ecosystem</u>	Standards Architects, Enterprise Architects, Business Analysts, Decision Makers. <u>Those involved in the design, development and deployment of SOA-based systems</u>	Service Providers, Service Consumers, Enterprise Architects, Decision Makers. <u>Those involved in governing, managing and securing SOA-based systems</u>
Concerns addressed	Understanding ecosystem constraints and contexts in which business can be conducted safely <u>predictably</u> and effectively.	Effective construction of SOA-based systems.	Creation and management of effective, equitable, and assured SOA-based processes. <u>Processes to ensure governance, management and security of SOA-based systems.</u>
Modeling Techniques used	UML class diagrams	UML class, sequence,, component, activity, communication, and composite structure diagrams	UML class and communication diagrams

Comment [PFB4]: Check goals

Comment [PFB5]: Check models and model names

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

1.4.1 ~~Ecosystem View~~ [Participation in a SOA Ecosystem viewpoint](#)

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

All stakeholders in the ecosystem have concerns addressed by this viewpoint. The primary concern for people is to ensure that they can conduct their business effectively

281 and safely in accordance with the SOA paradigm. The primary concern of decision
282 makers is the relationships between people and organizations using systems for which
283 they, as decision makers, are responsible but which they may not entirely own, and for
284 which they may not own all of the components of the system.
285 Given SOA's value in allowing people to access, manage and provide services across
286 ownership boundaries, we must explicitly identify those boundaries and the implications
287 of crossing them.

288 **1.4.2 System-ViewRealization of a SOA Ecosystem viewpoint**

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

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

295 **1.4.3 ProcessOwnership in a SOA Ecosystem Vviewpoint**

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

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

306 **1.5 Terminology**

307 The keywords "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT",
308 "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this
309 document are to be interpreted as described in [RFC2119].

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

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

314 **1.5.1 Usage of Terms**

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

321 1.6 References

322 1.6.1 Normative References

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

347 1.6.2 Non-Normative References

- 348 [BLOOMBERG/SCHMELZER] Jason Bloomberg and Ronald Schmelzer, *Service*
349 *Orient or Be Doomed!*, John Wiley & Sons: Hoboken, NJ, 2006.
- 350 [COX] D. E. Cox and H. Kreger, "Management of the service-oriented
351 architecture life cycle," "IBM Systems Journal" "44", No. 4, 709-
352 726, 2005
- 353 [ERA] A. Fattah, "Enterprise Reference Architecture," paper presented at
354 22nd Enterprise Architecture Practitioners Conference, London, UK,
355 April 2009.
- 356 [ITU-T Rec. X.700 | ISO/IEC 10746-3:1996(E)] Information processing systems—
357 Open Systems Interconnection—Basic Reference Model—Part 4:
358 Management Framework", International Telecommunication Union,
359 International Organization for Standardization and International
360 Electrotechnical Commission, Geneva, Switzerland, 1989.
- 361 [NEWCOMER/LOMOW] Eric Newcomer and Greg Lomow, *Understanding SOA with*
362 *Web Services*, Addison-Wesley: Upper Saddle River, NJ, 2005.

363 [OECD] Organization for Economic Cooperation and Development,
364 Directorate for Financial, Fiscal and Enterprise Affairs, OECD
365 Principles of Corporate Governance, SG/CG(99) 5 and 219, April
366 1999.

367 [TOGAF v9] *The Open Group Architecture Framework (TOGAF) Version 9*
368 *Enterprise Edition*, The Open Group, Doc Number: G091, February
369 2009.

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

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

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

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

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

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

390 **2 Architectural Goals and Principles**

391 | This section identifies the goals of [the this](#) Reference Architecture Foundation and the
392 architectural principles that underpin it.

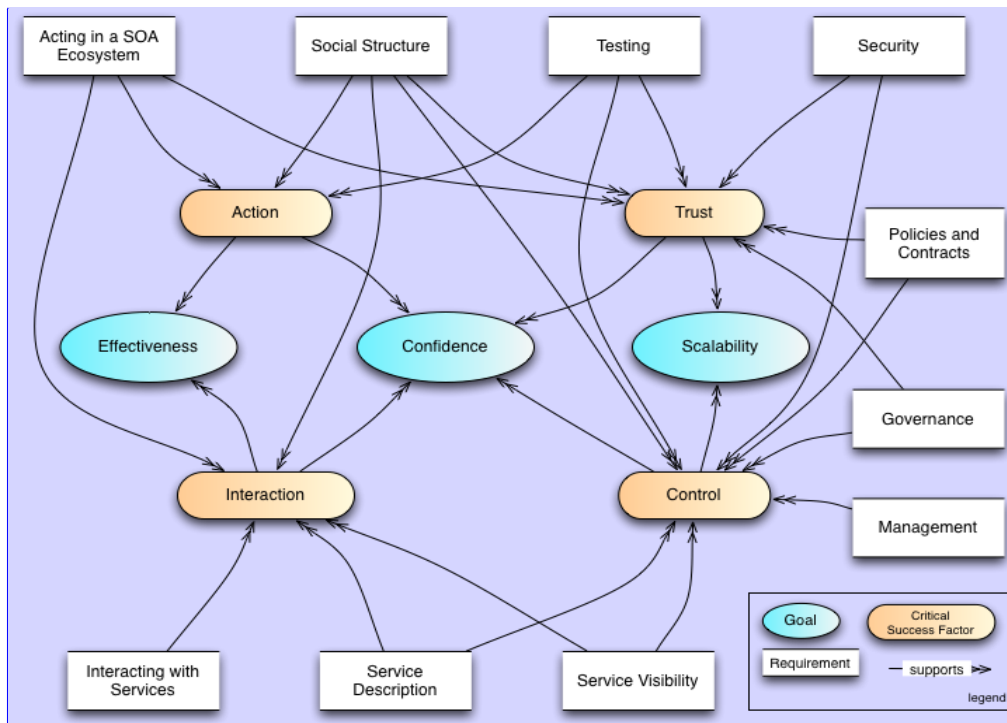
393 **2.1 Goals and Critical Success Factors of the Reference Architecture**
394 **Foundation**

395 There are three principal goals:

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

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

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



Comment [PFB6]: Need to make sure we still use these terms in the rest of the doc, and have not changed them. Also, double-check this for completeness after getting agreement on the section 3 text.

Figure 24 Critical Factors Analysis of the Reference Architecture

Figure 24 represents a Critical Factors Analysis (CFA) diagram demonstrating the relationship between the primary goals of this reference architecture, critical factors that determine the success of the architecture and individual elements that need to be modeled.

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

2.1.1 Goals

2.1.1.1 Effectiveness

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

The key factors that govern effectiveness from a participants' perspective are joint **action**, and **interaction** – especially interaction across ownership boundaries – with other participants in the ecosystem.

Comment [PFB7]: We need to relate action and joint action to resources or ownership boundaries in section 3.

2.1.1.2 Confidence

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

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

2.1.1.3 Scalability

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

2.1.2 Critical Success Factors

A critical success factor (CSF) is a property of the intended system, or a sub-goal that directly supports a goal and there is strong belief that without it the goal is unattainable. CSFs are not necessarily measurable in themselves. As illustrated in [Figure 2Figure 4](#), CSFs can be associated with more than one goal.

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

2.1.2.1 Action

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

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

2.1.2.2 Trust

The viability of a SOA ecosystem depends on participants being able to effectively measure the trustworthiness of the system and of participants. Trust is a private assessment of a participant's belief in the integrity and reliability of the SOA ecosystem (see Section [3.2.43.2.3](#)).

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

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

475 2.1.2.3 Interaction

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

480 2.1.2.4 Control

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

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

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

493 2.2 Principles of this ~~Reference~~ Reference Architecture Foundation

494 The following principles serve as core tenets that guided the evolution of this
495 Reference Architecture.

496 Technology Neutrality

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

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

506 Implications: This The Reference Architecture Foundation must be technology neutral,
507 meaning that we assume that technology will continue to evolve, and that
508 over the lifetime of this architecture that multiple, potentially competing
509 technologies will co-exist. Another immediate implication of technology
510 independence is that greater effort on the part of architects and other

511 decision makers to construct systems based on this architecture is
 512 needed.

513 **Parsimony**

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

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

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

523 **Distinction of Concerns**

524 Statement: Distinction of Concerns refers to the ability to cleanly identify and separate
 525 out the concerns of specific stakeholders in such a way that it is possible
 526 to create architectural models that reflect those stakeholders’ viewpoint. In
 527 this way, an individual stakeholder or a set of stakeholders that share
 528 common concerns only see those models that directly address their
 529 respective areas of interest.

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

536 Implications: This is a core tenet that drives this [Reference Architecture](#) to adopt the
 537 notion of architectural viewpoints and corresponding views. A *viewpoint*
 538 provides the formalization of the groupings of models representing one set
 539 of concerns relative to an architecture, while a *view* is the actual
 540 representation of a particular system. The ability to leverage an industry
 541 standard that formalizes this notion of architectural viewpoints and views
 542 helps us better ground these concepts for not only the developers of this
 543 [Reference Architecture](#) but also for its readers. The IEEE
 544 Recommended Practice for Architectural Description of Software-Intensive
 545 Systems [[ANSI/IEEE 1471-2000::ISO/IEC 42010-2007](#)] is the standard
 546 that serves as the basis for the structure and organization of this
 547 [Reference Architecture document](#).

548 **Applicability**

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

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

554 | Implications: ~~This~~[The](#) Reference Architecture [Foundation](#) needs to be relevant to the
555 problem of matching needs and capabilities under disparate domains of
556 ownership; to the concepts of “Intranet SOA” (SOA within the enterprise)
557 as well as “Internet SOA” (SOA outside the enterprise); to the concept of
558 “Extranet SOA” (SOA within the extended enterprise, i.e., SOA with
559 suppliers and trading partners); and finally, to “net-centric SOA” or
560 “Internet-ready SOA.”

3 Ecosystem View Participation in a SOA Ecosystem view

No man is an island

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

John Donne

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

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

The Ecosystem View Participation in a SOA Ecosystem view in the SOA-RAF focuses on the constraints and context in which people⁷ conduct business using a SOA-based system. By business we mean any shared activity entered into whose **objective** is to satisfy particular **needs** of each person.

The people actively participating in a SOA-based system, together with others who may potentially benefit from the services delivered by the system, together constitute the **stakeholders**. The stakeholders, the system and the environment (or context) within which they all operate, taken together forms the **SOA ecosystem**. That ecosystem may reflect the SOA-based activities within a particular enterprise or of a wider network of one or more enterprises and individuals. While a SOA-based system is essentially an IT concern, it is nonetheless a system engineered deliberately to be able to function in a SOA ecosystem. In this context, a service is the mechanism that brings a SOA-based

⁷ 'People' and 'person' must be understood as both human actors and 'legal persons', such as companies, who have rights and responsibilities similar to 'natural persons' (humans)

599 | [system capability together with stakeholder needs in the wider ecosystem. This is](#)
600 | [explored in more detail in Section 3.2.2. below.](#)

601 | Furthermore, this ~~ecosystem view~~[Participation in a SOA Ecosystem view](#) helps us
602 | understand the importance of execution context – the set of technical and business
603 | elements that allow interaction to occur in, and thus business to be conducted using, a
604 | SOA-based system.

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


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

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

623 | The main models in this view are:

- 624 | • the **Social Structure in a SOA Ecosystem Model** introduces the key elements
625 | that underlie the relationships between participants;
- 626 | • the **Action in a SOA Ecosystem Model** introduces the key concepts involved in
627 | actions, and shows how ownership, risk and transactions are key concepts in
628 | the SOA ecosystem.
- 629 | • The **Semantics in a SOA Ecosystem Model** focuses on the concepts that
630 | underlie meaning within the SOA ecosystem. It introduces the fundamental
631 | concept of proposition and shows how policies, facts and communication are all
632 | dependent on semantics.

Comment [PFB8]: This model, more than any needs to explain the Connection between humans and machines in the ecosystem. Maybe they all do, but somewhere we need to explicitly make this connection



Comment [PFB9]: Update model and view names

633 | Figure 32 Model elements described in the ~~ecosystem view~~[Participation in a SOA Ecosystem view](#)

3.1 Social Structure in a SOA Ecosystem Model

The actions undertaken by participants in a SOA ecosystem are performed in a *social context* that defines the relationships between the participants. That context is the social structure.

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




Figure 43 Social Structure

Social Structure

A social structure⁸ is a nexus of relationships amongst participants brought together for a specific purpose.

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

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

A social structure has a purpose – the overarching reason for which it exists. All social structures are established with implied or explicitly defined purpose. When explicit, this The purpose is usually defined as reflected in specific goals laid down in the social structure's constitution or other 'charter'.

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

Comment [PFB10]: Needs reworking. Remove 'enterprise' and 'market' and have an association between Social Structure and Goal; and another, 'Constitution defines Goal' (1..*)

⁸ Social structures are sometimes referred to as social institutions.

666 The RAF is concerned primarily with social structures that reflect relationships amongst
667 **participants** in SOA ecosystems-, [notably:](#)

668 [The SOA ecosystem is marked by two primary forms of social structure](#)

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

674 **Enterprise**

675 An enterprise is a social structure with internally established goals that reflect [its](#)
676 [a defined](#) purpose. It can act as a participant within other social structures,
677 including other enterprises.

678 **Peer**

679 A peer social structure is the locus of interaction between participants [with](#)
680 [individual goals](#) who are peers of one another.

681 Many interactions between participants take place within an enterprise social structure.
682 Depending on the scale and internal structure of the enterprise, these interactions can
683 cross ownership boundaries (an enterprise can itself be composed of enterprises).
684 However, interactions between participants within a peer social structure are inherently
685 across ownership boundaries.

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

691 **Constitution**

692 A constitution is a set of rules, written or unwritten, that [defines spell out the](#)
693 [purpose, goals, scope, and functioning of](#) a social structure.

694 Every social structure functions according to rules by which participants interact with
695 each other within the structure. In some cases, this is based on an explicit agreement,
696 in other cases participants behave as though they agree to the constitution without a
697 formal agreement. In still other cases, participants abide by the rules with some degree
698 of reluctance – this is an issue raised later on when we discuss governance in SOA-
699 based systems.

700 **3.1.1 Participants, Actors and Delegates**

701 As noted above, social structures have participants and stakeholders, some of whom
702 may be enterprises. They interact within the broad ecosystem. Actors operate within a
703 system and can be either **participants** (who are also stakeholders) – they have a stake
704 in the ecosystem; or **delegates** (human actors with no stake in the ecosystem or
705 automated agents), who act on behalf of a participant.




Figure 54 Actors, Participants and Delegates

Participant

~~A participant is a person⁹ who is both a stakeholder in the SOA ecosystem and an actor in the SOA-based system.~~

Stakeholder

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

Note: Not all stakeholders necessarily participate in the SOA ecosystem; indeed, the interest of **non-participant stakeholders** may be in realizing the benefits of a well-functioning ecosystem and not suffering unwanted consequences.

Actor

An actor is a person or automated agent capable of action.

~~The concept of actor encompasses many kinds of entities, human and corporate participants, even semi-autonomous computational agents.~~ An actor is either a participant (by definition, also a stakeholder) or a delegate.

Participant

A participant is a person¹⁰ who is both a stakeholder in the SOA ecosystem and an actor in the SOA-based system.

Delegate

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

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

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

Comment [PFB11]: We may need to be able to distinguish known and unknown stakeholders (such as potential consumers) From an architect's viewpoint, when building a specific system, they need to be aware of the intended and possible scope of their system, which networks will be involved, which ownership boundaries may be crossed, etc. How is such a task undertaken if there are unknown potential stakeholders?

Comment [PFB12]: Need to address this with the TC. Later we speak of "intent" etc. associated with action. Intent comes from a human, and sometimes is "encoded" in automation - but automation does not itself have intent. We need to be careful not to anthropomorphize automation (Turing notwithstanding), but at the same time, the relationship between the human entities and the automation is important - especially in the ecosystem perspective.

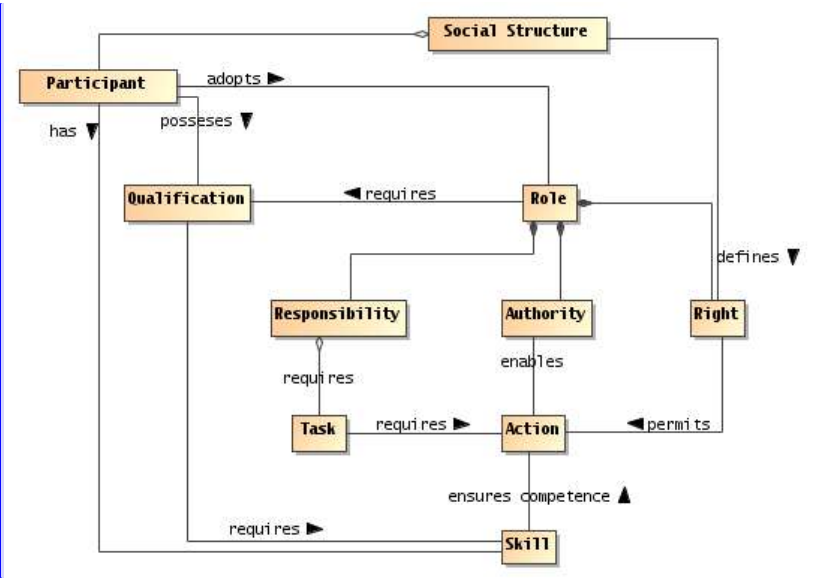
⁹ ~~Again, this can be a 'natural' or 'legal' person~~

¹⁰ ~~Again, this can be a 'natural' or 'legal' person~~

737 automation is performed by technology. If the actor is also a stakeholder in the
738 ecosystem, ~~on the other hand~~, then we use participant.
739 In order for a delegate to act on behalf of another person, they must be able to act and
740 have the authority to do so.

741 **3.1.2 Roles in Social Structures**

742 Social structures are abstractions: a social structure cannot directly perform actions –
743 only people or automated processes following the instructions of people can actually do
744 things. However, an actor may act on behalf of a social structure and certainly acts
745 within a social structure depending on the roles that the actor assumes.



Comment [PFB13]: We need one or two models, depending on whether we keep references to (human) Skill and its dependent concepts:
-One should cover the relationships around the concept of Skill (particular as a class of human Capability);
-A second should highlight the relationship between Skill as a Capability, and how it is brought to bear within specific action.
In any case, the current model is inappropriate and needs revision

746
747 Figure 65 Role in Social Structures

748 **Role**

749 A role is a type of relationship between a participant and a social structure that
750 reflects the rights, responsibilities, qualifications, and authorities of that
751 participant within the context of the social structure.

752 A role is not immutable and is often time-bound. A participant can have one or more
753 roles concurrently and may change them over time and in different contexts, even over
754 the course of a particular interaction. One participant with authority in the social
755 structure may formally *designate a role* for another participant, with associated rights
756 and responsibilities and that authority may even qualify a period during which this role
757 may be valid.

758 Conversely, someone who exhibits qualification and skill may *assume a role* without any
759 formal designation. For example, an office administrator who has demonstrated facility
760 with personal computers may be known as the 'goto' person for people who need help
761 with their computers.

762 While many roles are clearly identified, with appropriate names and definitions of
763 responsibilities, it is also entirely possible to separately bestow rights, [bestow or](#)
764 [assume](#) responsibilities and so on, often in a temporary fashion. For example, when a
765 company president delegates certain responsibilities on another person, this does not
766 imply that the other person has become company president.

767 **Right**

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

770 Rights can be constrained. For example, sellers might have a general right to refuse
771 service to potential customers but this right could be constrained so as to be exercised
772 only when based on certain criteria.

773 **Authority**

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

775 ~~Authority is constrained in terms of the kinds of actions that are authorized, and in terms~~
776 ~~of the skills and qualifications required to perform those actions and to be possessed by~~
777 ~~the persons invoking the authority (or on whom the authority rests).~~

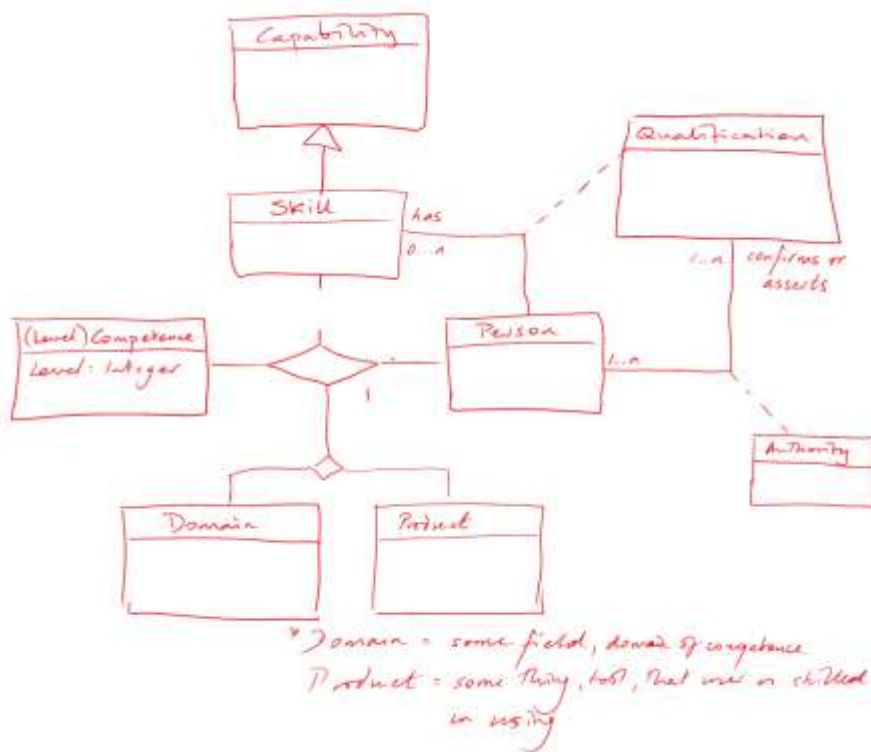
778 Rights, authorities, responsibilities and roles form the foundation for the security model
779 as well as contributing to the governance model in the 'Owning a SOA' View of the RAF.
780 Rights and responsibilities have similar structure to permissions and obligations; except
781 that the focus is from the perspective of the constrained participant rather than the
782 constrained actions.

783 **Responsibility**

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

786 Responsibility implies human [agencywillingness](#), which is why only participants, [and not](#)
787 [actors \(who can be non-human agents\)](#) are concerned, even if [their actionsthe](#)
788 [consequences of such responsibility](#) can impact ~~on~~ other (non-human) actors.

Comment [PFB14]: This is not shown in the UML model. Need another line in that.



Comment [PFB15]:
Ken: Needs explanation

Skill

A skill is a *human* capability.

Skills are typically defined in terms of requirements the competence of an individual (often measured to a certain level) to perform in a particular subject domain or with a particular tool or product. Such competence can therefore be used to ascertain whether a person is suitable for a given role: a role description may require that a participant possess a certain skill in order to assume that role.

Comment [PFB16]: Was "A skill..."

Qualification

A qualification is a public recognition that an actor is deemed by some authority to possess a particular skill.

In most cases a qualification involves the recognition within a particular social structure that the actor possesses the necessary set of skills that enable the actor to assume some role. The qualification may be granted for a specific period of time and may only be valid in certain contexts. A driver license may be granted for 5 or 10 years; it may be valid only for certain types of vehicle; it may only be valid in the country of issue or certain designated countries.

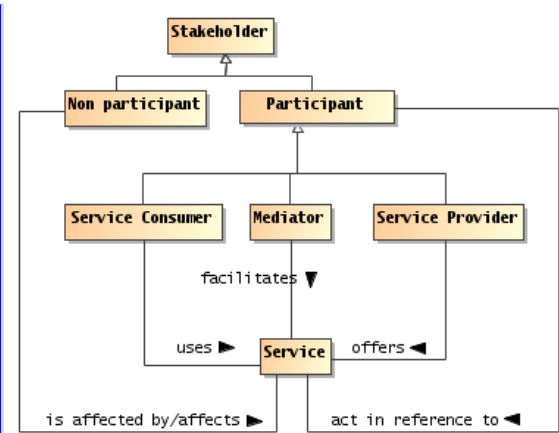
There is a distinction between a skill – the ability to act or capability – and a qualification – the recognition of that ability and, usually, the authority to act. Nonetheless, a qualification in itself is not a guarantee that the most appropriate and skilled person performs a given task. Likewise, and for any number of reasons, a person may be

810 assigned a role and given authority to act, irrespective of whether they have the
811 requisite skills or not.
812 For example, someone may have the skills to fly an airplane but not have a pilot's
813 license. Conversely, someone may have a pilot's license but, because of some
814 temporary cause (for example, illness), may be unable to fly a plane. In such a situation,
815 a person with the required skill but no qualification is preferable to a qualified person
816 who cannot perform the role or someone forced to assume that role without the required
817 skills.

818 3.1.2.1 Service Roles

819 In a SOA ecosystem, participants play one or more roles—in particular, offering and
820 using services. It is inherent to the SOA paradigm that a participant can play one or
821 more different roles in the ecosystem, including as a service consumer, a service
822 provider, a mediator, and so on, depending on the context. A participant may be playing
823 a role of a service provider in one relationship While playing simultaneously the role of a
824 consumer in another.
825

Comment [PFB17]: Heretical suggestion: do we need any of this, all relating to human attributes. What does it bring to an understanding of SOA ecosystems?



826
827 Figure 76 Service Roles

828 Participant roles in a Service

829 Service Provider

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

831 Service Consumer

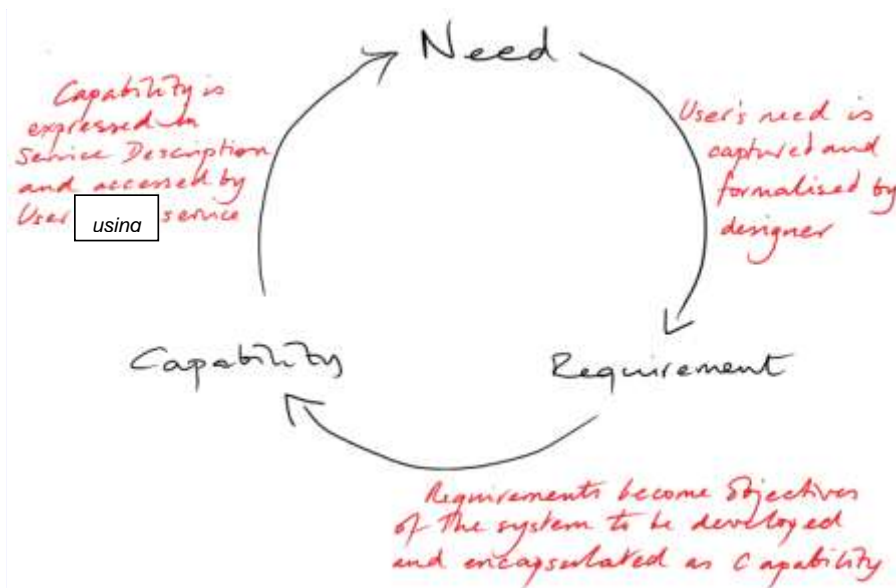
832 A service consumer is a role assumed by a participant who is interacting with a
833 service in order to fulfill a need.

834 Service Mediator

835 A service mediator is a role assumed by a participant to facilitate interaction and
836 connectivity in the offering and use of services~~who is facilitating the discovery,~~
837 ~~offer or use of services in some way.~~

Comment [PFB18]: Needs to be re-modelled, in line with TC discussions:
-Consumer, provider and (probably) mediator are all roles played by participants, not subclasses;
-Might want to model other roles, included roles played by a service vis-à-vis a participant (eg reg/rep)
We could also consider using other UML diagram types, such as the state diagram to help here (e.g. a service can perform different functions/roles when in different states)

838 It is a common understanding that service consumers typically initiate service
 839 interactions, although this is not necessarily true in all situations. As with service
 840 providers, several stakeholders may be involved in a service interaction supporting the
 841 consumer.



Comment [PFB19]: First sketch but needs some re-wording. Need gets mapped to requirements which get mapped to capabilities – however some of the capabilities get built, other capabilities already exist and are brought to bear via a service. Need to make this distinction or there will be confusion.
[Peter]: This first quick cut was done to emphasise the role of an ecosystem – to separate out the three distinct aspects (Needs, Requirements and Capabilities) and have stakeholders taking a stand for each. In contrast, systems are often limited to Requirements and Capabilities, not seeing wider stakeholder needs

842
 843 The roles of service provider and service consumer are often seen as symmetrical,
 844 which is not entirely correct. A consumer tends to express a 'Need' in non-formal terms:
 845 "I want to buy that book". The 'Need' has to be formalized and encapsulated by
 846 designers and developers as a 'Requirement'. This Requirement should then be
 847 reflected in the target service, as a 'Capability'. The Capability, when brought to bear
 848 using a service, delivers a 'Real World Effect' that *reflects* the capability and satisfies
 849 the need.

850 Service mediation by a participant can take many forms and may invoke and use other
 851 services ~~or components~~ in order to fulfill such mediation. For example, it might use a
 852 service registry in order to identify possible service partners; or it might use a filter that
 853 enhances another service by translating messages between English and Japanese.

Comment [PFB20]: This may need more work.

854 3.1.3 Shared State

855 State

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

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

861 For example, the total state of a lightbulb includes the temperature of the filament of the
 862 bulb. It also includes a great deal of other state – the composition of the glass, the dirt
 863 that is on the bulb's surface and so on. However, an actor may be primarily interested in

864 whether the bulb is 'on' or 'off' and not on the amount of dirt accumulated. That actor's
865 characterization of the state of the bulb reduces to the fact: 'bulb is now on'.

866 Many of the actions performed by people are inherently social in nature. Social
867 structures provide a context in which facts are given their meaning. For example, the
868 existence of a valid purchase order from a particular customer has a meaning that is
869 defined primarily by the company itself, ~~together with as well as~~ the wider social
870 structure (legal, political, etc.) that the company is ~~part of~~ embedded in.

871 Furthermore, social structures typically require formalized procedures to establish the
872 validity of particular facts. For example, the existence of an agreed contract typically
873 requires both parties to sign papers and to exchange those papers. If the signatures are
874 not performed correctly, or if the parties are not properly empowered to perform the
875 ritual, then it is as though nothing happened.

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

879 Other important classes of agreements include the policies adopted by an organization,
880 any agreements that it is holding for participants, and the assignment of roles to
881 participants within the organization.

882 ~~Many of the actions performed by people are inherently social in nature. The social~~
883 ~~context of an action is what gives it much of its meaning.~~

884 ~~For example, the total state of a lightbulb includes the temperature of the filament of the~~
885 ~~bulb. It also includes a great deal of other state — the composition of the glass, the dirt~~
886 ~~that is on the bulb's surface and so on. However, an actor may be primarily interested in~~
887 ~~whether the bulb is 'on' or 'off' and not on the amount of dirt accumulated. That actor's~~
888 ~~characterization of the state of the bulb reduces to the fact: 'bulb is now on'.~~

889 In a SOA ecosystem, there is a distinction between the set of facts about an entity that
890 only that entity can access — the so-called **Private State** — and the set of facts that
891 ~~might~~ may be accessible to other actors — the **Shared State**.

892 **Private State**

893 The private state of a actor is the entire state of an actor and that is knowable by,
894 and accessible to, only that actor.

895 **Shared State**

896 Shared state is that part of an actor's state that is knowable by, and may be
897 accessible to, other actors.

898 Note that shared state *does not* imply the state *is* accessible to *all* actors. It simply
899 refers to that subset of state that *may* be accessed by *other* actors through their
900 participation in **joint action**, defined below.

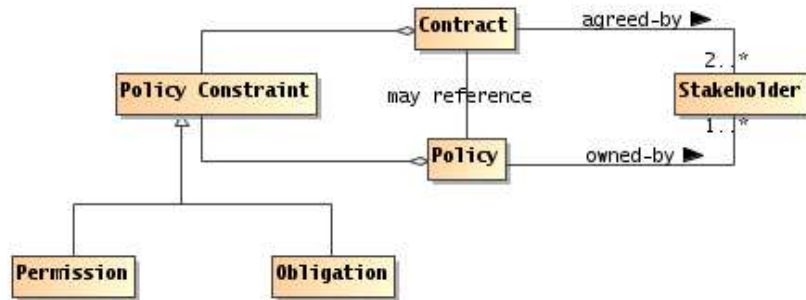
901 **3.1.4 Policies and Contracts**

902 As noted in the Reference Model, a **policy** represents some commitment and/or
903 constraint promulgated and enforced by a stakeholder and the stakeholder alone. A
904 **contract**, on the other hand, represents an agreement by two or more participants.
905 Enforcement of contracts may or may not be the responsibility of the parties to the

agreement but ~~will~~is usually ~~be~~ performed by a stakeholder in the ecosystem (public authority, legal system, etc.).

Commitment

~~A commitment is an objective explicitly stated by a participant.~~



Comment [PFB21]: Why 'policy constraint' rather than simply 'constraint'? I think this odel is very poor – Policy Constraint has multiple inheritance and strictly speaking is not a sub-class of either contract or Policy. Why not say, simply, that permissions and obligations <constrain> Policy?

Figure 87 Policies and Contracts

Policy

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

Policies can often be said to be about something – they have an object. For example, there may be policies about the use of a service. Policies have an **owner** – the stakeholder who asserts and takes responsibility for the policy. Thirdly, policies represent constraints – some measurable limitation on the state or behavior of the object of the policy, or of the behavior of the stakeholders of the policy.

Contract

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

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

Two important types of constraint that are relevant to a SOA ecosystem are Permission and ~~obligation~~Obligation.

Permission

A permission is a constraint that prescribes **actions** that an actor may (or may not) perform and/or the **states** the actor may (or may not) be in.

938 Note that permissions are distinct from ability and from authority. Authority refers to the
939 legitimate nature of an action as performed by an actor on behalf of a social structure
940 and ability refers to whether an actor has the capacity to perform the action, whereas
941 permission does not always involve acting on behalf of anyone.

942 **Obligation**

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

945 For example, once the service consumer and provider have entered into an agreement
946 to provide and consume a service, both participants incur obligations: the consumer is
947 obligated to pay for the service and the provider is obligated to provide the service –
948 based on the terms of the contract.

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

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


956 **3.1.5 Resource and Ownership**

957 ~~Fundamentally, we view ownership~~Ownership is defined as a relationship between a
958 stakeholder and a resource, where ~~someone (performing a role of the owner)~~ has
959 certain claims with respect to the resource.

