



Reference Architecture for Service Oriented Architecture Version 0.0

working-draft, May xx,2007

Specification URIs:

This Version:

[http://docs.oasis-open.org/soa-rm/\[additional path/filename\].html](http://docs.oasis-open.org/soa-rm/[additional path/filename].html)
[http://docs.oasis-open.org/soa-rm/\[additional path/filename\].pdf](http://docs.oasis-open.org/soa-rm/[additional path/filename].pdf)

Previous Version:

[http://docs.oasis-open.org/soa-rm/\[additional path/filename\].html](http://docs.oasis-open.org/soa-rm/[additional path/filename].html)
[http://docs.oasis-open.org/soa-rm/\[additional path/filename\].pdf](http://docs.oasis-open.org/soa-rm/[additional path/filename].pdf)

Latest Version:

[http://docs.oasis-open.org/soa-rm/\[additional path/filename\].html](http://docs.oasis-open.org/soa-rm/[additional path/filename].html)
[http://docs.oasis-open.org/soa-rm/\[additional path/filename\].pdf](http://docs.oasis-open.org/soa-rm/[additional path/filename].pdf)

Latest Approved Version:

[http://docs.oasis-open.org/soa-rm/\[additional path/filename\].html](http://docs.oasis-open.org/soa-rm/[additional path/filename].html)
[http://docs.oasis-open.org/soa-rm/\[additional path/filename\].pdf](http://docs.oasis-open.org/soa-rm/[additional path/filename].pdf)

Technical Committee:

OASIS Service Oriented Architecture Reference Model TC

Chair(s):

Francis G. McCabe

Editor(s):

Jeffrey A. Estefan, Jet Propulsion Laboratory, jeffrey.a.estefan@jpl.nasa.gov
Ken Laskey, MITRE Corporation, klaskey@mitre.org
Francis G. McCabe, Genietown, frankmccabe@mac.com
Danny Thornton, danny.thornton@soamodeling.org

Related work:

This specification is related to:

- OASIS Reference Model for Service Oriented Architecture

Abstract:

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

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

This architecture follows the recommended practice of describing architecture in terms of models, views, and viewpoints, as prescribed in IEEE1471. The RA has three main views: the Business via Service view which lays the foundation for conducting business in the context of Service Oriented Architecture; the Realizing Services view which addresses the requirements for constructing a Service Oriented Architecture; and the Owning Service Oriented Architecture view which focuses on the governance and management of SOA 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/tc_home.php?wg_abbrev=soa-rm).

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

Notices

Copyright © OASIS® 1993–2007. All Rights Reserved. OASIS trademark, IPR and other policies apply.

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 names "OASIS", [insert specific trademarked names and abbreviations here] are trademarks 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	6
1.1	What is a Reference Architecture	6
1.2	Service Oriented Architecture – An Ecosystems perspective	6
1.3	Relationship of the Reference Model	6
1.4	UML Modeling Notation	Error! Bookmark not defined.
1.5	Viewpoints and Views	7
1.6	Viewpoints of this Reference Architecture	8
1.6.1	Business via Services Viewpoint	8
1.6.2	Realizing a Service Oriented Architecture Viewpoint	9
1.6.3	Owning a Service Oriented Architecture Viewpoint	9
1.7	Terminology	9
2	Requirements	10
2.1	Goals of the Reference Architecture	10
2.1.1	Effectiveness	11
2.1.2	Assurance	11
2.1.3	Wide scale adoption	13
3	Business via Services View	15
3.1	Stakeholders and Participants Model	15
3.2	Needs and Capabilities Model	18
3.3	Resources Model	19
3.4	Social Structure Model	20
3.4.1	Shared state and social facts	21
3.4.2	Acting in a social context	23
3.4.3	Transactions and exchanges model	25
3.4.4	Roles in Social Structures	26
3.4.5	Governance and Social Structures	27
3.5	Tail Piece	Error! Bookmark not defined.
4	Realizing a Service Oriented Architecture View	29
4.1	Service Description Model	29
4.1.1	Components of Service Description	30
4.1.2	Implications of the Description Model	36
4.1.3	Service Description and Service Level Agreements	37
4.1.4	Consumer Description	38
4.1.5	Information Model	38
4.2	Service Visibility Model	41
4.2.1	Visibility to Business	41
4.2.2	Attaining Visibility	43
4.2.3	Achieving Awareness	43
4.2.4	Mechanisms for Attaining Visibility	47
4.3	Interacting with Services Model	48
4.3.1	Action Model	48
4.3.2	Process Model	51
4.3.3	Orchestration and Service Composition	52

4.3.4 Dependencies between services	52
4.4 Policies and Contracts	53
4.4.1 Policy concepts	Error! Bookmark not defined.
4.4.2 Policy/Contract Language	Error! Bookmark not defined.
4.4.3 IT Mechanisms Supporting Policies and Contracts	56
4.4.4 Policy/Contract Relationships	Error! Bookmark not defined.
4.4.5 Contracts and Service Level Agreements	Error! Bookmark not defined.
4.4.6 Policies and Contracts and Service Descriptions	Error! Bookmark not defined.
5 Owning Service Oriented Architectures View	59
5.1 Governance of Service Oriented Architectures	59
5.1.1 Why is explicit Governance important to SOA?	59
5.1.2 What are the concerns of the stakeholders?	60
5.1.3 Inputs to the decision process	61
5.1.4 Implementing SOA Governance	61
5.1.5 Governance Bone Yard	62
5.1.6 Ken's notes on Governance	62
5.2 Security Model	64
5.2.1 Security Concepts	64
5.2.2 Threat Model	65
5.2.3 Mitigation Model	66
5.3 Services as Managed Entities Model	68
5.3.1 Definitions	70
5.3.2 Management Capabilities	71
5.3.3 Management Contracts and Policies	72
5.3.4 Manageability & Instrumentation	73
5.3.5 Management Infrastructure	73
5.3.6 Service Life-cycle	74
5.3.7 Service Provisioning	74
6 References	75
6.1 Normative References	75
6.2 Non-Normative References	75
A. Acknowledgements	76
B. Critical Factors Analysis	77
B.1 Goals	77
B.1.1 Critical Success Factors	77
B.1.2 Requirements	77
B.1.3 CFA Diagrams	77
C. Revision History	78

1 Introduction

2 Service Oriented Architecture is an important paradigm that has gained significant attention within the
3 information technology (IT) and business communities. The OASIS Reference Model for SOA provides a
4 common language for understanding the important features of SOA but does not address the issues
5 involved in constructing a SOA-based system. This document focuses on this aspect of SOA; while
6 maintaining a similarly high-level approach as the Reference Model itself.

7 1.1 What is a Reference Architecture

8 A Reference Architecture is a description of the concepts and relationships in a domain that enables
9 stakeholders to identify how a set of requirements may be realized. It differs from a Reference Model in
10 that a Reference Model describes the important concepts in the domain focusing on what distinguishes
11 the elements of the domain; a Reference Architecture elaborates further on the model to show a more
12 complete picture that includes showing what is involved in realizing the modeled entities.

13 We identify three broad categories of requirements: how Service Oriented Architecture fits into the life of
14 users and stakeholders, how SOA-based systems may be realized effectively, and what is involved in
15 owning a SOA-based system.

16 It is possible to define Reference Architectures at many levels of detail or abstraction. In this Reference
17 Architecture we have followed a high-level technology neutral approach; while at the same time being
18 fully aware of the dominant technologies likely to be employed. In fact, the degree of abstraction in
19 modeling concepts in the Reference Architecture is very similar to that employed in the Reference Model
20 itself. We believe that this will serve two purposes: ensuring that the true value of the SOA approach can
21 be realized on any appropriate technology, and it permits our audience to focus on the important issues
22 without becoming over-burdened with unnecessary detail.

23 1.2 Service Oriented Architecture – An Ecosystems perspective

24 Many systems cannot be understood by a simple decomposition of their parts into subsystems. There are
25 too many interactions between the parts. There has been much recent research about "complex systems"
26 because it has many applications to large systems such as the economy, or the human brain, where one
27 cannot understand the larger picture just by putting the pieces together.

28 A service-oriented architecture shares many of the characteristics of such systems. From the perspective
29 of a complex system, a SOA is a network of independent services, machines, the people who operate,
30 affect, use, and govern those services as well as the suppliers of equipment and personnel to these
31 people and services. This includes any entity, animate or inanimate, that may affect, or be affected by
32 the system. With a system that large, it is clear that nobody is really "in control" or "in charge"; although
33 there are definite stakeholders involved, each of whom has some control and influence of the whole.

34 We have multiple perspectives view of this SOA architecture: corresponding to the ways in which SOA
35 must be understood by key stakeholders: by its users, by its constructors and by its owners.

36 1.3 Relationship to the Reference Model

37 This Reference Architecture takes the Reference Model as its starting point; however, it is somewhat
38 more concrete than the Reference Model.

39 As a result, some of the concepts that are identified within the Reference Model are further expanded on
40 in the Reference architecture; other concepts are 'unpacked' into different aspects of the eco-
41 system/machine. Furthermore, additional concepts are introduced which were not 'of the essence' for
42 describing Service but are required in order to have a Service Oriented Architecture.

43 One key concept that has been unpacked is that of Execution Context. Within the RM, the execution
44 context stood for all the aspects of an information system that are needed to facilitate interaction. A large

45 part of the goals of the Reference Architecture is to show how interaction is realized; as a result, the
46 concept of Execution Concept is not as pertinent within the Reference Architecture.

47 **1.4 Relationship to other Reference Architectures**

48 It is fully recognized that other SOA reference architectures have emerged in the industry, both from the
49 analyst community and the vendor/solution provider community. Some of these reference architectures
50 are at a sufficient level of abstraction away from specific implementation technologies while others are
51 based on a solution or technology “stack.” Still others use emerging middleware technologies such as the
52 Enterprise Service Bus (ESB) as the architectural foundation.

53 The Reference Architecture for SOA is an abstract realization of SOA—showing how a SOA can be built
54 while omitting any reference to specific concrete technologies. As with the Reference Model for SOA, the
55 Reference Architecture is primarily focused on large-scale distributed IT systems where the participants
56 may be legally separate entities. While it is quite possible for many aspects of the Reference Architecture
57 to be realized on quite different platforms, we do not dwell on such opportunities.

58 **1.5 Viewpoints, Views and Models**

59 **1.5.1 ANSI/IEEE Std 1471-2000**

60 This Reference Architecture for SOA loosely follows the ANSI/IEEE Std 1471-2000 Recommended
61 Practice for Architectural Description of Software-Intensive Systems [#]. An architectural description
62 conforming to the ANSI/IEEE Std 1471-2000 recommended practice is described by a clause that
63 includes the following six (6) elements:

- 64 1. Architectural description identification, version, and overview information
- 65 2. Identification of the system stakeholders and their concerns judged to be relevant to the
66 architecture
- 67 3. Specifications of each viewpoint that has been selected to organize the representation of the
68 architecture and the rationale for those selections
- 69 4. One or more architectural views
- 70 5. A record of all known inconsistencies among the architectural description's required constituents
- 71 6. A rationale for selection of the architecture

72 The ANSI/IEEE Std 1471-2000 defines an architectural description (AD) as “a collection of products to
73 document the architecture,” where architecture is defined as:

74 **Architecture**

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

77 A system stakeholder is “an individual, team, or organization (or classes thereof) with interests in, or
78 concerns relative to, a system,” where system is defined as:

79 **System**

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

81 A stakeholder concern (or care-about) should not be confused with a formal requirement. A concern is an
82 area or topic of interest. Within that concern, system stakeholders may have many different requirements.
83 In other words, something that is of interest or importance is not the same as something that is obligatory
84 or of necessity.

85 When describing architectures, it is important to identify stakeholder concerns and associate them with
86 viewpoints to insure that those concerns will be addressed in some manner by the models that comprise
87 the views on the architecture. The ANSI/IEEE Std 1471-2000 defines views and viewpoints as follows:

88 **View**

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

90 **Viewpoint**

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

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

96 It is important to note that viewpoints are independent of a particular system. In this way, the architect can
97 select a set of candidate viewpoints first, or create a set of candidate viewpoints, and then use those
98 viewpoints to construct specific views that will be used to organize the AD. A view, on the other hand, is
99 specific to a particular system. Therefore, the practice of creating an AD involves first selecting the
100 viewpoints and then using those viewpoints to construct specific views for a particular system or
101 subsystem. Note that the ANSI/IEEE Std 1471-2000 requires that all views correspond to exactly one
102 viewpoint. This helps maintain consistency among architectural views; a normative requirement of the
103 standard (see 5 above).

104 A view is comprised of one or more architecture models, where model is defined as:

105 **Model**

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

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

109 **1.5.2 UML Modeling Notation**

110 To help visualize structural and behavioral architectural concepts, it is useful to depict them using an
111 open standard visual modeling language. Although a myriad of architecture description languages exist
112 in practice, we have adopted the second generation Unified Modeling Language™ (UML®) known as UML
113 2 managed under the auspices of the Object Management Group™ (OMG™) as the primary viewpoint
114 modeling language. It should be noted that while UML 2 is used in this reference architecture,
115 formalization and recommendation of a UML Profile for SOA is beyond the scope of this specification.
116 Every attempt is made to utilize normative UML unless otherwise noted.

117 **1.6 Viewpoints of this Reference Architecture**

118 The Reference Architecture is partitioned into three views that conform to three primary viewpoints,
119 reflecting the main division of concerns noted above: the Business via Services viewpoint focuses on how
120 a SOA integrates with how people conduct their business; the Realizing a Service Oriented Architecture
121 viewpoint focuses on the salient aspects of building a SOA, and the Owning Service Oriented
122 Architectures viewpoint focuses on those aspects that relate to owning, managing and controlling a SOA.

123 The viewpoint specifications for each of the primary viewpoints of this reference architecture are
124 summarized in Table 1. Additional detail on each of the three viewpoints is further elaborated in the
125 following subsections. For this reference architecture, a one-to-one correspondence between viewpoints
126 and views is assumed.

Viewpoint Element	Viewpoint		
	<i>Business via Services</i>	<i>Realizing a SOA</i>	<i>Owning SOAs</i>
Main concepts	Captures what SOA means for people using it to conduct business	Deals with the requirements for constructing an SOA	Addresses issues involved in owning an SOA vs. using or building one
Stakeholders	People (using SOA)	Service Consumers, Service Providers	Service Providers
Concerns	Conduct business safely	Effective	Processes for

	and effectively	implementation using standard technologies	engaging in an SOA are effective, equitable, and safe
Modeling Techniques	UML Class diagrams	UML Class and Sequence diagrams	UML Class diagrams

127 Table 1 Viewpoint specifications for the OASIS Reference

128 **1.6.1 Business via Services Viewpoint**

129 The Business via Services viewpoint is intended to capture what using a SOA-based system means for
130 people using it to conduct their business. From this viewpoint, we are concerned with how SOA
131 integrates with and supports the service model from the perspective of the people who perform their tasks
132 and achieve their goals as mediated by Service Oriented Architectures. The Business via Services
133 viewpoint also sets the context and background for the other viewpoints in the Reference Architecture.

134 The stakeholders who have key roles in or concerns addressed by this viewpoint are *people* and the
135 primary concern for people is to ensure that they can use a SOA to conduct their business in a safe and
136 effective way. Given the public nature of the Internet, and the intended use of SOA to allow people to
137 access and provide services that cross ownership boundaries, it is necessary to be able to be somewhat
138 explicit about those boundaries and what it means to cross an ownership boundary.

139 The modeling techniques for expressing the visual models that comprise the associated view that
140 conforms to this viewpoint are UML models; principally, UML class diagrams.

141 **1.6.2 Realizing a Service Oriented Architecture Viewpoint**

142 The Realizing a Service Oriented Architecture Viewpoint focuses on the infrastructural elements that are
143 needed in order to support the discovery and interaction with services. From this viewpoint we are
144 concerned with the application of *normal* technologies available to system architects to realize the vision
145 of an SOA that may cross ownership boundaries. In particular, we are aware of the importance and
146 relevance of other standard specifications that may be used to facilitate the building of an SOA.

147 The modeling techniques for expressing the visual models that comprise the associated view that
148 conforms to this viewpoint are UML models; principally, UML class, sequence, and communication
149 diagrams.

150 **1.6.3 Owning a Service Oriented Architectures Viewpoint**

151 The Owning Service Oriented Architectures viewpoint addresses the issues involved in owning an SOA
152 as opposed to using one or building one. Many of these issues are not easily addressed by automation;
153 instead, they often involve people-oriented processes such as governance bodies.

154 The principal stakeholders who have key roles in or concerns addressed by this viewpoint are *service*
155 *providers*, *service consumers* and other non-participatory organizations charged with ensuring societal
156 goals – such as fair trade, public safety and so on. The primary concerns are that the processes for
157 engaging in an SOA are effective, equitable, and safe.

158 The modeling techniques for expressing the visual models that comprise the associated view that
159 conforms to this viewpoint are UML models; principally, UML class diagrams.

160 **1.7 Terminology**

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

164

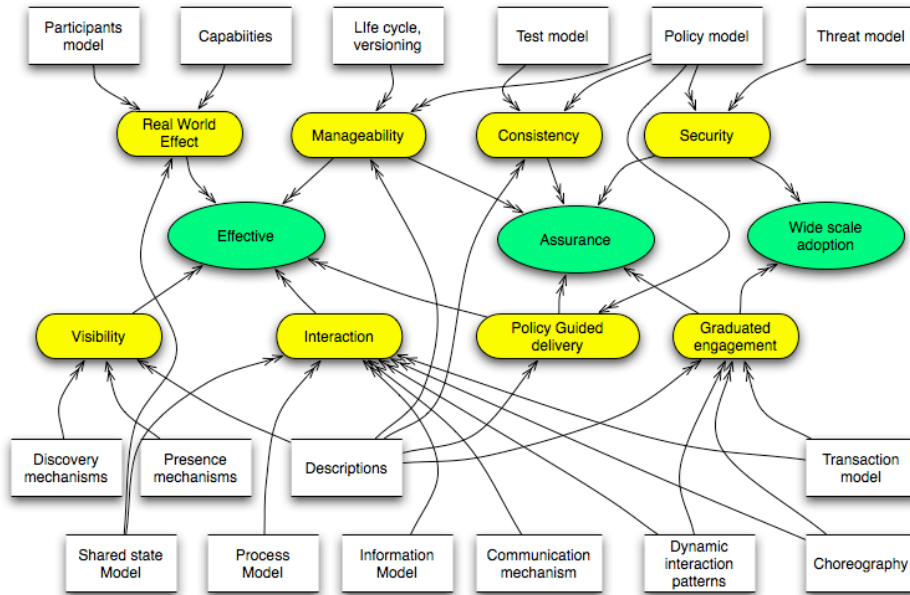
2 Requirements

165 A Reference Architecture is like an abstract machine. It is built to realize some function and it, in turn,
166 relies on a set of underlying capabilities that must be present for it to perform. In the case of the SOA RA,
167 its purpose is to enable a system to be a Service Oriented Architecture. The underlying capabilities are
168 the particular technologies that are used to realize the SOA; in particular technology choices such as Web
169 services technologies, implementation technologies are not part of an abstract RA.

170 The purpose of the RA is reflected in the set of requirements that the RA must satisfy. We can structure
171 these requirements into a set of goals, a set of critical success factors (CSFs) associated with these goals
172 and a set of requirements that are connected to the CSFs that ensure their satisfaction.

173 Note that not all of the requirements are mapped to solutions within the scope of this RA. Indeed, the RA
174 itself can be seen as generating a series of more explicit requirements for the realizing technology.

175 The overall requirements are illustrated in Figure 1.



176

177 *Figure 1 Critical Factors Analysis of the Reference Architecture*

178 The critical factors analysis (CFA) requirement technique and the diagram notation is summarized in
179 Appendix B.

180 2.1 Goals of the Reference Architecture

181 There are three principal goals of the Reference Architecture: that it shows how to build SOAs that
182 effectively meet the requirements of the stakeholders involved, that it does so with some assurance that
183 no harm is done, and that the architecture itself can be widespread.

184 2.1.1 Effectiveness

185 According to the RM, a SOA is a "paradigm for organizing and utilizing distributed capabilities that may be
186 under the control of different ownership domains". Enabling the offering and utilizing of capabilities across
187 ownership domains is the definition of effectiveness of Service Oriented Architecture.

188 For the architecture to be effective, it means that systems based on it can effectively bring the needs and
189 capabilities of participants together.

190 There are a number of critical success factors that must be addressed for this goal to be satisfied: it must
191 be possible for service providers and users to see each other, they must be able to interact in modern
192 systems, the effects of provided and using services must similarly be communicated and the SOA-based
193 system must itself be manageable with reasonable effort.

194 2.1.1.1 Visibility

195 Without service consumers and providers being able to "see" each other, they cannot interact with each
196 other. The RM goes on to identify three key concepts associated with visibility: awareness, willingness
197 and reachability. It also begins to define the basis for description which is a key to enabling visibility.
198 This document will elaborate on the makeup and structure of description, the means by which description
199 may be catalogued for discovery and retrieval, and the mechanisms for detecting the presence of a
200 service.

201 As part of the reachability aspect of visibility, it must be possible to determine if a service is active or not.
202 In general this is difficult to be certain of (a service might become unavailable the instant after you have
203 been informed that its available).

204 2.1.1.2 Interaction

205 Interaction between service providers and consumers is how a service is effected. Supporting interaction
206 is a major critical success factor for SOA. As is made clear in the rest of the architecture, interaction is
207 achieved in many distributed systems via message-based communication.

208 **Communication**

209 One key requirement for interaction between service participants is communication: defined as
210 the meaningful exchange of information that can be interpreted as part of a service interaction.

211 **Information model**

212 Successful interaction requires that the parties have agreement on the nature of the information
213 that is exchanged. This is captured in the information model associated with the service.

214 **Behavior model**

215 The dynamics of interaction must also be effectively agreed to by the parties in the interaction.

216 2.1.1.3 Real World Effect

217 The purpose of interacting with a service partner (both the service consumer and the provider) is to
218 achieve desired effects - the real world effect. The RWE captures the result of bringing a capability to
219 bear as a consequence of the interaction. How the RWE is described and realized in the RA is a critical
220 aspect of SOA.

221 2.1.2 Assurance

222 The RA should enable service providers and consumers to achieve their goals with the maximum
223 possibility of safety in the interaction. Note that we distinguish any assurance associated with the delivery
224 of a service from that resulting from the application of the service. The latter is beyond the scope of this
225 RA. Assurance of service is of the essence when providers and consumers are in different ownership
226 domains; and hence supporting the safety of that interaction is an important goal of the RA.

227 Some of the critical factors that concern safety of a system are how secure the system is whether
228 participants' expectations in using the system are consistent, whether stakeholders' policies can be
229 respected and whether the risk of interacting with the system is commensurate with expectations.

230 **2.1.2.1 Security**

231 In any system where there is exposure to multiple ownership domains, security is a paramount factor in
232 its success. Security can be addressed from several perspectives: the threats that must be addressed,
233 the mechanisms for mitigating those threats and the management of those mechanisms. In addition,
234 legitimate access to the system needs to be easy to avoid incentives to bypass security mechanisms.

235 The key threats are against the privacy of interaction, against the proper identification of systems and
236 people interacting with services, against the improper or compromised use of services, and against the
237 possible later repudiation of previous interactions.

238 Security should facilitate access to content and business logic by those who are entitled, that is,
239 continuous interruptions should not impede the functioning of the system.

