Reference Architecture Foundation for Service Oriented Architecture Version 1.0

Committee Specification Draft 03 / Public Review Draft 02

06 July 2011

Specification URIs:

This ~~Version~~version:

http://docs.oasis-open.org/soa-rm/soa-ra/v1.0/csprd02/soa-ra-v1.0-csprd02.pdf (Authoritative)
http://docs.oasis-open.org/soa-rm/soa-ra/v1.0/csprd02/soa-ra-v1.0-csprd02.html
http://docs.oasis-open.org/soa-rm/soa-ra/v1.0/csprd02/soa-ra-v1.0-csprd02.doc

Previous ~~Version~~version:

http://docs.oasis-open.org/soa-rm/soa-ra/v1.0/soa-ra-cd-02.pdf~~N/A~~

~~Latest Version:~~

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

Latest ~~Approved Version~~version:

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

Technical Committee:

OASIS Service Oriented Architecture Reference Model TC

Chair~~(s)~~:

~~Francis G. McCabe~~

~~Editor(s):~~

~~Jeff A. Estefan, Jet Propulsion Laboratory,~~ Ken Laskey (klaskey@mitre.org~~,~~), MITRE Corporation~~,~~

Editors:

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

Related work:

This specification is related to:

- OASIS Reference Model for Service Oriented Architecture
- ~~OASIS Reference Model for Service Oriented Architecture~~

Abstract:

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

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

~~This architecture~~The SOA-RAF follows the recommended practice of describing architecture in terms of models, views, and viewpoints, as prescribed in the ANSI/IEEE 1471~~ Std. This Reference Architecture~~-2000 (now ISO/IEC 42010-2007) Standard. The SOA-RAF is ~~principally targeted at~~of value to Enterprise Architects~~; however~~, Business and IT Architects as well as CIOs and other senior executives involved in strategic business and IT planning~~ should also find the architectural views and models described herein to be of value~~.

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

Status:

This document was last revised or approved by the ~~SOA~~OASIS Service Oriented Architecture Reference Model TC on the above date. The level of approval is also listed above. Check the "Latest ~~Version" or "Latest Approved Version~~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~~"Send A Comment"~~" button on the Technical Committee's web page at http://www.oasis-open.org/committees/soa-rm/~~.~~.

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

~~**The non-normative errata page for**~~ **Citation format:**

When referencing this specification ~~is located at~~the following citation format should be used:

[SOA-RAF]

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

 Copyright © OASIS Open 2011. All Rights Reserved.

Notices

Copyright © OASIS~~® 1993–2008.~~ Open 2011. All Rights Reserved.

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

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

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

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

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

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

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

The name "OASIS" is a trademark of OASIS~~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

 Copyright © OASIS Open 2011. All Rights Reserved.

 Copyright © OASIS Open 2011. All Rights Reserved.

 Copyright © OASIS Open 2011. All Rights Reserved.

Table of Figures

 Copyright © OASIS Open 2011. All Rights Reserved.

 Copyright © OASIS Open 2011. All Rights Reserved.

1 Introduction

Service Oriented Architecture (SOA) is an architectural paradigm that has gained significant attention within the information technology (IT) and business communities. The SOA ecosystem described in this document occupies the area between business and IT. It is neither wholly IT nor wholly business, but is of both worlds. Neither business nor IT completely own, govern and manage this SOA ecosystem. Both sets of concerns must be accommodated for the SOA ecosystem to fulfill its purposes.[1]

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

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

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

1.1 Context for Reference Architecture for SOA

1.1.1 What is a Reference Architecture?

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

It is possible to define reference architectures at many levels of detail or abstraction, and for many different purposes. ~~In fact, the reference architecture for one domain may represent a further specialization of another reference architecture, with additional requirements over those for which the more general reference architecture was defined.~~

A reference architecture ~~need~~is not~~ be~~ a concrete architecture; i.e., depending on the requirements being addressed by the reference architecture, it ~~may~~generally will not~~ be necessary to~~ completely specify all the technologies, components and their relationships in sufficient detail to enable direct implementation. ~~Such a concrete architecture may be valuable and necessary to ensure a successful implementation;~~

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

 Copyright © OASIS Open 2011. All Rights Reserved.

~~however, the detail necessary in concrete architectures may force technology choices that are not forced by the requirements per se, but by the technology choices available at the time.~~

1.1.2 What is this Reference Architecture?

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

~~When designing systems that are intended to be used across ownership boundaries over extended periods of time it is necessary to address not only how the system is to be constructed, but also how it integrates with the life of users of the system and what is involved in owning such a system. In effect, we take a total cost of ownership stance on the architecture of SOA-based systems.~~

While requirements are addressed more fully in Section 0, the SOA-RAF makes key assumptions~~ that we make in this Reference Architecture is~~ that SOA-based systems involve:

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

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

Such an environment as ~~a SOA~~described above is an *ecosystem* and, specifically in the context of SOA-based systems, is a **SOA ecosystem**. This concept of an ecosystem perspective of SOA is elaborated further in Section 0~~. Informally, our goal in this Reference Architecture is to show~~.

This SOA-RAF shows how Service Oriented Architecture fits into the life of users and stakeholders~~ in a SOA ecosystem~~, how SOA-based systems may be realized effectively, and what is involved in owning ~~such a SOA-based system. We believe that this approach will serve~~and managing them. This serves two purposes: ~~ensuring~~to ensure that SOA-based systems take account of the ~~true value~~specific constraints of a SOA ~~meeting the stated requirements can be realized using appropriate technology~~ecosystem, and ~~permitting~~to allow the audience to focus on the ~~important~~high-level issues without becoming over-burdened with ~~the~~ details of a particular implementation technology.

1.1.3 Relationship to the OASIS Reference Model for SOA

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

The ~~Reference Architecture~~SOA-RAF goes ~~a step~~ further ~~than the Reference Model~~ in that ~~we try to show~~it shows how ~~we might actually have~~ SOA-based systems~~.~~ can be realized – albeit in an abstract way. As noted above, SOA-based systems are better thought of as ~~ecosystems~~dynamic systems rather than stand-alone software products. Consequently, how they are used and managed is at least as important architecturally as how they are constructed.

~~In terms of approach, the primary difference between the Reference Model and this Reference Architecture is that the former focuses entirely on the distinguishing features of SOA; whereas this~~

 Copyright © OASIS Open 2011. All Rights Reserved.

~~document introduces concepts and architectural elements as needed in order to fulfill the core requirement of realizing SOA-based systems.~~

1.1.4 Relationship to other Reference Architectures

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

As with the Reference Model ~~for SOA, the~~, this Reference Architecture ~~for SOA~~ is primarily focused on large-scale distributed IT systems where the participants may be legally separate entities. ~~While it~~It is quite possible for many aspects of ~~the~~this Reference Architecture to be realized on quite different platforms~~, we do not dwell on such opportunities~~.

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

1.1.5 Expectations set by this Reference Architecture Foundation

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

~~In order to actually use, construct and manage SOA-based systems many additional design decisions and technology choices will need to be made. For example, we identify in this Reference Architecture a mode of interaction between service participants based on some form of message communication. The particular style of message communication, the transport technologies and the message encoding technologies are all important issues that are beyond the scope of this document. Similarly, the particular governance models used in a given application will need to be elaborated on and make concrete – for example, the exact committees and their jurisdictions would have to be set.~~

~~We believe that our approach will serve two purposes: ensuring that the true value of the SOA approach can be realized on any appropriate technology, and permitting our audience to focus on the important issues without becoming over-burdened with the details.~~

~~The primary contribution of this Reference Architecture is to make clear which technology and design choices are needed and what their purpose is. For example, we identify the role of participants and their relationships in terms of social structures. The specific organizations involved; how roles are designed and how the service interaction mechanisms determine the rights and responsibilities of the participants is also beyond our scope: we identify the need for the determination but not the specifics.~~

1.2 Service Oriented Architecture – An Ecosystems ~~perspective~~Perspective

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

 Copyright © OASIS Open 2011. All Rights Reserved.

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

Viewed as whole, a SOA ecosystem is a network of ~~independent services, machines, the people who operate, affect, use, and govern those~~ discrete processes and machines that, together with a community of people, creates, uses, and governs specific services as well as ~~the~~external suppliers of ~~equipment and personnel to these people and~~ resources required by those services.~~ This includes~~

In a SOA ecosystem there may not be any ~~entity, animate~~single person or ~~inanimate, that may affect or be affected by the system. With a system~~organization that ~~large, it is clear that nobody~~ is really "in control" or "in charge" of the whole ~~ecosystem;~~ although there are ~~definite~~identifiable stakeholders ~~involved, each of whom has some control and~~ who have influence ~~over~~ within the community and control over aspects of the overall system.

~~Instead of visualizing a SOA as a single complex machine, it is perhaps more productive to think of it as an ecosystem: a space where people, machines and services inhabit in order to further both their own objectives and the objectives of the larger community. In certain situations this may be a difficult psychological step for owners of so-called enterprise systems to take: after all, such owners may rightly believe that since they own the system they should also have complete control of it.~~

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

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

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

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

1.3 Viewpoints, Views and Models

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

1. ~~This Reference Architecture~~ The SOA-RAF uses and follows the ~~ANSI~~[2]~~/~~IEEE~~ Std 1471-2000 and ISO~~[3]~~/IEC~~[4]~~ 42010-2007 standard.~~ "Recommended Practice for Architectural Description of Software-Intensive Systems" **[ANSI/IEEE ~~Std~~ 1471~~,~~] and [ISO/IEC 42010]**. An architectural description conforming to ~~the ANSI/IEEE 1471-2000::ISO/IEC 42010-2007 recommended practice is described by a clause that includes~~this standard must include the following six (6) elements:
2. Architectural description identification, version, and overview information

[2] American National Standards Institute

[3] International Organization for Standardization

[4] International Electrotechnical Commission

 Copyright © OASIS Open 2011. All Rights Reserved.

3. Identification of the system stakeholder~~stakeholders~~s and their concerns judged to be relevant to the architecture
4. Specifications of each viewpoint that has been selected to organize the representation of the architecture and the rationale for those selections
5. One or more architectural views
6. A record of all known inconsistencies among the architectural description's required constituents
7. A rationale for selection of the architecture (in particular, showing how the architecture supports the identified stakeholders' concerns).

The ~~ANSI/IEEE 1471-2000::ISO/IEC 42010-2007~~standard defines the following terms[5]:

Architecture

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

Architectural Description

A collection of products that document the architecture.

System

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

System Stakeholder

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

A stakeholder's concern should not be confused with either a need or a formal requirement. A concern, as understood here, is an area or topic of interest. Within that concern, system stakeholders may have many different requirements. In other words, something that is of interest or importance is not the same as something that is obligatory or of necessity **[TOGAF ~~v8.1~~v9]**.

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

View

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

Viewpoint

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

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

It is important to note that viewpoints are independent of a particular system~~.~~ (or solutions). In this way, the architect can select a set of candidate viewpoints first, or create ~~a set of candidate~~new viewpoints, and then use those viewpoints to construct specific views that will be used to organize the architectural

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

 Copyright © OASIS Open 2011. All Rights Reserved.

description. A view, on the other hand, is specific to a particular system. Therefore, the practice of creating an architectural description involves first selecting the viewpoints and then using those viewpoints to construct specific views for a particular system or subsystem. Note that ~~ANSI/IEEE 1471-2000::ISO/IEC 42010-2007~~the standard requires that each view corresponds to exactly one viewpoint. This helps maintain consistency among architectural views~~;~~ which is a normative requirement of the standard.

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

Model

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

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

1.3.2 UML Modeling Notation

~~To~~An open standard modeling language is used to help visualize structural and behavioral architectural concepts~~, it is useful to depict them using an open standard visual modeling language.~~. Although many architecture description languages exist~~ in practice~~, we have adopted the Unified Modeling Language™ 2 (UML® 2) **[UML 2]** as the ~~primary~~ main viewpoint modeling language. ~~It~~ Normative UML is used unless otherwise stated but it should be noted that ~~while UML 2 is used in this Reference Architecture, formalization and recommendation of a UML Profile for SOA~~it can only partially describe the concepts in each model – it is ~~beyond the scope of this specification. Every attempt is made~~important to ~~utilize normative UML unless otherwise noted.~~read the text in order to gain a more complete understanding of the concepts being described in each section..

~~Figure 1 illustrates an annotated example of a UML class diagram that is used to represent a visual model depiction of the Resources Model in the Business via Services View (Section). The figure caption describes the UML semantics of this diagram.~~

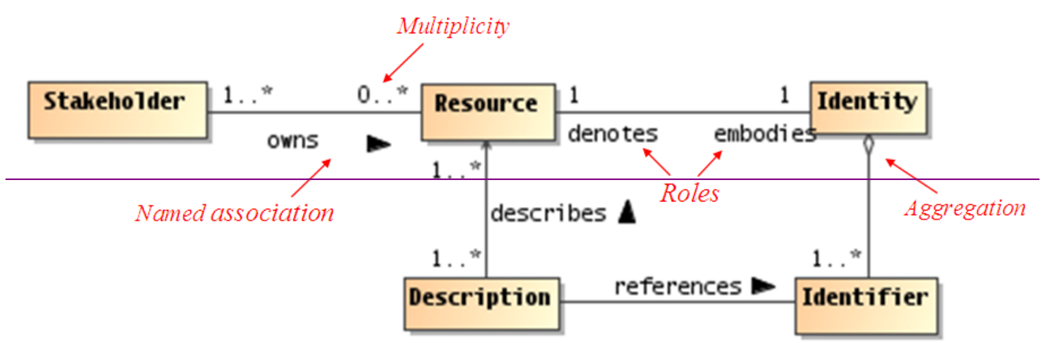

~~*Figure Example UML class diagram—Resources model.*~~

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

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

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

~~A stronger form of aggregation is known as composition, which involves using a filled-in diamond at the end of an association line (not shown in above diagram). For example, if the association between Identity and Identifier were a composition rather than an aggregation as shown, deleting Identity would also~~

 Copyright © OASIS Open 2011. All Rights Reserved.

~~delete any owned Identifiers. There is also an element of exclusive ownership in a composition relationship between classifiers, but this usually refers to specific instances of the owned classes (objects).~~

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

introduces the UML notation that is used in this document.

1.4 SOA-RAF Viewpoints ~~of this Reference Architecture~~

~~This Reference Architecture is partitioned into~~The RAF uses three views that conform to three ~~primary~~ viewpoints~~, reflecting the main division of concerns noted above: the viewpoint focuses on how people conduct their business using SOA-based systems; the viewpoint focuses on the salient aspects of building~~ : *Participation in a SOA Ecosystem*, *Realization of a SOA*~~, and the viewpoint focuses on those aspects that relate to owning, managing and controlling~~ *Ecosystem*, and *Ownership in a SOA*~~.~~

~~The viewpoint specifications for each of the primary viewpoints of this Reference Architecture are summarized in . Additional detail on each of the three viewpoints is further elaborated in the following subsections. For this Reference Architecture,~~ *Ecosystem*. There is a one-to-one correspondence between viewpoints and views (see Table 1~~is assumed.~~).

Viewpoint Element	Viewpoint		
	~~*Business via Services*~~Participation in a SOA Ecosystem	~~*Realizing Service Oriented Architectures*~~Realization of a SOA Ecosystem	~~*Owning Service Oriented Architectures*~~Ownership in a SOA Ecosystem
Main concepts covered	Captures what ~~SOA means~~is meant for people ~~using it~~ to ~~conduct business.~~participate in a SOA ecosystem.	~~Deals with the requirements for constructing a SOA.~~Captures what is meant to realize a SOA-based system in a SOA ecosystem.	~~Addresses issues involved in owning and managing a SOA.~~ Captures what is meant to own a SOA-based system in a SOA ecosystem
Stakeholders addressed	~~People (using SOA), Decision Makers, Enterprise Architects, Standards Architects and Analysts.~~All participants in the SOA ecosystem	~~Standards Architects, Enterprise Architects, Business Analysts, Decision Makers, Standards Architects and Analysts.~~Those involved in the design, development and deployment of SOA-based systems	~~Service Providers, Service Consumers, Decision Makers.~~Those involved in governing, managing, securing, and testing SOA-based systems
Concerns addressed	~~Conduct~~Understanding ecosystem constraints and contexts in which business ~~safely~~can be conducted predictably and effectively.	Effective construction of SOA-based systems.	Processes ~~for engaging in a SOA are effective, equitable~~to ensure governance, management, security, and ~~assured.~~testing of SOA-based systems.

 Copyright © OASIS Open 2011. All Rights Reserved.

Modeling Techniques used	UML class diagrams	UML class~~ and~~, sequence ~~diagrams~~, component, activity, communication, and composite structure diagrams	UML class and communication diagrams

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

1.4.1 ~~Business via Services~~Participation in a SOA Ecosystem Viewpoint

~~The Business via Services~~This viewpoint ~~is intended to capture what using a SOA-based system means~~captures a SOA ecosystem as an environment for people ~~using it~~ to conduct their business. We do not limit the applicability of ~~SOA-based systems~~such an ecosystem to commercial and enterprise systems. We use the term business to include any transactional activity ~~of interest to a user; especially activities shared by~~ between multiple users.

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

~~The~~All stakeholders ~~who~~in the ecosystem have ~~key roles in or~~ concerns addressed by this viewpoint~~ are decision makers and *people*.~~. The primary concern for people is to ensure that they can ~~use a SOA to~~ conduct their business effectively and safely in ~~a safe and effective way. For decision makers, their~~ accordance with the SOA paradigm. The primary concern ~~revolves around~~of decision makers is the relationships between people and organizations using systems ~~that the~~for which they, as decision makers, are responsible ~~for.~~

~~Given the public nature~~but which they may not entirely own, and for which they may not own all of the components of the ~~Internet, and the intended use of SOA to allow~~system.

Given SOA's value in allowing people to access, manage and provide services across ownership boundaries~~that cross ownership boundaries, it is necessary to be able to be somewhat explicit about~~, we must explicitly identify those boundaries and ~~what it means to cross an ownership boundary~~the implications of crossing them.

1.4.2 ~~Realizing Service Oriented Architectures~~Realization of a SOA Ecosystem Viewpoint

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

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

1.4.3 ~~Owning Service Oriented Architectures~~Ownership in a SOA Ecosystem Viewpoint

~~The Owning Service Oriented Architectures Viewpoint~~This viewpoint addresses the ~~issues~~concerns involved in owning ~~a SOA as opposed to using one or building one.~~and managing SOA-based systems within the SOA ecosystem. Many of these ~~issues~~concerns are not easily addressed by automation; instead, they often involve people-oriented processes such as governance bodies.

Owning a SOA-based system ~~involves~~ implies being able to manage an evolving system. ~~In our view, SOA-based systems are more like ecosystems than conventional applications; the challenges of owning and managing SOA-based systems are the challenges of managing~~It involves playing an active role in a wider ecosystem. ~~Thus, in this view, we are~~This viewpoint is concerned with how systems are managed effectively, how decisions are made and promulgated to the required end points~~, and~~; how to ensure that people may use the system effectively; and how the system can be protected against, and ~~that~~recover from consequences of, malicious ~~people cannot easily corrupt it for their own gain~~intent.

1.5 Terminology

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

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

~~Terms such as~~ The terms “SOA-RAF”, “this ~~“~~Reference Architecture” and “Reference Architecture Foundation” refer to this document, ~~and~~while “the Reference Model” ~~refer~~refers to the OASIS Reference Model for Service Oriented Architecture”. **[SOA-RM].**

1.5.1 Usage of Terms

Certain terms used in this document to denote concepts with formal definitions and are used with specific meanings. Where reference is made to a formally defined concept and the prescribed meaning is intended, we use a **bold font**. The first time these terms are used, they are also hyperlinked to their definition in the body of the text. Where a more colloquial or informal meaning is intended, these words are used without special emphasis.

1.6 References

1.6.1 Normative References

[ANSI/IEEE 1471] *IEEE Recommended Practice for Architectural Description of Software-Intensive Systems*, American National Standards Institute/Institute for Electrical and Electronics Engineers, September 21, 2000.

[ISO/IEC 42010] International Organization for Standardization and International Electrotechnical Commission, *System and software engineering — Recommended practice for architectural description of software-intensive systems*, July 15, 2007.

[RFC2119] S. Bradner, *Key words for use in RFCs to Indicate Requirement Levels*, http://www.ietf.org/rfc/rfc2119.txt, IETF RFC 2119, March 1997.

[SOA-RM] ~~C. M. MacKenzie, K. Laskey, F. McCabe, P. F. Brown, and R. Metz, (editors), “~~OASIS Standard, "Reference Model for Service Oriented Architecture 1.0, ~~OASIS Open,~~12 October ~~12,~~ 2006. http://docs.oasis-open.org/soa-rm/v1.0/soa-rm.pdf

[UML 2] *Unified Modeling Language: Superstructure*, Ver. 2.1.1, OMG Adopted Specification, OMG document formal/2007-02-05, Object Management Group, Needham, MA, February 5, 2007.

[WA] Architecture of the World Wide Web, W3C, 2004. http://www.w3.org/TR/webarch.

[WSA] David Booth, et al., "Web Services Architecture", W3C Working Group Note, World Wide Web Consortium (W3C) (Massachusetts Institute of Technology, European Research Consortium for Informatics and Mathematics, Keio University), February, 2004. http://www.w3.org/TR/2004/NOTE-ws-arch-20040211/

1.6.2 ~~[WA] Tim Berners Lee, *Design Issues*, W3C, 1996.~~ Non-Normative References

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

Copyright © OASIS Open 2011. All Rights Reserved.

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

[DCMI] Dublin Core reference

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

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

[ISO 11179] ISO 11179 reference

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

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

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

[TOGAF ~~v8.1~~v9] The Open Group Architecture Framework (TOGAF) ~~*8.1*~~Version 9 Enterprise Edition**,** The Open Group, Doc Number: ~~G051, December 19, 2003~~G091, February 2009**.**

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

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

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

~~**[WA]** Architecture of the World Wide Web, W3C, 2004. http://www.w3.org/TR/webarch.~~

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

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

~~**[RFC 4880]** Network Working Group, *OpenPGP Message Format*, 2007~~

~~**[WS-BPEL]** Web Services Business Process Execution Language Version 2.0 *http://docs.oasis-open.org/wsbpel/2.0/OS/wsbpel-v2.0-OS.html*~~

~~**[WS-CDL]** Web Services Choreography Description Language Version 1.0, *http://www.w3.org/TR/ws-cdl-10/*~~

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

 Copyright © OASIS Open 2011. All Rights Reserved.

2 Architectural Goals and Principles

~~In this~~This section~~, we identify both the goals of the architecture and the architectural principles that underlie our approach to the architecture.~~

~~In order to be clearer in setting~~ identifies the goals of this Reference Architecture~~, we~~ Foundation and the architectural principles that underpin it.

2.1 Goals and Critical Success Factors of the Reference Architecture Foundation

There are three principal goals:

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

There are four factors critical ~~factors analysis~~ to ~~identify~~ the ~~key~~achievement of these goals:

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

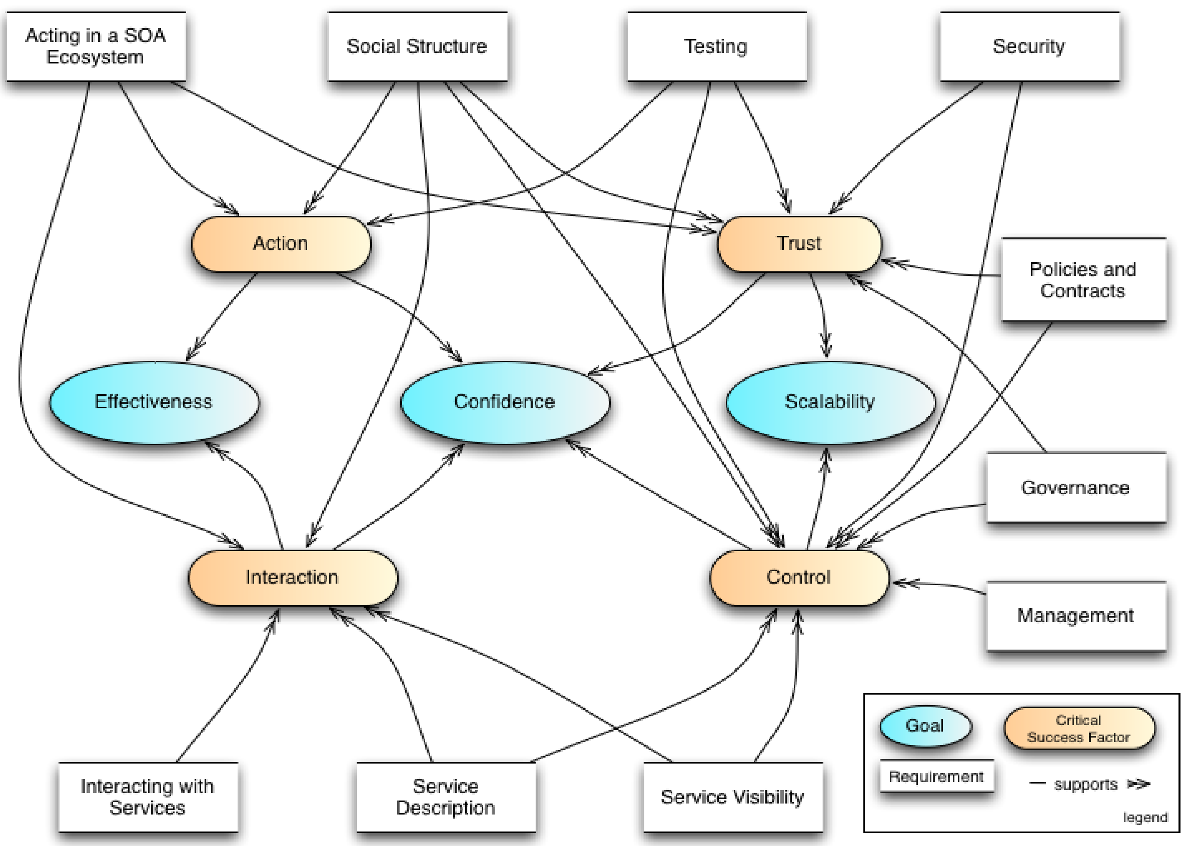

 Copyright © OASIS Open 2011. All Rights Reserved.

Figure 1 - Critical Factors Analysis of the Reference Architecture

Figure 1 represents a Critical Factors Analysis (CFA) diagram demonstrating the relationship between the primary goals of this reference architecture, critical factors that determine the success ~~factors and requirements~~ of ~~this~~the architecture~~.~~ and individual elements that need to be modeled.

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

2.1.1 Goals ~~of this Reference Architecture~~

~~Note that not all of the requirements are mapped to solutions within the scope of this Reference Architecture. Indeed, this document can be seen as generating a series of more explicit requirements for the realizing technology.~~

~~The overall requirements are illustrated in .~~

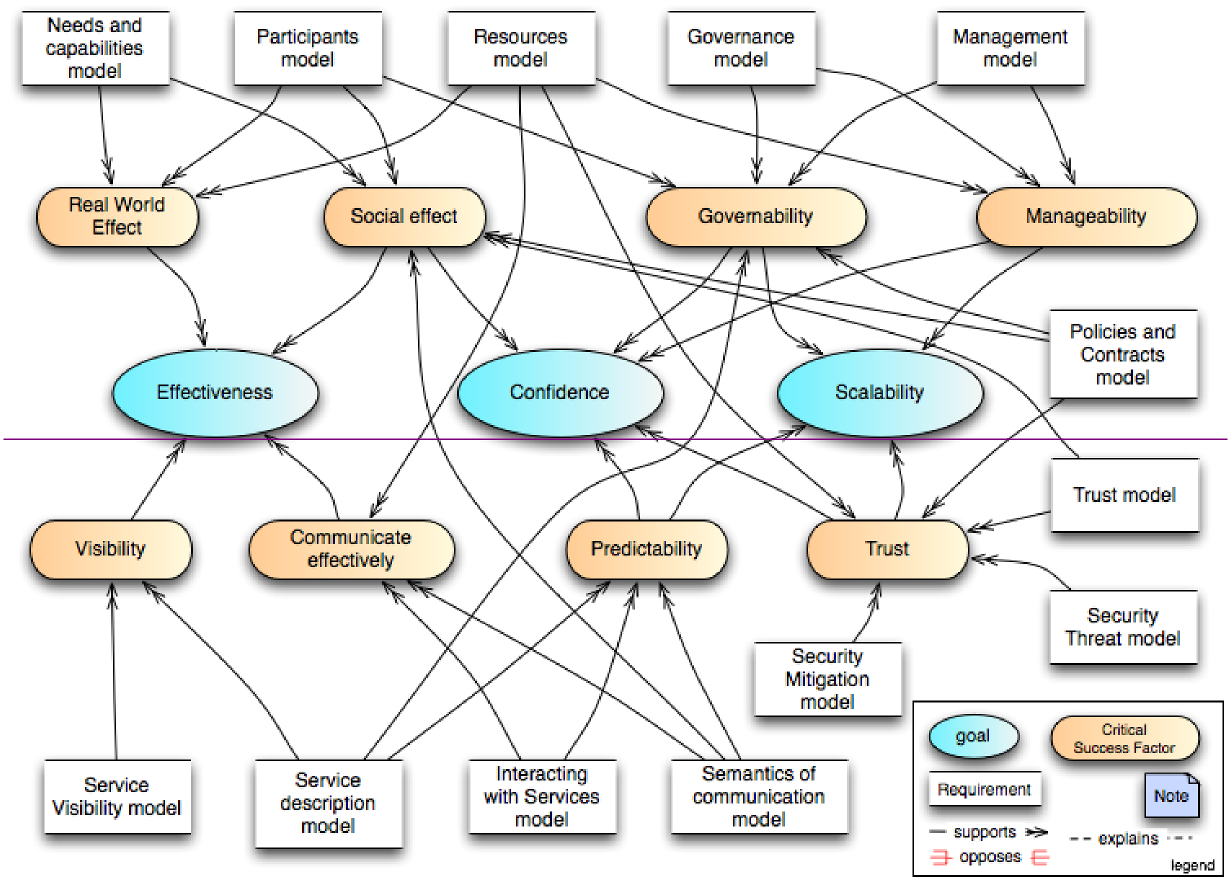

~~*Figure Critical Factors Analysis of the Reference Architecture*~~

~~There are three principal goals of this Reference Architecture:~~

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

 Copyright © OASIS Open 2011. All Rights Reserved.

2.1.1.1 Effectiveness

A primary purpose of ~~this architecture~~the SOA-RAF is to show ~~what is involved in~~how SOA-based systems ~~to~~ ensure that participants can use the facilities of the system to ~~get~~meet their needs ~~met. Of course, not all participants' needs can be met by interacting electronically; but~~ . This does not imply that every need has a SOA solution, but for those needs that can~~, can be met using~~ benefit, we look at what is needed to use the ~~framework of a~~ SOA ~~-based system~~ paradigm effectively.

The ~~critical~~key factors that ~~determine~~govern effectiveness from a participant's perspective are ~~visibility between the~~ actions undertaken– especially across ownership boundaries – with other participants~~, that they can communicate effectively,~~ in the ecosystem and ~~that actual real world effects and social effects can be realized. In addition, it is critical that the overall system is manageable and governable.~~

~~**Real World Effect**~~

~~It is of the essence that participants can use a SOA-based system~~lead to ~~realize actual effects in the world. This implies that the capabilities that are accessed as a result of service interaction are 'wired-up' so to speak, with the real world.~~

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

~~2.1.1.2 Social effect~~

~~Many, if not most, effects that are desired in the use of SOA-based systems are actually social effects more than physical effects. For example, opening a bank account is primarily about the relationship between a customer and a bank – the effect of the opened account is a change in the relationship between the customer and the bank.~~

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

~~2.1.1.3 Visibility~~

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

~~2.1.1.4 Communicate effectively~~

~~In order for there to be effective uses of capabilities and meeting of needs, it is critical that participants can not only see each other but can also interact with each other. The models that address this are the Interacting with Services model, the Resources model and the Semantics of~~ measurable results~~Communication model~~.

~~2.1.1.5~~2.1.1.2 Confidence

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

~~In addition to ensuring~~Confidence has many dimensions: confidence in the successful interactions with other participants, confidence in the assessment of trust, as well as confidence that ~~social effects are~~ the ecosystem is properly managed.

2.1.1.3 Scalability

The third goal of this reference architecture is scalability. In architectural terms, we determine scalability in terms of the smooth growth of complex systems as the number and complexity of services and

 Copyright © OASIS Open 2011. All Rights Reserved.

interactions between participants increases. Another measure of scalability is the ease with which interactions can cross ownership boundaries.

2.1.2 ~~captured, other~~ Critical Success Factors

A critical success factor (CSF) is a property of the intended system, or a sub-goal that directly supports a goal and there is strong belief that without it the goal is unattainable. CSFs are not necessarily measurable in themselves. As illustrated in Figure 1, CSFs can be associated with more than one goal.

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

~~2.1.1.6 Manageability and Governability~~

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

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

2.1.2.1 ~~The scope of management of SOA-based systems is constrained by the existence of multiple ownership domains.~~ Action

Participants' principal mode of participation in a SOA ecosystem is action; typically action in the interest of achieving some desired real world effect~~Management may include setting policies such as technology choices but may not, in some cases, include setting policies about the services that are offered.~~

. Understanding how action is related to SOA is thus critical to the paradigm.

~~2.1.1.7~~2.1.2.2 Trust

The viability of a SOA ecosystem depends on participants being able to effectively measure the trustworthiness of the system and of participants. Trust is a private assessment of a participant's belief in the integrity and reliability of the SOA ecosystem (see Section 3.1.4~~Trust itself is clearly a critical factor in ensuring confidence. Trust itself~~).

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

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

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

 Copyright © OASIS Open 2011. All Rights Reserved.

~~2.1.1.8~~ Interaction~~Predictability~~

2.1.2.3 ~~A factor that engenders confidence in any system~~

In order for a SOA ecosystem to function, it is ~~predictability. By predictability, we principally mean~~essential that the ~~expectations of~~means for participants ~~of SOA-based systems can be tied to the actual performance of those systems (what you see is what you get).~~

~~The primary means of ensuring predictability is effective descriptions: service descriptions document services,~~to interact with each other is available throughout the system. Interaction encompasses not only the ~~interacting with services model addresses expectations relating to how services are used~~mechanics and ~~the~~ semantics of ~~communications model addresses how meaning~~communication but also the means for discovering and ~~intent can be exchanged between participants~~offering communication.

2.1.2.4 Control

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

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

~~2.1.1.9~~1.1.1.1 The scope of management of SOA-based systems is constrained by the existence of multiple ownership domains. ~~Scalability~~

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

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

2.2 Principles of this Reference Architecture Foundation

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

~~Principle 1:~~ Technology Neutrality

Statement: Technology neutrality refers to independence from particular technologies.

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

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

 Copyright © OASIS Open 2011. All Rights Reserved.

architects and other decision makers to construct systems based on this architecture is needed.

~~**Principle 2:**~~ **Parsimony**

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

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

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

~~Principle 3:~~ ~~Separation~~Implications: Parsimoniously designed systems tend to have fewer but better targeted features.

Distinction of Concerns

Statement: ~~Separation~~Distinction of Concerns refers to the ability to cleanly ~~delineate~~identify and separate out the concerns of specific stakeholders in such a way that it is possible to create architectural models ~~in such a way that~~that reflect those stakeholders’ viewpoint. In this way, an individual stakeholder or a set of stakeholders that share common concerns only see those models that directly address their respective areas of interest. ~~This principle could just as easily be referred to as the Separation of Stakeholder Concerns principle, but the focus here is predominantly on loose coupling of models.~~

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

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

~~**Principle 4:**~~ **Applicability**

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

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

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

 Copyright © OASIS Open 2011. All Rights Reserved.

3 ~~Business via Services~~Participation in a SOA Ecosystem View

No man is an island

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

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

Together, these ideas describe an environment in which business functions (realized in the form of services) address business needs. Service implementations utilize capabilities to produce specific (real world) effects that fulfill those business needs. Both the people[6] using the services, and the capabilities themselves, may be distributed across ownership domains, with different policies and conditions of use in force– this environment is referred to as a **SOA Ecosystem**.