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

962 **3.1.5.1 Resource**

963 A resource is generally understood as an asset: it has value to someone. Key to this
964 concept in a SOA ecosystem is that a resource needs to be identifiable. We define
965 resource as follows:



Comment [PFB22]: Drop concept of 'Identity' from our model? In any case, this model is wrong.

966
967 Figure 98 Resources

968 **Resource**

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

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

973 Contracts, policies, obligations, and permissions which are codified, services and
974 capabilities, and SOA-based systems are all examples of resources. An *implied* policy,
975 contract, obligation or permission would not be a resource.

976 **Identifier**

977 An identifier is any sequence of characters that may be unambiguously resolved
978 to identifying a resource.

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

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

983 The ability to identify a resource is important in interactions to determine such things as
984 rights and authorizations, to understand what functions are being performed and what
985 the results mean, and to ensure repeatability or characterize differences with future
986 interactions. The specific subset of individual characteristics that are necessary and
987 sufficient in order to unambiguously identify a resource ~~will~~depends on the ecosystem
988 and/or specific interactions within a system. However, a SOA ecosystem needs to
989 *unambiguously* identify a resource at any moment and in any interaction, many of which
990 may not be predictable given the operation of systems across ownership boundaries.
991 The way to achieve this with the use of identifiers.

992

993 **Description**

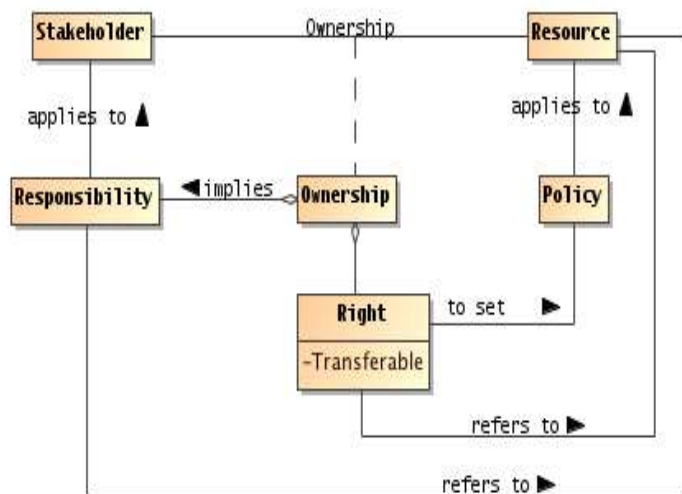
994 A description is a set of assertions about a resource.

995 Description as related to the SOA ecosystem is discussed in detail in Section 4.1.

996 **3.1.5.2 Ownership**

997 **Ownership**

998 Ownership is a particular set of claims, expressed as rights and responsibilities
999 that a stakeholder has in relation to a resource; It may include the right to transfer
1000 that ownership to another entity.



Comment [PFB23]: Poorly modeled and probably unnecessary.

Figure 109 Resource Ownership

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

A stakeholder who owns a resource may delegate some or all of these rights and responsibilities to others, but typically retains the right-responsibility to see that the delegated responsibilities-rights are not exercised as intended. There may also be joint ownership of a resource, where the responsibility is shared.

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

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

Ownership Boundary

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

3.2 Action in a SOA Ecosystem Model

At the core of participants' interest in a SOA ecosystem is the concept of action – participants act in order to further their goals. Critically, participants' actions may require

Comment [PFB24]: Need to show how action and joint action relate to resources
[Peter]: This is why ownership is key – and role of stakeholder as owner...

1028 use of resources that do *not belong to them*, i.e., that are outside of their ownership
1029 boundary.

1030 In this model we establish the key principles of action as an abstract concept. We
1031 elaborate on action in the context of a social context as joint action. Put simply, **joint**
1032 **actions** are simply coordinated **public actions** that involve more than one actor.

1033 Given that participants must communicate with each other we also show the role of
1034 communication in action and joint action.

1035 A key aspect of joint action revolves around the **trust** that both parties must exhibit in
1036 order to participate in joint actions. The willingness to act and a mutual understanding of
1037 both the information exchanged and the expected results is the particular focus of
1038 Section [3.2.43-2.3](#).

Comment [PFB25]: Do we need to be explicit?

Comment [PFB26]: Do we mean communicative action or something else here?

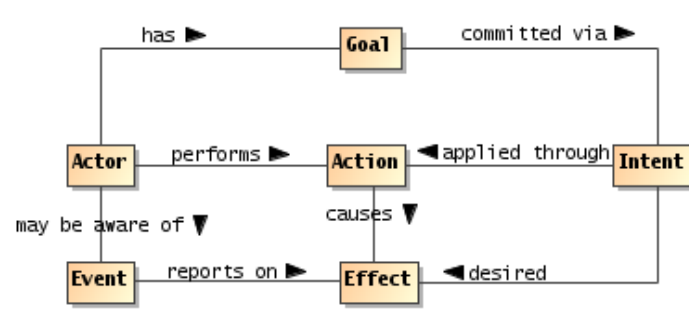
1039 **3.2.1 Action and Joint Action**

1040 Entities act in order to achieve their goals. However, the form of action that is of most
1041 interest within a SOA ecosystem is that involving interaction between more than one
1042 actor – **joint action**.

Comment [PFB27]: Reword

1043 **3.2.1.1 Action and Actors**

1044 As modeled in [Figure 11](#)~~Figure 40~~, **actions** are purposeful processes that actors
1045 engage in in order to achieve particular objectives.



Comment [PFB28]: Drop? Goal, Intent and Event have all been dropped as defined concepts

1046
1047 *Figure [11](#)~~40~~ Actions, Real World Effect and Events*

1048 **Action**

1049 An action is the application of **intent** to cause an **effect**.

1050 The aspect of action that distinguishes it from mere force or accident is that someone
1051 *intends* that the action achieves a desired objective or effect. ~~Whilst~~**While** this definition
1052 of action is very general, we are mostly concerned with actions that have specific effects
1053 on the SOA ecosystem – what we call **Real World Effects**. The ultimate real world
1054 effect of an action, however, may go beyond the initial intended effect.

1055 **Objective**

1056 An objective is the **real world effect** that a participant intends to achieve.

Comment [PFB29]: Shared state

1057 Objectives refer to **real world effects** that participants believe are achievable by a
1058 specific action or set of actions. In contrast, a goal is not linked with a specific action.

Comment [PFB30]: State change

1059 For example, someone may wish to have enough light to read a book. In order to satisfy
1060 that goal, the reader walks over to flip a light switch. The *objective* is to turn on the
1061 lamp, the *goal* is to be able to read. The real world effect is more light being available
1062 for me to read.

Comment [PFB31]: Change the state of the light bulb

1063 3.2.1.2 Actions and Real World Effects

1064 While an effect is any measurable change resulting from an action, a SOA ecosystem is
1065 concerned more specifically with real world effects.

1066 Real World Effect

1067 A real world effect is a measurable change to the shared state of one or more
1068 participants in an ecosystem.

1069 ~~Real world effect~~ This implies measurable change in the overall state of the SOA
1070 ecosystem ~~as a result of some action, although the change~~ ~~In particular, something~~
1071 ~~changed that~~ is primarily relevant to one or more participants ~~in the ecosystem.~~

1072 3.2.1.3 Joint Actions

1073 Joint actions are the foundation of interaction between participants in a SOA ecosystem.
1074 In this Reference Architecture Foundation, we see joint actions in at least two levels: as
1075 communication and as participants using and offering services.

1076 Joint Action

1077 A joint action is a coordinated set of actions involving the efforts of two or more
1078 actors to achieve a real world effect.

1079 Note that the effect of a joint action is *not* always equivalent to one or more effects of
1080 the individual actions of the participating actors, i.e., it may be more than the sum of the
1081 parts.

1082 Choreography

1083 ~~A choreography is a description of sequence and timing of individual actions in~~
1084 ~~order to successfully achieve one or more joint actions.~~

1085 ~~A choreography defines how individual actions performed by actors can be aggregated~~
1086 ~~together to denote joint actions.~~

1087 In order for multiple actors to participate in a joint action, they must each act according
1088 to their role within the joint action. This is achieved through communication and
1089 messaging, which in turn facilitate choreography and orchestration.

Comment [PFB32]: Do we need to link this back to the RM Process model? Maybe we should talk only about composability? Would this be easier?
What does the TC think?

1090 Communication – the formulation, transmission, receipt and interpretation of messages
1091 – is the foundation of all interaction within the SOA ecosystem, given the inherent
1092 separation – often across ownership boundaries – of actors in the system.

1093 Communication between actors requires that they play the roles of ‘sender’ or ‘receiver’
1094 of messages as appropriate to a particular action – although it is not necessarily
1095 required that they both be active simultaneously.

1096 An actor sends a message to another actor with the intent either

- 1097 • to communicate with other actors without the communication itself intending to
1098 cause a relevant real world effect – a **communicative action**; or

- 1099 • to establish joint action in order to deliver a capability ~~or part thereof which~~
1100 intentionally impacts the shared state and thus causes a real world effect – a
1101 **service action**.

1102 Communicative Action and Service Action are defined in more detail in Sections 3.2.4
1103 and 3.2.6 below.

1104 Different viewpoints ~~will~~ lead to different joint actions being interpreted as most
1105 important. For example, from the viewpoint of ecosystem **governance**, the integrity of
1106 the communicative action may be dominant; from the viewpoint of ecosystem security,
1107 the integrity of the service action may be dominant; the nature and fact of the
1108 established agreement may be dominant; from the viewpoint of ecosystem security, the
1109 communicative action may be dominant.

1110 The concept of joint action allows us to honor the fact that both parties in an interaction
1111 are required for there to be an actual effect; it allows us to separate out the different
1112 levels of the interaction into appropriate semantic layers; and it allows us to recombine
1113 those layers in potentially different ways whilst still achieving the intended real world
1114 effects of action in a SOA ecosystem.

1115 3.2.2 Needs and Capabilities

1116 Participants in a SOA ecosystem often need other participants to *do* something,
1117 leveraging a capability that they do not themselves possess. For example, a customer
1118 requiring a book may call upon a service provider to deliver the book. Likewise, the
1119 service provider needs the customer to pay for it.

1120 There is a reason that actors are engaged in this choreography process: different actors
1121 have different **needs** and have or apply different **capabilities** for satisfying those
1122 needs. It is core to the concept of a service. The SOA-RM defined a service as “the
1123 mechanism by which needs and capabilities are brought together”. This idea of services
1124 being a mechanism “between” needs and capabilities was introduced in order to
1125 emphasize capability as the notional or existing business functionality that would
1126 address well-defined needs.

1127 Business functionality

1128 Business functionality is a defined set of business-aligned tasks that provide
1129 recognizable business value to stakeholder ‘consumer’ and possibly others in the
1130 SOA ecosystem.

1131 The idea of a service in a SOA ecosystem combines business functionality with service
1132 implementation, including the artifacts needed and made available as IT resources.
1133 From the perspective of software developers, a SOA service enables the use of
1134 capabilities in an IT context. For the consumer, the service (combining business
1135 functionality and implementation) produces intended real world effects. They are not
1136 concerned with the underlying artifacts which make that delivery possible.

1137 A **need** is formalized as one or more **requirements** that must be fulfilled in order to
1138 achieve a stated goal.

1139 A **requirement** is a formal statement of a desired real world effect that, if achieved, will
1140 satisfy a need. This requirement can then be used to create a capability that in turn can
1141 be brought to bear to satisfy that need. ~~defined with one or more objectives and is~~

Comment [PFB33]: I am only adding the importance of the integrity of the respective actions

Comment [PFB34]: Gut feeling that this needs to come earlier in this section, together with the new (sketched) diagram proposed of Need > Requirement > Capability

1142 fulfilled by executing one or more joint actions to generate the real world effect to meet
1143 those objectives. In other words, a need is met or achieved by a capability brought to
1144 bear using a service. Both a the requirement, and the capability to fulfill it, is are
1145 expressed as in terms of a desired real world effect. An **objective**, on the other hand, is
1146 expressed as a desired and measurable change in the state of the ecosystem.

1147 Generally, a goal is a long term, broadly stated, desired state of the world that may
1148 be outcome that is, in practice, difficult to measure. On the other hand, an objective is a
1149 directly measurable and (preferably ideally) predictable outcome of a particular action or
1150 set of actions within an ecosystem.

1151 **Capability**

1152 A capability is an action or set of actions real world effect that a service provider
1153 is able to provide execute in order to provide a real world effect that responds to
1154 a service consumer's need.

1155 The Reference Model makes a distinction between a capability (as a potential to make a
1156 real world effect) and the ability of bringing that capability to bear (in a realized service)
1157 that delivers the real world effect.

1158 **3.2.3 Services Reflecting Business**

1159 The SOA paradigm often emphasizes the prescribed interface through which service
1160 interaction is accomplished. While this enables predictable integration in the sense of
1161 traditional software development, the prescribed interface alone does not guarantee that
1162 services will be composable into business solutions.

1163 **Business solution**

1164 A **business solution** is a set of defined interactions that combine implemented
1165 or notional business functionality in order to address a set of business needs.

1166 **Composability**

1167 **Composability** is the ability by which individual services providing defined
1168 business functionality can be combined to provide more complex business
1169 solutions

1170 Composability is important because many of the benefits of a SOA approach assume
1171 multiple uses for services, and multiple use requires the service deliver a business
1172 function that is applicable to multiple business solutions.

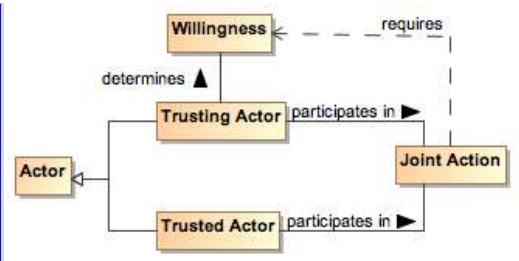
1173 To realize composability, capabilities must be identified that serve as building blocks for
1174 business solutions. In a SOA ecosystem, these building blocks are captured as services
1175 representing well-defined business functions, operating under well-defined policies and
1176 other constraints, and generating well-defined real world effects. These service building
1177 blocks should be relatively stable so as not to force repeated changes in the
1178 compositions that utilize them, but should also embody SOA attributes that readily
1179 support creating compositions that can be varied to reflect changing circumstances.

1180 The SOA paradigm emphasizes both composition of services and opacity of how a
1181 given service is implemented. With respect to opacity, the SOA-RM states that the
1182 service could carry out its described functionality through one or more automated and/or
1183 manual processes that themselves could invoke other available services.

1184 [Any composition can itself be made available as a service and the details of the](#)
1185 [business functionality, conditions of use, and effects are among the information](#)
1186 [documented in its service description.](#)
1187 [For services to be useful as composable building blocks in the SOA ecosystem, the](#)
1188 [services should, whenever possible, deliver capability that is applicable to multiple](#)
1189 [needs. Simply providing a Web Service interface for an existing IT artifact does not](#)
1190 [create opportunities for sharing business functions. Furthermore, the use of tools to](#)
1191 [auto-generate service software interfaces will not guarantee services than can](#)
1192 [effectively be used within compositions if the underlying code represents programming](#)
1193 [constructs rather than business functions. In such cases, services that tightly reflect the](#)
1194 [software details will be as brittle to change as the underlying code and will not exhibit](#)
1195 [the ill-defined characteristic of loose coupling.](#)

1196 **3.2.33.2.4 Trust and Risk**

1197 For a joint action to occur each actor must be able ~~to interact and the respective~~
1198 ~~stakeholders be able~~ and **willing** to participate ~~in the joint action~~. **Willingness** is the
1199 internal commitment of a human actor to carry out its part of a joint action.
1200 Important considerations in willingness are determining **trust** and **risk**.Willingness ~~on~~
1201 ~~the part of actors~~ to interact is not the same as a willingness to perform requested
1202 actions. For example, a service provider that rejects all attempts to perform some action
1203 may still be fully willing and engaged in interacting with the consumer.



1204
1205 Figure 1244 Trusting Actor and Willingness

1206 **Trust**

1207 Trust is a private assessment or internal perception of one participant that
1208 another entity will perform actions in accordance with an assertion regarding a
1209 desired real world effect.

1210 **Risk**

1211 Risk is a private assessment or internal perception that certain undesirable real
1212 world effects may result from action taken – or that the RWE might not meet
1213 certain criteria (e.g., performance).

1214 Trust is involved in all joint actions – it is necessary for *all* the actors involved ~~in a joint~~
1215 ~~action to trust each other~~ at least to the extent required for continuance of the joint
1216 action. The degree and nature of that trust ~~will is~~ likely ~~to~~ be different for each actor.

1217 An actor perceiving risk may take actions to mitigate the risk. At one extreme this will
1218 result in a refusal to interact. Alternately, it may involve adding protection – for example

Comment [PFB35]: Drop?

Comment [PFB36]: Not true of an orchestration/choreography. In the orch/cho, only the actors in a particular joint action need to trust each other, there does not have to be a transitive trust among all actors in that case.

1219 by using encrypted communication and/or anonymization – to reduce the perception of
1220 risk. Often standard procedures are put in place to increase trust and to mitigate risk.

1221 **3.2.3.13.2.4.1 Assessing Trust and Risk**

1222 The assessments of trust and risk are based on evidence available to the *trusting*
1223 *participant*. In general, participants will seek evidence from their private knowledge of
1224 the *trusted* actor as well as **evidence** of the **reputation** of the trusted actor.

1225 Trust is based on the confidence that one participant has accurately and sufficiently
1226 gathered and assessed evidence to a degree appropriate for the situation for which trust
1227 is being assessed.

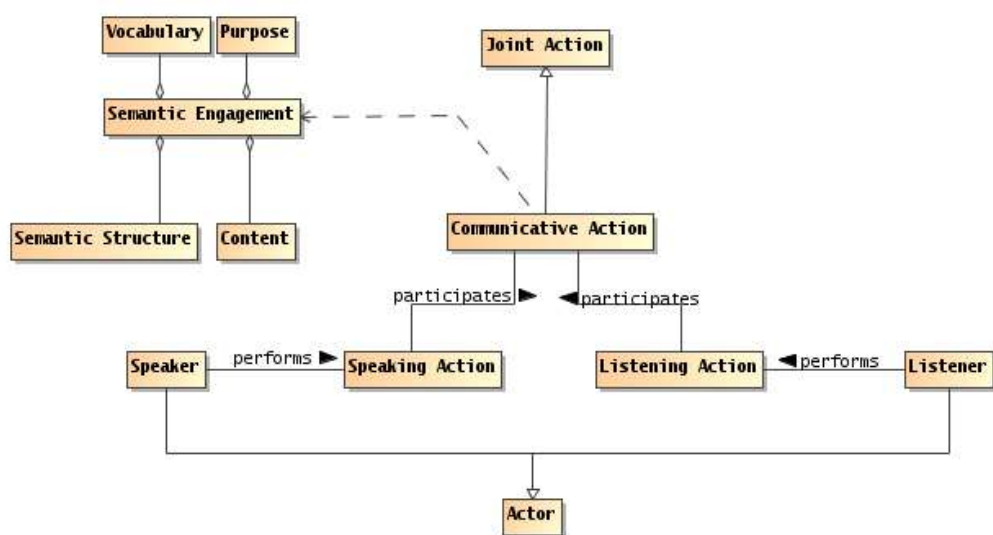
1228 Assessment of trust is rarely binary. An actor is not completely **trusted** or **untrusted**.
1229 There is typically some degree of uncertainty in the accuracy or completeness of the
1230 evidence or the assessment. Similarly, there is uncertainty in the amount and
1231 consequences of potential **risk**.

1232 The relevance of trust depends on the assessment of risk. If there is little or no
1233 perceived risk, then the degree of trust may not be relevant in assessing possible
1234 actions. For example, most people consider there to be an acceptable level of risk to
1235 privacy when using search engines, and submit queries without any sense of trust being
1236 considered.

1237 As perceived risk increases, the issue of trust becomes more of a consideration. For
1238 interactions with a high degree of risk, the trusting participant will typically require
1239 stronger or additional evidence when evaluating the balance between risk and trust.

1240 **3.2.43.2.5 Using Joint Action for Communication**

1241 We define communicative action as the action of message exchange:



1242
1243 Figure 1342 Communication as Joint Action

Comment [PFB37]: Should this be participant or actor? If we need to include human delegates, should we more precisely state 'human actor' to include this?

Comment [PFB38]: Drop? Or seriously redraft

1244 Communicative Action

1245 A communicative action is a joint action in which an actor communicates a
1246 message to one or more other actors.

1247 A communicative action has actors playing the role of **sender** and at least one
1248 **recipient**; all actors must perform their part in order for the communicative action to
1249 occur. ~~In addition, associated with the communicative action is the content of the~~
1250 ~~communication — the message.~~

1251 A given communicative action may have any number of **recipients**. In some situations,
1252 the sender may not be aware of the recipient. However, without both a sender and a
1253 recipient there is no communication.

1254 ~~Note that an In addition, a~~ Associated with the communicative action is the content of the
1255 communication — the message, which an actor receiving it a message not only acquires
1256 the message but must also be able to correctly interpret ~~it~~. The extent of that correct
1257 interpretation ~~will~~ depends on the role of the actor and the purpose of the
1258 communication.

1259 Even though communication is effected through action, it is not actually effective if the
1260 recipient cannot correctly interpret the message. However, interpretation can itself be
1261 characterized in terms of semantic engagement: the proper understanding of a
1262 message in a given context.

1263 We can characterize the necessary modes of interpretation in terms of a shared
1264 understanding of a common vocabulary and of the purpose of the communication. More
1265 formally, we can say that a communication has a combination of message and purpose.

1266 Interactions between service consumers and providers do not need to resemble human
1267 speech. Machine-machine communication is typically highly stylized in form, it may
1268 have particular forms and it may involve particular terms not found in everyday human
1269 interaction.

1270 3.2.6 Semantics and Semantic Engagement

1271 A SOA ecosystem is a space in which actors need to share understanding¹¹ as well as
1272 sharing actions. Indeed, such shared understanding is a pre-requisite to a joint action
1273 being carried out as intended. It is vital to a trusted and effective ecosystem. Semantics
1274 are therefore pervasive throughout SOA ecosystems and important in communicative
1275 actions described above, as well as a driver for policies and other aspects of the
1276 ecosystem.

1277 In order to arrive at shared understanding, an actor must effectively process and
1278 understand assertions in a manner appropriate to the particular context. An assertion, in
1279 general, is a measurable and explicit statement made by an actor. In a SOA ecosystem,
1280 in particular, assertions are concerned with the ‘what’ and the ‘why’ of the state of the
1281 ecosystem and its actors.

¹¹ We use a mechanical, Turing test-based approach to understanding here: if an actor behaves as though it understands an utterance then we assume that it does understand it.

1282 Understanding and interpreting those assertions allows other actors to know what may
1283 be expected of them in any particular joint action. An actor can potentially 'understand'
1284 an assertion in a number of ways, but it is specifically the process of arriving at a *shared*
1285 understanding that is important in the ecosystem. This process is semantic engagement
1286 by the actor with the SOA ecosystem. It can be instantaneous or progressively
1287 achieved. What is important is that there is the level of engagement appropriate to the
1288 particular context.

1289 Semantic Engagement

1290 Semantic engagement is the process by which an actor engages with a set of
1291 assertions based on that actor's interpretation and understanding of those
1292 assertions.

1293 Different actors have differing capabilities and requirements for understanding
1294 assertions. This is true for both human and non-human actors. For example, a purchase
1295 order process does not require that a message forwarding agent 'understands' the
1296 purchase order, but a processing agent does need to 'know' what to do with the order once
1297 received.

1298 The impact of any assertion can only be fully understood in terms of specific social
1299 contexts; contexts that necessarily include the actors that are involved. For example, a
1300 policy statement that governs the actions relating to a particular resource may have a
1301 different impact or purpose for the participant that owns the resource than for the actor
1302 that is trying to access it: the former understands the purpose of the policy as a
1303 statement of enforcement; and the latter understands it as a statement of constraint.

1304 3.2.53.2.7 Using Communication for Service Joint Action for Real World 1305 Effect

1306 Like communicative actions, **service actions** are inherently joint actions. Unlike
1307 communicative actions, however, service actions cause real world effects. This is a
1308 result of participants delivering and benefitting from a service. ~~there can be no service~~
1309 ~~action without both the service and the actor originating the action.~~

1310 **Service Action**

1311 A service action is a joint action in which ~~the~~ participants deliver ~~the a~~ service or
1312 a part thereof.

1313 ~~Service actions are inherently joint actions; they require both the entity performing the~~
1314 ~~action and the service itself to participate in the action.~~

1315 The two systems involved in SOA-based systems are the system of communication on
1316 the one hand and the system of services on the other.

1317 **3.3 Semantics in a SOA Ecosystem Model**

1318 ~~Semantics is important to the SOA ecosystem because it is pervasive: it is explicitly~~
1319 ~~important in the communication between actors, but is also a driver for policies, and~~
1320 ~~many other aspects of the ecosystem.~~

1321 ~~In particular, we are concerned with how an actor can effectively process and~~
1322 ~~understand assertions about the 'what' and the 'why' of the state of an ecosystem and its~~

Comment [PFB39]: We titled 3.2.4
"Using Joint Action for
Communication", so this seemed
appropriate

1323 actors. This process is one of semantic engagement by the actor with the SOA
1324 ecosystem.
1325 An assertion is a measurable and explicit statement made by an actor in the ecosystem
1326 and the set of assertions provide the basis of the actor's semantic engagement with the
1327 ecosystem.
1328 Any proposition can only be fully understood in terms of specific social contexts;
1329 contexts that necessarily include the actors that are involved. For example, a policy
1330 statement that governs the **actions** relating to a particular **resource** has a different
1331 meaning for the **participant** that **owns** the **resource** than for the **actor** that is trying to
1332 access it: the former interprets the policy as a statement of enforcement; and the latter
1333 interprets the policy as a statement of constraint.
1334 In addition, the ability of an **actor** to understand assertions is also very variable and
1335 context sensitive. For any given **actor**, for any given **action**—whether the **action** is
1336 private or is a **joint action**—the semantic engagement reflects the degree to which an
1337 actor 'engages' in order to understand certain assertions.
1338 Even within a single **joint action** the different actors involved may have very different
1339 means of understanding the activity. For example, an **actor** requesting a particular
1340 record from a database may understand the request in terms of accessing customer
1341 data. The database **actor** "understands" the same request as at best an SQL script to
1342 execute.

1343 **Semantic Engagement**

1344 **Semantic engagement** is the process by which an actor engages with a set of
1345 assertions based on that actor's interpretation and understanding of those
1346 assertions.

1347 Semantic engagement refers to how an actor engages with, or understands, any
1348 assertions that it may encounter. Different actors will have differing capabilities and
1349 requirements for understanding assertions. This is true for both human and non-human
1350 actors.

1351 For example, a purchase order process does not require that any 'message forwarding
1352 agent' understands the purchase order, but such understanding is required of the 'order
1353 processing agent'.

1354 Theoretically, an actor can act on an infinite set of assertions; Practically, however, the
1355 set of assertions the actor can access at any one time is likely to finite.

1356 **3.3.1 Private and Public semantics**

1357 A SOA ecosystem can be viewed as a space in which actors share understanding as
1358 well as sharing **actions**. As such, we need to be able to distinguish the shared meaning
1359 of utterances from the full or private meanings of those utterances. An important
1360 distinction here is that of public versus private semantics:

1361 **Public Semantics**

1362 The **public semantics** of an assertion is that subset of the possible
1363 interpretations of the assertion that is available to any observer by virtue of the
1364 observer's situation in a social structure.

1365 ~~The public semantics of an assertion is effectively the ‘ecosystem perspective’ of the~~
1366 ~~assertion.~~
1367 ~~Public semantics depends on communication in some form. If an assertion (whether it is~~
1368 ~~the content of a communicative action or some other assertion such as a policy~~
1369 ~~statement) is never communicated then the issue of that assertion’s public semantics is~~
1370 ~~somewhat moot—there are no observers. Of course, the most obvious observer of an~~
1371 ~~assertion that has been communicated is the intended recipient of the communication.~~
1372 ~~However, in general, the public semantics of an assertion would enable any observer to~~
1373 ~~make the same inferences.~~

1374 3.3 Architectural Implications

1375 3.3.1 Social structures

1376 A SOA ecosystem’s participants are organized into various forms of social structure.
1377 Not all social structures are hierarchical: a SOA ecosystem should be able to
1378 incorporate peer-to-peer forms of organization as well as hierarchic structures. In
1379 addition, it should be possible to identify and manage any constitutional agreements
1380 that define the social structures present in a SOA ecosystem.

- 1381
 - 1382 • Different social structures have different rules of engagement
 - 1383 ○ Techniques for expressing constitutions are important
 - 1384 • social structures have roles and members
 - 1385 ○ Techniques for identifying, managing members of social structures
 - 1386 ○ Techniques for describing roles and role adoption
 - 1387 • social structures may be complex
 - 1388 ○ Child social structures’ constitutions depend on their parent constitutions
 - 1389 • Social structures overlap and interact
 - 1390 ○ A given actor may be member of multiple social structures
 - 1391 ○ Social structures may be associated with different jurisdictions
 - 1392 ○ Social structures may involved in disputes with one another
 - 1393 ▪ Requiring conflict resolution

1393 3.3.2 The Importance of Action

1394 Participants participate in a SOA ecosystem in order to get their needs met. This
1395 involves action; both individual actions and joint actions.

1396 Any architectural realization of a SOA ecosystem should address:

- 1397 • How actions are modeled:
 - 1398 ○ Identifying the performer or agent of the action;
 - 1399 ○ the target of the action; and the
 - 1400 ○ verb of the action.

1401 Joint actions are actions involving multiple actors. Any explicit models of joint action
1402 should take into account

- 1403 • The choreography that defines the joint action.
- 1404 • The potential for joint actions to be multiply layered on top of each other

3.3.3 Communications as a Means of Mediating Action

Using message exchange for mediating action implies

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

3.3.4 Semantics

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

The operation of the SOA ecosystem is significantly enhanced if

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

3.3.5 Trust and Risk

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

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

1445
1446
1447
1448
1449
1450
1451
1452
1453
1454
1455
1456

3.3.6 Policies and Contracts

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

4 Realizing Service Oriented Architectures ViewRealization of a SOA Ecosystem view

Make everything as simple as possible but no simpler.

Albert Einstein

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

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




Figure 1443 Model Elements Described in the Realizing Service Oriented Architectures ViewRealization of a SOA Ecosystem view

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

4.1 Service Description Model

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

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

Comment [PFB40]: Should? Must?

1489 There are several points to note regarding the following discussion of service
1490 description:

- 1491 • The Reference Model states that one of the hallmarks of SOA is the large amount of
1492 associated description. The model presented below focuses on the description of
1493 services but it is equally important to consider the descriptions of the consumer,
1494 other participants, and needed resources other than services.
- 1495 • Descriptions are inherently incomplete but may be determined as *sufficient* when it is
1496 possible for the participants to access and use the described services based only on
1497 the descriptions provided. This means that, at one end of the spectrum, a description
1498 along the lines of “*That service on that machine*” may be sufficient for the intended
1499 audience. On the other extreme, a service description with a machine-process-able
1500 description of the semantics of its operations and real world effects may be required
1501 for services accessed via automated service discovery and planning systems.
- 1502 • Descriptions come with context, i.e. a given description comprises information
1503 needed to adequately support the context. For example, a list of items can define a
1504 version of a service, but for many contexts an indicated version number is sufficient
1505 without the detailed list. The current model focuses on the description needed by a
1506 service consumer to understand what the service does, under what conditions ~~will~~
1507 the service will do it, how well does the service do it, and what steps are needed by
1508 the consumer to initiate and complete a service interaction. Such information also
1509 enables the service provider to clearly specify what is being provided and the
1510 intended conditions of use.
- 1511 • Descriptions ~~will~~ change over time as, for example, the ingredients and nutrition
1512 information for food labeling continues to evolve. A requirement for transparency of
1513 transactions may require additional description for those associated contexts.
- 1514 • Description always proceeds from a basis of what is considered "common
1515 knowledge". This may be social conventions that are commonly expected or possibly
1516 codified in law. It is impossible to describe everything and it can be expected that a
1517 mechanism as far reaching as SOA will also connect entities where there is
1518 inconsistent "common" knowledge.
- 1519 • Descriptions will become the collection point of information related to a service or
1520 any other resource, but it ~~will-is~~ not necessarily ~~be~~ the originating point or the
1521 motivation for generating this information. In particular, given a SOA service as the
1522 access to an underlying capability, the service may point to some of the capability's
1523 previously generated description, e.g. a service providing access to a data store may
1524 reference update records that indicate the freshness of the data.
- 1525 • Descriptions of the provider and consumer are the essential building blocks for
1526 establishing the execution context of an interaction.

1527 These points emphasize that there is no one “right” description for all contexts and for
1528 all time. Several descriptions for the same subject may exist at the same time, and this
1529 emphasizes the importance of the description referencing source material maintained
1530 by that material's owner rather than having multiple copies that become out of synch
1531 and inconsistent.

Comment [PFB41]: Shall? Must?

1532 It may also prove useful for a description assembled for one context to cross-reference
1533 description assembled for another context as a way of referencing ancillary information
1534 without overburdening any single description. Rather than a single artifact, description
1535 can be thought of as a web of documents that enhance the total available description.

1536 | This Reference Architecture [Foundation](#) uses the term service description for
1537 consistency with the concept defined in the Reference Model. Some SOA literature
1538 treats the idea of a “service contract” as equivalent to service description. In ~~this~~
1539 [Reference Architecture](#) [the SOA-RAF](#), the term service description is preferred.
1540 Replacing service description with service contract implies just one side of the
1541 interaction is governing and misses the point that a single set of policies identified by a
1542 service description may lead to numerous contracts, i.e. service level agreements,
1543 leveraging the same description.

1544 **4.1.1 The Model for Service Description**

1545 | ~~Figure 15~~[Figure 14](#) shows Service Description as a subclass of the general Description
1546 class, where Description is a subclass of the resource class as defined in Section
1547 3.1.5.1. In addition, each resource is assumed to have a description. The following
1548 section discusses the relationships among elements of general description and the
1549 subsequent sections focus on service description itself. Other descriptions, such as
1550 those of participants, are important to SOA but are not individually elaborated in this
1551 document.

1552 **4.1.1.1 Elements Common to General Description**

1553 The general Description class is composed of a number of elements that are expected
1554 to be common among all specialized descriptions supporting a service oriented
1555 architecture. A registry often contains a subset of the description instance, where the
1556 chosen subset is identified as that which facilitates mediated discovery. Additional
1557 information contained in a more complete description may be needed to initiate and
1558 continue interaction.