240 **2.1.2.2 Consistency**

241 In order to effectively interact across ownership boundaries it is critical that the actual interaction matches
242 expectations for that interaction. The key to this predictability is adequate and explicit descriptions of all
243 facets of services; further supported by explicit policy statements.

244 **2.1.2.3 Explicit Policies**

245 Any machine is necessarily broad in its applicability and is often under-constrained. Policy statements
246 define the choices that a service provider and/or service consumer (or other stakeholder) makes. Access
247 to these policy choices is an important aspect of ensuring predictability and consistency.

248 A critical factor for SOA is that consistent policy can be specified, applied and enforced. Policies can be
249 applied to many aspects of an SOA-based system; here we focus on the role of policy in delivering
250 service functionality itself.

251 **2.1.2.4 Participants' roles**

252 The rights, obligations, roles and collaborative context of participants in service interactions must be
253 accurately represented in the architecture. This is important from the perspective of the RWE — which is
254 expressed in terms of the facts and commitments shared by service participants — and from the
255 perspective of the security threats that any SOA is subject to.

256 **2.1.2.5 Graduated engagement**

257 There is a fundamental principle in interacting across ownership boundaries that there be no "unpleasant
258 surprises". While we cannot eliminate such unexpected consequences we can insist on a model where
259 the expectations of the parties in an interaction are commensurate with their commitments; i.e., as we
260 gradually get deeper into an engagement with a service partner the commitments and expectations
261 become similarly deeper.

262 For example, if a service receives a sequence of bytes that it cannot understand, it should not assume
263 that its service has been validly invoked. Another example is the inadmissibility (in many jurisdictions) of
264 the so-called drive-by license agreement: the user must perform an explicit action signaling agreement for
265 it to be binding.

266 **2.1.2.6 Manageability**

267 Given that a large-scale SOA is likely to be populated with many thousands of services, managing them
268 becomes a critical factor for the assured delivery of services.

269 Manageability requires clear descriptions of who is responsible for the service and what responsibility
270 entails. This includes how a given service contributes and consumes resources as part of the SOA
271 ecosystem.

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

275 In the context of multiple ownership domains manageability may be less an issue of managing IT
276 infrastructure components (although that is very critical) used to realize services than managing the use
277 of services, managing the relationships between participants in their use of services, and permitting
278 management of services across ownership boundaries.

279 **2.1.3 Wide scale adoption**

280 It is an explicit goal of this work that the model it promotes will find widespread acceptance in Industry.
281 While we cannot guarantee wide-scale adoption: In addition to assurance issues mentioned above, we
282 identify a number of factors that will enhance the adoptability of SOA-based systems: applicability,
283 scalability, loose coupling, low cost of entry, reusability, technology independence and simplicity.

284 **2.1.3.1 Scalability**

285 Any architecture whose design principles are not effective across the wide range of possible scales from
286 a small intra-Enterprise system to a full-scale Internet-wide deployment is unlikely to find general
287 acceptance. Any given instantiation of the architecture need not scale to the full Internet; the RA itself
288 must be capable of such scale.

289 On the other hand, human ownership boundaries are not limited to single legal jurisdictions or geographic
290 regions.

- 291 • It should be possible for service providers and consumers to interact even across such
292 boundaries.
- 293 • It should be possible for services to be offered in ways that are sensitive to the local environment;
294 the same service may have different presentations in different locales.
- 295 • Reduced dependence on one particular natural language

296 This may have some surprising consequences. For example, although a natural language description of
297 some item may be easier to understand for a native speaker of that language, the language used can
298 become an unnecessary barrier for a non-native speaker. Formal descriptions, i.e., descriptions in a
299 formal notation, may be more difficult to understand, however, they are more neutral than descriptions in
300 any given natural language, and therefore are less likely to be misunderstood by people as well as by
301 automated processors.

302 **2.1.3.2 Loosely coupled**

303 A loosely coupled system is one in which the constraints on the interaction between components is
304 minimal: sufficient to permit interoperation without additional constraints that may be an artifact of
305 implementation technology. Each non-essential constraint may have negative impacts on the scalability of
306 systems based on the architecture.

307 **2.1.3.3 Understandability**

308 An architecture that is highly scalable, but which is complex and difficult to understand will not find
309 general acceptance.

310 **2.1.3.4 Reusability**

311 The extent to which services are reusable, descriptions, and other artifacts are reusable will be a critical
312 factor in promoting adoption. Without reusability there is no investment and no growth.

- 313 • Artifacts should be realized in appropriately compatible technology. There should not be unnecessary
314 technological dependencies.
- 315 • Pre-existing capabilities should be made accessible as services with the minimum of transformation.
- 316 • Services should be combinable in appropriate forms without inconsistent effect and undue effort.

317 **2.1.3.5 Low cost of entry**

318 A low cost of entry is an important critical success factor for wide-scale adoption.

319 The technology requirements for a given system should be commensurate with the complexity of that
320 system. Simple service systems should be simple to design and build; complex systems with complex
321 requirements should not be made more complex by the constraints of this architecture.

322 **2.1.3.6 Technology neutral**

323 The extent to which the RA is technology neutral will greatly effect its applicability in different
324 circumstances. The counterpoint to being technology neutral is the ability to operate across several
325 different kinds of platform. Enabling service providers and consumers to occupy different platforms is a
326 key CSF that will help to drive adoption of the RA.

327 **2.1.3.7 Simplicity**

328 A key factor in wide-scale adoption will be simplicity of design and realization. A complex architecture, or
329 one that is difficult to understand, would be an impediment to wide-scale adoption.

330 The hallmark of good design is simplicity. This is better expressed as sufficient complexity to satisfy
331 requirements while avoiding unnecessary complexity. History shows us that it is often better to err on the
332 side of reduced functionality if that promotes simplicity.

- 333 • Minimal assumptions
334 • Compositionality

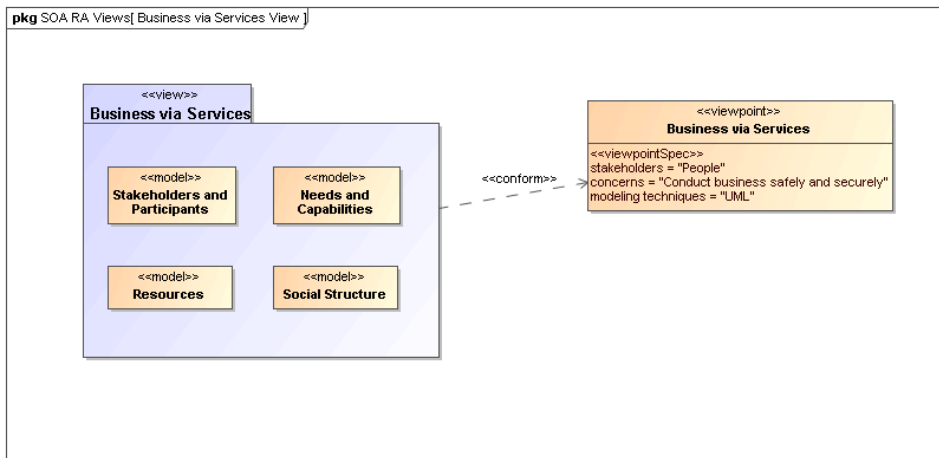
3 Business via Services View

335
336
337
338
339
340
341
342
343
344
345
346

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

347 The *Business via Services* view focuses on how Service Oriented Architecture fits into the life of users
348 and stakeholders. The function of SOA is to facilitate action in a community of people; where action is
349 characterized in terms of providing services and consuming services to realize mutually desirable real
350 world effects. Thus, our tasks in this view are to model the people involved—the participants and other
351 stakeholders—their goals and activities and the relevant relationships between people as they affect the
352 utility and safety of actions that are performed.

353 The models in this view include the Stakeholders and Participants Model, the Needs and Capabilities
354 Model, the Resources Model, and the Social Structure Model.

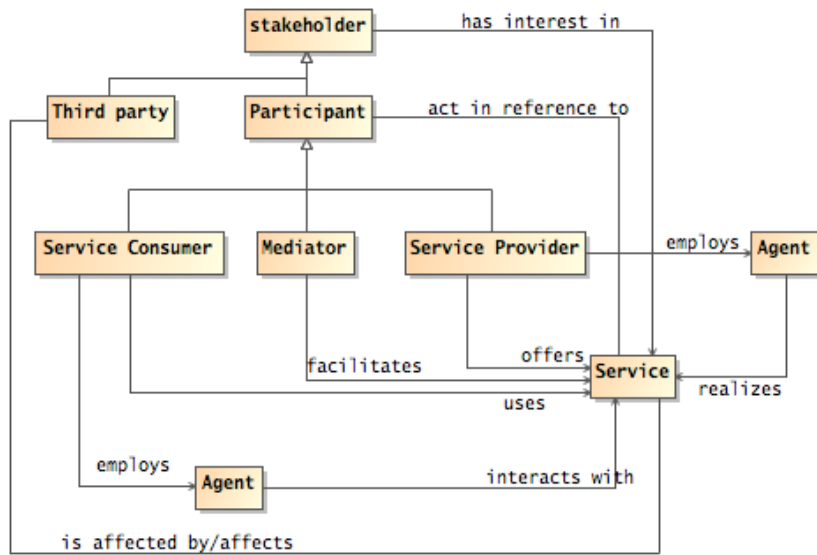


355
356

Figure 2 Model elements described in the Business via Services view

3.1 Stakeholders and Participants Model

357 A SOA is deployed in the context of human and non-human entities capable of action. In this section we
358 focus on the relationship between these ultimate actors and the services that they use and deploy.
359



360
361 *Figure 3 Service Participants*

362 **Stakeholder**

363 A stakeholder is a human, corporation or non-human agent that has an interest in the states of
364 services and/or the outcomes of service interactions.

365 Stakeholders do not necessarily participate in service interactions. For example, a government may have
366 an interest in the outcomes of commercial services deployed in a SOA without actively participating in the
367 interactions (the government may collect tax from one or more participants without being part of the
368 interaction itself).

369 **Participant**

370 A participant is a human, organization or non-human agent that has the capability and
371 requirement to act in the context of a Service Oriented Architecture.

372 A participant is an example of a stakeholder whose interests lie in the successful use of and
373 fulfillment of services. Note that we admit non-human agents as an extreme case: the normal
374 situation is where participants are either human or corporations. However, human participants
375 always require *representation* in an electronic system – they require agents.

376 It is convenient to classify service participants into service providers and service consumers. The reason
377 for this is twofold: an extremely common mode of interaction is where a provider participant offers some
378 functionality as a service and a consumer participant uses that service to achieve one of his or her goals.
379 Secondly, it helps to illustrate the dominant situation where the participants are not truly symmetric: they
380 each have different objectives and often have different capabilities. However, it should be noted that there
381 are patterns of interactions where it is not clear that the distinction between service provider and
382 consumer are valid. [Give example here.](#)

383 **Service Provider**

384 A service provider is a participant that offers a service that permits some capability to be used by
385 other participants.

386 In normal parlance, the service provider commonly refers to either the ultimate owner of the capability that
387 is offered or at least an agent acting as proxy for the owner. For example, an individual may own a

388 business capability but will enter into an agreement with another individual (the proxy) to provide SOA
389 access to that business -- so that the owner can focus on running the business itself.

390 Note that several kinds of stakeholders may be involved in provisioning a service. These include the
391 provider of the capability, an enabler that exposes it as a service, a mediator that translates and/or
392 manages the relationship between service consumers and the service, a host that offers support for the
393 service, a government that permits the service and/or collects taxes based on service interactions and so
394 on.

395 **Service Consumer**

396 A service consumer is a participant that interacts with a service in order to satisfy some desired
397 need. As with service providers, several stakeholders may be involved in a service interaction
398 supporting the consumer.

399 It is a common understanding that service consumers typically initiate service interactions. Again, this is
400 not necessarily true in all situations (for example, in publish-and-subscribe scenarios, a service consumer
401 may initiate an initial subscription, but thereafter, the interactions are initiated by publishers).

402 **Service mediator**

403 A service mediator is a participant that facilitates the offering or use of services in some way
404 without necessarily directly interacting with the service. There are many kinds of mediator, for
405 example a registry is a kind of mediator that permits providers and consumers to find each other.
406 Another example might be a filter service that enhances another service by encrypting and
407 decrypting messages. Yet another example of a mediator is a proxy broker that actively stands
408 for one or other party in an interaction.

409 **Agent**

410 An agent is any person or non-human entity that is capable of acting on behalf of a person or
411 organization.

412 **Artificial agent**

413 An artificial agent is a constructed entity that is used by people to enable them to offer, consumer
414 and otherwise participate in services. Common examples of artificial agents are software
415 applications that make use of services, hardware devices that embody an agent with a particular
416 mission, and enterprise systems that offer services.

417 In the context of SOA, since interaction between participants is mediated via the networks such as the
418 Internet, people require proxies to participate. Artificial agents that can have direct access to electronic
419 communications permit this.

420 We do not attempt to characterize artificial agents in terms of their internal architecture, computational
421 requirements or platforms here. Within the Service as Business view, an agent stands for whatever
422 information technology resources are required to facilitate the human use and other forms of participation
423 involving services.

424 **Third-party stakeholder**

425 A third party is any stakeholder who may be affected by the use or provisioning of services or
426 who has an interest in the outcome of service interactions.

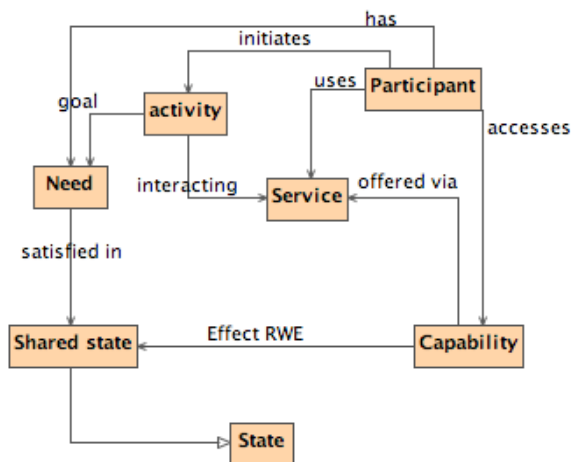
427 There are two main classes of such non-participatory stakeholders: innocent bystanders who are
428 materially affected by someone's use or provisioning of a service, and regulatory agencies who wish to
429 control in some way (such as by taxation) services.

430 At its most basic, a service is provided by a provider and used by a consumer in order to achieve a
431 change in the real world that meets a desired goal.

432 However, interactions between service participants and actions undertaken as a result can only be
433 understood in the context of other relationships between participants. I.e., in order to understand the
434 validity and consequences of a service interaction, it is necessary to understand the relative roles of the
435 participants (and stakeholders generally).

436 **3.2 Needs and Capabilities Model**

437 The motivation for participants interacting is the satisfaction of needs. Participants use and deploy
 438 services in order that they can get their needs met. From a consumer perspective, the need is often
 439 related to the role they represent in the social structure; for the provider, the need is to gain satisfaction,
 440 monetary or otherwise, for use of the service.



441
 442 *Figure 4 Needs and Capabilities*

443 **Capability**

444 A resource that may be used by a service provider to achieve a real world effect on behalf of a
 445 service consumer.

446 As noted in the RM, the Real World Effect is couched in terms of changes to the state that is shared by
 447 the participants in the service; in particular the public aspects of that state. In this Reference Architecture
 448 we further refine this notion in terms of social structures.

449 Thus, we might refer to a capability as being able to effect facts that have meaning within a social
 450 structure; i.e., to be able to modify the state of the social structure. This does not rule out physical effects
 451 of using a service -- for all effects are ultimately rooted in the physical world -- but that the highest
 452 interpretation of a capability is in terms of the social structures in which it is embedded.

453 For example, a book selling service may have the capability of delivering a book to a customer. From a
 454 purely physical perspective, there is not much to distinguish one book from another, or one physical
 455 location from another. What makes the book selling service particular, however, is the ability to deliver the
 456 work of a particular author to a particular customer. Furthermore, even if the book has been successfully
 457 delivered, if the customer fails to pay for the book (if, for example, the credit card used turned out to be
 458 fraudulent), then the customer does not own the book, and is obliged to return it. The concept of
 459 transferring the ownership of the book, which is the real capability offered by the book selling service, has
 460 no direct counterpart in the physical world.

461 Capabilities themselves have owners, making those owners stakeholders in the SOA.

462 By making a capability available for use, via the Service, the owners aim to satisfy their needs as well as
 463 the needs of other participants who use the service. The extent to which a capability is exposed via a
 464 service (or via multiple services) is just one of the choices that the owners of capabilities have.

465 **Need**

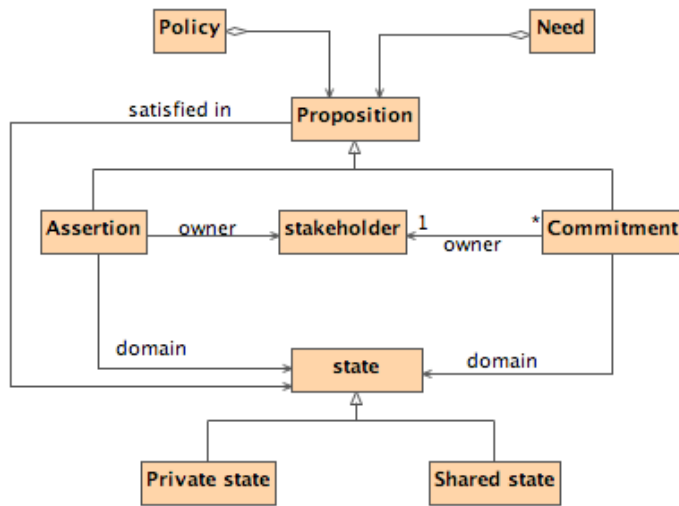
466 A need is a measurable requirement that a service participant is actively seeking to satisfy. The
467 aspects of a need are that it can be measured and that it belongs to a participant (more generally
468 any stakeholder).

469 A need is characterized by a proposition: an expression whose truth can be measured. However, needs
470 themselves are often visible only to the owner of the need. Furthermore, the extent to which a need is
471 captured in a formalizable way is likely to be very different in each situation.

472 Measuring expectations is not always easy, because while the actions of the service are private, only the
473 results are public. In addition, there is always the potential for unexpected consequences of actions.

474 As an example, let us consider the ability to withdraw money from a bank account. Suppose as part of its
475 implementation, the bank reports withdrawals over a certain amount to a credit agency. This potentially
476 could alter an individual's credit report. The service description for the bank withdrawal service has to
477 include this information so that a service consumer can accurately understand the actual real world effect.
478 Of course this could be dynamic if the bank uses different credit reporting services that have different
479 policies about credit ratings.

480 Figure 5 captures some of the key concepts and relationships involving needs. Since needs and policies
481 share some of the same characteristics we include policy in the model; however, policies themselves are
482 discussed in more detail in Section 4.4.




483
484 *Figure 5 Needs expressed using propositions*

485 A need is owned by a stakeholder. Many stakeholders are active participants in services, but not all.

486 **3.3 Resources Model**

487 As we noted above, a key relationship between many elements is that of ownership. The resource model
488 focuses on what it is that can be owned: resources.



489

490 *Figure 6 Resources*

491 Our model of resources is very simple, but is the foundation for modeling many of the things that a SOA
 492 deals in: information, physical resources and so on:

493 **Resource**

494 A resource is any entity of some perceived value, where the value may be in the function it
 495 performs or something intrinsic in its nature. For example, a diamond has value in its
 496 appearance, its physical properties, and its monetary worth; a data set has value in the
 497 information it carries. A resource that has identity associated with it – it can be identified – and
 498 has an owner.

499 Resources are not confined to physical entities: anything, real, virtual or abstract that may be owned is a
 500 resource. The key attributes of a resource are that it can be identified, and that it has an owner.

501 This definition of resource is a simplification and elaboration of the concept that underlies the Web
 502 architecture. Being more abstract, we do not require that the identity of a resource be in any particular
 503 form (although in practice, many resource identifiers are URIs), nor do we require resources to have
 504 representations. However, we do require resources to have owners.

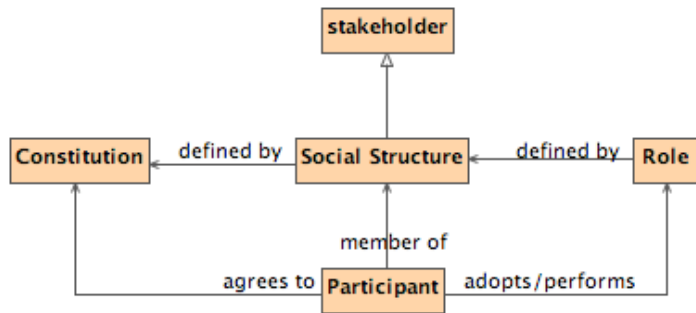
505 One important class of resources in this architecture are the capabilities that underlie services. In this
 506 case, we can say that a capability is a resource that can cause an effect in the world. For example, a light
 507 bulb is a resource that when activated gives off light; a book is a resource that when read allows one to
 508 gain knowledge from its content. Other examples of resources are services themselves, descriptions of
 509 entities (a kind of meta-resource), IT infrastructure elements used to deliver services, contracts and
 510 policies, and so on.

511 **3.4 Social Structure Model**

512 The actions undertaken by participants, whether mediated by services or in some other way, are normally
 513 performed in the context of a social context which defines the meaning of the actions themselves. We can
 514 formalize that context as a **social structure**: the embodiment of a particular social context.

515

516 The social structure model is important to defining and understanding the implications of crossing
 517 ownership boundaries; it is the foundation for an understanding of security in SOA and also provides the
 518 context for determining how SOAs can be effectively managed and governed.



519

520 *Figure 7 Social Structure*

521 **Social Structure**

522 A social structure (sometimes identified as social institutions) embodies some of the cultural
 523 aspects that characterize the relationships and actions among a group of participants.

524 Social structures are often, but not necessarily, aligned with organizations. For example, a meeting of
 525 like-minded fellows by a watering hole may be an informally defined social structure (with its own rules)
 526 that is not connected with any organization per se.

527 The richness of social structures reflects the richness of human culture itself. However, in the context of
 528 the Reference Architecture, we are concerned primarily with social structures that are embodied in legal
 529 and quasi-legal frameworks; i.e., they have some rules that are commonly understood.


530 So, for example, a corporation is a common kind of social structure, as is a fishing club. At the other
 531 extreme, the legal frameworks of entire countries and regions also count as social structures.

532 It is not necessarily the case that the social structures involved in a service interaction are explicitly
 533 identified by the participants. For example, when a customer buys a book over the Internet, the social
 534 structure that defines the validity of the transaction is often the legal framework of the region that the book
 535 vendor belongs to. This legal jurisdiction qualification is typically buried in the fine print of the service
 536 description.

537 **3.4.1 Shared state and social facts**

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

542 Social facts are inherently public or shared: they only have meaning in the context of the social structure
 543 and to the participants in the social structure.



544

545 *Figure 8 Shared state and social facts*

546 **Shared state**

547 The set of facts and commitments that manifest themselves to service participants as a result of
548 interacting with a service.

549 Note that a participant has only a partial view of the shared state in a system. Furthermore, the participant
550 will have internal state that is not accessible to other participants directly. However, elements of the
551 shared state are in principle accessible to participants even if a given participant does not have access to
552 all elements at any given time.

553 **Social fact**

554 A social fact is an element of the state of a social structure that is sanctioned by that social
555 structure. For example, the existence of a valid purchase order with a particular customer has a
556 meaning that is defined primarily by the company itself.