The role of a service in the SOA context is to enable effective business solutions in this environment. Any technology system created to deliver a service in such an environment is referred to as a **SOA-based system**. SOA is thus a paradigm that guides the identification, design, implementation (i.e., organization), and utilization of such services.

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

The *Participation in a SOA Ecosystem* view in the SOA-RAF focuses on the constraints and context in which people conduct business using a SOA-based system. By business we mean any shared activity entered into whose **objective** is to satisfy particular **needs** of each participant. The OASIS SOA RM defines SOA as "a paradigm for organizing and utilizing distributed capabilities that may be under the control of different ownership domains." To put it another way, to effectively employ the SOA paradigm, the architecture must take into account the fact and implications of different ownership domains, and how

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

 Copyright © OASIS Open 2011. All Rights Reserved.

best to organize and utilize capabilities that are distributed across those different ownership domains. These are the main architectural issues that the Participating in a SOA Ecosystem view tries to address.

The subsections below expand on the completely abstract reference model by identifying more fully and with more specificity what challenges need to be addressed in order to successfully apply the SOA paradigm. Although this section does not provide a specific recipe, it does identify the important things that need to be thought about and resolved within an ecosystem context.

The people actively participating in a SOA-based system, together with others who may potentially benefit from the services delivered by the system, together constitute the **stakeholders**. The stakeholders, the system and the environment (or context) within which they all operate, taken together forms the **SOA ecosystem**. That ecosystem may reflect the SOA-based activities within a particular enterprise or of a wider network of one or more enterprises and individuals.Although a SOA-based system is essentially an IT concern, it is nonetheless a system engineered deliberately to be able to function in a SOA ecosystem. In this context, a service is the mechanism that brings a SOA-based system capability together with stakeholder needs in the wider ecosystem. This is explored in more detail in Section 3.2.2 below.

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

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

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

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

The main models in this view are:

- the **Social Structure in a SOA Ecosystem Model** introduces the key elements that underlie the relationships between participants~~The *Business via Services*~~ and that must be considered as pre-conditions in order to effectively bring needs and capabilities together across ownership boundaries;
- the **Action in a SOA Ecosystem Model** introduces the key concepts involved in service actions, and shows how joint action and real-world effect are what is being aimed for in a SOA ecosystem..

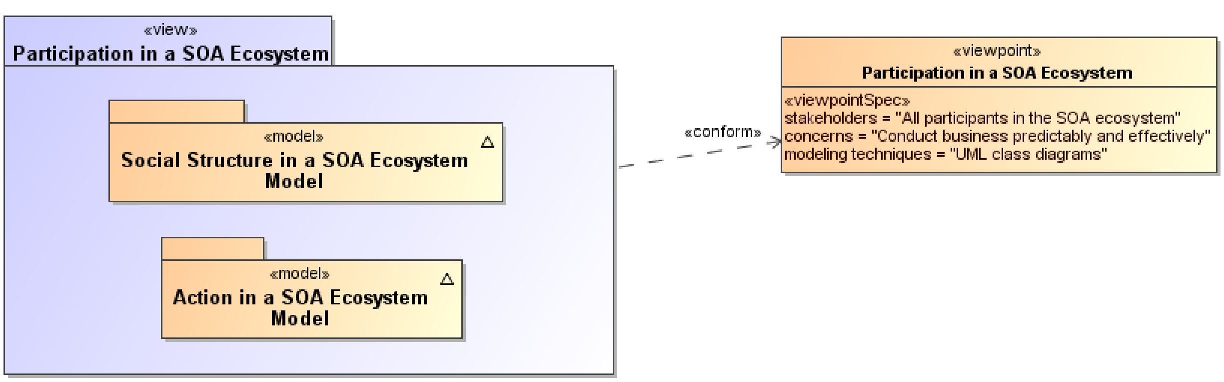

 Copyright © OASIS Open 2011. All Rights Reserved.

~~View focuses on what a SOA-based system means for people using it to conduct their business.[7] The mode of business in a SOA-based system is characterized in terms of providing services and consuming services to realize mutually desirable real world effects.~~

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

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

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

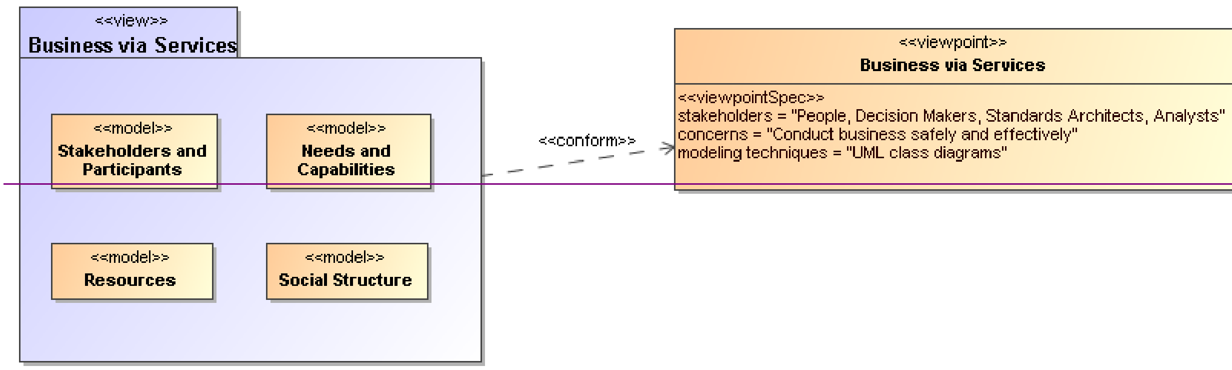

Figure 2 - Model elements described in the ~~Business via Services~~Participation in a SOA Ecosystem view

3.1 ~~Stakeholders and Participants~~Social Structure in a SOA Ecosystem Model

The actions undertaken by participants in a SOA ecosystem are performed in a *social context* that defines the relationships between the participants. That context is the social structure~~A SOA-based system is deployed in the context of human and non-human entities capable of action. In this section we focus on the relationship between these ultimate actors and the services that they use and deploy.~~

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

 Copyright © OASIS Open 2011. All Rights Reserved.

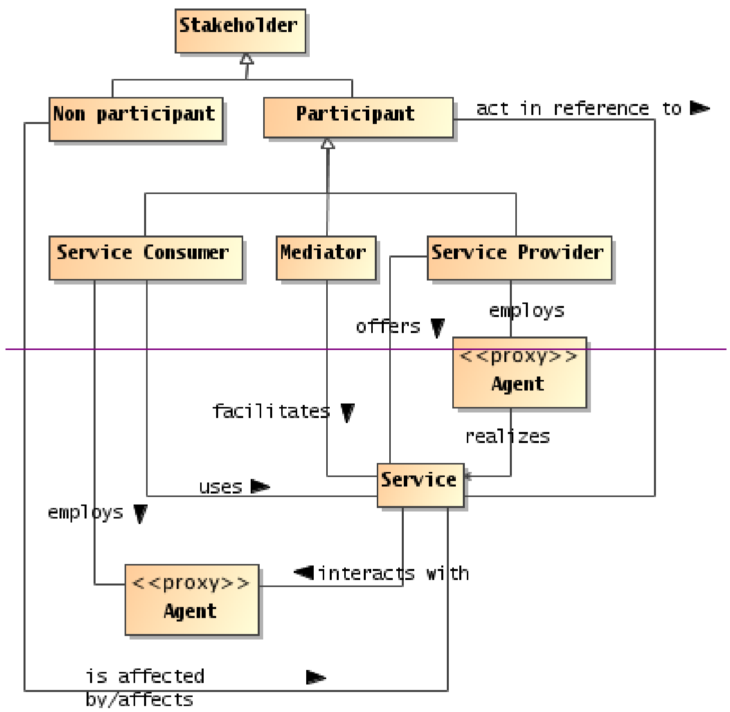

~~*Figure Service Participants*~~

~~**Stakeholder**~~

~~A stakeholder is ~~. In order to achieve success in applying the SOA paradigm, the overall social structure in which the SOA effort is to be undertaken must be taken into consideration. Ownership boundaries and their implications can only be understood and addressed within the context of the larger social structure within which they exist and the nature of the relationships between the different participants in that structure.

> The primary function of the Social Structure Model is to explain the relationships between an individual participant~~entity, human or non-human, or organization of entities that has an interest in the states of services and/or the outcomes of service interactions.~~

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

~~**Participant**~~

> ~~A participant is a stakeholder that has the capability to act in the context of a SOA-based system.~~

~~A participant is a stakeholder whose interests lie in the successful use of and fulfillment of services. However, human participants always require *representation* in an electronic system – they require agents. Note that we admit non-human agents that have no identifiable representative as an extreme case: the normal situation is where participants are either human or organizations.~~

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

~~**Service Provider**~~

Copyright © OASIS Open 2011. All Rights Reserved.

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

~~In normal parlance, the service provider commonly refers to either the ultimate owner of the capability that is offered or at least an agent acting as proxy for the owner. For example, an individual may own a business capability but will enter into an agreement with another individual (the proxy) to provide SOA access to that business -- so that the owner can focus on running the business itself.~~

~~Note that several kinds of stakeholders may be involved in provisioning a service. These include but are not limited to the provider of the capability, an enabler~~ and the social context of that participant~~exposes it as a service, a mediator that translates and/or manages the relationship between service consumers and the service, a host that offers support for the service, a government that permits the service and/or collects taxes based on service interactions.~~

~~Service Consumer~~

> ~~A service consumer is a participant that interacts with a service in order to access a capability to address a need.~~

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

~~Service mediator~~

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

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

~~Agent~~

> ~~An agent is any entity that is capable of acting on behalf of a person or organization.~~

~~In order for people to be able to offer, consume and otherwise participate in services, they require the use of an agent capable of directly interacting with electronic communications – a service agent. Common examples are software applications that make use of services, hardware devices that embody an agent with a particular mission, and enterprise systems that offer services.~~

~~We do not attempt to characterize service agents in terms of their internal architecture, computational requirements or platforms here.~~

~~Non-participant stakeholder~~

> ~~A non-participant is any stakeholder who may be affected by the use or provisioning of services or who has an interest in the outcome of service interactions but does not directly participate in and may not be aware of the interactions.~~

~~There are two main classes of such non-participatory stakeholders: third parties who are affected by someone's use or provisioning of a service, and regulatory agencies who wish to control the outcome of service interactions in some way (such as by taxation).~~

~~Resources~~ . The~~Model~~

~~In many instances it is important to be able to~~ model ~~the assets that stakeholders may have access to. The Reference Architecture itself has many instances of such resources; for example service descriptions, services themselves and the capabilities that underlie services are all resources.~~

 Copyright © OASIS Open 2011. All Rights Reserved.

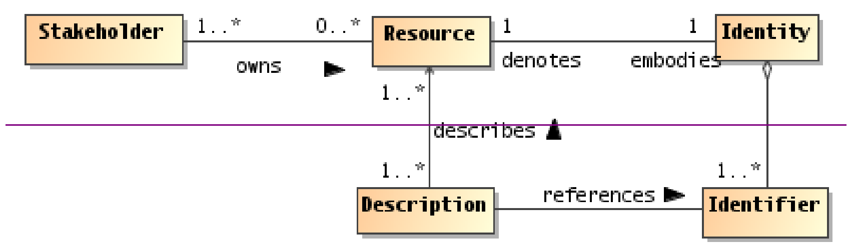

~~*Figure Resources model*~~

~~Our model of resources is very simple, but is the foundation for modeling many of the things that a SOA-based system deals in such as information, services, capabilities, descriptions, policies and contracts.~~

~~3.1.1.11.1.1.1 Resource~~

> ~~A resource is any entity of some perceived value, where the value may be in the function it performs or something intrinsic in its nature. may vary over time.~~

~~A resource has identity and it has an owner. A resource may have more than one identifier, but any well-formed identifier should unambiguously resolve to the intended resource.~~

~~An important class of resource is the class of capabilities that underlie services. For example, a light bulb is a resource that when activated gives off light; a book is a resource that when read allows one to gain knowledge from its content. Other examples of resources are services themselves, descriptions of entities (a kind of meta-resource), IT infrastructure elements used to deliver services, contracts and policies, and so on.~~

~~Identity~~

> ~~Identity is the collection of individual characteristics by which a thing or person is recognized or known. In this architecture, we further restrict this to the collection of identifiers by which a person or thing is known.~~

~~Identity is an important, if abstract, concept. For example, in ensuring that a user is authenticated, the role of the authentication process is to validate the identity of the person that is attempting to gain access to a resource.~~

~~Identifier~~

> ~~An identifier is any block of data – such as a string – that is associated with a particular identity.~~

~~It is good practice to use globally unique identifiers; for example globally unique IRIs. However, the primary requirement of an identifier is that it can be used to uniquely disambiguate the indicated resource from other resources.~~

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

~~3.1.2 Ownership Model~~

~~Understanding what it means to own something it important when we use an SOA-based system to exchange value. Ownership is~~ also ~~important in understanding the various kinds of obligations participants may enter into. Fundamentally, we view ownership as a relationship between a stakeholder and a resource, where the owner has certain rights over the resource (note not necessarily absolute rights).~~

~~Ownership~~

> ~~Ownership is a relationship between an entity, a resource and a set of rights and responsibilities. When an entity owns a resource, the entity has the right to exercise the rights over the resource and may transfer ownership to another entity.~~

 Copyright © OASIS Open 2011. All Rights Reserved.

~~In addition, owning a resource brings with it a set of responsibilities. The nature of these responsibilities will vary with the resource and the nature of the ownership; but typically, if the use of a resource harms someone, then the owner of the resource will be held responsible.~~

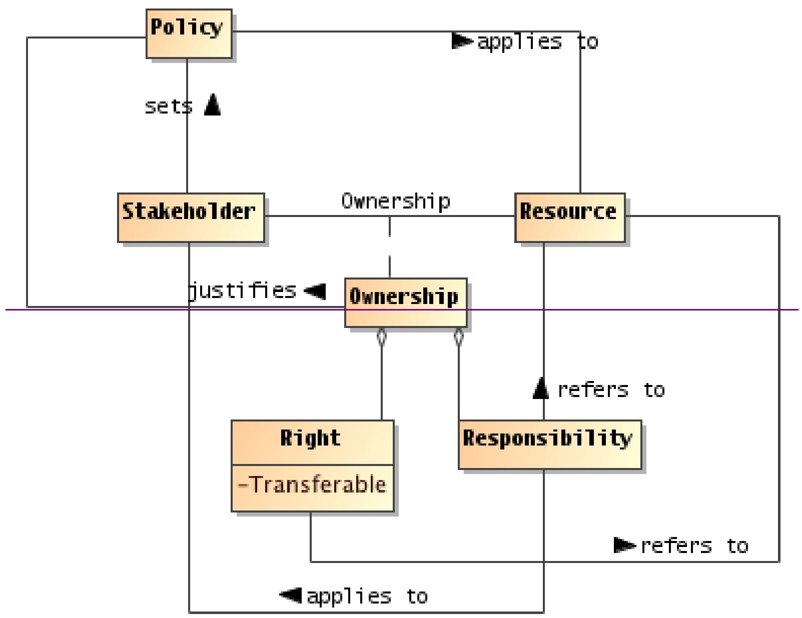

~~*Figure Resource Ownership Model*~~

~~To own a resource implies taking responsibility for creating, maintaining, and if it is to be available to others, provisioning the resource. One who owns a resource may delegate any of these functions to others, but still has the responsibility to see the function is done. There may also be joint ownership of a resource, where the responsibility is shared.~~

~~Ownership is rarely absolute, rarely involves complete control over the resource. In reality, ownership is normally constrained to a particular set of rights. For example, one stakeholder may own the rights to deploy a capability as a service, another may own the rights to the profits that result from using the capability, and yet another may own the rights to use the service! However, a crucial property that distinguishes ownership from merely renting is the right to transfer ownership to another person or organization.~~

~~3.2 Needs and Capabilities Model~~

~~The motivation for participants interacting is the satisfaction of needs. From a consumer perspective, the motivation for interacting with a service is to satisfy a business objective, which in turn, is often related to the role they represent in the social structure; for the provider, the need is to gain satisfaction, monetary or otherwise, for other participants' use of the service.~~

 Copyright © OASIS Open 2011. All Rights Reserved.

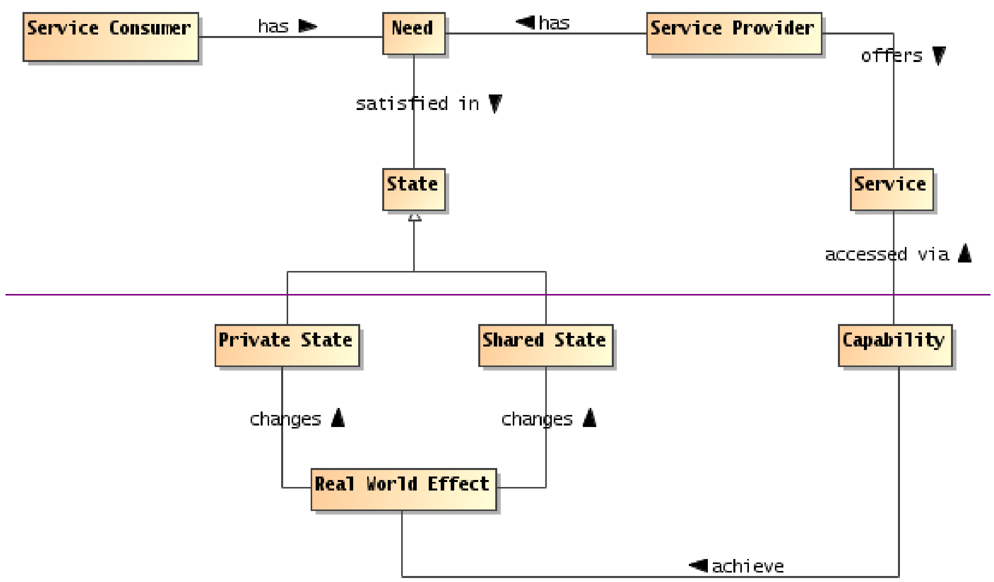

~~*Figure Needs and Capabilities*~~

~~**Capability**~~

> ~~A capability is a resource that may be used by a service provider to achieve a real world effect on behalf of a service consumer.~~

~~The model in show that there is an inherent indirection between needs and having them satisfied. Both needs and the effects of using capabilities are expressed in terms of state: a need is expressed as a condition on the desired state and the Real World Effect of using capabilities is a change in the state of the world.~~

~~As noted in the Reference Model, the Real World Effect is couched in terms of changes to the state that is shared by the participants in the service; in particular the public aspects of that state. In this Reference Architecture we further refine this notion in terms of changes in the social facts that are mandated by social structures – see Section .~~

~~By making a capability available for use, via the Service, the owners aim to address their needs as well as the needs of other participants who use the service. The extent to which a capability is exposed via a service (or via multiple services) is controlled by the owner of the capability.~~

helps in~~Need~~

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

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

~~A need is characterized by a proposition – see Section . However, the extent to which a need is captured in a formal way is likely to be very different in each situation.~~

~~3.3 Social Structure Model~~

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

~~The social structure model is important to~~ defining and understanding the implications of crossing ownership boundaries~~ownership boundaries; it~~. It is, for example, the foundation for ~~an~~ understanding ~~of~~ security~~ in SOA and also provides the context for determining how SOA-based systems can be effectively managed and governed~~, governance and management in the SOA ecosystem.

Copyright © OASIS Open 2011. All Rights Reserved.

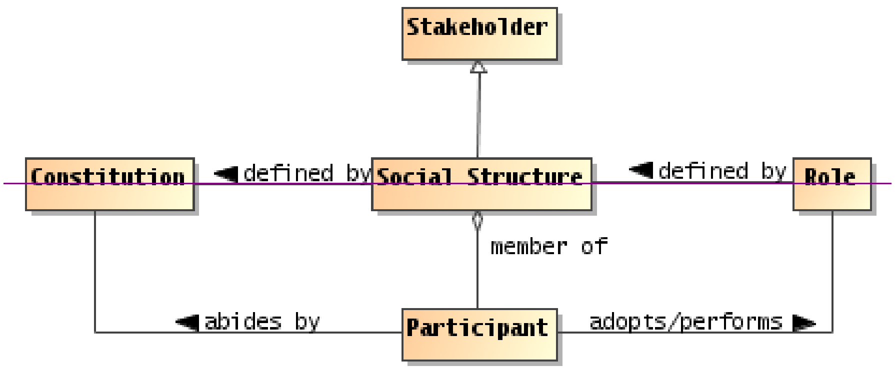

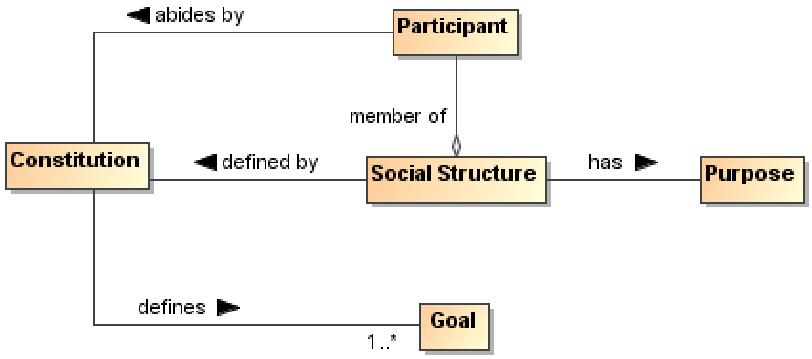

Figure 3~~ Social Structure~~

 - *Social Structure*

Social Structure

A social structure[8]~~social structure (sometimes identified as social institutions) embodies some~~ is a nexus of ~~the cultural aspects that characterize the~~ relationships amongst participants brought together for a specific purpose.

A social structure represents a collection of participants and is established with an implied or explicitly defined purpose. The purpose is usually reflected in specific goals laid down in the social structure's constitution~~actions~~ or other 'charter'.

A social structure may have any number of participants and a large number of different relationships may exist among ~~a group of participants~~participants. The organizing principle for these relationships is the social structure's purpose. In addition, a given participant can be a member of multiple social structures. Thus, there may be interaction among social structures, sometimes resulting in disagreements when the premises of the social structures do not align.

~~In the Reference Architecture, we are concerned primarily with social structures that reflect the anticipated participants in SOA-based systems; these are often embodied in legal and quasi-legal frameworks; i.e., they have some rules that are commonly understood.~~

[8] Social structures are sometimes referred to as social institutions.

Copyright © OASIS Open 2011. All Rights Reserved.

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

~~It~~The RAF is ~~not necessarily the case that the~~concerned primarily with social structures ~~involved in a service interaction are explicitly identified by the~~that reflect relationships amongst **participants** in SOA ecosystems, notably:

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

Enterprise

> An enterprise is a social structure with an identifiable leadership structure, and that has internally established goals that reflect a defined purpose. It can act as a participant within other social structures, including other enterprises and is represented by members of its leadership structure.

Peer Group

> A peer group is a social structure with no discernable leadership structure, that may or may not have internally established goals, but is identiable as the locus of interaction between participants with individual goals seeking common outcomes and who are considered peers of one another.

Many interactions between participants take place within social structures. Depending on the scale and internal structure of an enterprise social structure, these interactions may or may not cross ownership boundaries (an enterprise can itself be composed of sub-enterprises). However, interactions between participants within a peer social structure inherently cross ownership boundaries~~.~~.

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

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

Constitution

> A constitution is ~~an agreement shared by~~ a ~~group~~set of ~~participants~~rules, written or unwritten, that ~~defines a social structure.~~
>
> ~~The primary~~ spell out the purpose ~~of the constitution is to define the roles of participants in the institution~~, goals, scope, and functioning of a social structure~~how to establish the regulations that define the legal actions. The regulations of the social structure effectively define how those assertions and commitments that are relevant to the social structure are created.~~.

~~A constitution may be explicitly written down or it may be only partially written.~~

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

~~A constitution is an agreement. It is abided to by the participants in the~~Every social structure functions according to rules by which participants interact with each other within the structure. In some cases, this is based on an explicit agreement, in other cases participants behave as though they agree to the constitution~~constitution~~ without a formal agreement. ~~For example, when a new employee joins a company, he or she is often required to sign an employment contract. That contract defines key aspects of the relationship between the new employee and the company.~~ In still other ~~situations the act~~cases, participants abide by the rules with some degree of reluctance, such as governance of SOA-based

 Copyright © OASIS Open 2011. All Rights Reserved.

systems, covered below. In all cases, the constitution may change over time, in those cases of implicit agreement ~~is less formal and less clearly established~~the change can occur quickly.

~~3.3.1~~3.1.1 ~~Shared State~~Participants, Actors and ~~social facts~~Delegates

Social structure~~Most~~s have **Stakeholders** – people – some of whom may be enterprises. They interact within the broad SOA ecosystem. **Actors** on the other hand operate strictly within a SOA-based system.

There is also the concept of **Participant** which is particularly important as it reflects the hybrid role of a person who is both a Stakeholder in the ecosystem (and thus primarily concerned with expressing needs and seeing those needs fulfilled) *and* an Actor in the System (and thus directly involved with system-level activity.

A stakeholder can be either a participant (and thus also an actor with a specific functional role in a SOA-based system); or a non-participant – someone who, without participating, nonetheless has something at stake within the ecosystem.

An actor can be either a **participant** (and thus also a stakeholder with a stake in the ecosystem); or a **delegate** – a human actor with no stake in the specific action delegated or some automated agent – acting on behalf of a participant.

The hybrid role of Participant provides a bridge between the wider (real-world) ecosystem – the world of the stakeholder – and the more specific (usually technology-focused) system – the world of the actor.

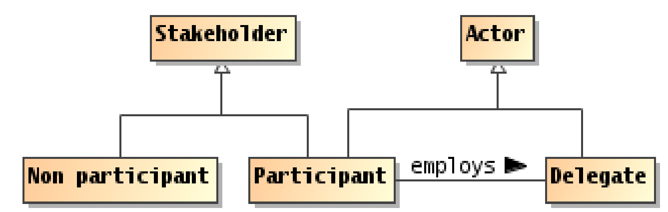

Figure 4 - Actors, Participants and Delegates

Stakeholder

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

Note: Not all stakeholders necessarily participate in the SOA ecosystem; indeed, the interest of non-participant stakeholders may be in realizing the benefits of a well-functioning ecosystem and not suffering unwanted consequences. They can not all or always be identified in advance but due account is often taken of such stakeholder types, including potential customers, beneficiaries, affected third parties, as well as potential "negative stakeholders" who might deliberately seek a negative impact on the ecosystem (such as hackers or criminals).

Actor

An actor is a human or non-human agent capable of action within a SOA-based system.

Participant

A participant is a person[9] who is both a stakeholder in the SOA ecosystem and an actor in the SOA-based system.

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

 Copyright © OASIS Open 2011. All Rights Reserved.

Delegate

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

A delegate can be a person or an automated or semi-automated agent.

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

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

In order for a delegate to act on behalf of another person, they must be able to act and have the authority to do so.

3.1.2 Roles in Social Structures

Social structures are abstractions: a social structure cannot directly perform actions – only people or automated processes following the instructions of people can actually do things. However, an actor may act on behalf of a social structure and certainly acts within a social structure depending on the roles that the actor~~people and most of the important aspects of a person's state are inherently social in~~ assumes and the nature of the relationships between the concerned parties or stakeholders.

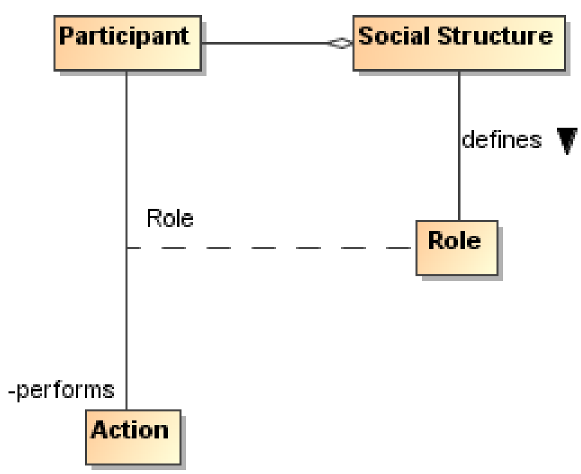

Figure 5 - Role in Social Structures

Role

A role is a type of relationship between a participant and the actions that the participant may perform (or is allowed to perform) within a social structure.

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

One participant with appropriate authority in the social structure may formally *designate a role* for another participant, with associated rights and responsibilities, and that authority may even qualify a period during which the designated role may be valid. In addition, while many roles are clearly identified, with appropriate names and definitions of responsibilities, it is also possible to separately bestow rights, bestow or assume responsibilities and so on, often in a temporary fashion. For example, when a

 Copyright © OASIS Open 2011. All Rights Reserved.

company president delegates certain responsibilities on another person, this does not imply that the other person has become company president. Likewise, a company president may bestow on someone else her role during a period of time that she is on vacation or otherwise unreachable, with the understanding that she will re-assume the role when she returns from vacation.

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

Authority

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

Right

> A right is a predetermined permission conferred upon an actor to perform some action or assume a role in relation to the social structure.

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

Responsibility

> A responsibility is a predetermined obligation on a participant to perform some action or assume a role in relation to other participants.

Responsibility implies human agency, which is why only participants, as opposed to all actors (who can be non-human agents) are concerned. This applies even if the consequences of such responsibility can impact other (human and non-human) actors. Having authority often implies having responsibility.

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

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

Two important types of constraints that are relevant to a SOA ecosystem are Permission and Obligation.

Permission~~an action is what gives it much of its meaning. We call actions in society social~~

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

Note that permissions are distinct from ability and from authority. Authority refers to the legitimate nature of an action as performed by an actor on behalf of a social structure. Ability refers to whether an actor has the capacity to perform the action. Permission does not always involve acting on behalf of anyone, nor does it imply or require the capacity to perform the action.

Obligation

> An obligation is a constraint that prescribes the actions that an actor must (or must not) perform and/or the states the actor must (or must not) attain or maintain.

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

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

 Copyright © OASIS Open 2011. All Rights Reserved.

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

3.1.2.1 Service Roles

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

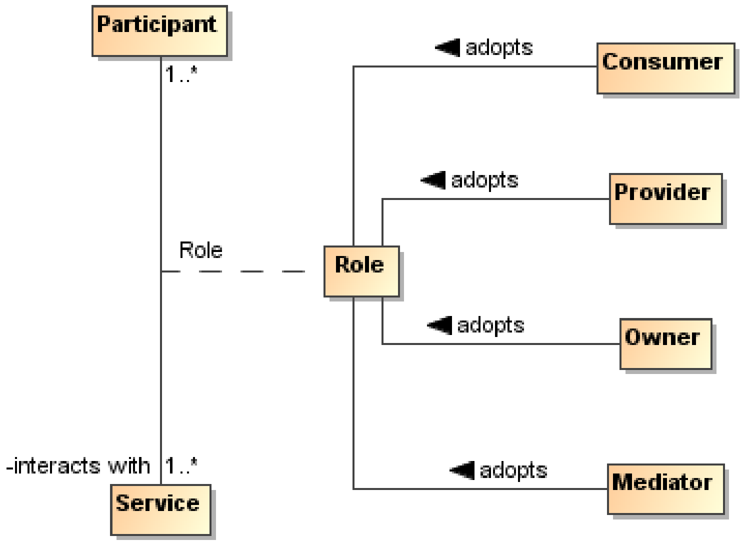

Figure 6 - Participant Roles in a Service

Provider

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

Consumer

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

Mediator

A mediator is a role assumed by a participant to facilitate interaction and connectivity in the offering and use of services.

Owner

An owner is a role assumed by a participant who is claiming and exercising ownership over a service.

It is a common understanding that service interactions are typically initiated by service consumers, although this is not necessarily true in all situations. Additionally, as with service providers, several stakeholders may be involved in a service interaction supporting a given consumer.

The roles of service provider and service consumer are often seen as symmetrical, which is also not entirely correct. A consumer tends to express a 'Need' in non-formal terms: "I want to buy that book". The type of 'Need' that a service is intended to fulfill has to be formalized and encapsulated by designers and developers as a 'Requirement'. This Requirement should then be reflected in the target service, as a 'Capability'that, when accessed via a service, delivers a 'Real World Effect' to an arbitrary user: "The chosen book is ordered for the user." It thus satisfies the need that has been defined for an archetypal user. Specific and particular users may not experience a need exactly as captured by the service: "I don't want to pay that much for the book", "I wanted an eBook version", etc. There can therefore be a process

 Copyright © OASIS Open 2011. All Rights Reserved.

of implicit and explicit negotiation between the user and the service, aimed at finding a 'best fit' between the user's specific need and the capabilities of the service that are available and consistent with the service provider's offering. This process may continue up until the point that the user is able to accept what is on offer as being the best fit and finally 'invokes' the service. 'Execution context' has thus been established. This is explored in more detail later on. Service mediation by a participant can take many forms and may invoke and use other services in order to fulfill such mediation. For example, it might use a service registry in order to identify possible service partners; or, in our book-buying example, it might provide a price comparison service, suggest alternative suppliers, different language editions or delivery options.

3.1.3 Resource and Ownership

3.1.3.1 Resource

~~and those facts that are~~ A resource is generally understood ~~in a society social facts. It is often the case that social actions give rise to social facts~~as an asset: it has value to someone. Key to this concept in a SOA ecosystem is that a resource needs to be identifiable.

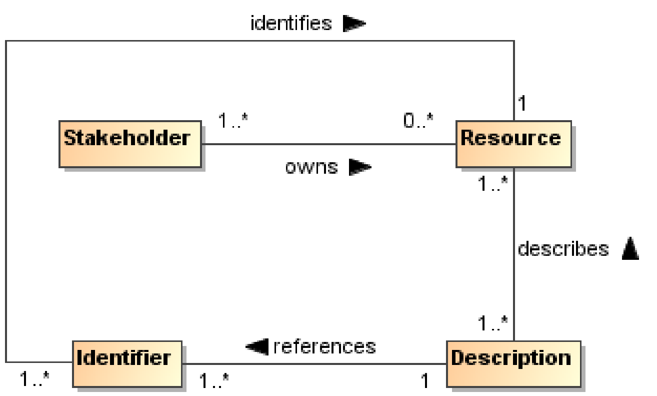

Figure 7 - Resources

Resource~~Compared~~

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

A resource~~facts about~~ may be identifiable by different methods but within a SOA ecosystem a resource must have at least one well-formed identifier that may be unambiguously resolved to the ~~natural world, social facts~~intended resource.

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

Identifier

~~inherently abstract: they only have meaning in the~~ An identifier is any sequence of characters that may be unambiguously resolved to identifying a particular resource.

Identifiers typically require a context ~~of a social structure.~~in order to establish the connection with the resource. In a SOA ecosystem, it is good practice to use globally unique identifiers; for example globally unique Internationalized Resource Identifiers (IRIs).

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

The ability to identify a resource is important in interactions to determine such things as rights and authorizations, to understand what functions are being performed and what the results mean, and to ensure repeatability or characterize differences with future interactions. Many interactions within a SOA

 Copyright © OASIS Open 2011. All Rights Reserved.

ecosystem take place across ownership boundaries and the combination of interactions can be unpredictable. Identifiers provide the means for all resources important to a given SOA system to be *unambiguously* identifiable at any moment and in any interaction.

3.1.3.2 Ownership

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

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

Ownership

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

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

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

A crucial property that distinguishes ownership from a more limited ***right to use*** is the right to transfer rights and responsibilities totally and irrevocably to another stakeholder. When a stakeholder uses a resource but does not own the resource, that stakeholder may not transfer the right to use the resource to a third stakeholder. The owner of the resource maintains the rights and responsibilities of being able to authorize other stakeholders to use the owned resource.

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

Ownership Boundary

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

In a SOA ecosystem, providers and consumers of services may be, or may be acting on behalf of, different owners, and thus the interaction between the provider and the consumer of a given service will necessarily cross an ownership boundary. It is important to identify these ownership boundaries in a SOA ecosystem, as successfully crossing them requires the elements identified in the following sections be addressed. Addressing the elements identified in the following sections is referred to in the OASIS SOA RM as establishing the execution context.

3.1.4 Trust and Risk

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

 Copyright © OASIS Open 2011. All Rights Reserved.