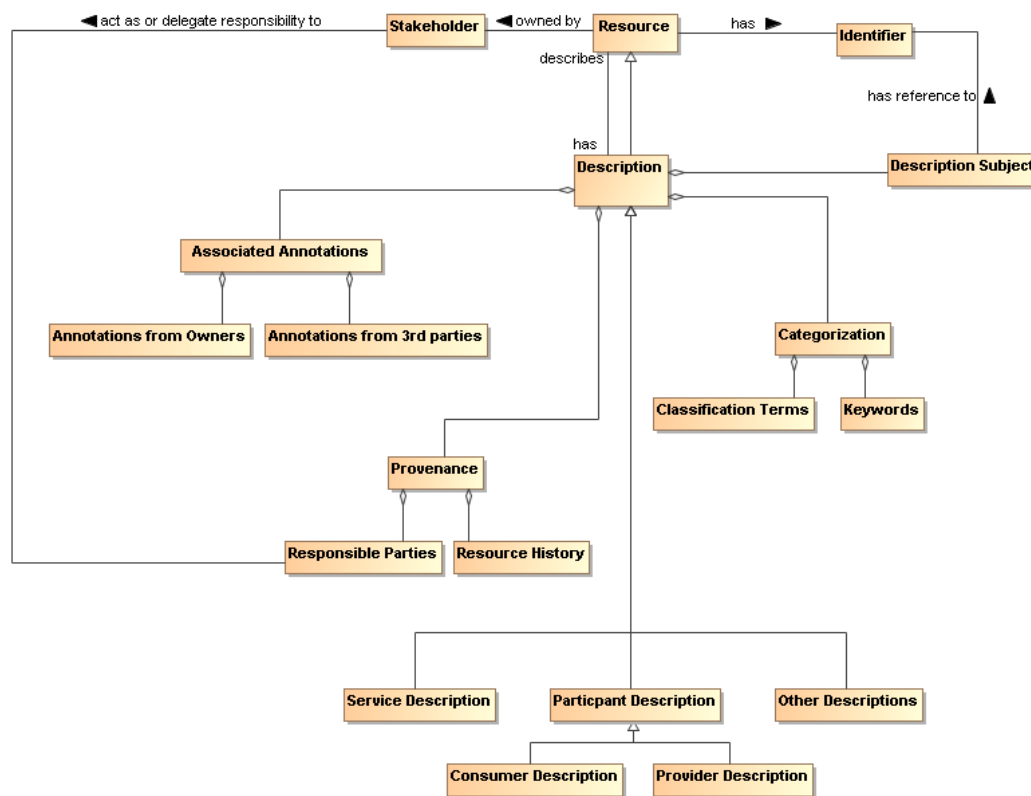


Figure 15-44 General Description

4.1.1.1.1 Description Subject

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

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

4.1.1.1.2 Provenance

While the resource Identifier provides the means to know which subject and subject description are being considered, Provenance as related to the Description class provides information that reflects on the quality or usability of the subject. Provenance specifically identifies the entity (human, defined role, organization, ...) that assumes responsibility for the resource being described and tracks historic information that

1578 establishes a context for understanding what the resource provides and how it has
1579 changed over time. Responsibilities may be directly assumed by the stakeholder who
1580 owns a resource or the Owner may designate Responsible Parties for the various
1581 aspects of maintaining the resource and provisioning it for use by others. There may be
1582 more than one entity identified under Responsible Parties; for example, one entity may
1583 be responsible for code maintenance while another is responsible for provisioning of the
1584 executable code. The historical aspects may also have multiple entries, such as when
1585 and how data was collected and when and how it was subsequently processed, and as
1586 with other elements of description, may provide links to other assets maintained by the
1587 resource owner.

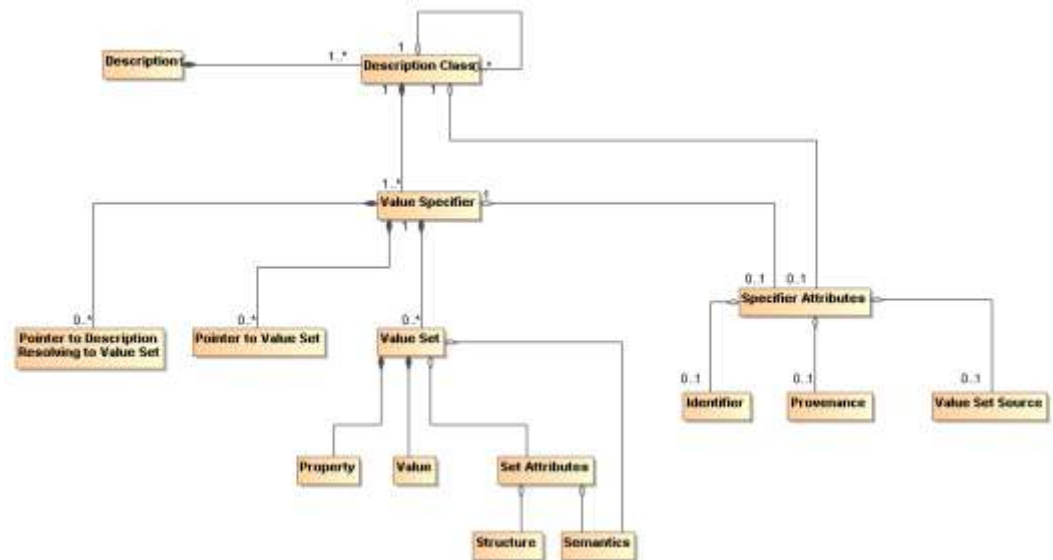
1588 4.1.1.1.3 Keywords and Classification Terms

1589 A traditional element of description has been to associate the resource being described
1590 with predefined keywords or classification taxonomies that derive from referenceable
1591 formal definitions and vocabularies. This Reference Architecture [Foundation](#) does not
1592 prescribe which vocabularies or taxonomies may be referenced, nor does it limit the
1593 number of keywords or classifications that may be associated with the resource. It
1594 does, however, state that a normative definition SHOULD be referenced, whether that
1595 be a representation in a formal ontology language, a pointer to an online dictionary, or
1596 any other accessible source. See Section 4.1.1.2 for further discussion on associating
1597 semantics with assigned values.

1598 4.1.1.1.4 Associated Annotations

1599 The general description instance may also reference associated documentation that is
1600 in addition to that considered necessary in this model. For example, the owner of a
1601 service may have documentation on best practices for using the service. Alternately, a
1602 third party may certify a service based on their own criteria and certification process;
1603 this may be vital information to other prospective consumers if they were willing to
1604 accept the certification in lieu of having to perform another certification themselves.
1605 Note, while the examples of Associated Documentation presented here are related to
1606 services, the concept applies equally to description of other entities.

1607 4.1.1.2 Assigning Values to Description Instances
1608



1609
1610 Figure 16-45 Representation of a Description

1611 Figure 15Figure-14 shows the template for a general description but individual
1612 description instances depend on the ability to associate meaningful values with the
1613 identified elements. Figure 16Figure-15 shows a model for a collection of information
1614 that provides for value assignment and traceability for both the value meaning and the
1615 source of a value. The model is not meant to replace existing or future schema or other
1616 structures that have or will be defined for specific implementations, but it is meant as
1617 guidance for the information such structures need to capture to generate sufficient
1618 description. It is expected that tools will be developed to assist the user in populating
1619 description and auto-filling many of these fields, and in that context, this model provides
1620 guidance to the tool developers.

1621 In Figure 16Figure-15 each class has an associated value specifier or is made up of
1622 components that will eventually resolve to a value specifier. For example, Description
1623 has several components, one of which is Categorization, which would have an
1624 associated a value specifier.

1625 A value specifier consists of

- 1626 • a collection of value sets with associated property-value pairs, pointers to such value
1627 sets, or pointers to descriptions that eventually resolve to value sets that describe
1628 the component; and
- 1629 • attributes that qualify the value specifier and the value sets it contains.

1630 The qualifying attributes for the value specifier include

- 1631 • an optional identifier that would allow the value set to be defined, accessed, and
1632 reused elsewhere;

Comment [PFB42]: Shall? Must?

1633 • provenance information that identifies the party (individual, role, or organization) that
 1634 has responsibility for assigning the value sets to any description component;

1635 • an optional source of the value set, if appropriate and meaningful, e.g. if a particular
 1636 data source is mandated.

1637 If the value specifier is contained within a higher-level component, (such as Service
 1638 Description containing Service Functionality), the component may inherit values from
 1639 the attributes from its container.

1640 Note, provenance as a qualifying attribute of a value specifier is different from
 1641 provenance as part of an instance of Description. Provenance for a service identifies
 1642 those who own and are responsible for the service, as described in Section 3.
 1643 Provenance for a value specifier identifies who is responsible for choosing and
 1644 assigning values to the value sets that comprise the value specifier. It is assumed that
 1645 granularity at the value specifier level is sufficient and provenance is not required for
 1646 each value set.

1647 The value set also has attributes that define its structure and semantics.

1648 • The semantics of the value set property should be associated with a semantic
 1649 context conveying the meaning of the property within the execution context, where
 1650 the semantic context could vary from a free text definition to a formal ontology.

1651 • For numeric values, the structure would provide the numeric format of the value and
 1652 the “semantics” would be conveyed by a dimensional unit with an identifier to an
 1653 authoritative source defining the dimensional unit and preferred mechanisms for its
 1654 conversion to other dimensional units of like type.

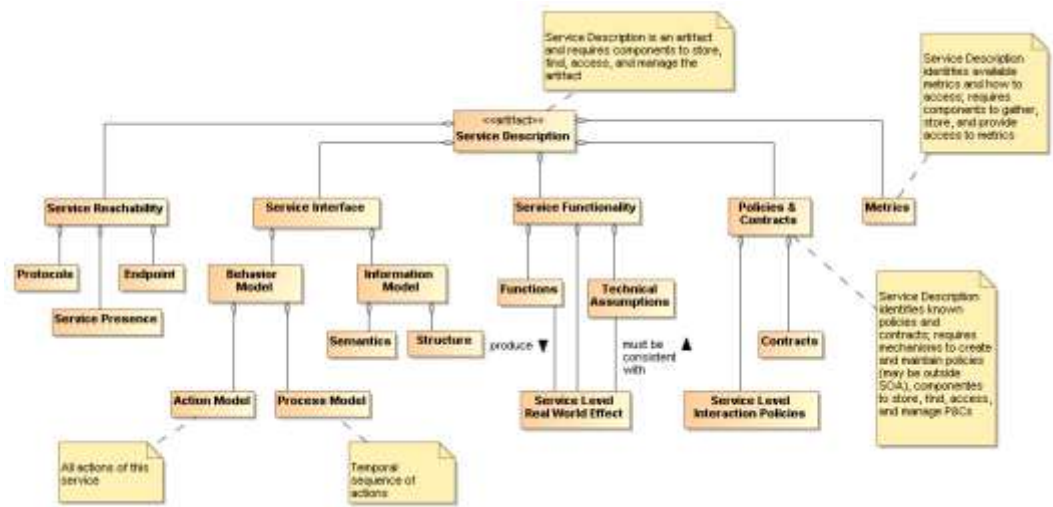
1655 • For nonnumeric values, the structure would provide the data structure for the value
 1656 representation and the semantics would be an associated semantic model.

1657 • For pointers, architectural guidelines would define the preferred addressing scheme.

1658 The value specifier may indicate a default semantic model for its component value sets
 1659 and the individual value sets may provide an override.

1660 The property-value pair construct is introduced for the value set to emphasize the need
 1661 to identify unambiguously both what is being specified and what is a consistent
 1662 associated value. The further qualifying of Structure and Semantics in the Set
 1663 Attributes allows for flexibility in defining the form of the associated values.

1664 4.1.1.3 Model Elements Specific to Service Description



1665
1666 Figure 1746 Service Description

1667 The major elements for the Service Description subclass follow directly from the areas
1668 discussed in the Reference Model. Here, we discuss the detail shown in Figure 1746
1669 and the purpose served by each element of service description.

1670 Note, the intent in the subsections that follow is to describe how a particular element,
1671 such as the service interface, is reflected in the service description, not to elaborate on
1672 the details of that element. ~~Other sections of the Reference Model and this Reference~~
1673 ~~Architecture describe the “physics” of each element whereas the service description~~
1674 ~~subsections will only touch on the meta aspects.~~

1675 4.1.1.3.1 Service Interface

1676 As noted in the Reference Model, the service interface is the means for interacting with
1677 a service. For this Reference Architecture the SOA-RAF and as shown in Section 4.3
1678 the service interface will support an exchange of messages, where

- 1679
- 1680 • the message conforms to a referenceable message exchange pattern (MEP),
 - 1681 • the message payload conforms to the structure and semantics of the indicated
1682 information model,
 - 1683 • the messages are used to denote events or actions against the service, where
1684 the actions are specified in the action model and any required sequencing of
actions is specified in the process model.

Comment [PFB43]: Description?
Or prescription? Does? Shall?
Must?

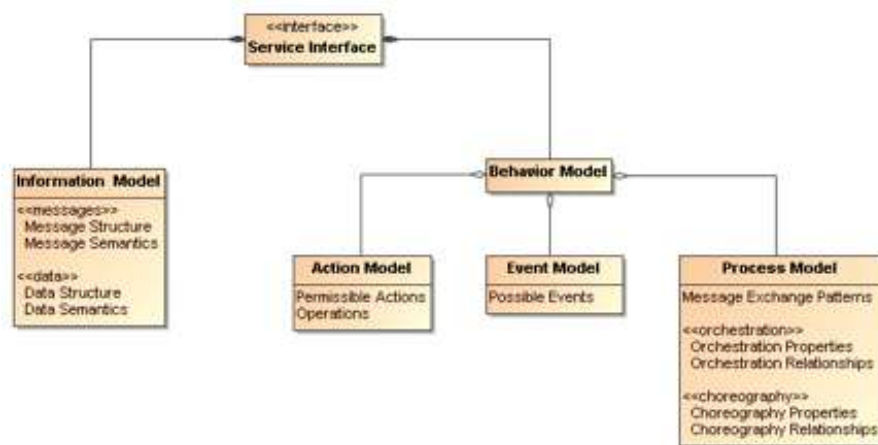


Figure 18-47 Service Interface

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

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

4.1.1.3.2 Service Reachability

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

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

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

4.1.1.3.3 Service Functionality

While the service interface and service reachability are concerned with the mechanics of using a service, service functionality and performance metrics (discussed in Section 4.1.1.3.4) describe what can be expected when interacting with a service. Service Functionality, shown in Figure 17Figure 16 as part of the overall Service Description model and extended in Figure 19Figure 18, is an unambiguous expression of service function(s) and the real world effects of invoking the function. The Functions likely represent business activities in some domain that produce the desired real world effects.

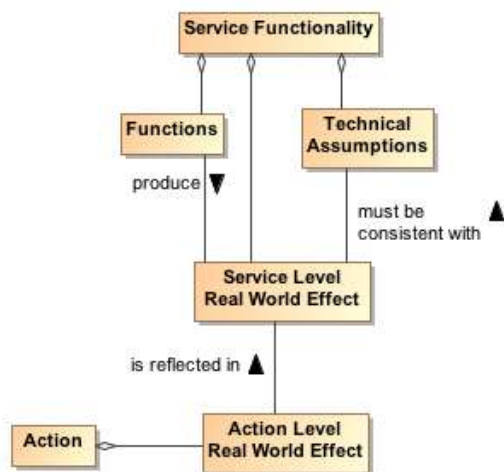


Figure 19-48 Service Functionality

The Service Functionality may also be constrained by Technical Assumptions that underlie the effects that can result. Technical assumptions 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 assumptions are likely related to the underlying capability accessed by the service. In any case, the real world effects must be consistent with the Technical Assumptions.

Comment [PFB44]: How can an assumption lead to a constraint?

In Figure 17-46 and Figure 19-48, we specifically refer to Service Level and Action Level real world effects.

Service Level Real World Effect

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

Action Level Real World Effect

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

Comment [PFB45]: Delete? RWE is not a change of information only shared state

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

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

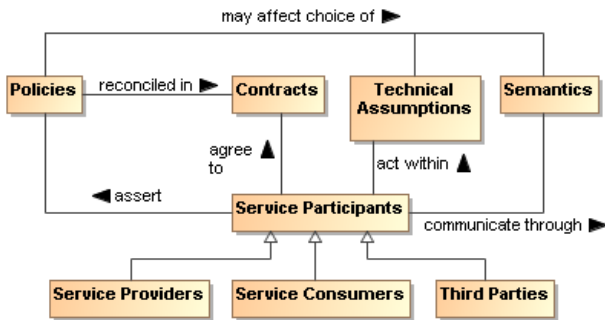
1739 4.1.1.3.4 Policies and Contracts, Metrics, and Compliance Records

1740 Policies prescribe the conditions and constraints for interacting with a service and
1741 impact the willingness to continue visibility with the other participants. Whereas
1742 technical assumptions are statements of “physical” fact, policies are subjective
1743 assertions made by the service provider (sometimes as passed on from higher
1744 authorities).

1745 The service description provides a central location for identifying what policies have
1746 been asserted by the service provider. The specific representation of the policy, e.g. in
1747 some formal policy language, is likely done outside of the service description and the
1748 service description would reference the normative definition of the policy.

1749 Policies may also be asserted by other service participants, as illustrated by the model
1750 shown in Figure 20Figure 19. Policies that are generally applicable to any interaction
1751 with the service are likely to be asserted by the service provider and included in the
1752 Policies and Contracts section of the service description. Conversely, policies that are
1753 asserted by specific consumers or consumer communities would likely be identified as
1754 part of a description’s Annotations from 3rd parties (see Section 4.1.1.4) because
1755 these would be specific to those parties and not a general aspect of the service being
1756 described.

Comment [PFB46]: What does this actually mean?



1757
1758 Figure 20Figure 19 Model for Policies and Contracts as related to Service Participants

1759 In Figure 17Figure 16 and Figure 21Figure 20, we specifically refer to Service Level
1760 Interaction Policies. In a similar manner to that discussed for Service Level vs. Action
1761 Level Real World Effects in Section 4.1.1.3.3, individual Actions may have associated
1762 policies stating conditions for performing the action, but these must be reflected in and
1763 be consistent with the policies made visible at the service level and thus the description
1764 of the service as a whole. The relationship to Action Level Policies and the implications
1765 on defining the scope of a service are discussed in Section 4.1.2.1.




Figure 2120 Action-Level and Service-Level Policies

As noted in Figure 20Figure 19, the policies asserted may affect the allowable Technical Assumptions that can be embodied in services or their underlying capabilities and may affect the semantics that can be used. For example of the former, there may be a policy that specifies the surge capacity to be accommodated by a server, and a service that designs for a smaller capacity would not be appropriate to use. For the latter, a policy may require that only services using a community-sponsored vocabulary can be used. Contracts are agreements among the service participants. The contract may reconcile inconsistent policies asserted by the participants or may specify details of the interaction. Service level agreements (SLAs) are one commonly used category of contracts.

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

The definition and later enforcement of policies and contracts are predicated on the existence of metrics; the relationships among the relevant concepts are shown in the model in Figure 22Figure 24. Performance Metrics identify quantities that characterize the speed and quality of realizing the real world effects produced using the SOA service; in addition, policies and contracts may depend on nonperformance metrics, such as whether a license is in place to use the service. Some of these metrics reflect the underlying capability, e.g. a SOA service cannot respond in two seconds if the underlying capability is expected to take five seconds to do its processing; some metrics reflect the implementation of the SOA service, e.g. what level of caching is present to minimize data access requests across the network.

Comment [PFB47]: Idem as for policy

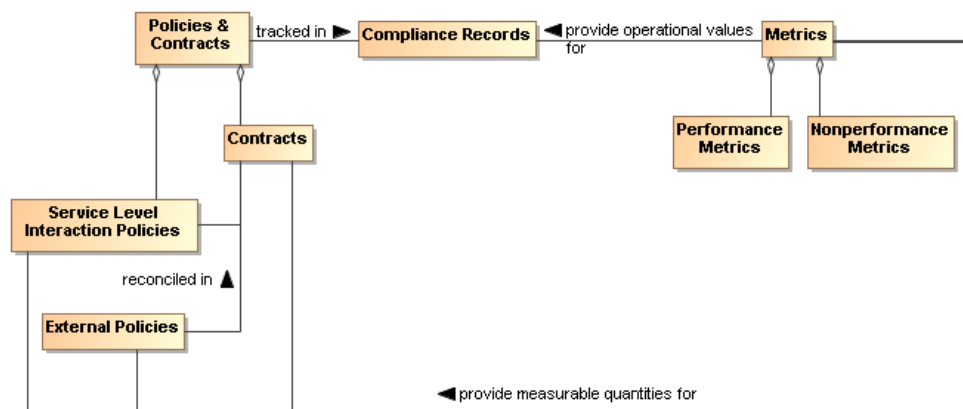


Figure 2224 Policies and Contracts, Metrics, and Compliance Records

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

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

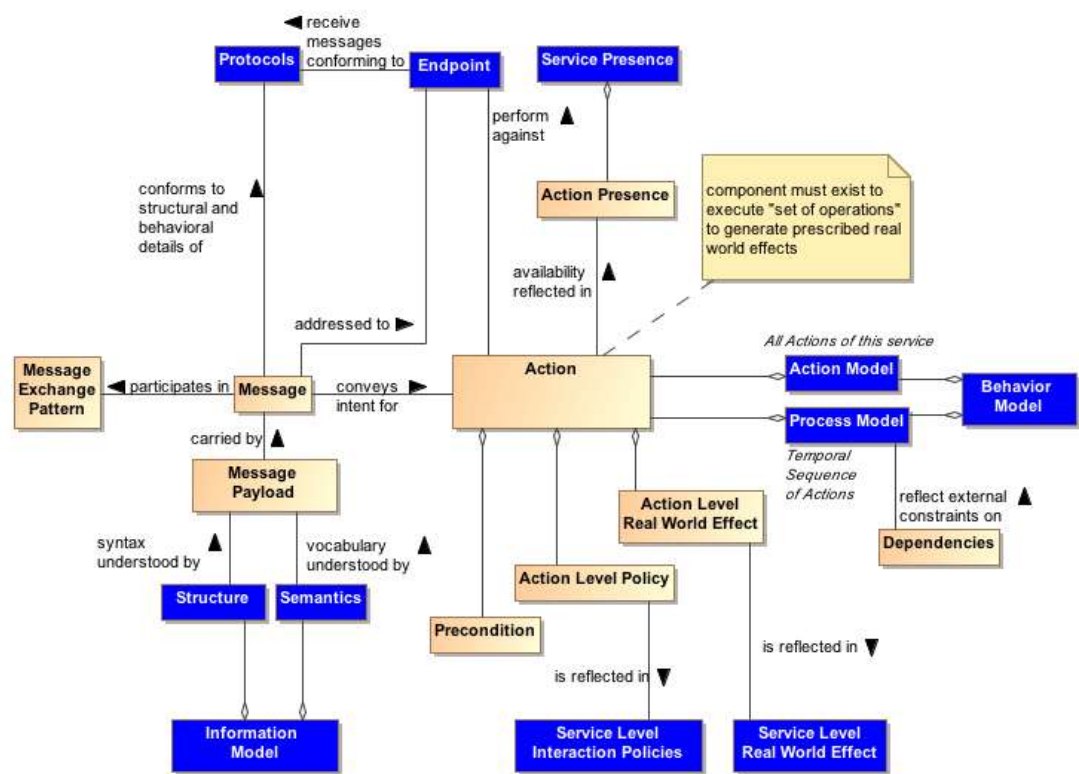
Note, even when policies relate the perspective of a single participant, policy compliance can be measured and policies may be enforceable without contractual agreement with other participants. This should be reflected in the policy, contract, and compliance record information maintained in the service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

4.1.2 Use Of Service Description

4.1.2.1 Service Description in support of Service Interaction

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

1823 groundwork for what can occur, whereas service interaction defines the specifics
1824 through which occurrences are realized.



1825
1826 Figure 23 Relationship Between Action and Service Description Components

1827 Figure 23 combines the models in the subsections of Section 4.1.1 to
1828 concisely relate action and the relevant components of Service Description. The
1829 purpose of Figure 23 is to demonstrate that the components of service
1830 description go beyond arbitrary documentation and form the critical set of information
1831 needed to define the what and how of action. In Figure 23, the leaf nodes from
1832 Figure 17 are shown in blue.

1833 action is invoked via a Message where the structure and behavioral details of the
1834 message conform to an identified Protocol and is directed to the address of the
1835 identified endpoint, and the message payload conforms to the service Information
1836 Model.

1837 The availability of an action is reflected in the Action Presence and each Action
1838 Presence contributes to the overall Service Presence; this is discussed further in
1839 Section 4.2.2.3. Each action has its own endpoint and also its own protocols associated

1840 with the endpoint¹² and to what extent, e.g. current or average availability, there is
1841 presence for the action through that endpoint. The endpoint and service presence are
1842 also part of the service description.

1843 An action may have preconditions where a Precondition is something that needs to be
1844 in place before an action can occur, e.g. confirmation of a precursor action. Whether
1845 preconditions are satisfied is evaluated when someone tries to perform the action and
1846 not before. Presence for an action means someone can initiate it and is independent of
1847 whether the preconditions are satisfied. However, the successful completion of the
1848 action may depend on whether its preconditions were satisfied.

1849 Analogous to the relationship between actions and preconditions, the Process Model
1850 may imply Dependencies for succeeding steps in a process, e.g. that a previous step
1851 has successfully completed, or may be isolated to a given step. An example of the
1852 latter would be a dependency that the host server has scheduled maintenance and
1853 access attempts at these times would fail. Dependencies related to the process model
1854 do not affect the presence of a service although these may affect whether the business
1855 function successfully completes.

1856 The conditions under which an action can be invoked may depend on policies
1857 associated with the action. The Action Level Policies MUST be reflected in the Service
1858 Level Interaction Policies because such policies may be critical to determining whether
1859 the conditions for use of the service are consistent with the policies asserted by the
1860 service consumer. The service level interaction policies are included in the service
1861 description.

1862 Similarly, the result of invoking an action is one or more real world effects, and the
1863 Action Level Real World Effects MUST be reflected in the Service Level Real World
1864 Effect included in the service description. The unambiguous expression of action level
1865 policies and real world effects as service counterparts is necessary to adequately
1866 understand what constitutes the service interaction.

1867 An adequate service description MUST provide a consumer with information needed to
1868 determine if the service policies and the (business) functions and service-level real
1869 world effects are of interest and there is nothing in the technical assumptions that
1870 preclude use of the service.

1871 Note at this level, the business functions are not concerned with the action or process
1872 models. These models are detailed separately.

1873 The service description is not intended to be isolated documentation but rather an
1874 integral part of service use. Changes in service description SHOULD immediately be
1875 made known to consumers and potential consumers.

1876 4.1.2.1.1 Description and Invoking Actions Against a Service

1877 At this point, let us assume the descriptions were sufficient to establish willingness; see
1878 Section 4.2.2.2. [Figure 23](#) indicates the service endpoint establishes where to

¹² 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.

1879 actually carry out the interaction. This is where we start considering the action and
1880 process models.

1881 The action model identifies the multiple actions a user can perform against a service
1882 and the user would perform these in the context of the process model as specified or
1883 referenced under the Service Interface portion of Service Description. For a given
1884 business function, there is a corresponding process model, where any process model
1885 may involve multiple actions. From the above discussion of model elements of
1886 description we may conclude (1) actions have reachability information, including
1887 endpoint and presence, (2) presence of service is some aggregation of presence of its
1888 actions, (3) action preconditions and service dependencies do not affect presence
1889 although these may affect successful completion.

1890 Having established visibility, the interaction can proceed. Given a business function, the
1891 consumer knows what will be accomplished (the service functionality), the conditions
1892 under which interaction will proceed (service policies and contracts), and the process
1893 that must be followed (the process model). The remaining question is how does the
1894 description information for structure and semantics enable interaction.

1895 We have established the importance of the process model in identifying relevant actions
1896 and their sequence. Interaction proceeds through messages and thus it is the syntax
1897 and semantics of the messages with which we are here concerned. A common
1898 approach is to define the structure and semantics that can appear as part of a message;
1899 then assemble the pieces into messages; and, associate messages with actions.

1900 Actions make use of structure and semantics as defined in the information model to
1901 describe its legal messages.

1902 The process model identifies actions to be performed against a service and the
1903 sequence for performing the actions. For a given action, the Reachability portion of
1904 description indicates the protocol bindings that are available, the endpoint
1905 corresponding to a binding, and whether there is presence at that endpoint. The
1906 interaction with actions is through messages that conform to the structure and
1907 semantics defined in the information model and the message sequence conforming to
1908 the action's identified MEP. The result is some portion of the **real world effect that will**
1909 **need to must** be assessed and/or processed (e.g. if an error exists, that part that covers
1910 the error processing would be invoked).

Comment [PFB48]: Shared state?

1911 4.1.2.1.2 The Question of Multiple Business Functions

1912 Action level effects and policies MUST be reflected at the service level for service
1913 description to support visibility.

1914 It is assumed that a SOA service represents an identifiable business function to which
1915 policies can be applied and from which desired business effects can be obtained. While
1916 contemporary discussions of SOA services and supporting standards do not constrain
1917 what actions or combinations of actions can or should be defined for a service, **this**
1918 **Reference Architecture the SOA-RAF** considers the implications of service description in
1919 defining the range of actions appropriate for an individual SOA service.

1920 Consider the situation if a given SOA service is the container for multiple independent
1921 (but loosely related) business functions. These are not multiple effects from a single
1922 function but multiple functions with potentially different sets of effects for each function.

1923 A service can have multiple actions a user may perform against it, and this does not
1924 change with multiple business functions. As an individual business function corresponds
1925 to a process model, so multiple business functions imply multiple process models. The
1926 same action may be used in multiple process models but the aggregated service
1927 presence would be specific to each business function because the components being
1928 aggregated ~~will likely~~ may be different between process models. In summary, for a
1929 service with multiple business functions, each function has (1) its own process model
1930 and dependencies, (2) its own aggregated presence, and (3) possibly its own list of
1931 policies and real world effects.

1932 A common variation on this theme is for a single service to have multiple endpoints for
1933 different levels of quality of service (QoS). Different QoS imply separate statements of
1934 policy, separate endpoints, possibly separate dependencies, and so on. One could say
1935 the QoS variation does not require this because there can be a single QoS policy that
1936 encompasses the variations. and all other aspects of the service would be the same
1937 except for the endpoint used for each QoS. However, the different aspects of policy at
1938 the service level would need to be mapped to endpoints, and this introduces an
1939 undesirable level of coupling across the elements of description. In addition, it is
1940 obvious that description at the service level can become very complicated if the number
1941 of combinations is allowed to grow.

1942 One could imagine a service description that is basically a container for action
1943 descriptions, where each action description is self contained; however, this would lead
1944 to duplication of description components across actions. If common description
1945 components are factored, this either is limited to components common across all
1946 actions or requires complicated tagging to capture the components that often but do not
1947 universally apply.

1948 If a provider cannot describe a service as a whole but must describe every action, this
1949 leads to the situation where it may be extremely difficult to construct a clear and concise
1950 service description that can effectively support discovery and use without tedious logic
1951 to process the description and assemble the available permutations. In effect, if
1952 adequate description of an action begins to look like description of a service, it may be
1953 best to have it as a separate service.

1954 Recall, more than one service can access the same underlying capability, and this is
1955 appropriate if a different real world effect is to be exposed. Along these lines, one can
1956 argue that different QoS are different services because getting a response in one
1957 minute rather than one hour is more than a QoS difference; it is a fundamental
1958 difference in the business function being provided.

1959 As a best practice, a criteria for whether a service is appropriately scoped may be the
1960 ease or difficulty in creating an unambiguous service description. A consequence of
1961 having tightly-scoped services is there will be a greater reliance on combining services,
1962 i.e. more fundamental business functions, to create more advanced business functions.
1963 This is consistent with the principles of service oriented architecture and is the basic
1964 position of the Reference Architecture, although not an absolute requirement.
1965 Combining services increases the reliance on understanding and implementing the
1966 concepts of orchestration, choreography, and other approaches yet to be developed;
1967 these are discussed in more detail in section 4.4 Interacting with Services.

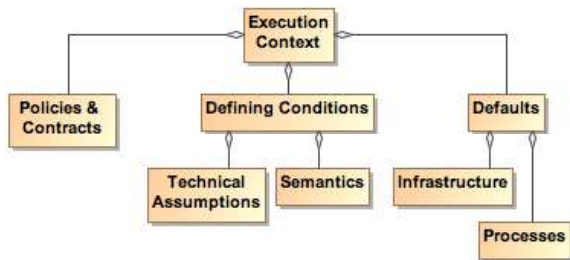
Comment [PFB49]: Assumption?
Fact? Prescription?

1968 **4.1.2.1.3 Service Description, Execution Context, and Service Interaction**

1969 The service description MUST provide sufficient information to support service visibility,
1970 including the willingness of service participants to interact. However, the corresponding
1971 descriptions for providers and consumers may both contain policies, technical
1972 assumptions, constraints on semantics, and other technical and procedural conditions
1973 that must be aligned to define the terms of willingness. The agreements which
1974 encapsulate the necessary alignment form the basis upon which interactions may
1975 proceed – in the Reference Model, this collection of agreements and the necessary
1976 environmental support establish the execution context.

1977 To illustrate the concept of the execution context, consider a Web-based system for
1978 timecard entry. For an employee onsite at an employer facility, the execution context
1979 requires a computer connected to the local network and the employee must enter their
1980 network ID and password. Relevant policies include that the employee must maintain
1981 the most recent anti-virus software and virus definitions for any computer connected to
1982 the network.

1983 For the same employee connecting from offsite, the execution context specifies the
1984 need for a computer with installed VPN software and a security token to negotiate the
1985 VPN connection. The execution context also includes proxy settings as needed to
1986 connect to the offsite network. The employee must still comply with the requirements for
1987 onsite computers and access, but the offsite execution context includes additional items
1988 before the employee can access the same underlying capability and realize the same
1989 real world effects, i.e. the timecard entries.