557 **Commitment**

558 A commitment is a social fact about the future: in the future some fact will be true and a
559 participant has the current responsibility of ensuring that that fact will indeed be true. A
560 commitment to deliver some good is a classic example of a fact about the future.

561 Other important classes of social facts include the policies adopted by an organization, any agreements
562 that it is holding for participants, and the assignment of participants to roles within the organization. The
563 social facts that are understood in the context of a social structure define the shared state that is
564 referenced in Figure 5.

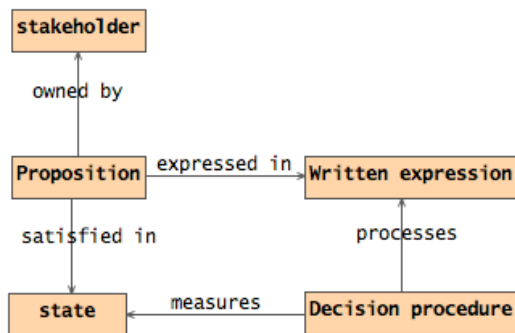
565 Social facts and commitments are inherently abstract; however, facts have the property of being verifiable
566 (technically, a social fact can be verified to determine if it is satisfied in the social context). If as a result of
567 interacting with a service, a buyer incurs the obligation of paying for some good or service, this obligation
568 (and the discharge of it) is measurable (perhaps by further interactions with the same or other services).

569 **3.4.1.1 Measuring social facts**

570 We make considerable use of the term **proposition** and related concepts. Propositions are the basis of
571 policies, agreements and contracts, ownership, social facts, shared state and many other elements of the
572 architecture.

573 **Proposition**

574 A proposition is an expression, normally in a language that has a well-defined written form, that
575 expresses some property of the world from the perspective of a stakeholder. The truth of the
576 proposition may be measured – using a decision procedure – by examining the world and
577 checking that the proposition and the world are consistent with each other.



578
579 *Figure 9 Propositions*

580 There are two kinds of propositions that relate to needs (and policies), assertions and promises.

581 **Assertion**

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

584 **Promise**

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

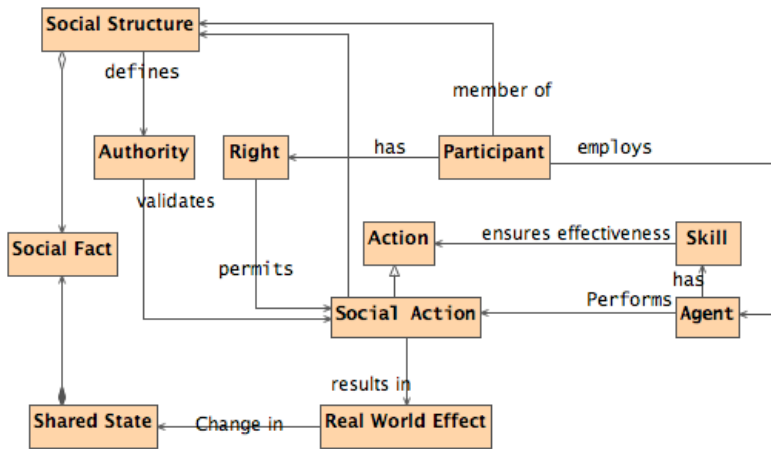
587 For example, an airline may report its record in on-time departures for its various flights. This is a claim
588 made by the airline which may (or may not be) verified. The same airline may promise that some
589 percentage of its flights depart within 5 minutes of their scheduled departure. The truth of this promise
590 depends on the effectiveness of the airline in meeting its commitments.

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

594 **3.4.2 Acting in a social context**

595 The essence of SOA is *action at a distance*: service participants interact with each other, possibly
596 remotely, in order to act. There is always a desire to achieve an effect, an effect in the real world. Of
597 course, there are many possible effects that are desirable and undesirable; we cannot, in general,
598 completely characterize all of the effects of interacting with services.

599 In the context of SOA, actions are primarily social in nature -- one participant is asking another to do
600 something -- and goal oriented -- the purpose of interacting with a service is to satisfy a need; by
601 attempting to ensure that a remote entity applies its capabilities to the need.



602
603 *Figure 10 Acting within Social Structures*

604 **Real World Effect**

605 The result of a participant performing an action in response to a service interaction.

606 **Action**

607 An action is the application of a capability on a target resource by an active entity (otherwise
608 known as the agent of the action) in order to achieve an effect

609 The application of intent by a participant (or agent) to achieve a real world effect.

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

613 When participants interact with each other they are performing **joint actions**:

614 **Joint Action**

615 The application of intent by two or more participants to achieve a real world effect.

616 Intent is at the heart of many social activities: it represents an agent's relationship to one or more of its
617 goals:

618 **Intent**

619 The relationship between an agent and its goals that signifies a commitment by the agent to
620 achieve that goal.

621 Many, if not most, instances of Real World Effect involve acting in the context of a social structure; i.e.,
622 the effect desired is the establishment of one or more social facts.

623 More formally, we can model the concept of a social action:

624 **Social action**

625 A social action is an action that results in a change in the state of a social structure by
626 establishing one or more new social facts. A social action consists of a physical action together
627 with an appropriate authority.

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

631 Social facts typically require some kind of ritual to establish: the action itself is physical, its interpretation
632 is social. For example, the existence of an agreed contract typically requires both parties to sign papers
633 and to exchange those papers. If the ritual is not performed correctly, or if the parties are not properly
634 empowered to perform the ritual, then it is as though nothing happened.

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

637 For example, the hiring of a new employee is an action that is defined by the hiring company (and not, for
638 example, by the president of another company). For a hiring to be valid, it is often the case that specific
639 business processes must be followed, with key actions to be performed only by suitably authorized
640 personnel (such as the company CEO).

641 **Right**

642 A right is a predetermined permission that permits the role player to perform some action or adopt
643 a stance in relation to the social structure and other role players. For example, in most
644 circumstances, sellers have a right to refuse service to potential customers; but may only do so
645 based on certain criteria.

646 **Authority**

647 The right to act on behalf of an organization or another person. Usually, this is constrained terms
648 of the kinds of actions that are authorized, and in terms of the necessary skills and qualifications
649 of the persons invoking the authority.

650 In fact, any entity may authorize another entity to act as its agent. Often the actions that are so authorized
651 are restricted in some sense. In the case of human organizations, the only way that they can act is via an
652 agent.

653 One of the primary benefits of formalizing the relationships between people in terms of groups,
654 corporations, legal entities and so on, is that it allows greater efficiencies in the operation of society.
655 However, corporations, governments and even society, are abstractions: a government is not a person
656 that can perform actions -- only people can actually do things.

657 For example, a fishing club is an abstraction that is important to its members. The club, however, cannot
658 act physically in the world. On the other hand, a person who is appropriately empowered by the fishing

659 club can, and so, when that person writes a cheque and mails it to the telephone company, that action
660 counts as though the fishing club has paid its bills.

661 An artificial agent is somewhere between a person and a corporation: non-human agents definitely can
662 perform actions; however, given foreseeable technology, it would not be reasonable to expect an agent to
663 take responsibility for its actions: instead some combination of the designer, builder, deployer, owner of
664 the agent is ultimately responsible for the actions of non-human agents.

665 **Skill**

666 A skill is a competence or capability to achieve some real world effect. Skills are typically
667 associated with roles in terms of requirements: a given role description may require that the role
668 player has a certain skill.

669 **3.4.3 Transactions and exchanges model**

670 An important class of joint action is the **business transaction**, or **contract exchange**.

671 **Business transaction**

672 A business transaction is a joint action engaged in by two or more participants in which
673 **resources** are exchanged.

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

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

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

681 **Business agreement**

682 A business agreement is an agreement entered into by two or more partners that constrains their
683 future behaviors and permitted states. A business agreement is typically associated with business
684 transactions: the transaction is guided by the agreement and an agreement can be the result of a
685 transaction.

686 Business transactions often have a well defined life-cycle: a negotiation phase in which the terms of the
687 transaction are discussed, an agreement action which establishes the commitment to the transaction, an
688 action phase in which the agreed-upon items are exchanged (they may need to be manufactured before
689 they can be exchanged), and a termination phase in which there may be long-term commitments by both
690 parties but no particular actions required (e.g., if the exchanged goods are found to be defective, then
691 there is likely a commitment to repair or replace them).

692 From an architectural perspective, the business transaction often represents the top-most mode of
693 interpretation of service interactions. When participants interact in a service, they exchange information,
694 perform actions that have an effect in the world, and so on. These exchanges can often (always?) be
695 interpreted as realizing part of, and in support of, business transactions.

696 **Business process**

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

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

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

706 **Process Choreography**

707 The description of the possible interactions that may take place between two or more participants
708 to fulfill an objective.

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

713 **3.4.4 Roles in Social Structures**

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

715 **Role**

716 A role is an identified relationship between a participant and a social structure that defines the
717 rights, responsibilities, qualifications, and authorities of that participant within the context of the
718 social structure.

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

722 **Role player**

723 In the discussions below we refer to the role player as a participant that is acting in a given role.
724 We refer to this as the participant adopting a role.

725 A role player may adopt one or more roles; and have zero or more skills and qualifications.

726 Note that, while many roles are clearly identified, with appropriate names and definitions of the
727 responsibilities, it is also entirely possible to separately bestow rights, responsibilities and so on; usually
728 in a temporary fashion. For example, when a CEO delegates the responsibility of ensuring that the
729 company accounts are correct to the CTO, this does not imply that the CTO is adopting the full role of
730 CFO.

731 Rights, authorities, responsibilities and roles form the foundation for the security architecture of the
732 Reference Architecture. We should be able to trace back any particular security policy to the appropriate
733 relationships between the participants involved, the actions they are performing (or states that they are in)
734 and the Social Structure governing the policy.

735 **Responsibility**

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

738 For example, a role player adopting the role of secretary of a standards group is obliged to ensure that all
739 the minutes of the various meetings are properly recorded; and members of certain standards groups are
740 obliged to declare any pre-existing IP claims that may be relevant to the work of the groups.

741 Rights and responsibilities have similar structure to permissive and obligation policies; except that the
742 focus is from the perspective of the constrained participant rather than the constrained actions.

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

745 For example, what actually happens when an issuing agency determines the status of some stakeholder
746 (i.e., person, corporate entity or non-human agent) is that some person performs some action (such as
747 signing a certificate) that has the effect of qualifying the stakeholder. For that action to be valid, the
748 person signing the certificate has to be empowered by the agency and must be acting within his or her
749 authority.

750 **Qualification**

751 A qualification is a public determination by an issuing authority that a stakeholder has achieved
752 some state. The issuing authority may require some successful actions on the part of the

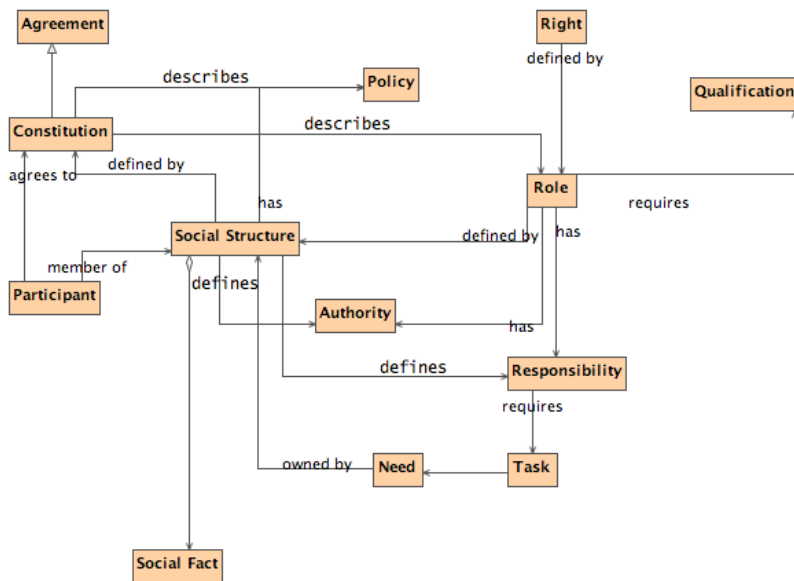
753 stakeholder (such as demonstrating some skills). The qualification may have constraints attached
 754 to it; for example, the certification may be time limited.

755 For example, someone may have the skills to fly an airplane but not have a pilot's license. Conversely,
 756 someone may have a pilot license, but because of some temporary cause be incapable of flying a plane
 757 (they may be ill for example).

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

760 3.4.5 Governance and Social Structures

761 Given that SOA mediates an important aspect of people's relationships, it follows that there are
 762 commitments entered into by participants that require enforcement by the community and that the SOA
 763 itself must reflect the requirements of the community itself.



764
 765 *Figure 11 Social Structures and Governance*

766 Both of these are aspects of the governance of Service Oriented Architecture.
 767 The key elements of our model that relate to governance are the constitution of the social structure, the
 768 policies of the social structure, authority in a social structure, and the associated mechanisms of
 769 enforcement.

770 **Constitution**

771 A constitution is an agreement which defines a social structure. The primary purpose of the
 772 constitution is to define the roles of participants in the institution, and how to establish the
 773 regulations that define the legal actions. The regulations of the social structure effectively define
 774 how those assertions and commitments that are relevant to the social structure are created.

775 For example, a company's constitution is normally called the "Articles of Association". A company's
 776 articles define the officers of the company, their rights and responsibilities and the purpose of the
 777 company. It will often also declare what the rules are for resolving conflicts.

778 A constitution is an agreement, and is also a social fact itself. It is agreed to by the participants in the
 779 social structure. For example, when a new employee joins a company, he or she is often required to sign

780 an employment contract. That contract defines key aspects of the relationship between the new employee
781 and the company.

782 With few exceptions, social structures are embedded in other social structures. One result of this is that
783 the institution's constitution is often viewable as a social fact in one or more outer social structure. For
784 example, the Articles of Association of a company is considered a legal document that supports the legal
785 fact of existence of the company -- by the legal jurisdiction of the company.

786 The main exception to this is, of course, the agreement that defines the constitution of a country. Notably,
787 for most people who are born into the country, its constitution is one that they often do not explicitly agree
788 to. However, it is universal for people who are naturalizing their citizenship to be required to explicitly
789 agree to the constitution of their new country.

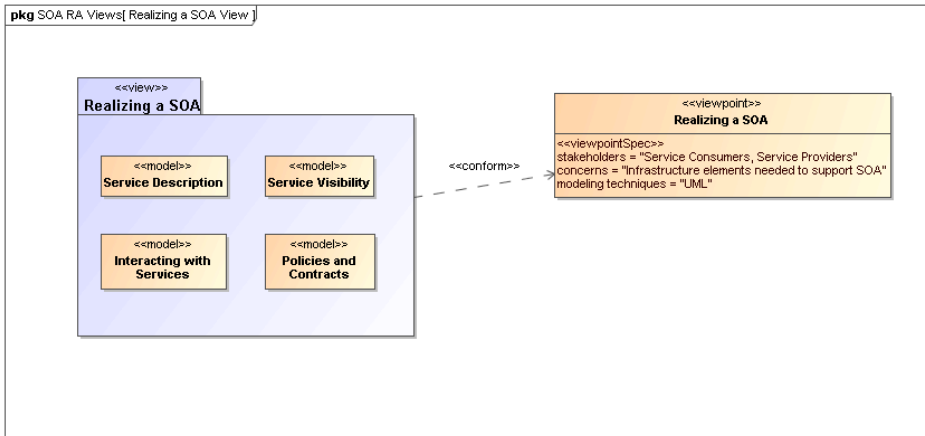
790 **4 Realizing a Service Oriented Architecture View**

791
792
793

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

794 The *Realizing a Service Oriented Architecture* view focuses on the infrastructural elements that are
795 needed in order to support the discovery and interaction with services. The key questions asked are
796 "What are services, what support is needed and how are they realized?"

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



799
800 *Figure 12 Model elements described in the Realizing a Service Oriented Architecture view*

801 **4.1 Service Description Model**

802 SOA depends on a wide variety of descriptions to characterize the needs and capabilities it can facilitate
803 connecting. Description elements, such as those indicating the real world effects produced by a service
804 and those desired by the consumers, provide the basis for determining the match between consumers
805 and providers. Policies and attributes that are needed to evaluate policy compliance are also important
806 elements to determine the conditions under which interactions may be initiated and continue to
807 completion, and description can inform as to which policies may or must be applied.

808 For SOA to enable efficient connectivity between providers and consumers, descriptions must provide
809 sufficient information to achieve visibility between the provider and consumer and to support continued
810 interaction. The information provided by description may be augmented during the interaction. For
811 example, the interaction may reach a point where message exchanges must be encrypted; it may or may
812 not be important that the description indicate that at some point encrypted messages may be required.
813 The critical point is that this additional information becomes available during the interaction and neither
814 the provider nor the consumer is required to have undocumented a priori details about the other, including
815 details of their needs and capabilities, in order for interaction to be initiated or proceed.

816 Several points to make:

- 817 • This model currently focuses on the description of services but it is equally important to consider the
818 descriptions of the consumer and possibly other participants.
- 819 • Descriptions are inherently incomplete. The necessary elements of description depend on the
820 context. The intent of "standard" description sets is to capture "essential" information, i.e. that most

- 821 likely to be needed. It should be understood that what is considered essential will change over time
822 as, for example, the ingredients and nutrition information for food labeling. A requirement for
823 transparency of transactions may require additional description for those associated contexts.
- 824 • Description always proceeds from a basis of what is considered "common knowledge". This may be
825 social conventions that are commonly expected or possibly codified in law. It is impossible to describe
826 everything and it can be expected that a mechanism as far reaching as SOA will also connect entities
827 where there is inconsistent "common" knowledge.
 - 828 • The basis for determining when a description is *sufficient* is quite simple: is it possible for the intended
829 audience of participants to use the descriptions to provide and access services that are offered and
830 used by them. This means that, at one end of the spectrum, a description along the lines of "*That*
831 *service on that machine*" may be sufficient for the intended audience. On the other extreme, a service
832 description with a machine-process-able description of the semantics of its operations and real world
833 effect may be required for services accessed via automated service discovery and planning systems.
834 Generally, somewhere in between these extremes is effective.
 - 835 • Descriptions of the provider and consumer are the essential building blocks for establishing the
836 execution context of an interaction.

837 4.1.1 Components of Service Description

838 A service description is an artifact, usually document-based, that defines or references the information
839 needed to use a service. This includes not only the information and behavior models associated with a
840 service and is needed to define the service interface but also includes information needed to decide
841 whether the service is appropriate for the current needs of the service consumer. Thus, the description
842 will also include information on data associated with service reachability, service functionality, and the
843 policies and contracts associated with a service. More specifically, a service description must convey the
844 following:

- 845 • Service Reachability - The ability for service participants to locate and interact with one another.
846 Reachability includes the address to access the service (i.e., the endpoint), an indication of whether
847 the service is currently available (i.e., the service presence), and the protocols needed to
848 communicate with the service. The service presence may include a static representation of
849 conditions, a representation that is updated at regularly defined intervals, or a dynamic means to
850 assess the current service availability.
- 851 • Service Functionality - An unambiguous expression of service function(s) and the real world effects of
852 invoking the function. The Functions may be expressed as natural language text, reference to an
853 existing taxonomy of functions, or reference to a more formal knowledge capture providing richer
854 description and context. This portion of description should also include technical assumptions,
855 dependencies, or limitations that underlie the effects that can result. [THERE IS OBVIOUSLY A
856 CONNECTION BETWEEN FUNCTIONS (VERBS) AND RWE (NOUNS). DO WE NEED BOTH FUNCTIONS AND
857 EFFECTS EXPLICITLY DESCRIBED?]
- 858 • Interaction Policies & Contracts - Information for prospective consumers pertaining to conditions or
859 constraints when interacting with a service. Whereas technical assumptions, dependencies, and
860 limitations described under the Service Functionality are statements of "physical" fact, policies are
861 subjective assertions made by the service provider (sometimes as passed on from higher authorities)
862 and contracts are agreements on policies between the service provider and some consumer or
863 consumer community. For example, a policy can be that the consumer must have purchased a
864 particular license to use the service, but a contract can identify a specific consumer community who
865 operates under a separately negotiated license. Policies & Contracts are often associated with
866 Performance Metrics through Service Level Agreements, as discussed below.
- 867 • Information Model - A definition of the data model required to exchange information with a service
868 and invoke actions defined in the behavior model. The Information Model includes both the Structure
869 and Semantics of the information exchange. Messages are constructed that conform to the
870 structures defined in the Information Model, and the pattern for using individual messages follows that
871 defined by the Message Exchange Pattern.

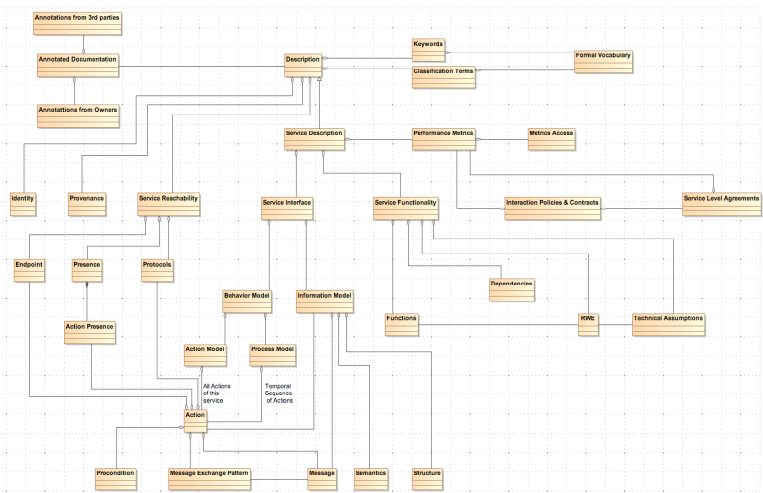
- 872 • Behavior Model - The characterization of the responses to and the temporal dependencies between
873 the actions that can be performed on a service. The Behavior Model is discussed in more detail
874 below.
- 875 – Action Model - The description of actions that may be performed against a service.
- 876 – Process Model - The temporal relationships between service actions and events
- 877 [ACTIONS SEEM TO CORRESPOND TO WSDL OPERATIONS AND I THINK I NEED
878 TO MAKE THIS CONNECTION. WSDL DOESN'T HAVE THE EQUIVALENT OF A
879 PROCESS MODEL, AND I THINK THIS WOULD BE EXPECTED TO BE HANDLED BY
880 SOME REFERENCED BPEL ORCHESTRATION. SO THE DIFFICULTY IS HOW TO
881 KEEP THIS FROM GETTING TOO DEEPLY INTO WS AND HOW TO AVOID THE
882 FACT THAT THERE ARE TWO RELEVANT OVERLAPPING BUT NOT IDENTICAL
883 WSDL VERSIONS.]

884 4.1.2 The Model for Service Description

885 These aspects of description and their relationships are among the elements of description are shown in
886 Figure 13. In this figure, Service Description is shown as a subclass of a general Description class.
887 Participant description is another subclass that is not elaborated here. The mechanism by which values
888 are associated with description is illustrated in Figure 14.

889 Any description should have associated identity – so a description instance can be referenced – and
890 provenance – so the entity responsible for the subject being described can be identified. In addition, the
891 description may associate its subject with predefined keywords or classification taxonomies that derive
892 from reference-able formal definitions and vocabularies.