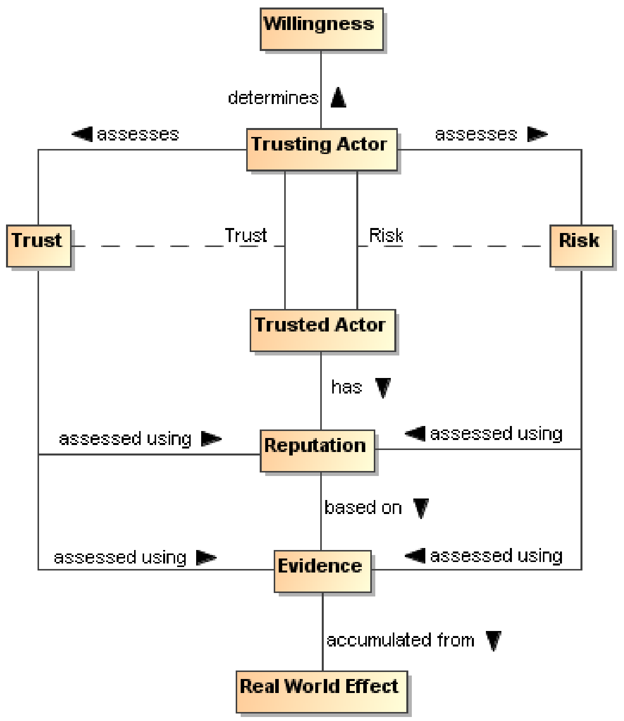

Figure 8 - Willingness and Trust

Willingness

> Willingness is the internal commitment of a human actor to carry out its part of an interaction.

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

Trust

> Trust is a private assessment or internal perception of one actor that another actor will perform actions in accordance with an assertion regarding a desired real world effect.

Risk

> Risk is a private assessment or internal perception of the likelihood that certain undesirable real world effects will result from actions taken and the consequences or implications of such.

Trust is involved in all interactions – it is necessary for *all* participants (consumers, providers, mediators) involved in a given interaction to trust all involved actors, at least to the extent required for continuance of the interaction. The degree and nature of that trust is likely to be different for each actor, most especially when those actors are in different ownership boundaries.

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

3.1.4.1 Assessing Trust and Risk

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

 Copyright © OASIS Open 2011. All Rights Reserved.

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

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

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

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

3.1.5 Policies and Contracts

As noted in the Reference Model, a **policy** represents some commitment and/or constraint advertised and enforced by a stakeholder and that stakeholder alone. A **contract**, on the other hand, represents an agreement by two or more participants. Enforcement of

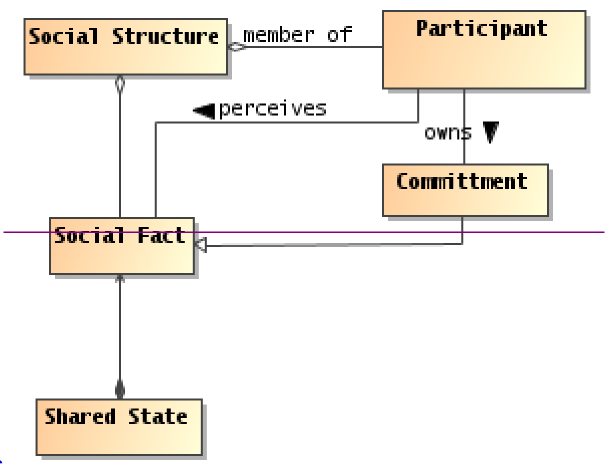

contracts

may or may not be the responsibility of the parties to the agreement but is usually performed by a stakeholder in the ecosystem (public authority, legal system, etc.).

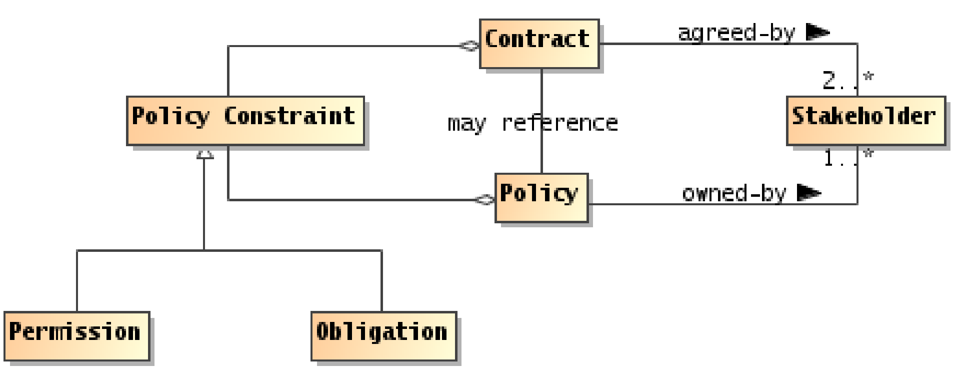

Figure 9 - Policies~~Shared State~~ and ~~Social Facts~~Contracts

Copyright © OASIS Open 2011. All Rights Reserved.

Policy

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

Policies can often be said to be about something – they have an object. For example, there may be policies about the use of a service. Policies have an **owner** – the stakeholder who asserts and takes responsibility for the policy. Note that the policy owner may or may not be the owner of the object of the policy. Thirdly, policies represent constraints – some measurable limitation on the state or behavior of the object of the policy, or of the behavior of the stakeholders owning the policy.

Contract

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

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

Both policies and contracts imply a desire to see constraints respected and enforced. Policies are owned by service providers – individual (or aggregate) stakeholders – and contracts are owned by both service providers and consumers – the parties to the contract; these stakeholders are responsible for ensuring that any constraints in the policy~~Shared State~~

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

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

~~**Social Fact**~~

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

~~Social facts typically require some kind of ritual to establish: the action itself is physical, its~~ or contract are enforced, although the actual enforcement may be delegated to a different mechanism. A contract does not necessarily oblige the contracting parties to act (for example to use a service) but it does constrain how they act if and when the condition covered by the contract occurs (for example, when a service is invoked and used).

Communication ~~interpretation is social. For example, the existence of an agreed contract typically requires both parties to sign papers and to exchange those papers. If the signatures are not performed correctly, or if the parties are not properly empowered to perform the ritual, then it is as though nothing happened.~~

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

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

~~**Commitment**~~

 Copyright © OASIS Open 2011. All Rights Reserved.

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

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

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

~~3.4 Acting in a Social Context~~

~~3.4.1 Actions, Real World Effect and Events~~

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

~~Action~~

~~Action is the application of intent by a participant (or agent) to achieve a real world effect.~~

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

~~An action may have preconditions where a precondition is something that needs to be in place before an action can occur, e.g. confirmation of a precursor action. One important class of such preconditions are the conditions associated with security: authentication and authorization of the participants attempting actions.~~

~~shows a model of how actions are associated with agents that perform actions, the results of performing actions and how actions are associated with intention.~~

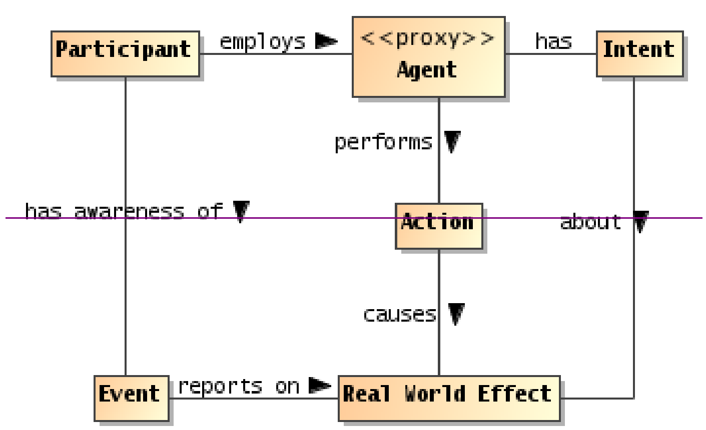

~~*Figure Actions, Real World Effect and Events Model*~~

~~Real World Effect~~

~~A Real World Effect is the changes in the state of the world as a result of a participant performing an action in response to a service interaction.~~

~~The result of performing an action is, in the expected case, something changes in the world. This is the Real World Effect of performing the action. Many, if not most, instances of Real World Effect involve acting in the context of a social structure; i.e., the effect desired is the establishment of one of more social facts.~~

~~Changes in the world can be *reported* by means of events:~~

 Copyright © OASIS Open 2011. All Rights Reserved.

~~Event~~

> ~~An event is an occurrence that at least one participant has an interest in being aware of.~~

~~In the case of this Reference Architecture, a key class of events is that which reflects the effects of actions that have been performed – i.e., we are especially interested in events that report on Real World Effects of actions.~~

~~In effect, an event is the corollary to action: in a public arena, joint actions result in changes to the world; these changes are manifested as events that participants in the arena have an awareness of.~~

~~A key feature of action that distinguishes it from mere force or accident is that someone or something intended the action to occur. Intent represents an agent's relationship to one or more of its goals:~~

~~Intent~~

> ~~Intent is the relationship between an agent and its goals that signifies a commitment by the agent to achieve that goal.~~

~~An agent's intent in performing an action is to further one or more of the agent's goals.~~

~~3.4.2 Social Actions~~

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

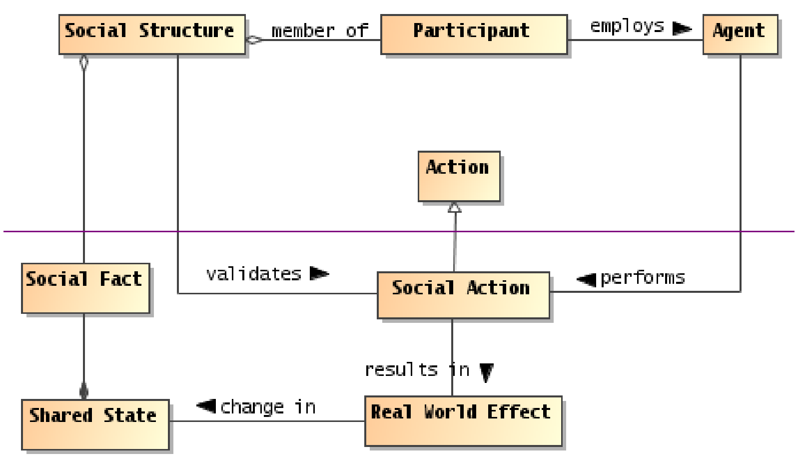

~~Figure Acting within Social Structures~~

~~Social Action~~

> ~~A social action is an action which is defined primarily by the effect it has on the relationship between participants and state of a social structure by establishing one or more new social facts. A social action consists of a physical action together with an appropriate authority.~~

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

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

~~3.4.3 Interaction as Joint Action~~

~~When participants interact with services they are conducting actions that are inherently collaborative and joint in nature: there is no dance without a partner.~~

Copyright © OASIS Open 2011. All Rights Reserved.

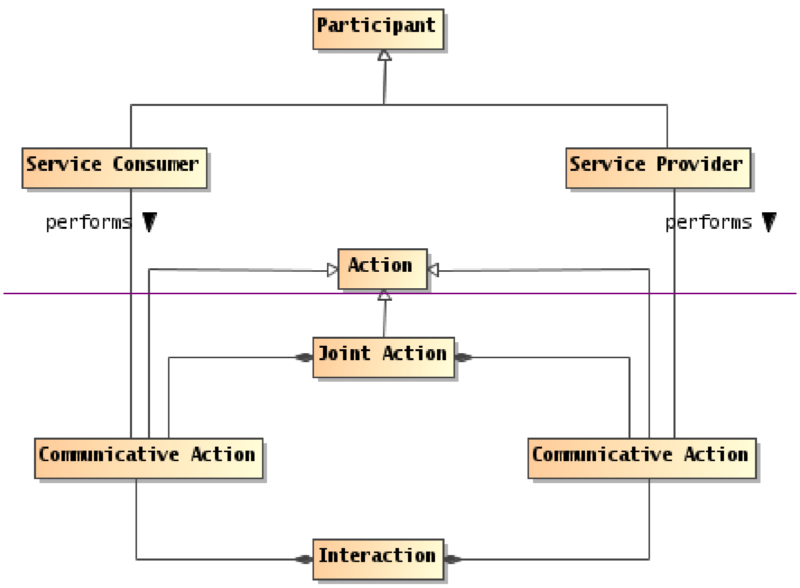

~~3.4.3.11.1.1.1 Figure Service Interaction as Joint Action~~

~~Every action that Is part of an interaction between a service consumer and a service is inherently a *joint action* – involving both participants. Just as action is the foundation of an individual's actions in the context of SOA-based systems, interactions are characterized by joint actions:~~

~~**Joint Action**~~

> ~~A joint action is an action involving the efforts of two or more participants to achieve a real world effect.~~

~~Joint actions are actions that inherently require two or more participants in order to properly relate the activities to the participants' intentions. Typically, a joint action involves two participants in communicative actions – one participant speaking and the other listening.~~

~~Joint actions are the foundation for understanding interaction between participants in a SOA-based system. It is not possible for there to be interaction between service providers and consumers without the participants engaging in a series of joint actions – typically joint communicative actions.~~

~~3.4.4 Semantics of Communication Model~~

~~Interaction is a form of communication. In this Reference Architecture, we use *messages* as the medium of interaction between service participants. Messages are exchanged that represent actions, and messages are exchanged that represent the reporting of events. In this model, we outline one way that this can be modeled effectively – in terms of shared vocabularies, shared semantics and shared understanding of communicated intent.~~

~~Since service consumers and providers are not directly acting against each other, they must do so indirectly – primarily by means of some form of communication. Speaking to someone is an action; if the speech conveys a request or a pronouncement of some kind, the former actions are used as vehicles to convey the true actions. Thus in , we see **Action** appear twice – once in modeling the communicative actions needed to support interaction and once as the intended or conveyed action.~~

 Copyright © OASIS Open 2011. All Rights Reserved.

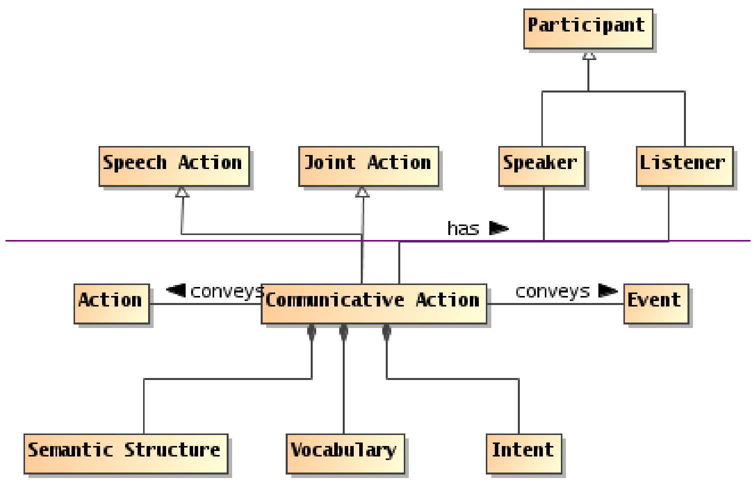

~~*Figure Semantics of Communication Model*~~

~~Communicative Action~~

~~3.4.5~~3.1.6

Communication

A communication is a process of reaching mutual understanding, in which participants not only exchange information as messages but share the meaning of this information.

A communication~~Communicative actions are joint actions where service participants communicate with each other. A Communicative Action has a speaker and a Listener; each of whom~~ involves at least one actor in the role of **sender** and at least one other actor in the role of **recipient.**All actors must perform their ~~part for the communicative action to occur.~~

~~**Semantic Structure**~~

~~A communicative action has an aspect which conveys the meaning of the content being communicated. Typically, a semantic structure takes the form of a proposition which is either true, false or intended to be true or false.~~

~~The concept of semantic structure is quite abstract. However, in many cases involving machines, the semantic structure will be conveyed as some form of highly regular tree structure, with a well defined method for interpreting the structure. For example, an invoice will often follow pre-established standards for communicating invoices.~~

~~**Intent**~~

~~The purpose of the communicative action is its **intent**. The intent, together with the semantic structure convey either an action – such as a request from a service consumer to the service – or an event – which typically reports on the results of previous communicative acts.~~

~~**Vocabulary**~~

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

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

 Copyright © OASIS Open 2011. All Rights Reserved.

A communication involves a message, which an actor receiving must be able to correctly interpret. The extent of that correct interpretation depends on the role of the actor and the purpose of the communication.

A communication is not effective unless the recipient can correctly interpret the message (or at least, that part of it which is relevant to the participant). However, interpretation can itself be characterized in terms of semantic engagement: the proper understanding of a message in a given context.

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

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

3.1.7 Semantics and Semantic Engagement

A SOA ecosystem is a space in which actors need to share understanding[10] as well as sharing actions. Indeed, such shared understanding is a pre-requisite to a joint action being carried out as intended. It is vital to a trusted and effective ecosystem. Semantics are therefore pervasive throughout SOA ecosystems and important in communicative actions described above, as well as a driver for policies and other aspects of the ecosystem.

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

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

Semantic Engagement

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

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

The impact of any assertion can only be fully understood in terms of specific social contexts that necessarily include the actors that are involved. For example, a policy statement that governs the actions

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

 Copyright © OASIS Open 2011. All Rights Reserved.

relating to a particular resource may have a different impact or purpose for the participant that owns the resource than for the actor that is trying to access it: the former understands the purpose of the policy as a statement of enforcement - the latter understands it as a statement of constraint.

3.2 Action in a SOA Ecosystem Model

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

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

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

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

Also, it is important to note that both actions and RWE are 'fractal' in nature, in the sense that they can often be broken down into more and more granularity depending on how they are examined and what level of detail is important.

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

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

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

A key aspect of joint action revolves around the trust that both parties must exhibit in order to participate in the joint action. The willingness to act and a mutual understanding of both the information exchanged and the expected results is the particular focus of Sections 3.1.4 and 3.1.7.

3.2.1 Needs, Requirements and Capabilities

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

There is a reason that participants are engaged in this activity: different participants have different **needs** and have or apply different **capabilities** for satisfying them.These are core to the concept of a service.

 Copyright © OASIS Open 2011. All Rights Reserved.

The SOA-RM defines a service as “the mechanism by which needs and capabilities are brought together”. This idea of services being a mechanism “between” needs and capabilities was introduced in order to emphasize capability as the notional or existing business functionality that would address a well-defined need. Service is therefore the *implementation* of such business functionality *such that it is accessible* through a well-defined interface. A capability that is isolated (i.e., it is inaccessible to potential consumers) is emphatically not a service.

Business functionality

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

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

In a SOA context, the consumer (as a stakeholder) expresses a need (“I want to buy a book”) and looks to an appropriate service to fulfill that need and assesses issues such as the trustworthiness, intent and willingness of a particular provider. This ecosystem communication continues up to the point when the consumer is ready to act. The consumer (as an actor now) will then interact with a provider by invoking a service (for example, ordering the book using an online bookseller) and engaging in relevant actions (validating the purchase, submitting billing and delivery details) within the system with a view to achieving the desired Real World Effect (having the book delivered).

Need

~~However, any communication requires the core elements outlined in this model: some form of shared vocabulary, a shared basis for understanding communications, and a shared basis for establishing the intentions of participants.~~

~~3.4.6 Transactions and Exchanges Model~~

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

~~**Business Transaction**~~

> ~~A business transaction is a joint action engaged in by two or more participants in which the real world effect is an increase in apparent value to the participants.~~

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

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

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

~~**Business Agreement**~~

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

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

~~From an architectural perspective, the business transaction often represents the top-most mode of interpretation of service interactions. When participants interact in a service, they exchange information~~

 Copyright © OASIS Open 2011. All Rights Reserved.

~~and perform actions that have an effect in the world. These exchanges can be interpreted as realizing part of, and in support of, business transactions.~~

~~**Business Process**~~

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

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

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

> ~~Process~~ A need is a general statement expressed by a stakeholder of something deemed necessary. It may be formalized as one or more **requirements** that must be fulfilled in order to achieve a stated goal.

Requirement

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

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

~~Choreography~~

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

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

Capability

> A capability is an ability to deliver a real world effect.

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

3.2.2 Services Reflecting Business

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

Business solution

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

Composability

> **Composability** is the ability to combine individual services, each providing defined business functionality, so as to provide more complex business solutions.

To achieve composability, capabilities must be identified that serve as building blocks for business solutions. In a SOA ecosystem, these building blocks are captured as services representing well-defined business functions, operating under well-defined policies and other constraints, and generating well-defined real world effects. These service building blocks should be relatively stable so as not to force

 Copyright © OASIS Open 2011. All Rights Reserved.

repeated changes in the compositions that utilize them, but should also embody SOA attributes that readily support creating compositions that can be varied to reflect changing circumstances.

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

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

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

3.2.3 Action, Communication and Joint Action

In general terms, entities act in order to achieve their goals. However, the form of action that is of most interest within a SOA ecosystem is that involving interaction across ownership boundaries (between more than one actor) – **joint action.**

3.2.3.1 Action and Actors

Action

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

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

Objectives refer to real world effects that participants believe are achievable by a specific action or set of actions that deliver appropriate changes in shared state. In contrast, a goal is not expressed in terms of specific action but rather in terms of desired end state.

For example, someone may wish to have enough light to read a book. In order to satisfy that goal, the reader walks over to flip a light switch. The *objective* is to change the state of the light bulb, by turning on the lamp, whereas the *goal* is to be able to read. The *real world effect* is more light being available to enable the person to read.

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

Real World Effect

> A real world effect is a measurable change to the shared state of pertinent entities, relevant to and experienced by specific stakeholders of an ecosystem.

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

3.2.3.2 Communication and Joint Actions

In this Reference Architecture Foundation, we are concerned with two levels of activity: as communication and as participants engaged in joint actions to use and offer services.

 Copyright © OASIS Open 2011. All Rights Reserved.

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

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

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

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

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

Joint Action

> Joint action is the coordinated set of actions involving the efforts of two or more actors to achieve an effect.

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

Different viewpoints lead to either communication or joint action as being considered most important. For example, from the viewpoint of ecosystem security, the integrity of the communications may be dominant; from the viewpoint of ecosystem governance, the integrity of the joint action may be dominant.

3.2.4 State, Shared State and Real-World Effect

State

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

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

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

In a SOA ecosystem, there is a distinction between the set of facts about an entity that only that entity can access – the so-called Private State – and the set of facts that may be accessible to other actors in the SOA-based system – the public or Shared State.

Private State

> The private state is that part of of an entity’s state that is knowable by, and accessible to, only that entity.

Shared State

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

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

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

 Copyright © OASIS Open 2011. All Rights Reserved.

3.3 Architectural Implications

3.3.1 Social structures

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

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

3.3.2 Resource and Ownership

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

- Mechanism for assigning and guaranteeing uniqueness of globally unique identifiers
- Identifying the extent of the enterprise over which the identifier needs to be understandable and unique
- Mechanism and framework for ensuring the long-livedness of identifiers (i.e., they cannot just change arbitrarily)

3.3.3 Policies and Contracts

- Policies are constraints
  - Policies MUST be expressed
  - Constraints MUST be enforceable
  - Manage,emt of potentially large numbers of policies MUST be achievable
- Policies have owners
  - Policies SHOULD be established by social structures.
- Policies may not be consistent with one another
  - Policy conflict resolution techniques MUST exist and be in place
- Agreements are constraints agreed to
  - Contracts SHOULD be enforced by mechanisms of the social structure

3.3.4 Communications as a Means of Mediating Action

Using message exchange for mediating action implies

- Ensuring correct identification of the structure of messages:
  - Identifying the syntax of the message;
  - Identifying the vocabularies used in the communication
  - Identifying the higher-level structure of the communication, such as policy assertion, contract enforcement, etc.
- A principal objective of communication is to mediate action

 Copyright © OASIS Open 2011. All Rights Reserved.

    - Messages convey actions and events
    - Receiving a message is an action, but is not the same action as the action conveyed by the message
    - Actions are associated with objectives of the actors involved
        - Explicit representation of objectives may facilitate automated processing of messages
    - An actor agreeing to adopt an objective becomes responsible for that objective

3.3.5 Semantics

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

The operation of the SOA ecosystem is significantly enhanced if

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

3.3.6 Trust and Risk

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

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

3.3.7 Needs, Requirements and Capabilities

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

- Architecture needs to
    - Take into account existing capabilities available as services

3.3.8 The Importance of Action

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

Any architectural realization of a SOA ecosystem should address:

- How actions are modeled:
    - Identifying the performer or agent of the action;
    - the target of the action; and the
    - verb of the action.

Any explicit models of joint action should take into account

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

 Copyright © OASIS Open 2011. All Rights Reserved.

4 Realization of a SOA Ecosystem view

~~3.4.71.1.1 Roles in Social Structures~~

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

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

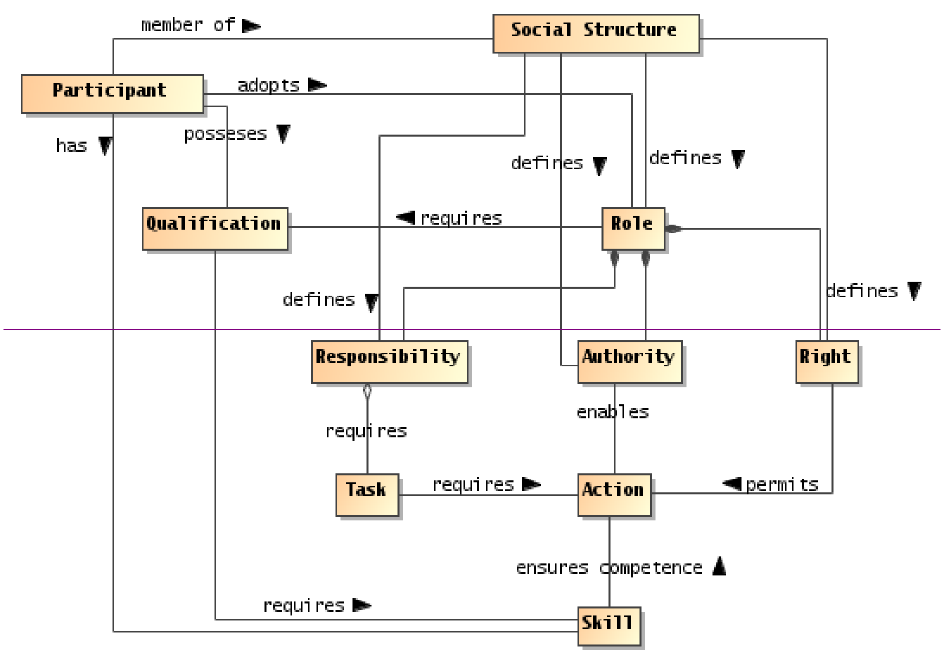

~~*Figure Roles, Rights and Responsibilities Model*~~

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

~~**Role**~~

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

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

~~A participant may adopt one or more roles; and have zero or more skills and qualifications. For example, a participant adopting the role of secretary of a standards group is obliged to ensure that all the minutes of the various meetings are properly recorded; and members of certain standards groups are obliged to declare any pre-existing IP claims that may be relevant to the work of the groups.~~

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

Copyright © OASIS Open 2011. All Rights Reserved.

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

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

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

~~Right~~

> ~~A right is a predetermined permission that permits an agent to perform some action or adopt a stance in relation to the social structure and other agents. For example, in most circumstances, sellers have a right to refuse service to potential customers; but may only do so based on certain criteria.~~

~~Authority~~

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

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

~~Responsibility~~

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

~~Skill~~

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

~~Qualification~~

> ~~A qualification is a public determination by an issuing authority that a stakeholder has achieved some state. The issuing authority may require some successful actions on the part of the stakeholder (such as demonstrating some skills). The qualification may have constraints attached to it; for example, the certification may be time limited.~~

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

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

~~3.5 Governance and Social Structures~~

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

 Copyright © OASIS Open 2011. All Rights Reserved.

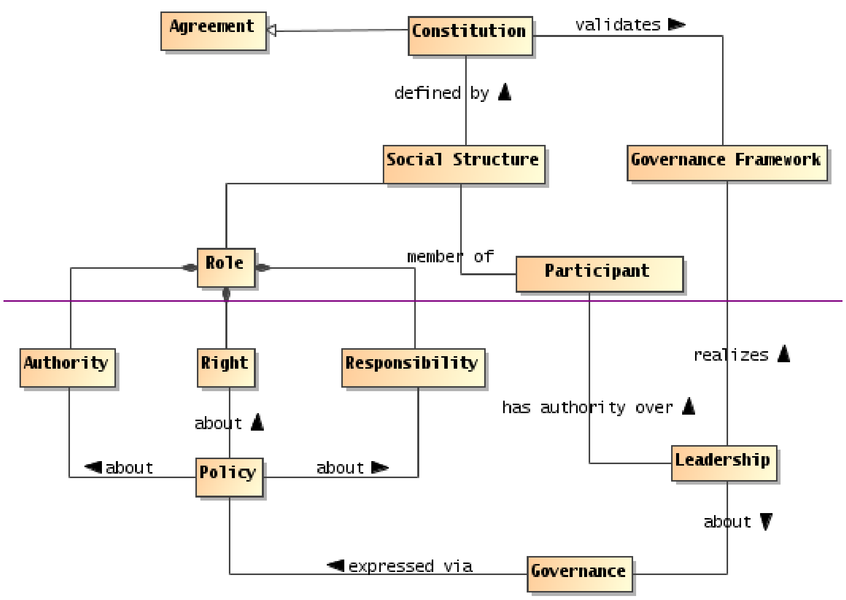

~~*Figure Social Structures and Governance*~~

~~Both of these are aspects of the governance of Service Oriented Architecture.~~

~~The key elements of our model that relate to governance are the constitution of the social structure, the policies of the social structure, authority in a social structure, and the associated mechanisms of enforcement.~~

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

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

~~3.6 Proposition Model~~

~~The Reference Architecture makes use of descriptions of entities and states in the world. For example, we talk about a need being satisfied in Section , a policy being enforced in Section a service description in Section .~~

~~In order to be able to relate a description with the entity that it being described we need the description to be verifiable relative to the entity. The proposition model identifies the key components that can support the verifiability of descriptions.~~

~~**Proposition**~~

Copyright © OASIS Open 2011. All Rights Reserved.

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

~~In principle, the truth of a proposition must be verifiable – using a decision procedure – by examining the world and checking that the proposition and the world are consistent with each other.[11]~~

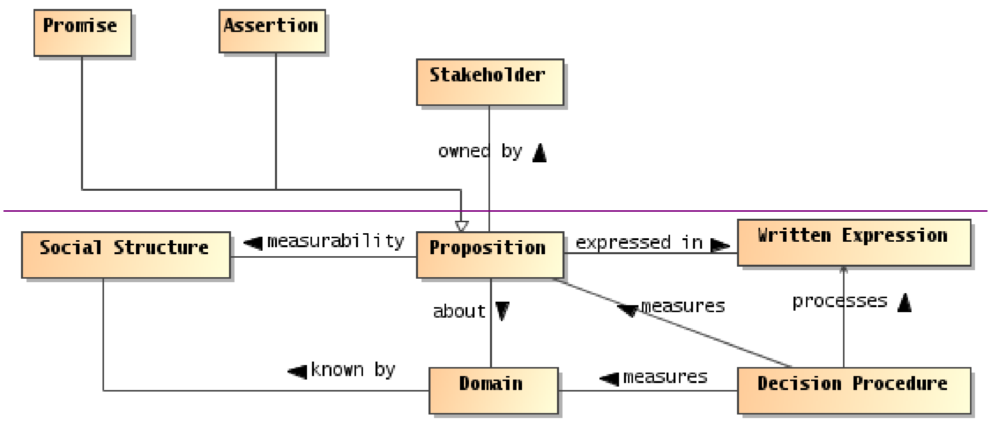

~~Figure Propositions~~

~~Decision Procedure~~

> ~~A process for determining whether an expression is true, or is satisfied, in the world.~~

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

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

~~Note that not all description languages have a decision procedure. However, for the uses to which we put the concept of proposition: policies, service descriptions, and so on, we require that the descriptive language have a decision procedure.~~

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

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

 Copyright © OASIS Open 2011. All Rights Reserved.

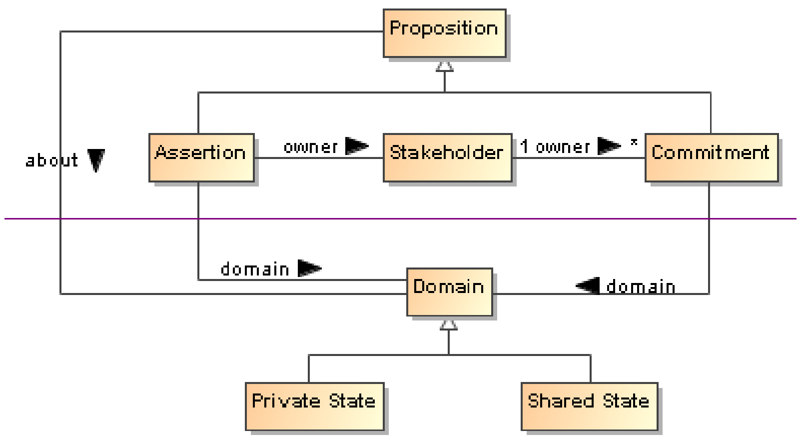

~~*Figure Assertions and Promises*~~

~~**Assertion**~~

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

~~**Promise**~~

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

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

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

 Copyright © OASIS Open 2011. All Rights Reserved.

~~4 Realizing Service Oriented Architectures View~~

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

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

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

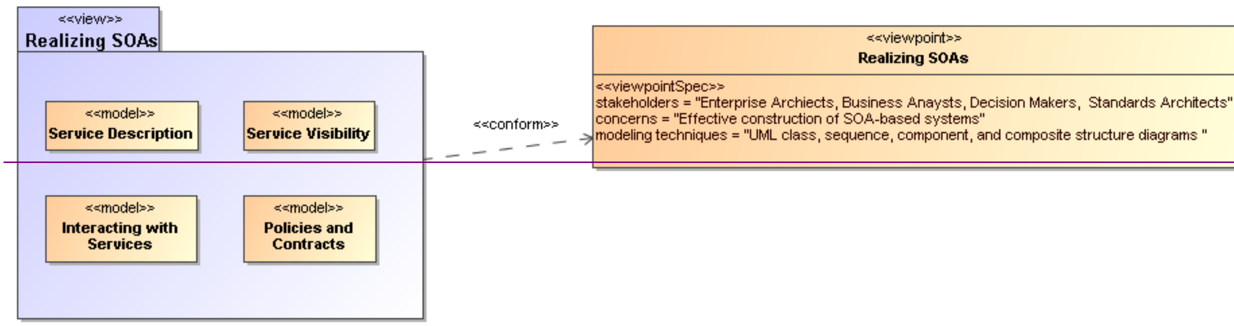

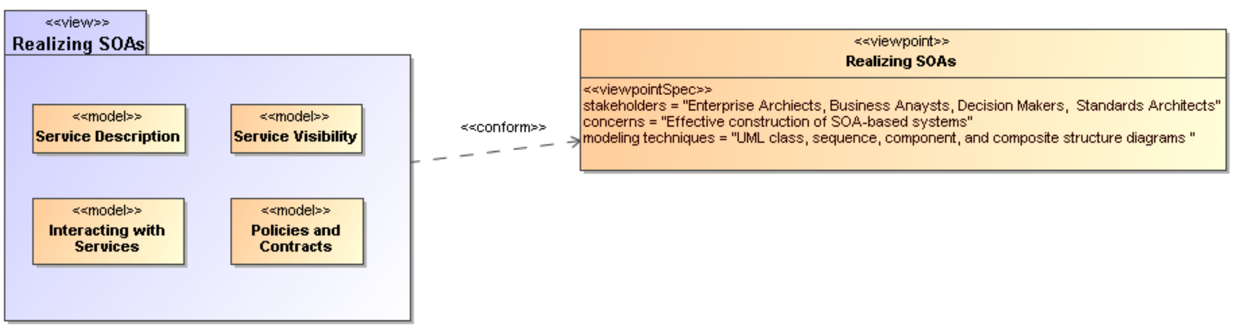

Figure 10 - Model Elements Described in the ~~Realizing~~Realization of a SOA Ecosystem view

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

4.1 Service Description Model

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

 Copyright © OASIS Open 2011. All Rights Reserved.