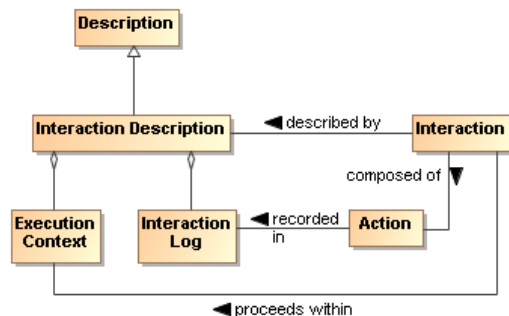
1990
1991 *Figure 2423 Execution Context*

1992 [Figure 24](#)[Figure 23](#) shows a few broad categories found in execution context. These are
1993 not meant to be comprehensive. Other items may need to be included to collect a
1994 sufficient description of the interaction conditions. Any other items not explicitly noted in
1995 the model but needed to set the environment SHOULD be included in the execution
1996 context.

1997 While the execution context captures the conditions under which interaction can occur,
1998 it does not capture the specific service invocations that do occur in a specific interaction.
1999 A service interaction as modeled in [Figure 23](#)[Figure 22](#) introduces the concept of an
2000 Interaction Description which is composed of both the Execution Context and an
2001 Interaction Log. The execution context specifies the set of conditions under which the
2002 interaction occurs and the interaction log captures the sequence of service interactions
2003 that occur within the execution context. This sequence should follow the Process Model

2004 but can include details beyond those specified there. For example, the Process Model
 2005 may specify an action that results in identifying a data source, and the identified source
 2006 is used in a subsequent action. The Interaction Log would record the specific data
 2007 source used.

2008 The execution context can be thought of as the container in which the interaction occurs
 2009 and the interaction log captures what happens inside the container. This combination is
 2010 needed to support auditability and repeatability of the interactions.



2011
 2012 | Figure 2524 Interaction Description

2013 SOA allows flexibility to accomplish repeatability or reusability. One benefit of this is that
 2014 a service can be updated without disrupting the user experience of the service. So,
 2015 Google can improve their ranking algorithm without notifying the user about the details
 2016 of the update.

2017 However, it may also be vital for the consumer to be able to recreate past results or to
 2018 generate consistent results in the future, and information such as what conditions, which
 2019 services, and which versions of those services are used is indispensable in retracing
 2020 one's path. The interaction log is a critical part of the resulting real world effects
 2021 because it defines how the effects were generated and possibly the meaning of
 2022 observed effects. This increases in importance as dynamic composability becomes
 2023 more feasible. In essence, a result has limited value if one does not know how it was
 2024 generated.

2025 The interaction log SHOULD be a detailed trace for a specific interaction, and its reuse
 2026 is limited to duplicating that interaction. An execution context can act as a template for
 2027 identical or similar interactions. Any given execution context MAY define the conditions
 2028 of future interactions.

2029 Such uses of execution context imply (1) a standardized format for capturing execution
 2030 context and (2) a subclass of general description could be defined to support visibility of
 2031 saved execution contexts. The specifics of the relevant formats and descriptions are
 2032 beyond the scope of this [Reference Architecture document](#).

2033 A service description is unlikely to track interaction descriptions or the constituent
 2034 execution contexts or interaction logs that include mention of the service. However, as
 2035 appropriate, linking to specific instances of either of these could be done through
 2036 associated annotations.

2037 4.1.3 Relationship to Other Description Models

2038 | While the representation shown in [Figure 16](#)~~Figure 15~~ is derived from considerations
2039 | related to service description, it is acknowledged that other metadata standards are
2040 | relevant and should, as possible, be incorporated into this work. Two standards of
2041 | particular relevance are the Dublin Core Metadata Initiative (DCMI) and ISO 11179,
2042 | especially Part 5.

2043 | When the service description (or even the general description class) is considered as
2044 | the DCMI “resource”, [Figure 16](#)~~Figure 15~~ aligns nicely with the DCMI resource model.
2045 | While some differences exist, these are mostly in areas where DCMI goes into detail
2046 | that is considered beyond the scope of the current Reference Architecture. For
2047 | example, DCMI defines classes of “shared semantics” whereas this Reference
2048 | Architecture [Framework](#) considers that an identification of relevant semantic models is
2049 | sufficient. Likewise, the DCMI “description model” goes into the details of possible
2050 | syntax encodings whereas for the Reference Architecture [Framework](#) it is sufficient to
2051 | identify the relevant formats.

2052 | With respect to ISO 11179 Part 5, the metadata fields defined in that reference may be
2053 | used without prejudice as the properties in [Figure 16](#)~~Figure 15~~. Additionally, other
2054 | defined metadata sets may be used by the service provider if the other sets are
2055 | considered more appropriate, i.e. it is fundamental to this [Reference Architecture](#) to
2056 | identify the need and the means to make vocabulary declarations explicit but it is
2057 | beyond the scope to specify which vocabularies are to be used. In addition, the
2058 | identification of domain of the properties and range of the values has not been included
2059 | in the current Reference Architecture discussion, but the text of ISO 11179 Part 5 can
2060 | be used consistently with the model prescribed in this document.

2061 | Description as defined [in the context of this Reference Architecture here](#) considers a
2062 | wide range of applicability and support of the principles of service oriented architecture.
2063 | Other metadata models can be used in concert with the model presented here because
2064 | most of these focus on a finer level of detail that is outside the present scope, and so
2065 | provide a level of implementation guidance that can be applied as appropriate.

2066 4.1.4 Architectural Implications

2067 | The description of service description indicates numerous architectural implications on
2068 | the SOA ecosystem:

- 2069 | • ~~Description will~~[It](#) changes over time and its contents will reflect changing needs and
2070 | context. This requires the existence of:
 - 2071 | ○ mechanisms to support the storage, referencing, and access to normative
2072 | definitions of one or more versioning schemes that may be applied to identify
2073 | different aggregations of descriptive information, where the different schemes
2074 | may be versions of a versioning scheme itself;
 - 2075 | ○ configuration management mechanisms to capture the contents of the each
2076 | aggregation and apply a unique identifier in a manner consistent with an
2077 | identified versioning scheme;
 - 2078 | ○ one or more mechanisms to support the storage, referencing, and access to
2079 | conversion relationships between versioning schemes, and the mechanisms
2080 | to carry out such conversions.

- 2081 • Description makes use of defined semantics, where the semantics may be used for
2082 categorization or providing other property and value information for description
2083 classes. This requires the existence of:
- 2084 ○ semantic models that provide normative descriptions of the utilized terms,
2085 where the models may range from a simple dictionary of terms to an ontology
2086 showing complex relationships and capable of supporting enhanced
2087 reasoning;
 - 2088 ○ mechanisms to support the storage, referencing, and access to these
2089 semantic models;
 - 2090 ○ configuration management mechanisms to capture the normative description
2091 of each semantic model and to apply a unique identifier in a manner
2092 consistent with an identified versioning scheme;
 - 2093 ○ one or more mechanisms to support the storage, referencing, and access to
2094 conversion relationships between semantic models, and the mechanisms to
2095 carry out such conversions.
- 2096 • Descriptions include reference to policies defining conditions of use and optionally
2097 contracts representing agreement on policies and other conditions. This requires the
2098 existence of (as also enumerated under governance):
- 2099 ○ descriptions to enable the policy modules to be visible, where the description
2100 includes a unique identifier for the policy and a sufficient, and preferably a
2101 machine processible, representation of the meaning of terms used to describe
2102 the policy, its functions, and its effects;
 - 2103 ○ one or more discovery mechanisms that enable searching for policies that
2104 best meet the search criteria specified by the service participant; where the
2105 discovery mechanism ~~will have~~has access to the individual policy
2106 descriptions, possibly through some repository mechanism;
 - 2107 ○ accessible storage of policies and policy descriptions, so service participants
2108 can access, examine, and use the policies as defined.
- 2109 • Descriptions include references to metrics which describe the operational
2110 characteristics of the subjects being described. This requires the existence of (as
2111 partially enumerated under governance):
- 2112 ○ the infrastructure monitoring and reporting information on SOA resources;
 - 2113 ○ possible interface requirements to make accessible metrics information
2114 generated or most easily accessed by the service itself;
 - 2115 ○ mechanisms to catalog and enable discovery of which metrics are available
2116 for a described resources and information on how these metrics can be
2117 accessed;
 - 2118 ○ mechanisms to catalog and enable discovery of compliance records
2119 associated with policies and contracts that are based on these metrics.
- 2120 • Descriptions of the interactions are important for enabling auditability and
2121 repeatability, thereby establishing a context for results and support for understanding
2122 observed change in performance or results. This requires the existence of:
- 2123 ○ one or more mechanisms to capture, describe, store, discover, and retrieve
2124 interaction logs, execution contexts, and the combined interaction
2125 descriptions;

- 2126 ○ one or more mechanisms for attaching to any results the means to identify
2127 and retrieve the interaction description under which the results were
2128 generated.
- 2129 • Descriptions may capture very focused information subsets or can be an aggregate
2130 of numerous component descriptions. Service description is an example of an likely
2131 aggregate for which manual maintenance of all aspects the whole would not be
2132 feasible. This requires the existence of:
- 2133 ○ tools to facilitate identifying description elements that are to be aggregated to
2134 assemble the composite description;
- 2135 ○ tools to facilitate identifying the sources of information to associate with the
2136 description elements;
- 2137 ○ tools to collect the identified description elements and their associated
2138 sources into a standard, referenceable format that can support general
2139 access and understanding;
- 2140 ○ tools to automatically update the composite description as the component
2141 sources change, and to consistently apply versioning schemes to identify the
2142 new description contents and the type and significance of change that
2143 occurred.
- 2144 • Descriptions provide up-to-date information on what a resource is, the conditions for
2145 interacting with the resource, and the results of such interactions. As such, the
2146 description is the source of vital information in establishing willingness to interact
2147 with a resource, reachability to make interaction possible, and compliance with
2148 relevant conditions of use. This requires the existence of:
- 2149 ○ one or more discovery mechanisms that enable searching for described
2150 resources that best meet the criteria specified by a service participant, where
2151 the discovery mechanism will have has access to individual descriptions,
2152 possibly through some repository mechanism;
- 2153 ○ tools to appropriately track users of the descriptions and notify them when a
2154 new version of the description is available.

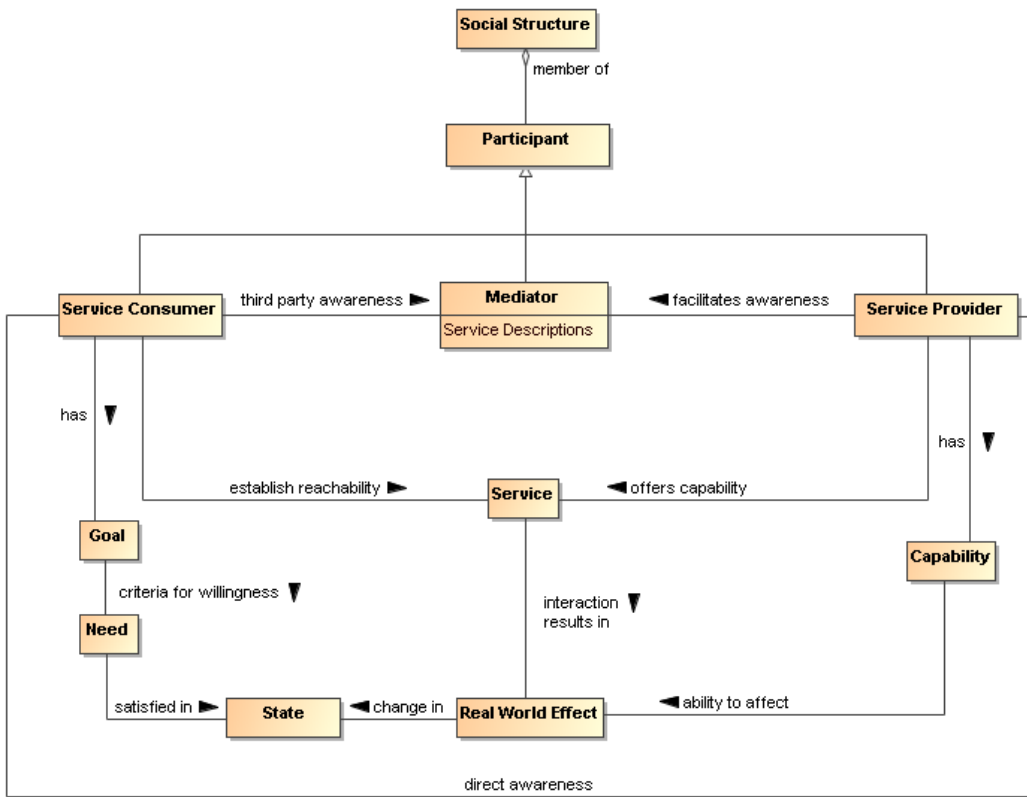
2155 4.2 Service Visibility Model

2156 One of the key requirements for participants interacting with each other in the context of
2157 a SOA is achieving visibility: before services can interoperate, the participants have to
2158 be visible to each other using whatever means are appropriate. The Reference Model
2159 analyzes visibility in terms of awareness, willingness, and reachability. In this section,
2160 we explore how visibility may be achieved.

2161 4.2.1 Visibility to Business

2162 The relationship of visibility to the SOA ecosystem encompasses both human social
2163 structures and automated IT mechanisms. Figure 26Figure-25 depicts a business
2164 setting that is a basis for visibility as related to the social structure Model in the Service
2165 Ecosystem ViewParticipation in a SOA Ecosystem view (see Section **Error! Reference**
2166 **source not found.**). Service consumers and service providers may have direct
2167 awareness or mediated awareness where mediated awareness is achieved through
2168 some third party. A consumer's willingness to use a service is reflected by the
2169 consumer's presumption of satisfying goals and needs based on the description of the

2170 service. Service providers offer capabilities that have real world effects that result in a
2171 change in state of the consumer. Reachability of the service by the consumer leads to
2172 interactions that change the state of the consumer. The consumer can measure the
2173 change of state to determine if the claims made by description and the real world effects
2174 of consuming the service meet the consumer's needs.
2175



2176
2177 Figure 2625 Visibility to Business

2178 Visibility and interoperability in a SOA ecosystem requires more than location and
2179 interface information. A meta-model for this broader view of visibility is depicted in
2180 Section 4.1. In addition to providing improved awareness of service capabilities through
2181 description of information such as reachability, behavior models, information models,
2182 functionality, and metrics, the service description may contain policies valuable for
2183 determination of willingness to interact.

2184 A mediator of service descriptions may provide event notifications to both consumers
2185 and providers about information relating to service descriptions. One example of this
2186 capability is a publish/subscribe model where the mediator allows consumers to
2187 subscribe to service description version changes made by the provider. Likewise, the
2188 mediator may provide notifications to the provider of consumers that have subscribed to
2189 service description updates.

2190 Another important business capability in a SOA environment is the ability to narrow
2191 visibility to trusted members within a social structure. Mediators for awareness may
2192 provide policy based access to service descriptions allowing for the dynamic formation
2193 of awareness between trusted members.

2194 4.2.2 Visibility

2195 Attaining visibility is described in terms of steps that lead to visibility. While there can be
2196 many contexts for visibility within a single social structure, the same general steps can
2197 be applied to each of the contexts to accomplish visibility.

2198 Attaining SOA visibility requires

- 2199 • service description creation and maintenance,
- 2200 • processes and mechanisms for achieving awareness of and accessing descriptions,
- 2201 • processes and mechanisms for establishing willingness of participants,
- 2202 • processes and mechanisms to determine reachability.

2203 Visibility may occur in stages, i.e. a participant can become aware enough to look or ask
2204 for further description, and with this description, the participant can decide on
2205 willingness, possibly requiring additional description. For example, if a potential
2206 consumer has a need for a tree cutting (business) service, the consumer can use a web
2207 search engine to find web sites of providers. The web search engine (a mediator) gives
2208 the consumer links to relevant web pages and the consumer can access those
2209 descriptions. For those prospective providers that satisfy the consumer's criteria, the
2210 consumer's willingness to interact increases. The consumer ~~likely-may~~ contacts several
2211 tree services to get detailed cost information (or arrange for an estimate) and may ask
2212 for references (further description). ~~Likely, t~~ The consumer ~~will-is likely to~~ establish full
2213 visibility and proceed with ~~the~~ interaction with ~~a-the~~ tree service who mutually
2214 establishes visibility.

2215 4.2.2.1 Awareness

2216 A service participant is aware of another participant if it has access to a description of
2217 that participant with sufficient completeness to establish the other requirements of
2218 visibility.

2219 Awareness is inherently a function of a participant; awareness can be established
2220 without any action on the part of the target participant other than the target providing
2221 appropriate descriptions. Awareness is often discussed in terms of consumer
2222 awareness of providers but the concepts are equally valid for provider awareness of
2223 consumers.

2224 Awareness can be decomposed into the creation of descriptions, making them
2225 available, and discovering the descriptions. Discovery can be initiated or it can be by
2226 notification. Initiated discovery for business may require formalization of the required
2227 capabilities and resources to achieve business goals.

2228 Achieving awareness in a SOA can range from word of mouth to formal service
2229 descriptions in a standards-based registry-repository. Some other examples of
2230 achieving awareness in a SOA are the use of a web page containing description
2231 information, email notifications of descriptions, and document based descriptions.

2232 A mediator as discussed for awareness is a third party participant that provides
2233 awareness to one or more consumers of one or more services. Direct awareness is
2234 awareness between a consumer and provider without the use of a third party.

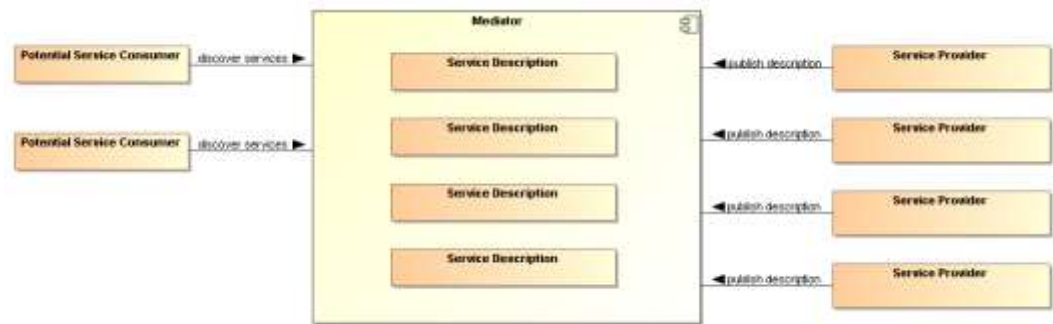
2235 Direct awareness may be the result of having previously established an execution
2236 context, or direct awareness may include determining the presence of services and then
2237 querying the service directly for description. As an example, a priori visibility of some
2238 sensor device may provide the means for interaction or a query for standardized sensor
2239 device metadata may be broadcast to multiple locations. If acknowledged, the service
2240 interface for the device may directly provide description to a consumer so the consumer
2241 can determine willingness to interact.

2242 The same medium for awareness may be direct in one context and may be mediated in
2243 another context. For example, a service provider may maintain a web site with links to
2244 the provider's descriptions of services giving the consumers direct awareness to the
2245 provider's services. Alternatively, a community may maintain a mediated web site with
2246 links to various provider descriptions of services for any number of consumers. More
2247 than one mediator may be involved, as different mediators may specialize in different
2248 mediation functions.

2249 Descriptions may be formal or informal. Section 4.1, provides a comprehensive model
2250 for service description that can be applied to formal registry/repositories used to
2251 mediate visibility. Using consistent description taxonomies and standards based
2252 mediated awareness helps provide more effective awareness.

2253 **4.2.2.1.1 Mediated Awareness**

2254 Mediated awareness promotes loose coupling by keeping the consumers and services
2255 from explicitly referring to each other and the descriptions. Mediation lets interaction
2256 vary independently. Rather than all potential service consumers being informed on a
2257 continual basis about all services, there is a known or agreed upon facility or location
2258 that houses the service description.



2259
2260 Figure 2726 Mediated Service Awareness

2261 In Figure 27Figure 26, the potential service consumers perform queries or are notified in
2262 order to locate those services that satisfy their needs. As an example, the telephone
2263 book is a mediated registry where individuals perform manual searches to locate
2264 services (i.e. the yellow pages). The telephone book is also a mediated registry for
2265 solicitors to find and notify potential customers (i.e. the white pages).

2266 In mediated service awareness for large and dynamic numbers of service consumers
2267 and service providers, the benefits typically far outweigh the management issues
2268 associated with it. Some of the benefits of mediated service awareness are

- 2269 • Potential service consumers have a known location for searching thereby eliminating
2270 needless and random searches
- 2271 • Typically a consortium of interested parties (or a sufficiently large corporation) signs
2272 up to host the mediation facility
- 2273 • Standardized tools and methods can be developed and promulgated to promote
2274 interoperability and ease of use.

2275 However, mediated awareness can have some risks associated with it:

- 2276 • A single point of failure. If the central mediation service fails then a ~~potentially~~ large
2277 number of service providers and consumers ~~will be~~ potentially adversely affected.
- 2278 • A single point of control. If the central mediation service is owned by, or controlled
2279 by, someone other than the service consumers and/or providers then the latter may
2280 be put at a competitive disadvantage based on policies of the discovery provider.

2281 A common mechanism for mediated awareness is a registry-repository. The registry
2282 stores links or pointers to service description artifacts. The repository in this example is
2283 the storage location for the service description artifacts. Service descriptions can be
2284 pushed (publish/subscribe for example) or pulled from the register-repository mediator.

2285 The registry is like a card catalog at the library and a repository is like the shelves for
2286 the books. Standardized metadata describing repository content can be stored as
2287 registry objects in a registry and any type of content can be stored as repository items in
2288 a repository. The registry may be constructed such that description items stored within
2289 the mediation facility repository ~~will have~~ has intrinsic links in the registry while
2290 description items stored outside the mediation facility ~~will~~ have extrinsic links in the
2291 registry.

2292 When independent but like SOA IT mechanisms interoperate with one another, the IT
2293 mechanisms may be referred to as federated.

2294 **4.2.2.1.2 Awareness in Complex Social Structures**

2295 Awareness applies to one or more communities within one or more social structures
2296 where a community consists of at least one description provider and one description
2297 consumer. These communities may be part of the same social structure or be part of
2298 different ones.

2299 In ~~Figure 28~~ Figure 27, awareness can be within a single community, multiple
2300 communities, or all communities in the social structure. The social structure can
2301 encourage or restrict awareness through its policies, and these policies can affect
2302 participant willingness. The information about policies should be incorporated in the
2303 relevant descriptions. The social structure also governs the conditions for establishing
2304 contracts, the results of which will be reflected in the execution context if interaction is to
2305 proceed.

Comment [PFB50]: MAY?

Comment [PFB51]: MUST?

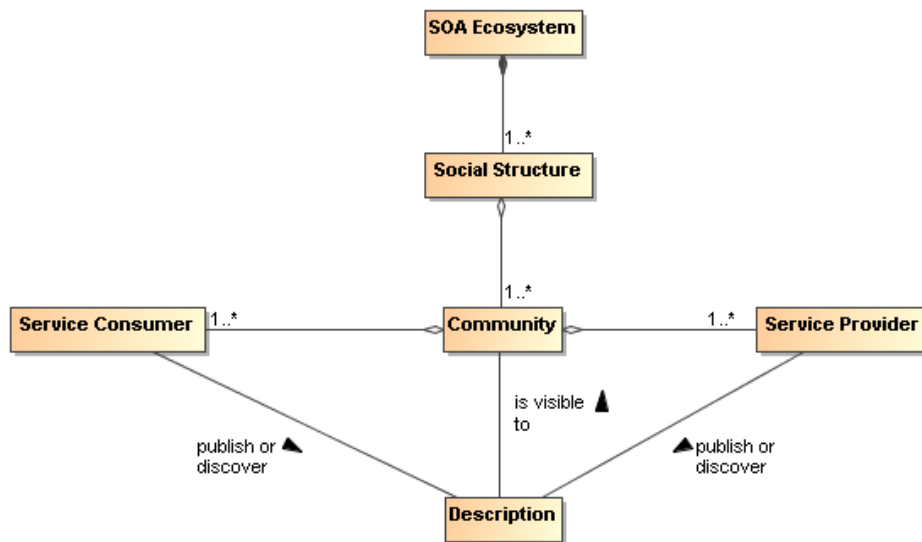


Figure 2827 Awareness in a SOA Ecosystem

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


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

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

4.2.2.2 Willingness

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

2329



2330

2331

2332

2333

2334

2335

2336

2337

2338

2339

2340

2341

2342

2343

2344

2345

2346

2347

2348

2349

2350

2351

2352

2353

2354

2355

2356

2357

2358

Figure 29: Business, Description and Willingness

Figure 29 relates elements of the [Service Ecosystem View](#) [Participation in a SOA Ecosystem view](#), and elements from the Service Description Model to willingness. By having a willingness to interact within a particular social structure, the social structure provides the participant access to capabilities based on conditions the social structure finds appropriate for its context. The participant can use these capabilities to satisfy goals and objectives as specified by the participant's needs.

In Figure 29, information used to determine willingness is defined by Description. Information referenced by Description may come from many sources. For example, a mediator for descriptions may provide 3rd party annotations for reputation. Another source for reputation may be a participant's own history of interactions with another participant.

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

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

4.2.2.3 Reachability

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

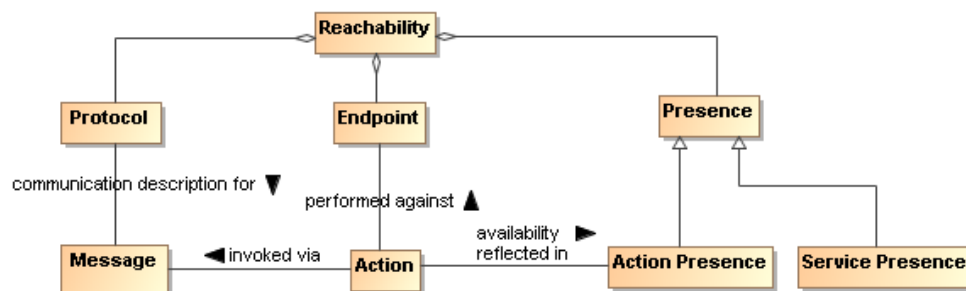


Figure 3029 Service Reachability

Endpoint

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

Protocol

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

Presence

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

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

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

4.2.3 Architectural Implications

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

- Mechanisms providing support for awareness **will likely** have the following minimum capabilities:
 - creation of Description, preferably conforming to a standard Description format and structure;

- 2391 ○ publishing of Description directly to a consumer or through a third party
- 2392 mediator;
- 2393 ○ discovery of Description, preferably conforming to a standard for Description
- 2394 discovery;
- 2395 ○ notification of Description updates or notification of the addition of new and
- 2396 relevant Descriptions;
- 2397 ○ classification of Description elements according to standardized classification
- 2398 schemes.
- 2399 • In a SOA ecosystem with complex social structures, awareness may be provided for
- 2400 specific communities of interest. The architectural mechanisms for providing
- 2401 awareness to communities of interest ~~will~~ require support for:
 - 2402 ○ policies that allow dynamic formation of communities of interest;
 - 2403 ○ trust that awareness can be provided for and only for specific communities of
 - 2404 interest, the bases of which is typically built on keying and encryption
 - 2405 technology.
- 2406 • The architectural mechanisms for determining willingness to interact ~~will~~ require
- 2407 support for:
 - 2408 ○ verification of identity and credentials of the provider and/or consumer;
 - 2409 ○ access to and understanding of description;
 - 2410 ○ inspection of functionality and capabilities;
 - 2411 ○ inspection of policies and/or contracts.
- 2412 • The architectural mechanisms for establishing reachability ~~will~~ require support for:
 - 2413 ○ the location or address of an endpoint;
 - 2414 ○ verification and use of a service interface by means of a communication
 - 2415 protocol;
 - 2416 ○ determination of presence with an endpoint which may only be determined at
 - 2417 the point of interaction but may be further aided by the use of a presence
 - 2418 protocol for which the endpoints actively participate.

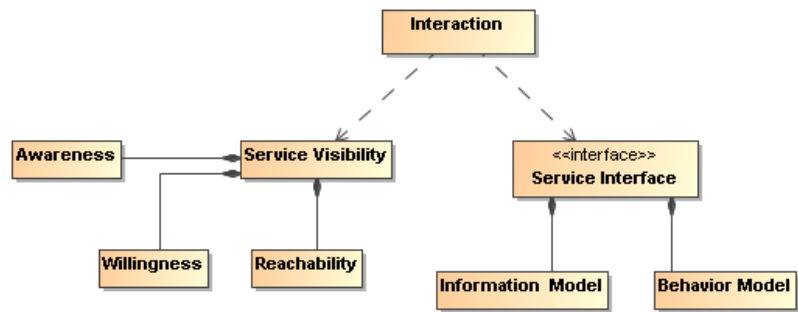
2419 4.3 Interacting with Services Model

2420 Interaction is the activity involved in using a service to access capability in order to
 2421 achieve a particular desired real world effect, where real world effect is the actual *result*
 2422 of using a service. An interaction can be characterized by a sequence of actions.
 2423 Consequently, interacting with a service, i.e. performing actions against the service—
 2424 usually mediated by a series of message exchanges—involves actions performed by
 2425 the service. Different modes of interaction are possible such as modifying the shared
 2426 state of a resource. Note that a participant (or delegate acting on behalf of the
 2427 participant) can be the sender of a message, the receiver of a message, or both.

2428 4.3.1 Interaction Dependencies

2429 Recall from the Reference Model that service visibility is the capacity for those with
 2430 needs and those with capabilities to be able to interact with each other, and that the
 2431 service interface is the means by which the underlying capabilities of a service are
 2432 accessed. Ideally, the details of the underlying service implementation are abstracted
 2433 away by the service interface. [Service] interaction therefore has a direct dependency
 2434 on the visibility of the service as well as its implementation-neutral interface (see [Figure](#)

2435 | 31Figure 30). Service visibility is composed of awareness, willingness, and reachability
2436 and service interface is composed of the information and behavior models. Service
2437 visibility is modeled in Section 4.2 while service interface is modeled in Section 4.1.

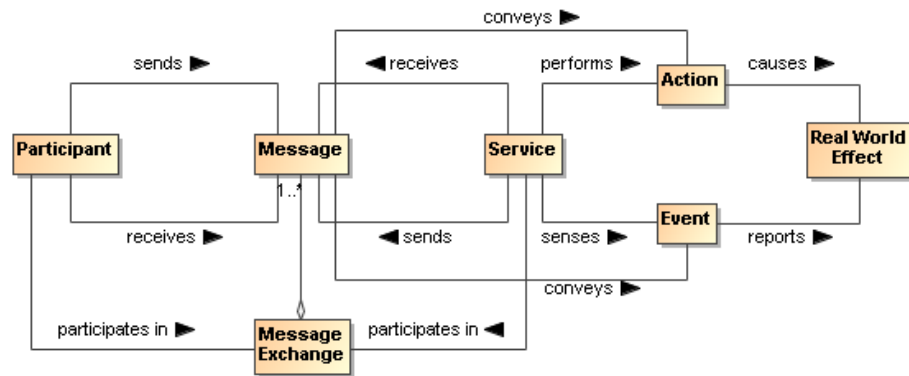


2438
2439 | Figure 3130 Interaction dependencies.

2440 | **4.3.2 Actions and Events**

2441 | For purposes of this Reference Architecture the SOA-RAF, the authors have committed
2442 to the use of message exchange between service participants to denote actions
2443 performed against and by the service, and to denote events that report on real world
2444 effects that are caused by the service actions. A visual model of the relationship
2445 between these concepts is shown in Figure 3234.

Comment [PFB52]: Don't like this indirect formulation – we took them out throughout s3



2446
2447 | Figure 3234 A "message" conveys either an action or an event.

2448 | A message conveys either an action or an event. In other words, both actions and
2449 events, realized by the SOA services, are denoted by the messages. The Reference
2450 Model states that the action model characterizes the “permissible set of actions that
2451 may be invoked against a service.” We extend that notion here to include events as
2452 part of the event model and that messages denote either actions or notification of
2453 events.

2454 | In Section **Error! Reference source not found.**, we saw that participants interact with
2455 each other in order to perform actions. An action is not itself the same thing as the
2456 result of performing the action. When an action is performed against a service, the real
2457 world effect that results is reported in the form of notification of events.

2458 4.3.3 Message Exchange

2459 *Message exchange* is the means by which service participants (or their delegates)
2460 interact with each other. There are two primary modes of interaction: joint actions that
2461 cause real world effects, and notification of events that report real world effects.¹³

2462 A message exchange is used to affect an action when the messages contain the
2463 appropriately formatted content that should be interpreted as joint action and the
2464 delegates involved interpret the message appropriately.

2465 A message exchange is also used to communicate event notifications. An event is an
2466 occurrence that is of interest to some participant; in our case when some real world
2467 effect has occurred. Just as action messages ~~will~~ have formatting requirements, so ~~will~~
2468 do event notification messages. In this way, the Information Model of a service must
2469 specify the syntax (structure), and semantics (meaning) of the action messages and
2470 event notification messages as part of a service interface. It must also specify the
2471 syntax and semantics of any data that is carried as part of a payload of the action or
2472 event notification message. The Information Model is described in greater detail in the
2473 Service Description Model (see Section 4.1).

2474 In addition to the Information Model that describes the syntax and semantics of the
2475 messages and data payloads, exception conditions and error handling in the event of
2476 faults (e.g., network outages, improper message formats, etc.) must be specified or
2477 referenced as part of the Service Description.

2478 When a message is interpreted as an action, the correct interpretation typically requires
2479 the receiver to perform a set of operations. These *operations* represent the sequence
2480 of actions (often private) a service must perform in order to validly participate in a given
2481 joint action.

2482 Similarly, the correct consequence of realizing a real world effect may be to initiate the
2483 reporting of that real world effect via an event notification.

2484 Message Exchange

2485 The means by which joint actions and event notifications are coordinated by
2486 service participants (or delegates).