893 The general description instance may also reference associated documentation that is in addition to that
894 considered necessary in this model. For example, the owner of a service may have documentation on
895 best practices for using the service. Alternately, a third party may certify a service based on their own
896 criteria and certification process, but such a certification may be vital information to other prospective
897 consumers if they were willing to accept the certification in lieu of having to perform another certification
898 themselves. Note, while the examples of Associated Documentation are related to services, the concept
899 applies equally to description of other entities.



900
901 *Figure 13 Service Description*

902 The Service Description subclass may also be composed of information related to Performance Metrics.
903 As with many quantities, the performance metrics are not themselves defined by the Service Description
904 but metric values or the means to access such values may be an important part of description. It is from
905 such performance metrics that Service Level Agreements may be defined.

906 4.1.3 Service Description in support of Service Interaction

907 If we assume we have awareness, i.e. access to relevant descriptions, the service participants must still
908 establish willingness and presence to ensure full visibility [ref visibility section] and interact with the
909 service. The agreements that establish conditions for willingness are collected in the execution context of
910 the interaction. The execution context can be thought of as a series of answers to the questions of why
911 would the participants be willing to interact and whether such interaction is possible.

912 From a description standpoint, a consumer would show interest in a service if the service functionality is
913 what is needed and the service policies are at least worth pursuing if not immediately acceptable. By
914 saying functionality is of interest, we are saying the (business) functions and real world effects (RWE) are
915 of interest and there is nothing in the dependencies or technical assumptions that is a showstopper. Note
916 at this level, the business functions are not concerned with the action or process models. These models
917 get into the nuts and bolts of making the business function happen and will be dealt with at that level later.

918 A service can result in more than one RWE from a business function. In addition, there can be multiple
919 dependencies for the service to successfully complete its functions and there can be numerous policies
920 that provide conditions that may affect willingness. For a service with a single business function (see
921 extension to multiple business function variations below) identifying interaction policies and applicable
922 contracts for the service as a whole and dependencies, technical assumptions, and real world effects as
923 part of service functionality is sufficient.

924 At this point, let us assume the descriptions were sufficient to establish willingness; the details of working
925 this out are considered elsewhere. Figure 12 indicates the service endpoint establishes where to go to
926 actually carry out the interaction. This is where we have to start considering the action and process
927 models.

928 We may have multiple actions a user can perform against a service and the user would perform these in
929 the context of the process model. For a given business function, there is a corresponding process model,
930 where any process model may involve multiple actions. Each action has its own endpoint and also its own
931 protocols associated with the endpoint¹ and whether there is presence for the action through that
932 endpoint. How presence through any endpoint relates to presence of the service is still an open question
933 but likely presence of a service is an aggregation of the presence of the service's actions, and the service
934 level may aggregate to some degraded or restricted presence if some action presence is not confirmed.
935 For example, if error processing actions are not available, the service can still provide required
936 functionality if no error processing is needed. This implies Reachability in some local sense for each
937 action and Reachability also applying at the service/business function level.

938 An action may have preconditions where a precondition is something that needs to be in place before an
939 action can occur, e.g. confirmation of a precursor action. Whether preconditions are satisfied is evaluated
940 when someone tries to perform the action and not before. Presence for an action means someone can
941 initiate it and is independent of whether the preconditions are satisfied.

942 A service may have dependencies. As stated above, the presence of a service is some aggregate of the
943 presence of its actions. A dependency does not affect the presence of a service although it may affect
944 whether the business function successfully completes.

945 In summary to this point, (1) actions has reachability information, including endpoint and presence, (2)
946 presence of service is some aggregation of presence of its actions, (3) action preconditions and service
947 dependencies do not affect presence although these may affect successful completion.

948 Having established visibility, the interaction can proceed. Given a business function, the consumer knows
949 what will be accomplished (the service functionality), the conditions under which interaction will proceed
950 (service policies and contracts), and the process that must be followed (the process model). Given the
951 process model, the consumer knows which actions need to be performed; given the action, the consumer
952 knows the endpoint and protocol to be used and whether there is presence for the action. The remaining
953 question is how does the description information for structure and semantics enable interaction.

954 In the discussion above, we indicate the importance of the process model in identifying relevant actions
955 and their sequence. Interaction with the actions are through messages and thus it is the syntax and
956 semantics of the messages with which we are concerned. There seems to be a number of ways to
957 approach this but the common way now² is to define the structure and semantics that can appear as part
958 of a message and then assemble the pieces into messages and associate messages with actions.
959 Actions make use of structure and semantics as defined in the information model to describe its legal
960 messages. In addition, MEP defines sequencing and use of messages for a given action.

961 So to continue from above, the process model identifies actions I need to perform against a service and
962 the action sequence. I check to see what protocol bindings are available and then check for the endpoint
963 and (possibly) whether there is presence at that endpoint. The interaction with actions is through
964 messages that conform to the structure and semantics defined in the information model and the message
965 sequence conforming to the action's identified MEP. The result is some portion of the RWE initially
966 examined in the service description (e.g. only that part that covers the processing error generated).

¹ 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. [While I do not intend to make this WSDL specific, it should have an obvious mapping.]

² WSDL defines syntax through <types> and SAWSDL proposes to add a pointer to semantics at various places in a WSDL document.

Ken Laskey 5/16/07 5:17 PM
Formatted: Font:9 pt

Ken Laskey 5/16/07 5:29 PM
Formatted: Font:9 pt

967 **4.1.3.1 The question of multiple business functions**

968 It is assumed a service provides a well-defined business function. That is simple. Could (should) it also
969 optionally provide variations of the business functions, e.g. different qualities of service. [I HAVE
970 PREVIOUSLY ARGUED NOT BUT LET'S SEE WHERE THIS TAKES US. I WOULD BE INTERESTED TO HAVE OTHER
971 VARIATION EXAMPLES OTHER THAN QoS FOR WHICH VARIATIONS MAKE MORE SENSE THAN SEPARATE SERVICES
972 (BUT DRAWING A BLANK AT THE MOMENT).]

973 As noted, a service can have more than one RWE from a business function. There can also be multiple
974 dependencies. If there is more than one business function, then we probably need to identify whether a
975 given dependency or RWE is connected with the service as a whole or one of the business functions. (If
976 there is one business function, relating to the service or its business function is one and the same.)

977 **WOULD WE SAY THAT BY DEFINITION ALL DEPENDENCIES, RWE, AND TECHNICAL ASSUMPTIONS RELATE TO ALL**
978 **BUSINESS FUNCTION VARIATIONS AND SO CAN BE CONNECTED AT SERVICE FUNCTIONALITY? ARE THERE**
979 **REASONABLE COUNTEREXAMPLES?**

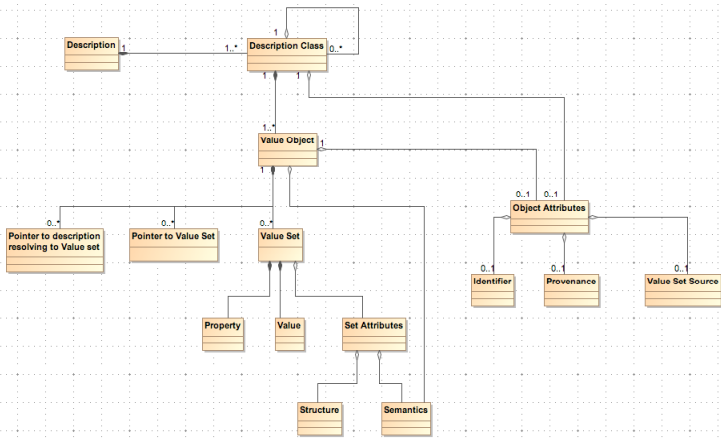
980 We also have policies connected with a service. If we have multiple business functions, does it make
981 sense that some policies relate to only a subset of the functions? One could would say the QoS variation
982 does not require this because we can have a single QoS policy that encompasses the variations. **[ARE**
983 **THERE COUNTER-EXAMPLES TO SAYING POLICIES RELATE TO THE SERVICE AND THE EXPRESSION OF POLICY AT**
984 **THE SERVICE LEVEL CAN ENCOMPASS POLICY VARIATIONS ASSOCIATED WITH THE BUSINESS FUNCTION**
985 **VARIATIONS?]**

986 As noted above, we may have multiple actions a user can perform against a service and this does not
987 change with multiple functions. The idea of presence, especially an aggregated determination of
988 presence, would likely be affected. A business function corresponds to a process model, so **multiple**
989 **business functions imply multiple process models** because either the process is different or the specific
990 used in multiple process

991 **s function** because the

992 **rocess model, (2) its own**

993
994
995



996

997 *Figure 14 Representation of a Description Class*

998 Figure 13 shows the template for a service description but individual description instances depend on the
999 ability to associate meaningful values; this is described through the structures shown in Figure 14. As
1000 shown, each class is represented by a value object (N.B. this is a temporary term that should be changed
1001 to something less computer code oriented) or is made up by components that will eventually resolve to a
1002 value object. For example, Description has several components, one of which is Identity. Identity will be
1003 represented by a value object.

1004 A value object consists of

- 1005 • a collection of value sets with associated property-value pairs, pointers to such value sets, or pointers
1006 to descriptions that eventually resolve to value sets that describe the component; and
- 1007 • attributes that qualify the value sets.

1008 The qualifying attributes include

- 1009 • an optional identifier that would allow the value set to be defined, accessed, and reused elsewhere;
- 1010 • provenance information that identifies the party (individual, role, or organization) that has
1011 responsibility for assigning the value sets to any description component;
- 1012 • an optional source of the value set, if appropriate and meaningful, e.g. if a particular data source is
1013 mandated.

1014 If the value object is contained within a higher-level component, (such as Service Description containing
1015 Service Functionality), the component may inherit values for the attributes from its container.

1016 Note, provenance as a qualifying attribute of a value object is different from provenance as part of a
1017 general instance of Description or, more specifically, a service description. Provenance for a service
1018 identified who “owns” the service, i.e. who is responsible for its creation, maintenance, and provisioning.
1019 Provenance for a value object identifies who is responsible for choosing and assigning values to the value
1020 sets that comprise the value object. It is assumed that granularity at the value object level is sufficient and
1021 provenance is not required for each value set.

1022 The value set also has attributes that define its syntax and semantics.

- 1023 • The semantics of the value set property should be associated with a semantic model conveying the
1024 meaning of the property within the context for use, where the semantic model could vary from a free
1025 text definition to a formal ontology.
- 1026 • For numeric values, the syntax would provide the numeric format of the value and the “semantics”
1027 would be conveyed by a dimensional unit with an identifier to an authoritative source defining the
1028 dimensional unit and preferred mechanisms for its conversion to other dimensional units of like type.
- 1029 • For nonnumeric values, the syntax would provide the data structure for the value representation and
1030 the semantics would be an associated semantic model.
- 1031 • For pointers, architectural guidelines would define the preferred addressing scheme.

1032 The value object may indicate a default semantic model for its component value sets and the individual
1033 value sets may provide an override.

1034 **4.1.5 Relationship to Other Description Models**

1035 While the representation shown in Figure 14 is derived from considerations related to service description,
1036 it is acknowledged that other metadata standards are relevant and should, as possible, be incorporated
1037 into this work. Two standards of particular relevance are the Dublin Core Metadata Initiative (DCMI) and
1038 ISO 11179, especially Part 5.

1039 When the service description (or even the general description class) is considered as the DCMI
1040 “resource”, Figure 14 aligns nicely with the DCMI resource model. While some differences exist, these
1041 are mostly in areas where DCMI goes into detail that is considered beyond the scope of the current
1042 Reference Architecture. For example, DCMI defines classes of “shared semantics” whereas for the
1043 Reference Architecture, it is sufficient to prescribe that an identification of relevant semantic models is
1044 sufficient. Likewise, the DCMI “description model” goes into the details of possible syntax encodings
1045 whereas for the Reference Architecture it is sufficient to identify the relevant formats.

1046 With respect to ISO 11179 Part 5, the metadata fields defined in that reference may be used without
1047 prejudice as the properties in Figure 14 above. Additionally, other defined metadata sets may be used by
1048 the service provider if the other sets are considered more appropriate, i.e. it is fundamental to this
1049 Reference Architecture to identify the need and the means to make vocabulary declarations explicit but it
1050 is beyond the scope to specify which vocabularies are to be used. In addition, the identification of domain
1051 of the properties and range of the values has not been included in the current Reference Architecture
1052 discussion, but the text of ISO 11179 Part 5 can be used consistently with the model prescribed in this
1053 document.

1054 4.1.6 Implications of the Description Model

1055 There are numerous implications that follow from a consideration of the description model shown in
1056 Figure 13.

1057 1. The overall description model is applied to the service and not the components of the service. For
1058 example, the Action Model identifies numerous actions that can be performed against a service and the
1059 Process Model defines the order in which the actions are performed, but the real world effects are defined
1060 for the service and not the individual actions. Similarly, numerous policies may be associated with a
1061 service, but individual policies are not associated with each action. Thus, a SOA service must represent
1062 an identifiable business function to which policies can be applied and from which desired business effects
1063 can be obtained.

1064 1a. Specifying a model where policies are not associated at the action level but only at the service level
1065 goes against many of the discussions of SOA use but it is not obvious that there is really a significant
1066 body of practice that effectively uses such distributed information. If I apply policies at the action level,
1067 how do I describe the service for use when it may have hidden policies embedded anywhere in its
1068 structure? Can I ever describe a service or must I describe every operation/action/endpoint (o/a/e)?
1069 Does this say every o/a/e should be its own service? Does it say the service description has to be
1070 structured in such a way that it describes details of its every o/a/e? Does a service description basically
1071 become the structure that holds all the o/a/e descriptions? What does this say about discovery? Is there
1072 any value in discovering a service or do I discover o/a/e? Does this get me back to every o/a/e is a
1073 service?

1074 1b. For this model, actions are assumed to correspond to WSDL operations – that is not necessary but
1075 it is sometimes useful. An o/a/e is an action with which some external entity must interact in order for the
1076 RWE to be realized. The sequence of such actions is defined in the process model. Only actions needed
1077 to realize the RWE are actions of that service. The RWE should be kept well scoped so a service does
1078 not do a number of unrelated, albeit useful, things. For example, while stock quotes are nice and weather
1079 reports are nice, you wouldn't include these as o/a/e for one service. I would argue that different QoS are
1080 different services. Note, more than one service can access the same capability, and this is appropriate if
1081 a different RWE is provided. I would argue that getting a response in one minute rather than one hour is
1082 more than a QoS difference; it is a fundamental difference in the business function I am receiving.

1083 1c. That all said, RWE in Figure 13 connects to both Functionality and Action, and Policies connect to
1084 Actions. What does that mean? For RWE, it is simple: the functionality is described in terms of the RWE
1085 realized and Actions against the service result in the RWE. However, the RWE can be described at the
1086 service level and the consumer typically does not need to know which actions in particular are needed for
1087 which RWE. This, however, may need to be discussed in regard to availability as questioned in item 6
1088 following. **HOWEVER, IF IN GENERAL THE CONSUMER DOESN'T NEED TO KNOW WHICH**
1089 **ACTION(S) LEAN TO A PARTICULAR RWE, IS THIS LINE NEEDED IN THE SERVICE DESCRIPTION**
1090 **DIAGRAM?**

1091 1.d Back now to policies and their connection with actions, policy alternatives should result in different
1092 RWE and this relationship should be unambiguously expressed at the service level. Analogous to what
1093 was noted in 1.c., the specific policy relationship to a particular action is not something the consumer
1094 typically needs to know, although there may be an availability issue to consider. **SO, IF IN GENERAL**
1095 **THE CONSUMER DOESN'T NEED TO KNOW WHICH POLICIES RELATE TO WHICH ACTION(S), IS**
1096 **THE LINE NEEDED IN THE SERVICE DESCRIPTION DIAGRAM?**

1097 1.e The next question comes in at item 4 below when we take up preconditions.

1098 3.1. A few words on identity and privacy [likely to be moved elsewhere]

1099 Identity needs to be sufficient so one can evaluate provenance when deciding on use of a service but not
1100 so complete as to violate privacy. For example, I may make my office phone number publicly available
1101 but I do not want to do the same with my cell phone number. For a participant description, there will be
1102 the publicly available information, and Web Services would provide an interesting mechanism for
1103 attribute-based access control of additional information. So my cell phone number could be accessed by
1104 someone in the role of emergency medical personnel. In a less dramatic fashion, someone who is a
1105 registered user of my service could get access to me for priority support.

1106 **4.1.7 Identity and versioning**

1107 Identity requires an unambiguous identifier. In some cases, a new identifier could be assigned for each
1108 version, but this would get tedious, would break interfaces when there were updates, and violates Web
1109 architecture which states there should be no proprietary decomposition of a URI to extract additional
1110 information. Assuming a resource has one identifier that covers multiple versions, then the description
1111 should also unambiguously identify the version value and the version definition from which the values
1112 derives. For example, a versioning scheme of (i.j.k) can indicate

1113 * increment i when enhancements do not guarantee backward compatibility;

1114 * increment j for enhancements for which backward compatibility is guaranteed for any previous j for
1115 the same i;

1116 * increment k for bug fixes with no compatibility impacts.

1117 From the perspective of a flexible SOA architecture, it is not important and should probably not be
1118 codified within the general SOA description as to what are the details of an applied versioning scheme but
1119 it should be clear what scheme is being used and what is available to the consumer to decipher a value
1120 derived from the scheme. Similar principles will be used when discussing attributes such as status.

1121 <A discussion of how version designation/value should be impacted by change of version of any of its
1122 components (or any descriptions?) still needs to be discussed.>

1123 **4.1.8 Service Description and Service Level Agreements.**

1124 Service Level Agreements (SLAs) receive a great deal of attention from many of those interested in SOA
1125 because they allude to metrics that can be the basis for charging models and thus the economic models
1126 for SOA. SLAs also have a comfortable feel as the business agreements through which business
1127 services can be procured. In its extreme, SLAs encompass all aspects of service description because it
1128 would prescribe the details of service use and the conditions under which performance metrics are
1129 obtained and analyzed. Unfortunately, the expanded role of SLA as service description does not support
1130 the primary discriminators of services for which contracts of use are not in place and results in massive
1131 redundancies if the characteristics of the service must be repeated for the enforcement functions
1132 surrounding such SLAs.

1133 In the current model, SLAs have a more limited function that derives from the idea of associated policies
1134 and agreements on policies that eventually form the basis for service interactions proceeding within an
1135 execution context. To begin, providers and consumers both have policies and possibly technical
1136 assumptions (i.e. physical, representational, or system-level constraints) that must be aligned in order for
1137 interactions to proceed among service participants. The agreements which encapsulate the necessary
1138 alignment form the basis upon which interactions may proceed – in the SOA Reference Model, this
1139 collection of agreements and the necessary environmental support establish the execution context. Note
1140 that the policies and constraints are properties of the participants and are the basis for any future
1141 interactions among for new or alternate participants but the final contracts/agreements are less reusable
1142 and more specific to a particular execution context.

1143 Where then do SLAs fit in? In the current model, SLAs are the agreed upon values that performance
1144 metrics are intended to satisfy for any contract. The metrics may apply to a single instance of an
1145 execution context, over several reuses of a single or prescribed set of execution contexts, or at the
1146 extreme, an average over all interactions. The metrics can be cumulative or scoped to a particular period
1147 of time. The SLA is the prescribed value associated with an identified Performance Metric, where the
1148 Performance Metrics are fundamental to the service. Using such a definition, SLAs become valuable as

1149 measurable targets and can be modified as appropriate without affecting more permanent and more
1150 broadly relevant elements of the service description.

1151 **4.1.9 Consumer Description**

1152 Illustration - Consumer Description

1153

1154 What's not listed in the SOA RM but has to be done for an implementation is Service Description
1155 versioning and version compatibility.

1156

1157 <DETAILS OF CONSTRUCTING AN INFORMATION MODEL GOES ELSEWHERE BUT NOT SURE
1158 WHERE. HOW TO REFERENCE INSTANCES OF MODELS OR MODEL DEFINITIONS GOES IN
1159 DESCRIPTION.>

1160 **4.1.10 Information Model**

1161 Within the context of SOA, the information model is a characterization of the information that is
1162 associated with the use of a service. The information model describes the structure, format, and meaning
1163 of information and data that may be exchanged with a service as well as prescribing what information
1164 needs to be provided to the service in order to access its capabilities and interpret responses.

1165 **4.1.10.1 Data-Level Information Model**

1166

1167 Layering

1168

1169 Data Formats, Elements and Definitions

1170

1171 Schema

1172

1173 **4.1.10.2 Message Level Information Model**

1174 The message level information architecture can be divided into several areas—the content type of
1175 messages exchanged within the architecture, the type of packages in which they are enclosed, the
1176 metadata associated with messages, the topic space in which they are exchanged, and the security
1177 information carried by them.

1178 **4.1.10.2.1 Message Types**

1179 Messages should be classified into types based upon their content. The type of a message will usually
1180 indicate how it is processed by the receiver. Messages of differing types are most commonly
1181 distinguished by the format and syntax of their content, but can also be distinguished by the format of the
1182 message envelope, a type name, and, for example, a URL endpoint, an applicable XML schema, etc., or
1183 even by source or destination applications.

1184 The taxonomy of message type names should be defined such that it aids in interpretation of the
1185 message content. Message types generated by a particular service or subsystem should begin with a
1186 common prefix or namespace specifier. Some message types may fall into a natural hierarchy, for
1187 example, [SPECIFY AN EXAMPLE]. The taxonomy for message types could embed this hierarchy within
1188 the message type name.

1189 The metadata accompanying a message should indicate its type. The message type name should be
1190 adequate for locating and interpreting all registry, schema, and data dictionary information describing the
1191 message content. This information should be published in the metadata registry-repository.

1192 <THIS SHOULD BE HANDLED IN INTERACTION. AGAIN, HOW TO REFERENCE INSTANCES OF
1193 MODELS OR MODEL DEFINITIONS GOES IN DESCRIPTION.>

1194 4.1.10.2.2 Message Topic

1195 In some circumstances it may be relevant to identify the topic of subject of a message. For example, *this*
1196 *message is about your invoice #324510*. This allows processing of messages based on topic identifiers
1197 rather than functional end-point: for example, *invoices go here and purchase requests go there*.

1198 4.1.10.2.3 Topic Space Division

1199 In publish-and-subscribe (pub/sub) messaging, topics, queues, and/or messaging endpoints divide up the
1200 subject space. For the purposes of this reference architecture, all such divisions will be referred to as
1201 topic divisions. NOTE: Message exchange patterns (MEPs) such as pub/sub are described in the
1202 Interacting with Services model of the Realizing Services view.

1203 Topics need names that uniquely identify them to the participants and logically divide up the message
1204 space. The granularity of these is also important. Too many divisions can be hard to manage and may
1205 require subscribers to listen to many topics, or require publishers to duplicate messages on many topics.
1206 Too few divisions can result in subscribers getting far more messages than they need. (Metadata
1207 keywords in the message header should be used to mitigate the latter problem by allowing subscribers to
1208 filter messages). In general, the following guidelines apply:

1209 Topic division should be consumer-oriented. If all the primary consumers of messages generated by a
1210 given publisher divide messages up by country ID and by no other classification, then topics should be
1211 created for the various countries involved, even if this violates other guidelines for topic division.

- 1212 • The fewer different types of messages delivered on a single topic, the easier the consumer's job
1213 tends to be.
- 1214 • Topic divisions should be chosen such that they will balance load without resulting in unnecessary
1215 topic proliferation.
- 1216 • All topics, queues, and endpoints specific to a particular organization or common service should be
1217 named using a common prefix or top-level namespace, for example "ORG2."