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

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

- ~~SOA-RM~~The Reference Model states that one of the hallmarks of SOA is the large amount of associated description. The model presented below focuses on the description of services but it is equally important to consider the descriptions of the consumer, other participants~~participants,~~, and needed resource~~resources~~s other than services.
- Descriptions are inherently incomplete but may be determined as *sufficient* when it is possible for the participants~~participants~~ to access and use the described services based only on the descriptions provided. This means that, at one end of the spectrum, a description along the lines of “That service on that machine” may be sufficient for the intended audience. On the other extreme, a service description with a machine-process-able description of the semantics of its operations and real world effect~~real world effect~~s may be required for services accessed via automated service discovery and planning systems.
- Descriptions come with context, i.e. a given description comprises information needed to adequately support the context. For example, a list of items can define a version of a service, but for many contexts an indicated version number is sufficient without the detailed list. The current model focuses on the description needed by a service consumer~~will~~ to understand what the service does, under what conditions he service will do it, how well the service does it, and what steps are needed by the consumer to initiate and complete a service interaction. Such information also enables the service provider to clearly specify what is being provided and the intended conditions of use.
- Descriptions change over time as, for example, the ingredients and nutrition information for food labeling continues to evolve. A requirement for transparency of transactions may require additional description for those associated contexts.
- Description always proceeds from a basis of what is considered "common knowledge". This may be social conventions that are commonly expected or possibly codified in law. It is impossible to describe everything and it can be expected that a mechanism as far reaching as SOA will also connect entities where there is inconsistent "common" knowledge.
- Descriptions ~~will~~ become the collection point of information related to a service or any other resource~~resource,~~, but it ~~will~~is not necessarily~~ be~~ the originating point or the motivation for generating this information. In particular, given a SOA service as the access to an underlying capability, the service may point to some of the capability’s previously generated description, e.g. a service providing access to a data store may ~~reference update records that indicate the freshness of the data. As another example, it is more maintainable for description to reference the information maintained by an individual who is designated a Responsible Party (see Section) than to require the update of every instance where the individual is so designated~~also have access to information indicating the freshness of the data.

- ~~Descriptions of the provider and consumer are the essential building blocks for establishing the execution context of an interaction.~~

These points emphasize that ~~descriptions are assembled with respect to some context and~~ there is no one “right” description for all contexts and for all time. Several descriptions for the same subject may exist at the same time, and this emphasizes the importance of the description referencing source material maintained by that material’s owner rather than having multiple copies that become out of synch and inconsistent.

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

This Reference Architecture Foundation uses the term service description for consistency with the concept defined in ~~SOA-RM.~~the Reference Model. Some ~~of the current~~ SOA literature ~~speaks to~~ treats the idea of a ~~"~~“service contract~~"~~” as ~~effectively the~~ equivalent~~, although the details of what comprises the~~ to service description~~/contract may vary. The~~. In the SOA-RAF, the term service description is preferred ~~because policies are an element of~~. Replacing the term “service description~~ for any resource and~~” with

 Copyright © OASIS Open 2011. All Rights Reserved.

the ~~agreement on policies between service participants may be thought of as a contract. Saying ~~term “service contract~~ for the service description”~~ implies that just one side of the interaction is governing and misses the point that a single set of policies identified by a service description may lead to numerous contracts~~contracts,~~, i.e. service level agreements, leveraging the same description~~. Indeed, these agreements establish the execution context of the service interaction and are not a fundamental attribute of the service itself~~.

4.1.1 The Model for Service Description

Figure 11 shows Service Description ~~modeled ~~as a subclass of the general Description class, where Description is a subclass of the resource~~Resource~~ class as defined in ~~section~~Section 3.1.3.1. In addition, each resource~~Resource~~ is assumed to have a description. The following section discusses the relationships among elements of general description and the subsequent sections focus on service description~~ itself. Note, other~~. Other descriptions, such as those of participants~~participants,~~, are important to SOA but are not individually elaborated in this document.

4.1.1.1 ~~Model ~~Elements Common to General Description

The general Description class is composed of a number of elements that are expected to be common among all ~~specialized ~~descriptions supporting a service oriented architecture.

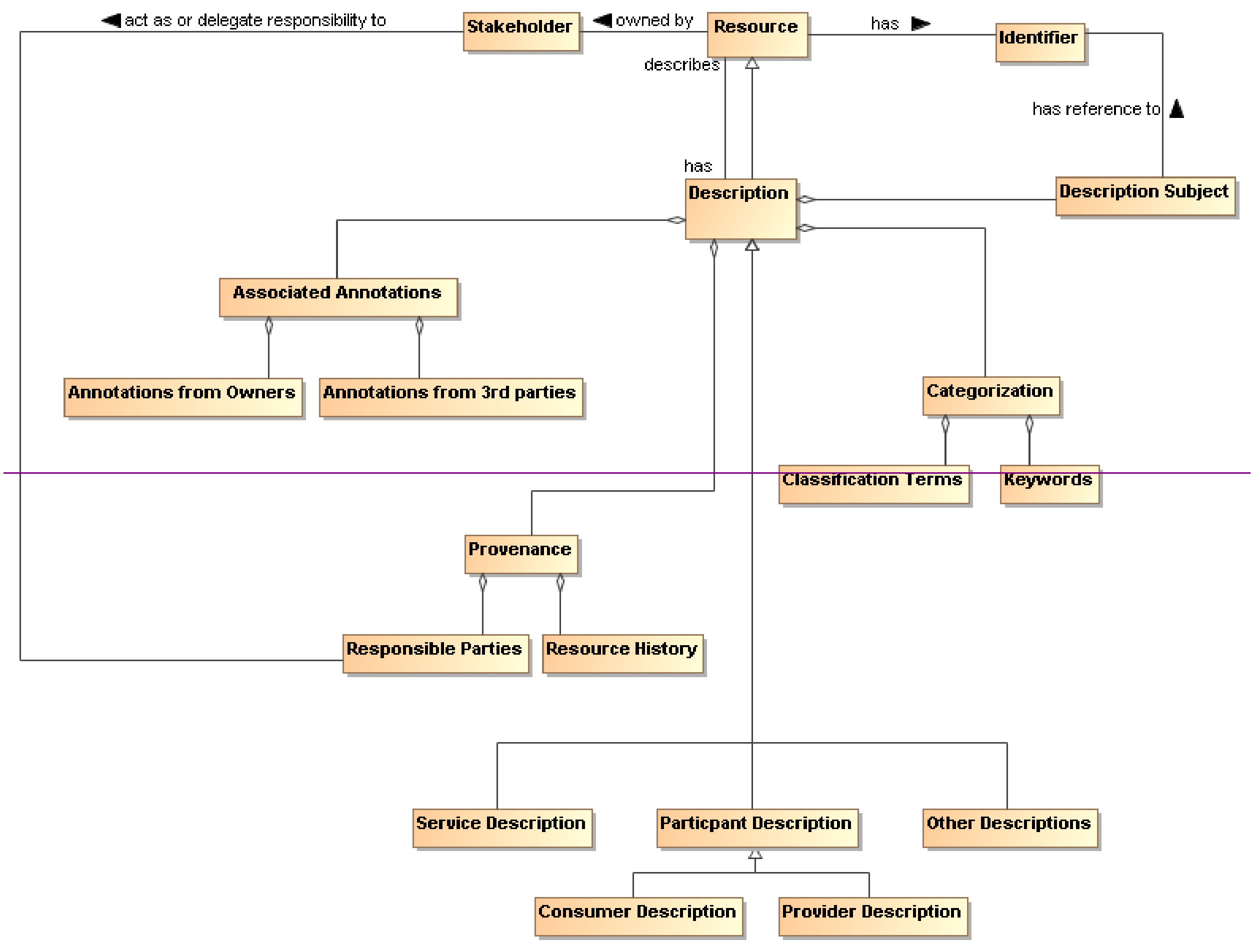

~~*Figure General Description Model*~~

~~4.1.1.1.1 Description Subject~~

~~The subject of a description is a Resource. The value assigned to the Description Subject class may be of any form that provides understanding of what constitutes the Resource, but it is often in human-~~

 Copyright © OASIS Open 2011. All Rights Reserved.

~~readable text. The Description Subject MUST also reference the Resource~~ A registry often contains a~~Identifier~~ ~~of the resource it describes so it can unambiguously identify the subject of each description instance.~~

~~As a Resource, Description also has an identifier with a unique value for each description instance. The description instance provides vital information needed to both establish visibility of the resource and to support its use in the execution context for the subsequent interaction. The identifier of the description instance allows the description itself to be referenced for discussion, access, or reuse of its content.~~ ~~While some~~ subset of the description instance ~~may be entered in a registry to support mediated~~, where the chosen subset is identified as that which facilitates discovery ~~of the description subject, the entire description instance will provide the~~. Additional information contained in a more complete description may be needed to initiate and continue interaction ~~with the subject~~.

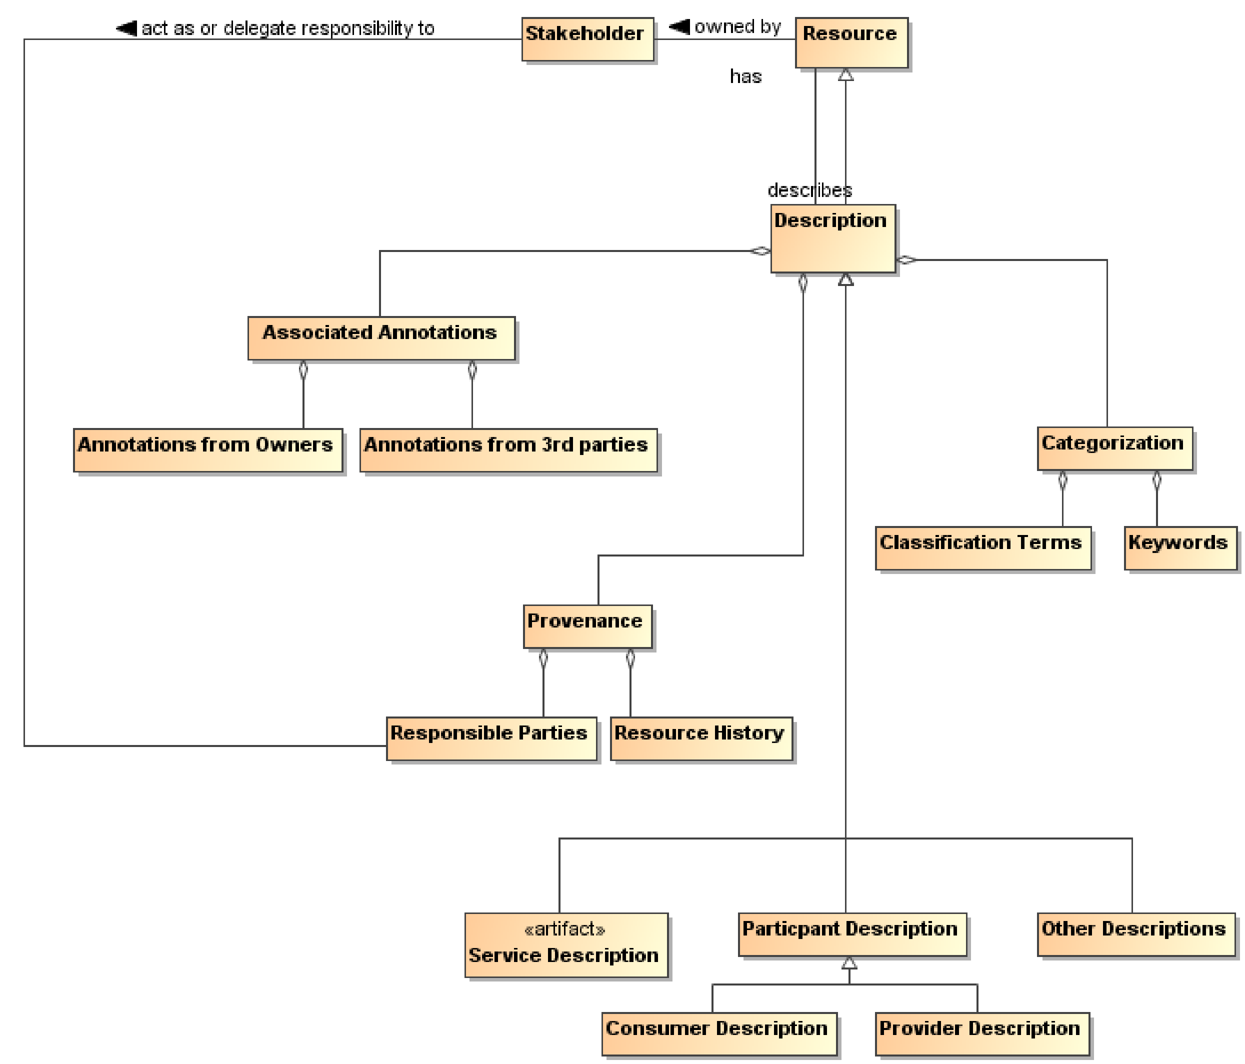

Figure 11 - General Description

~~4.1.1.1.2~~4.1.1.1.1 Provenance

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

 Copyright © OASIS Open 2011. All Rights Reserved.

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

~~4.1.1.1.3~~4.1.1.1.2 Keywords and Classification Terms

A traditional element of description has been to associate the resource~~resource~~ being described with predefined keywords or classification taxonomies that derive from referenceable formal definitions and vocabularies. This Reference Architecture Foundation does not prescribe which vocabularies or taxonomies may be referenced, nor does it limit the number of keywords or classifications that may be associated with the resource~~resource.~~. It does, however, state that a normative definition of any terms or keywords SHOULD be referenced, whether that be a representation in a formal ontology language, a pointer to an online dictionary, or any other accessible source. See Section 0 for further discussion on associating semantics with assigned values.

~~4.1.1.1.4~~4.1.1.1.3 Associated Annotations

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

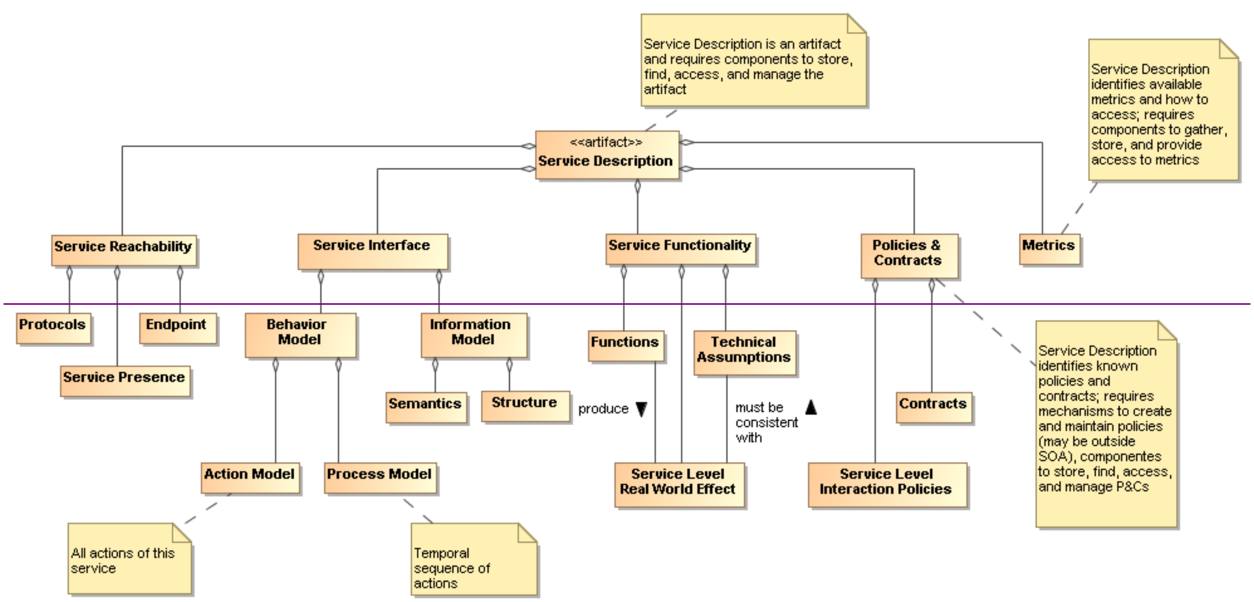

~~*Figure Service Description Model*~~

~~4.1.1.2 Model Elements Specific to Service Description~~

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

Copyright © OASIS Open 2011. All Rights Reserved.

~~4.1.1.2.11.1.1.1.1 Service Interface~~

~~As noted in the Reference Model, the service interface is the means for interacting with a service. For this reference architecture and as shown in Section the service interface will support an exchange of messages, where~~

- ~~the message conforms to a referenceable message exchange pattern (MEP),~~
- ~~the message payload conforms to the structure and semantics of the indicated information model,~~
- ~~the messages are used to invoke actions against the service, where the actions are specified in the action model and any required sequencing of actions is specified in the process model.~~

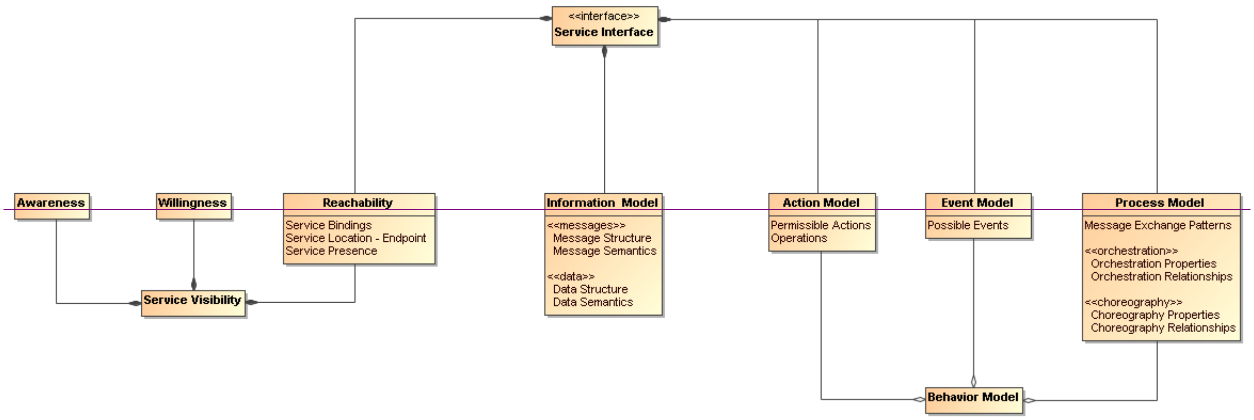

~~Figure Service Interface Model~~

~~4.1.1.2.21.1.1.1.1~~ ~~These aspects of messages are discussed in more detail in~~ Section ~~Service Reachability~~

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

~~In the present context, an endpoint is a referenceable entity, processor, or resource against which one can perform an action.~~[12] ~~As applied in general to an action, the endpoint is the conceptual location where one applies an action; with respect to service description, it is the actual address where a message is sent.~~

[12] This definition of endpoint is consistent with WS-Addressing (http://www.w3.org/TR/2006/REC-ws-addr-core-20060509/) but generalized for any action, not exclusively those implemented as Web Services.

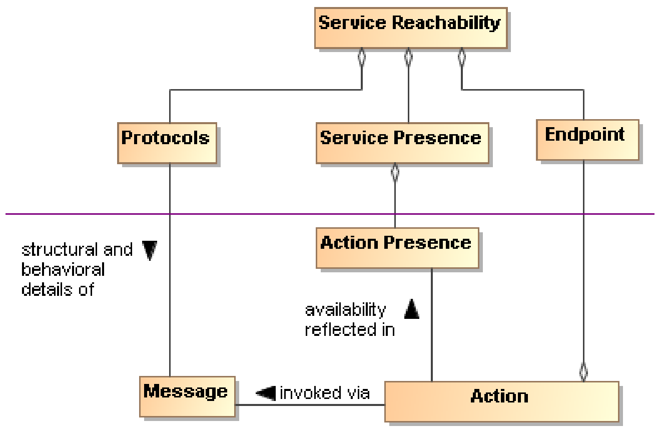

~~*Figure Service Reachability model*~~

~~In addition, the service description should provide information on service presence or on a means of establishing this presence. Presence for either an action or a service may include a static representation of availability or there may be a dynamic means to assess the current availability. The relationship between service presence and the presence of the individual actions that can be invoked is discussed under Establishing Reachability in Section .~~

~~4.1.1.2.3 Service Functionality~~

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

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

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

~~4.1.1.2.4 Policies and Contracts, Metrics, and Compliance Records~~

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

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

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

 Copyright © OASIS Open 2011. All Rights Reserved.

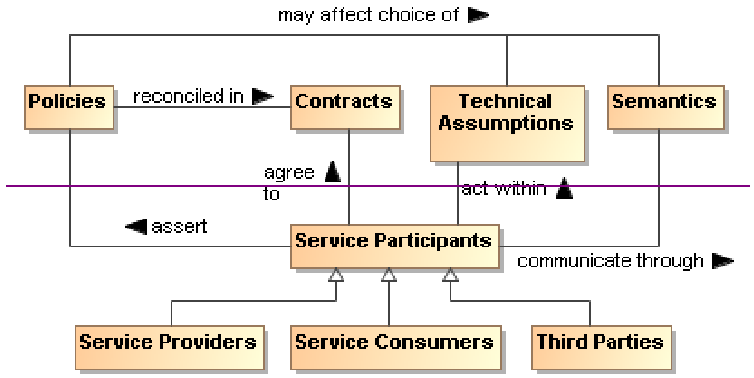

~~*Figure Model for Policies and Contracts as related to Service Participants*~~

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

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

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

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

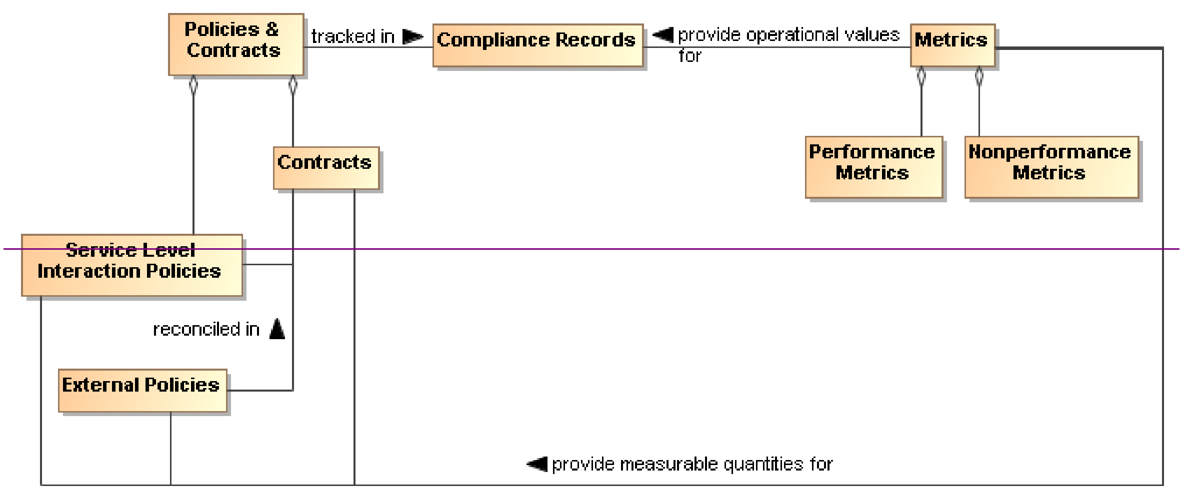

~~*Figure Model relating Policies and Contracts, Metrics, and Compliance Records*~~

~~As with many quantities, the actual performance metrics are not themselves defined by this Service Description because it is not known *a priori* which metrics are being collected by the services, the SOA~~

 Copyright © OASIS Open 2011. All Rights Reserved.

~~infrastructure, or other resources that participate in the SOA interactions. However, the service description SHOULD provide a placeholder (possibly through a link to an externally compiled list) for identifying which metrics are available and how these can be accessed.~~

~~The use of metrics to evaluate compliance is discussed in Section . The results of compliance evaluation SHOULD be maintained in compliance records and the means to access the compliance records SHOULD be included in the Policies and Contracts portion of the service description.~~

~~Note, even though policies are from the perspective of a single participant, policy compliance can be measured and policies may be enforceable even if there is not contractual agreement with other participants. This should be reflected in the policy, contract, and compliance record information maintained in the service description.~~

 Copyright © OASIS Open 2011. All Rights Reserved.

~~4.1.2 Use Of Service Description~~

~~4.1.2.1~~4.1.1.2 Assigning Values to Description Instances

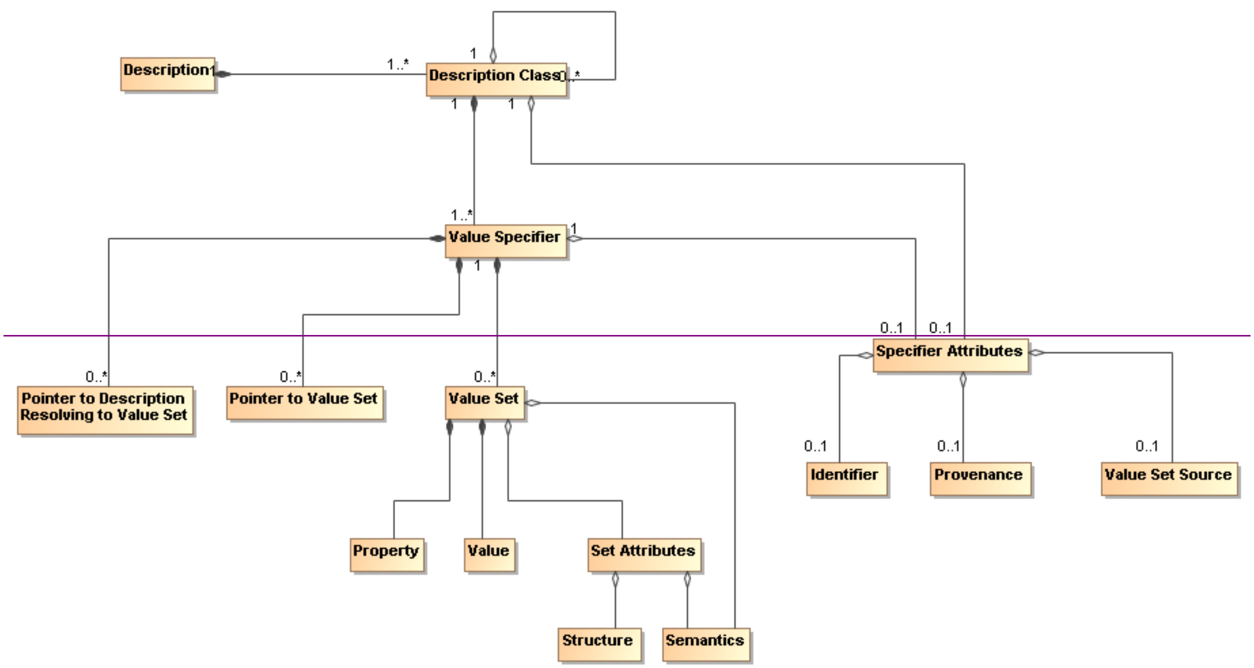

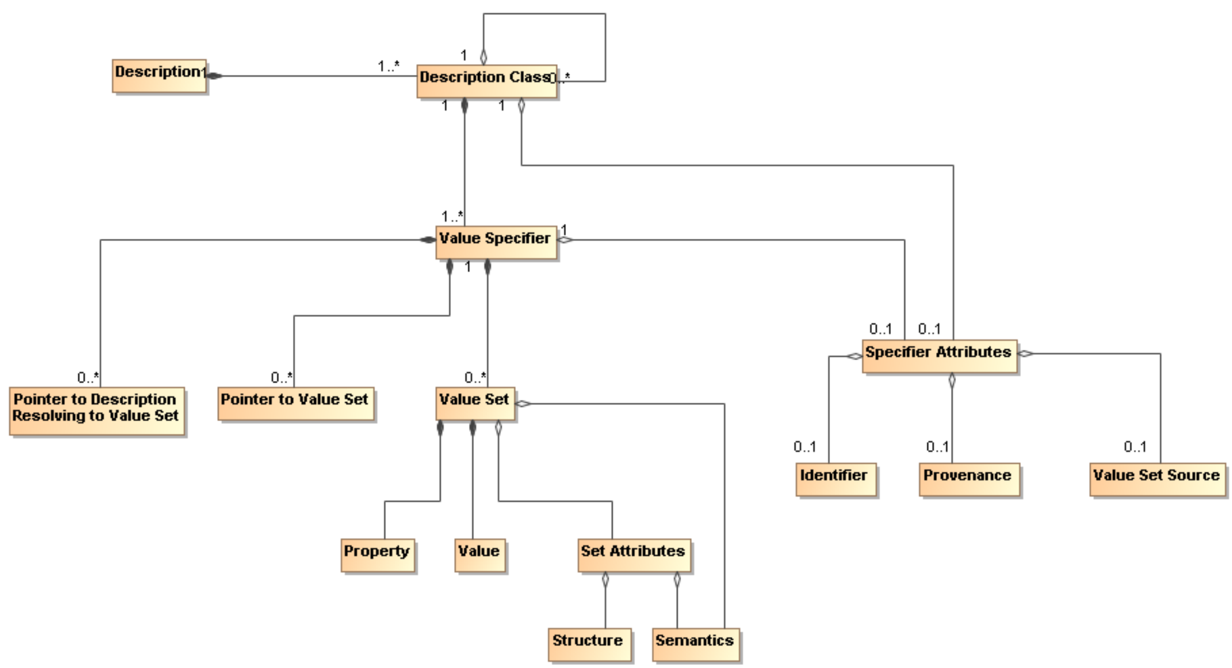

Figure 12 - Representation of a Description ~~Class~~

Figure 11 shows the template for a general description, but individual description instances depend on the ability to associate meaningful values with the identified elements. Figure 12 shows a model for a collection of information that provides for value assignment and traceability for both the ~~value~~ meaning and the source of a value. The model is not meant to replace existing or future schema or other structures that have or will be defined for specific implementations, but it is meant as guidance for the information such structures need to capture to generate sufficient description. It is expected that tools will

 Copyright © OASIS Open 2011. All Rights Reserved.

be developed to assist the user in populating description and ~~autofilling~~auto-filling many of these fields, and in that context, this model provides guidance to the tool developers.

~~For the model in~~In Figure 12, each class ~~is represented by a~~has an associated value specifier or is made up ~~by~~of components that~~ will~~ eventually resolve to a value specifier. For example, Description has several components, one of which is Categorization, which would ~~be represented by a~~have an associated value specifier.

A value specifier consists of

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

The qualifying attributes for the value specifier include

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

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

Note, provenance as a qualifying attribute of a value specifier is different from provenance as part of an instance of Description. Provenance for a service identifies those who own and are responsible for the service, as described in Section 3.1.3. Provenance for a value specifier identifies who is responsible for choosing and assigning values to the value sets that comprise the value specifier. It is assumed that granularity at the value specifier level is sufficient and provenance is not required for each value set.

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

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

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

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

 Copyright © OASIS Open 2011. All Rights Reserved.

4.1.1.3 Model Elements Specific to Service Description

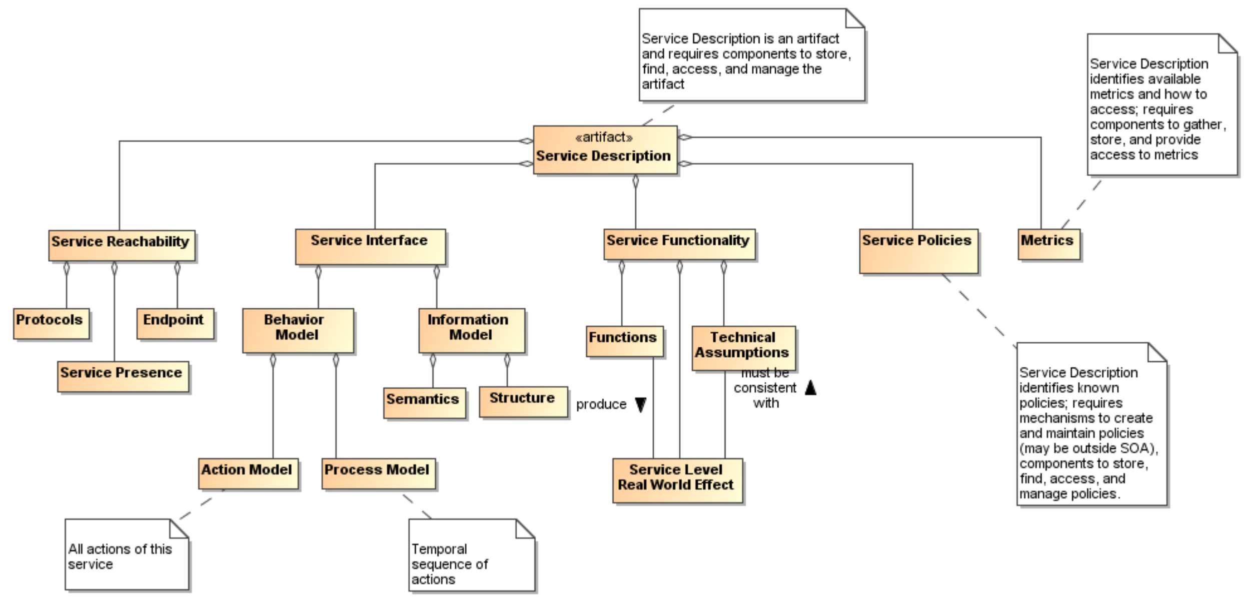

Figure 13 - Service Description

The major elements for the Service Description subclass follow directly from the areas discussed in the Reference Model. Here, we discuss the detail shown in Figure 13 and the purpose served by each element of service description.

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

4.1.1.3.1 Service Interface

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

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

 Copyright © OASIS Open 2011. All Rights Reserved.

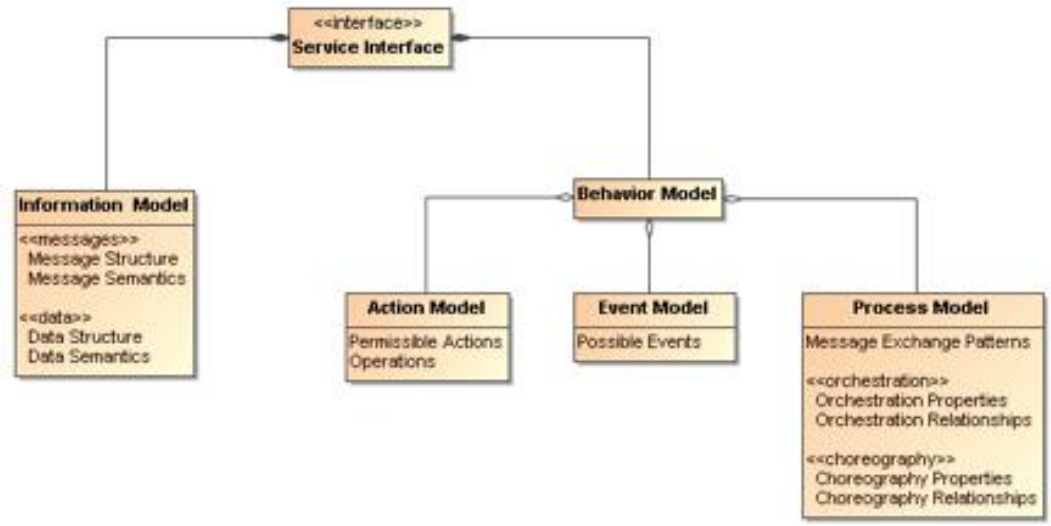

Figure 14 - Service Interface

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

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

4.1.1.3.2 Service Reachability

Service reachability, as modeled in Section 0 enables service participants to locate and interact with one another. To support service reachability, the service description should indicate the endpoints to which a service consumer can direct messages to invoke actions and the protocol to be used for message exchange using that endpoint.

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

4.1.1.3.3 Service Functionality

While the service interface and service reachability are concerned with the mechanics of using a service, service functionality and performance metrics (discussed in Section 4.1.1.3.4) describe what can be expected as a result of interacting with a service. Service Functionality, shown in Figure 13 as part of the overall Service Description model and extended in Figure 15, is a clear expression of service function(s) and the real world effects of invoking the function. The Functions represent business activities in some domain that produce the desired real world effects.

Copyright © OASIS Open 2011. All Rights Reserved.