2487 Operations

2488 The sequence of actions a service must perform in order to validly participate in a
2489 given joint action.

2490 4.3.3.1 Message Exchange Patterns (MEPs)

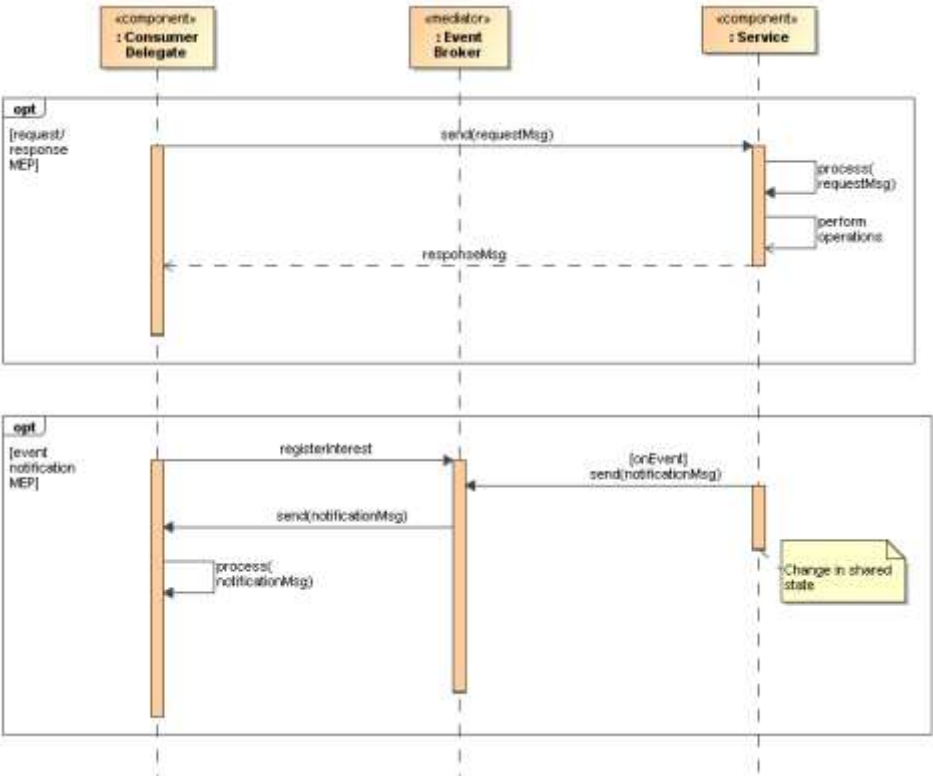
2491 ~~As stated earlier, this Reference Architecture~~The SOA-RAF commits to the use of
2492 message exchange to denote actions against the services, and to denote notification of
2493 events that report on real world effects that arise from those actions.

2494 Based on these assumptions, the basic temporal aspect of service interaction can be
2495 characterized by two fundamental message exchange patterns (MEPs):

¹³ The notion of “joint” in joint action implies that you have to have a speaker *and* a listener in order to interact.

2496
2497
2498
2499
2500
2501
2502
2503
2504
2505
2506

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



2507
2508
2509
2510
2511
2512
2513
2514
2515
2516

Figure 3332 Fundamental SOA message exchange patterns (MEPs)

In the UML sequence diagram shown in Figure 3332 it is assumed that the service participants (consumer and provider) have delegated message handling to hardware or software delegates acting on their behalf. In the case of the service consumer, this is represented by the *Consumer Delegate* component. In the case of the service provider, the delegate is represented by the *Service* component. The message interchange model illustrated represents a logical view of the MEPs and not a physical view. In other words, specific hosts, network protocols, and underlying messaging system are not shown as these tend to be implementation specific.

2517 Although such implementation-specific elements are considered outside the scope of
2518 this [Reference Architecture document](#), they are important considerations in modeling the
2519 SOA execution context. Recall from the Reference Model that the *execution context* of a
2520 service interaction is “the set of infrastructure elements, process entities, policy
2521 assertions and agreements that are identified as part of an instantiated service
2522 interaction, and thus forms a path between those with needs and those with
2523 capabilities.”

2524 4.3.3.2 Request/Response MEP

2525 In a request/response MEP, the Consumer Delegate component sends a request
2526 message to the Service component. The Service component then processes the
2527 request message. Based on the content of the message, the Service component
2528 performs the service operations. Following the completion of these operations, a
2529 response message is returned to the Consumer Delegate component. The response
2530 could be that a step in a process is complete, the initiation of a follow-on operation, or
2531 the return of requested information.¹⁴

2532 Although the sequence diagram shows a *synchronous* interaction (because the sender
2533 of the request message, i.e., Consumer Delegate, is blocked from continued processing
2534 until a response is returned from the Service) other variations of request/response are
2535 valid, including *asynchronous* (non-blocking) interaction through use of queues,
2536 channels, or other messaging techniques.

2537 What is important to convey here is that the request/response MEP represents action,
2538 which causes a real world effect, irrespective of the underlying messaging techniques
2539 and messaging infrastructure used to implement the request/response MEP.

2540 4.3.3.3 Event Notification MEP

2541 An event is made visible to interested consumers by means of an event notification
2542 message exchange that reports a real world effect; specifically, a change in shared
2543 state between service participants. The basic event notification MEP takes the form of a
2544 one-way message sent by a notifier component (in this case, the Service component)
2545 and received by components with an interest in the event (here, the Consumer Delegate
2546 component).

2547 Often the sending component may not be fully aware of all the components that **will**
2548 receive the notification; particularly in so-called publish/subscribe (“pub/sub”) situations.
2549 In event notification message exchanges, it is rare to have a tightly-coupled link
2550 between the sending and the receiving component(s) for a number of practical reasons.
2551 One of the most common is the potential for network outages or communication

¹⁴ 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.

interrupts that can result in loss of notification of events. Therefore, a third-party mediator component is often used to decouple the sending and receiving components . Although this is typically an implementation issue, because this type of third-party decoupling is so common in event-driven systems, [we felt that for this Reference Architecture, it was-is warranted for use in modeling this type of message exchange in the SOA-RAF](#). This third-party intermediary is shown in [Figure 33](#) as an Event Broker mediator. As with the request/response MEP, no distinction is made between synchronous versus asynchronous communication, although asynchronous message exchange is illustrated in the UML sequence diagram depicted in [Figure 33](#).

4.3.4 Composition of Services

Composition of services is the act of aggregating or “composing” a single service from one or more other services. A simple model of service composition is illustrated in [Figure 34](#).




Figure 34 Simple model of service composition.

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

It is possible for a service composition to be opaque from one perspective and transparent from another. For example, a service may appear to be a single service from the Consumer’s Delegate’s perspective, but is transparently composed of one or more services from a service management perspective. A Service Management Service needs to be able to have visibility into the composition in order to properly manage the dependencies between the services used in constructing the composite service—including managing the service’s lifecycle. The subject of services as management entities is described and modeled in the [Owning Service-Oriented Architectures Ownership in a SOA Ecosystem](#) View of [this Reference Architecture](#) the SOA-RAF and [will-is not be](#) further elaborated in this section. The point to be made here is that there can be different levels of opaqueness or transparency when it comes to visibility of service composition.

2589 Services can be composed in a variety of ways including direct service-to-service
2590 interaction by using programming techniques, or they can be aggregated by means of a
2591 scripting approach that leverages a service composition scripting language. Such
2592 scripting approaches are further elaborated in the following sub-sections on service-
2593 oriented business processes and collaborations.

2594 4.3.4.1 Service-Oriented Business Processes

2595 The concepts of business processes and collaborations in the context of transactions
2596 and exchanges across organizational boundaries are described and modeled as part of
2597 the [Service Ecosystem View](#) [Participation in a SOA Ecosystem view](#) of this [Reference](#)
2598 [Architecture](#) (see Section 3). Here, we focus on the belief that the principle of
2599 composition of services can be applied to business processes and collaborations. Of
2600 course, business processes and collaborations traditionally represent complex, multi-
2601 step business functions that may involve multiple participants, including internal users,
2602 external customers, and trading partners. Therefore, such complexities cannot simply
2603 be ignored when transforming traditional business processes and collaborations to their
2604 service-oriented variants.

2605 Business Processes

2606 Business processes are a set of one or more linked activities that are performed
2607 to achieve a certain business outcome.

2608 Service orientation as applied to business processes (i.e., “service-oriented business
2609 processes”) means that the aggregation or composition of all of the abstracted activities,
2610 flows, and rules that govern a business process can themselves be abstracted as a
2611 service **[BLOOMBERG/SCHMELZER]**.

2612 When business processes are abstracted in this manner and accessed through SOA
2613 services, all of the concepts used to describe and model composition of services that
2614 were articulated in Section 4.3.4 apply. There are some important differences from a
2615 composite service that represents an abstraction of a business process from a
2616 composite service that represents a single-step business interaction. As stated earlier,
2617 business processes have temporal properties and can range from short-lived processes
2618 that execute on the order of minutes or hours to long-lived processes that can execute
2619 for weeks, months, or even years. Further, these processes may involve many
2620 participants. These are important considerations for the consumer of a service-oriented
2621 business process and these temporal properties must be articulated as part of the meta-
2622 level aspects of the service-oriented business process in its Service Description, along
2623 with the meta-level aspects of any sub-processes that may be of use or need to be
2624 visible to the service consumer.

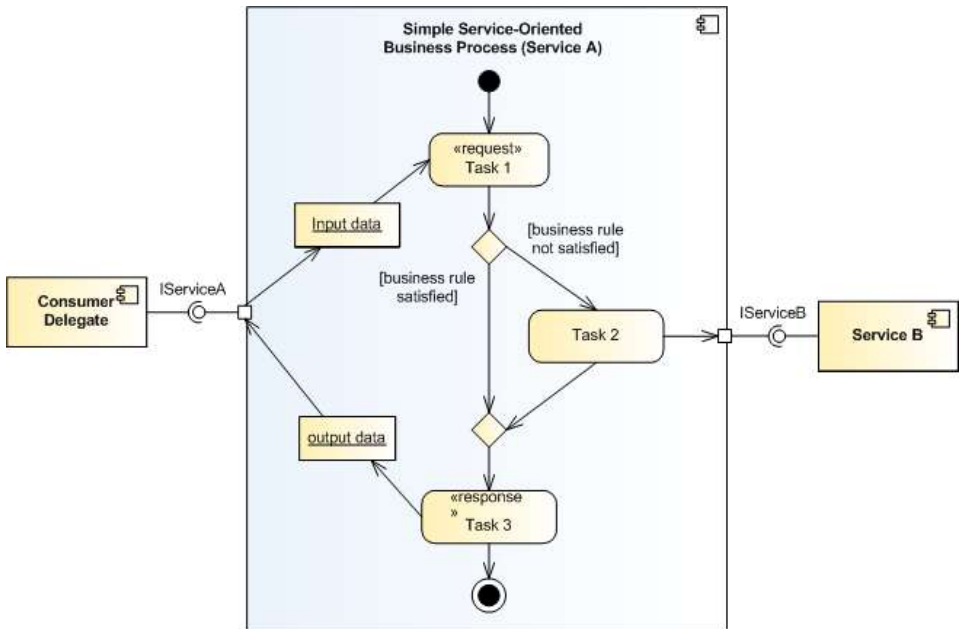
2625 In addition, a workflow activity represents a unit of work that some entity acting in a
2626 described role (i.e., role player) is asked to perform. Activities can be broken down into
2627 steps with each step representing a task for the role player to perform. A technique that
2628 is used to compose service-oriented business processes that are hierarchical (top-
2629 down) and self-contained in nature is known as *orchestration*.

2630 Orchestration

2631 A technique used to compose service-oriented business processes that are
2632 executed and coordinated by an actor acting as “conductor.”

2633 An orchestration is typically implemented using a scripting approach to compose
2634 service-oriented business processes. This typically involves use of a standards-based
2635 orchestration scripting language. In terms of automation, an orchestration can be
2636 mechanized using a business process orchestration engine, which is a hardware or
2637 software component (delegate) responsible for acting in the role of central
2638 conductor/coordinator responsible for executing the flows that comprise the
2639 orchestration.

2640 A simple generic example of such an orchestration is illustrated in [Figure 35](#)[Figure 34](#).



2641
2642 *Figure 35**34* Abstract example of orchestration of service-oriented business process.

2643 Here, we use a UML activity diagram to model the simple service-oriented business
2644 process as it allows us to capture the major elements of business processes such as
2645 the set of related tasks to be performed, linking between tasks in a logical flow, data that
2646 is passed between tasks, and any relevant business rules that govern the transitions
2647 between tasks. A task is a unit of work that an individual, system, or organization
2648 performs and can be accomplished in one or more steps or subtasks. While subtasks
2649 can be readily modeled, they are not illustrated in the orchestration model In [Figure](#)
2650 [35](#)[Figure 34](#)..

2651 This particular example is based on a request/response MEP and captures how one
2652 particular task (Task 2) actually utilizes an externally-provided service, Service B. The
2653 entire service-oriented business process is exposed as Service A that is accessible via
2654 its externally visible interface, IServiceA.

2655 Although not explicitly shown in the orchestration model above, it is assumed that there
2656 exists a software or hardware component, i.e., orchestration engine that executes the
2657 process flow. Recall that a central concept to orchestration is that process flow is

2658 coordinated and executed by a single conductor delegate; hence the name
2659 “orchestration.”

2660 4.3.4.2 Service-Oriented Business Collaborations

2661 Business collaborations typically represent the interaction involved in executing
2662 business transactions, where a business transaction is defined in the [Service
Ecosystem View Participation in a SOA Ecosystem view](#) as “a joint action engaged in by
2663 two or more participants in which resources are exchanged” (see Section [3.2.43.2.3](#)).
2664

2665 It is important to note that business collaborations represent “peer”-style interactions; in
2666 other words, peers in a business collaboration act as equals. This means that unlike
2667 the orchestration of business processes, there is no single or central entity that
2668 coordinates or “conducts” a business collaboration. These peer styles of interactions
2669 typically occur between trading partners that span organizational boundaries.

2670 Business collaborations can also be service-enabled. For purposes of this Reference
2671 Architecture [Foundation](#), we refer to these as “service-oriented business collaborations.”
2672 Service-oriented business collaborations do not necessarily imply exposing the entire
2673 peer-style business collaboration as a service itself but rather the collaboration uses
2674 service-based interchanges.

2675 The technique that is used to compose service-oriented business collaborations in
2676 which multiple parties collaborate in a peer-style as part of some larger business
2677 transaction by exchanging messages with trading partners and external organizations
2678 (e.g., suppliers) is known as *choreography* [NEWCOMER/LOMOW].

2679 Choreography

2680 A technique used to characterize service-oriented business collaborations based
2681 on ordered message exchanges between peer entities in order to achieve a
2682 common business goal.

2683 Choreography differs from orchestration primarily in that each party in a business
2684 collaboration describes its part in the service interaction. Note that choreography as we
2685 have defined it here should not be confused with the term *process choreography*, which
2686 is defined in the [Service Ecosystem View Participation in a SOA Ecosystem view](#) as “the
2687 description of the possible interactions that may take place between two or more
2688 participants to fulfill an objective.” This is an example of domain-specific nomenclature
2689 that often leads to confusion and why we are making note of it here.

2690 A simple generic example of a choreography is illustrated in [Figure 36](#)~~Figure 35~~

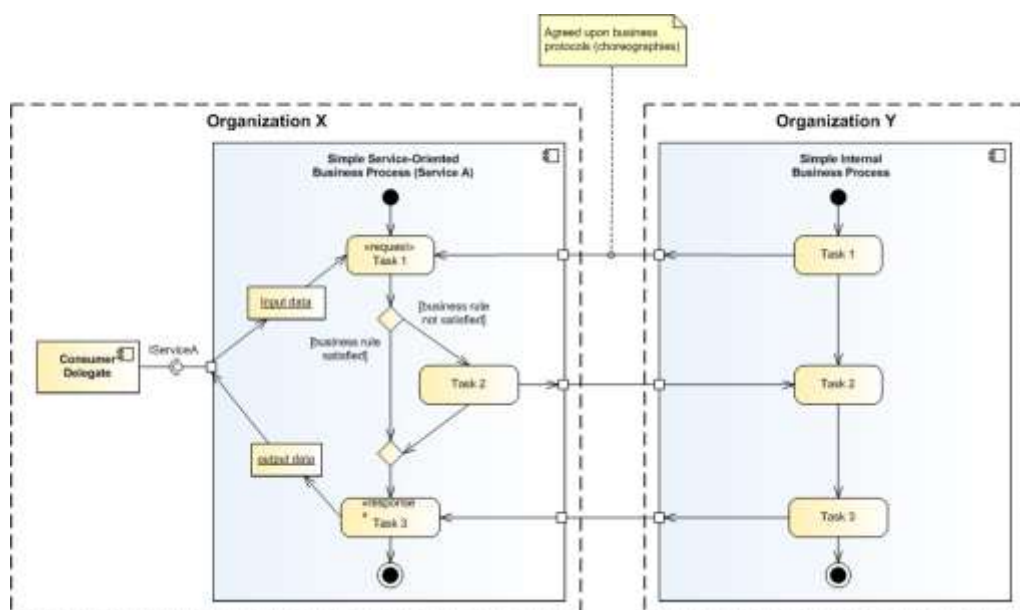


Figure 36-35 Abstract example of choreography of service-oriented business collaboration.

This example, which is a variant of the orchestration example illustrated earlier in Figure 35, adds trust boundaries between two organizations; namely, Organization X and Organization Y. It is assumed that these two organizations are peer entities that have an interest in a business collaboration, for example, Organization X and Organization Y could be trading partners. Organization X retains the service-oriented business process Service A, which is exposed to internal consumers via its provided service interface, IServiceA. Organization Y also has a business process that is involved in the business collaboration; however, for this example, it is an internal business process that is not exposed to potential consumers either within or outside its organizational boundary.

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

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

4.3.5 Architectural Implications of Interacting with Services

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

- A well-defined service Information Model that:

- 2716 ○ describes the syntax and semantics of the messages used to denote actions
2717 and events;
- 2718 ○ describes the syntax and semantics of the data payload(s) contained within
2719 messages;
- 2720 ○ documents exception conditions in the event of faults due to network outages,
2721 improper message/data formats, etc.;
- 2722 ○ is both human readable and machine processable;
- 2723 ○ is referenceable from the Service Description artifact.
- 2724 • A well-defined service Behavior Model that:
 - 2725 ○ characterizes the knowledge of the actions invokes against the service and
2726 events that report real world effects as a result of those actions;
 - 2727 ○ characterizes the temporal relationships and temporal properties of actions
2728 and events associated in a service interaction;
 - 2729 ○ describe activities involved in a workflow activity that represents a unit of
2730 work;
 - 2731 ○ describes the role (s) that a role player performs in a service-oriented
2732 business process or service-oriented business collaboration;
 - 2733 ○ is both human readable and machine processable;
 - 2734 ○ is referenceable from the Service Description artifact.
- 2735 • Service composition mechanisms to support orchestration of service-oriented
2736 business processes and choreography of service-oriented business collaborations
2737 such as:
 - 2738 ○ Declarative and programmatic compositional languages;
 - 2739 ○ Orchestration and/or choreography engines that support multi-step
2740 processes as part of a short-lived or long-lived business transaction;
 - 2741 ○ Orchestration and/or choreography engines that support compensating
2742 transactions in the presences of exception and fault conditions.
- 2743 • Infrastructure services that provides mechanisms to support service interaction,
2744 including but not limited to:
 - 2745 ○ mediation services such as message and event brokers, providers, and/or
2746 buses that provide message translation/transformation, gateway
2747 capability, message persistence, reliable message delivery, and/or
2748 intelligent routing semantics;
 - 2749 ○ binding services that support translation and transformation of multiple
2750 application-level protocols to standard network transport protocols;
 - 2751 ○ auditing and logging services that provide a data store and mechanism to
2752 record information related to service interaction activity such as message
2753 traffic patterns, security violations, and service contract and policy
2754 violations
 - 2755 ○ security services that abstract techniques such as public key
2756 cryptography, secure networks, virus protection, etc., which provide
2757 protection against common security threats in a SOA ecosystem;
 - 2758 ○ monitoring services such as hardware and software mechanisms that both
2759 monitor the performance of systems that host services and network traffic
2760 during service interaction, and are capable of generating regular
2761 monitoring reports.

- 2762 • A layered and tiered service component architecture that supports multiple message
2763 exchange patterns (MEPs) in order to:
- 2764 ○ promote the industry best practice of separation of concerns that facilitates
2765 flexibility in the presence of changing business requirements;
 - 2766 ○ promote the industry best practice of separation of roles in a service
2767 development lifecycle such that subject matter experts and teams are
2768 structured along areas of expertise;
 - 2769 ○ support numerous standard interaction patterns, peer-to-peer interaction
2770 patterns, enterprise integration patterns, and business-to-business
2771 integration patterns.

2772 4.4 Policies and Contracts Model

2773 A common phenomenon of many machines and systems is that the scope of potential
2774 behavior is much broader than is actually needed for a particular circumstance. This is
2775 especially true of a system as powerful as a SOA ecosystem. As a result, the behavior
2776 and performance of the system tend to be under-constrained by the implementation;
2777 instead, the actual behavior is expressed by means of policies of some form. Policies
2778 define the choices that stakeholders make; these choices are used to guide the actual
2779 behavior of the system to the desired behavior and performance.

2780 As noted in Section 3.1.4 a policy is a constraint of some form that is promulgated by a
2781 stakeholder who has the responsibility of ensuring that the constraint is enforced. In
2782 contrast, contracts are **agreements** between participants. However, like policies, it is a
2783 necessary part of contracts that they are enforceable.

2784 While responsibility for enforcement may differ, both contracts and policies share a
2785 common characteristic – there is a **constraint** that must be enforced. In both cases the
2786 mechanisms needed to enforce policy constraints ~~will be largely~~ are likely to be identical;
2787 in this model we focus on the issues involved in representing policies and contracts and
2788 on some of the principles behind their enforcement.

2789 4.4.1 Policy and Contract Representation

2790 A **policy constraint** is a specific kind of constraint: the ontology of policies and
2791 contracts includes the core concepts of permission, obligation, owner, subject. In
2792 addition, it may be necessary to be able combine policy constraints and to be able to
2793 resolve policy conflicts.

2794 4.4.1.1 Policy Framework

2795 Policy Framework

2796 A policy framework is a language in which policy constraints may be expressed.

2797 A policy framework combines a syntax for expressing policy constraints together with a
2798 decision procedure for determining if a policy constraint is satisfied.




Figure 37.36 Policies and Contracts

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

Logical Framework

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

Policy Ontology

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

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

Included in many policy ontologies are the basic signals of permissions and obligations. Some policy frameworks are sufficiently constrained that there is not possibility of representing an obligation; in which case there is often no need to 'call out' the distinction between permissions and obligations.

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

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

Policy Owner

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

Policy Subject

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

Policy Constraint

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

2831 Policy Object

2832 A policy object is an identifiable state, action or resource that is potentially
2833 constrained by the policy.

2834 Permission

2835 A ~~permission~~ constraint governs the ability of an actor to perform an action or to
2836 enter some specified ~~state~~.

2837 Note that permissions are distinct from ~~ability~~ and from authority. Authority refers to the
2838 legitimate nature of an action, whereas permission refers to the right to perform the
2839 action.

2840 Obligation

2841 An ~~obligation~~ constraint governs the requirement that a participant or other actor
2842 should perform an action or maintain some specified ~~state~~.

2843 For example, once the service consumer and provider have entered into an agreement
2844 to provide and consume a service, both participants incur obligations: the consumer is
2845 obligated to pay for the service and the provider is obligated to provide the service.

2846 Obligations to maintain state may range from a requirement to maintain a minimum
2847 balance on an account through a requirement that a service provider 'remember' that a
2848 particular service consumer is logged in.

2849 Obligations and permissions have a positive form and a negative form. A positive
2850 permission refers to something that a ~~policy subject~~ may do, a negative permission
2851 refers to something the ~~policy subject~~ may not do.

2852 These definitions are replicated from Section 3.1.4.

2853 4.4.2 Policy and Contract Enforcement

2854 The enforcement of policy constraints has to address two core problems: how to
2855 enforce the atomic policy constraints, and how to enforce combinations of policy
2856 constraints. In addition, it is necessary to address the resolution of policy conflicts.

2857 4.4.2.1 Enforcing Simple Policy Constraints

2858 The two primary kinds of policy constraint – permission and obligation – naturally lead to
2859 different styles of enforcement. A permission constraint must typically be enforced *prior*
2860 to the policy subject invoking the **policy object**. On the hand, an obligation constraint
2861 must typically be enforced post-facto through some form of auditing process and
2862 remedial action.

2863 For example, if a communications policy required that all communication be encrypted,
2864 this is enforceable at the point of communication: any attempt to communicate a
2865 message that is not encrypted can be blocked.

2866 Similarly, an obligation to pay for services rendered is enforced by ensuring that
2867 payment arrives within a reasonable period of time. Invoices are monitored for prompt
2868 (or lack of) payment.

2869 The key concepts in enforcing both forms of policy constraint are the policy decision and
2870 the policy enforcement.

2871 **Policy Decision**

2872 A policy decision is a determination as to whether a given policy constraint is
2873 satisfied or not.

2874 A policy decision is effectively a measurement of some state – typically a portion of the
2875 SOA ecosystem’s **shared state**. This implies a certain *timeliness* in the measuring: a
2876 measurement that is too early or is too late does not actually help in determining if the
2877 policy constraint is satisfied appropriately.

2878 **Policy Enforcement**

2879 A policy enforcement is the use of a mechanism to limit the behavior and/or state
2880 of policy subjects to comply with a policy decision.

2881 A policy enforcement implies the use of some mechanism to ensure compliance with a
2882 policy decision. The range of mechanisms is completely dependent on the kinds of
2883 atomic policy constraints that the policy framework may support. As noted above, the
2884 two primary styles of constraint – permission and **obligation** —~~will~~ lead to different
2885 styles of enforcement.

2886 **4.4.2.2 Enforcing Policy Combinations**

2887 Enforcing policy combinations is primarily an elaboration of enforcing simple policy
2888 constraints. The process of policy decisions is enhanced to allow a measurement to
2889 involve combinations of policy constraints and the process of policy enforcement may
2890 need to be enhanced to coordinate the enforcement of multiple policy constraints
2891 simultaneously.

2892 **4.4.2.3 Conflict Resolution**

2893 Whenever it is possible that more than one policy constraint applies in a given situation,
2894 there is the potential that the policies themselves are not mutually consistent. For
2895 example, a policy that requires communication to be encrypted and a policy that
2896 requires an administrator to read every communication ~~are likely to be in~~ conflict with
2897 each other – the two policies cannot both be satisfied.

2898 In general, with sufficiently rich policy frameworks, it is not possible to always resolve
2899 policy conflicts automatically. However, a reasonable approach is to augment the policy
2900 decision process with simple policy conflict resolution rules; with the potential for
2901 *escalating* a policy conflict to human adjudication.

2902 **Policy Conflict**

2903 A policy conflict exists between two or more policies in a policy decision process
2904 if it is not possible to satisfy all the policies that apply.

2905 **Policy Conflict Resolution**

2906 A policy conflict resolution rule is a way of determining which policy should
2907 prevail in a policy conflict.

2908 The inevitable consequence of policy conflicts is that it is not possible to guarantee that
2909 all policies are satisfied at all times. This, in turn, implies a certain *flexibility* in the
2910 application of policy constraints: they will not always be honored.

2911 4.4.3 Architectural Implications

2912 The key choices that must be made in a system of policies center around the policy
2913 framework and policy enforcement mechanisms

- 2914 • There SHOULD be a standard policy framework that is adopted across the SOA
2915 ecosystem:
 - 2916 ○ This framework MUST permit the expression of simple policy constraints
 - 2917 | ○ ~~This framework~~The framework MAY allow (to a varying extent) the
2918 combination of policy constraints, including
 - 2919 • Both positive and negative constraints
 - 2920 • Conjunctions and disjunctions of constraints
 - 2921 • The quantification of constraints
 - 2922 ○ The framework MUST at least allow the policy subject and the policy object to
2923 be identified as well as the policy constraint.
 - 2924 ○ The framework MAY allow further structuring of policies into modules,
2925 inheritance between policies and so on.
- 2926 • There SHOULD be mechanisms that facilitate the application of policies:
 - 2927 ○ There SHOULD be mechanisms that allow policy decisions to be made,
2928 consistent with the policy frameworks and with the state of the SOA
2929 ecosystem.
 - 2930 ○ There SHOULD be mechanisms to enforce policy decisions
 - 2931 • There SHOULD be mechanisms to support the measurement of
2932 whether certain policy constraints are satisfied or not, or to what
2933 degree they are satisfied.
 - 2934 • Such enforcement mechanisms MAY include support for both
2935 permission-style constraints and obligation-style constraints.
 - 2936 • Enforcement mechanisms MAY support the simultaneous enforcement
2937 of multiple policy constraints across multiple points in the SOA
2938 ecosystem.
 - 2939 ○ There SHOULD be mechanisms to resolve policy conflicts
 - 2940 • This MAY involve escalating policy conflicts to human adjudication.
 - 2941 ○ There SHOULD be mechanisms that support the management and
2942 promulgation of policies.

5 Owning Service Oriented ArchitecturesOwnership in a SOA Ecosystem View

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

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

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

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

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

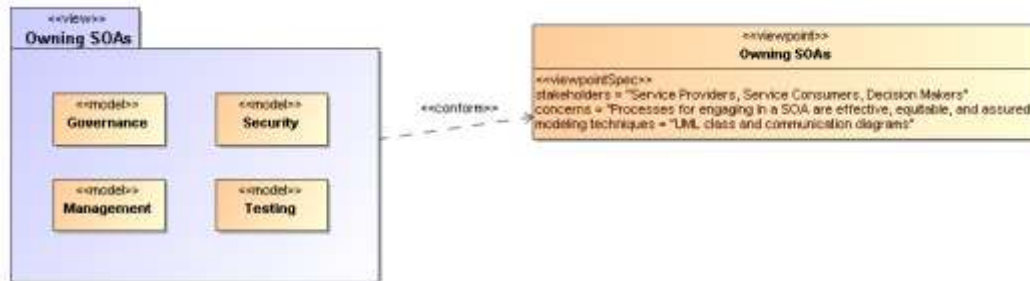


Figure 3837 Model Elements Described in the Owning Service Oriented ArchitecturesOwnership in a SOA
Ecosystem View

The following subsections present models of these functions.

5.1 Governance Model

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

2974 5.1.1 Understanding Governance

2975 5.1.1.1 Terminology

2976 Governance is about making decisions that are aligned with the overall organizational
2977 strategy and culture of the enterprise. **[Gartner]** It specifies the decision rights and
2978 accountability framework to encourage desirable behaviors **[Weill/Ross-MIT Sloan**
2979 **School]** towards realizing the strategy and defines incentives (positive or negative)
2980 towards that end. It is less about overt control and strict adherence to rules, and more
2981 about guidance and effective and equitable usage of resources to ensure sustainability
2982 of an organization's strategic objectives. **[TOGAF v8.1]**

2983 To accomplish this, governance requires organizational structure and processes and
2984 must identify who has authority to define and carry out its mandates. It must address
2985 the following questions: 1) what decisions must be made to ensure effective
2986 management and use?, 2) who should make these decisions?, and 3) how will these
2987 decisions be made and monitored? , and (4) how will these decisions be
2988 communicated? The intent is to achieve goals, add value, and reduce risk.

2989 Within a single ownership domain such as an enterprise, generally there is a hierarchy
2990 of governance structures. Some of the more common enterprise governance structures
2991 include corporate governance, technology governance, IT governance, and architecture
2992 governance **[TOGAF v8.1]**. These governance structures can exist at multiple levels
2993 (global, regional, and local) within the overall enterprise.

2994 It is often asserted that SOA governance is a specialization of IT governance as there is
2995 a natural hierarchy of these types of governance structures; however, the focus of SOA
2996 governance is less on decisions to ensure effective management and use of IT as it is
2997 to ensure effective management and use of SOA-based systems. Certainly, SOA
2998 governance must still answer the basic questions also associated with IT governance,
2999 i.e., who should make the decisions, and how these decisions will be made and
3000 monitored.

3001 5.1.1.2 Relationship to Management

3002 There is often confusion centered on the relationship between governance and
3003 management. As described earlier, governance is concerned with decision making.
3004 Management, on the other hand, is concerned with execution. Put another way,
3005 governance describes the world as leadership wants it to be; management executes
3006 activities that intends to make the leadership's desired world a reality. Where
3007 governance determines who has the authority and responsibility for making decisions
3008 and the establishment of guidelines for how those decisions should be made,
3009 management is the actual process of making, implementing, and measuring the impact
3010 of those decisions **[Loeb]**. Consequently, governance and management work in
3011 concert to ensure a well-balanced and functioning organization as well as an ecosystem
3012 of inter-related organizations. In the sections that follow, we elaborate further on the
3013 relationship between governance and management in terms of setting and enforcing
3014 service policies, contracts, and standards as well as addressing issues surrounding
3015 regulatory compliance.

3016 5.1.1.3 Why is SOA Governance Important?

3017 One of the hallmarks of SOA that distinguishes it from other architectural paradigms for
3018 distributed computing is the ability to provide a uniform means to offer, discover, interact

3019 with and use capabilities (as well the ability to compose new capabilities from existing
3020 ones) all in an environment that transcends domains of ownership. Consequently,
3021 ownership, and issues surrounding it, such as obtaining acceptable terms and
3022 conditions (T&Cs) in a contract, is one of the primary topics for SOA governance.
3023 Generally, IT governance does not include T&Cs, for example, as a condition of use as
3024 its primary concern.

3025 Just as other architectural paradigms, technologies, and approaches to IT are subject to
3026 change and evolution, so too is SOA. Setting policies that allow change management
3027 and evolution, establishing strategies for change, resolving disputes that arise, and
3028 ensuring that SOA-based systems continue to fulfill the goals of the business are all
3029 reasons why governance is important to SOA.

3030 5.1.1.4 Governance Stakeholders and Concerns

3031 As noted in Section **Error! Reference source not found.** the participants in a service
3032 interaction include the service provider, the service consumer, and other interested or
3033 unintentional third parties. Depending on the circumstances, it may also include the
3034 owners of the underlying capabilities that the SOA services access. Governance must
3035 establish the policies and rules under which duties and responsibilities are defined and
3036 the expectations of participants are grounded. The expectations include transparency
3037 in aspects where transparency is mandated, trust in the impartial and consistent
3038 application of governance, and assurance of reliable and robust behavior throughout the
3039 SOA ecosystem.

3040 5.1.2 A Generic Model for Governance

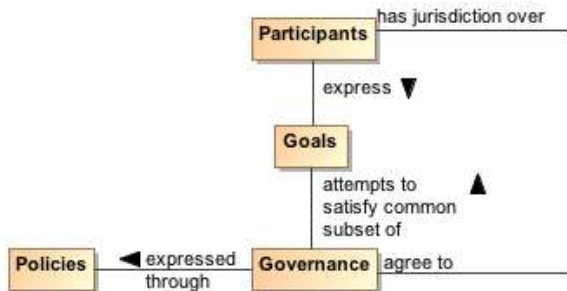
3041 Governance

3042 Governance is the prescribing of conditions and constraints consistent with
3043 satisfying common goals and the structures and processes needed to define and
3044 respond to actions taken towards realizing those goals.