1218 The first division of topics within the common top-level namespace should be by message type. In other
1219 words, topic names should follow the naming pattern <ORG-prefix>.<ORG-message-type> If further
1220 qualification is required, more specific topic divisions can be made, and additional suffixes appended to
1221 the topic name.

- 1222 • Temporary endpoints should be used wherever possible and appropriate.

1223 4.1.10.2.4 Topic Discovery

1224 As more messaging topics are made available across organizational boundaries as well as the enterprise,
1225 a means to topic discovery is required. Topic discovery will follow the other information discovery
1226 techniques described in the Service Visibility model. That is, information about the topic will be made
1227 discoverable in agreed upon-compliant metadata catalogs that are searchable using standardized
1228 mechanisms.

1229 Depending on the type of messaging employed, the metadata used to describe topics could be specific to
1230 an organization or more general across multiple organizations. For instance, if it is a topic that is used to
1231 communicate information about business-to-business transactions for say, purchase orders, the topic
1232 metadata could provide information first about the organization itself and further break down the
1233 information to describe the specific information for that organization, in this case purchase orders.

1234 The metadata catalog will provide information on the topic of interest including a number of other pieces
1235 of descriptive metadata that can be searched for including for example, quality of service, security
1236 requirements, etc.

1237 **4.1.10.2.5 Metadata Strategy**

1238 Messages should be accompanied by, or have references to, appropriate metadata. If references are
1239 used, they should be globally resolvable, for example, references to the metadata registry-repository or
1240 globally available URIs. Metadata should describe the message type and pedigree, as well as allowing
1241 subscribers to filter messages appropriately. The primary means for achieving this is through message
1242 headers. Again, this type of metadata should be registered in the metadata registry-repository.

1243 **4.1.10.2.6 Message Headers and Pedigree**

1244 A transport-independent metadata header for each message should identify

- 1245 • The message type
- 1246 • The source and destination of the message
- 1247 • The creation or last modification time of the message
- 1248 • The security properties of the message
- 1249 • The purpose of the message
- 1250 • The expiration time of the message
- 1251 • Any other properties used for filtering by subscribers

1252 While some of this information may be placed into transport specific headers, a standard content header
1253 should also either be present in the “body” of the message, for transport independence, or extractable
1254 from the transport-specific header in a non-transport dependent format. Destination should be included
1255 only as required by the addressing mechanism of the messaging transport, and should then be included
1256 only in the message header and not in the body. This is because, in pub/sub messaging systems,
1257 decoupling of senders from receivers is a key feature—publishers do not have to know where subscribers
1258 are.

1259 The following guidelines apply to metadata content:

- 1260 • Standard metadata keywords and/or tags (e.g., XML tags) should be defined and registered for use in
1261 message content headers
- 1262 • The standard service description documentation should indicate which metadata items are required in
1263 the message header, and which are optional
- 1264 • Message producers may need to place producer or purpose-specific keywords/tags in the message
1265 header as well. These may be metadata keywords used for a particular type of message, event,
1266 organization, subsystem, etc. To avoid naming collisions, these metadata keywords should adhere to
1267 a naming standard that uses prefixes or namespaces to identify their purpose
- 1268 • For maximum transport independence, keywords/tags should be alphanumeric, may include “_”, and
1269 should start with a letter.
- 1270 • It is best if metadata values are confined to simply data types supported by the various schema and
1271 messaging technology standards available (e.g., W3C XML schema or otherwise).

1272 **4.1.10.2.7 Message Format**

1273 A message should generally have the structure shown in the UML package diagram shown below.
1274 Specifics of the syntax and semantics are determined by the transport protocol, and the type of the
1275 message envelope itself is also determined by the transport.

1276 Messages are divided into: 1) the transport header, 2) the transport body, and message attachments. The
1277 transport header includes transport housekeeping, routing, and security information, as well as some
1278 metadata keywords/tags. The transport body contains the data to be transmitted. The transport body is
1279 further divided into two blocks: 3) the transport independent content header and 4) the producer content.

1280 The following guidelines are best practices for formatting messages:

- 1281 • The transport header should harbor routing and security information, and all metadata required by the
1282 specific transport.

- 1283 • The content header should contain required metadata keywords or tags, although they may be
1284 duplicated in the transport header if needed (for instance, for message filtering).
- 1285 • The producer content should be interpretable given the message type indicator included in the
1286 metadata.
- 1287 • Producer content that is text should use a universal data standard whenever possible (e.g., XML).
1288 Producer content that is standard should be defined by a standard schema (e.g., W3C XML schema).
- 1289 • Attachments should be identified using a standard designation that will allow the receiver to interpret
1290 the attachment (e.g., IETF Multipurpose Internet Mail Extensions)(MIME)).
- 1291 • Attachments that are not self-describing should be described by the metadata in the content and/or
1292 transport headers.

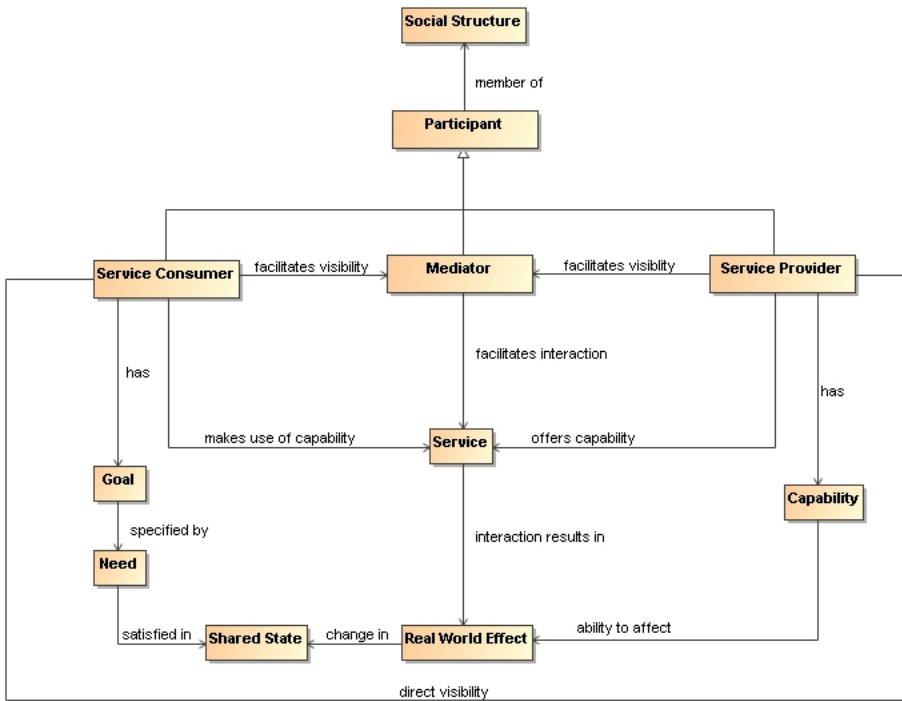
1293 **4.2 Service Visibility Model**

1294 One of the key requirements for participants interacting with each other in the context of an SOA is
1295 achieving visibility: before you can talk to someone, you have to know that they exist.

1296 **4.2.1 Visibility to Business**

1297 The relationship of visibility to the SOA ecosystem encompasses both human social structures and
1298 automated IT mechanisms. It is within the social structure where governance of SOA visibility takes
1299 place, the automated mechanisms for SOA visibility provide greater flexibility in how SOA service visibility
1300 is attained.

1301 Figure 15 depicts a business setting that is a basis for visibility as related to the Business Via Services
1302 View (see Section 3.2). Service consumers have goals specified by needs. Service providers have
1303 capabilities. Mediated or direct visibility between consumers and producers facilitates interactions that
1304 lead to satisfaction of needs. The forms of mediation for visibility provide increased flexibility for
1305 consumer/producer interoperability within complex social structures.

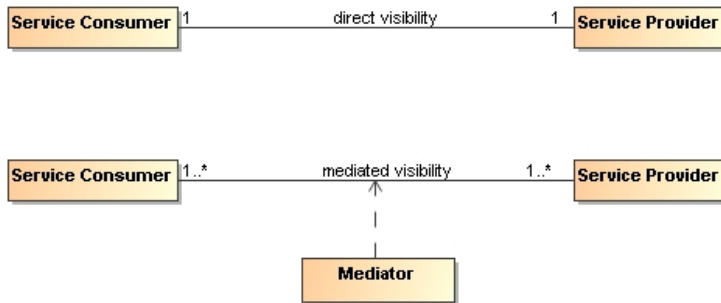


1306

1307 *Figure 15 Visibility to Business*

1308 The IT mechanisms for SOA visibility can play a much more prominent role in the social context of
 1309 computing interoperability than previous computing architectures allowed. Visibility and interoperability in
 1310 a SOA ecosystem requires more than location and interface information, or the traditional Application
 1311 Programming Interface (API). A meta-model for this broader view of visibility is depicted in the Section
 1312 4.1 Service Description Model. In addition to providing improved awareness of service capabilities the
 1313 Service Description can also be valuable for determination of willingness to interact, references to policies
 1314 and contracts in the Service Description for example.

1315 Figure 16 provides a distinction between direct visibility and mediated visibility where direct is a one to
 1316 one relationship between the provider and consumer and mediated is a one to many relationship between
 1317 providers and consumers.



1318

1319 *Figure 16 Direct and mediated visibility*

1320 Another important business capability in a SOA environment is the ability to narrow visibility to trusted
1321 members within a social structure, often referred to as Communities of Interest (COI) in government
1322 sectors. SOA standards and vendor products have been developed and continue to be developed to help
1323 automate the process of becoming aware, determining willingness, and establishing reachability for
1324 service interoperability between complex government and business social structures.

1325 **4.2.2 Attaining Visibility**

1326 Attaining visibility is described in terms of steps that lead to visibility. While there can be many contexts
1327 for visibility within a single social structure, the same general steps can be applied to each of the contexts
1328 to accomplish visibility.

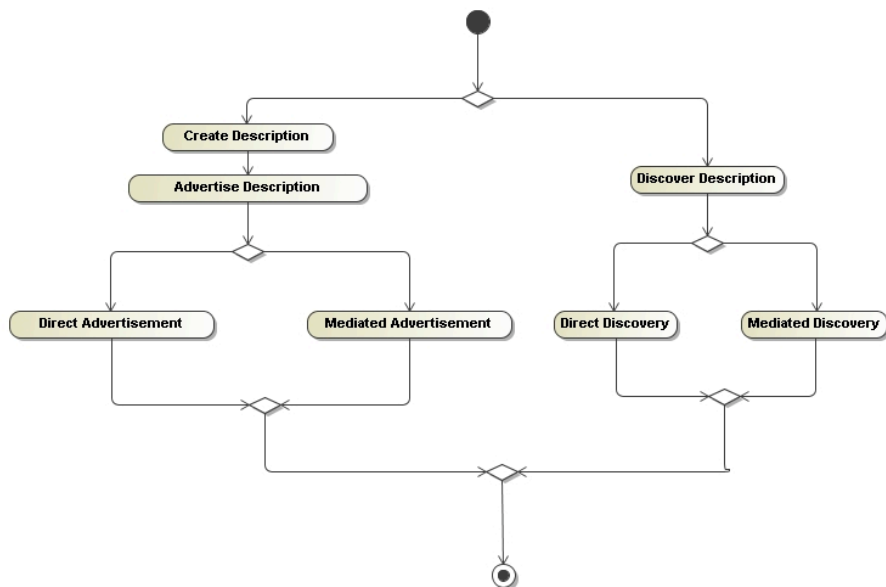
1329 Attaining SOA visibility requires

- 1330 • service description creation and maintenance,
- 1331 • putting description in a place where others can become aware of it,
- 1332 • mechanisms for achieving awareness of description,
- 1333 • mechanisms and/or processes for establishing willingness of participants,
- 1334 • mechanisms to determine reachability.

1335 **4.2.3 Achieving Awareness**

1336 A service participant is aware of another participant if it has access to a description of that participant with
1337 sufficient completeness to establish the other requirements of visibility.

1338 Awareness is a joint activity between two or more participants. Awareness is a joint activity between two
1339 or more participants. Awareness can be decomposed into the creation of description, advertisement of
1340 description, and discovery of description. Awareness is often discussed in terms of consumer awareness
1341 of providers but the concepts are equally valid for provider awareness of consumers. Prior to awareness
1342 in a SOA, a participant creates a description that captures qualifications to advertise to other participants.
1343 Discovery in the Service Visibility Model is the act of a consumer discovering a service description or a
1344 service provider discovering a likely consumer's description. Discovery can be initiated or it can be by
1345 notification. Initiated discovery for business may require formalization of the required capabilities and
1346 resources to achieve business goals. Figure 17 depicts the activities for achieving awareness.



1347
1348 *Figure 17 Achieving Awareness*

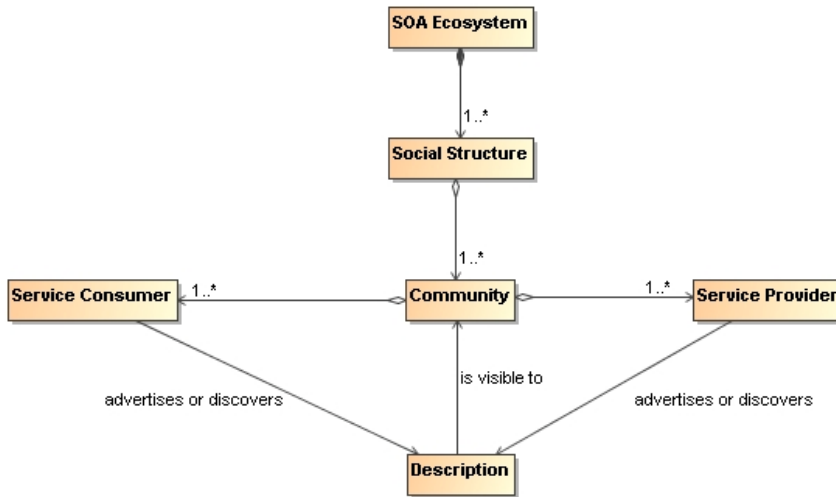
1349 Descriptions may be formal or informal. This SOA RA provides a comprehensive Service Description that
1350 can be applied to formal registry/repositories used to mediate visibility. Using consistent description
1351 taxonomies and standards based mediated visibility helps provide more effective awareness.

1352 **4.2.3.1 Awareness in Complex Social Structures**

1353 Joint awareness applies to one or more communities within one or more social structures where a
1354 community consists of at least one description provider and one description consumer. In Figure 17, joint
1355 awareness can be between a single community, multiple communities, or all communities in the social
1356 structure. The social structure can provide governance of awareness where the governance rules
1357 translate into policies and/or contracts which can then be incorporated into human processes and
1358 automated IT policy/contract mechanisms. The IT policy/contract mechanisms can be used by visibility
1359 access mechanisms to provide awareness between communities. The IT mechanisms for awareness
1360 may incorporate trust mechanisms to assure awareness between trusted communities. For example,
1361 government organizations will often want to limit awareness of an organization's services to specific
1362 communities of interest.

1363 Another common business model for awareness is maximizing awareness to communities within the
1364 social structure, the traditional market place business model. A centralized mediator often arises as a
1365 provider for this global visibility, a gatekeeper of visibility so to speak. For example, Google is a
1366 centralized mediator for accessing information on the web. As another example, television networks have
1367 centralized entities providing a level of visibility to communities that otherwise could not be achieved
1368 without going through the television network.

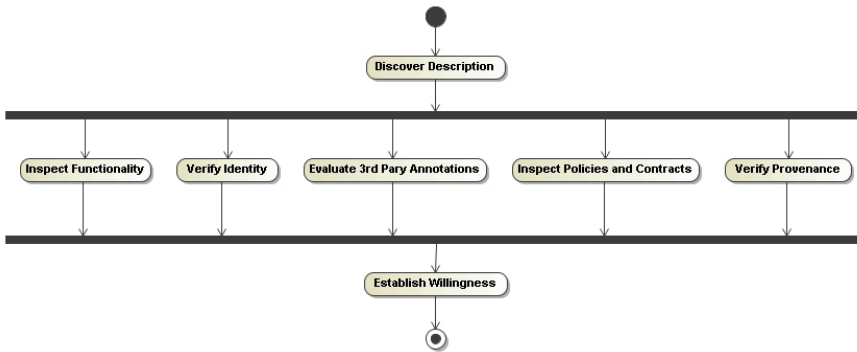
1369



1370
1371 *Figure 18 Joint Awareness*


1372 **4.2.3.2 Establishing Willingness**

1373 Having achieved awareness, participants use descriptions to help determine their willingness to interact
1374 with another participant. Both awareness and willingness are established prior to consumer/provider
1375 interaction. The activities in Figure 19, or a subset there of, can be performed to help establish
1376 willingness;



1377
1378 *Figure 19 Establishing Willingness*

1379 Figure 20 relates elements of the Business via Services View, and elements from the Service Description
1380 Model to willingness. By having a willingness to interact within a particular social structure, the social
1381 structure provides the participant access to capabilities that help satisfy the participant's goals and
1382 objectives as specified by needs. In Figure 20, Information used to establish willingness is defined by
1383 Description. Information referenced by Description may come from many sources. For example, a
1384 mediator for descriptions may provide 3rd party annotations for reputation. Another source for reputation
1385 may be a participant's own history of interactions with another participant.




1386

1387 *Figure 20 Business, Description and Willingness*

1388 Walking through elements referenced by Description, a participant will inspect functionality for potential
 1389 satisfaction of needs. Identity is associated with any participant, however, identity may or may not be
 1390 verified. If available, participant reputation may be a deciding factor for unwillingness to interact. Policies
 1391 and contracts referenced by the description may be particularly important to determine the agreements
 1392 and commitments required for business interactions. Provenance may be used for verification of
 1393 authenticity of a resource.

1394 **4.2.3.3 Determining Reachability**

1395 Reachability involves knowing the service location, service interface, and availability of a service. Figure
 1396 21 lists activities involved to determine reachability.



1397

1398 *Figure 21 Determining Reachability*

1399 **Location**

1400 A Location is the electronic address to where messages are sent. It is the information needed by
 1401 a participant's message delivery mechanism to send a message. Location is referenced from the
 1402 Service Description in the Service Description Model.

1403 **Interface**

1404 Interface verification involves determination of compatible communication protocols, compatible
 1405 message exchange capabilities, and service interface version. Interface is referenced from the
 1406 Service Description in the Service Description Model.

1407 **Presence**

1408 Presence is determined when a service can be reached at a particular point in time. Presence
 1409 may not be known in many cases until the act of interaction begins. To overcome this problem,

1410 IT mechanisms may make use of presence protocols to provide the current up/down status of a
1411 service.

1412 4.2.3.3.1 Re-establishing Reachability

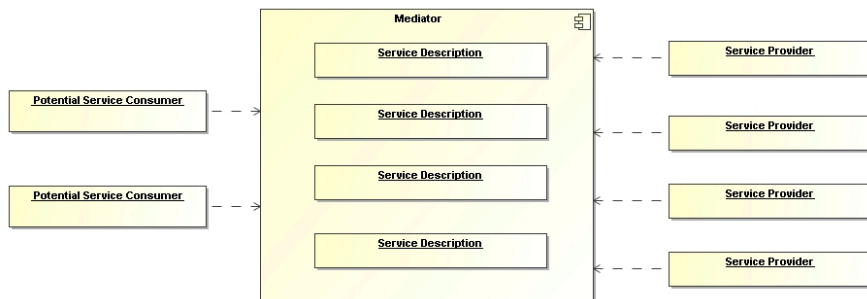
1413 After reachability has been established, there may be times when participants need to re-establish
1414 reachability such as when a service fails and a new location and version for the service needs to be
1415 determined. For SOA, both location and version are important for re-establishing reachability. Multiple
1416 versions of a service may be in operation for backward compatibility. A Domain Name Service (DNS)
1417 lookup for service location may not be sufficient for re-establishing service reachability after a failure.
1418

1419 4.2.4 Mechanisms for Attaining Visibility

1420 Attaining visibility in a SOA can range from word of mouth to formal Service Descriptions in a standards
1421 based registry/repository. Another example of attaining visibility in a SOA is the use of a web page
1422 containing service description information. In this case, the web page is the mediator for visibility. To
1423 gain the greatest degree of flexibility and interoperability in an ecosystem of services, standards based
1424 mediated visibility will most likely be employed for the advertisement and discovery of standards based
1425 service descriptions.

1426 4.2.4.1 Mediated Visibility

1427 Mediation promotes loose coupling by keeping the consumers and services from explicitly referring to
1428 each other and the descriptions. Mediation lets interaction vary independently. Rather than all potential
1429 service consumers being informed on a continual basis about all services, there is a known or agreed
1430 upon facility or location that houses the service description.



1431
1432 *Figure 22 Mediated Service Visibility*

1433 In Figure 22, the potential service consumers perform queries or are notified in order to locate those
1434 services that satisfy their needs. As an example, the telephone book is a mediated registry where
1435 individuals perform manual searches to locate services (i.e. the yellow pages). The telephone book is
1436 also a Mediated Registry for solicitors to find and notify potential customers (i.e. the white pages).

1437 In Mediated Service Visibility for large and dynamic numbers of service consumers and service providers,
1438 the benefits typically far outweigh the management issues associated with it. Some of the benefits of
1439 Mediated Service Visibility are

- 1440 • Potential service consumers have a known location for searching thereby eliminating needless and
1441 random searches
- 1442 • Typically a consortium of interested parties (or a sufficiently large corporation) signs up to host the
1443 mediation facility
- 1444 • Standardized tools and methods can be developed and promulgated to promote interoperability and
1445 ease of use.

1446 However, mediated visibility can have some risks associated with it:

- 1447 • A single point of failure. If the central visibility service fails then a potentially large number of service
- 1448 providers and consumers will be adversely affected.
- 1449 • A single point of control. If the central visibility service is owned by, or controlled by, someone other
- 1450 than the service consumers and/or providers then the latter may be put at a competitive disadvantage
- 1451 based on policies of the discovery provider.

1452 Mediation can also apply to service description metadata. There may be a standard vocabulary or

1453 mediated interaction among known metadata vocabularies.

1454 4.2.4.2 Service Registries and Repositories


1455 While there can be several mechanisms for service visibility in a SOA, a common mechanism for

1456 mediation in the industry is a registry. Figure 23 depicts a mediation facility containing a registry and a

1457 repository. The registry stores links or pointers to service description artifacts. The repository in this

1458 example is the storage location for the service description artifacts. Service descriptions can be pushed

1459 (publish/subscribe for example) or pulled from the register-repository mediator.



1460
1461 *Figure 23 Mediated Registry/Repository*

1462 The registry is like a card catalog at the library and a repository is like the shelves for the books.

1463 Standardized metadata describing repository content can be stored as registry objects in a registry and

1464 any type of content can be stored as repository items in a repository.

1465 4.3 Interacting with Services Model

1466 Interaction is the activity involved in making use of a capability offered in order to achieve a particular

1467 desired "real world effect", where real world effect is the actual "result" of using a service (as opposed to

1468 merely the "capability" offered by a service provider). An activity can be characterized by a sequence of

1469 actions. Consequently, interacting with a service involves performing "actions" against the service, usually

1470 through a series of information exchanges (e.g., messages), although other modes of interaction are

1471 possible such as modifying the shared state of a resource. Note that a participant (or agent acting on

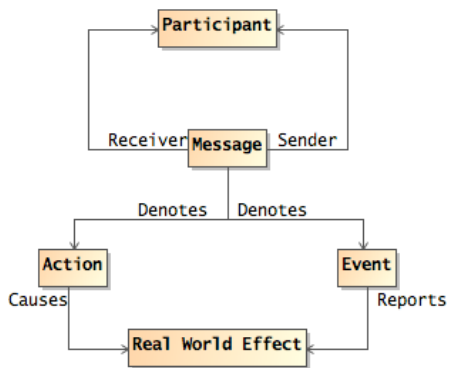
1472 behalf of the participant) can be the sender of a message, the receiver of a message, or both.