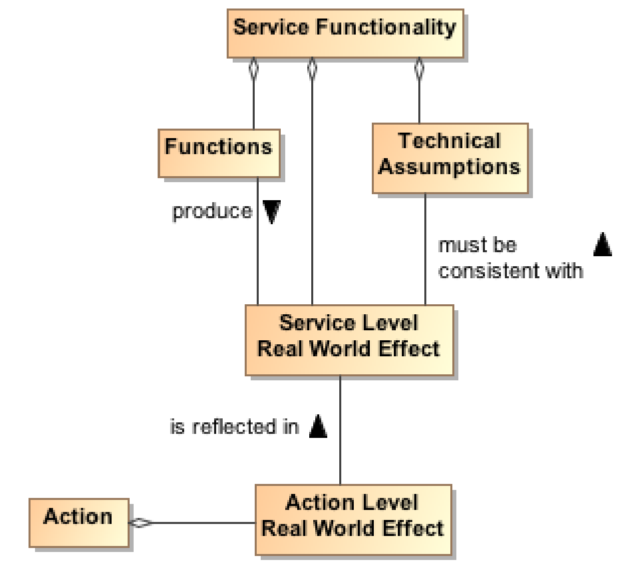

Figure 15 - Service Functionality

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

In Figure 13 and Figure 15, we specifically refer to Service Level and Action Level real world effects.

Service Level Real World Effect

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

Action Level Real World Effect

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

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

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

4.1.1.3.4 Service Policies, Metrics, and Compliance Records

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

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

 Copyright © OASIS Open 2011. All Rights Reserved.

Policies may also be asserted by other service participants, as illustrated by the model shown in Figure 16. Policies that are generally applicable to any interaction with the service are asserted by the service provider and included in the Service Policies section of the service description.

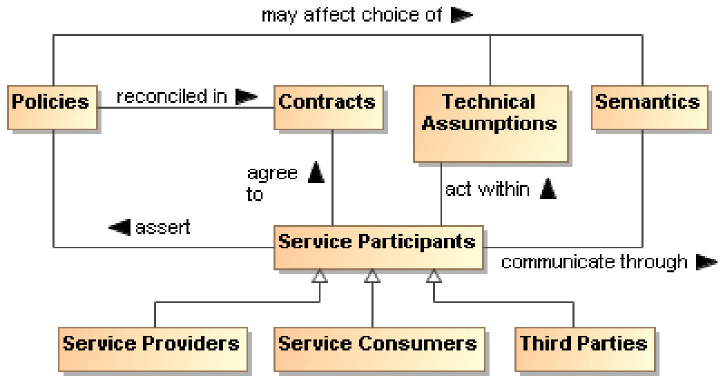

Figure 16 - Model for Policies and Contracts as related to Service Participants

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

As noted in Figure 16, the policies asserted may be reflected as Technical Constraints that available services or their underlying capabilities must be capable of meeting; it may similarly affect the semantics that can be used. For example of the former, there may be a policy that specifies the surge capacity to be accommodated by a server, but a service that is not designed to make use of the larger server capacity would not satisfy the intent of the policy and would not be appropriate to use. For the latter, a policy may require that only services that support interaction via a community-sponsored vocabulary can be used.

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

The definition and later enforcement of policies and contracts are predicated on the potential for measurement; the relationships among the relevant concepts are shown in the model in Figure 17. Performance Metrics identify quantities that characterize the speed and quality of realizing the real world effects produced using the SOA service; in addition, policies and contracts may depend on nonperformance metrics, such as whether a license is in place to use the service. Some of these metrics reflect the underlying capability, e.g. a SOA service cannot respond in two seconds if the underlying capability is expected to take five seconds to do its processing; some metrics reflect the SOA service, e.g. the additional overhead introduced when making data access requests across the network.

Copyright © OASIS Open 2011. All Rights Reserved.

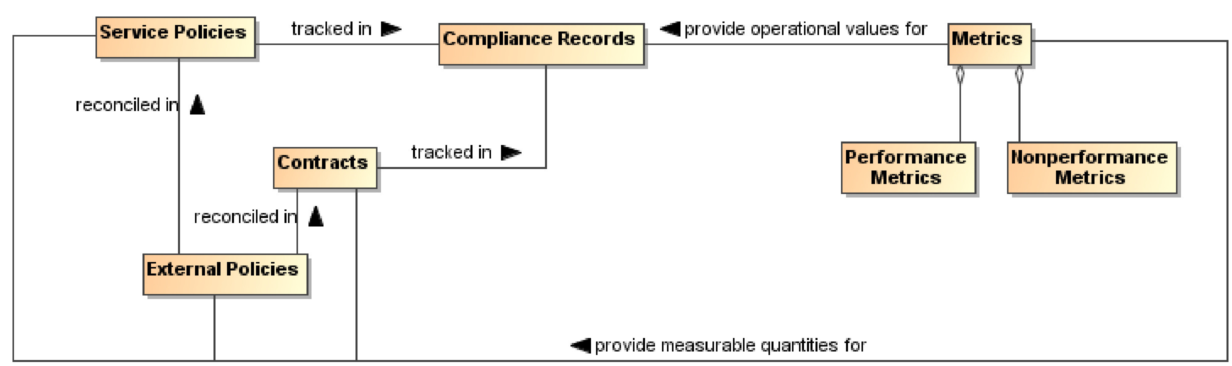

Figure 17 - Policies and Contracts, Metrics, and Compliance Records

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

The use of metrics to evaluate compliance is discussed in Section 4.1.1.3.4. The results of compliance evaluation SHOULD be maintained in compliance records and the means to access the compliance records MAY be included in the Service Policies portion of the service description. For example, the description may be in the form of static information (e.g. over the first year of operation, this service had a 91% availability), a link to a dynamically generated metric (e.g. over the past 30 days, the service has had a 93.3% availability), or access to a dynamic means to check the service for current availability (e.g., a ping). The relationship between service presence and the presence of the individual actions that can be invoked is discussed under Reachability in Section 0.

Note, even when policies relate to the perspective of a single participant, policy compliance can be measured and policies may be enforceable without contractual agreement with other participants. While certain elements of contracts and contract compliance are likely private, public aspects of compliance should be reflected in the compliance record information referenced in the service description.

4.1.2 Use of Service Description

~~4.1.2.2~~4.1.2.1 Service Description in support of Service Interaction

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

 Copyright © OASIS Open 2011. All Rights Reserved.

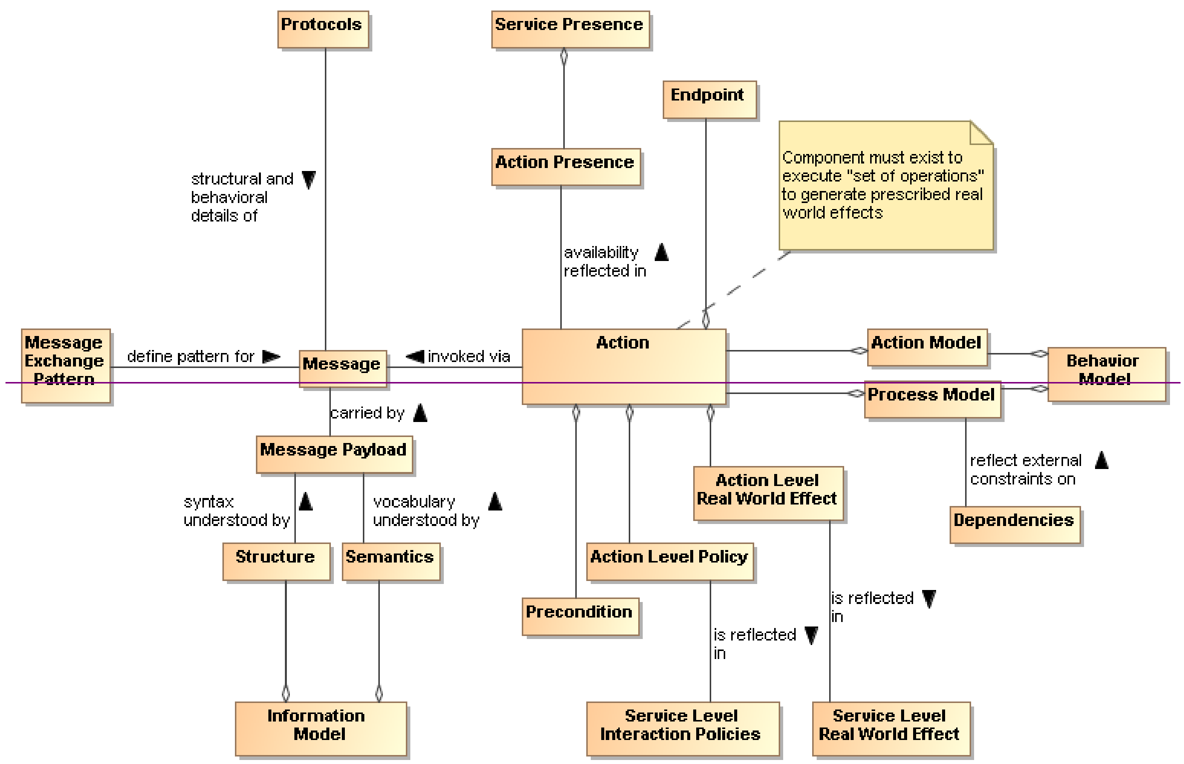

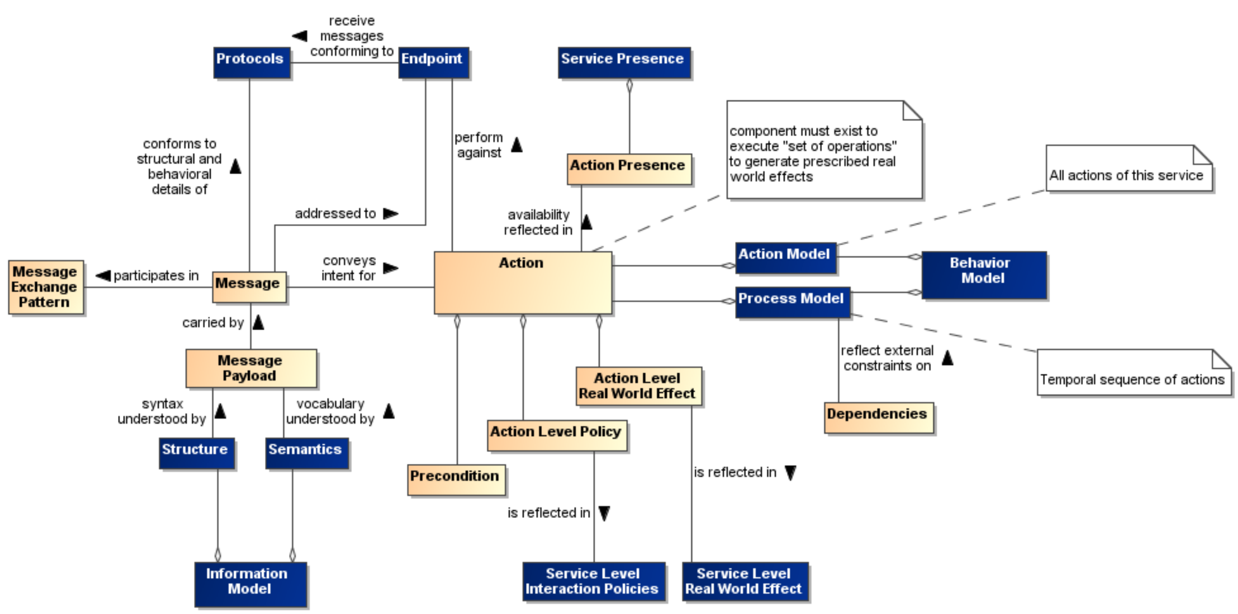

Figure 18~~ Model Showing~~ - Relationship ~~Between~~ between Action and Components of Service Description ~~Components~~Modelx

Figure 18~~EDITOR'S NOTE:~~

~~ONE QUESTION IS WHETHER THE MODEL SHOULD SHOW THE "SAME" ACTION AS POSSIBLY BEING INVOKED THROUGH THE SAME MESSAGE BUT USING A DIFFERENT PROTOCOL AT A DIFFERENT ENDPOINT AND THERE BEING A RELATIONSHIP BETWEEN Endpoint AND Protocol. AGAIN, THIS MAY NOT BE PART OF THE SERVICE DESCRIPTION SECTION BUT OF THE DISCUSSION OF A MODEL FOR ACTION ELSEWHERE.~~

combines the ~~Service Interface model of~~models in the subsections of Section 4.1.1~~ and the Service Reachability model of~~ to concisely relate action~~Action~~ and the relevant components of the Service

 Copyright © OASIS Open 2011. All Rights Reserved.

Description model. The purpose of Figure 18 is to demonstrate that the components of service description go beyond arbitrary documentation and form the critical set of information needed to define the what and how of action. In Figure 18, the leaf nodes from Figure 13~~Service Description.~~ are shown in blue.

Action is typically invoked via a Message where the structure and behavioral details of the message conform to an identified Protocol~~,~~ and is directed to the address of the identified endpoint, and the message payload conforms to the service Information Model~~, and the message sequencing follows an identified Message Exchange Pattern. The protocol, information model, and message exchange pattern are identified in the service description~~.

The availability of an action~~action~~ is reflected in the Action Presence and each Action Presence contributes to the overall Service Presence; this is discussed further in Section 0~~.~~. Each action~~action~~ has its own endpoint and also its own protocols associated with the endpoint[13] and to ~~what~~some extent, e.g. current or average availability, there is presence for the action~~action~~ through that endpoint. The endpoint and service presence are also part of the service description.

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

~~Presence of a service is an aggregation of the presence of the service's actions, and the service level may aggregate to some degraded or restricted presence if some action presence is not confirmed. For example, if error processing actions are not available, the service can still provide required functionality if no error processing is needed. This implies reachability relates to each action as well as applying to the service/business as a whole.~~

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

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

Similarly, the result of invoking an action~~action~~ is one or more real world effect~~real world effects~~s, and ~~the~~any Action Level Real World Effects MUST be reflected in the Service Level Real World Effect included in the service description. ~~If~~ The unambiguous expression of action level policies and real world effect~~real world effects at the action level are not unambiguously expressible at the~~s as service ~~level, then the service description becomes inadequate for expressing conditions for use or results of using the service, and the understanding of~~counterparts is necessary to adequately describe what constitutes ~~a~~the service interaction~~ is called into doubt~~.

[13] 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.

 Copyright © OASIS Open 2011. All Rights Reserved.

~~From a~~An adequate service description ~~standpoint,~~MUST provide a consumer ~~would show interest in a service~~with information needed to determine if the service ~~functionality is what is needed and the service~~ policies~~ are at least worth pursuing if not immediately acceptable. By saying functionality is of interest, we are saying~~, the (business) functions, and service-level real world effects are of interest, and there is nothing in the technical ~~assumptions~~constraints that preclude use of the service.

Note at ~~this ~~the service level, the business functions are not concerned with the action or process models. These models ~~get into the nuts and bolts of making the business function happen and will be dealt with at that level later.~~are detailed separately.

The service description is not intended to be isolated documentation but rather an integral part of service use. ~~The initial use of any service should be based on information contained in the service description, and changes in service description should be pushed to known consumers. Thus, changes would not be introduced that later are captured in perpetually out-of-date documentation but rather reference to the service description should be an integral part of service use. This idea is consistent with checking the service endpoint before invoking a service action, but use of service description information should be more intrinsic than merely for a DNS-type function~~Changes in service description SHOULD immediately be made known to consumers and potential consumers.

~~4.1.2.2.1~~4.1.2.1.1 Description and Invoking Actions Against a Service

At this point, let us assume the descriptions were sufficient to establish willingness; see Section 4.2.2.2~~.~~. Figure 18 indicates the service endpoint establishes where to ~~go to ~~actually carry out the interaction. This is where we ~~have to ~~start considering the action~~action~~ and process models.

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

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

~~In the discussion above, we indicate~~We have established the importance of the process model in identifying relevant actions and their sequence. Interaction ~~with the actions are~~proceeds through messages and thus it is the syntax and semantics of the messages with which we are here concerned. ~~There seems to be a number of ways to ~~A common approach ~~this but the common way now ~~is to define the structure and semantics that can appear as part of a message~~ and~~; then assemble the pieces into messages; and, associate messages with actions. Actions make use of structure and semantics as defined in the information model to describe its legal messages.~~ In addition, the message exchange pattern defines sequencing and use of messages for a given action.~~

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

 Copyright © OASIS Open 2011. All Rights Reserved.

~~4.1.2.2.2~~4.1.2.1.2 The ~~question~~Question of ~~multiple business functions~~Multiple Business Functions

~~The service description model discussed above applies to the service and not the components of the service. For example, the~~ Action ~~Model identifies numerous actions that can be performed against a service and the Process Model defines the order in which the actions are performed, but the real world~~ level effects ~~are defined for the service and not the individual actions. Similarly, numerous~~ and policies ~~may be associated with a service, but policies at the action level must~~MUST be reflected at the service level for service description to support visibility.

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

~~To begin, consider~~Consider the situation if a given SOA service is the ~~container~~mechanism for access to multiple independent (but ~~possibly~~ loosely related) business functions. ~~Note, this is~~These are not multiple effects from a single function but multiple functions with potentially different sets of effects for each function. ~~As noted above, a~~A service can have multiple actions a user ~~can~~may perform against it, and this does not change with multiple business functions. ~~An~~As an individual business function corresponds to a process model, so multiple business functions imply multiple process models ~~because either the process is different or the specific action performed for some process step is different.~~. The same action~~action~~ may be used in multiple process models but the aggregated service presence would be specific to each business function because the components being aggregated ~~will likely~~may be different between process models. In summary, for a service with multiple business functions, each function has (1) its own process model and dependencies, (2) its own aggregated presence, and (3) possibly its own list of policies and real world effect~~real world effects.~~ s.

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

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

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

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

As a best practice, a criteria for whether a service is appropriately scoped may be the ease or difficulty in creating an unambiguous service description. A consequence of having tightly-scoped services is there will likely be a greater reliance on combining services, i.e. more fundamental business functions, to create more advanced business functions. This is consistent with the principles of service oriented architecture and is the basic position of the Reference Architecture, although not an absolute requirement. Combining services increases the reliance on understanding and implementing the concepts of orchestration,

 Copyright © OASIS Open 2011. All Rights Reserved.

choreography, and other approaches yet to be developed; these are discussed in more detail in section 4.4 Interacting with Services.

~~4.1.2.2.3~~4.1.2.1.3 Service Description, Execution Context, and Service Interaction

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

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

For the same employee connecting from offsite, the execution context specifies the need for a computer with installed VPN software and a security token to negotiate the VPN connection. The execution context also includes proxy settings as needed to connect to the offsite network. The employee must still comply with the requirements for onsite computers and access, but the offsite execution context includes additional items before the employee can access the same underlying capability and realize the same

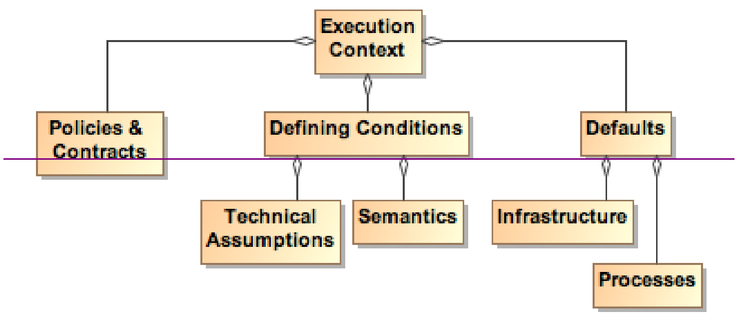

real world effect

s, i.e. the timecard entries.

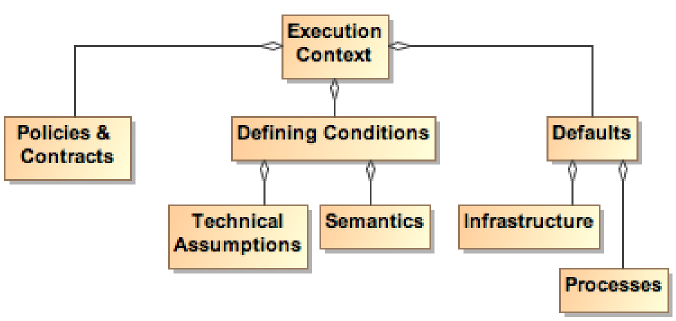

Figure 19 - Execution Context ~~model~~

Figure 19 shows a ~~number of contributors to the execution context. These~~ few broad categories found in execution context. These are not meant to ~~include any disconnects that could get in the way of interoperability and successful interactions, but other~~be comprehensive. Other items may need to be included to ~~collect~~provide a sufficient description of the interaction conditions. Any other items not explicitly noted in the model but needed to set the environment ~~would also be a candidate for including in the execution context. However, as noted in the Reference Model, it is not possible to describe everything and so a set of information items as potentially extensive as the execution context will never~~

 Copyright © OASIS Open 2011. All Rights Reserved.

~~be complete in every detail. As with the service description, the goal is to be sufficiently complete for the task at hand.~~SHOULD be included in the execution context.

While the execution context captures the conditions under which interaction can occur, it does not capture the specific service invocations that do occur in a specific interaction. A service interaction as modeled in Figure 20 introduces the concept of an Interaction Description which is composed of both the Execution Context and an Interaction Log. The execution context specifies the set of conditions under which the interaction occurs and the interaction log captures the sequence of service interactions that occur within the execution context. This sequence should follow the Process Model but can include details beyond those specified there. For example, the Process Model may specify an action that results in identifying a data source, and the identified source is used in a subsequent action. The Interaction Log would record the specific data source used.

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

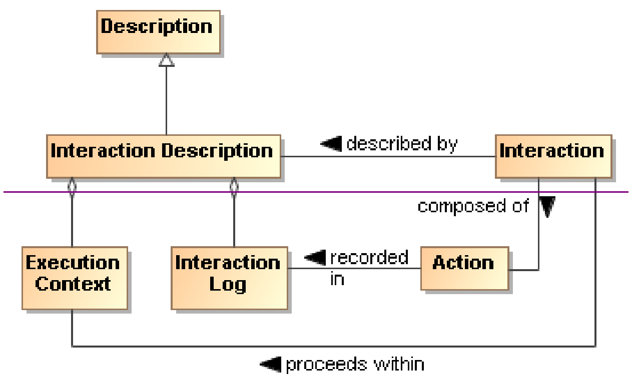

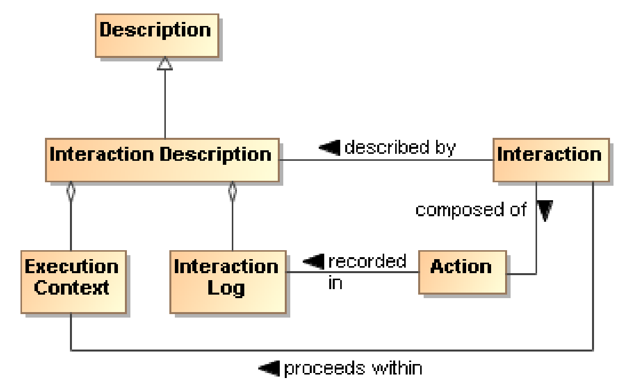

Figure 20 ~~Service~~ - Interaction ~~model~~Description

~~With respect to repeatability,~~ SOA allows ~~for a great deal of~~ flexibility to accomplish both repeatability and ~~one of its benefits is that services and their underlying capabilities~~reusability. In facilitating reusability, a service can be updated without ~~disturbing~~disrupting the ~~consumers.~~ user experience of the service. So, ~~for example,~~ Google can improve their ranking algorithm ~~in a manner transparent to the typical user~~ without notifying the user ~~being concerned with~~about the details of the update. ~~Indeed, improvements in Google often depend on the user being unaware of updates because that allows Google to adapt to content providers trying to game the ranking algorithms.~~

However, it may also be vital for the consumer to be able to recreate past results or to generate consistent results in the future, and information such as what conditions, which services, and which versions of those services ~~are~~were used is indispensible in retracing one's path. The interaction log is a critical part of the resulting real world effect~~real world effects~~s because it defines how the effects were generated and possibly the meaning of observed effects. This increases in importance as dynamic

 Copyright © OASIS Open 2011. All Rights Reserved.

composability becomes more feasible. In essence, a result has limited value if one does not know how it was generated.

The interaction log ~~is~~SHOULD be a detailed trace for a specific interaction, and its reuse is limited to duplicating that interaction. ~~On the other hand, an~~An execution context can ~~be reusable for the same participants using the same services or it can~~ act as a template for ~~those items to consider for~~ identical or similar interactions. ~~A previous~~Any given execution context ~~could provide a starting point for defining~~MAY define the conditions of future interactions~~, either between the same consumer and provider or by like-minded consumers and providers attempting to carry out similar tasks.~~.

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

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

4.1.3 Relationship to Other Description Models

While the representation shown in Figure 12 is derived from considerations related to service description, it is acknowledged that other metadata standards are relevant and should, as possible, be incorporated into this work. Two standards of particular relevance are the Dublin Core Metadata Initiative (DCMI) **[DCMI]** and ISO 11179~~,~~, **[ISO 11179]**, especially Part 5.

When the service description (or even the general description class) is considered as the DCMI "resource", Figure 12 aligns nicely with the DCMI resource model. While some differences exist, these are mostly in areas where DCMI goes into detail that is considered beyond the scope of the current Reference Architecture. For example, DCMI defines classes of "shared semantics" whereas ~~for the~~this Reference Architecture~~, it is sufficient to prescribe~~ Framework considers that an identification of relevant semantic models is sufficient. Likewise, the DCMI "description model" goes into the details of possible syntax encodings whereas for the Reference Architecture Framework it is sufficient to identify the relevant formats.

With respect to ISO 11179 Part 5, the metadata fields defined in that reference may be used without prejudice as the properties in Figure 12~~ above.~~. Additionally, other defined metadata sets may be used by the service provider if the other sets are considered more appropriate, i.e. it is fundamental to this ~~Reference Architecture~~reference architecture to identify the need and the means to make vocabulary declarations explicit but it is beyond the scope to specify which vocabularies are to be used. In addition, the identification of domain of the properties and range of the values has not been included in the current Reference Architecture discussion, but the text of ISO 11179 Part 5 can be used consistently with the model prescribed in this document.

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

4.1.4 Architectural Implications

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

- It changes~~Description will change~~ over time and its contents will reflect changing needs and context. This requires the existence of:
    - mechanisms to support the storage, referencing, and access to normative definitions of one or more versioning schemes that may be applied to identify different aggregations of descriptive information, where the different schemes may be versions of a versioning scheme itself;

 Copyright © OASIS Open 2011. All Rights Reserved.

    - configuration management mechanisms to capture the contents of~~ the~~ each aggregation and apply a unique identifier in a manner consistent with an identified versioning scheme;
    - one or more mechanisms to support the storage, referencing, and access to conversion relationships between versioning schemes, and the mechanisms to carry out such conversions.
- Description makes use of defined semantics, where the semantics may be used for categorization or providing other property and value information for description classes. This requires the existence of:
    - semantic models that provide normative descriptions of the utilized terms, where the models may range from a simple dictionary of terms to an ontology showing complex relationships and capable of supporting enhanced reasoning;
    - mechanisms to support the storage, referencing, and access to these semantic models;
    - configuration management mechanisms to capture the normative description of each semantic model and to apply a unique identifier in a manner consistent with an identified versioning scheme;
    - one or more mechanisms to support the storage, referencing, and access to conversion relationships between semantic models, and the mechanisms to carry out such conversions.
- Descriptions include reference to policies defining conditions of use~~ and optionally contracts representing agreement on policies and other conditions.~~. In this sense, policies are also resources that need to be visible, discoverable, and accessible. This requires the existence of (as also enumerated under governance):
    - ~~descriptions to enable the policy modules to be visible, where the description includes~~description of policies, including a unique identifier for the policy and a sufficient, and preferably a machine processible, representation of the meaning of terms used to describe the policy~~policy,~~, its functions, and its effects;
    - one or more discovery mechanisms that enable searching for policies that best meet the search criteria specified by the service participant~~participant;~~; where the discovery mechanism ~~will have~~has access to the individual policy~~policy~~ descriptions, possibly through some repository mechanism;
    - accessible storage of policies and policy~~policy~~ descriptions, so service participants~~participants~~ can access, examine, and use the policies as defined.
- Descriptions include references to metrics which describe the operational characteristics of the subjects being described. This requires the existence of (as partially enumerated under governance):
    - the infrastructure monitoring and reporting information on SOA resources;
    - possible interface requirements to make accessible metrics information generated~~ or most easily accessed by the service itself~~;
    - mechanisms to catalog and enable discovery of which metrics are available for a described resources and information on how these metrics can be accessed;
    - mechanisms to catalog and enable discovery of compliance records associated with policies and contracts~~contracts~~ that are based on these metrics.
- Descriptions of the interactions are important for enabling auditability and repeatability, thereby establishing a context for results and support for understanding observed change in performance or results. This requires the existence of:
    - one or more mechanisms to capture, describe, store, discover, and retrieve interaction logs, execution contexts, and the combined interaction descriptions;
    - one or more mechanisms for attaching to any results the means to identify and retrieve the interaction description under which the results were generated.
- Descriptions may capture very focused information subsets or can be an aggregate of numerous component descriptions. Service description is an example of ~~a likely~~an aggregate for which manual maintenance of ~~all aspects~~the whole would not be feasible. This requires the existence of:
    - tools to facilitate identifying description elements that are to be aggregated to assemble the composite description;
    - tools to facilitate identifying the sources of information to associate with the description elements;

 Copyright © OASIS Open 2011. All Rights Reserved.

- ○ tools to collect the identified description elements and their associated sources into a standard, referenceable format that can support general access and understanding;
  - ○ tools to automatically update the composite description as the component sources change, and to consistently apply versioning schemes to identify the new description contents and the type and significance of change that occurred.
- ~~Descriptions provide up-to-date information on what a resource is, the conditions for interacting with the resource, and the results of such interactions. As such, the~~The description is the source of vital information in establishing willingness to interact with a resource~~resource,~~, reachability to make interaction possible, and compliance with relevant conditions of use. This requires the existence of:
  - ○ one or more discovery mechanisms that enable searching for described resource~~resources~~s that best meet the criteria specified by a service participant~~participant, where the discovery mechanism will have access to individual descriptions, possibly through some repository mechanism;~~;
  - ○ tools to appropriately track users of the descriptions and notify them when a new version of the description is available.

4.2 Service Visibility Model

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

4.2.1 Visibility to Business

The relationship of visibility to the SOA ecosystem encompasses both human social structure~~social structures~~s and automated IT mechanisms. Figure 21 depicts a business setting that is a basis for visibility as related to the social structure~~Social Structure~~ Model in the ~~Business Via Services View~~Participation in a SOA Ecosystem view (see Section 3.1~~). Service consumers~~). Service consumers and service providers may have direct awareness or mediated awareness where mediated awareness is achieved through some third party. A consumer's willingness to use a service is reflected by the consumer's presumption of satisfying goals and needs based on the service description~~ of the service.~~. Service providers offer capabilities that have real world effect~~real world affects~~s that result in a change in state~~ of the consumer.~~. Reachability of the service by the consumer ~~leads~~may lead to interactions that change the state of the ~~consumer.~~SOA ecosystem.   The consumer can measure the change of state to determine if the claims made by description and the real world effect~~real world effects~~s of consuming the service meet the consumer's needs.

 Copyright © OASIS Open 2011. All Rights Reserved.

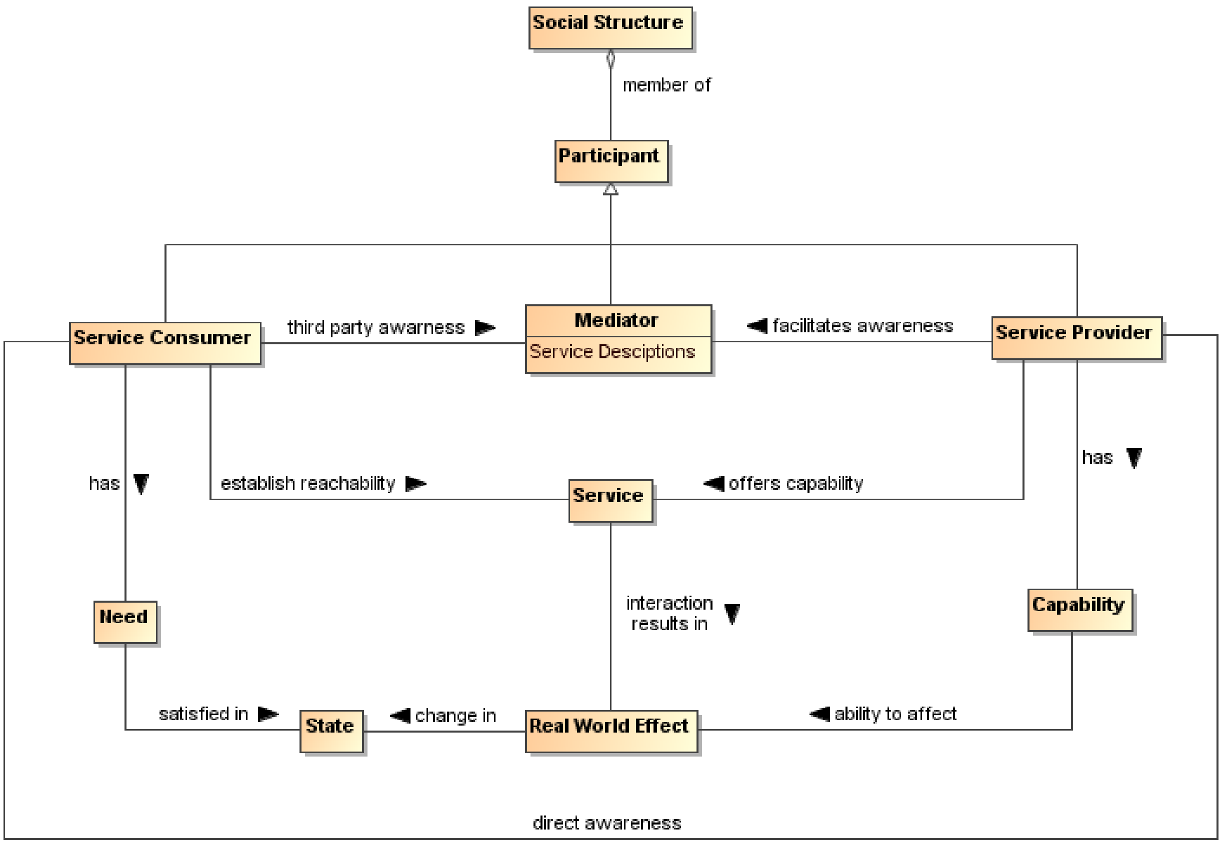

Figure 21 - Visibility to Business ~~Model~~

Visibility and interoperability in a SOA ecosystem requires more than location and interface information~~,~~ ~~or the traditional Application Programming Interface (API).~~. A meta-model for this broader view of visibility is depicted in Section 4.1. In addition to providing improved awareness of service capabilities through description of information such as reachability, behavior models, information models, functionality, and metrics, the service description may contain policies valuable for determination of willingness to interact.

~~Another important business capability in a SOA environment is the ability to narrow visibility to trusted members within a social structure, often referred to as Communities of Interest (COI) in government sectors. Mediators for awareness may provide policy based access to~~A mediator using service descriptions~~, allowing for the dynamic formation of awareness between members of a COI.~~

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

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

4.2.2 Visibility

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

Attaining SOA visibility requires

 Copyright © OASIS Open 2011. All Rights Reserved.

- service description creation and maintenance,
- processes and mechanisms for achieving awareness of and accessing descriptions,
- processes and mechanisms for establishing willingness of participants~~participants,~~,
- processes and mechanisms to determine reachability.

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

4.2.2.1 ~~Achieving~~ Awareness

~~A~~ An important means for a service participant~~participant is~~ to be aware of another participant~~participant if it has~~ is to have access to a description of that participant~~participant with~~ and for the description to have sufficient completeness to establish the other requirements of visibility.

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

Awareness can be decomposed into: creating the~~ creation of~~ descriptions, making them available, and discovering the descriptions. ~~Discovery in the Service Visibility Model is the process where a consumer discovers a service description or a service provider discovers a likely consumer's description.~~ Discovery can be initiated or it can be by notification. Initiated discovery for business may require formalization of the required capabilities and resource~~resources~~s to achieve business goals. ~~and depict a typical process for achieving awareness.~~

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

~~*Figure Publishing Description*~~

A mediator ~~as discussed~~ for awareness is a third party participant~~participant~~ that provides awareness to one or more consumers of one or more services. ~~See Section , for an overview of participants.~~ Direct awareness is awareness between a consumer and provider without the use of a third party. A registry/repository can act as a mediator; a Web page displaying similar information can also be considered a mediator.