3045 The following is a generic model of governance represented by segmented models that
3046 begin with motivation and proceed through measuring compliance. It is not ~~meant to be~~
3047 ~~an-all-encompassing~~ ~~treatise on governance~~ but a focused subset that captures the
3048 aspects necessary to describe governance for SOA. It ~~is also does~~ not ~~meant to~~ imply
3049 that practical application of governance is a single, isolated instance of these models; in
3050 ~~fact reality~~, there ~~are likely may be~~ hierarchical ~~chains of governance that apply~~ and
3051 ~~possibly~~ parallel chains ~~of governance~~ that ~~govern deal with~~ different aspects or focus
3052 on different goals. This is discussed further in section 5.1.2.5. The defined models are
3053 simultaneously applicable to each of the overlapping instances.

3054 A given enterprise may already have portions of these models in place. To a large
3055 extent, the models shown here are not specific to SOA; discussions on direct
3056 applicability begin in section 5.1.3.

3057 5.1.2.1 Motivating Governance



3058
3059 | Figure 3938 Motivating governance model

3060 An organizational domain such as an enterprise is made up of participants who may be
3061 individuals or groups of individuals forming smaller organizational units within the
3062 enterprise. The overall business strategy should be consistent with the Goals of the
3063 participants; otherwise, the business strategy would not provide value to the participants
3064 and governance towards those ends becomes difficult if not impossible. This is not to
3065 say that an instance of governance will-simultaneously satisfies all the goals of all the
3066 participants; rather, the goals of any governance instance must sufficiently satisfy a
3067 useful subset of each participant's goals so as to provide value and ensure the
3068 cooperation of all the participants.

3069 A policy is the formal characterization of the conditions and constraints that governance
3070 deems as necessary to realize the goals which it is attempting to satisfy. Policy may
3071 identify required conditions or actions or may prescribe limitations or other constraints
3072 on permitted conditions or actions. For example, a policy may prescribe that
3073 safeguards must be in place to prevent unauthorized access to sensitive material. It
3074 may also prohibit use of computers for activities unrelated to the specified work
3075 assignment. Policy is made operational through the promulgating and implementing of
3076 Rules and Regulations (as defined in section 5.1.2.3).

3077 As noted in section 4.4.2, policy may be asserted by any participant or on behalf of the
3078 participant by its organization. Part of the purpose of governance is to arbitrate among
3079 diverse goals of participants and diverse policies articulated to realize those goals. The
3080 intent is to form a consistent whole that allows governance to minimize ambiguity about
3081 its purpose. While resolving all ambiguity would be an ideal, it is unlikely that all
3082 inconsistencies will be identified and resolved before governance becomes operational.

3083 For governance to have effective jurisdiction over participants, there must be some
3084 degree of agreement by each-all participants that they will abide by the governance
3085 mandates. A minimal degree of agreement often presages participants who “slow-roll” if
3086 not actively reject complying with Policies that express the specifics of governance.

5.1.2.2 Setting Up Governance

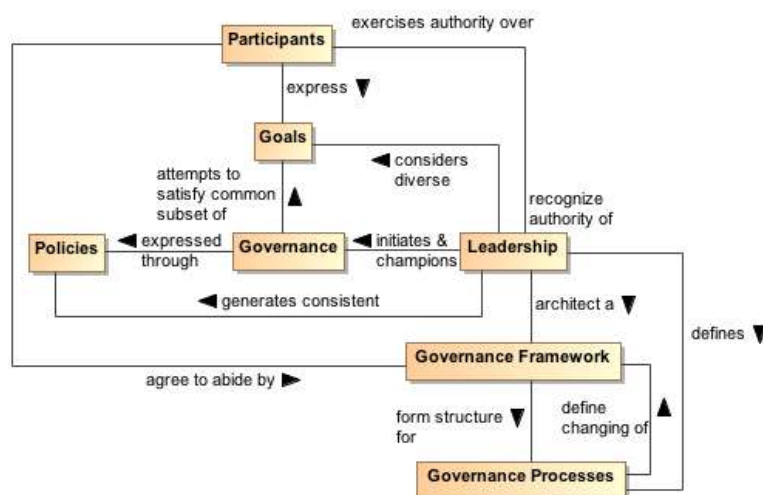


Figure 4039 Setting up governance model

Leadership

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

Governance Framework

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

Governance Processes

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

As noted earlier, governance requires an appropriate organizational structure and identification of who has authority to make governance decisions. In [Figure 40](#)[Figure 39](#), the entity with governance authority is designated the Leadership. This is someone, possibly one or more of the participants, that participants recognize as having authority for a given purpose or over a given set of issues or concerns.

The Leadership is responsible for prescribing or delegating a working group to prescribe the Governance Framework that forms the structure for Governance Processes which define how governance is to be carried out. This does not itself define the specifics of how governance is to be applied, but it does provide an unambiguous set of procedures

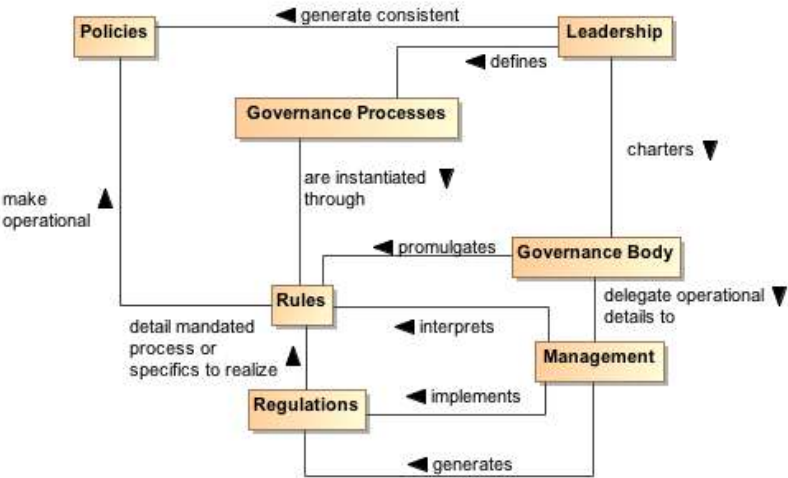
3114 that should ensure consistent actions which participants agree are fair and account for
3115 sufficient input on the subjects to which governance ~~will be~~ applied.

3116 The participants may be part of the working group that codifies the Governance
3117 Framework and Processes. When complete, the participants must acknowledge and
3118 agree to abide by the products generated through application of this structure.

3119 The Governance Framework and Processes are often documented in the charter of a
3120 body created or designated to oversee governance. This is discussed further in the
3121 next section. Note that the Governance Processes should also include those necessary
3122 to modify the Governance Framework itself.

3123 An important function of Leadership is not only to initiate but also be the consistent
3124 champion of governance. Those responsible for carrying out governance mandates
3125 must have Leadership who makes it clear to participants that expressed Policies are
3126 seen as a means to realizing established goals and that compliance with governance is
3127 required.

3128 **5.1.2.3 Carrying Out Governance**



3129
3130 *Figure 4140 Carrying out governance model*

3131 **Rule**

3132 A Rule is a prescribed guide for carrying out activities and processes leading to
3133 desired results, e.g. the operational realization of policies.

3134 **Regulation**

3135 A Regulation is a mandated process or the specific details that derive from the
3136 interpretation of Rules and lead to measureable quantities against which
3137 compliance can be measured.

3138 To carry out governance, Leadership charts a Governance Body to promulgate the
3139 Rules needed to make the Policies operational. The Governance Body acts in line with
3140 Governance Processes for its rule-making process and other functions. Whereas
3141 Governance is the setting of Policies and defining the Rules that provide an operational

context for Policies, the operational details of governance ~~are likely~~may be delegated by the Governance Body to Management. Management generates Regulations that specify details for Rules and other procedures to implement both Rules and Regulations. For example, Leadership could set a Policy that all authorized parties should have access to data, the Governance Body would promulgate a Rule that PKI certificates are required to establish identity of authorized parties, and Management can specify a Regulation of who it deems to be a recognized PKI issuing body. In summary, Policy is a predicate to be satisfied and Rules prescribe the activities by which that satisfying occurs. A number of rules may be required to satisfy a given policy; the carrying out of a rule may contribute to several policies being realized.

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

5.1.2.4 Ensuring Governance Compliance

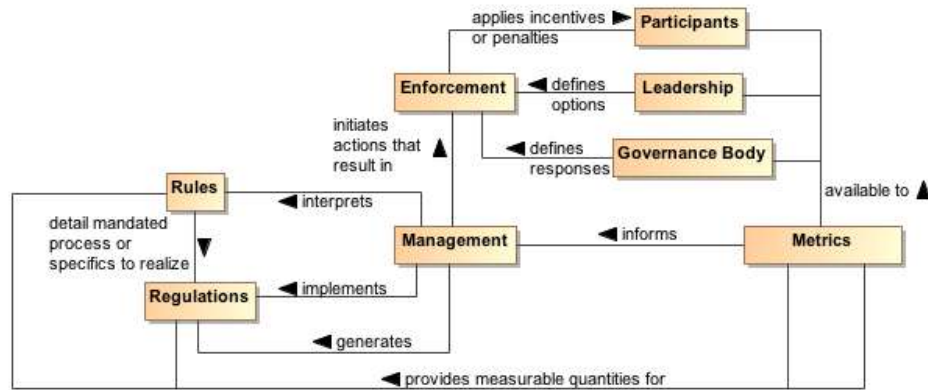


Figure 4244 Ensuring governance compliance model

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

The Leadership in its relationship with participants ~~will have~~has certain options that can be used for Enforcement. A common option may be to effect future funding. The Governance Body defines specific enforcement responses, such as what degree of

3175 compliance is necessary for full funding to be restored. It is up to Management to
3176 identify compliance shortfalls and to initiate the Enforcement process.
3177 Note, enforcement does not strictly need to be negative consequences. Management
3178 can use Metrics to identify exemplars of compliance and Leadership can provide
3179 options for rewarding the participants. ~~It is likely t~~The Governance Body ~~that~~ defines
3180 awards or other incentives.

3181 5.1.2.5 Considerations for Multiple Governance Chains

3182 As noted in section 5.1.2, instances of the governance model often occur as a tiered
3183 arrangement, with governance at some level delegating specific authority and
3184 responsibility to accomplish a focused portion of the original level's mandate. For
3185 example, a corporation may encompass several lines of business and each line of
3186 business governs its own affairs in a manner that is consistent with and contributes to
3187 the goals of the parent organization. Within the line of business, an IT group may be
3188 given the mandate to provide and maintain IT resources, giving rise to IT governance.

3189 In addition to tiered governance, there ~~are likely to may~~ be multiple governance chains
3190 working in parallel. For example, a company making widgets ~~likely~~ has policies intended
3191 to ensure they make high quality widgets and make an impressive profit for their
3192 shareholders. On the other hand, Sarbanes-Oxley is a parallel governance chain in the
3193 United States that specifies how the management must handle its accounting and
3194 information that needs to be given to its shareholders. The parallel chains may just be
3195 additive or may be in conflict and require some harmonization.

3196 Being distributed and representing different ownership domains, a SOA participant ~~is~~
3197 ~~falls likely~~ under the jurisdiction of multiple governance domains simultaneously and
3198 may individually need to resolve consequent conflicts. The governance domains may
3199 specify precedence for governance conformance or it may fall to the discretion of the
3200 participant to decide on the course of actions they believe appropriate.

3201 5.1.3 Governance Applied to SOA

3202 5.1.3.1 Where SOA Governance is Different

3203 SOA governance is often discussed in terms of IT governance, but rather than a parent-
3204 child relationship, ~~Figure 43~~Figure 42 shows the two as siblings of the general
3205 governance described in section 5.1.2. There are obvious dependencies and a need for
3206 coordination between the two, but the idea of aligning IT with business already
3207 demonstrates that resource providers and resource consumers must be working
3208 towards common goals if they are to be productive and efficient. While SOA governance
3209 ~~will be~~is shown to be active in the area of infrastructure, it is a specialized concern for
3210 having a dependable platform to support service interaction; a ~~host range~~ of traditional
3211 IT issues is ~~considered to be therefore~~ out of scope ~~of this document~~. A SOA
3212 governance plan for an enterprise will not ~~of itself~~ resolve shortcomings with the
3213 enterprise's IT governance.

3214 Governance in the context of SOA is that organization of services: that promotes their
3215 visibility; that facilitates interaction among service participants; and that directs that the
3216 results of service interactions are those real world effects as described within the

3217 service description and constrained by policies and contracts as assembled in the
3218 execution context.

3219 SOA governance must specifically account for control across different ownership
3220 domains, i.e. all the participants may not be under the jurisdiction of a single
3221 governance authority. However, for governance to be effective, the participants must
3222 agree to recognize the authority of the Governance Body and must operate within the
3223 Governance Framework and through the Governance Processes so defined.

3224 SOA governance must account for interactions across ownership boundaries, which
3225 likely may also imply across enterprise governance boundaries. For such situations,
3226 governance emphasizes the need for agreement that some Governance Framework
3227 and Governance Processes have jurisdiction, and the governance defined must satisfy
3228 the Goals of the participants for cooperation to continue. A standards development
3229 organization such as OASIS is an example of voluntary agreement to governance over
3230 a limited domain to satisfy common goals.

3231 The specifics discussed in the figures in the previous sections are equally applicable to
3232 governance across ownership boundaries as it is within a single boundary. There is a
3233 charter agreed to when participants become members of the organization, and this
3234 charter sets up the structures and processes that will be followed. Leadership may be
3235 shared by the leadership of the overall organization and the leadership of individual
3236 groups themselves chartered per the Governance Processes. There are
3237 Rules/Regulations specific to individual efforts for which participants agree to local
3238 goals, and Enforcement can be loss of voting rights or under extreme circumstances,
3239 expulsion from the group.

3240 Thus, the major difference for SOA governance is an appreciation for the cooperative
3241 nature of the enterprise and its reliance on furthering common goals if productive
3242 participation is to continue.

3243 5.1.3.2 What Must be Governed

3244 An expected benefit of employing SOA principles is the ability to quickly bring resources
3245 to bear to deal with unexpected and evolving situations. This requires a great deal of
3246 confidence in the underlying capabilities that can be accessed and in the services that
3247 enable the access. It also requires considerable flexibility in the ways these resources
3248 can be employed. Thus, SOA governance requires establishing confidence and trust
3249 while instituting a solid framework that enables flexibility, indicating a combination of
3250 strict control over a limited set of foundational aspects but minimum constraints beyond
3251 those bounds.

3252

Comment [PFB53]: MUST,
SHOULD?

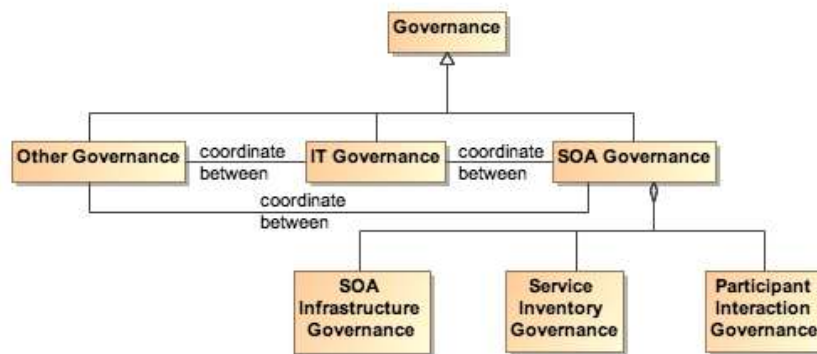


Figure 4342 Relationship among types of governance

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

Comment [PFB54]: Is this the same as ‘a SOA-based System’? If so, which is the preferred term?

5.1.3.2.1 Governance of SOA Infrastructure

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

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

5.1.3.2.2 Governance of the Service Inventory

Given an infrastructure in which other SOA services can operate, a key governance issue is which SOA services to allow in the ecosystem. The major concern SHOULD be a definition of well-behaved services, where the required behavior will likely inherit their

Comment [PFB55]: “required”...likely.. ???

3284 | characteristics from experiences with distributed computing but ~~will~~ also evolve with
3285 SOA experience. A major requirement for ensuring well-behaved services is collecting
3286 sufficient metrics to know how the service affects the SOA infrastructure and whether it
3287 complies with established infrastructure policies.

3288 | Another common concern of service approval is whether there ~~will be~~ is a possibility of
3289 duplication of function by multiple services. Some governance models talk to a tightly
3290 controlled environment where a primary concern is to avoid any service duplication.
3291 Other governance models talk to a market of services where the consumers have wide
3292 choices. For the latter, it is anticipated that the better services will emerge from market
3293 consensus and the availability of alternatives will drive innovation.

3294 | ~~It is likely that a~~ Some combination of control and openness will emerge, possibly with a
3295 different appropriate balance for different categories of use. For SOA governance, the
3296 issue is less which services are approved but rather ensuring that sufficient description
3297 is available to support informed decisions for appropriate use. Thus, SOA governance
3298 SHOULD concentrate on identifying the required attributes to adequately describe a
3299 service, the required target values of the attributes, and the standards for defining the
3300 meaning of the attributes and their target values. Governance may also specify the
3301 processes by which the attribute values are measured and the corresponding
3302 certification that some realized attribute set may imply.

3303 For example, unlimited access for using a service may require a degree of life cycle
3304 maturity that has demonstrated sufficient testing over a certain size community.
3305 Alternately, the policy may specify that a service in an earlier phase of its life cycle may
3306 be made available to a smaller, more technically sophisticated group in order to collect
3307 the metrics that would eventually allow the service to advance its life cycle status.

3308 This aspect of governance is tightly connected to description because, given a well-
3309 behaved set of services, it is the responsibility of the consumer (or policies promulgated
3310 by the consumer's organization) to decide whether a service is sufficient for that
3311 consumer's intended use. The goal is to avoid global governance specifying criteria that
3312 are too restrictive or too lax for the local needs of which global governance has little
3313 insight.

3314 Such an approach to specifying governance allows independent domains to describe
3315 services in local terms while still having the services available for informed use across
3316 domains. In addition, changes to the attribute sets within a domain can be similarly
3317 described, thus supporting the use of newly described resources with the existing ones
3318 without having to update the description of all the legacy content.

3319 **5.1.3.2.3 Governance of Participant Interaction**

3320 Finally, given a reliable services infrastructure and a predictable set of services, the
3321 third aspect of governance is prescribing what is required during a service interaction.

3322 Governance would specify adherence to service interface and service reachability
3323 parameters and would require that the result of an interaction **MUST** correspond to the
3324 real world effects as contained in the service description. Governance would ensure
3325 preconditions for service use are satisfied, in particular those related to security aspects
3326 such as user authentication, authorization, and non-repudiation. If conflicts arise,

3327 governance would specify resolution processes to ensure appropriate agreements,
3328 policies, and conditions are met.

3329 It would also rely on sufficient monitoring by the SOA infrastructure to ensure services
3330 remain well-behaved during interactions, e.g. do not use excessive resources or exhibit
3331 other prohibited behavior. Governance would also require that policy agreements as
3332 documented in the execution context for the interaction are observed and that the
3333 results and any after effects are consistent with the agreed policies. ~~It is likely that in~~
3334 ~~this area the g~~Governance ~~will~~ focus on more contractual and legal aspects rather than
3335 the precursor descriptive aspects. SOA governance may prescribe the processes by
3336 which SOA-specific policies are allowed to change, but there are ~~likely probably~~ more
3337 business-specific policies that will be governed by processes outside SOA governance.

Comment [PFB56]: SHOULD?

3338 5.1.3.3 Overarching Governance Concerns

3339 There are numerous governance related concerns whose effects span the three areas
3340 just discussed. One is the area of standards, how these are mandated, and how the
3341 mandates may change. The Web Services standards stack is an example of relevant
3342 standards where a significant number are still under development. In addition, while
3343 there are notional scenarios that guide what standards are being developed, the fact
3344 that many of these standards do not yet exist precludes operational testing of their
3345 adequacy or effectiveness as a necessary and sufficient set.

3346 That said, standards are critical to creating a SOA ecosystem where SOA services can
3347 be introduced, used singularly, and combined with other services to deliver complex
3348 business functionality. As with other aspects of SOA governance, the Governance
3349 Body should identify the minimum set felt to be needed and rigorously enforce that that
3350 set be used where appropriate. The Governance Body ~~must~~ take care to expand and
3351 evolve the mandated standards in a predictable manner and with sufficient technical
3352 guidance that new services ~~will be~~are able to coexist as much as possible with the old,
3353 and changes to standards do not cause major disruptions.

3354 Another area that may see increasing activity as SOA expands ~~will be~~is additional
3355 regulation by governments and associated legal institutions. New laws are ~~likely that will~~
3356 ~~may~~ deal with transactions which are service based, possibly including taxes on the
3357 transactions. Disclosures laws ~~are likely to~~may mandate certain elements of description
3358 so both the consumer and provider act in a predictable environment and are protected
3359 from ambiguity in intent or action. Such laws are ~~likely to~~spawn rules and regulations
3360 that will influence the metrics collected for evaluation of compliance.

Comment [PFB57]: MUST?

3361 5.1.3.4 Considerations for SOA Governance

3362 The Reference Architecture definition of a loosely coupled system is one in which the
3363 constraints on the interactions between components is minimal: sufficient to permit
3364 interoperation without additional constraints that may be an artifact of implementation
3365 technology. While governance experience for standalone systems provides useful
3366 guides, we must be careful not to apply constraints that would preclude the flexibility,
3367 agility, and adaptability we expect to realize from a SOA ecosystem.

3368 One of the strengths of SOA is it can make effective use of diversity rather than
3369 requiring monolithic solutions. Heterogeneous organizations can interact without
3370 requiring each conforms to uniform tools, representation, and processes. However, with

3371 this diversity comes the need to adequately define those elements necessary for
3372 consistent interaction among systems and participants, such as which communication
3373 protocol, what level of security, which vocabulary for payload content of messages. The
3374 solution is not always to lock down these choices but to standardize alternatives and
3375 standardize the representations through which an unambiguous identification of the
3376 alternative chosen can be conveyed. For example, the URI standard specifies the URI
3377 string, including what protocol is being used, what is the target of the message, and how
3378 may parameters be attached. It does not limit the available protocols, the semantics of
3379 the target address, or the parameters that can be transferred. Thus, as with our
3380 definition of loose coupling, it provides absolute constraints but minimizes which
3381 constraints it imposes.

3382 There is not a one-size-fits-all governance but a need to understand the types of things
3383 governance ~~will be~~ called ~~up~~ on to do in the context of the goals of SOA. ~~It is likely that~~
3384 ~~s~~Some communities ~~will may~~ initially desire and require very stringent governance
3385 policies and procedures while other ~~will~~ see need for very little. Over time, best
3386 practices will evolve, ~~likely~~ resulting in some consensus on a sensible minimum and,
3387 except in extreme cases where it is demonstrated to be necessary, a loosening of strict
3388 governance toward the best practice mean.

3389 A question of how much governance may center on how much time governance
3390 activities require versus how quickly is the system being governed expected to respond
3391 to changing conditions. For large single systems that take years to develop, the
3392 governance process could move slowly without having a serious negative impact. For
3393 example, if something takes two years to develop and the steps involved in governance
3394 take two months to navigate, then the governance can go along in parallel and may not
3395 have a significant impact on system response to changes. Situations where it takes as
3396 long to navigate governance requirements as it does to develop a response are
3397 examples where governance may need to be reevaluated as to whether it facilitates or
3398 inhibits the desired results. Thus, the speed at which services are expected to appear
3399 and evolve needs to be considered when deciding the processes for control. The
3400 added weight of governance should be appropriate for overall goals of the application
3401 domain and the service environment.

3402 Governance, as with other aspects of any SOA implementation, should start small and
3403 be conceptualized in a way that keeps it flexible, scalable, and realistic. A set of useful
3404 guidelines would include:

- 3405 • Do not hardwire things that will inevitably change. For example, develop a
3406 system that uses the representation of policies rather than code the policies into
3407 the implementations.
- 3408 • Avoid setting up processes that demo well for three services without considering
3409 how ~~it will~~~~they may~~ work for 300. Similarly, consider whether the display of
3410 status and activity for a small number of services will also be effective for an
3411 operator in a crisis situation looking at dozens of services, each with numerous,
3412 sometimes overlapping and sometimes differing activities.
- 3413 • Maintain consistency and realism. A service solution responding to a natural
3414 disaster cannot be expected to complete a 6-week review cycle but be effective
3415 in a matter of hours.

3416 5.1.4 Architectural Implications of SOA Governance

3417 The description of SOA governance indicates numerous architectural requirements on
3418 the SOA ecosystem:

- 3419 • Governance is expressed through policies and assumes multiple use of focused
3420 policy modules that can be employed across many common circumstances. This
3421 requires the existence of:
 - 3422 ○ descriptions to enable the policy modules to be visible, where the
3423 description includes a unique identifier for the policy and a sufficient, and
3424 preferably a machine process-able, representation of the meaning of
3425 terms used to describe the policy, its functions, and its effects;
 - 3426 ○ one or more discovery mechanisms that enable searching for policies that
3427 best meet the search criteria specified by the service participant; where
3428 the discovery mechanism will have access to the individual policy
3429 descriptions, possibly through some repository mechanism;
 - 3430 ○ accessible storage of policies and policy descriptions, so service
3431 participants can access, examine, and use the policies as defined.
- 3432 • Governance requires that the participants understand the intent of governance,
3433 the structures created to define and implement governance, and the processes to
3434 be followed to make governance operational. This requires the existence of:
 - 3435 ○ an information collection site, such as a Web page or portal, where
3436 governance information is stored and from which the information is always
3437 available for access;
 - 3438 ○ a mechanism to inform participants of significant governance events, such
3439 as changes in policies, rules, or regulations;
 - 3440 ○ accessible storage of the specifics of Governance Processes;
 - 3441 ○ SOA services to access automated implementations of the Governance
3442 Processes
- 3443 • Governance policies are made operational through rules and regulations. This
3444 requires the existence of:
 - 3445 ○ descriptions to enable the rules and regulations to be visible, where the
3446 description includes a unique identifier and a sufficient, and preferably a
3447 machine process-able, representation of the meaning of terms used to
3448 describe the rules and regulations;
 - 3449 ○ one or more discovery mechanisms that enable searching for rules and
3450 regulations that may apply to situations corresponding to the search
3451 criteria specified by the service participant; where the discovery
3452 mechanism will have access to the individual descriptions of rules and
3453 regulations, possibly through some repository mechanism;
 - 3454 ○ accessible storage of rules and regulations and their respective
3455 descriptions, so service participants can understand and prepare for
3456 compliance, as defined.
 - 3457 ○ SOA services to access automated implementations of the Governance
3458 Processes.

- 3459 • Governance implies management to define and enforce rules and regulations.
3460 Management is discussed more specifically in section **Error! Reference source**
3461 **not found.**, but in a parallel to governance, management requires the existence
3462 of:
- 3463 ○ an information collection site, such as a Web page or portal, where
3464 management information is stored and from which the information is
3465 always available for access;
 - 3466 ○ a mechanism to inform participants of significant management events,
3467 such as changes in rules or regulations;
 - 3468 ○ accessible storage of the specifics of processes followed by management.
- 3469 • Governance relies on metrics to define and measure compliance. This requires
3470 the existence of:
- 3471 ○ the infrastructure monitoring and reporting information on SOA resources;
 - 3472 ○ possible interface requirements to make accessible metrics information
3473 generated or most easily accessed by the service itself.

3474 5.2 Security Model

3475 Security is one aspect of confidence – the confidence in the integrity, reliability, and
3476 confidentiality of the system. In particular, security focuses on those aspects of
3477 assurance that involve the accidental or malign intent of other people to damage or
3478 compromise trust in the system and on the availability of SOA-based systems to
3479 perform desired capability.

3480 Security

3481 Security concerns the set of mechanisms for ensuring and enhancing trust and
3482 confidence in the SOA ecosystem.

3483 Providing for security for Service Oriented Architecture is somewhat different than for
3484 other contexts; although many of the same principles apply equally to SOA and to other
3485 systems. The fact that SOA embraces crossing ownership boundaries makes the issues
3486 involved with moving data more visible.

3487 As well as securing the movement of data within and across ownership boundaries,
3488 security often revolves around resources: the need to guard certain resources against
3489 inappropriate access – whether reading, writing or otherwise manipulating those
3490 resources.

3491 Any comprehensive security solution must take into account the people that are using,
3492 maintaining and managing the SOA. Furthermore, the relationships between them must
3493 also be incorporated: any security assertions that may be associated with particular
3494 interactions originate in the people that are behind the interaction.

3495 We analyze security in terms of the social structures that define the legitimate
3496 permissions, obligations and roles of people in relation to the system, and mechanisms
3497 that must be put into place to realize a secure system. The former are typically captured
3498 in a series of security policy statements; the latter in terms of security *guards* that
3499 ensure that policies are enforced.

3500 How and when to apply these derived security policy mechanisms is directly associated
3501 with the assessment of the *threat model* and a *security response model*. The threat
3502 model identifies the kinds of threats that directly impact the message and/or application
3503 of constraints, and the response model is the proposed mitigation to those threats.
3504 Properly implemented, the result can be an acceptable level of risk to the safety and
3505 integrity of the system.

3506 5.2.1 Secure Interaction Concepts

3507 We can characterize secure interactions in terms of key security concepts [ISO/IEC
3508 27002]: confidentiality, integrity, authentication, authorization, non-repudiation, and
3509 availability. The concepts for secure interactions are well defined in other standards
3510 and publications. The security concepts here are not defined but rather related to the
3511 SOA ecosystem perspective of [this reference architecture foundation](#) [the SOA-RAF](#).

3512 5.2.1.1 Confidentiality

3513 Confidentiality concerns the protection of privacy of participants in their interactions.
3514 Confidentiality refers to the assurance that unauthorized entities are not able to read
3515 messages or parts of messages that are transmitted.

3516 Note that confidentiality has degrees: in a completely confidential exchange, third
3517 parties would not even be aware that a confidential exchange has occurred. In a
3518 partially confidential exchange, the identities of the participants may be known but the
3519 content of the exchange obscured.

3520 5.2.1.2 Integrity

3521 Integrity concerns the protection of information that is exchanged – either from
3522 unauthorized writing or inadvertent corruption. Integrity refers to the assurance that
3523 information that has been exchanged has not been altered.

3524 Integrity is different from confidentiality in that messages that are sent from one
3525 participant to another may be obscured to a third party, but the third party may still be
3526 able to introduce his own content into the exchange without the knowledge of the
3527 participants.

3528 [Figure 44](#)~~Figure 43~~ applies confidentiality and integrity to communicative action.

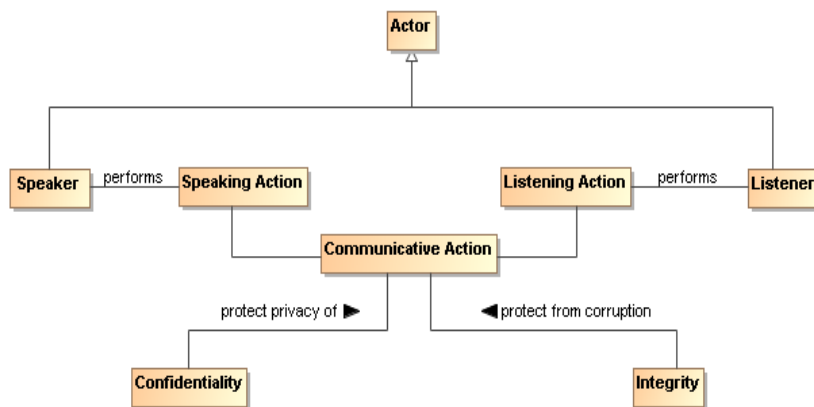


Figure 443 Confidentiality and Integrity

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

5.2.1.3 Authentication

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

Figure 45 applies authentication to the identity of participants.




Figure 454 Authentication

5.2.1.4 Authorization

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

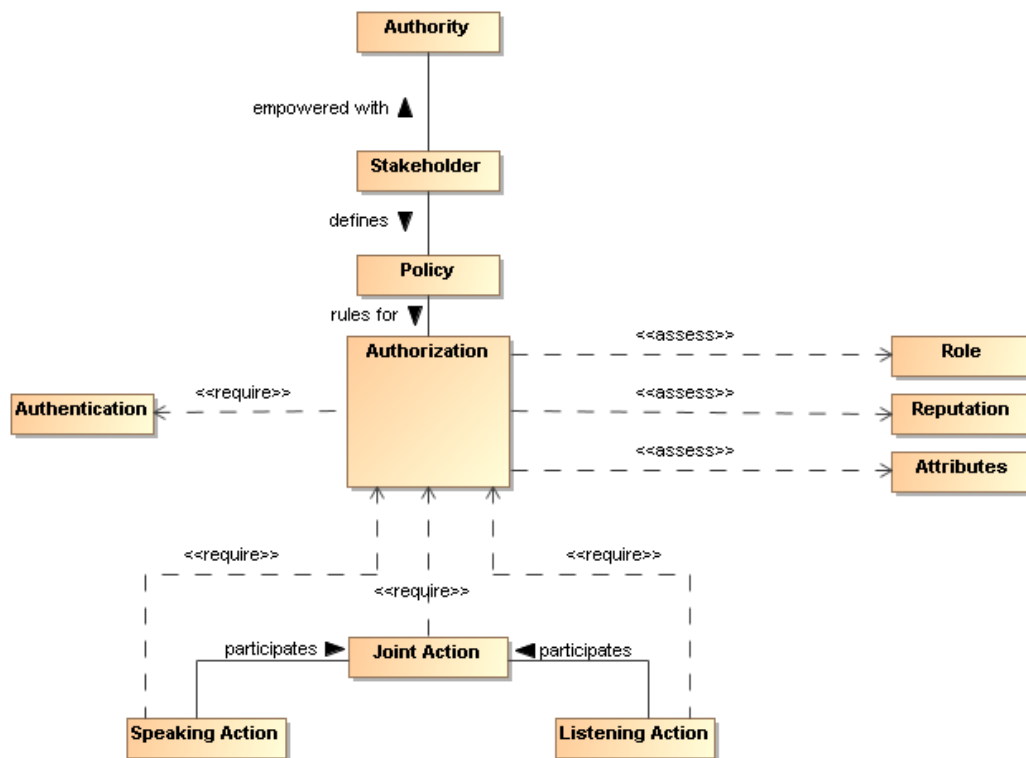


Figure 4645 Authorization

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

The role of policy for security is to permit stakeholders to express their choices. In ~~Figure 46~~Figure 45, a policy is a written constraint and the role, reputation, and attribute assertions are evaluated according to the constraints in the authorization policy. A combination of security mechanisms and their control via explicit policies can form the basis of an authorization solution.