1473 4.3.1 Action Model

1474 For purposes of this SOA reference architecture, the authors have committed to the use of "message"

1475 exchange between service participants to denote actions against the services that "cause" a real world

1476 effect, and to denote events that "report" on real world effects that arise from those actions.



1477

1478 Figure 24 A "message" denotes either an action or an event.

1479 A "message" denotes either an action or an event. In other words, actions and events are realized
1480 through messages.

1481 4.3.1.1 Message Exchange

1482 Message exchange is the "means" by which service participants (or their agents) interact with each other.
1483 There are two primary modes of interaction: joint actions (see Section 3.4.2) and communicating real
1484 world effects.

1485 A message exchange is used to effect an action when the messages contain the appropriately formatted
1486 content that should be interpreted as an action and the agents involved interpret the message
1487 appropriately (you have to have a speaker and a listener).

1488 A message exchange is also used to communicate event notifications. An event is a report of an
1489 occurrence that is of interest to some participant; in our case when some real world effect has occurred.
1490 Just as action messages will have formatting requirements, so will event notification messages.

1491 When a message is interpreted as an action, the correct interpretation typically requires the receiver to
1492 perform a set of operations. These *operations* represent the sequence of (private) actions a service must
1493 perform in order to validly participate in a given joint action.

1494 Similarly, the correct consequence of realizing a real world effect may be to initiate the reporting of that
1495 real world effect via an event notification.

1496 Message Exchange

1497 The means by which joint actions and event notifications are coordinated by service participants
1498 (or agents).

1499 Operations

1500 The sequence of (private) actions a service must perform in order to validly participate in a given
1501 joint action.

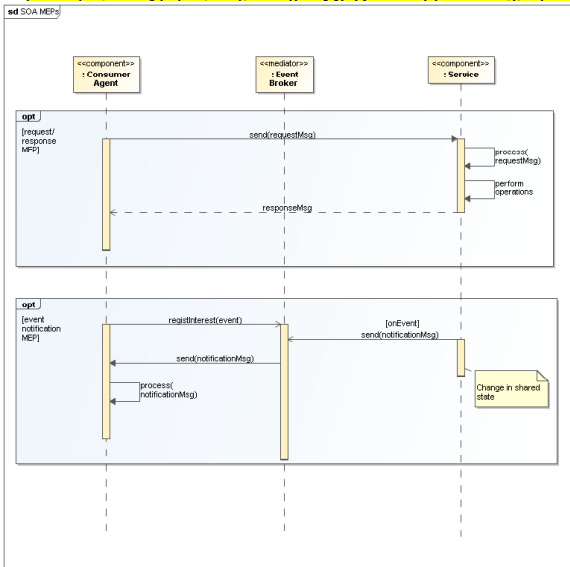
1502 4.3.1.2 Message Exchange Patterns (MEPs)

1503 As stated earlier, this reference architecture commits to the use of message exchange to denote actions
1504 against the services, and to denote events that report on real world effects that arise from those actions.

1505 Because message exchange denotes actions against services, service interaction patterns can be
1506 characterized by a common set of message exchange patterns (MEPs):

- 1507 • Request/response to represent action
- 1508 • Event notification to represent event

1509 This is by no means a complete list of all possible MEPs but it does represent those that are most
 1510 commonly used in exchange of information both within organizations and across organizational
 1511 model, which is referenced
 1512



1513
 1514 *Figure 25 Set of Common SOA Message Exchange Patterns (MEPs)*

1515 The interaction (sequence) diagram reflected in Figure 25 shows four different service interaction
 1516 "options" based on the stated guard condition (e.g., [request/response MEP]). Introduced in UML 2,
 1517 these fragments are known as "interaction fragments" and the "opt" keyword on the interacting fragment
 1518 frame is known as an "interaction operand". In this case, the interaction operand opt stands for "option."
 1519 The way to read these options is by reviewing the guard condition, which is shown in between brackets
 1520 "[guard]" (e.g., [request/response MEP]), and if the condition is true, then this particular interaction
 1521 fragment is executed.

1522 In the interaction diagram shown in Figure 25, it is assumed that the service participants (consumer and
 1523 provider) have delegated message handling to hardware or software agents acting on their behalf. The
 1524 message interchange model illustrated represents a logical view of the MEPs and not a physical view. In
 1525 other words, specific hosts, network protocols, and underlying messaging system are not shown as these
 1526 tend to be implementation specific. While such implementation-specific elements are typically considered
 1527 outside the scope of reference architecture such as this SOA-RA, they are important considerations in
 1528 modeling the "execution context", a subject to be addressed in additional detail in a subsequent section.

1529 **4.3.1.2.1 Request/Response MEP**

1530 In a request/response MEP, the Consumer Agent sends a request message to the Provider Agent. The
1531 Provider Agent (which is the hardware or software component that actually implements the service) then
1532 processes the request message. Based on the content of the message, the Provider Agent performs the
1533 service operations. Following the completion of these operations, a response message is returned to the
1534 Consumer Agent. This type of MEP is considered a "synchronous interaction" because the sender of the
1535 request message (i.e., Consumer Agent) is blocked from continued processing until a response is
1536 returned from the Provider Agent.

1537 **4.3.1.2.2 Event Notification MEP**

1538 An event is realized by means of an event notification message exchange. The basic Event Notification
1539 MEP takes the form of a one-way message sent by a notifier agent and received by agents with an
1540 interest in the event. Often the sending agent may not be fully aware of all the agents that will receive the
1541 notification; particularly in so-called pub/sub situations.

1542 **4.3.1.3 Effects of actions**

1543 Have to connect the message with the change of state that represents the real world effect of an action.

1544 **4.3.2 Process Model**

1545 The Process Model could be translated into a machine processable artifact that could be used like a
1546 schema for message flow validation of the single service. The Process Model contains the Message
1547 Exchange Patterns (MEPS) for the service. The Process Model could also be used by Service
1548 Orchestration for the automation of higher order service interactions.

1549 Getting things done sometimes requires performing more than one action.

1550 From WS-CDL :

1551 - "Business Process Languages layer": describes the execution logic of Web Services based applications
1552 by defining their control flows (such as conditional, sequential, parallel and exceptional execution) and
1553 prescribing the rules for consistently managing their non-observable data

1554 - "Choreography layer": describes collaborations of participants by defining from a global viewpoint their
1555 common and complementary observable behavior, where information exchanges occur, when the jointly
1556 agreed ordering rules are satisfied.


1557 Beyond the need to specify how to combine services to accomplish some business is the need for the
1558 SOA infrastructure to run its own internal processes. An implemented SOA needs a "traffic cop"
1559 process/service to

- 1560 • accept requests,
- 1561 • decide what needs to be done with the request,
- 1562 • route the request or derived requests to other services,
- 1563 • collect the results of other services,
- 1564 • decide on next actions based on behavior model and responses to routed requests,
- 1565 • continue this until the initial request is satisfied or terminated,
- 1566 • package and send results to receiver designated by original requester (where receiver possibly *not*
- 1567 the requester).

1568 The traffic cop will make use of known compositions to make sure routing is done properly and for error
1569 recovery. Whatever the traffic cop decides, it will probably be captured internally as a service
1570 composition and corresponding behavior model. This is likely a useful format for logging and future
1571 audits. It also gives a basis for evaluating levels of service and identifying bottlenecks.

1572 **4.3.3 Orchestration and Service Composition**

1573 Composition – a service visible to a service consumer via a single interface and described via a single
1574 service description is composed of more than one component service, where each component service is
1575 visible to the containing service via a single interface and described via a single service description.
1576 Figure 1 shows a service A that relies on two other services in its implementation. The service consumer
1577 does not know that Services B and C are used, or whether they are used in serial or parallel, or their
1578 operations succeed or fail. The consumer only cares about the success or failure of Service A.



1579
1580 *Figure 26 Service Composition*

1581
1582 "[need to define business processes and their execution (typically called "orchestration" but the term is
1583 not used in the 5 May 2006 WS-BPEL draft) and choreography]"

1584 ----
1585 The combining of services needs to be looked at in the context of not just how you would specify the
1586 combination of more atomic services into a higher level one, but also the behavior models and the
1587 orchestration/choreography that results in generating real world effects. The behavior models must be
1588 consistent (NEEDS TO BE DEFINED) or the combination of services would not make sense.

1589 **4.3.4 Dependencies between services**

1590 Transitive effects of policies

1591 Preconditions

1592 Configuration dependencies

1593 Not all dependencies may be explicitly listed

1594 Service parameters

1595 What needs to be done about dependencies

1596 2. The Service Interface is made up of the Action Model, i.e. things you can do with a service, and the
1597 Process Model, i.e. order you must follow. Here, the actions are only those that result in service's
1598 documented real world effects and not general maintenance/management functions, such as update and
1599 delete. As noted in Figure 1, the action model comprises all the available actions internal to the service,
1600 however, the Process Model may refer to actions of other services. In such a case, the external services
1601 would be listed under Dependencies.

1602 3. The process model implies a maintaining of state. [QUESTION: WHERE IN THE RA WILL WE TALK
1603 ABOUT STATE? TO WHAT EXTENT DO WE FOLLOW A WEB MODEL AND PASS ALL STATE
1604 INFORMATION AROUND VS. HAVING AN INSTANCE OF THE PROCESS TO WHICH THE
1605 CONSUMER NEEDS TO RECONNECT TO MAKE USE OF SAVED STATE?] State will be considered in
1606 the model of Interactions and any appropriate references will be added to Service Description.

1607 4. Whereas dependencies are identified at the service level as part of Service Functionality,
1608 preconditions may be associated with individual actions. Some preconditions are likely satisfied internal to
1609 the service by complying with the process model, but other preconditions may depend on external
1610 conditions that will need to be checked before the action can be successfully performed. [QUESTION: IS
1611 THE RELATIONSHIP BETWEEN ACTIONS IN THE PROCESS MODEL AND PRECONDITIONS JUST
1612 WHETHER INTERNALLY OR EXTERNALLY SATISFIED? DOES THE PROCESS MODEL AND
1613 ACTIONS MADE AVAILABLE TO BE PART OF IT PRIMARILY CONCENTRATE ON WHAT THE
1614 SERVICE FINDS SO IMPORTANT THAT IT NEEDS TO BE HANDLED INTERNALLY BUT WITH
1615 CONSUMER INTERACTION THROUGH THE VARIOUS ACTIONS?] A Dependency is what has to be

1616 there to do something and a Precondition is exactly what that something had to do or accomplish. For
1617 example, a Dependency could be I need access to a service to do currency conversion; the Precondition
1618 for another action is that I got a value from that service so I can proceed with my business. Handling of
1619 preconditions will also be part of the Interaction model and any appropriate references will be added to
1620 Service Description. In response to the question here, Rex noted we will probably need to talk to both
1621 internally satisfied and externally satisfied preconditions. But as with policies and RWE, it is unclear if
1622 service description needs to have the granularity of specifying this to the action level. Again, look at this
1623 from the perspective of discovery: if you specify detail at the action level and you feel this detail is really
1624 necessary, how does it reflect at the service level or is it just a surprise you find when you try to use the
1625 service? **FOR SERVICE DESCRIPTION, IT MAY BE MOST APPROPRIATE TO HAVE ACTION AS A
1626 LEAF NODE AND HAVE RWE ATTACH TO THE PROCESS MODEL AS THE SEQUENCE OF WHAT
1627 YOU DO TO GET EACH RWE.**

1628 5. [likely to require significant massaging, most likely elsewhere] There are numerous complications
1629 with the way Web Services are often described. A service that supports more than one independent
1630 action or set of actions appears to correspond to a Web Service with multiple operations. We have
1631 already said that functionality, policies, and the like are service descriptors and not action descriptors.
1632 Thus, we are restricting the concept of WS operations to actions that support the overall function and not
1633 different ways of doing the same function. For example, the common WS scenario of different operations
1634 providing different qualities of services would be invalid, and these would be different services rather than
1635 operations of the same service. However, having the service checks (through a consumer action) to see if
1636 the requestor is eligible for enhanced QoS before allowing the action to make use of the enhanced QoS is
1637 valid. **THIS IS BEATEN TO DEATH IN WHAT ALREADY BEEN ADDED IN THE ITEMS ABOVE. IN
1638 SUMMARY, IT NEEDS TO BE CONSISTENT WITH THE POLICY AND RWE LINES TO ACTION OR
1639 PROCESS MODEL.**

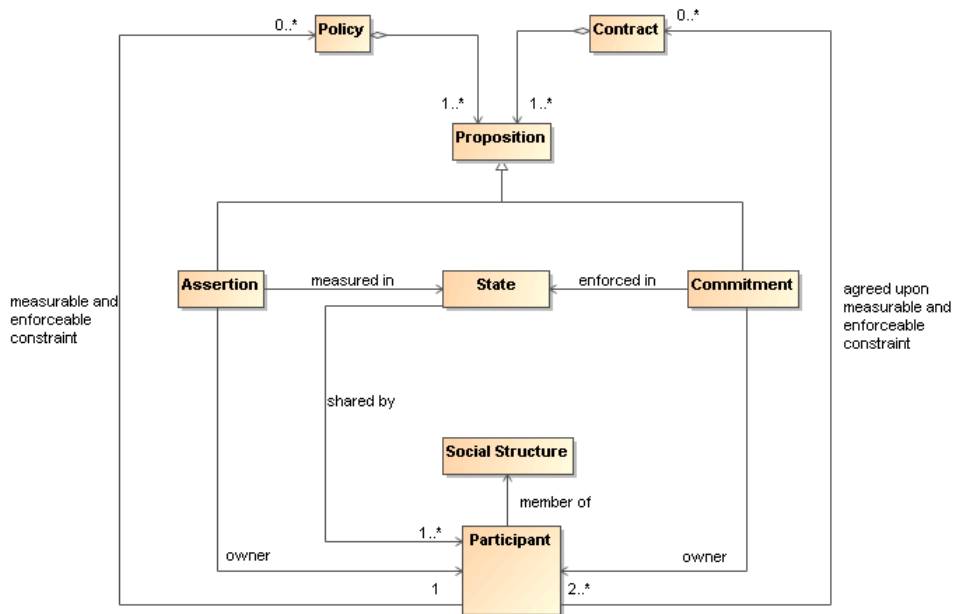
1640 6. **[QUESTION: WE HAVE RESOLVED THAT ACTIONS (AND BY EXTENSION, WS OPERATIONS)
1641 ARE COVERED BY DESCRIPTION AT THE SERVICE LEVE. HOW DOES AN ACTION NOT BEING
1642 AVAILABLE AFFECT THE SERVICE'S OVERALL AVAILABILITY? WHAT DOES THIS INDICATE
1643 ABOUT PRESENCE AS INCLUDED UNDER SERVICE REACHABILITY?]**
1644

1645 **4.4 Policies and Contracts**

1646 A policy is the representation of a constraint or condition on the use, deployment, or description of an
1647 owned entity as defined by any participant. A contract is a representation of an agreement between two
1648 or more participants.

1649 Core aspects of contracts and policies are the constraint assertions, the owners, the measurement and
1650 enforcement of the policy or contract. An assertion may be an expression of a policy and/or a contract.
1651 Assertions are enforceable and measurable statements about the way a service is realized.

1652 In Section **Error! Reference source not found.**, measurable assertions are characterized as
1653 propositions - an expression of some property of the world whose truth can be measured by examining
1654 the world and checking that the expression and the world are consistent with each other.



1655
1656 *Figure 27 Distinguishing between policies and contracts*

1657
1658 Figure 27 derives from the [Business via Services View](#). Policies and contracts are an aggregation of
1659 propositions. Both are measurable and enforceable, however contracts are agreed upon by two or more
1660 participants while a policy may not have been agreed to.

1661 In a business context, contracts are legally binding agreements between two or more parties. A contract
1662 is formed when there is an offer that is duly made and the offer is accepted and there is evidence that
1663 indicates there was a tangible exchange of value between the two parties.

1664 The measurability and enforcement of propositions may include many indirectly related participants within
1665 the social structure. Dispute resolutions, for example, may involve courts. Policy and contract IT
1666 mechanisms support automated governance and management within the SOA ecosystem to improve
1667 governance and management efficiency. Advances in IT standards and technologies dictate the level of
1668 automation achieved in support of SOA governance and management.

1669 Providing automated mechanisms to enforce policies and contracts can help the social structure operate
1670 more efficiently and also enable the social structure to operate at higher levels of abstraction.

1671 Understanding the complete environment which policies and contracts apply in a SOA requires
1672 understanding of the processes surrounding policies and contracts in the social structure, the IT
1673 mechanisms that support automated enforcement of policies and contracts, and the traversal from/to the
1674 social structure to/from the IT policy automation mechanisms.


1675 From the IT perspective, high level policies and contracts are translated into low level rules and
1676 measurable properties. For low level rules and measurable properties, both contracts and policies are
1677 likely to be enforced by the same type of IT policy mechanisms.

1678 Policies and contracts have wide applicability within the Reference Architecture. They are used to
1679 express security policies, service policies, relationships and constraints within the social structures that
1680 encapsulate service participants, management of services and many other instances. The enforcement of
1681 a policy or contract may be a part of the SOA computing environment or it may be handled outside of the
1682 SOA computing environment. The RA is concerned with the underlying principles and IT mechanisms

1683 that support enforceable and measurable contracts and policies in the widest range of situations for a
1684 SOA.

1685 4.4.1 Policy and Contract Relationships

1686 Figure 4-17 depicts relationships between policies and contracts. A contract may include references to
1687 policies and other contracts while a policy may include references to contracts and other policies. For
1688 example, a contract may reference a set of policies and a policy may prioritize certain contracts over
1689 others.



1690

1691 *Figure 4-17 Policy/Contract Relationship*

1692 Policy

1693 For IT, policies are measurable and enforceable rules/assertions for IT mechanisms that define
1694 the choices in behavior of a system.

1695 Contract

1696 Contracts are the set of rules/assertions that define the agreements under which service
1697 functionality is delivered.

1698 A policy may result in the application of different choice of behavior at a particular juncture when
1699 compared to a contract. However, this could be captured in a process model that makes use of the same
1700 policy/contract IT mechanisms. For example, a contract between an Internet Service Provider (ISP) and
1701 a consumer specifying maximum bandwidth usage may result in an automatic increase in the consumer's
1702 bill if that bandwidth is exceeded. An Internet Service Provider (ISP) policy specifying maximum
1703 bandwidth usage may result in the purchase of additional resources by the ISP if monitored bandwidth is
1704 exceeded.

1705 4.4.2 Policies and Contracts Life Cycle

1706 The genesis of a policy or contract will likely start with a governance or management process, see section
1707 [5.1 Governance of Service Oriented Architectures](#) and section [5.3 Services as Managed Entities Model](#).
1708 The governance and management process may use formal and standardized policy languages and
1709 employ the use of common IT mechanisms to maximize computing interoperability and compliance. IT
1710 mechanisms for policies and contracts are discussed further in section [4.4.2 IT Mechanisms Supporting](#)
1711 [Policies and Contracts](#).

1712 4.4.3 Policy Types

1713 When discussing IT policies, two prevalent policy types are **access control policies** and **event-**
1714 **condition-action policies**. Access control policies are constraints applied to authenticated users.
1715 Access control policies ask permission to perform an action. There can be many access control policy
1716 models applied in a SOA, the definition of specific access control models is beyond the scope of this
1717 reference architecture. Event-condition-action policies specify actions to be taken when certain events
1718 occur. Event-condition-action policies are also commonly referred to as obligation policies.

1719 Policies and contracts can contain a mix of permissions and obligations, but policies tend to be
1720 permission oriented and contracts tend to be obligation oriented. The mechanisms for enforcing a
1721 permission-oriented constraint is prevention at the point of action. The mechanisms for enforcing

1722 obligations is assurance of compliance at the point of action. For example, there may be a policy for the
1723 types of customer transactions logged. At the point of action for logging customer transactions, a request
1724 is made to determine if logging should be done. There may be a contract for the types of customer
1725 transactions audited. In this case, auditing requires both logging of the transaction as well as assurance
1726 of compliance with the auditing requirements.

1727 **4.4.4 Policy and Contract Specification**

1728 The language used to describe policies and contracts inevitably constrains the forms and types of policies
1729 and contracts expressible in the description. Formal policy language definitions are outside the scope of
1730 this specification. For formal policy languages, standard specifications such as XACML and WS-Policy
1731 may be referenced. Policy/Contract descriptions may be associated with a service through the Service
1732 Description as defined the section [4.1 Service Description Model](#). Policy enforcement points and policy
1733 decision points can interpret policy descriptions expressed in a consumer/provider agreed upon language
1734 to make policy decisions.

1735 Regardless of the language used to describe policies and contracts, there are certain aspects that must
1736 be captured in any system for representing policies and contracts: how to describe atomic policy
1737 constraints, how to handle the composition of policies, how to resolve conflict between policies and how
1738 to realize enforcement of policy constraints.

1739 **4.4.4.1 Policy Composition**

1740 Multiple policies may be defined for one or more services in one or more ownership domains. The
1741 application of policies and contracts over distributed services requires the ability to compose one or more
1742 policies into an overarching policy. The composition of policies may be implemented as a hierarchy or
1743 nesting and/or it can be implemented as intersections and unions of sets.

1744 **4.4.4.2 Conflict Resolution**

1745 The analysis of policy rules may result in conflicts between the policy rules. There can be many causes
1746 for policy conflicts such as conflicting policy rules between ownership domains or policy language
1747 specifications that do not convert to first order predicate logic for IT policy mechanisms. This can cause
1748 policy decision results to be indeterminate. Policy administration mechanisms may provide conflict
1749 resolution capabilities prior to the storage/distribution of policies. At run time, conflicts may resolve to
1750 higher authorities inside and outside the SOA IT mechanisms.

1751 **4.4.4.3 Delegation of Policy**

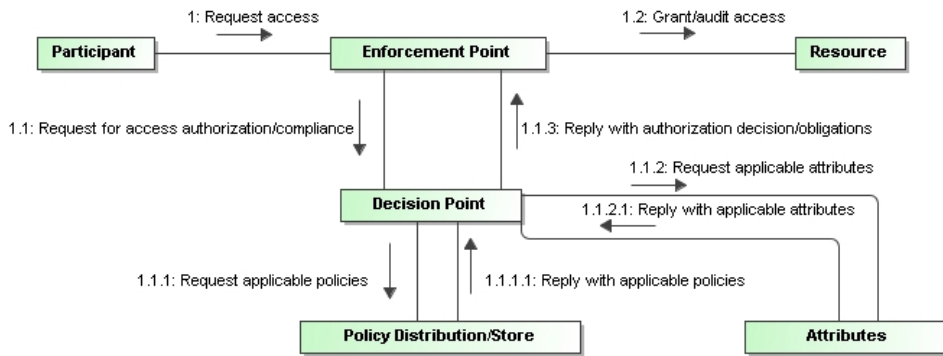
1752 Policy authorization may be delegated to agents acting on behalf of a client to enable decentralized policy
1753 administration and/or policy enforcement. This allows policies to be administered and/or enforced in a
1754 hierarchical fashion. Policies may also be transferred to an agent or resource to effectively allow that
1755 agent or resource to separate from an ownership domain. The agent or resource may join another
1756 ownership domain or rejoin the same ownership domain at a later time.