Direct awareness may be the result of having previously established an execution context~~ and possibly indicates successful~~, or direct awareness may include determining the presence of services and then querying the service directly for description. As an example, a priori visibility of some sensor device may provide the means for interaction ~~has occurred in the past.~~ or a query for standardized sensor device metadata may be broadcast to multiple locations. If acknowledged, the service interface for the device may directly provide description to a consumer so the consumer can determine willingness to interact.

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

 Copyright © OASIS Open 2011. All Rights Reserved.

~~*Figure Discovering Description*~~

~~There may be numerous methods to facilitate discovery. For example, descriptions could be discovered by browsing a web site, querying a public registry, or via email notifications.~~

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

4.2.2.1.1 Mediated Awareness

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

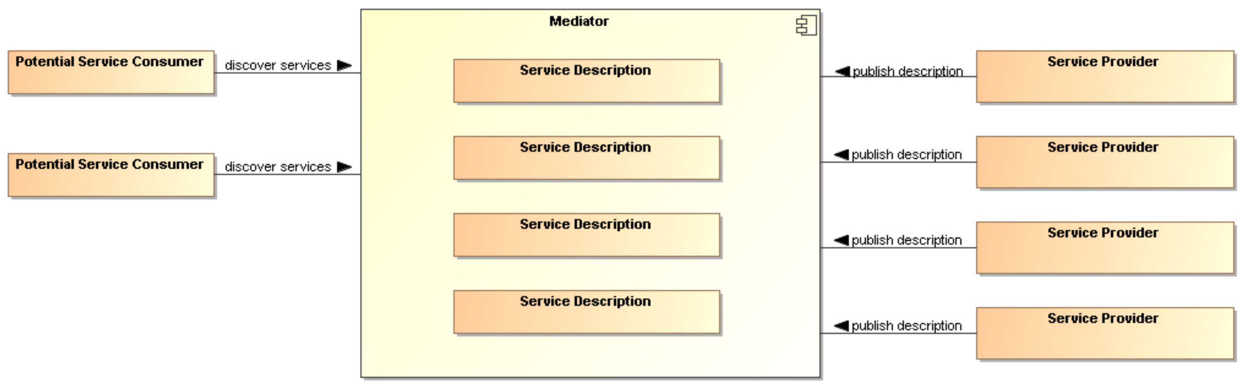

Figure 22 - Mediated Service Awareness

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

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

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

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

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

Copyright © OASIS Open 2011. All Rights Reserved.

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

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

4.2.2.1.2 Awareness in Complex Social Structures

~~4.2.2.1.2~~1.1.1.1.1 ~~Awareness in Complex Social Structures~~

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

In Figure 23, awareness can be between consumers and providers within a single community, multiple communities, or all communities in the social structure~~social structure.~~. The social structure~~social structure~~ can encourage or restrict awareness through its policies, and these policies can affect participant~~participant~~ willingness. The information about policies should be incorporated in the relevant descriptions. ~~The information about policies should be incorporated in the relevant descriptions.~~ The social structure~~social structure~~ also governs the conditions for establishing contracts~~contracts,~~, the results of which ~~will be~~are reflected in the execution context if interaction is to proceed.

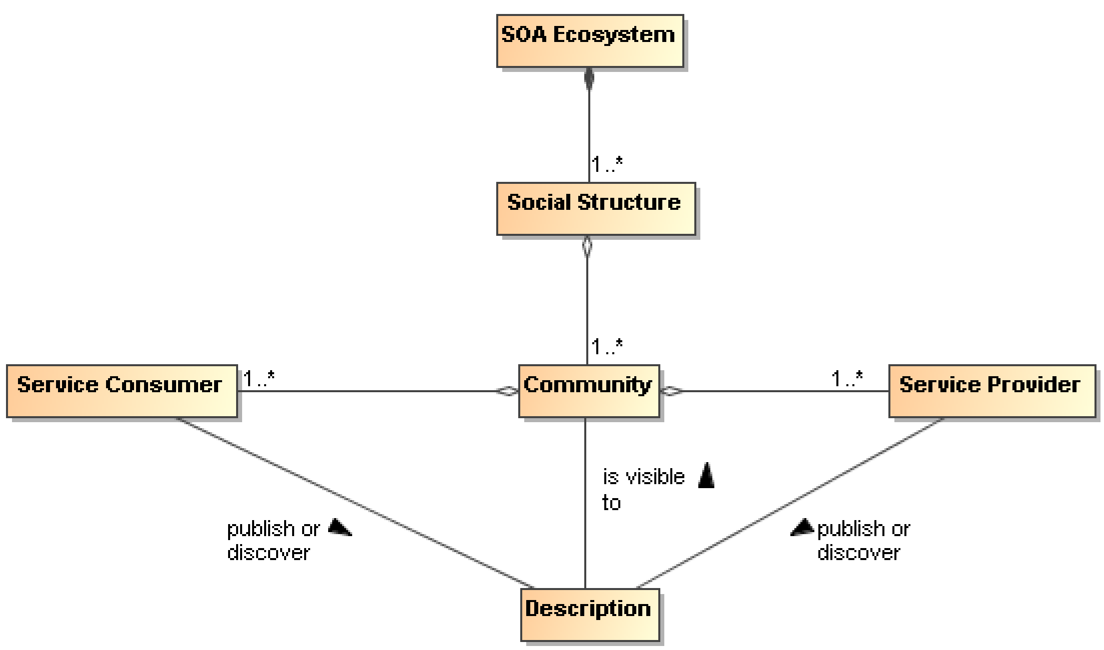

Figure 23 - Awareness ~~In~~in a SOA Ecosystem

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

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

 Copyright © OASIS Open 2011. All Rights Reserved.

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

4.2.2.2 ~~Determining~~ Willingness

Having achieved awareness, participants~~participants~~ use descriptions to help determine their willingness to interact with another participant~~participant.~~. Both awareness and willingness are determined prior to consumer/provider interaction. ~~The activities in , or a subset there of, can be performed to help determine willingness.~~

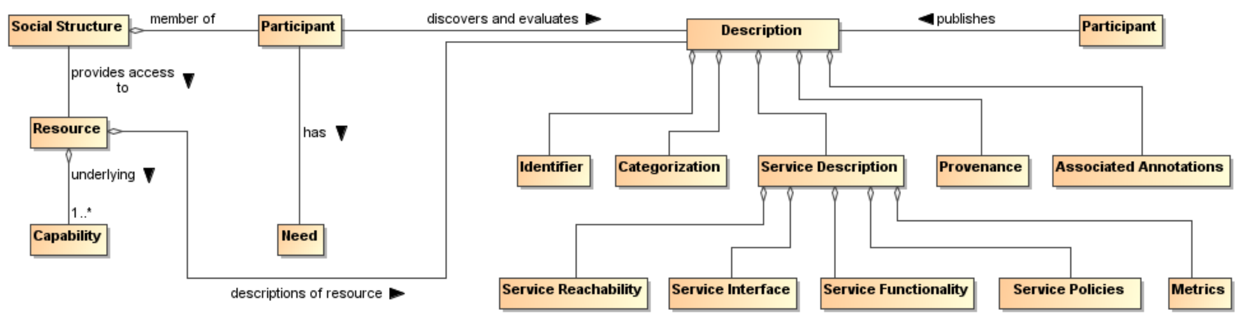

Figure 24 ~~Determining~~ ~~Willingness~~

~~In any given process to determine willingness, one or more of the transitions or flows depicted above may be executed. For example, in a particular service interaction, it may be important to inspect policies and to verify provenance; another interaction may call for evaluating 3rd party annotations in addition.~~

~~Figure~~ Business, Description and Willingness

Figure 24 relates elements of the ~~Business via Services View~~*Participation in a SOA Ecosystem* view, and elements from the Service Description Model to willingness. By having a willingness to interact within a particular social structure~~social structure,~~, the social structure provides the participant~~participant~~ access to capabilities based on conditions the social structure~~social structure~~ finds appropriate for its context. The participant~~participant~~ can use these capabilities to satisfy goals and objectives as specified by the participant~~participant's~~'s needs.

In Figure 24, information used to determine willingness is defined by Description. Information referenced by Description may come from many sources. For example, a mediator for descriptions may provide 3rd party annotations for reputation. Another source for reputation may be a participant~~participant's~~'s own history of interactions with another participant~~participant.~~.

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

~~4.2.2.3 Establishing Reachability~~

~~Reachability involves knowing the service endpoint, service interface, and presence of a service. lists activities involved to establish reachability. For reachability, service descriptions should include sufficient data to enable a service consumer and service provider to interact with each other. At a minimum,~~

 Copyright © OASIS Open 2011. All Rights Reserved.

~~service descriptions should include information about the location of the service and the service interface. The subject of access control and other process model type activities to establish a connection are left for the Interacting with Services Model.~~

~~*Figure Establishing Reachability*~~

~~Endpoint~~

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

~~**Interface**~~

> ~~Interface verification involves determination of compatible communication protocols, compatible message exchange capabilities, and service interface version.~~

~~**Presence**~~

~~Presence is established when a service can be reached at a particular point in time. Presence may not be known in many cases until the act of interaction begins. To overcome this problem, IT mechanisms may make use of presence protocols to provide the current up/down status of a service.~~

~~Service reachability enables service participants to locate and interact with one another. Each action may have its own endpoint and also its own protocols associated with the endpoint[14] and whether there is presence for the action through that endpoint. Presence of a service is an aggregation of the presence of the service's actions, and the service level may aggregate to some degraded or restricted presence if some action presence is not confirmed. For example, if error processing actions are not available, the service can still provide required functionality if no error processing is needed. This implies reachability relates to each action as well as applying to the service/business as a whole~~

~~After reachability has been established, there may be times when participants need to re-establish reachability such as when a service fails and a new location and version for the service needs to be determined. Disconnected operations is another example for re-establishment of reachability. For SOA, both endpoint location and service interface version are important for re-establishing reachability. For example, multiple versions of a service may be in operation for backward compatibility. A Domain Name Service (DNS) lookup for service location may not be sufficient for re-establishing service reachability after a failure.~~

~~4.2.3 Mechanisms for Attaining Visibility~~

~~While there can be many mechanisms for service visibility in a SOA, this section covers some examples of those mechanisms.~~

~~4.2.3.1 Mechanisms for Awareness~~

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

---

[14] 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.

 Copyright © OASIS Open 2011. All Rights Reserved.

~~A common mechanism for mediated awareness in the industry is a registry-repository. depicts a mediation facility containing a registry and a repository. The registry stores links or pointers to service description artifacts. The repository in this example is the storage location for the service description artifacts. Service descriptions can be pushed (publish/subscribe for example) or pulled from the register-repository mediator.~~

~~*Figure Mediated Registry-Repository*~~

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

~~When like SOA IT mechanisms interoperate with one another, the IT mechanisms may be referred to as federated. An example use of federation is combining different domains of knowledge as in .~~

~~*Figure Federated Registry-Repository*~~

~~4.2.3.2 Mechanisms for Willingness~~

Mechanisms that aid in determining willingness make use of the artifacts referenced by descriptions of services. Mechanisms for establishing willingness could be as simple as rendering service description information for human consumption to automated evaluation of functionality, policies, and contracts~~contracts~~ by a rules engine. The rules engine for determining willingness could operate as a policy decision procedure~~policy decision point~~ as defined in Section 1.1.1.

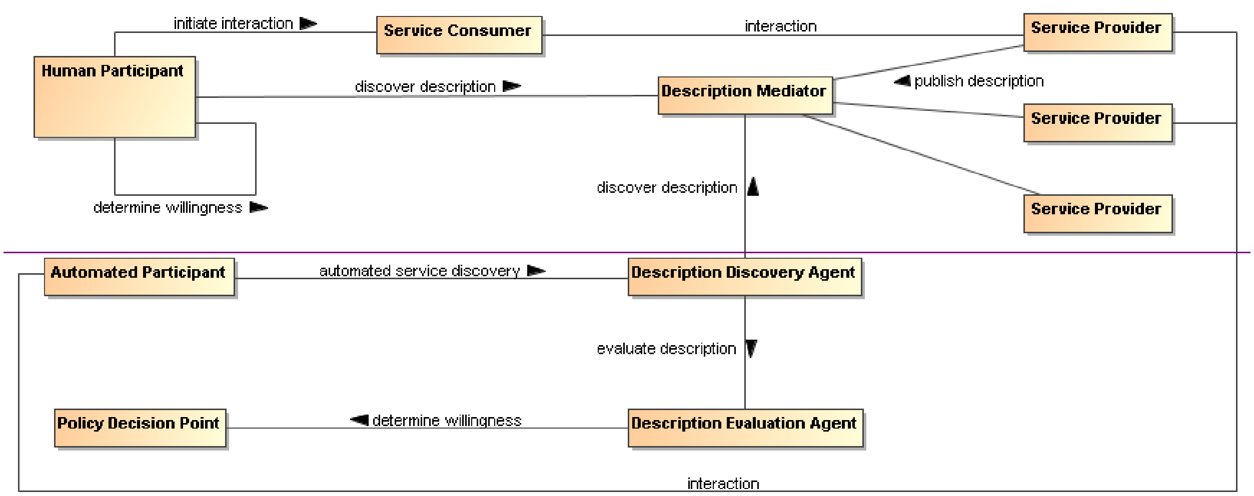

~~*Figure Mechanisms for Willingness*~~

~~Figure 25 is an example of manual determination of willingness by a human participant and one possible example of automated determination of willingness. For functionality that may be provided by the Enterprise Service Bus see Section . For models explaining the Policy Decision Point see Section .~~

~~4.2.3.3~~4.2.2.3 ~~Mechanisms for~~ Reachability

Reachability ~~mechanisms will often begin with a tool that is capable of reading service description interfaces and generating a client capable of interacting with the provider's service. The establishment of~~ involves knowing the endpoint, protocol, and presence ~~occurs when~~ of a service. At a minimum, reachability requires information about the location of the service and the ~~client has started interactions with~~protocol describing the means of communication.

 Copyright © OASIS Open 2011. All Rights Reserved.

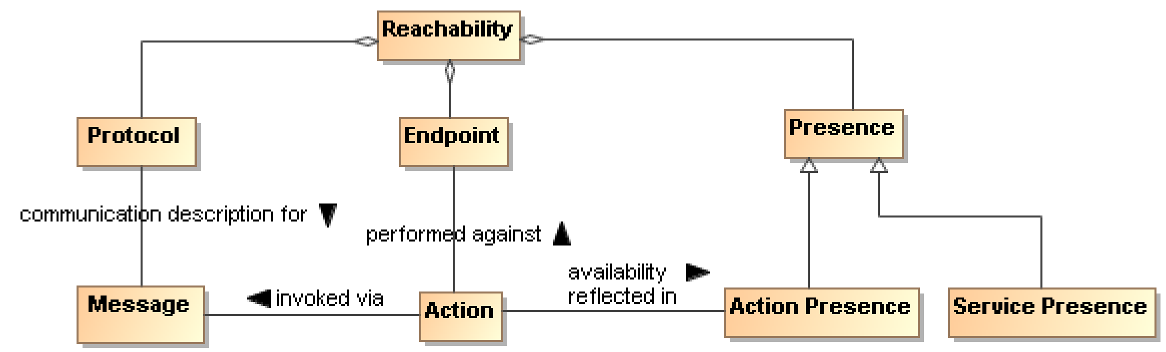

Figure 25 - Service Reachability

Endpoint

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

Protocol~~provider's~~

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

Presence~~. Expected~~

Presence is the measurement of reachability of a service ~~operating times may~~at a particular point in time.

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

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

~~4.2.4~~4.2.3 Architectural Implications

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

- Mechanisms providing support for awareness ~~will likely~~ have the following minimum capabilities:
  - creation of Description, preferably conforming to a standard Description format and structure;
  - publishing of Description directly to a consumer or through a third party mediator;
  - discovery of Description, preferably conforming to a standard for Description discovery;
  - notification of Description updates or notification of the addition of new and relevant Descriptions;
  - classification of Description elements according to standardized classification schemes.
- In a SOA ecosystem with complex social structure~~social structures~~s, awareness may be provided for specific communities of interest.   The architectural mechanisms for providing awareness to communities of interest ~~will~~ require support for:
  - policies that allow dynamic formation of communities of interest;
  - trust that awareness can be provided for and only for specific communities of interest, the bases of which is typically built on keying and encryption technology.

 Copyright © OASIS Open 2011. All Rights Reserved.

- The architectural mechanisms for determining willingness to interact ~~will~~ require support for:
    - verification of identity and credentials of the provider and/or consumer;
    - access to and understanding of description;
    - inspection of functionality and capabilities;
    - inspection of policies and/or contracts~~contracts.~~.
- The architectural mechanisms for establishing reachability ~~will~~ require support for:
    - the location or address of an endpoint;
    - verification and use of a service interface ~~which includes~~by means of a communication ~~protocols, message exchange capabilities, and service interface version~~protocol;
    - determination of presence with an endpoint which may only be determined at the point of interaction but may be further aided by the use of a presence protocol for which the endpoints actively participate.

4.3 Interacting with Services Model

Interaction is the ~~use of~~activity involved in using a service to access capability in order to achieve a particular desired real world effect~~real world effect,~~, where real world effect is the actual result of using a service. An interaction can be characterized by a sequence of communicative actions. Consequently, interacting with a service~~ involves performing actions against ~~, i.e. participating in joint action with the service~~, —~~usually ~~through~~mediated by a series of ~~information~~ message exchanges~~ (e.g., messages), although other modes of interaction are possible such as modifying—~~involves individual actions performed by both the ~~shared state of a resource.~~service and the consumer.[15] Note that a participant~~participant~~ (or delegate~~agent~~ acting on behalf of the participant~~participant)~~) can be the sender of a message, the receiver of a message, or both.

4.3.1 Interaction Dependencies

Recall from the Reference Model that service visibility is the capacity for those with needs and those with capabilities to be able to interact with each other, and that the service interface is the means by which the underlying capabilities of a service are accessed. Ideally, the details of the underlying service implementation are abstracted away by the service interface. [Service] interaction therefore has a direct dependency on the visibility of the service as well as its implementation-neutral interface (see Figure 26). Service visibility is composed of awareness, willingness, and reachability and service interface is composed of the information and behavior models. Service visibility is modeled in Section 4.2 while service interface is modeled in Section 4.1.

[15] In order for multiple actors to participate in a joint action, they must each act according to their role within the joint action. For SOA-based systems, this is achieved through a message exchange style of communication. The concept of "joint action" is further described in Section **Error! Reference source not found.**.

 Copyright © OASIS Open 2011. All Rights Reserved.

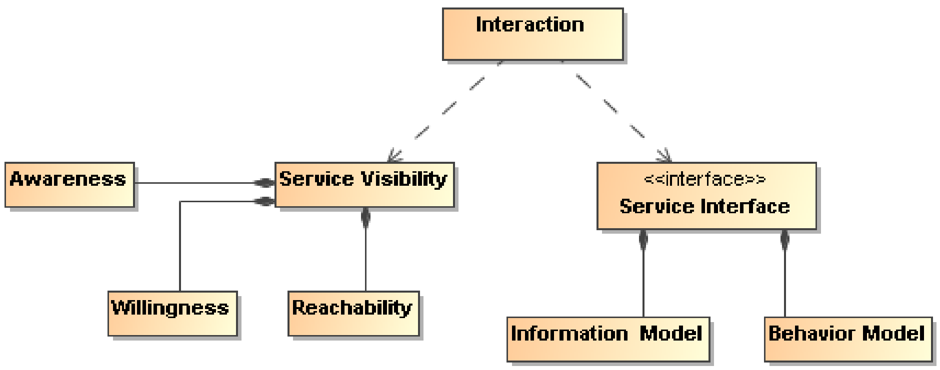

Figure 26~~For purposes of this SOA Reference Architecture, the authors have committed to the use of~~ - Interaction dependencies

4.3.2 Actions and Events

The SOA-RAF uses message exchange between service participants~~participants~~ to denote actions performed against and by the ~~services that *cause* a real world effect~~service, and to denote event~~events~~s that report on real world effect~~real world effects~~s that ~~arise from those~~ are caused by the service actions. A visual model of the relationship between these concepts is shown in Figure 27.

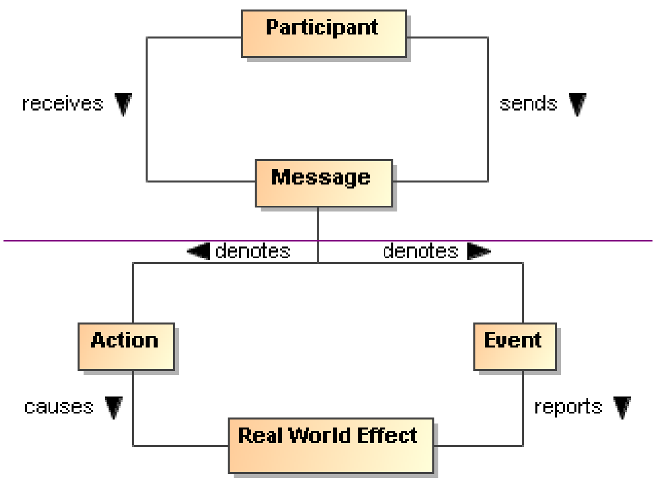

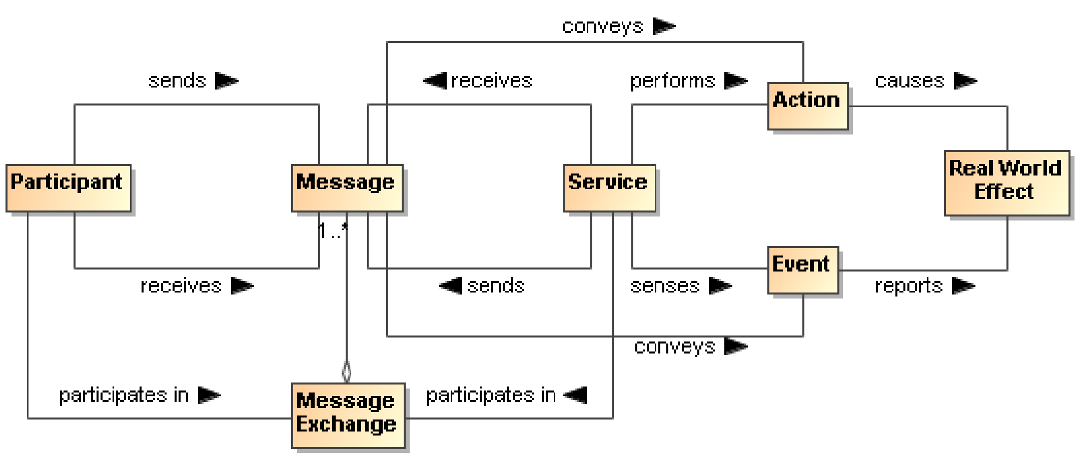

Figure 27 - A "message" denotes either an action or an event~~.~~

 Copyright © OASIS Open 2011. All Rights Reserved.

~~A *Message* denotes either an action or an event. In other words, both~~Both actions and event~~events are~~s, realized ~~through~~by the SOA services, are denoted by the messages. The~~ OASIS~~ Reference Model states that the ~~Action Model~~action model characterizes the "permissible set of actions that may be invoked against a service." We extend that notion here to include event~~events~~s as part of the ~~action~~event model and that messages ~~denote either~~are intended for invoking actions or for notification of event~~events.~~s.

~~4.3.1 Actions and Events~~

In Section 3.2.3~~,~~ we saw that participants~~participants~~ interact with each other in order to ~~perform~~participate in joint actions. ~~An action~~ A joint action is not itself the same thing as the result of ~~performing~~ the joint action~~action.~~. When ~~an action~~a joint action is ~~performed against~~ participated in with a service, the real world effect~~real world effect~~ that results ~~is~~ may be reported in the form of ~~events (see Section).~~ an event notification.

~~In this Reference Architecture, we use *messages* and *message exchange* to denote both actions and results of actions.~~

~~**Message Exchange**~~

~~4.3.2~~4.3.3 Message Exchange

Message exchange is the means by which service participants~~participants~~ (or their delegate~~agents~~s) interact with each other. There are two primary modes of interaction: joint actions that cause real world effect~~real world effects,~~s and notification of event~~events~~s that report real world effects.[16]

A message exchange is used to affect an action~~action~~ when the messages contain the appropriately formatted content~~ that should be interpreted as joint~~, are directed towards a particular action ~~and~~in accordance with the action model, and the delegate~~agents~~s involved interpret the message appropriately.

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

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

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

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

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

 Copyright © OASIS Open 2011. All Rights Reserved.

Message Exchange

~~Message Exchange~~

The means by which joint ~~actions~~action and event notifications are coordinated by service participants~~participants~~ (or delegate~~agents~~s).

Operations

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

~~4.3.2.1~~4.3.3.1 Message Exchange Patterns (MEPs)

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

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

- Request/response to represent how actions cause a real world effect~~real world effect~~
- Event notification to represent how event~~events~~s report a real world effect~~real world effect~~

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

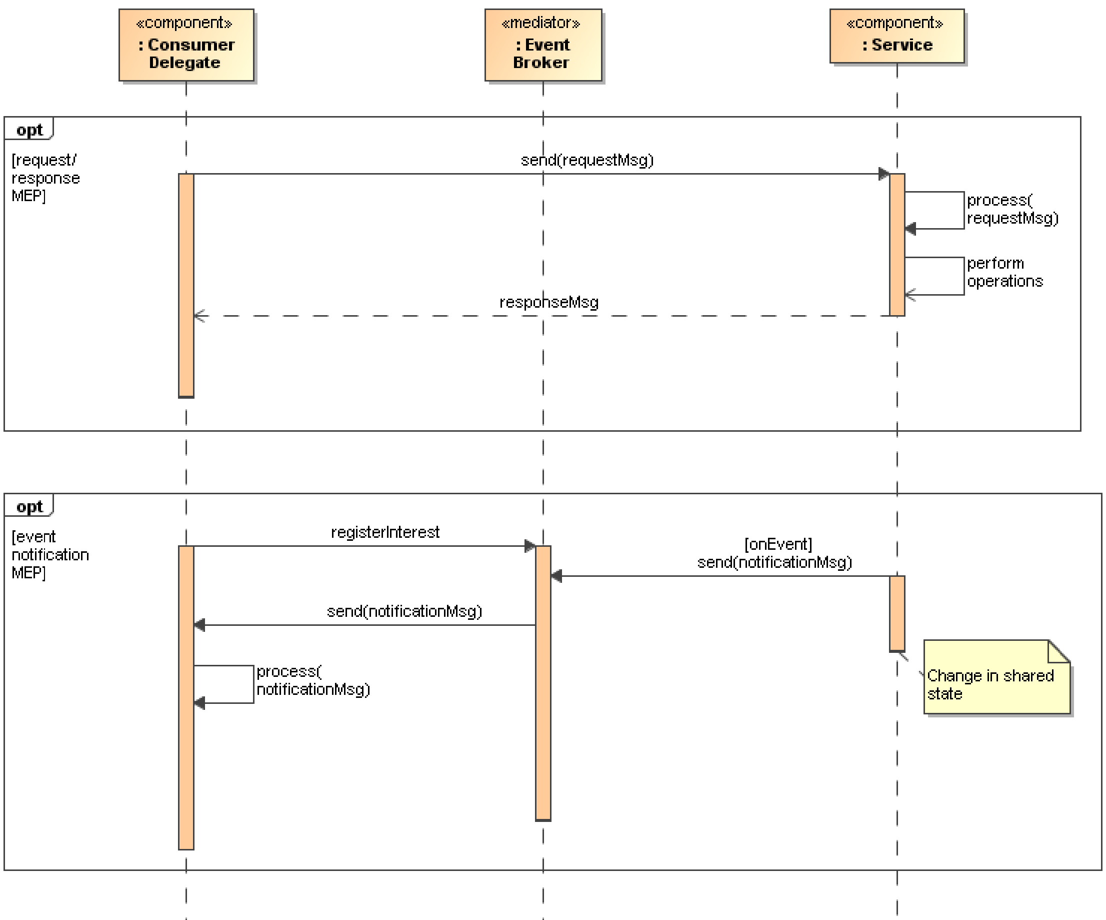

Figure 28 - Fundamental SOA message exchange patterns (MEPs)

Recall from the ~~OASIS~~ Reference Model that the Process Model characterizes “the temporal relationships between and temporal properties of actions and event~~events~~s associated with interacting with the service.” Thus, MEPs are a key element of the Process Model. The meta-level aspects of the Process

 Copyright © OASIS Open 2011. All Rights Reserved.

Model (just as with the Action Model) are provided as part of the Service Description Model (see Section 4.1).

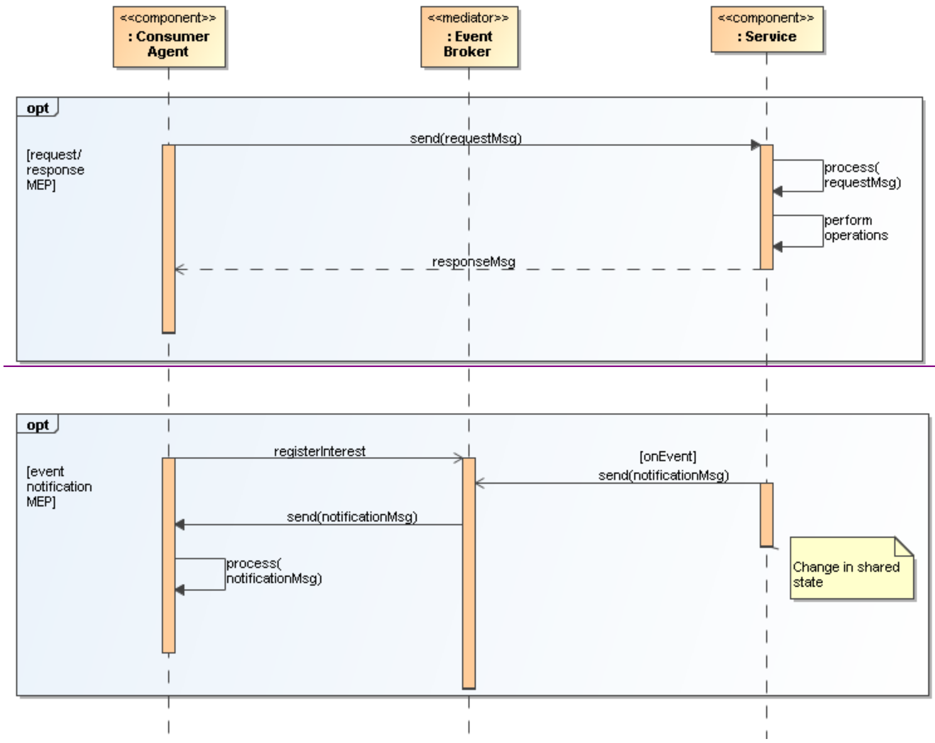

~~*Figure Fundamental SOA message exchange patterns (MEPs)*~~

In the UML sequence diagram shown in Figure 28 it is assumed that the service participants~~participants~~ (consumer and provider) have delegated message handling to hardware or software ~~agents~~delegates acting on their behalf. In the case of the service consumer~~service consumer,~~, this is represented by the *Consumer ~~Agent~~Delegate* component. In the case of the service provider, the delegate~~agent~~ is represented by the *Service* component. The message interchange model illustrated represents a logical view of the MEPs and not a physical view. In other words, specific hosts, network protocols, and underlying messaging system are not shown as these tend to be implementation specific. Although such implementation-specific elements are considered outside the scope of this ~~Reference Architecture~~document, they are important considerations in modeling the SOA execution context. Recall from the Reference Model that the *execution context* of a service interaction is “the set of infrastructure elements, process entities, policy assertions and agreements that are identified as part of an instantiated service interaction, and thus forms a path between those with needs and those with capabilities.”

~~4.3.2.2~~4.3.3.2 Request/Response MEP

In a request/response MEP, the Consumer ~~Agent~~Delegate component sends a request message to the Service component. The Service component then processes the request message. Based on the content of the message, the Service component performs the service ~~operations.~~operation and the associated private actions. Following the completion of these operations, a response message is

 Copyright © OASIS Open 2011. All Rights Reserved.

returned to the Consumer ~~Agent~~Delegate component. The response could be that a step in a process is complete, the initiation of a follow-on operation, or the return of requested information.[17]

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

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

~~4.3.2.3~~4.3.3.3 Event Notification MEP

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

Often the sending ~~agent~~component may not be fully aware of all the ~~agents~~components that ~~will~~wish to receive the notification; particularly in so-called publish/subscribe ("pub/sub") situations. In event notification message exchanges, it is rare to have a tightly-coupled link between the sending and the receiving ~~agent~~component(s) for a number of practical reasons. One of the most common needs for pub/sub messaging is the potential for network outages or communication interrupts that can result in loss of notification of event~~events.~~ s. Therefore, a third-party ~~agent~~mediator component is ~~usually~~often used ~~that serves as an intermediary that may have the ability to store event notification messages and serves~~ to decouple the sending and ~~received agents.~~receiving components.

Although this is typically an implementation issue, because this type of third-party decoupling is so common in event-driven systems, ~~we felt that for this Reference Architecture,~~ it ~~was~~is warranted for use in modeling this type of message exchange~~.~~ in the SOA-RAF. This third-party intermediary is shown in Figure 28 as an Event Broker mediator. As with the request/response MEP, no distinction is made between synchronous versus asynchronous communication, although asynchronous message exchange is illustrated in the UML sequence diagram depicted in Figure 28~~.~~ .

~~4.3.3~~4.3.4 Composition of Services

Composition of services is the act of aggregating or "composing" a single service from one or more other services. ~~Before we provide an architectural~~A simple model of service composition~~, it~~ is illustrated in Figure 29~~important that we distinguish two fundamentally different types of services, *atomic services* and *composite services*.~~.

~~**Atomic Service**~~

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

[17] There are cases when a response is not always desired and this would be an example of a "one-way" MEP. Similarly, while not shown here, there are cases when some type of "callback" MEP is required in which the consumer agent is actually exposed as a service itself and is able to process incoming messages from another service.

 Copyright © OASIS Open 2011. All Rights Reserved.

~~**Composite Service**~~

> ~~A service visible to a service consumer (or agent) via a single interface and described via a single service description that is the aggregation or composition of one or more other services. These other services can be atomic services, other composite services, or a combination of both.~~~~[18]~~

~~From the consumer's point of view, the distinction is, of course, mostly irrelevant. The consumer still interacts with a composite service via a single interface and utilizes the meta-level information about the composite service provided by a single Service Description. Nevertheless, there are important dependencies that need to be considered in services that utilize other services such as propagation of policy constraints, security profiles, etc.~~

~~A simple model of service composition is illustrated in~~

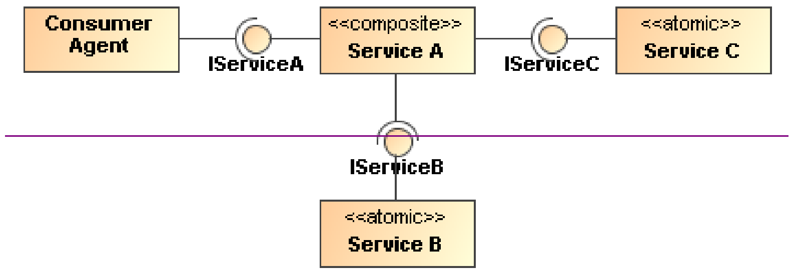

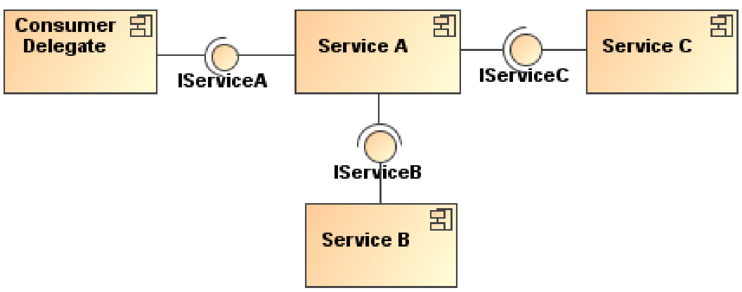

Figure 29 - Simple model of service composition ~~*("public" composition).*~~

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

It is possible for a service composition to be opaque from one perspective and transparent from another. For example, a service may appear to be a single service from the Consumer's Delegate's~~Consumer~~

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

 Copyright © OASIS Open 2011. All Rights Reserved.