5.2.1.5 Non-repudiation

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

3565 5.2.1.6 Availability

3566 Availability concerns the ability of systems to use and offer the services for which they
3567 were designed. One of the threats against availability is the so-called denial of service
3568 attack in which attackers attempt to prevent legitimate access to the system.

3569 We differentiate here between general availability – which includes aspects such as
3570 systems reliability – and availability as a security concept where we need to respond to
3571 active threats to the system.

3572 5.2.2 Where SOA Security is Different

3573 The core security concepts are fundamental to all social interactions. The evolution of
3574 sharing information using a SOA requires the flexibility to dynamically secure computing
3575 interactions in a computing ecosystem where the owning social groups, roles, and
3576 authority are constantly changing as described in section 5.1.3.1.

3577 SOA policy-based security can be more adaptive for a computing ecosystem than
3578 previous computing technologies allow for, and typically involves a greater degree of
3579 distributed mechanisms.

3580 Standards for security, as is the case with all aspects of SOA, play a large role in
3581 flexible security on a global scale. SOA security may also involve greater auditing and
3582 reporting to adhere to regulatory compliance established by governance structures.

3583 5.2.3 Security Threats

3584 There are a number of ways in which an attacker may attempt to compromise the
3585 security of a system. The two primary sources of attack are third parties attempting to
3586 subvert interactions between legitimate participants and an entity that is participating but
3587 attempting to subvert its partner(s). The latter is particularly important in a SOA where
3588 there may be multiple ownership boundaries and trust boundaries.

3589 The threat model lists some common threats that relate to the core security concepts
3590 listed in Section 5.2.1. Each technology choice in the realization of a SOA can
3591 potentially have many threats to consider.

3592 Message alteration

3593 If an attacker is able to modify the content (or even the order) of messages that
3594 are exchanged without the legitimate participants being aware of it then the
3595 attacker has successfully compromised the security of the system. In effect, the
3596 participants may unwittingly serve the needs of the attacker rather than their own.

3597 An attacker may not need to completely replace a message with his own to
3598 achieve his objective: replacing the identity of the beneficiary of a transaction
3599 may be enough.

3600 Message interception

3601 If an attacker is able to intercept and understand messages exchanged between
3602 participants, then the attacker may be able to gain advantage. This is probably
3603 the most commonly understood security threat.

3604 Man in the middle

3605 In a man-in-the-middle attack, the legitimate participants believe that they are
3606 interacting with each other; but are in fact interacting with the attacker. The
3607 attacker attempts to convince each participant that he is their correspondent;
3608 whereas in fact he is not.

3609 In a successful man-in-the-middle attack, legitimate participants ~~will offend~~ not
3610 have an ~~an true~~accurate understanding of the state of the other participants. The
3611 attacker can use this to subvert the intentions of the participants.

3612 Spoofing

3613 In a spoofing attack, the attacker convinces a participant that he is really
3614 someone else – someone that the participant would normally trust.

3615 Denial of service attack

3616 In a denial of service (DoS) attack, the attacker attempts to prevent legitimate
3617 users from making use of the service. A DoS attack is easy to mount and can
3618 cause considerable harm: by preventing legitimate interactions, or by slowing
3619 them down enough, the attacker may be able to simultaneously prevent
3620 legitimate access to a service and to attack the service by another means.

3621 A variation of the DoS attack is the Distributed Denial of Service attack. In a
3622 DDoS attack the attacker uses multiple agents to the attack the target. In some
3623 circumstances this can be extremely difficult to counteract effectively.

3624 One of the features of a DoS attack is that it does not require valid interactions to
3625 be effective: responding to invalid messages also takes resources and that may
3626 be sufficient to cripple the target.

3627 Replay attack

3628 In a replay attack, the attacker captures the message traffic during a legitimate
3629 interaction and then replays part of it to the target. The target is persuaded that a
3630 similar transaction to the previous one is being repeated and it ~~will respond~~s as
3631 though it were a legitimate interaction.

3632 A replay attack may not require that the attacker understand any of the individual
3633 communications; the attacker may have different objectives (for example
3634 attempting to predict how the target would react to a particular request).

3635 False repudiation

3636 In false repudiation, a user completes a normal transaction and then later
3637 attempts to deny that the transaction occurred. For example, a customer may
3638 use a service to buy a book using a credit card; then, when the book is delivered,
3639 refuse to pay the credit card bill claiming that *someone else* must have ordered
3640 the book.

3641 5.2.4 Security Responses

3642 Security goals are never absolute: it is not possible to guarantee 100% confidentiality,
3643 non-repudiation, etc. However, a well designed and implemented security response
3644 model can ensure acceptable levels of security risk. For example, using a well-designed

3645 cipher to encrypt messages may make the cost of breaking communications so great
3646 and so lengthy that the information obtained is valueless.

3647 Performing threat assessments, devising mitigation strategies, and determining
3648 acceptable levels of risk are the foundation for an effective process to mitigating threats
3649 in a cost-effective way.¹⁵ The choice in hardware and software to realize a SOA will be
3650 the basis for threat assessments and mitigation strategies. The stakeholders of a
3651 specific SOA implementation should determine acceptable levels of risk based on threat
3652 assessments and the cost of mitigating those threats.

Comment [PFB58]: the basis? Or a basis?

3653 5.2.4.1 Privacy Enforcement

3654 The most efficient mechanism to assure confidentiality is the encryption of information.
3655 Encryption is particularly important when messages must cross trust boundaries;
3656 especially over the Internet. Note that encryption need not be limited to the content of
3657 messages: it is possible to obscure even the existence of messages themselves
3658 through encryption and 'white noise' generation in the communications channel.

3659 The specifics of encryption are beyond the scope of this architecture. However, we are
3660 concerned about how the connection between privacy-related policies and their
3661 enforcement is made.

3662 A policy enforcement point for enforcing privacy may take the form of an automatic
3663 function to encrypt messages as they leave a trust boundary; or perhaps simply
3664 ensuring that such messages are suitably encrypted.

3665 Any policies relating to the level of encryption being used would then apply to these
3666 centralized messaging functions.

3667 5.2.4.2 Integrity Protection

3668 To protect against message tampering or inadvertent message alteration, and to allow
3669 the receiver of a message to authenticate the sender, messages may be accompanied
3670 by a digital signature. Digital signatures provide a means to detect if signed data has
3671 been altered. This protection can also extend to authentication and non-repudiation of a
3672 sender.

3673 A common way a digital signature is generated is with the use of a private key that is
3674 associated with a public key and a digital certificate. The private key of some entity in
3675 the system is used to create a digital signature for some set of data. Other entities in the
3676 system can check the integrity of the signed data set via signature verification
3677 algorithms. Any changes to the data that was signed will cause signature verification to
3678 fail, which indicates that integrity of the data set has been compromised.

3679 A party verifying a digital signature must have access to the public key that corresponds
3680 to the private key used to generate the signature. A digital certificate contains the public

¹⁵ In practice, there are perceptions of security from all participants regardless of ownership boundaries. Satisfying security policy often requires asserting sensitive information about the message initiator. The perceptions of this participant about information privacy may be more important than actual security enforcement within the SOA for this stakeholder.

3681 key of the owner, and is itself protected by a digital signature created using the private
3682 key of the issuing Certificate Authority (CA).

3683 **5.2.4.3 Message Replay Protection**

3684 To protect against replay attacks, messages may contain information that can be used
3685 to detect replayed messages. The simplest requirement to prevent replay attacks is that
3686 each message that is ever sent is unique. For example, a message may contain a
3687 message ID, a timestamp, and the intended destination.

3688 By storing message IDs, and comparing each new message with the store, it becomes
3689 possible to verify whether a given message has been received before (and therefore
3690 should be discarded).

3691 The timestamp may be included in the message to help check for message freshness.
3692 Messages that arrive after their message ID could have been cleared (after receiving
3693 the same message some time previously) may also have been replayed. A common
3694 means for representing timestamps is a useful part of an interoperable replay detection
3695 mechanism.

3696 The destination information is used to determine if the message was misdirected or
3697 replayed. If the replayed message is sent to a different endpoint than the destination of
3698 the original message, the replay could go undetected if the message does not contain
3699 information about the intended destination.

3700 In the case of messages that are replies to prior messages, it is also possible to include
3701 seed information in the prior messages that is randomly and uniquely generated for
3702 each message that is sent out. A replay attack can then be detected if the reply does
3703 not embed the random number that corresponds to the original message.

3704 **5.2.4.4 Auditing and Logging**

3705 False repudiation involves a participant denying that it authorized a previous interaction.
3706 An effective strategy for responding to such a denial is to maintain careful and complete
3707 logs of interactions which can be used for auditing purposes. The more detailed and
3708 comprehensive an audit trail is, the less likely it is that a false repudiation would be
3709 successful.

3710 The countermeasures assume that the non-repudiation tactic (e.g. digital signatures) is
3711 not undermined itself. For example, if private key is stolen and used by an adversary,
3712 even extensive logging cannot assist in rejecting a false repudiation.

3713 Unlike many of the security responses discussed here, it is likely that the scope for
3714 automation in rejecting a repudiation attempt is limited to careful logging.

3715 **5.2.4.5 Graduated engagement**

3716 The key to managing and responding to DoS attacks is to be careful in the use of
3717 resources when responding to interaction. Put simply, a system has a choice to respond
3718 to a communication or to ignore it. In order to avoid vulnerability to DoS attacks a
3719 service provider should be careful not to commit resources beyond those implied by the
3720 current state of interactions; this permits a graduation in commitment by the service
3721 provider that mirrors any commitment on the part of service consumers and attackers
3722 alike.

3723 5.2.5 Architectural Implications of SOA Security

3724 Providing SOA security in an ecosystem of governed services has the following
3725 implications on the policy support and the distributed nature of mechanisms used to
3726 assure SOA security:

- 3727 • Security expressed through policies have the same architectural implications as
3728 described in Section 4.4.3 for policies and contracts architectural implications.
- 3729 • Security policies require mechanisms to support security description
3730 administration, storage, and distribution.
- 3731 • Service descriptions supporting security policies should:
 - 3732 ○ have a meta-structure sufficiently rich to support security policies;
 - 3733 ○ be able to reference one or more security policy artifacts;
 - 3734 ○ have a framework for resolving conflicts between security policies.
- 3735 • The mechanisms that make-up the execution context in secure SOA-based
3736 systems should:
 - 3737 ○ provide protection of the confidentiality and integrity of message
3738 exchanges;
 - 3739 ○ be distributed so as to provide centralized or decentralized policy-based
3740 identification, authentication, and authorization;
 - 3741 ○ ensure service availability to consumers;
 - 3742 ○ be able to scale to support security for a growing ecosystem of services;
 - 3743 ○ be able to support security between different communication technologies;
- 3744 • Common security services include:
 - 3745 ○ services that abstract encryption techniques;
 - 3746 ○ services for auditing and logging interactions and security violations;
 - 3747 ○ services for identification;
 - 3748 ○ services for authentication;
 - 3749 ○ services for authorization;
 - 3750 ○ services for intrusion detection and prevention;
 - 3751 ○ services for availability including support for quality of service
3752 specifications and metrics.

3753 5.3 Management Model

3754 Management

3755 Management is the control of the use, configuration, and availability of resources
3756 in accordance with the policies of the stakeholders involved.

3757 There are three separate but linked domains of interest within the management of SOA-
3758 based systems. The first and most obvious is the management and support of the
3759 resources that are involved in any complex system – of which SOA-based systems are
3760 excellent examples. The second is the promulgation and enforcement of the policies
3761 and contracts agreed to by the stakeholders in SOA-based systems. The third domain is

3762 the management of the relationships of the participants in SOA-based systems – both to
3763 each other and to the services that they use and offer.

3764 There are many artifacts in a large system that may need management. As soon as
3765 there is the possibility of more than one instance of a thing, the issue of managing those
3766 things becomes relevant. Historically, systems management capabilities have been
3767 organized by the following functional groups known as “FCAPS” functions (based on
3768 ITU-T Rec. M.3400 (02/2000), “TMN Management Functions”): Fault management,
3769 configuration management, account management, performance and security
3770 management.


3771 In the context of SOA we see many possible resources that may require management:
3772 services, service descriptions, service capabilities, policies, contracts, roles,
3773 relationships, security, and infrastructure elements. In addition, given the ecosystem
3774 nature of SOA, it is also potentially necessary to manage the business relationships
3775 between participants in the SOA.

3776 Managing systems that may be used across ownership boundaries raises issues that
3777 are not normally present when managing a system within a single ownership domain.
3778 For example, care is required managing a service when the owner of the service, the
3779 provider of the service, the host of the service and access mediators to the service may
3780 all belong to different stakeholders. In addition, it may be important to allow service
3781 consumers to communicate their requirements to the service provider so that they are
3782 satisfied in a timely manner.

3783 A given service may be provided and consumed in more than one version. Version
3784 control of services is important both for service providers and service consumers (who
3785 may need to ensure certainty in the version of the service they are interacting with).

3786 In fact, managing a service has quite a few similarities to using a service: suggesting
3787 that we can use the service oriented model to manage SOA-based systems as well as
3788 provide them. A management service would be distinguished from a non-management
3789 service more by the nature of the capabilities involved (i.e., capabilities that relate to
3790 managing services) than by any intrinsic difference.

3791 In this model, we show how the SOA framework may apply to managing services as
3792 well as using and offering them. There are, of course, some special considerations that
3793 apply to service management which we bring out: namely that we ~~will be~~ managing
3794 the life-cycle of services, managing any service level attributes, managing
3795 dependencies between services and so on.



3796
3797 | Figure 4746 Managing resources in a SOA

3798 The core concept in management is that of a manageability capability:

3799 **Manageability Capability**

3800 The manageability capability of a resource is the capability that allows it to be
3801 managed with respect to some property. Note that manageability capabilities are
3802 not necessarily part of the managed entities themselves.

3803 Manageability capabilities are the core resources that management systems use to
3804 manage: each resource that may be managed in some way has a number of aspects
3805 that may be managed. For example, a service's life-cycle may be manageable, as may
3806 its Quality of Service parameter; a policy may also be managed for life-cycle but Quality
3807 of Service would not normally apply.

3808 **Life-cycle manageability**

3809 A manageability capability associated with a resource that permits the life cycle
3810 of the resource to be managed. As noted above, the life-cycle manageability
3811 capability of a resource is unlikely to reside within the resource itself (you cannot
3812 tell a system that is not running to start itself).

3813 The life-cycle management of a resource typically refers to how the resource is created,
3814 how it is destroyed and what dependencies ~~there might~~may exist that must be
3815 simultaneously managed.

3816 **Configuration manageability**

3817 A capability that permits the configuration of resources to be managed. Service
3818 configuration, in particular, may be complex in cases where there are
3819 dependencies between services and other resources.

3820 **Event monitoring manageability**

3821 Managing the reporting of events and faults is one of the key lower-level
3822 manageability capabilities.

3823 **Accounting manageability**

3824 A capability associated with resources that allows for the use of those resources
3825 to be measured and accounted for. This implies that not only can the *use* of
3826 resources be properly measured, but also that those *using* those resources also
3827 be properly identified.

3828 Accounting for the use of resources by participants in the SOA supports the
3829 proper budgeting and allocation of funding by participants.

3830 **Quality of service manageability**

3831 A manageability capability associated with a resource that permits any quality of
3832 service associated with the resource to be managed. Classic examples of this
3833 include bandwidth requirements and offerings associated with a service.

3834 **Business performance manageability**

3835 A manageability capability that is associated with services that permits the
3836 service's business performance to be monitored and managed. In particular, if
3837 there are business-level service level agreements that apply to a service, being
3838 able to monitor and manage those SLAs is an important role for management
3839 systems.

3840 | Building support for arbitrary business monitoring is ~~likely to be~~ challenging.
3841 | However, given a *measure* for determining a service's compliance to business
3842 | service level agreements, management systems can monitor that performance in
3843 | a way that is entirely similar to other management tasks.

3844 | **Policy manageability**

3845 | Where the policies associated with a resource may be complex and dynamic, so
3846 | those policies themselves may require management. The ability to manage those
3847 | policies (such as promulgating policies, retiring policies and ensuring that policy
3848 | decision points and enforcement points are current) is a management function.

3849 | In the particular case of policies, there is a special relationship between
3850 | management and policies. Just like other artifacts, policies require management
3851 | in a SOA. However, much of management is about *applying* policies also: where
3852 | governance is often about what the policies regarding artifacts and services
3853 | should be, a key management role is to ensure that those policies are
3854 | consistently applied.

3855 | **Management Service**

3856 | A management service is a service that manages other services and resources.

3857 | **Management Policy**

3858 | A management policy is a policy whose topic is a management topic. Just as with
3859 | other aspects of a SOA, the management of resources within the SOA may be
3860 | governed by management policies, contracts (such as SLAs).

3861 | In a deployed system, it may well be that different aspects of the management of a
3862 | given service are managed by different management services. For example, the life-
3863 | cycle management of services often involves managing dependencies between
3864 | services and resource requirements. Managing quality of service is often very specific to
3865 | the service itself; for example, quality of service attributes for a video streaming service
3866 | are quite different to those for a banking system.

3867 | There are additional concepts of management that often also apply to IT management:

3868 | **Systems management**

3869 | Systems management refers to enterprise-wide maintenance and administration
3870 | of distributed computer systems.

3871 | **Network management**

3872 | Network management refers to the maintenance and administration of large-
3873 | scale networks such as computer networks and telecommunication networks.
3874 | Systems and network management execute a set of functions required for
3875 | controlling, planning, deploying, coordinating, and monitoring the distributed
3876 | computer systems and the resources of a network.

3877 | However, for the purposes of this Reference Architecture [Foundation](#), while recognizing
3878 | their importance, we do not focus on systems management or network management.

3879 | - the specific identifier is not prescribed by this [Reference Architecture](#) but the
3880 | structure and semantics of the identifier must be indicated for the identifier value to be
3881 | properly used. For example, part of identity may include version identification.

3882 For this, the configuration management plan or similar document from which the version
3883 number is derived must be identified.

3884 **5.3.1 Management and Governance**

3885 The primary role of governance in the context of SOA is to allow the stakeholders in the
3886 SOA to be able to negotiate and set the key policies that govern the running of the
3887 system. Recall that in an ecosystems perspective, the goal is less to have complete
3888 fine-grained control but more to enable the individual participants to work together.
3889 Policies that are set at the governance of a SOA ~~will~~ tend to focus on the rules of
3890 engagement between participants – what kind of interacts are permissible, how to
3891 resolve disputes, and so on.
3892 While governance may be primarily focused on setting policies, management is more
3893 focused on realization and enforcement of policies.

3894 **5.3.2 Management Contracts and Policies**

3895 As we noted above, management can often be viewed as the application of contracts
3896 and policies to ensure the smooth running of the SOA. Policies play an important part
3897 in managing systems both as artifacts that need to be managed and as the guiding
3898 constraints to determine how the SOA should be managed.

3899 **5.3.2.1 Policies**

3900 "Although provision of management capabilities enables a service to become
3901 manageable, the extent and degree of permissible management are defined in
3902 management policies that are associated with the services. Management policies are
3903 used to define the obligations for, and permissions to, managing the service." [WSA]
3904 On the other hand, a policy without any means of enforcing it is vacuous. In the case of
3905 management policy, we rely on a management infrastructure to realize and enforce
3906 management policy.

3907 **5.3.3 Management Infrastructure**

3908 In order for a service or other resource to be manageable there must be a
3909 corresponding manageability capability that can effect that management. The
3910 particulars of this capability ~~will~~ vary ~~somewhat~~ depending on the nature of the
3911 capability. For example, a service life-cycle manageability capability requires the ability
3912 to start a service, to stop the service, and potentially to pause the service. Conversely,
3913 in order to manage document-like artifacts, such as service descriptions, the capability
3914 of storing the artifacts, controlling access to those artifacts, allowing updates of the
3915 artifacts to be deployed are all important capabilities for managing them.
3916 Elements of a basic service management infrastructure should include the following
3917 characteristics:
3918 • Integrate with existing security services
3919 • Monitoring
3920 • Heartbeat and Ping

- 3921 • Alerting
- 3922 • Pause/Restore/Restart Service Access
- 3923 • Logging, Auditing, Non-Repudiation
- 3924 • Runtime Version Management
- 3925 • Complement other infrastructure services (discovery, messaging, mediation)
- 3926 • Message Routing and Redirection
- 3927 • Failover
- 3928 • Load-balancing
- 3929 • QoS, Management of Service Level Objects and Agreements
- 3930 • Availability
- 3931 • Response Time
- 3932 • Throughput
- 3933 • Fault and Exception Management

3934 **5.3.4 Service Life-cycle**

3935 Managing a service's life cycle involves managing the establishment of the service,
3936 managing its steady-state performance, and managing its termination. The most
3937 obvious feature of this is that a service cannot manage its own life cycle (imagine asking
3938 a non-functioning service to start). Another important consideration is that services may
3939 have resource requirements that must be established at various points in the services'
3940 life cycles. These dependencies may take the form of other services being established;
3941 possibly even services that are not exposed by the service's own interface.

3942 **5.4 SOA Testing Model**

3943 *Program testing can be used to show the presence of bugs,*
3944 *but never to show their absence!*
3945 Edsger Dijkstra

3946 Testing for SOA combines the typical challenges of software testing and certification
3947 with the additional needs of accommodating the distributed nature of the resources, the
3948 greater access of a more unbounded consumer population, and the desired flexibility to
3949 create new solutions from existing components over which the solution developer has
3950 little if any control. The purpose of testing is to demonstrate a required level of reliability,
3951 correctness, and effectiveness that enable prospective consumers to have adequate
3952 confidence in using a service. Adequacy is defined by the consumer based on the
3953 consumer's needs and context of use. As the Dijkstra quote points out, absolute
3954 correctness and completeness cannot be proven by testing; however, for SOA, it is
3955 critical for the prospective consumer to know what testing has been performed, how it
3956 has been performed, and what were the results.

3957 **5.4.1 Traditional Software Testing as Basis for SOA Testing**

3958 SOA services are largely software artifacts and can leverage the body of experience
3959 that has evolved around software testing. IEEE-829 specifies the basic set of software

3960 test documents while allowing flexibility for tailored use. As such, the document
 3961 structure can also provide guidance to SOA testing.

3962 IEEE-829 covers test specification and test reporting through use of the following
 3963 document types:

- 3964 • *Test plan* documenting the scope (what [will-is to](#) be tested, both which entity and
 3965 what features of the entity), the approach (how it [will-beis](#) tested), and the needed
 3966 resources (who [will-dedoes](#) the testing, for how long), with details contained in the:
- 3967 • *Test design specification*: features to be tested, test conditions (e.g. test cases,
 3968 test procedures needed) and expected results (criteria for passing test); entrance
 3969 and exit criteria
- 3970 • *Test case specification*: test data used for input and expected output
- 3971 • *Test procedure specification*: steps required to run the test, including any set-up
 3972 preconditions
- 3973 • *Test item transmittal* to identify the test items being transmitted for testing
- 3974 • *Test log* to record what occurred during test, i.e. which tests run, who ran, what
 3975 order, what happened
- 3976 • *Test incident report* to capture any event that happened during test which requires
 3977 further investigation
- 3978 • *Test summary* as a management report summarizing test run and results,
 3979 conclusions

3980 In summary, IEEE-829 captures (1) what was tested, (2) how it was tested, e.g. the test
 3981 procedure used, and (3) the results of the test.

3982 5.4.1.1 Types of Testing

3983 There are numerous aspects of testing that, in total, work to establish that an entity is
 3984 (1) built as required per policies and related specifications prescribed by the entity's
 3985 owner, and (2) delivers the functionality required by its intended users. This is often
 3986 referred to as verification and validation.

3987 Policies, as described in Section 4.4, that are related to testing may prescribe but are
 3988 not limited to the business processes to be followed, the standards with which an
 3989 implementation must comply, and the qualifications of and restrictions on the users. In
 3990 addition to the functional requirements prescribing what an entity does, there may also
 3991 be non-functional performance and/or quality metrics that state how well the entity does
 3992 it. The relation of these policies to SOA testing is discussed further below.

3993 The identification of policies is the purview of governance (section 5.1) and the assuring
 3994 of compliance (including response to noncompliance) with policies is a matter for
 3995 management (section **Error! Reference source not found.**).

3996 5.4.1.2 Range of Test Conditions

3997 Test conditions and expected responses are detailed in the test case specification. The
 3998 test conditions should be designed to cover the areas for which the entity's response
 3999 must be documented and may include:

4000 • nominal conditions

4001 • boundaries and extremes of expected conditions
4002 • breaking point where the entity has degraded below a certain level or has
4003 otherwise ceased effective functioning
4004 • random conditions to investigate unidentified dependencies among combinations
4005 of conditions
4006 • errors conditions to test error handling
4007 The specification of how each of these conditions should be tested for SOA resources,
4008 including the infrastructure elements of the SOA ecosystem, is beyond the scope of this
4009 [Reference Architecture document](#) but is an area that ~~will~~ evolves along with operational
4010 SOA experience.

4011 **5.4.1.3 Configuration Management of Test Artifacts**

4012 The test item transmittal provides an unambiguous identification of the entity being
4013 tested, thus REQUIRING that the configuration of the entity is appropriately tracked and
4014 documented. In addition, the test documents (such as those specified by IEEE-829)
4015 MUST also be under a documented and appropriately audited configuration
4016 management process, as should other resources used for testing. The description of
4017 each artifact would follow the general description model as discussed in section 4.1.1.1;
4018 in particular, it would include a version number for the artifact and reference to the
4019 documentation describing the versioning scheme from which the version number is
4020 derived.

4021
4022 [EDITOR'S NOTE: TO WHAT EXTENT SHOULD CM BE EXPLICITLY INCLUDED IN THE MANAGEMENT
4023 SECTION?]

4024 **5.4.2 Testing and the SOA Ecosystem**

4025 [EDITOR'S NOTE: THE EMPHASIS THOUGH MUCH OF THE RA IS THE LARGER ECOSYSTEM BUT WE NEED
4026 WORDS IN SECTION 3 TO ACKNOWLEDGE THE EXISTENCE OF THE ENTERPRISE AND THAT AN
4027 ENTERPRISE (AS COMMONLY INTERPRETED) IS LIKELY MORE CONSTRAINED AND MORE PRECISELY
4028 DESCRIBED FOR THE CONTEXT OF THE ENTERPRISE. THE ECOSYSTEM PERSPECTIVE, THOUGH, IS
4029 STILL APPLICABLE FOR THE FOLLOWING REASONS:

- 4030
- 4031 1. A GIVEN ENTERPRISE MAY COMPRISE NUMEROUS CONSTITUENT ENTERPRISES THAT
4032 RESEMBLE THE INDEPENDENT ENTITIES DESCRIBED FOR THE ECOSYSTEM. AN ENTERPRISE
4033 MAY ATTEMPT TO REDUCE VARIATIONS AMONG THE CONSTITUENTS BUT THE [ECOSYSTEM](#)
4034 [VIEW PARTICIPATION IN A SOA ECOSYSTEM VIEW](#) ENABLES SOA TO BENEFIT THE ENTERPRISE
4035 WITHOUT REQUIRING THE ENTERPRISE ISSUES TO BE FULLY RESOLVED.
 - 4036 2. RESOURCES SPECIFICALLY MOTIVATED BY THE CONTEXT OF THE ENTERPRISE CAN BE MORE
4037 READILY USED IN A DIFFERENT CONTEXT IF ECOSYSTEM CONSIDERATIONS ARE INCLUDED AT
4038 AN EARLY STAGE. THE CHANGE IN A CONTEXT MAY BE A FUNDAMENTAL CHANGE IN THE
4039 ENTERPRISE OR THE NEWLY DISCOVERED APPLICABILITY OF ENTERPRISE RESOURCES TO USE
4040 OUTSIDE THE ENTERPRISE.

4041
4042 IN THIS [REFERENCE ARCHITECTURE DOCUMENT](#), REFERENCE TO THE SOA ECOSYSTEM APPLIES BUT
4043 WITH POSSIBLY LESS GENERALITY TO AN ENTERPRISE USE OF SOA.]

4044 Testing of SOA artifacts for use in the SOA ecosystem differs from traditional software
4045 testing for several reasons. First, a highly touted benefit of SOA is to enable
4046 unanticipated consumers to make use of services for unanticipated purposes.
4047 Examples of this could include the consumer using a service for a result that was not

4048 considered the primary one by the provider, or the service may be used in combination
4049 with other services in a scenario that is different from the one considered when
4050 designing for the initial target consumer community. It is unlikely that a new consumer
4051 will push the services back to anything resembling the initial test phase to test the new
4052 use, and thus additional paradigms for testing are necessary. Some testing may
4053 depend on the availability of test resources made available as a service outside the
4054 initial test community, while some testing is likely to be done as part of limited use in the
4055 operational setting. The potential responsibilities related to such "consumer testing" is
4056 discussed further below.

4057 Secondly, in addition to consumers who interact with a service to realize the described
4058 real world effects, the developer community is also intended to be a consumer. In the
4059 SOA vision of reuse, the developer ~~will~~ [composes](#) new solutions using existing services,
4060 where the existing services provides access to some desired real world effects that are
4061 needed by the new solution. The new solution is a consumer of the existing services,
4062 enabling repeated interactions with the existing services playing the role of reusable
4063 components. Note, those components are used at the locations where they individually
4064 reside and are not typically duplicated for the new solution. The new solution may itself
4065 be offered as a SOA service, and a consumer of the service composition representing
4066 the new solution may be totally unaware of the component services being used. (See
4067 section 4.3.4 for further discussion on service compositions.)

4068 Another difference from traditional testing is that the distributed, unbounded nature of
4069 the SOA ecosystem makes it unlikely to have an isolated test environment that
4070 duplicates the operational environment. A traditional testing approach often makes use
4071 of a test system that is identical to the eventual operational system but isolated for
4072 testing. After testing is successfully completed, the tested entity would be migrated to
4073 the operational environment, or the test environment may be delivered as part of the
4074 system to become operational. This is not feasible for the SOA ecosystem as a whole.

4075 SOA services must be testable in the environment and under the conditions that can be
4076 encountered in the operational SOA ecosystem. As the ecosystem is in a state of
4077 constant change, so some level of testing is continuous through the lifetime of the
4078 service, leveraging utility services used by the ecosystem infrastructure to monitor its
4079 own health and respond to situations that could lead to degraded performance. This
4080 implies the test resources must incorporate aspects of the SOA paradigm, and a
4081 category of services may be created to specifically support and enable effective
4082 monitoring and continuous testing for resources participating in the SOA ecosystem.

4083 While SOA within an enterprise may represent a more constrained and predictable
4084 operational environment, the composability and unanticipated use aspects are highly
4085 touted within the enterprise. The expanded perspective on testing may not be as
4086 demanding within an enterprise but fuller consideration of the ecosystem enables the
4087 enterprise to be more responsive should conditions change.

4088 5.4.3 Elements of SOA Testing

4089 IEEE-829 identifies fundamental aspects of testing, and many of these should carry
4090 over to SOA testing: in particular, the identification of what is to be tested, how it is to be
4091 tested, and by whom the testing is to be done. While IEEE-829 identifies a suggested

4092 | document tree, the availability of these documents in the SOA ecosystem is ~~an~~
4093 | ~~additional matter of concern that will be~~ discussed below.

4094 | **5.4.3.1 What is to be Tested**

4095 | The focus of this discussion is the SOA service. It is recognized that the infrastructure
4096 | components of any SOA environment are likely to also be SOA services and, as such,
4097 | ~~will fall~~s under the same testing guidance. Other resources that contribute to a SOA
4098 | environment may not be SOA services, but ~~will be~~are expected to satisfy the intent if not
4099 | the letter of guidance presented here. Specific differences for such resources are as
4100 | yet largely undefined and further elaboration is beyond the scope of ~~this Reference~~
4101 | ~~Architecture~~the SOA-RAF.

4102 | The following discussion often focuses on a singular SOA service but it is implicit that
4103 | any service may be a composite of other services. As such, testing the functionality of a
4104 | composite service may effectively be testing an end-to-end business process that is
4105 | being provided by the composite service. If new versions are available for the
4106 | component services, appropriate end-to-end testing of the composite may be required
4107 | in order to verify that the composite functionality is still adequately provided. The level
4108 | of required testing of an updated composite ~~will depend~~s on policies of those providing
4109 | the service, policies of those using the service, and mission criticality of those
4110 | depending on the service results.

4111 | The SOA service to be tested MUST be unambiguously identified as specified by its
4112 | applicable configuration management scheme. Specifying such a scheme is beyond
4113 | the scope of ~~this Reference Architecture~~the SOA-RAF other than to say the scheme
4114 | should be documented and itself under configuration management.

4115 | **5.4.3.1.1 Origin of Test Requirements**

4116 | In the Service Description model (Figure 21), the aspects of a service that need to be
4117 | described are:

- 4118 | • the service functionality and technical assumptions that underlie the functionality;
- 4119 | • the policies that describe conditions of use;
- 4120 | • the service interface that defines information exchange with the service;
- 4121 | • service reachability that identifies how and where message exchange is to occur;
- 4122 | and
- 4123 | • metrics access for any participant to have information on how a service is
- 4124 | performing.

4125 | Service testing must provide adequate assurance that each of these aspects is
4126 | operational as defined.

4127 | The information in the service description comes from different sources. The
4128 | functionality is defined through whatever process identifies needs and the community
4129 | for which these needs ~~will be~~are addressed. The process may be ad hoc as serves the
4130 | prospective service owner or strictly governed, but defining the functionality is an
4131 | essential first step in development. It is also an early and ongoing focus of testing to
4132 | ensure the service accurately reflects the described functionality and the described
4133 | functionality accurately addresses the consumer needs.

4134 Policies define the conditions of development and conditions of use for a service and
4135 are typically specified as part of the governance process. Policies constraining service
4136 development, such as coding standards and best practices, require appropriate testing
4137 and auditing during development to ensure compliance. While the governance process
4138 | ~~will-identifiesy~~ development policies, these are likely to originate from the technical
4139 community responsible for development activities. Policies that define conditions of use
4140 often define business practices that service owners and providers or those responsible
4141 for the SOA infrastructure want followed. These policies are initially tested during
4142 service development and are continuously monitored during the operational lifetime of
4143 the service.