1757 **4.4.5 IT Mechanisms Supporting Policies and Contracts**

1758 The common policy architectural elements that are provided in this section are based on the minimal
1759 mechanisms required to provide policy guided delivery across distributed services within an ownership
1760 domain and across ownership domains. The same mechanisms can provide compliance assurance
1761 and/or auditing of contractual obligations between participants.

1762

1763 **4.4.5.1 Basic Standards Based Policy and Contract Elements**



1764
1765 *Figure 28 Basic Standards-based Policy/Contract IT Elements*

1766
1767 *Figure 29 Interacting agents in the context of a communications policy*

1768 **Resource**

1769 A resource is any entity that has a name and an owner; see Section 3.3.

1770 **Attributes**

1771 Attributes are named values that define characteristics of participants, resources, actions, or the
1772 environment.

1773 **Decision Point**

1774 The decision point evaluates participant requests against relevant policies/contracts and
1775 attributes to render an authorization decision. The decision point provides a measurement for an
1776 assertion. The decision point renders an authorization decision in the form of permit, deny,
1777 indeterminate, not applicable, or a set of obligations. A decision point may obtain an
1778 authorization decision from a computing mechanism or from outside the computing system,
1779 decisions by humans through workflow for example.

1780 **Enforcement point**

1781 The enforcement point enforces and assures the decision point decisions. In a Service Oriented
1782 Architecture, one policy or contract may be applicable to multiple distributed services. Due to the
1783 distributed nature of a SOA, the enforcement or auditing of authorizations is attributed to an
1784 enforcement point that is separate from the decision point. The enforcement point is responsible
1785 for protecting access and determining access compliance to one or more resources. When
1786 attempting to access a resource, the enforcement point sends a description of the attempted
1787 access to a decision point. The decision point evaluates the request against its available
1788 policies/contracts and produces an authorization decision that is returned to the enforcement
1789 point. Like the decision point, an enforcement point may require a means of enforcement
1790 outside the computing system.

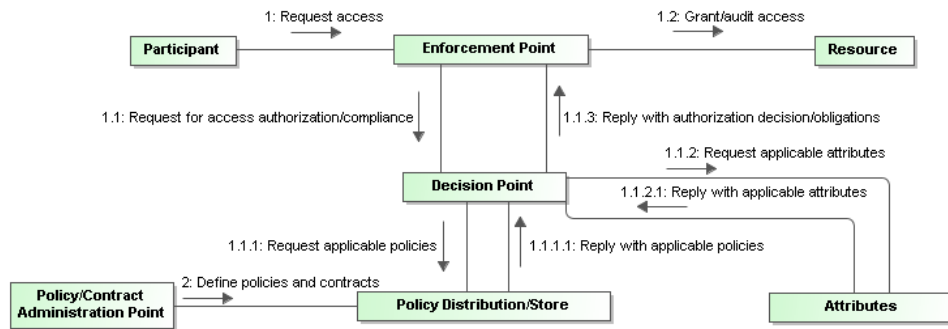
1791 **Policy Distribution/Store**

1792 The Policy Distribution/Store distributes policy to decision points or stores policies for retrieval by
1793 decision points.

1794 **4.4.5.2 Policy/Contract Administration**

1795

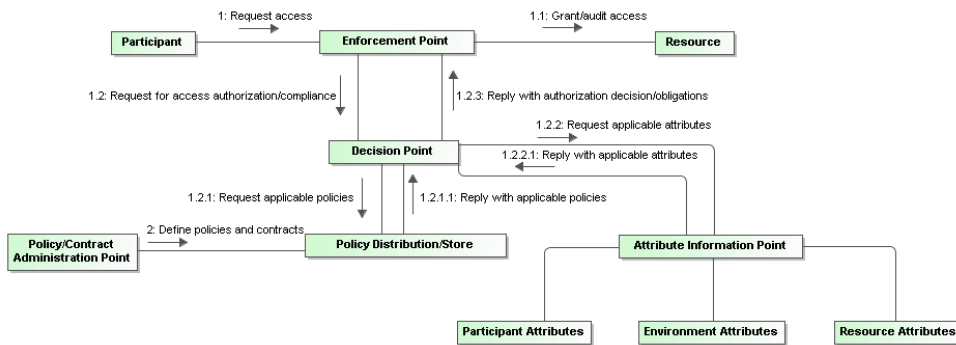
1796 A Policy/Contract administration point can provide many enterprise SOA policy/contract administration capabilities but the end result of the administration point is to store or distribute policy/contract updates.
 1797



1798
 1799 *Figure 30 Policy/Contract Administration Point*

1800 **4.4.5.3 Policies/Contracts and Attributes**

1801 There are many possible approaches to the management and application of policy/contract attributes. A
 1802 commonality of the application of attributes is an attribute collection point that collects and forwards
 1803 attributes to a Decision Point. The SOA RA references this collection point as an Attribute Information
 1804 Point. Illustration 4 depicts handling resource, environment, and participant attributes.



1805
 1806 *Figure 31 Policies/Contracts and Attributes*

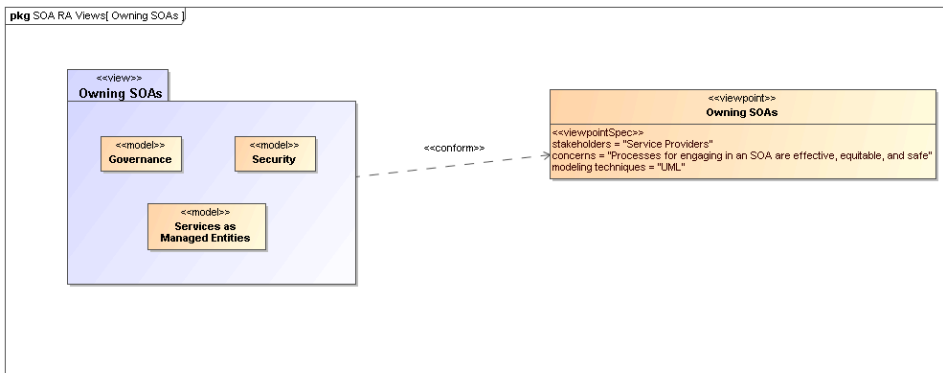
1807
1808
1809
1810
1811
1812
1813
1814
1815
1816

5 Owing Service Oriented Architectures View

*In the absence of policy-based governance,
organizations will operate as unruly collection of
factions that pull in opposing directions.*
Paul A. Strassmann

The Owing Service Oriented Architectures view focuses on what is involved in owning a SOA-based system. [NEED A BIT MORE DESCRIPTIVE TEXT HERE]

The key models in this view are the SOA Governance Model, Security Model, and Services as Managed Entities Model.



1817
1818

Figure 32 Model elements described in the Owing Service Oriented Architectures view

1819

5.1 Governance of Service Oriented Architectures

1820 Given the importance that people attach to the outcomes of their actions, including those mediated by an
1821 SOA, it becomes important to be able to manage the relationship that stakeholders have with their SOA.
1822 Governance is fundamentally about how decisions that are pertinent to the adoption, use and evolution of
1823 an SOA are arrived at and who has the decision rights to make such decisions. As such, we view the
1824 enactment of decisions as being primarily a management concern rather than a governance concern.

1825

5.1.1 Why is explicit Governance important to SOA?

1826 Just as anarchy is also a form of government, so an SOA without an explicit governance structure is also
1827 subject to governance. An explicit model for managing the relationship promotes SOA as an equitable
1828 platform for people to conduct their work.

1829 One of the hallmarks of SOA as compared to other paradigms for distributed computing is the
1830 acknowledgment of the importance of ownership boundaries. Ownership, and issues around it, is one of
1831 the primary topics for governance.

1832 No technology platform is static; likewise, an SOA is also subject to change and evolution. Managing that
1833 evolution, establishing strategies for change, resolving disputes that arise, and ensuring that the SOA
1834 continues to fulfill the goals of the business are all reasons why governance is important to SOA.

1859 need for very little. Over time, best practices will evolve, likely resulting in some consensus on a sensible
1860 minimum and, except in extreme cases where it is demonstrated to be necessary, a loosening of strict
1861 governance toward the best practice mean.

1862 <I was looking over Bob Ellinger's paper on Governance and would be interested to find out how the
1863 organizational structure described there maps to one of the archetypes above.. Seemed like the closest fit
1864 was the IT Duopoly. My only concern was that that we do need to factor in that all orgs involved in a SOA
1865 implementation may not fit that mold>

1866 Ownership boundaries

1867 * Dispute resolution capability

1868 * Ongoing management of environment

1869 * Dealing with issues that go bump in the night

1870 * Providing guidance on the "right" way of doing things

1871 * Who makes the decisions?

1872 * Guidelines vs. Best Practices vs. Laws

1873 * How do you handle exceptions to policies/guidelines?

1874 * What is the feedback loop via which current exceptions may end up becoming future policy/guideline?

1875 * Who tests to make sure that policies are being conformed to?

1876 * Distinguishing between manual vs. automated processes

1877 * How are the contracts between providers and consumers managed?

1878 * More...

1879 **5.1.3 Inputs to the decision process**

1880 * Policies and Contracts

1881 * Best practices

1882 * Architectural principles

1883 * Government regulations

1884 * Laws

1885 * Organizational rules

1886 **5.1.4 Implementing SOA Governance**

1887 Establish processes used by stakeholders

1888 Establish specific groups, committees and boards with responsibilities for different aspects of governance

1889 Setting standards

1890 Courts, juries and executioners

1891 Processes for making decisions

1892 Establishing a constitution

1893 The realization of what is considered "right" for participants in a SOA will be as varied as the participants
1894 themselves and will be based on their shared expectations as well as factors such as architectural
1895 principles, best practices, government regulations, laws, organizational rules. At a high level, it requires
1896 the stakeholders to:

- 1897 • Define stakeholder goals and strategies
- 1898 • Create the organizational structure with the appropriate the decision rights (The Consitution)
- 1899 • Formulate polices and decision making processes that are appropriate to the domain
- 1900 • Define the standards to be applied to the domain
- 1901 • Define the metrics that need to be collected to ensure that policy enforcement

- 1902 • Put mechanisms into place that provides for the enforcement of policies and the ability to collect
- 1903 metrics
- 1904 • Implement feedback and adjudication mechanisms that can adjust the existing policies as needed
- 1905 • Execute and refine on an ongoing basis

1906 **5.1.5 Governance Bone Yard**

1907 <I am not sure if we have adequately emphasized the incentive/dis-incentive aspect of governance in this>

1908 <One of the items that has come up on various conversations is the concept of a Center of Excellence

1909 when it comes to SOA. I am leery of associating it with the "Board" as that association has a tendency to

1910 imbue it with a significant amount of authority. In certain circumstances it simply may not have any but

1911 may simply be a place for that has influence (based on expertise etc.) and not authority. Not sure if this is

1912 something that should be addressed under the governance umbrella>

- 1913 * Governance must address the entire lifecycle of services
- 1914 * Creation
- 1915 * Testing
- 1916 * Provisioning
- 1917 * Utilization
- 1918 * Operation
- 1919 * Changes/Enhancements
- 1920 * Versioning
- 1921 * Retirement
- 1922 * What IT infrastructure is needed to support SOA Governance?
- 1923 * Repository
- 1924 * Policy Management
- 1925 * Metadata Management
- 1926 * Contract Management
- 1927 * Service Management Systems
- 1928 * Service Mediation Systems

1929 **5.1.6 Ken's notes on Governance**

1930 What does it mean to have governance across ownership boundaries? One aspect is management, and

1931 per an earlier email suggestion, a management section has been added to the wiki draft to begin to reflect

1932 this separation. In general, governance reflects what some authority wants to happen, e.g. policies, and

1933 management provides the details and mechanisms by which policy becomes reality. Much of the

1934 management-focused material in the current governance write-up will likely migrate to the management

1935 section, including the appropriate parts dealing with life cycle considerations.

1936

1937 In general, governance should be a function of what one wants to accomplish, and thus while SOA should

1938 leverage existing structures and best practices, it should not adopt approaches developed for single

1939 systems in a single ownership domain if these would significantly inhibit or even preclude the benefits we

1940 expect from SOA. Governance for SOA, both development and enforcement, is likely to parallel

1941 governance for traditional commerce. This leads to the following conclusions:

1942

1943 1. There will be a range of governance depending on the perceived needs of the participants. In a free

1944 market, a dominant mechanism is the satisfaction of the consumer, i.e. if consumers do not find sufficient

1945 value in an offering, the product is not used and will either be modified to better serve its intended

1946 audience or it will disappear. One can see this with numerous consumer products and with shareware on

1947 the Web. There is little if any governance, and this will likely serve similar situations for SOA where
1948 experience with fitness for use is the dominant governance mechanism.

1949

1950 Even with the market, there are situations where market feedback is not considered sufficient in terms of
1951 speed, precision, or need to mitigate effects. This is seen where there are health and safety
1952 considerations, such as advance approval of new drugs. Alternately, there are intermediate situations, for
1953 example the automobile industry, where the market is the dominant governance mechanism but a third
1954 party, i.e. some level of government, intervenes where health or safety is an issue. Further discussion on
1955 this is included under item 3 below.

1956

1957 The conclusion then is there is not a one-size-fits-all governance but a need to understand the types of
1958 things governance will be called on to do in the context of the goals of SOA. It is likely that some
1959 communities will initially desire and require very stringent governance policies and procedures while other
1960 will see need for very little. Over time, best practices will evolve, likely resulting in some consensus on a
1961 sensible minimum and, except in extreme cases where it is demonstrated to be necessary, a loosening of
1962 strict governance toward the best practice mean.

1963

1964 2. Whatever level of governance is chosen, it must have effective enforcement, including collection of and
1965 access to information needed for enforcement. At a basic level, this requires a relatively free flow of
1966 information on consumer experience so prospective consumers can determine whether a given resource
1967 available through a SOA implementation provides described functionality and robustness consistent with
1968 consumer expectations. Again, the need for enforcement depends on the importance, e.g. life criticality,
1969 of the resource. If a resource is free of charge and is generally available to provide some useful but non-
1970 critical function, its mere availability may be sufficient and little if any governance or associated
1971 enforcement is necessary. For something where there is advance arrangement for the service, e.g. a
1972 subscription service, information is likely needed to document availability and form the basis for penalties
1973 as prescribed by applicable enforceable policy. If the service is important, reporting may be more critical
1974 than penalties because there will be an imperative for understanding and fixing any problem that occurs.

1975

1976 The conclusion then is enforcement is likely dependent on available information (metrics?) and the
1977 enforcement mechanism should be consistent with the level of governance perceived as needed. Some
1978 aspects of enforcement fall under management.

1979

1980 3. Regulatory governance likely to evolve to reflect perceived needs of stakeholders, including non-
1981 participatory stakeholders and regulatory governance of the Commons.

1982

1983 Governance as grounds for mediation of differences between participants.

1984 Governance to codify consensus behavior between participants.

1985 Governance to protect participants.

1986 Governance to protect non-participants from side effects.

1987 Governance to protect the Commons.

1988

1989 SOA provides an interesting example where we are trying to prescribe governance for something that in
1990 many cases does not yet exist. We are trying to use past experience to deal with anticipated needs and
1991 requirements. From a historical context, governance concepts evolved first to provide grounds for
1992 mediation of differences between participants, and later to codify consensus behavior between
1993 participants and protect the participants from damaging behavior of one of the parties. The enforcement
1994 mechanism could be any agreed upon third party who had been given authority over the participants or
1995 later some governmental body that also generated the policy to be enforced. In any case, the
1996 enforcement mechanism was answerable to that body's stakeholders. When government is the
1997 enforcement mechanism, the stakeholders include not only the immediate participants but other non-

1998 participatory stakeholders who may be affected by side effects of the primary interaction. The RM gives
1999 examples of real world effects that go beyond the immediate public actions, such as a change in credit
2000 rating after getting a loan for a large purchase. In a society, perceived effects on the well being of the
2001 general population (the non-participatory stakeholders) often lead to additional policies and enforcement,
2002 such as environmental standards. The intent of such governance is to protect and regulate behavior that
2003 affects the Commons, the resources under the ownership or protection of society as a whole.

2004

2005 The conclusion then is governance by third parties, whether through government or other agreed upon
2006 organization, is likely to develop and have significant effect on the overall governance of SOA.

2007 **5.2 Security Model**

2008 Providing for security in the context of Service Oriented Architecture is not especially different to other
2009 contexts. However, the fact that SOA embraces crossing ownership boundaries and the fact that we aim
2010 to explicitly relate the IT architecture with the human architecture (see Business via Services) makes it
2011 possible to give a more complete accounting of security.

2012 Any comprehensive security solution must take into account the people that are using, maintaining and
2013 managing the SOA. Furthermore, the relationships between them must also be incorporated: any security
2014 assertions that may be associated with particular interactions originate in the people that are behind the
2015 interaction.

2016 Concepts such as constitutions, roles, and authority within social structures play an important part in the
2017 establishment of ownership and trust boundaries within and between social structures. These in turn
2018 provide a sound rationale for determining what security policies should be applied and how.

2019 Security is associated with a *threat model* and a *security response model*. The threat model identifies the
2020 kinds of threats that are the concern of security specialists, and the response model is the basis for
2021 responding to those threats to provide assurance in the safety and integrity of the system.

2022 **5.2.1 Security Concepts**

2023 Security is one aspect of assurance – the confidence in the integrity and reliability of the system. In
2024 particular, security focuses on those aspects of assurance that involve the malign intent of other people.
2025 We can characterize security in terms of five key concepts: confidentiality, integrity, authentication,
2026 authorization, and non-repudiation.

2027 **Confidentiality**

2028 Confidentiality concerns the protection of privacy of participants in their interactions.

2029 Confidentiality refers to the assurance that unauthorized entities are not able to read messages or
2030 documents that are interchanged.

2031 Note that confidentiality has degrees: in a completely confidential exchange, third parties would
2032 not even be aware that a confidential exchange has occurred. In a partially confidential exchange,
2033 the identities of the participants may be known but the content of the exchange obscured.

2034 **Integrity**

2035 Integrity concerns the protection of information that is exchanged from corruption. Integrity refers
2036 to the assurance that information that has been exchanged has not been tampered with.

2037 Integrity is different from confidentiality in that messages that are sent from one participant to
2038 another may be obscured to a third party, but the third party may still be able to introduce his own
2039 content into the exchange without the knowledge of the participants.

2040 **Authentication**

2041 Authentication concerns the identity of the participants in an exchange. Authentication refers to
2042 the means by which one participant can be assured of the identity of other participants.

2043 **Authorization**

2044 Authorization concerns the legitimacy of the interaction. Authorization refers to the means by
2045 which one participant may be assured that the information and actions that are exchanged are
2046 valid and may be acted on.

2047 **Non-repudiation**

2048 Non-repudiation concerns the accountability of participants. Non-repudiation refers to the means
2049 by which a participant may not, at a later time, successfully deny having participated in the
2050 interaction or having performed the actions as reported by other participants.

2051 **Availability**

2052 Availability concerns the ability of systems to use and offer the services for which they were
2053 designed. One of the threats against availability is the so-called denial of service attack in which
2054 attackers attempt to prevent legitimate access to the system.

2055 We differentiate here between general availability – which includes aspects such as systems
2056 reliability – and availability as a security concept where we need to respond to active threats to
2057 the system.

2058 Note that these security concepts are never absolute: it is not possible to guarantee 100% confidentiality,
2059 non-repudiation, etc. *However, a well designed and implemented security response model can ensure*
2060 *that the costs of abrogating security are greater than the potential benefits of having done so.* For
2061 example, using a well-designed cipher to encrypt messages may make the cost of breaking
2062 communications so great and so lengthy that the information obtained is valueless.

2063 While confidentiality and integrity can be viewed as primarily the concerns of the direct participants in an
2064 interaction, authentication and authorization and non-repudiation imply the participants are acting within a
2065 broader social structure.

2066 **5.2.2 Threat Model**

2067 There are a number of ways in which an attacker may attempt to compromise the security of a system.
2068 The two primary sources of attack are third parties attempting to subvert interactions between legitimate
2069 participants and an entity that is participating but attempting to subvert its partner(s). The latter is
2070 particularly important in an SOA where there may be multiple ownership boundaries and trust boundaries.

2071 **Message alteration**

2072 If an attacker is able to modify the content (or even the order) of messages that are exchanged
2073 without the legitimate participants being aware of it then the attacker has successfully
2074 compromised the security of the system. In effect, the participants may unwittingly serve the
2075 needs of the attacker rather than their own.

2076 An attacker may not need to completely replace a message with his own to achieve his objective:
2077 replacing the identity of the beneficiary of a transaction may be enough.

2078 **Message interception**

2079 If an attacker is able to intercept and understand messages exchanged between participants,
2080 then the attacker may be able to gain advantage. This is probably the most commonly understood
2081 security threat.

2082 **Man in the middle**

2083 In a man in the middle attack, the legitimate participants believe that they are interacting with
2084 each other; but are in fact interacting with the attacker. The attacker attempts to convince each
2085 participant that he is their correspondent; whereas in fact he is not.

2086 In a successful man-in-the-middle attack, legitimate participants will often not have a true
2087 understanding of the state of the other participants. The attacker can use this to subvert the
2088 intentions of the participants.

2089 **Spoofing**

2090 In a spoofing attack, the attacker convinces a participant that he is really someone else –
2091 someone that the participant would normally trust.

2092 **Denial of service attack**

2093 In a denial of service attack, the attacker attempts to prevent legitimate users from making use of
2094 the service. A DoS attack is easy to mount and can cause considerable harm: by preventing
2095 legitimate interactions, or by slowing them down enough, the attacker may be able to
2096 simultaneously prevent legitimate access to a service and to attack the service by another
2097 means.

2098 A variation of the DoS attack is the **Distributed Denial of Service** attack. In a DDoS attack the
2099 attacker uses multiple agents to attack the target. In some circumstances this can be
2100 extremely difficult to counteract effectively.

2101 One of the features of a DoS attack is that it does not require valid interactions to be effective:
2102 responding to invalid messages also takes resources and that may be sufficient to cripple the
2103 target.

2104 **Replay attack**

2105 In a replay attack, the attacker captures the message traffic during a legitimate interaction and
2106 then replays part of it the target. The target is persuaded that a similar transaction to the previous
2107 one is being repeated and it will respond as though it were a legitimate interaction.

2108 A replay attack may not require that the attacker understand any of the individual
2109 communications; the attacker may have different objectives (for example attempting to predict
2110 how the target would react to a particular request).

2111 **Repudiation**

2112 In a repudiation attack, the attacker completes a normal transaction and then later attempts to
2113 deny that the transaction occurred. For example, a customer may use a service to buy a book
2114 using a credit card; then, when the book is delivered, refuse to pay the credit card bill claiming
2115 that *someone else* must have ordered the book.

2116 **5.2.3 Mitigation Model**

2117 Responding to security threats in a coherent and consistent way is key to mitigating those threats in a
2118 cost-effective way. We identify a few elements of the architecture that may form the basis of a
2119 comprehensive security response model for SOA.

2120 We structure the security mitigation model into a number of different elements: an association between
2121 policies and security elements, mechanisms intended to support privacy and integrity, mechanisms
2122 intended to support authority, mechanisms intended to support obligation-style policies and mechanisms
2123 intended to resist DoS attacks.