~~Agent's~~ perspective, but is transparently composed of one or more services from a service management perspective. A Service Management ~~Service~~capability needs to be able to have visibility into the composition in order to properly manage the dependencies between the services used in constructing the composite service—including managing the service's lifecycle. The subject of services as management entities is described and modeled in the ~~Owning Service Oriented Architectures~~*Ownership in a SOA Ecosystem* View of ~~this Reference Architecture~~the SOA-RAF and ~~will~~is not ~~be~~ further elaborated ~~here.~~in this section. The point to be made here is that there can be different levels of opaqueness or transparency when it comes to visibility of service composition.

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

~~4.3.3.1~~4.3.4.1 Service-Oriented Business Processes

The concepts of business processes and collaborations in the context of transactions and exchanges across organizational boundaries are described and modeled as part of the ~~Business via Services View~~*Participation in a SOA Ecosystem* view of this ~~Reference Architecture~~reference architecture (see Section 3). Here, we focus on the belief that the principle of composition of services can be applied to business processes and collaborations. Of course, business processes and collaborations traditionally represent complex, multi-step business functions that may involve multiple participants~~participants,~~, including internal users, external customers, and trading partners. Therefore, such complexities cannot simply be ignored when transforming traditional business processes and collaborations to their service-oriented variants.

Business Processes

> Business processes are ~~comprised of~~ a set of ~~coherent~~one or more linked activities that~~, when~~ are performed ~~in a logical sequence over a period of time and with appropriate rules applied, result in~~to achieve a certain business outcome.

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

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

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

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

Orchestration

 Copyright © OASIS Open 2011. All Rights Reserved.

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

An orchestration is typically implemented using a scripting approach to compose service-oriented business processes. This typically involves use of a standards-based orchestration scripting language. ~~An example of such a language is the Web Services Business Process Execution Language (WS-BPEL) **[WS-BPEL].**~~ In terms of automation, an orchestration can be mechanized using a business process orchestration engine, which is a hardware or software component (delegate~~(agent)~~) responsible for acting in the role of central conductor/coordinator responsible for executing the flows that comprise the orchestration.

A simple generic example of such an orchestration is illustrated in Figure 30.

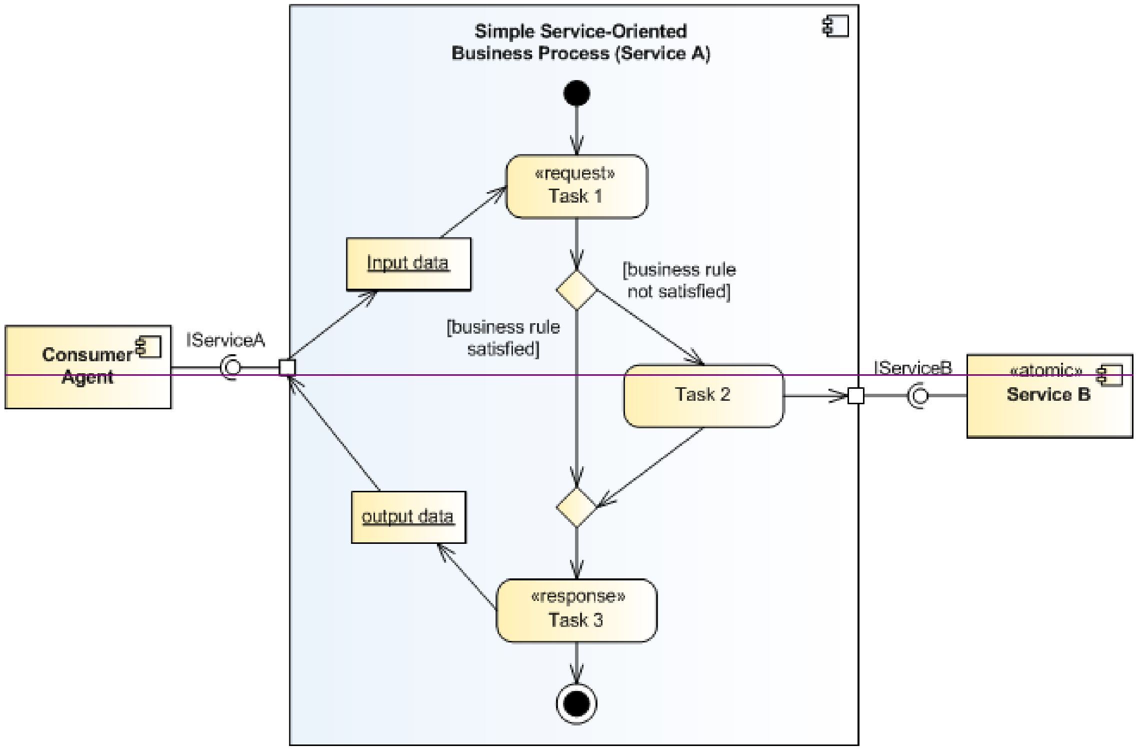

Copyright © OASIS Open 2011. All Rights Reserved.

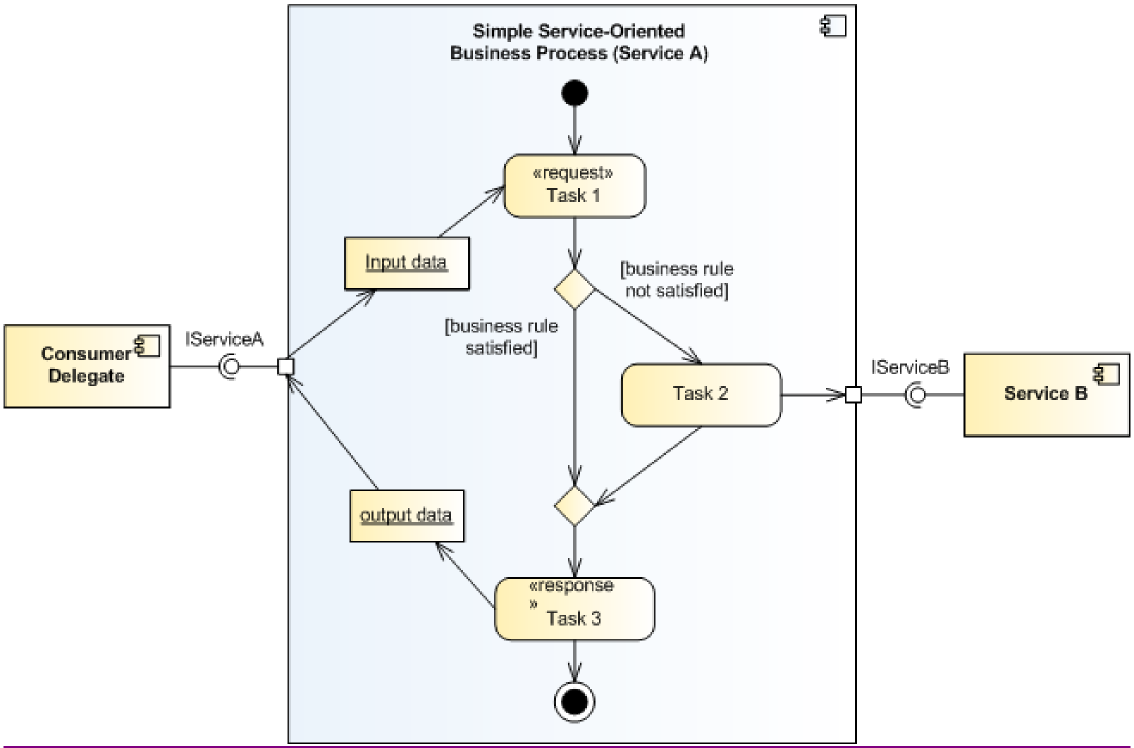

Figure 30 - Abstract example of orchestration of service-oriented business process~~.~~

Here, we use a UML activity diagram to model the simple service-oriented business process as it allows us to capture the major elements of business processes such as the set of related tasks to be performed, linking between tasks in a logical flow, data that is passed between tasks, and any relevant business rules that govern the transitions between tasks. A task is a unit of work that an individual, system, or organization performs and can be accomplished in one or more steps or subtasks. While subtasks can be readily modeled, they are not illustrated in the orchestration model in Figure 30.

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

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

~~4.3.3.2~~4.3.4.2 Service-Oriented Business Collaborations

~~Turning our attention to business collaborations we note that business~~ Business collaborations typically represent the interaction involved in executing business transaction~~business transactions, where a *business transaction* is defined in the Business via Services View as “a joint action engaged in by two or more participants in which resources are exchanged” (see Section).~~s.

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

~~Similar to service-enablement of business processes, business~~Business collaborations can also be service-enabled. For purposes of this Reference Architecture Foundation, we refer to these ~~types of business collaborations~~ as “service-oriented business collaborations.” ~~Of course, unlike service-oriented business processes, the concept of service~~Service-oriented business collaborations ~~does~~do not

 Copyright © OASIS Open 2011. All Rights Reserved.

necessarily imply exposing the entire peer-style business collaboration as a service itself but rather the collaboration uses service-based interchanges.

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

Choreography

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

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

~~As is the case of an orchestration, a choreography is typically implemented by using a scripting approach to composing service-oriented business collaborations. This typically involves use of a standards-based choreography scripting language. An example of such a language is the Web Services Choreography Description Language~~ ~~**[WS-CDL]**~~.

A simple generic example of a choreography is illustrated in Figure 31.

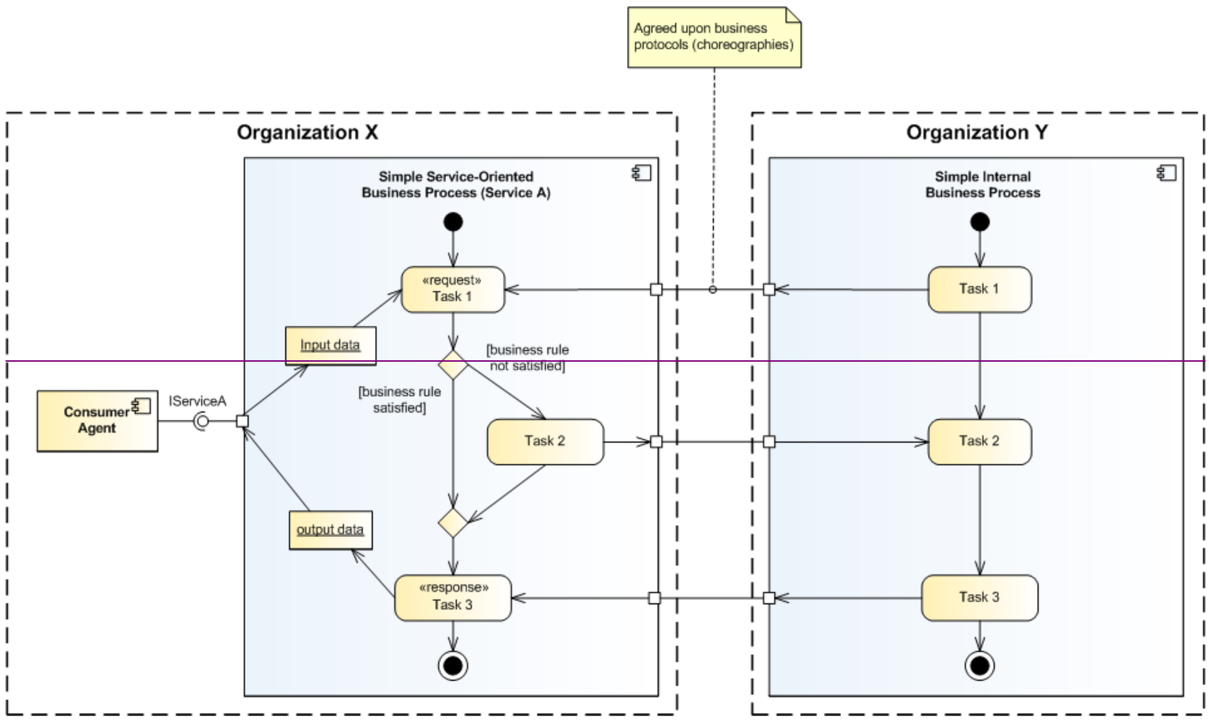

Copyright © OASIS Open 2011. All Rights Reserved.

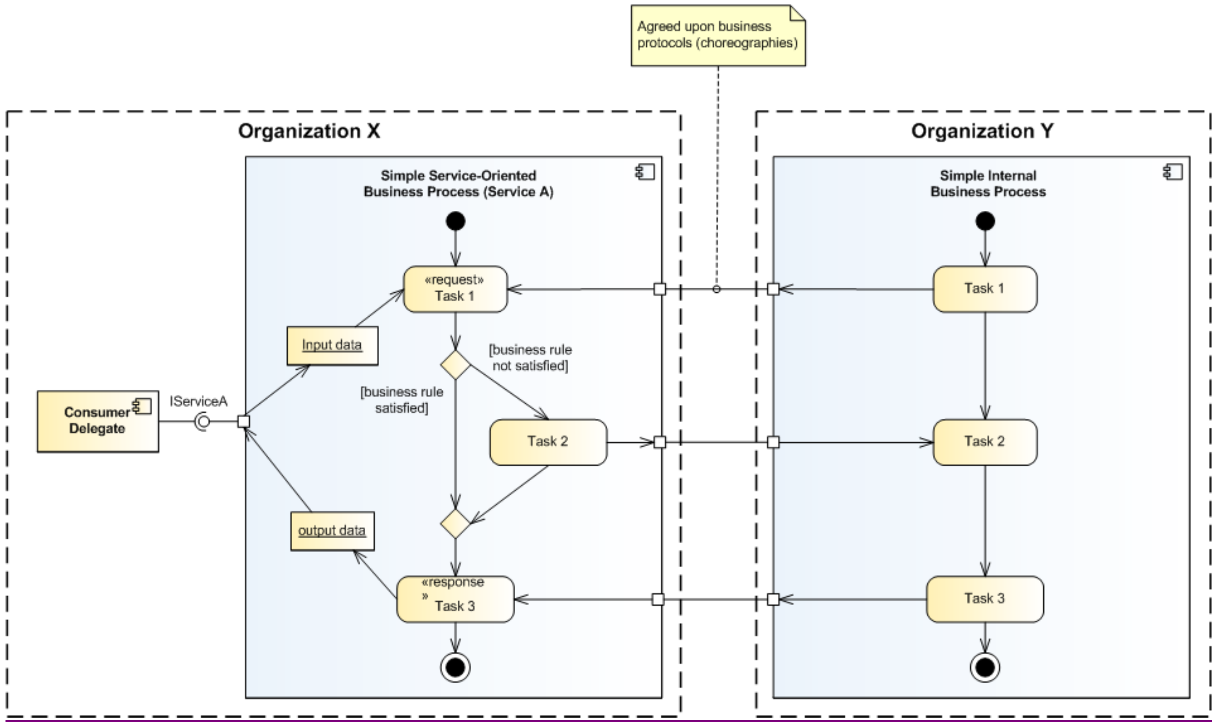

Figure 31 - Abstract example of choreography of service-oriented business collaboration~~.~~

This example, which is a variant of the orchestration example illustrated earlier in Figure 30 adds trust boundaries between two organizations; namely, Organization X and Organization Y. It is assumed that these two organizations are peer entities that have an interest in a business collaboration, for example, Organization X and Organization Y could be trading partners. Organization X retains the service-oriented business process Service A, which is exposed to internal consumers via its provided service interface, IServiceA. Organization Y also has a business process that is involved in the business collaboration; however, for this example, it is an internal business process that is not exposed to potential consumers either within or outside its organizational boundary.

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

~~If,~~In a peer-style business collaboration~~ in which visibility into and use of each participating organization's internal service-oriented business processes was necessary as part of an end-to-end business transaction, then it would be desirable to select~~, a choreography scripting language must be capable of describing the coordination of those service-oriented processes that ~~would support interaction between different orchestration engines that spans~~cross organizational boundaries.~~ WS-CDL is an example~~

Architectural Implications of ~~such a language.~~

~~4.4~~1.1 Interacting with ~~Policies and Contracts Model~~

4.3.5 Services

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

- A well-defined service Information Model that:
    - describes the syntax and semantics of the messages used to denote actions and events;
    - describes the syntax and semantics of the data payload(s) contained within messages;

 Copyright © OASIS Open 2011. All Rights Reserved.

    - documents exception conditions in the event of faults due to network outages, improper message/data formats, etc.;
    - is both human readable and machine processable;
    - is referenceable from the Service Description artifact.
- A well-defined service Behavior Model that:
    - characterizes the knowledge of the actions invokes against the service and events that report real world effects as a result of those actions;
    - characterizes the temporal relationships and temporal properties of actions and events associated in a service interaction;
    - describe activities involved in a workflow activity that represents a unit of work;
    - describes the role (s) that a role player performs in a service-oriented business process or service-oriented business collaboration;
    - is both human readable and machine processable;
    - is referenceable from the Service Description artifact.
- Service composition mechanisms to support orchestration of service-oriented business processes and choreography of service-oriented business collaborations such as:
    - Declarative and programmatic compositional languages;
    - Orchestration and/or choreography engines that support multi-step processes as part of a short-lived or long-lived business transaction;
    - Orchestration and/or choreography engines that support compensating transactions in the presences of exception and fault conditions.
- Infrastructure services that provides mechanisms to support service interaction, including but not limited to:
    - mediation services such as message and event brokers, providers, and/or buses that provide message translation/transformation, gateway capability, message persistence, reliable message delivery, and/or intelligent routing semantics;
    - binding services that support translation and transformation of multiple application-level protocols to standard network transport protocols;
    - auditing and logging services that provide a data store and mechanism to record information related to service interaction activity such as message traffic patterns, security violations, and service contract and policy violations
    - security services that provide centralized authorization and authentication support, etc., which provide protection against common security threats in a SOA ecosystem;
    - monitoring services such as hardware and software mechanisms that both monitor the performance of systems that host services and network traffic during service interaction, and are capable of generating regular monitoring reports.
- A layered and tiered service component architecture that supports multiple message exchange patterns (MEPs) in order to:
    - promote the industry best practice of separation of concerns that facilitates flexibility in the presence of changing business requirements;

promote the industry best practice of separation of roles~~As described in the Reference Model, a policy is the representation of a constraint or condition on the use, deployment, or description of an owned entity as defined by any participant. A contract is a representation of an agreement between two or more participants. Technically, the only difference between a policy and a contract is the agreement between two or more parties to a contract and the enforceability of a policy by one party on other parties.~~

~~In Section , Policies and contracts are discussed in the context of the Business via Services View with generalizations about IT mechanisms in support of the view. Section breaks down a core aspect of policies, a proposition, and provides the basis for the IT mechanisms discussed in Section . Section concludes with some general policy and contract principles common to SOA policies.~~

~~4.4.1 Automating Support for Policies and Contracts~~

~~Policy and contract IT mechanisms support automated governance and management within the SOA ecosystem to improve governance and management efficiency. Understanding the complete environment which policies and contracts apply in a SOA requires understanding of the processes surrounding policies and contracts in the social structure, the IT mechanisms that support automated~~

~~enforcement of policies and contracts, and the traversal from/to the social structure to/from the IT policy automation mechanisms. The architecture SHOULD provide mechanisms to enforce policies and contracts to ensure efficient operations consistent with the goals of the social structure.~~

~~derives from Section , . Core aspects of policies and contracts are the propositions, the owners, and the measurement and enforcement of the policy or contract. In Section , , measurable assertions and commitments are characterized as propositions - an expression of some property of the world whose truth can be measured by examining the world and checking that the expression and the world are consistent with each other. Assertions are claims about current state while commitments are agreements to future state.~~

~~Figure Distinguishing between policies and contracts~~

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

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

~~The measurability and enforcement of propositions may include many indirectly related participants within the social structure. Dispute resolutions, for example, may involve courts.~~

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

~~Policies and contracts have wide applicability within the Reference Architecture. They are used to express security policies, service policies, relationships and constraints within the social structures that encapsulate service participants, management of services and many other instances. The enforcement of a policy or contract may be a part of the SOA-based computing environment or it may be handled outside of the SOA-based computing environment. The Reference Architecture is concerned with the underlying IT mechanisms and principles that support enforceable and measurable contracts and policies in the widest range of situations for a SOA.~~

- ~~Policy and~~ in a service development lifecycle such that subject matter experts and teams are structured along areas of expertise;
- support numerous standard interaction patterns, peer-to-peer interaction patterns, enterprise integration patterns, and business-to-business integration patterns.

~~4.4.2 Contract Types~~

~~depicts assertions and commitments as an aggregation of measurable constraints. We can analyze policy and contract constraints in a number of dimensions: positive constraints vs. negative constraints; and permission-style vs. obligation-style constraints.~~

 Copyright © OASIS Open 2011. All Rights Reserved.

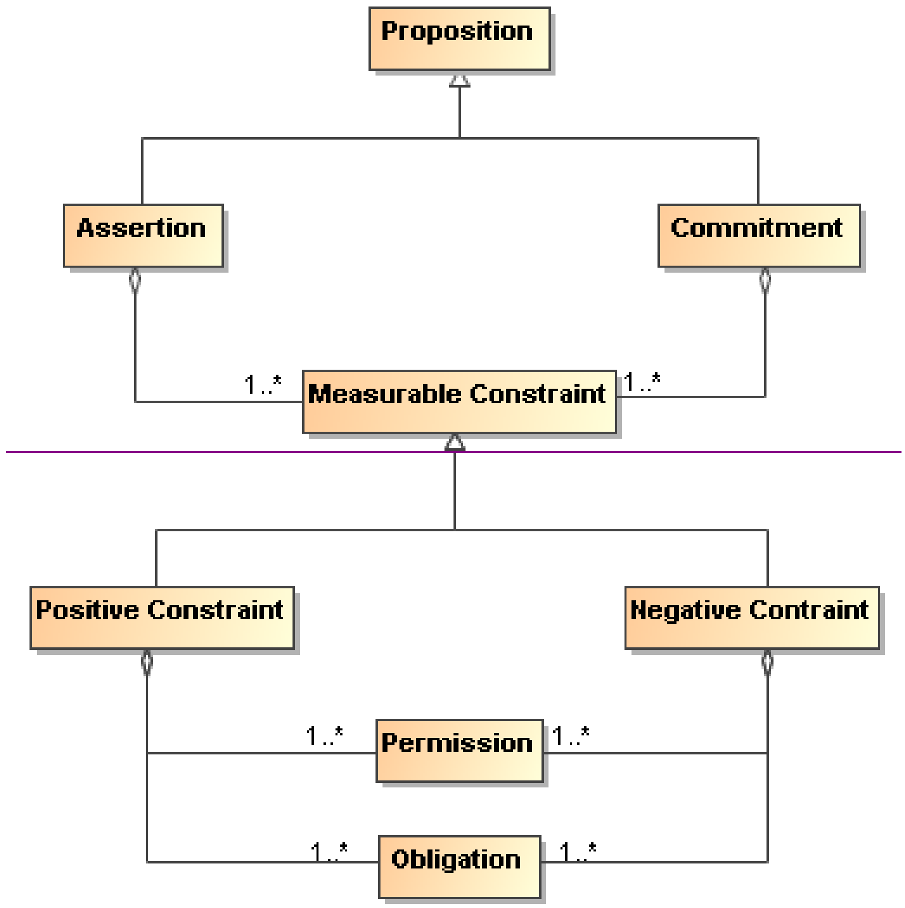

~~Figure Policy and Contract Constraints~~

~~Positive constraints are about the things that you may/should do and negative constraints are about the things that you should not do. A permission-style constraint is about the right to access some resource or perform some action; an obligation-style constraint is about the requirement to perform some action or maintain the state of a resource.~~

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

~~Permission-style constraints are often checkable a-priori: before the intended action or access is completed the current permission constraints may be applied to deny the access if necessary. However, obligation-style constraints can normally only be verified post-priori. Permission constraints are sometimes referred to as access control policies given the preponderance of security-related policies in many applications. One use of obligation constraints is for metrics collection and compliance.~~

~~Policies and contracts can contain a mix of permissions and obligations, and, in sufficiently rich policy management frameworks, can be combined in interesting ways: for example, you may be obliged to give permission to certain actions; or you may be permitted to enter into obligations (this is the core of the right to enter into contracts).~~

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

 Copyright © OASIS Open 2011. All Rights Reserved.

~~4.4.3 IT Mechanisms Supporting Policies and Contracts~~

4.4 Policies and Contracts Model

A common phenomenon of many machines and systems is that ~~they are~~ the scope of potential behavior is much broader ~~in their potential~~ than is actually needed for a particular circumstance. This is especially true of a system as powerful as a SOA ecosystem. As a result, the behavior and performance of the system tend to be under-constrained by the implementation~~. Policy statements~~; instead, the actual behavior is expressed by means of policies of some form. Policies define the choices that stakeholder~~a service provider and/or service consumer (or other stakeholder) makes~~s make; these choices are used to guide the actual behavior of the system to the desired behavior and performance.

As noted in Section 3.1.5, a policy is a constraint of some form that is promulgated by a stakeholder who has the responsibility of ensuring that the constraint is enforced. In contrast, contracts are **agreements** between participants. However, like policies, it is a necessary part of contracts that they are enforceable.

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

4.4.1 Policy and Contract Representation

A **policy constraint** is a specific kind of constraint: the ontology of policies and contracts includes the core concepts of permission, obligation~~While there are many possible approaches to the realization of policy/contracts for a SOA, one approach based on current policy standardization efforts is depicted in this section. The common policy architectural elements that are provided in this section are based on the minimal mechanisms required to provide policy guided delivery across distributed services within an ownership domain and across ownership domains.~~

~~4.4.3.1 Permission Based Policy and Contract Mechanisms~~

~~For IT mechanisms, policies and contracts are measurable and enforceable rules that define choices in the behavior of a system. Contracts are the set of rules that define the agreements under which service functionality is delivered. depicts mechanisms in support of permission style policy requests where the measurement of rules occurs in decision procedures identified by a Decision Point mechanism in the diagram.~~

~~Figure Permission Policy Mechanisms~~

~~Policy/Contract Administration Point~~

~~A Policy/Contract Administration Point is the mechanism for a SOA that allows a participant to administer policies for storage and/or distribution. There can be many enterprise SOA policy/contract administration capabilities and the Policy/Contract Administration Point is a generalization for any of these type of capabilities.~~

~~Policy Distribution/Repository~~

~~The Policy Distribution/Repository distributes policy to decision points or stores policies for retrieval by decision points.~~

~~Attribute Information Point~~

~~The Attribute Information Point is responsible for collecting and forwarding attributes to the Decision Point. Attributes are named values that define characteristics of participants, resources, actions, or the environment. Attributes are defined in the Service Description Model in Section .~~

~~Audit Point~~

~~In Figure 33, the Audit Point is any mechanism that records participant actions requiring permission decisions or records the measurement results for obligations discussed in Section .~~

 Copyright © OASIS Open 2011. All Rights Reserved.

~~An auditing mechanism may store audited information and/or provide event notifications of audited information. Auditing may be used for activities like forensic investigation and regulatory compliance.~~

~~**Resource**~~

~~A resource is any entity of some perceived value. Resources are defined in the Resource Model in Section .~~

~~**Decision Point**~~

~~The Decision Point evaluates participant requests against relevant policies/contracts and attributes to render a permission decision. The Decision Point provides a measurement for an assertion. The Decision Point generally renders a permission decision in the form of permit, deny, indeterminate, not applicable, or a set of obligations. A Decision Point may obtain a permission decision from a computing mechanism or from outside the computing system, decisions by people through workflow for example.~~

~~**Enforcement Point**~~

~~The Enforcement Point enforces and assures the Decision Point decisions and obligations. In a Service Oriented Architecture, one policy or contract may be applicable to multiple distributed services. Due to the distributed nature of a SOA, the enforcement of permission decisions is attributed to an Enforcement Point that is separate from the Decision Point. One Decision Point can provide decisions for many distributed Enforcement Points.~~

~~For permission decisions, the Enforcement Point often performs enforcement in the form of protecting access and determining access compliance to one or more resources. When attempting to access a resource, the Enforcement Point sends a description of the attempted access to a Decision Point. The Decision Point evaluates the request against its available policies/contracts and produces a permission decision that is returned to the Enforcement Point. Like the Decision Point, an Enforcement Point may require a means of enforcement outside the computing system.~~

~~4.4.3.2 Obligation Based Policy and Contract Mechanisms~~

~~In , the Enforcement Point creates or uses a mechanism for measuring policy obligations. Just as it is the responsibility of the Enforcement Point to ensure permission decisions, it is the responsibility of the Enforcement Point to ensure that policy obligations are met. This may require a one time measurement or ongoing monitoring of the obligation. For example, there may be the contractual obligation to allocate a certain level of bandwidth for a customer's transactions. The contractual obligation may also require ongoing monitoring to ensure the customer's transactions do not exceed allotted bandwidth and if exceeded, the provider may happily levy exorbitant over usage fees.~~

~~While depicts measurement of obligations based on an access request, the Enforcement Point may acquire policy obligations independent of permission requests from other participants. To provide a real-world analogy, a consciences taxicab owner may have a policy that taxis not operate when the roads are icy. At the start of a working day, the roads are clear but the forecast is for possible icy conditions later in the day. A dispatcher, a designated Enforcement Point, asks the owner, a Decision Point, whether they should send taxicabs out for the day. The owner says yes as long as the weather reports do not indicate there could be icy roads. The dispatcher checks a website which provides registry listings of service providers that provide reports for local road conditions. The dispatcher chooses a local traffic reporting service, a Measurement Point, that will send traffic reports via email about the road conditions. The dispatcher goes on with his job not worried about checking weather conditions, correctly or incorrectly relying on the email notification to meet the taxicab company's obligation as to the safety of its drivers.~~

~~Figure~~ , owner, subject. In addition, it may be necessary to be able combine policy constraints and to be able to resolve policy conflicts.

4.4.1.1 Policy Framework

Policy Framework

A policy framework is a language in which ~~Obligation~~ ~~Policy Mechanisms~~

 Copyright © OASIS Open 2011. All Rights Reserved.

~~**Measurement Point**~~

~~The Measurement Point identifies mechanisms for measuring and monitoring policy obligations.~~

~~The Measurement Point in receives and responds to the Enforcement Point requests to measure policy obligations. The Measurement Point may also audit and provide event notifications of obligation measurements.~~

~~In , the Measurement Point can be used to collect metrics and report those metrics to the Audit Point. Metrics may be used to verify compliance either in an automated fashion or at a later point in time. If compliance is automated, then the Measurement Point may adjust the behavior of the system in accordance with compliance policies or contracts.~~

~~4.4.4 Policy and Contract Principles~~

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

~~4.4.4.1 Policies and Contracts Goals~~

~~Policies SHOULD reflect the goals of governance or management processes, see Section and section . The governance and management processes SHOULD use formal and standardized policy languages to enable the widest possible understanding and use of stated policies and contracts, and architecture components SHOULD be available to enable compliance.~~

~~4.4.4.2 Policy and Contract Specification~~

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

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

- ~~how to describe atomic policy constraints~~
- ~~how to nest policy constraints allowing for abstractions and refinements of a policy constraint~~

~~how to reference~~ policy constraints ~~allowing for the reuse of a policy constraint~~may be expressed.

~~how to define alternative~~A policy framework combines a syntax for expressing policy constraints together with a decision procedure for determining if a policy constraint is satisfied.

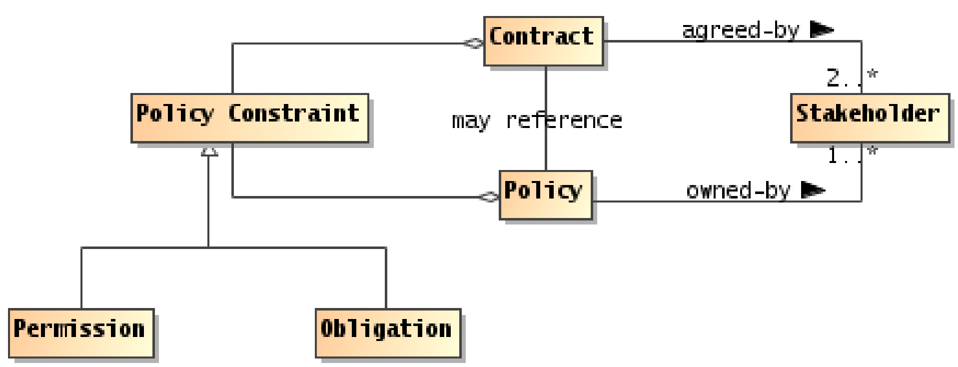

Figure 32~~the selection~~ - Policies and Contracts

We can characterize (caricature) a policy framework in terms of ~~compatible~~a logical framework and an ontology of policies. The policy ontology details specific kinds of policy constraints ~~between the consumer~~

 Copyright © OASIS Open 2011. All Rights Reserved.

~~and provider~~that can be expressed; and the logical framework is a 'glue' that allows us to express combinations of policies.

Logical Framework

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

Policy Ontology

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

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

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

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

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

Policy Owner

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

Policy Subject

A policy subject is an actor~~policy versioning~~

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

Policy Constraint

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

Policy Object

A policy object is an identifiable state, action or resource~~modules~~

~~4.4.4.3 Policy Composition~~

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

4.4.2 Policy and Contract Enforcement

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

4.4.2.1 Enforcing Simple Policy Constraints

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

 Copyright © OASIS Open 2011. All Rights Reserved.

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

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

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

Policy Decision

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

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

Policy Enforcement

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

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

~~4.4.4.4~~4.4.2.2 Conflict Resolution

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

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

Policy Conflict

> A policy conflict exists between ~~the policy rules. There can be many causes for~~ two or more policy constraints in a policy decision process if the satisfaction of one or more policy constraints leads directly to the violation of one or more other policy constraints.

Policy Conflict Resolution

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

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

4.4.3 Architectural Implications

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

- There SHOULD be a standard policy framework that is adopted across ownership domains ~~and~~ within the SOA ecosystem:
  - This framework MUST permit the expression of simple policy ~~language specifications that do not convert~~ constraints
  - The framework MAY allow (to ~~first order predicate logic for IT~~a varying extent) the combination of policy constraints, including

 Copyright © OASIS Open 2011. All Rights Reserved.

        - Both positive and negative constraints
        - Conjunctions and disjunctions of constraints
        - The quantification of constraints
    - The framework MUST at least allow the policy subject and the policy object to be identified as well as the policy constraint.
    - The framework MAY allow further structuring of policies into modules, inheritance between policies and so on.
- There SHOULD be mechanisms~~. This can cause policy decision results to be indeterminate. Policy administration~~ that facilitate the application of policies:
    - There SHOULD be mechanisms ~~may provide conflict resolution capabilities prior to the storage/distribution of policies. At run time, conflicts may propagate to higher authorities inside or outside the SOA-based IT~~that allow policy decisions to be made, consistent with the policy frameworks.
    - There SHOULD be mechanisms~~.~~ to enforce policy decisions
        - There SHOULD be mechanisms to support the measurement of whether certain policy constraints are satisfied or not, or to what degree they are satisfied.

~~4.4.4.5~~ Such enforcement mechanisms MAY include support for both permission~~Delegation of Policy~~

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

~~4.4.5 Architectural Implications~~

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

- ~~Policy and Contract language specifications will typically provide support for the following capabilities:~~
    - ~~expression of assertion and commitment policy constraints;~~
    - ~~expression of positive and negative policy~~-style constraints~~;~~
        - ~~expression of permission~~ and obligation ~~policy~~-style constraints~~;~~.
    - ~~nesting of~~Enforcement mechanisms MAY support the simultaneous enforcement of multiple policy constraints ~~allowing for abstractions and refinements of a policy constraint;~~
    - ~~definition of alternative policy constraints to allow for the selection of compatible policy constraints for a consumer and provider;~~
    - ~~composition of policies to combine one or more policies.~~
        - ~~Policy and contract mechanisms in a SOA~~ across multiple points in the SOA ecosystem ~~will require the following capabilities:~~.
    - ~~decision procedures which must~~There SHOULD be ~~able~~mechanisms to ~~measure and render decisions on constraints;~~
    - ~~enforcement of decisions;~~
    - ~~measurement and notification of obligation constraints;~~
    - ~~auditability of decisions, enforcement, and obligation measurements;~~
    - ~~administration of~~resolve policy ~~and contract language artifacts;~~
    - ~~storage of policies and contracts;~~
    - ~~distribution of policies/contracts;~~
    - ~~conflict resolution or elevation of~~ conflicts ~~in policy rules;~~
        - ~~delegation of~~ This MAY involve escalating policy ~~authority~~conflicts to ~~agents acting on behalf of a client;~~human adjudication.
    - ~~decision procedures capable of incorporating roles and/or attributes for rendered decisions.~~

 Copyright © OASIS Open 2011. All Rights Reserved.