4144 The testing of the service interface and service reachability are often related but
4145 essentially reflect different motivations and needs. The service interface is specified as
4146 a joint product of the service owners and providers who define service functionality, the
4147 prospective consumer community, the service developer, and the governance process.
4148 The semantics of the information model must align with the semantics of those who
4149 consume the service in order for there to be meaningful exchange of information. The
4150 structure of the information is influenced by the consumer semantics and the
4151 requirements and constraints of the representation as interpreted by the service
4152 developer. The service process model that defines actions which can be performed
4153 against a service and any temporal dependencies derive from the defined functionality
4154 and may be influenced by the development process. Any of these constraints may be
4155 identified and expressed as policy through the governance process.

4156 Service reachability conditions are the purview of the service provider who identifies the
4157 service endpoint and the protocols recognized at the endpoint. These may be
4158 constrained by governance decisions on how endpoint addresses may be allocated and
4159 what protocols should be used.

4160 While the considerations for defining the service interface derive from several sources,
4161 testing of the service interface is more straightforward and isolated in the testing
4162 process. At any point where the interface is modified or exposes a new resource, the
4163 message exchange should be monitored both to ensure the message reaches its
4164 intended destination and it is parsed correctly once received. **Once an interface has**
4165 | **been shown to function properly, it is unlikely ~~it will~~ fail later unless something**
4166 **fundamental to the service changes.**

4167 The service interface is also tested when the service endpoint changes. Testing of the
4168 endpoint ensures message exchange can occur at the time of testing and the initial
4169 testing shows the interface is being processed properly at the new endpoint.
4170 Functioning of a service endpoint at one time does not guarantee it is functioning at
4171 another time, e.g. the server with the endpoint address may be down, making testing of
4172 service reachability a continual monitoring function through the life of the service's use
4173 of the endpoint. Also, while testing of the service endpoint is a necessary and most
4174 commonly noted part of the test regiment, it is not in itself sufficient to ensure the other
4175 aspects of testing discussed in this section.

4176 Finally, governance is impossible without the collection of metrics against which service
4177 behavior can be assessed. Metrics are also a key indicator for consumers to decide if a
4178 service is adequate for their needs. For instance, the average response time or the
4179 recent availability can be determining factors even if there are no rules or regulations

Comment [PFB59]: A very bold assumption!

4180 promulgated through the governance process against which these metrics are
4181 assessed. The available metrics are a combination of those expected by the consumer
4182 community and those mandated through the governance process. The total set of
4183 metrics will evolve over time with SOA experience. Testing of the services that gather
4184 and provide access to the metrics will follow testing as described in this section, but for
4185 an individual service, testing will ensure that the metrics access indicated in the service
4186 description is accurate.

4187 The individual test requirements highlight aspects of the service that testing must
4188 consider but testing must establish more than isolated behavior. The emphasis is the
4189 holistic results of interacting with the service in the SOA environment. Recall that the
4190 execution context is the set of agreements between a consumer and a provider that
4191 define the conditions under which service interaction occurs. The agreements are
4192 expected to be predominantly the acceptance of the standard conditions as enumerated
4193 by the service provider, but it may include the identification of alternate conditions that
4194 will govern the interaction.

4195 For example, the provider may prefer a policy where it can sell the contact information
4196 of its consumers but will honor the request of a consumer to keep such information
4197 private. The identification of the alternate privacy policy is part of the execution context,
4198 and it is the application of and compliance with this policy that operational monitoring
4199 will attempt to measure. The collection of metrics showing this condition is indeed met
4200 when chosen is considered part of the ongoing testing of the service.

4201 Other variations in the execution context also require monitoring to ensure that different
4202 combinations of conditions perform together as desired. For example, if a new privacy
4203 policy takes additional resources to apply, this may affect quality of service and
4204 propagate other effects. These could not be tested during the original testing if the
4205 alternate policy did not exist at that time.

4206 **5.4.3.1.2 Testing Against Non-Functional Requirements**

4207 Testing against non-functional requirements constitutes testing of business usability of
4208 the service. In a marketplace of services, non-functional characteristics may be the
4209 primary differentiator between services that produce essentially the same real world
4210 effects.

4211 As noted in the previous section, non-functional characteristics are often associated
4212 with policies or other terms of use and may be collected in service level contracts
4213 offered by the service providers. Non-functional requirements may also reflect the
4214 network and hardware infrastructure that support communication with the service, and
4215 changes may impact quality of service. The service consumer and even the service
4216 provider may not be aware of all such infrastructure changes but the changes may
4217 manifest in shared states that impact the usability of the service.

4218 In general, a change in the non-functional requirements results in a change to the
4219 execution context, but as with any collection of information that constitutes a
4220 description, the execution context is unable to explicitly capture all non-functional
4221 requirements that may apply. A change in non-functional requirements, whether
4222 explicitly part of the execution context or an implicit contributor, may require retesting of
4223 the service even if its functionality and the implementation of the functionality has not
4224 changed. Depending on the circumstances, retesting may require a formal recertifying

4225 of end-to-end behavior or more likely will be part of the continuous monitoring that
4226 applies throughout the service lifetime.

4227 **5.4.3.1.3 Testing Content and the Interests of Consumers**

4228 As noted in section 5.4.1.1, testing may involve verification of conformance with respect
4229 to policies and technical specifications and validation with respect to sufficiency of
4230 functionality to meet some prescribed use. It may also include demonstration of
4231 performance and quality aspects. For some of these items, such as demonstrating the
4232 business processes followed in developing the service or the use of standards in
4233 implementing the service, the testing or relevant auditing is done internal to the service
4234 development process and follows traditional software testing and quality assurance. If it
4235 is believed of value to potential consumers, information about such testing could be
4236 included in the service description. However, it is not required that all test or
4237 compliance artifacts be available to consumers, as many of the details tested may be
4238 part of the opacity of the service implementation.

4239 Some aspects of the service being tested will reflect directly on the real world effects
4240 realized through interaction with the service. In these cases, it is more likely that testing
4241 results will be directly relevant to potential consumers. For example, if the service was
4242 designed to correspond to certain elements of a business process or that a certain
4243 workflow is followed, testing should verify that the real world effects reflect that the
4244 business process or workflow were satisfactorily captured.

4245 The testing may also need to demonstrate that specified conditions of use are satisfied.
4246 For example, policies may be asserted that require certain qualifications of or impose
4247 restrictions on the consumers who may interact with the service. The service testing
4248 must demonstrate that the service independently enforces the policies or it provides the
4249 required information exchanges with the SOA ecosystem so other resources can ensure
4250 the specified conditions.

4251 The completeness of the testing, both in terms of the features tested and the range of
4252 parameters for which response is tested, depends on the context of expected use: the
4253 more critical the use, the more complete the testing. There are always limits on the
4254 resources available for testing, if nothing else than the service must be available for use
4255 in a finite amount of time.

4256 This again emphasizes the need for adequate documentation to be available. If the
4257 original testing is very thorough, it may be adequate for less demanding uses in the
4258 future. If the original testing was more constrained, then well-documented test results
4259 establish the foundation on which further testing can be defined and executed.

4260 **5.4.3.2 How Testing is to be Done**

4261 Testing should follow well-defined methodologies and, if possible, should reuse test
4262 artifacts that have proven generally useful for past testing. For example, IEEE-829
4263 notes that test cases are separated from test designs to allow for use in more than one
4264 design and to allow for reuse in other situations. In the SOA ecosystem, description of
4265 such artifacts, as with description of a service, enables awareness of the item and
4266 describes how the artifact may be accessed or used.

4267 As with traditional testing, the specific test procedures and test case inputs are
4268 important so the tests are unambiguously defined and entities can be retested in the
4269 future. Automated testing and regression testing may be more important in the SOA
4270 ecosystem in order to re-verify a service is still acceptable when incorporated in a new
4271 use. For example, if a new use requires the services to deal with input parameters
4272 outside the range of initial testing, the tests could be rerun with the new parameters. If
4273 the testing resources are available to consumers within the SOA ecosystem, the testing
4274 as designed by test professionals could be consumed through a service accessed by
4275 consumers, and their results could augment those already in place. This is discussed
4276 further in the next section.

4277 **5.4.3.3 Who Performs the Testing**

4278 As with any software, the first line of testing is unit testing done by software developers.
4279 It is likely that initial testing will be done by those developing the software but may also
4280 be done independently by other developers. For SOA development, unit testing is likely
4281 confined to a development sandbox isolated from the SOA ecosystem.

4282 SOA testing will differ from traditional software testing in that testing beyond the
4283 development sandbox must incorporate aspects of the SOA ecosystem, and those
4284 doing the testing must be familiar with both the characteristics and responses of the
4285 ecosystem and the tools, especially those available as services, to facilitate and
4286 standardize testing. Test professionals will know what level of assurance must be
4287 established as the exposure of the service to the ecosystem and ecosystem to the
4288 service increases towards operational status. These test professionals may be internal
4289 resources to an organization or may evolve as a separate discipline provided through
4290 external contracting.

4291 As noted above, it is unlikely that a complete duplicate of the SOA ecosystem will be
4292 available for isolated testing, and thus use of ecosystem resources will manifest as a
4293 transition process rather than a step change from a test environment to an operational
4294 one. This is especially true for new composite services that incorporate existing
4295 operational services to achieve the new functionality. The test professionals will need to
4296 understand the available resources and the ramifications of this transition.

4297 As with current software development, a stage beyond work by test professionals will
4298 make use of a select group of typical users, commonly referred to as beta testers, to
4299 report on service response during typical intended use. This establishes fitness by the
4300 consumers, providing final validation of previously verified processes, requirements, and
4301 final implementation.

4302 In traditional software development, beta testing is the end of testing for a given version
4303 of the software. However, although the initial test phase can establish an appropriate
4304 level of confidence consistent with the designed use for the initial target consumer
4305 community, the operational service will exist in an evolving ecosystem, and later
4306 conditions of use may differ from those thought to be sufficient during the initial testing.
4307 Thus, operational monitoring becomes an extension of testing through the service
4308 lifetime. This continuous testing will attempt to ensure that a service does not consume
4309 an inordinate amount of ecosystem resources or display other behavior that degrades
4310 the ecosystem, but it will not undercover functional errors that may surface over time.

4311 As with any software, it is the responsibility of the consumers to consider the
4312 reasonableness of solutions in order to spot errors in either the software or the way the
4313 software is being used. This is especially important for consumers with unanticipated
4314 uses that may go beyond the original test conditions. It is unlikely the consumers will
4315 initiate a new round of formal testing unless the new use requires a significantly higher
4316 level of confidence in the service. Rather the consumer becomes a new extension to
4317 the testing regiment. Obvious testing would include a sanity check of results during the
4318 new use. However, if the details of legacy testing are associated with the service
4319 through the service description and if testing resources are available through automated
4320 testing services, then the new consumers can rerun and extend previous testing to
4321 include the extended test conditions. If the test results are acceptable, these can be
4322 added to the documentation of previous results and become the extended basis for
4323 future decisions by prospective consumers on the appropriateness of the service. If the
4324 results are not acceptable or in some way questionable, the responsible party for the
4325 service or testing professionals can be brought in to decide if remedial action is
4326 necessary.

4327 5.4.3.4 How Testing Results are Reported

4328 For any SOA service, an accurate reporting of the testing a service has undergone and
4329 the results of the testing is vital to consumers deciding whether a service is appropriate
4330 for intended use. Appropriateness may be defined by a consumer organization and
4331 require specific test regiments culminating in a certification; appropriateness could be
4332 established by accepting testing and certifications that have been conferred by others.

4333 The testing and certification information should be identified in the service description.
4334 Referring to the general description model of [Figure 15](#)~~Figure 14~~, tests conducted by or
4335 under a request from the service owner (see ownership in section **Error! Reference**
4336 **source not found.**) would be captured under Annotations from Owners. Testing done
4337 by others, such as consumers with unanticipated uses, could be associated through
4338 Annotations from 3rd Parties. The annotations should clearly indicate what was tested,
4339 how the testing was done, who did the testing, and the testing results. The clear
4340 description of each of these artifacts and of standardized testing protocols for various
4341 levels of sophistication and completeness of testing would enable a common
4342 understanding and comparison of test coverage. It will also make it more
4343 straightforward to conduct and report on future testing, facilitating the maintenance of
4344 the service description.

4345 Consumer testing and the reporting of results raises additional issues. While stating
4346 who did the testing is mandatory, there may be formal requirements for authentication of
4347 the tester to ensure traceability of the testing claims. In some circumstances, persons
4348 or organizations would not be allowed to state testing claims unless the tester was an
4349 approved entity. In other cases, ensuring the tester had a valid email may be sufficient.
4350 In either case, it would be at the discretion of the potential consumer to decide what
4351 level of authentication was acceptable and which testers are considered authoritative in
4352 the context of their anticipated use.

4353 Finally, in a world of openly shared information, we would see an ever-expanding set of
4354 testing information as new uses and new consumers interact with a service. In reality,
4355 these new uses may represent proprietary processes or classified use that should only

4356 be available to authorized parties. Testing information, as with other elements of
4357 description, may require special access controls to ensure appropriate access and use.

4358 **5.4.4 Testing SOA Services**

4359 Testing of SOA services should be consistent with the SOA paradigm. In particular,
4360 testing resources and artifacts should be visible in support of service interaction
4361 between providers and consumers, where here the interaction is between the testing
4362 resource and the tester. In addition, the idea of opacity of the implementation should
4363 limit the details that need to be available for effective use of the test resources. Testing
4364 that requires knowledge of the internal structure of the service or its underlying
4365 capability should be performed as part of unit testing in the development sandbox, and
4366 should represent a minimum level of confidence before the service begins its transition
4367 to further testing and eventual operation in the SOA ecosystem.

4368 **5.4.4.1 Progression of SOA Testing**

4369 Software testing is a gradual exercise going from micro inspection to testing macro
4370 effects. The first step in testing is likely the traditional code reviews. SOA
4371 considerations would account for the distributed nature of SOA, including issues of
4372 distributed security and best practices to ensure secure resources. It would also set the
4373 groundwork for opacity of implementation, hiding programming details and simplifying
4374 the use of the service.

4375 Code review is likely followed by unit testing in a development sandbox isolated from
4376 the operational environment. The unit testing is done with full knowledge of the service
4377 internal structure and knowledge of resources representing underlying capabilities. It
4378 tests the interface to ensure exchanged messages are as specified in the service
4379 description and the messages can be parsed and interpreted as intended. Unit testing
4380 also verifies intended functionality and that the software has dealt correctly with internal
4381 dependencies, such as structure of a file system or access to other dedicated
4382 resources.

4383 Some aspects of unit testing require external dependencies be satisfied, and this is
4384 often done using mock objects to substitute for the external resources. In particular, it
4385 will likely be necessary to include mocks of existing operational services, both those
4386 provided as part of the SOA infrastructure and services from other providers.

4387 **Service Mock**

4388 A service mock is an entity that mimics some aspect of the performance of an
4389 operational service without committing to the real world effects that the
4390 operational service would produce.

4391 Mocks are discussed in detail in sections 5.4.4.3 and 5.4.4.4.

4392 After unit testing has demonstrated an adequate level of confidence in the service, the
4393 testing must transition from the tightly controlled environment of the development
4394 sandbox to an environment that more clearly resembles the operational SOA ecosystem
4395 or, at a minimum, the intended enterprise. While sandbox testing will use simple mocks
4396 of some aspects of the SOA environment, such as an interface to a security service
4397 without the security service functionality, the dynamic nature of SOA makes a full
4398 simulation infeasible to create or maintain. This is especially true when a new

4399 composite service makes use of operational services provided by others. Thus, at
4400 some point before testing is complete, the service will need to demonstrate its
4401 functionality by using resources and dealing with conditions that only exist in the full
4402 ecosystem or the intended enterprise. Some of these resources may still provide test
4403 interfaces -- more on this below -- but the interfaces will be accessible using the SOA
4404 environment and not just implemented for the sandbox.

4405 At this stage, the opacity of the service becomes important as the details of interacting
4406 with the service now rely on correct use of the service interface and not knowledge of
4407 the service internals. The workings of the service will only be observable through the
4408 real world effects realized through service interactions and external indications that
4409 conditions of use, such as user authentication, are satisfied. Monitoring the behavior of
4410 the service will depend on service interfaces that expose internal monitoring or provide
4411 required information to the SOA infrastructure monitoring function. The monitoring
4412 required to test a new service is likely to have significant overlap with the monitoring the
4413 SOA infrastructure includes to monitor its own health and to identify and isolate
4414 behavior outside of acceptable bounds. This is exactly what is needed as part of
4415 service testing, and it is reasonable to assume that the ecosystem transition includes
4416 use of operational monitoring rather than solely dedicated monitoring for each service
4417 being tested.

4418 Use of SOA monitoring resources during the explicit testing phase sets the stage for
4419 monitoring and a level of continual testing throughout the service lifetime.

4420 **5.4.4.2 Testing Traditional Dependencies vs. Service Interactions**

4421 A SOA service is not required to make use of other operational services beyond what
4422 may be required for monitoring by the ecosystem infrastructure. The service can
4423 implement hardcoded dependencies which have been tested in the development
4424 sandbox through the use of dedicated mocks. While coordination may be required with
4425 real data sources during integration testing, the dependencies can be constrained to
4426 things that can be tested in a more traditional manner. Policies can also be set to
4427 restrict access to pre-approved users, and thus the question of unanticipated users and
4428 unanticipated uses can be eliminated. Operational readiness can be defined in terms of
4429 what can be proven in isolated testing. While all this may provide more confidence in
4430 the service for its designed purpose, such a service will not fully participate in the
4431 benefits or challenges of the ecosystem. This is akin to filling a swimming pool with sea
4432 water and having someone in the pool say they are swimming in the ocean.

4433 In considering the testing needed for a fully participating service, consider the example
4434 of a new composite service that combines the real world effects and complies with the
4435 conditions of use of five existing operational services. The developer of the composite
4436 service does not own any of the component services and has limited, if any, ability to
4437 get the distributed owners to do any customization. The developer also is limited by the
4438 principle of opacity to information comprising the service description, and does not know
4439 internal details of the component services. The developer of the composite service
4440 must use the component services as they exist as part of the SOA environment,
4441 including what is provided to support testing by new users. This introduces
4442 requirements for what is needed in the way of service mocks.

4443 5.4.4.3 Use of Service Mocks

4444 Service mocks enables the tested service to respond to specific features of an
4445 operational service that is being used as a component. It allows service testing to
4446 proceed without needing access to or with only limited engagement with the component
4447 service. Mocks can also mimic difficult to create situations for which it is desired to test
4448 the new service response. For composite services using multiple component services,
4449 mocks may be used in combination to function for any number of the components.
4450 Note, when using service mocks, it is important to remember that it is not the
4451 component service that is being tested (although anomalous behavior may be
4452 uncovered during testing) but the use of the component in the new composite.

4453 Individual service mocks can emphasize different features of the component service
4454 they represent but any given mock does not have to mimic all features. For example, a
4455 mock of the service interface can echo a sent message and demonstrate the message
4456 is reaching its intended destination. A mock could go further and parse the sent
4457 message to demonstrate the message not only reached its destination but was
4458 understood. As a final step, the mock could report back what actions would have been
4459 taken by the component service and what real world effects would result. If the
4460 response mimicked the operational response, functional testing could proceed as if the
4461 real world effect actually occurred.

4462 There are numerous ways to provide mock functionality. The service mock could be a
4463 simulation of the operational service and return simulated results in a realistic response
4464 message or event notification. It is also possible for the operational service to act as its
4465 own mock and simply not execute the commit stage of its functionality. The service
4466 mock could use a combination of simulation and service action without commit to
4467 generate a report of what would have occurred during the defined interaction with the
4468 operational service.

4469 As the service proceeds through testing, mocks should be systematically replaced by
4470 the component resources accessed through their operational interfaces. Before beta
4471 testing begins, end-to-end testing, i.e. proceeding from the beginning of the service
4472 interaction to the resulting real world results, should be accomplished using component
4473 resources via their operational interfaces.

4474 5.4.4.4 Providers of Service Mocks

4475 In traditional testing, it is often the test professionals who design and develop the
4476 mocks, but in the distributed world of SOA, this may not be efficient or desirable.

4477 In the development sandbox, it is likely the new service developer or test professionals
4478 working with the developer will create mocks adequate for unit testing. Given that most
4479 of this testing is to verify the new service is performing as designed, it is not necessary
4480 to have high fidelity models of other resources being accessed. In addition, given
4481 opacity of SOA implementation, the developer of the new service may not have
4482 sufficient detailed knowledge of a component service to build a detailed mock of the
4483 component service functionality. Sharing existing mocks at this stage may be possible
4484 but the mocks would need to be implemented in the sandbox, and for simple models it
4485 is likely easier to build the mock from scratch.

4486 As testing begins its transition to the wider SOA environment, mocks may be available
4487 as services. For existing resources, it is possible that an Open Source model could
4488 evolve where service mocks of available functions can be catalogued and used during
4489 initial interaction of the tested service and the operational environment. Widely used
4490 functions may have numerous service mocks, some mimicking detailed conditions
4491 within the SOA infrastructure. However, the Open Source model is less likely to be
4492 sufficient for specialty services that are not widely used by a large consumer
4493 community.

4494 The service developer is probably best qualified for also developing more detailed
4495 service mocks or for mock modes of operational services. This implies that in addition
4496 to their operational interfaces, services will routinely provide test interfaces to enable
4497 service mocks to be used as services. As noted above, a new service developer
4498 wanting to build a mock of component services is limited to the description provided by
4499 the component service developer or owner. The description typically will detail real
4500 world effects and conditions of use but will not provide implementation details, some of
4501 which may be proprietary. Just as important in the SOA ecosystem, if it becomes
4502 standard protocol for developers to create service mocks of their own services, a new
4503 service developer is only responsible for building his own mocks and can expect other
4504 mocks to be available from other developers. This reduces duplication of effort where
4505 multiple developers would be trying to build the same mocks from the same insufficient
4506 information. Finally, a service developer is probably best qualified to know when and
4507 how a service mock should be updated to reflect modified functionality or message
4508 exchange.

4509 It is also possible that testing organizations will evolve to provide high-fidelity test
4510 harnesses for new services. The harnesses would allow new services to plug into a test
4511 environment and would facilitate accessing mocks of component services. However, it
4512 will remain a constant challenge for such organizations to capture evolving uses and
4513 characteristics of service interactions in the real SOA environment and maintain the
4514 fidelity and accuracy of the test systems.

4515 **5.4.4.5 Fundamental Questions for SOA Testing**

4516 In order for the transition to the SOA operational environment to proceed, it is necessary
4517 to answer two fundamental questions:

- 4518 • Who provides what testing resources for the SOA operational environment, e.g.
4519 mocks of interfaces, mocks of functionality, monitoring tools?
- 4520 • What testing needs to be accomplished before operational environment
4521 resources can be accessed for further testing?

4522 The discussion in section 5.4.4.4 notes various levels of sophistication of service mocks
4523 and different communities are likely to be responsible for different levels. Section
4524 5.4.4.4 advocates a significant role for service developers, but there needs to be
4525 community consensus that such mocks are needed and that service developers will
4526 agree to fulfilling this role. There is also a need for consensus as to what tools should
4527 be available as services from the SOA infrastructure.

4528 As for use of the service mocks and SOA environment monitoring services, practical
4529 experience is needed upon which guidelines can be established for when a new service

4530 has been adequately tested to proceed with a greater level of exposure with the SOA
4531 environment. Malfunctioning services could cause serious problems if they cannot be
4532 identified and isolated. On the other hand, without adequate testing under SOA
4533 operational conditions, it is unlikely that problems can be uncovered and corrected
4534 before they reach an operational stage.

4535 As noted in section 5.4.4.2, some of these questions can be avoided by restricting
4536 services to more traditional use scenarios. However, such restriction will limit the
4537 effectiveness of SOA use and the result will resemble the constraints of traditional
4538 integration activities we are trying to move beyond.

4539 **5.4.5 Architectural Implications for SOA Testing**

4540 The discussion of SOA Testing indicates numerous architectural implications on the
4541 SOA ecosystem:

- 4542 • The distributed, boundary-less nature of the SOA ecosystem makes it
4543 infeasible to create and maintain a single mock of the entire ecosystem to
4544 support testing activities.
- 4545 • A standard suite of monitoring services needs to be defined, developed,
4546 and maintained. This should be done in a manner consistent with the evolving
4547 nature of the ecosystem.
- 4548 • Services should provide interfaces that support access in a test mode.
- 4549 • Testing resources must be described and their descriptions must be
4550 catalogued in a manner that enables their discovery and access.
- 4551 • Guidelines for testing and ecosystem access need to be established and
4552 the ecosystem must be able to enforce those guidelines asserted as policies.
- 4553 • Services should be available to support automated testing and regression
4554 testing.
- 4555 • Services should be available to facilitate updating service description by
4556 anyone who has performed testing of a service.

4557 6 Conformance

4558 | This Reference Architecture [Framework](#) is an abstract architectural description of
4559 | Service Oriented Architecture, which means that it is especially difficult to construct
4560 | tests for conformance to the architecture. In addition, conformance to an architectural
4561 | specification does not, by itself, guarantee any form of interoperability between multiple
4562 | implementations.

4563 | However, it *is* possible to decide whether or not a given architecture is conformant to an
4564 | architectural description such as this one. In discussions of conformance we use the
4565 | term **target architecture** to identify the (typically concrete) architecture that may be
4566 | viewable as conforming to the abstract principles outlined in this [Reference](#)
4567 | [Architecture](#) document.

4568 Target Architecture

4569 | A target architecture is an architectural description of a system that is intended to
4570 | be viewed as conforming to ~~this Reference Architecture~~ [the SOA-RAF](#).

4571 | While we cannot guarantee interoperability between target architectures (or more
4572 | specifically between applications and systems residing within the ecosystems of those
4573 | target architectures), interoperability between target architectures is promoted by
4574 | conformance to this Reference Architecture [Framework](#) as it reduces the semantic
4575 | impedance mismatch between the different ecosystems.

4576 | The primary measure of conformance is whether given concepts as described in ~~this~~
4577 | [Reference Architecture](#) document have corresponding ~~ingene~~ [concepts with-in](#) the target
4578 | architecture. Such a correspondence MUST honor the relationships identified within this
4579 | document for the target architecture to be considered conforming.

4580 | For example, in Section 3.1.5.1 we identify resource as a key concept. A resource is
4581 | associated with an owner and a number of identifiers. For a target architecture to
4582 | conform to ~~this Reference Architecture~~ [the SOA-RAF](#), it must be possible to find
4583 | corresponding concepts of resource, identifier and owner within the target architecture:
4584 | say *entity*, *token* and *user*. Furthermore, the relationships between *entity*, *token* and
4585 | *user* MUST mirror the relationships between resource, identifier and owner
4586 | appropriately.

4587 | Clearly, such correspondence is simpler if the terminology within the target architecture
4588 | is identical to that in ~~this Reference Architecture~~ [the SOA-RAF](#). But so long as the
4589 | 'graph' of concepts and relationships is consistent, that is all that is required for the
4590 | target architecture to conform to this Reference Architecture [Framework](#).

4591 | [EDITOR'S NOTE: The conformance section is not complete]

4592

4593 A. Acknowledgements

4594 The following individuals have participated in the [work of the technical committee](#)
4595 [responsible for](#) creation of this specification and are gratefully acknowledged:

4596 Participants:

4597 [Chris Bashioum, MITRE Corporation](#)

4598 Rex Brooks, Individual Member

4599 Peter Brown, Individual Member

4600 Scott Came, Search Group Inc.

4601 Joseph Chiusano, Booz Allen Hamilton

4602 Robert Ellinger, Northrop Grumman Corporation

4603 David Ellis, Sandia National Laboratories

4604 Jeff A. Estefan, Jet Propulsion Laboratory

4605 Don Flinn, Individual Member

4606 Anil John, Johns Hopkins University

4607 Ken Laskey, MITRE Corporation

4608 Boris Lublinsky, Nokia Corporation

4609 Francis G. McCabe, Individual Member

4610 Christopher McDaniels, USSTRATCOM

4611 Tom Merkle, Lockheed Martin Corporation

4612 Jyoti Namjoshi, Patni Computer Systems Ltd.

4613 Duane Nickull, Adobe Inc.

4614 James Odell, Associate

4615 Michael Poulin, Fidelity Investments

4616 Michael Stiefel, Associate

4617 Danny Thornton, Northrop Grumman

4618 Timothy Vibbert, Lockheed Martin Corporation

4619 Robert Vitello, New York Dept. of Labor

4620 [The committee would particularly like to underline the significant contributions made by](#)
4621 [Rex Brooks, Jeff Estefan, Ken Laskey, Boris Lublinsky, Frank McCabe, Michael Poulin](#)
4622 [and Danny Thornton](#)

4623
4624
4625
4626
4627
4628
4629
4630
4631
4632
4633
4634
4635
4636
4637
4638
4639
4640
4641
4642
4643
4644
4645
4646
4647
4648
4649
4650
4651
4652
4653
4654
4655
4656
4657

B. Glossary Index of Defined Terms

The first page number refers to the first use of the term. The second, where necessary, refers to the page where the term is formally defined.

- Action
- Action Level Real World Effect
- Actor
- Architecture
- Architectural Description
- Authority
- Business Processes
- Capability
- Choreography
- Commitment
- Communicative Action
- Constitution
- Contract
- Delegate
- Description
- Endpoint
- Enterprise
- Governance
- Governance Framework
- Governance Processes
- Identifier
- Identity
- Joint Action
- Leadership
- Life-cycle manageability
- Logical Framework
- Management
- Management Policy
- Management Service
- Manageability Capability
- Message Exchange
- Model**

Comment [PFB60]: List to be checked (some definitions have been deleted in the meantime) and referenced

4658 Obligation
4659 Objective
4660 Operations
4661 Orchestration
4662 Ownership
4663 Ownership Boundary
4664 Participant
4665 Peer
4666 Permission
4667 Policy
4668 Policy Conflict
4669 Policy Conflict Resolution
4670 Policy Constraint
4671 Policy Decision
4672 Policy Enforcement
4673 Policy Framework
4674 Policy Object
4675 Policy Ontology
4676 Policy Owner
4677 Policy Subject
4678 Presence
4679 Private State
4680 Protocol
4681 Public Semantics
4682 Qualification
4683 Real World Effect
4684 Regulation
4685 Resource
4686 Responsibility
4687 Right
4688 Risk
4689 Role
4690 Rule
4691 Security
4692 Semantic Engagement
4693 Service Action
4694 Service Consumer

4695 Service Level Real World Effect
4696 Service Mediator
4697 Service Provider
4698 Shared State
4699 Skill
4700 Social Structure
4701 Stakeholder
4702 State
4703 System
4704 System Stakeholder
4705 Trust
4706 View
4707 Viewpoint

B.C. The Unified Modeling Language, UML

Error! Reference source not found. illustrates an annotated example of a UML class diagram that is used to represent a visual model depiction of the Resources Model in the [Service Ecosystem View Participation in a SOA Ecosystem view](#) (Section 3).

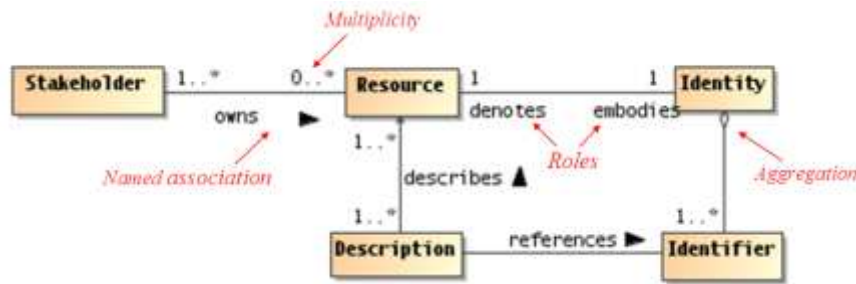


Figure 4847 Example UML class diagram—Resources.

Lines connecting boxes (classifiers) represent associations between things. An association has two roles (one in each direction). A role can have cardinality, for example, one or more ("1..*") stakeholders own zero or more ("0..*") resources. The role from classifier A to B is labeled closest to B, and vice versa, for example, the role between resource to Identity can be read as resource embodies Identity, and Identity denotes a resource.

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

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

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

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

Comment [PFB61]: Bad example and poorly described. 'Multiplicity'? The term is Cardinality. The aggregation symbol is wrong (missing the -> at the opposite end to the diamond). Inheritance symbol should be shown. Should we include n-ary associations, as these are used in new proposed Skill model...

4740 **C.D. Critical Factors Analysis**

4741 A critical factors analysis (CFA) is an analysis of the key properties of a project. A CFA
4742 is analyzed in terms of the goals of the project, the critical factors that will lead to its
4743 success and the measurable requirements of the project implementation that support
4744 the goals of the project. CFA is particularly suitable for capturing quality attributes of a
4745 project, often referred to as “non-functional” or “other-than-functional” requirements: for
4746 example, security, scalability, wide-spread adoption, and so on. As such, CFA
4747 complements rather than attempts to replace other requirements capture techniques.

4748 **C.1D.1 Goals**

4749 A goal is an overall target that you are trying to reach with the project. Typically, goals
4750 are hard to measure by themselves. Goals are often directed at the potential consumer
4751 of the product rather than the technology developer.

4752 **Critical Success Factors**

4753 A critical success factor (CSF) is a property, sub-goal that directly supports a goal and
4754 there is strong belief that without it the goal is unattainable. CSFs themselves are not
4755 necessarily measurable in themselves.

4756 **Requirements**

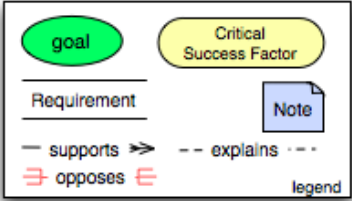
4757 A requirement is a specific measurable property that directly supports a CSF. The key
4758 here is measurability: it should be possible to unambiguously determine if a requirement
4759 has been met. While goals are typically directed at consumers of the specification,
4760 requirements are focused on technical aspects of the specification.

4761 **CFA Diagrams**

4762 It can often be helpful to illustrate graphically the key concepts and relationships
4763 between them. Such diagrams can act as effective indices into the written descriptions
4764 of goals etc., but is not intended to replace the text.

4765 The legend:

4766



4767 illustrates the key elements of the graphical notation. Goals are written in round ovals,
4768 critical success factors are written in round-ended rectangles and requirements are
4769 written using open-ended rectangles. The arrows show whether a
4770

4771 CSF/goal/requirement is supported by another element or opposed by it. This highlights
4772 the potential for conflict in requirements.