2124 **5.2.3.1 Policies for security**

2125 Mechanisms are not the same as solutions; a combination of security mechanisms and their control via
2126 explicit policies can form the basis of a solution. Elsewhere in the architecture policies are used to
2127 express routing constraints, business constraints and information processing constraints. Security policies
2128 are used to marry stakeholders' choices with mechanisms to enforce security.

2129 Security policies are not equivalent to security. However, they are very important as the expression of
2130 choices that can be used by security mechanisms to enforce security.

2131 The role of a machine readable security policy is to permit, on the one hand, stakeholders to express their
2132 choices; and, on the other hand, to act as instructions for security enforcement mechanisms.

2133 **5.2.3.2 Privacy Enforcement**

2134 The most efficient mechanism to enforce privacy is the encryption of information. Encryption is particularly
2135 important when messages must cross trust boundaries; especially the Internet. Note that encryption need
2136 not be limited to the content of messages: it is possible to obscure even the existence of messages
2137 themselves through encryption and 'white noise' generation in the communications channel.

2138 The specifics of encryption are beyond the scope of this architecture. However, we are concerned about
2139 how the connection between privacy-related policies and their enforcement is made. In Section XX, we
2140 show how policies in general are enforced using a combination of Policy Decision Points (PDP) and
2141 Policy Enforcement Points (PEP).

2142 A PEP for enforcing privacy may take the form of an automatic function to encrypt messages as they
2143 leave a trust boundary; or perhaps simply ensuring that such messages are suitably encrypted. If it is
2144 important to completely disguise message traffic then some central function to generate encrypted 'white
2145 noise' when no messages are being transmitted does typically require a centralized facility.

2146 Any policies relating to the level of encryption being used would then apply to these centralized
2147 messaging functions.

2148 **5.2.3.3 Integrity**

2149 To protect against tampering, and to allow the receiver of a message to authenticate the sender,
2150 messages may be accompanied by a digital signature. Digital signatures provide a means to detect if
2151 signed data has been altered.

2152 A digital signature can be generated with the use of a private key that is associated with a public key and
2153 a digital certificate. The private key of some entity in the system is used to create a digital signature for
2154 some set of data. Other entities in the system can check the integrity of the signed data set via signature
2155 verification algorithms. Any changes to the data that was signed will cause signature verification to fail,
2156 which indicates that integrity of the data set has been compromised.

2157 A party verifying a digital signature must have access to the public key that corresponds to the private key
2158 used to generate the signature. A digital certificate contains the public key of the owner, and is itself
2159 protected by a digital signature created using the private key of the issuing Certificate Authority (CA).

2160 **5.2.3.4 Message Replay Protection**

2161 To protect against replay attacks, messages may contain information that can be used to detect replayed
2162 messages. The simplest requirement to prevent replay attacks is that each message that is ever sent is
2163 unique. For example, a message may contain a message ID, a timestamp, the intended destination.

2164 By caching message IDs, and comparing each new message with the cache, it becomes possible to
2165 verify whether a given message has been received before (and therefore should be discarded).

2166 The timestamp may be included in the message to help check for message freshness. Messages that
2167 arrive after their message ID could have been cleared (after receiving the same message some time
2168 previously) may also have been replayed. A common means for representing timestamps is a useful part
2169 of an interoperable replay detection mechanism.

2170 The destination information is used to determine if the message was misdirected or replayed. If the
2171 replayed message is sent to a different endpoint than the destination of the original message, the replay
2172 could go undetected if the message does not contain information about the intended destination.

2173 In the case of messages that are replies to prior messages, it is also possible to include seed information
2174 in the prior messages that is randomly and uniquely generated for each message that is sent out. A
2175 replay attack can then be detected if the reply does not embed the random number that corresponds to
2176 the original message.

2177 **5.2.3.5 Trust, Social Structures and Identity**

2178 Trust is an assertion as to the behavior of participants in relation to each other. In terms of security
2179 assurance, trust often refers to the confidence that target systems may have as to the identity and validity
2180 of a participant as they interact with the system. However, in general, trust is a far larger topic.

2181 There are various kinds of trust domain: at the infrastructure level, a trust domain may refer to the
2182 networking equipment that is under the control of the owners of a SOA and is used to propagate
2183 communication. At an application level, a trust domain may refer to a social structure (see Section 3.4)
2184 within which members have previously established a certain degree of trust.

2185 Generally, there are special procedures necessary to communicate across trust domains: for example,
2186 participants may need to present credentials to participate in a trust domain. Once authenticated,
2187 credentials would typically not be needed to continue within that trust domain.

2188 The connection between policies and trust domains is similar to that for privacy: when a participant
2189 wishes to perform an action that requires access to a trust domain, depending on the policies that are in
2190 place, he/she must provide suitable credentials to the PEP before continuing the interaction.

2191 In the context of a SOA that is used by many people, there may not be a single repository for information
2192 that can justify trust. Often different aspects of trust are managed by different entities. For example, a
2193 corporate directory might be used to verify the employment of an individual, whereas a bank would be
2194 used to verify their credit worthiness and a government agency used to verify their residency.

2195 Together, the various entities that provide corroboration of an individual's identity and trustworthiness
2196 form a *web of trust*. Webs of trust need not be functionally organized: third parties who are known to both
2197 may also be used to facilitate trust. Webs of trust have some promise in permitting the efficient scaling of
2198 large SOA-based systems. Of course, a complex and long *trust chain* is likely to be more fragile and less
2199 trustworthy (sic) than a simple one.

2200 **5.2.3.6 Authority, Social Structures and Authorization**

2201 The authority held by that participant often determines the validity of actions that a participant engages in.
2202 As noted in Section 3.4.2, that authority is always in relation to a particular social structure.

2203 In the context of SOA, the meeting point of action, policy, and validity is often at the service itself. When a
2204 participant attempts an action against another participant, the latter may require that the former is
2205 properly authorized. (For example, when opening a bank account, it is often required that the customer
2206 has appropriate residency status in the bank's country.)

2207 A PEP for enforcing authorization policies would normally be attached to the service that offers the
2208 capability. Note that for other reasons, it may not be advisable to *embed* such a PEP within the service
2209 capability itself.

2210 The core task of any authorization PEP is to verify that a requested action is valid for the participant;
2211 given the identified role that the participant has within the social structure that validates the action.

2212 **5.2.3.7 Auditing and logging**

2213 A non-repudiation attack involves a participant denying that it authorized a previous interaction. An
2214 effective strategy for responding to such a denial is to maintain careful and complete audits of
2215 interactions. The more detailed and comprehensive an audit trail is, the less likely it is that a false
2216 repudiation would be successful.

2217 Unlike many of the security responses discussed here, it is likely that the scope for automation in
2218 rejecting a repudiation attempt is limited to careful logging.

2219 **5.2.3.8 Graduated engagement**

2220 The key to managing and responding to DoS attacks is to be careful in the use of resources when
2221 responding to interaction. Put simply, a system has a choice to respond to a communication or to ignore
2222 it. In order to avoid vulnerability to DoS attacks a service provider should not commit to any interaction to
2223 a significantly greater extent than the service consumers.

2224 **5.3 Services as Managed Entities Model**

2225 Managing systems that may be used across ownership boundaries raises issues that may not normally
2226 be present in managing a system within a single ownership domain. For example, how is the
2227 management of a service to be arranged when the owner of the service, the provider of the service, the

2228 host of the service and access mediators to the service may all belong to different stakeholders.
2229 Furthermore, how may a service customer communicate his or her requirements to the service provider
2230 so that they are satisfied in a timely manner.

2231 In fact, managing a service has quite a few similarities to using a service: suggesting that we can use the
2232 service oriented model to manage SOA systems as well as provide them. A management service would
2233 be distinguished from a non-management service more by the nature of the capabilities involved (i.e.,
2234 capabilities that relate to managing services) than by any intrinsic difference.

2235 In this model we show how the SOA framework may apply to managing services as well as using and
2236 offering them. There are, of course, some special considerations that apply to service management which
2237 we bring out: namely that we will be managing the life-cycle of services, managing any service level
2238 attributes, managing dependencies between services and so on.

2239 As in other views of SOA, policies and contracts also are important in managing systems. It has long
2240 been known that a systematic management policy framework may be very helpful in managing distributed
2241 systems, and SOA is no exception.

2242 Critical SOA Management Topics:

2243 Distributed, independently generated parts that need to interact.

2244 Notions of Mgt:

2245 - management infrastructure (manage things/resources) *current emphasis (more technical)

2246 - manage participant relationships (and other services used to do this) (manage actions/uses)

2247 (more social)

2248 Service lifecycle (see telemanagement group as a reference).

2249 Hook into governance (day-to-day activities of gov.)

2250 Governance proposes, management disposes

2251 Service Level Agreements

2252 Policies and contracts

2253 Policy generation Is a governance function, policy implementation is a management function

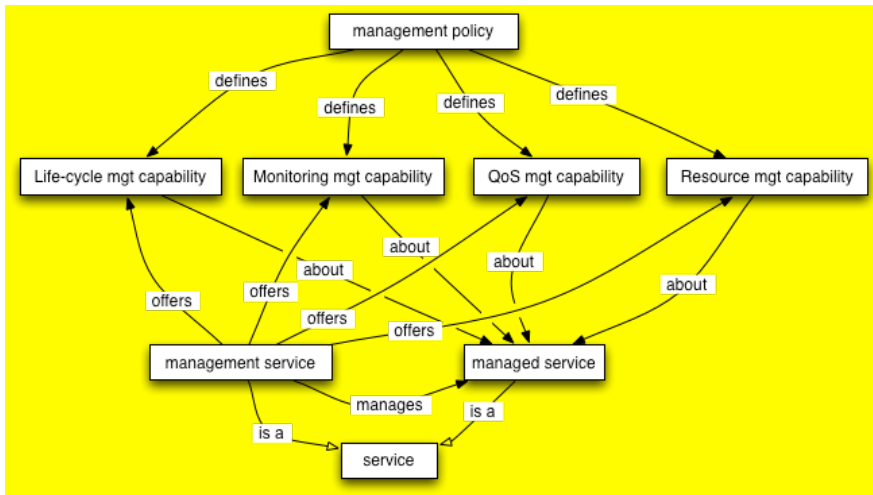
2254

2255 Services may be managed. In a service oriented management scheme we may use services to manage
2256 other services.

2257 A managed service has a number of its aspects that are managed. Furthermore, a management service
2258 may offer management of certain aspects of other services. Collectively, these aspects are known as
2259 manageability capabilities. (Ed. Note: something of a mouthful).

2260 In a deployed system, it may well be that different aspects of the management of a given service are
2261 managed by different management services. For example, the life-cycle management of services often
2262 involves managing dependencies between services and resource requirements. Managing quality of
2263 service is often very specific to the service itself; for example, quality of service attributes for a video
2264 streaming service are quite different to those for a banking system.

2265



2266
2267 *Figure 34 Services and Managed Services*

2268 **5.3.1 Definitions**

2269 **Managed service**

2270 A service that may be managed

2271 **Management service**

2272 A service that manages other services

2273 **Management Policy**

2274 A policy whose topic is a management topic

2275 **Manageability capability**

2276 An aspect of a service that may be managed by a management service

2277 **Lifecycle Manageability**

2278 A manageability capability that permits a management service to control the lifecycle of a
2279 managed service. Lifecycle management actions include starting services, stopping them,
2280 pausing them and so on.

2281 **Quality of Service Manageability**

2282 A manageability capability that permits a management service to control quality of service
2283 aspects of a management service.

2284 **Monitoring Manageability**

2285 A manageability capability

2286

2287 SOA is by definition a "distributed" paradigm; therefore, from an information technology (IT) perspective,
2288 a managed distributed system architecture is needed to fully realize the potential of SOA. Distributed
2289 capabilities facilitated by the SOA paradigm may be under different ownership domains. This suggests
2290 that there is no theoretical limit as to how widely dispersed—geographically or otherwise—participants in
2291 a SOA environment can be so long as there exists a means for service participants to communicate. The
2292 prospect of having to support a highly distributed system architecture poses significant challenges from a
2293 systems and network management point of view, and requires the introduction of a specialized form of

2294 management known as "service management". The bottom line is that a SOA must be managed to be
2295 effective.

2296 **Systems management**

2297 Systems management refers to enterprise-wide maintenance and administration of distributed
2298 computer systems.

2299 **Network management**

2300 Network management refers to the maintenance and administration of large-scale networks such
2301 as computer networks and telecommunication networks. Systems and network management
2302 execute a set of functions required for controlling, planning, deploying, coordinating, and
2303 monitoring the distributed computer systems and the resources of a network.

2304 **Service management**

2305 Service management, which is the subject of this reference architecture view, refers to the
2306 management and administration of service-based resources through a set of activities and
2307 capabilities that continuously monitor, control, coordinate, and report on the qualities and usage
2308 of these resources. Examples of service qualities include health qualities, or common Qualities of
2309 Service (QoS) attributes such as availability and performance, and accessibility. Examples of
2310 service usage that may be monitored or controlled include frequency, duration, scope, functional
2311 extent, and access authorization. Ultimately, service management is about insuring that
2312 acceptable levels of service quality meet the needs of the service consumer.

2313 **5.3.2 Management Capabilities**

2314 Historically, systems management capabilities have been organized by the following functional groups
2315 known as "FCAPS" functions (based on the ITU-T Rec. X.700 | ISO/IEC 7498-4:1989(E) standard):

- 2316 • Fault Management
- 2317 • Configuration Management
- 2318 • Accounting Management
- 2319 • Performance Management
- 2320 • Security Management

2321

2322 From a service management perspective, each of these functional groups can be leveraged and defined
2323 for purposes of this SOA reference architecture as follows (in concert with ITU-T Rec. X.700 | ISO/IEC
2324 7498-4:1989(E)):

2325 **Fault Management**

2326 Encompasses fault detection, isolation and the correction of abnormal operation of the SOA
2327 environment. Faults cause SOA distributed systems to fail to meet their operational objectives
2328 and they may be persistent or transient. Faults manifest themselves as particular events (e.g.,
2329 errors) in the operation of a distributed system. Error detection provides capabilities to recognize
2330 faults. Fault management includes functions to a) maintain and examine error logs, b) accept
2331 and act upon error detection notifications, c) trace and identify faults, d) carry out sequences of
2332 diagnostic tests, and e) correct faults. For purposes of this reference architecture, monitoring
2333 functions such as service status and alerting are included in this functional group.

2334

2335 **Accounting Management**

2336 Enables charges to be established for the use of resources in the SOA environment, and for
2337 costs to be identified for the use of those resources. Accounting management includes functions
2338 to a) inform service consumers of costs incurred or resources consumed, b) enable accounting
2339 limits to be set and tariff schedules to be associated with the use of resources, and c) enable
2340 costs to be combined where multiple resources are invoked to achieve a given objective

2341 (resulting in a real-world effect). For purposes of this reference architecture, related accounting
2342 functions such as metering and billing fall into this category.

2343 **Configuration Management**

2344 Identifies, exercises control over, collects data from and provides data to SOA distributed
2345 systems for the purpose of preparing for, initializing, starting, providing for the continuous
2346 operation of, and terminating services. Configuration management includes functions to a) set
2347 the parameters that control the routine operation of the SOA distributed system, b) associate
2348 names with managed resources and sets of managed resources, c) initialize and close down
2349 managed resources, d) collect information on demand about the current condition of the SOA
2350 distributed system, e) obtain announcements of significant changes in the condition of the SOA
2351 distributed system, and f) change the configuration of the SOA distributed system. For purposes
2352 of this reference architecture, related configuration management functions of service versioning
2353 and service provisioning (i.e., supplying of services) is included in this functional category.

2354 **Performance Management**

2355 Enables the behavior of resources in the SOA environment and the effectiveness of service-
2356 oriented activities to be evaluated. Performance management includes functions to a) gather
2357 statistical information, b) maintain and examine logs of system state histories, c) determine
2358 system performance under natural and artificial conditions, and d) alter system modes of
2359 operation for the purpose of conducting performance management activities. Measurements
2360 gathered as part of performance management are used to compare against service level
2361 agreements (SLAs).

2362 **Security Management**

2363 Support the application of security policies by means of functions which include a) the creation,
2364 deletion and control of security services and mechanisms, b) the distribution of security-related
2365 information, and c) the reporting of security-relevant events. A more detailed treatment on the
2366 topic of security is provided in the Security View of this SOA reference architecture.

2367 **5.3.3 Management Contracts and Policies**

2368

2369 **5.3.3.1 SLA management (ex of contract), etc.**

2370

2371 QoS Attributes:

2372

2373 Quality of Service capabilities will enable graceful degradation, fault tolerance, high reliability, and
2374 bounded deterministic behavior. Typical QoS attributes are:

2375

2376 **Availability**

2377 Probability for service availability including such factors as MTBF (H/W and S/W) and MTTR

2378 **Accessibility**

2379 Probability of successful service instantiation when required

2380 **Scalability**

2381 Probability to successfully serve requests independent of load

2382 **Integrity**

2383 Measurement of interaction correctness with respect to the source versus probabilistic
2384 requirement

2385 **Performance**
2386 Measurement of round trip service request throughput and latency versus requirement

2387 **Reliability**
2388 The probability of being able to maintain a service at specified service quality

2389 **Regulatory**
2390 The probability of conforming to rules, standards, service level agreements

2391

2392 **5.3.3.2 Policies**

2393 "Although provision of management capabilities enables a service to become manageable, the extent and
2394 degree of permissible management are defined in management policies that are associated with the
2395 services. Management policies are used to define the obligations for, and permissions to, managing the
2396 service." [WSA]

2397

2398 **Will come back to this..**

2399

2400 **Relate to policies, i.e., "policies are also intended as a vehicle to express SLAs."**

2401

2402 **5.3.4 Manageability & Instrumentation**

2403

2404 **5.3.5 Management Infrastructure**

2405

2406 Elements of a basic service management infrastructure should include the following characteristics:

2407

- 2408 • Integrate with existing security services
- 2409 • Monitoring
- 2410 • Heartbeat and Ping
- 2411 • Alerting
- 2412 • Pause/Restore/Restart Service Access
- 2413 • Logging, Auditing, Non-Repudiation
- 2414 • Runtime Version Management
- 2415 • Complement other infrastructure services (discovery, messaging, mediation)

2416

2417 * Message Routing and Redirection

2418 * Failover

2419 * Load-balancing

2420

2421 * QoS, Management of Service Level Objects and Agreements

2422 * Availability

2423 * Response Time

2424 * Throughput

2425
2426 • Fault and Exception Management
2427

2428 **5.3.6 Service Life-cycle**

2429 **5.3.7 Service Provisioning**

2430 Requirements on a management system should be to manage the services and not the infrastructure.

2431 6 References

2432 6.1 Normative References

- 2433 [ANSI/IEEE Std 1471-2000] *IEEE Recommended Practice for Architectural Description of Software-*
2434 *Intensive Systems*, American National Standards Institute/Institute for Electrical
2435 and Electronic Engineers, September 21, 2000.
- 2436 [RFC2119] S. Bradner, *Key words for use in RFCs to Indicate Requirement Levels*,
2437 <http://www.ietf.org/rfc/rfc2119.txt>, IETF RFC 2119, March 1997.
- 2438 [SOA-RM] C. M. MacKenzie, K. Laskey, F. McCabe, P. F. Brown, and R. Metz, (editors),
2439 "Reference Model for Service Oriented Architecture 1.0, OASIS Open, October
2440 12, 2006.
- 2441 [WSA] David Booth, et al., "Web Services Architecture", W3C Working Group Note,
2442 World Wide Web Consortium (W3C) (Massachusetts Institute of Technology,
2443 European Research Consortium for Informatics and Mathematics, Keio
2444 University), February, 2004

2445 6.2 Non-Normative References

- 2446 [COX] D. E. Cox and H. Kreger, "Management of the service-oriented architecture life
2447 cycle," *"IBM Systems Journal" "44"*, No. 4, 709-726, 2005
- 2448 [DEEPAK] Deepak Kakadia, et al., "Enterprise Management Systems Part I: Architectures
2449 and Standards," *"Sun BluePrints™ Online"*, Sun Microsystems, Inc., Santa
2450 Clara, CA, April, 2002.
- 2451 [ITU-T Rec. X.700 | ISO/IEC 10746-3:1996(E)] *Information processing systems—Open Systems*
2452 *Interconnection—Basic Reference Model—Part 4: Management Framework*,
2453 International Telecommunication Union, International Organization for
2454 Standardization and International Electrotechnical Commission, Geneva,
2455 Switzerland, 1989.
- 2456 [OECD] Organization for Economic Cooperation and Development, Directorate for
2457 Financial, Fiscal and Enterprise Affairs, *OECD Principles of Corporate*
2458 *Governance*, SG/CG(99) 5 and 219, April 1999.
- 2459 [TOGAF] *The Open Group Architecture Framework (TOGAF) 8.1 Enterprise Edition*,
2460 *The Open Group, Doc Number: G051, December 19, 2003.*
- 2461 [WEILL] Harvard Business School Press, *IT Governance: How Top Performers Manage*
2462 *IT Decision Rights for Superior Results*, Peter Weill and Jeanne W. Ross, 2004
2463

2464

A. Acknowledgements

2465 The following individuals have participated in the creation of this specification and are gratefully
2466 acknowledged:

2467 **Participants:**

2468 [Participant Name, Affiliation | Individual Member]

2469 [Participant Name, Affiliation | Individual Member]

2470

2471

B. Critical Factors Analysis

2472
2473
2474

A critical factors analysis (CFA) is an analysis of the key properties of a project. A CFA is analyzed in terms of the goals of the project, the critical factors that will lead to its success and the measurable requirements of the project implementation that support the goals of the project.

2475

B.1 Goals

2476
2477
2478

A goal is an overall target that you are trying to reach with the project. Typically, goals are hard to measure by themselves. Goals are often directed at the potential consumer of the product rather than the technology developer.

2479

B.1.1 Critical Success Factors

2480
2481
2482

A critical success factor (CSF) is a property, sub-goal that directly supports a goal and there is strong belief that without it the goal is unattainable. CSFs themselves are not necessarily measurable in themselves.

2483

B.1.2 Requirements

2484
2485
2486
2487

A requirement is a specific measurable property that directly supports a CSF. The key here is measurability: it should be possible to unambiguously determine if a requirement has been met. While goals are typically directed at consumers of the specification, requirements are focused on technical aspects of the specification.

2488

B.1.3 CFA Diagrams

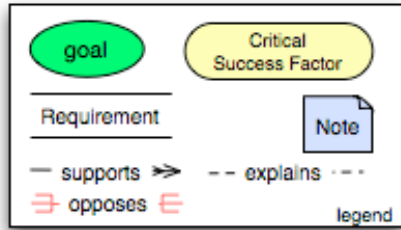
2489
2490
2491

It can often be helpful to illustrate graphically the key concepts and relationships between them. Such diagrams can act as effective indices into the written descriptions of goals etc., but is not intended to replace the text.

2492

The legend:

2493



2494

2495

2496

2497

2498

illustrates the key elements of the graphical notation. Goals are written in round ovals, critical success factors are written in round-ended rectangles and requirements are written using open-ended rectangles. The arrows show whether a CSF/goal/requirement is supported by another element or opposed by it. This highlights the potential for conflict in requirements.

2499

C. Revision History

2500 [optional; should not be included in OASIS Standards]

2501

Re vis ion	D a t e	Ed ito r	Changes Made
Disp	Dis	Dis	DisplayTe

2502

2503