- ~~Owning Service Oriented Architectures~~There SHOULD be mechanisms that support the management and promulgation of policies.

 Copyright © OASIS Open 2011. All Rights Reserved.

5 Ownership in a SOA Ecosystem 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 *Owning Service Oriented Architectures*~~*Governments are instituted among Men, deriving their just power from the consent of the governed*
American Declaration of Independence

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

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

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

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

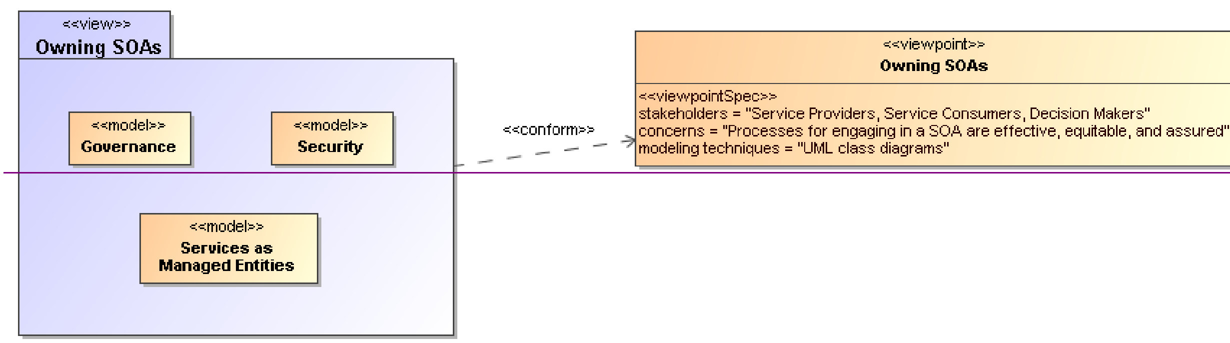

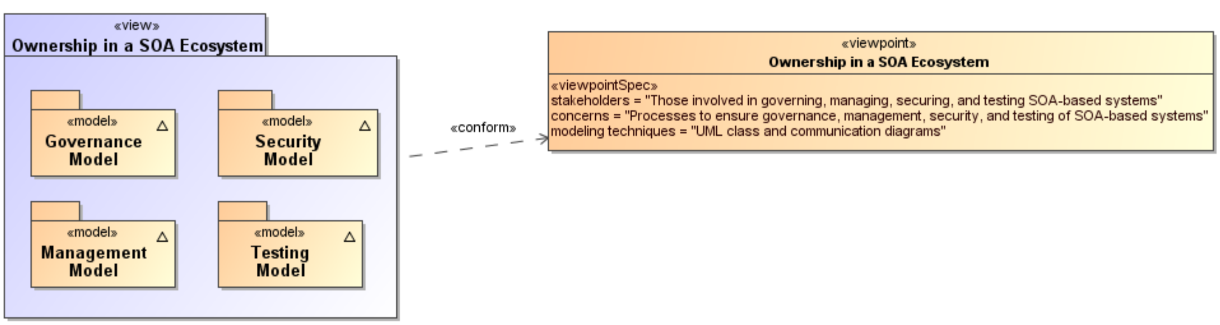

Figure 33 - Model ~~elements described in the Owning Service Oriented Architectures view~~Elements Described in the Ownership in a SOA Ecosystem View

The following subsections present models of these functions.

 Copyright © OASIS Open 2011. All Rights Reserved.

5.1 Governance Model

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

5.1.1 Understanding Governance

5.1.1.1 Terminology

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

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

1. ~~1)~~ what decisions must be made to ensure effective management and use?,
2. ~~2)~~ who should make these decisions?,
3. ~~and 3)~~ how will these decisions be made and monitored? , and
4. how will these decisions be communicated?

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

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

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

5.1.1.2 Relationship to Management

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

5.1.1.3 Why is SOA Governance Important?

One of the hallmarks of SOA that distinguishes it from other architectural paradigms for distributed computing is the ability to provide a uniform means to offer, discover, interact with and use capabilities (as well the ability to compose new capabilities from existing ones) all in an environment that transcends

 Copyright © OASIS Open 2011. All Rights Reserved.

domains of ownership~~ownership.~~. Consequently, ownership~~ownership,~~, and issues surrounding it, such as obtaining acceptable terms and conditions (T&Cs) in a contract, is one of the primary topics for SOA governance. Generally, IT governance does not include T&Cs, for example, as a condition of use as its primary concern.

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

5.1.1.4 Governance Stakeholders and Concerns

As noted in Section 3.1.1~~,~~ the participants~~participants~~ in a service interaction include the service provider, the service consumer~~service consumer,~~, and other interested or unintentional third parties. Depending on the circumstances, it may also include the owners of the underlying capabilities that the SOA services access. Governance must establish the policies and rules under which duties and responsibilities~~responsibilities~~ are defined and the expectations of participants~~participants~~ are grounded. The expectations include transparency in aspects where transparency is mandated, trust in the impartial and consistent application of governance, and assurance of reliable and robust behavior throughout the SOA ecosystem.

5.1.2 A Generic Model for Governance

Governance

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

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

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

5.1.2.1 Motivating Governance

 Copyright © OASIS Open 2011. All Rights Reserved.

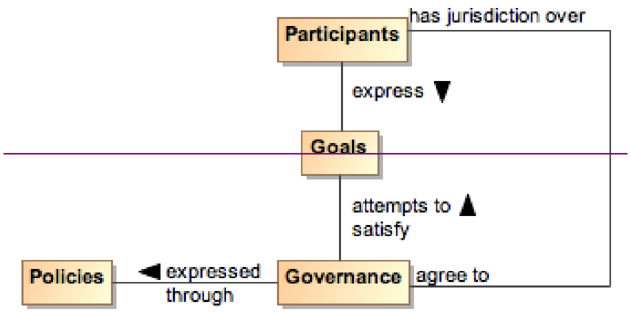

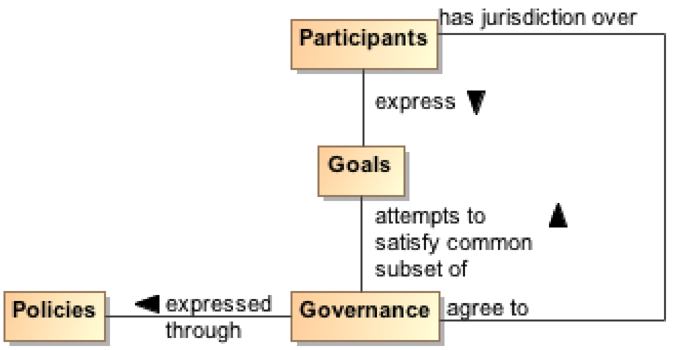

Figure 34 - Motivating ~~governance model~~Governance

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

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

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

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

 Copyright © OASIS Open 2011. All Rights Reserved.

5.1.2.2 Setting Up Governance

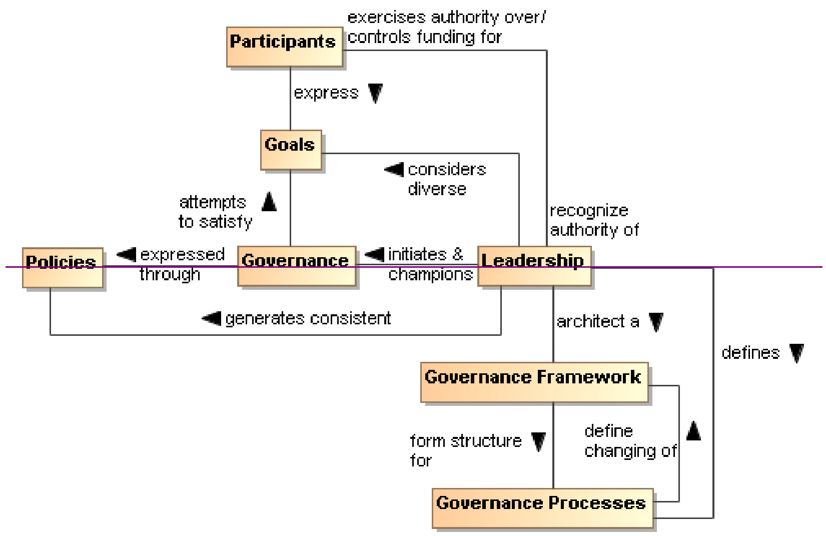

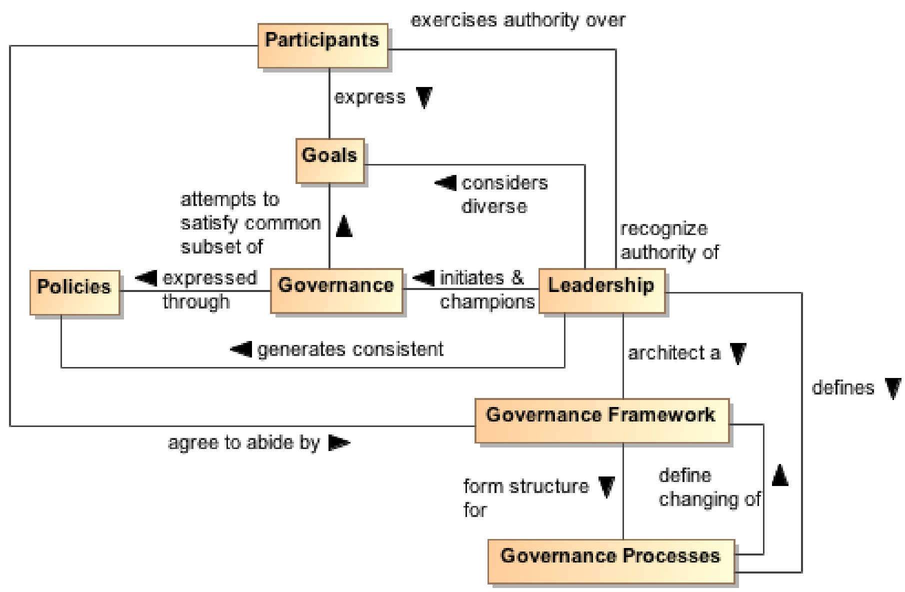

Figure 35 - Setting ~~up~~ Up Governance

Leadership

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

Governance Framework

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

Governance Processes

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

 Copyright © OASIS Open 2011. All Rights Reserved.

As noted earlier, governance requires an appropriate organizational structure and identification of who has authority to make governance decisions. In Figure 35~~the above figure,~~, the entity with governance authority~~authority~~ is designated the Leadership. This is someone, possibly one or more of the participants~~ that Participants~~, which participants recognize as having authority for a given purpose~~authority and who typically has some control~~ or over ~~the Participants~~a given set of issues or concerns.

The Leadership is responsible for prescribing or delegating a working group to prescribe the Governance Framework that forms the structure for Governance Processes ~~that~~which define how governance is to be carried out. This does not itself define the specifics of how governance is to be applied, but it does provide an unambiguous set of procedures that should ensure consistent actions which participants~~Participants~~ agree are fair and account for sufficient input on the subjects to which governance ~~will be~~is applied. ~~Note that the Governance Processes should also include those necessary to modify the Governance Framework itself. The Governance Processes are likely reviewed and agreed to by the Participants.~~

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

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

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

5.1.2.3 Carrying Out Governance

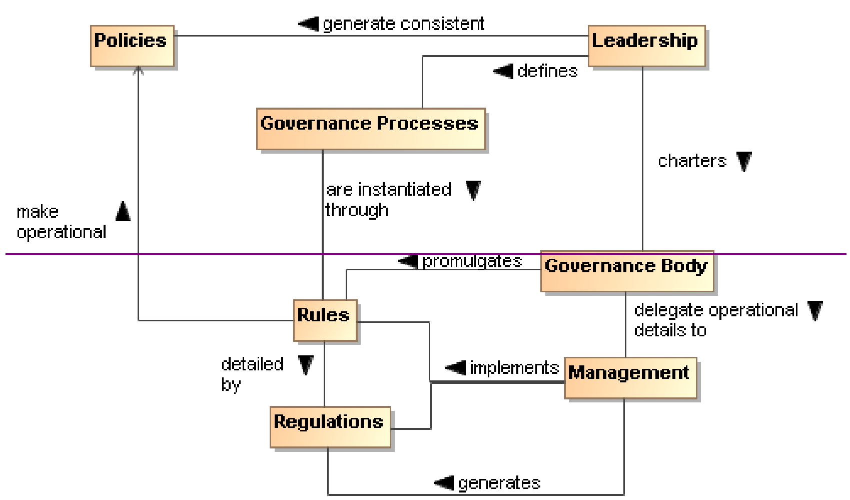

Copyright © OASIS Open 2011. All Rights Reserved.

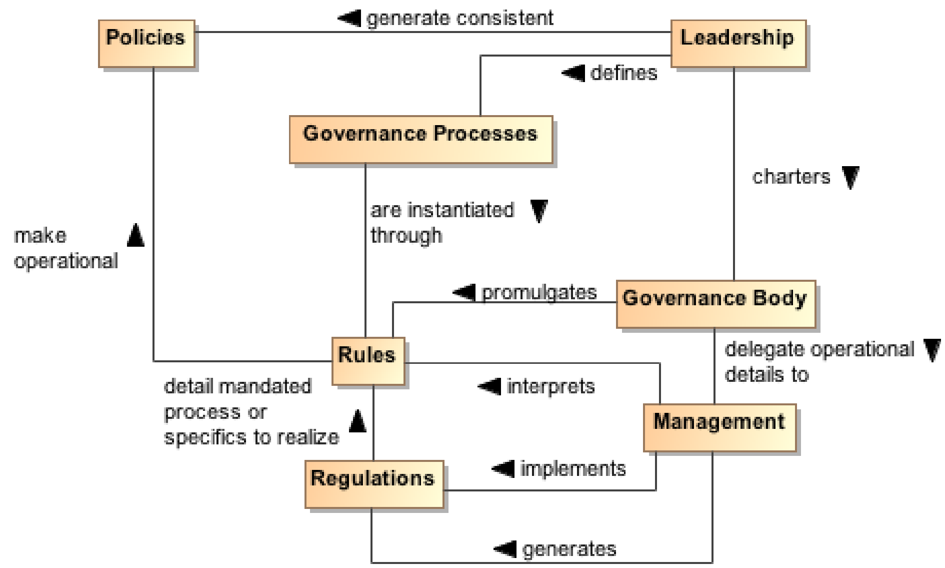

Figure 36 - Carrying Out Governance ~~Model~~

Rule

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

Regulation

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

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

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

Copyright © OASIS Open 2011. All Rights Reserved.

5.1.2.4 Ensuring ~~governance compliance~~Governance Compliance

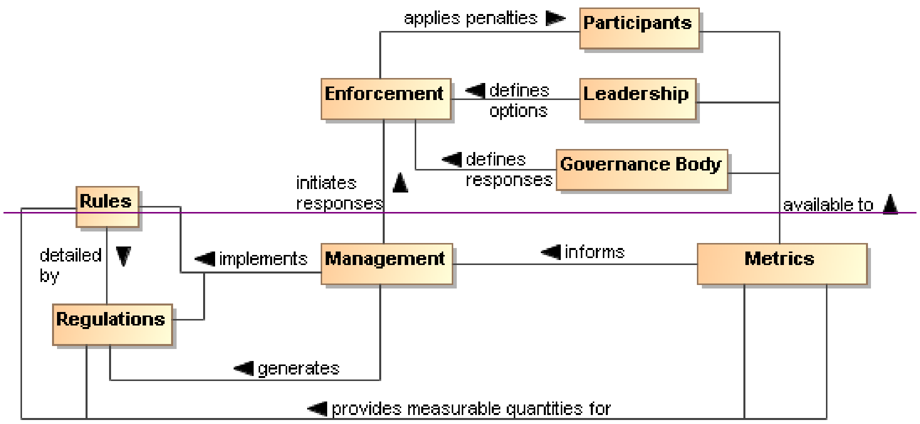

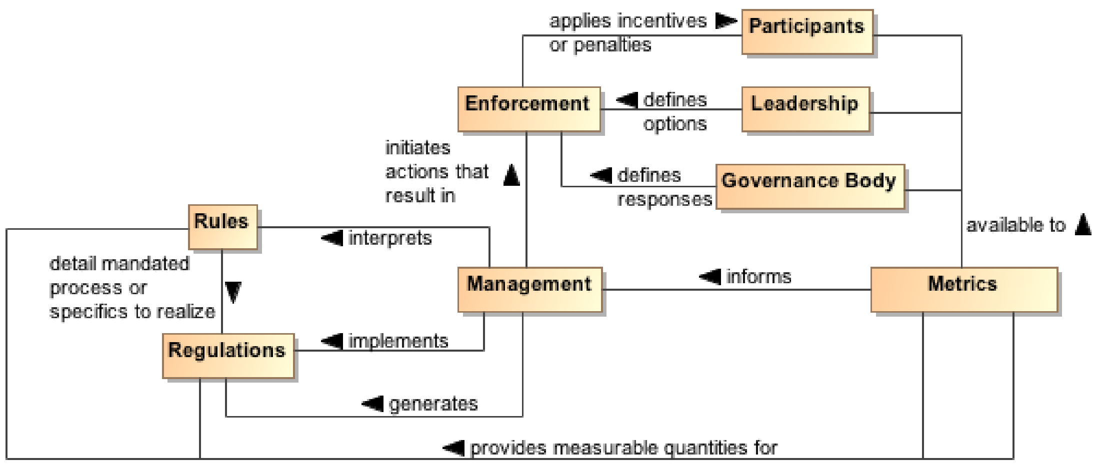

Figure 37 ~~-~~ Ensuring ~~governance compliance model~~Governance Compliance

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

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

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

5.1.2.5 Considerations for Multiple Governance Chains

As noted in section 5.1.2, instances of the governance model often occur as a tiered arrangement, with governance at some level delegating specific authority and responsibility to accomplish a focused portion of the original level's mandate. For example, a corporation may encompass several lines of business and each line of business governs its own affairs in a manner that is consistent with and contributes to the

 Copyright © OASIS Open 2011. All Rights Reserved.

goals of the parent organization. Within the line of business, an IT group may be given the mandate to provide and maintain IT resources, giving rise to IT governance.

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

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

5.1.3 Governance Applied to SOA

5.1.3.1 Where SOA Governance is Different

~~5.1.3~~1.1.1 ~~Governance Applied to SOA~~

~~5.1.3.1~~1.1.1.1 ~~Where SOA Governance is Different~~

SOA governance is often discussed in terms of IT governance, but rather than a parent-child relationship, Figure 38 shows the two as siblings of the general governance described in section 5.1.2. There are obvious dependencies and a need for coordination between the two, but the idea of aligning IT with business already demonstrates that resource providers and resource consumers must be working towards common goals if they are to be productive and efficient. While SOA governance is shown to be active in the area of infrastructure, it is a specialized concern for having a dependable platform to support service interaction; a range of traditional IT issues is therefore out of scope of this document. A SOA governance plan for an enterprise will not of itself resolve shortcomings with the enterprise's IT governance.

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

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

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

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

The specifics discussed in the figures in the previous sections are equally applicable to governance across ~~. A standards development organization such as OASIS is an example of voluntary agreement to governance over a limited domain to satisfy common goals.~~

 Copyright © OASIS Open 2011. All Rights Reserved.

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

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

5.1.3.2 What Must be Governed

~~Participants agree to local goals, and Enforcement can be loss of voting rights or under extreme circumstances, expulsion from the group.~~

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

~~5.1.3.2~~1.1.1.1 ~~What Must be Governed~~

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

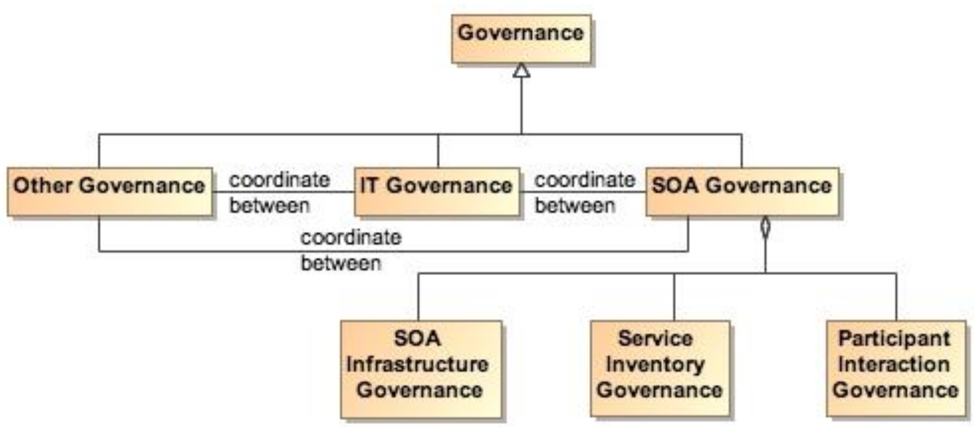

Figure 38 - Relationship Among Types of Governance

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

- SOA infrastructure – the “plumbing” that provides utility functions that enable and support the use of the service
- Service inventory – the requirements on a service to permit it to be accessed within the infrastructure

Participant interaction – the consistent expectations with which all ~~Thus, SOA governance requires establishing confidence and trust while instituting a solid framework that enables flexibility, indicating a combination of strict control over a limited set of foundational aspects but minimum constraints beyond those bounds.~~

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

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

5.1.3.2.1 Governance of SOA ~~infrastructure~~Infrastructure

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

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

5.1.3.2.2 Governance of the ~~service inventory~~Service Inventory

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

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

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

For example, unlimited access for using a service may require a degree of life cycle maturity that has demonstrated sufficient testing over a certain size community. Alternately, the policy may specify that a

 Copyright © OASIS Open 2011. All Rights Reserved.

service in an earlier phase of its life cycle may be made available to a smaller, more technically sophisticated group in order to collect the metrics that would eventually allow the service to advance its life cycle status.

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

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

5.1.3.2.3 Governance of ~~participant interaction~~Participant Interaction

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

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

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

5.1.3.3 Overarching ~~governance concerns~~Governance Concerns

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

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

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

 Copyright © OASIS Open 2011. All Rights Reserved.

5.1.3.4 Considerations for SOA Governance

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

~~SOA governance must work effectively across ownership boundaries. Thus, there are likely to be multiple governance chains working in parallel. For example, a company making widgets likely has policies intended to ensure they make high quality widgets and make an impressive profit for their shareholders. On the other hand, Sarbanes-Oxley is a parallel governance chain in the United States that specifies how the management must handle its accounting and information that needs to be given to its shareholders. The parallel chains may just be additive or may be in conflict and require some harmonization.~~

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

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

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

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

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

 Copyright © OASIS Open 2011. All Rights Reserved.

5.1.4 Architectural Implications of SOA Governance

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

- Governance is expressed through policies and assumes multiple use of focused policy modules that can be employed across many common circumstances. This requires the existence of:
    - descriptions to enable the policy modules to be visible, where the description includes a unique identifier for the policy and a sufficient, and preferably a machine process-able, representation of the meaning of terms used to describe the policy, its functions, and its effects;
    - one or more discovery mechanisms that enable searching for policies that best meet the search criteria specified by the service participant~~participant;~~, where the discovery mechanism will have access to the individual policy descriptions, possibly through some repository mechanism;
    - accessible storage of policies and policy descriptions, so service participants~~participants~~ can access, examine, and use the policies as defined.
- Governance requires that the participants~~participants~~ understand the intent of governance, the structures created to define and implement governance, and the processes to be followed to make governance operational. This requires the existence of:
    - an information collection site, such as a Web page or portal, where governance information is stored and from which the information is always available for access;
    - a mechanism to inform participants~~participants~~ of significant governance event~~events~~s, such as changes in policies, rules, or regulations;
    - accessible storage of the specifics of Governance Processes;
    - SOA services to access automated implementations of the Governance Processes
- Governance policies are made operational through rules and regulations. This requires the existence of:
    - descriptions to enable the rules and regulations to be visible, where the description includes a unique identifier and a sufficient, and preferably a machine process-able, representation of the meaning of terms used to describe the rules and regulations;
    - one or more discovery mechanisms that enable searching for rules and regulations that may apply to situations corresponding to the search criteria specified by the service participant~~participant;~~, where the discovery mechanism will have access to the individual descriptions of rules and regulations, possibly through some repository mechanism;
    - accessible storage of rules and regulations and their respective descriptions, so service participants~~participants~~ can understand and prepare for compliance, as defined.
    - SOA services to access automated implementations of the Governance Processes.
- Governance implies management to define and enforce rules and regulations. Management is discussed more specifically in section 1.1, but in a parallel to governance, management requires the existence of:
    - an information collection site, such as a Web page or portal, where management information is stored and from which the information is always available for access;
    - a mechanism to inform participants~~participants~~ of significant management event~~events~~s, such as changes in rules or regulations;
    - accessible storage of the specifics of processes followed by management.
- Governance relies on metrics to define and measure compliance. This requires the existence of:
    - the infrastructure monitoring and reporting information on SOA resources;
    - possible interface requirements to make accessible metrics information generated or most easily accessed by the service itself.

5.2 Security Model

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

 Copyright © OASIS Open 2011. All Rights Reserved.

Security

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

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

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

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

~~However, the fact that we aim to explicitly relate the IT architecture with the human architecture (see) makes it possible to give a more complete accounting of security. In effect, an analysis of the social structures in place around a SOA-based system forms a backdrop and context for security.~~

~~Concepts such as constitutions, roles, and authority within social structures play an important part in the establishment of ownership and trust boundaries within and between social structures.~~

~~In addition, security often revolves around *resources*: the need to guard certain resources against inappropriate access – whether reading, writing or otherwise manipulating those resources. The basic resource model that informs our discussion is outlined in Section .~~

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

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

5.2.1 ~~Security~~Secure Interaction Concepts

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

5.2.1.1 Confidentiality

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

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

5.2.1.2 Integrity

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

 Copyright © OASIS Open 2011. All Rights Reserved.

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

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

5.2.1.3 Authentication

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

Figure 39 - Authentication

applies authentication to the identity of participants.

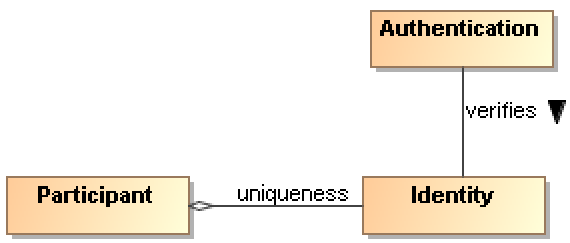

Figure 39 - Authentication

5.2.1.4 Authentication

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

Copyright © OASIS Open 2011. All Rights Reserved.

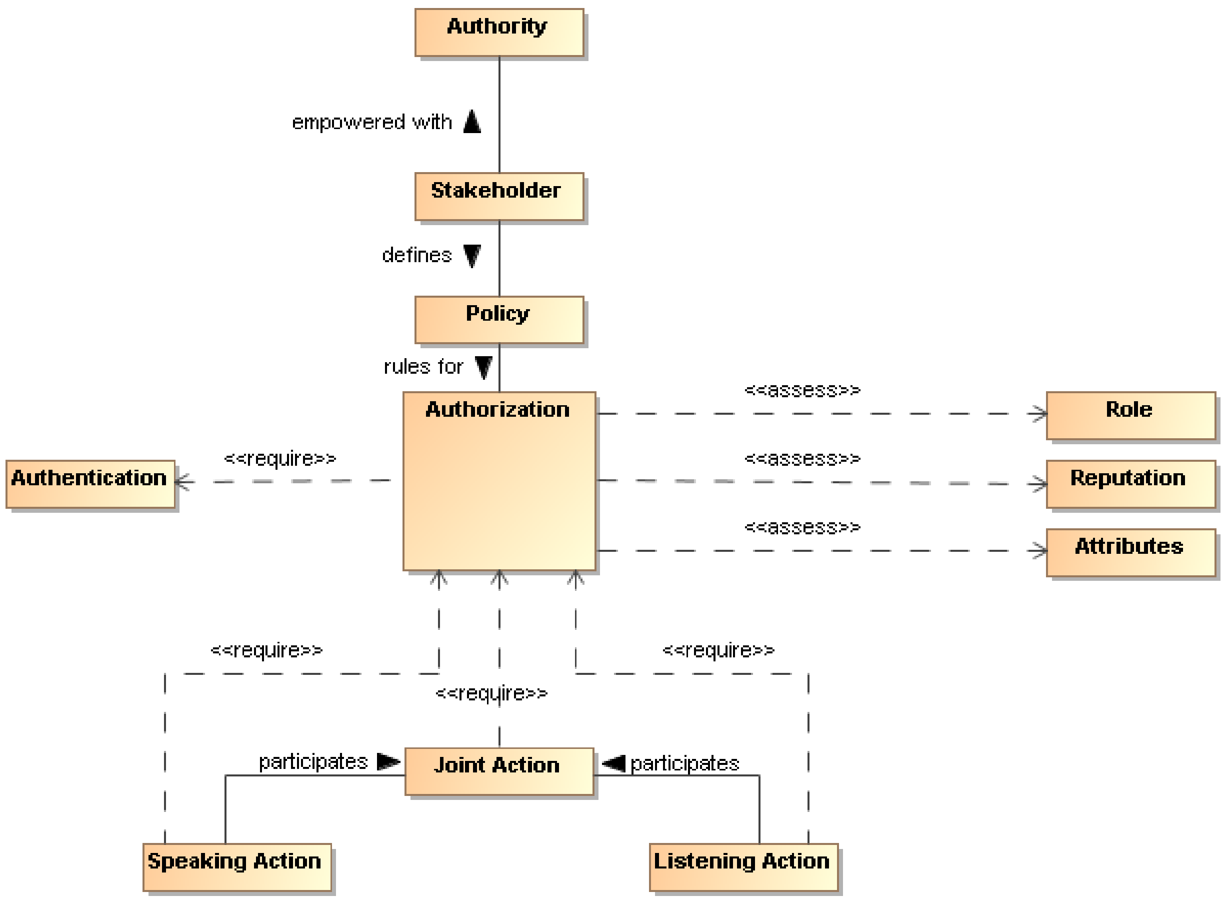

Figure 40 - Authorization

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

The role of policy for security is to permit stakeholders to express their choices. In Figure 40, a policy is a written constraint and the role, reputation, and attribute assertions are evaluated according to the constraints in the authorization policy. A combination of security mechanisms and their control via explicit policies can form the basis of an authorization solution.

5.2.1.5 Non-repudiation

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

~~5.2.1.3~~5.2.1.6 Availability

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

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

 Copyright © OASIS Open 2011. All Rights Reserved.

~~**Authentication**~~

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

~~**Authorization**~~

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

~~5.2.1.41.1.1.1 Non-repudiation~~

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

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

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

5.2.2 Where SOA Security is Different

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

~~SOA is primarily about action and events. This model focuses on the issues around these concepts more than simple data exchange.~~

SOA policy-based security can be more adaptive for a computing ecosystem than previous computing technologies allow for, and typically involves a greater degree of distributed mechanisms. ~~Section provides one example of distributed policy-based computing mechanisms that may be present as part of the realization of SOA security. Distributed security mechanisms allow for centralized identity and policy services as well as centralized or decentralized authentication and authorization services.~~

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

~~5.2.3 Trust Model~~

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

~~models trust in terms of a participant, the participant's identity and credentials, and the participant's authorization to perform an action.~~

 Copyright © OASIS Open 2011. All Rights Reserved.

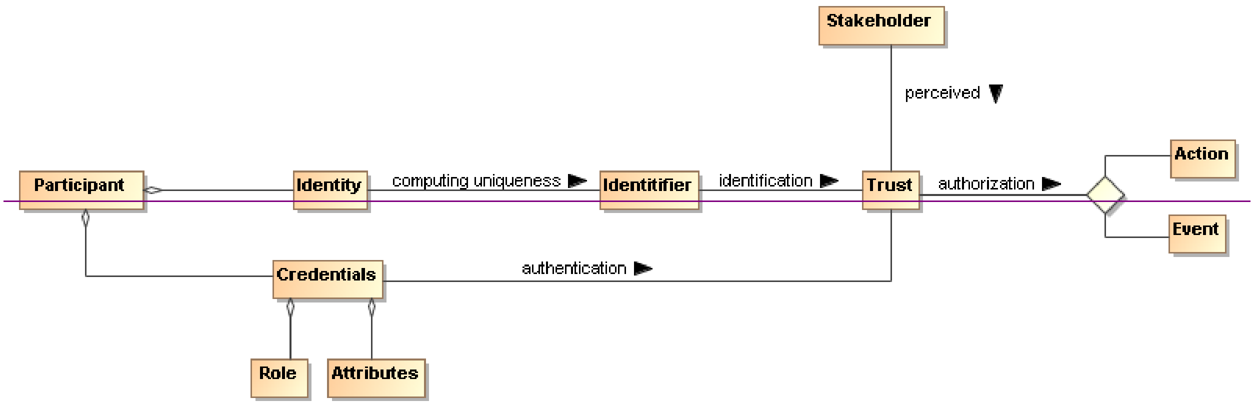

~~*Figure Trust Model*~~

~~**Trust**~~

> ~~Trust is the relationship, as perceived by a stakeholder, between a participant and a set of actions and events, which concerns the legitimacy of the agent's actions and reported events.~~

~~**Credentials**~~

> ~~The role and/or set of attributes a stakeholder uses to determine authorization to actions.~~

~~Trust is not easily modeled as a single number or other scalar value. The motivation for this definition of trust is to allow us to distinguish the purpose of the trust as well as the degree of trust. For example, one may trust a stranger to hold a space in a queue for the Cinema, but one would typically not trust that same person to hold one's car keys for a fortnight's vacation.~~

~~5.2.3.1 Trust Domain~~

~~The Trust Domain in models abstract concepts behind the formation of policy-based trusted social groups.~~

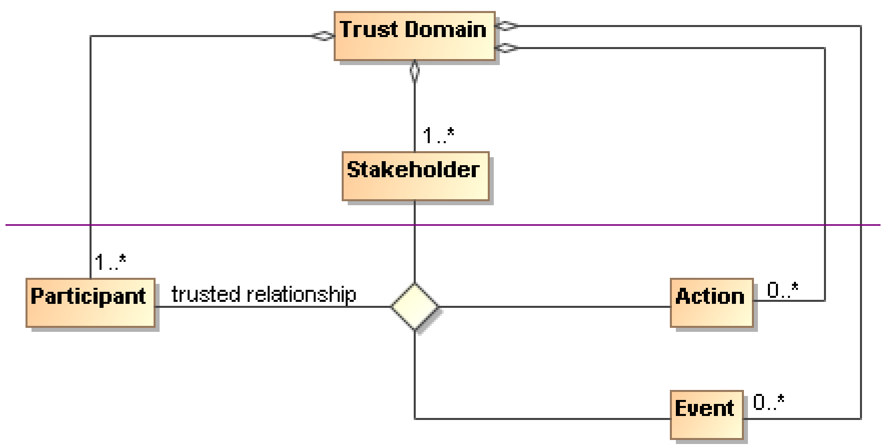

Figure ~~Trust Domain~~

~~**Trust Domain**~~

> ~~An abstract space of actions which all share a common trust requirement; i.e., all participants that perform any of the actions must be in the same trust relationship.~~

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

 Copyright © OASIS Open 2011. All Rights Reserved.

~~5.2.3.2 Centralized and Decentralized Trust Authority~~

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

~~Trust domains will require a centralized and/or decentralized authentication and authorization authority to form trust relationships. An example of a centralized authority might be a governing body that requires regulatory compliance for all participants performing a specific action. A decentralized trust authority gives individual participant's more authority to authenticate and authorize actions and events.~~

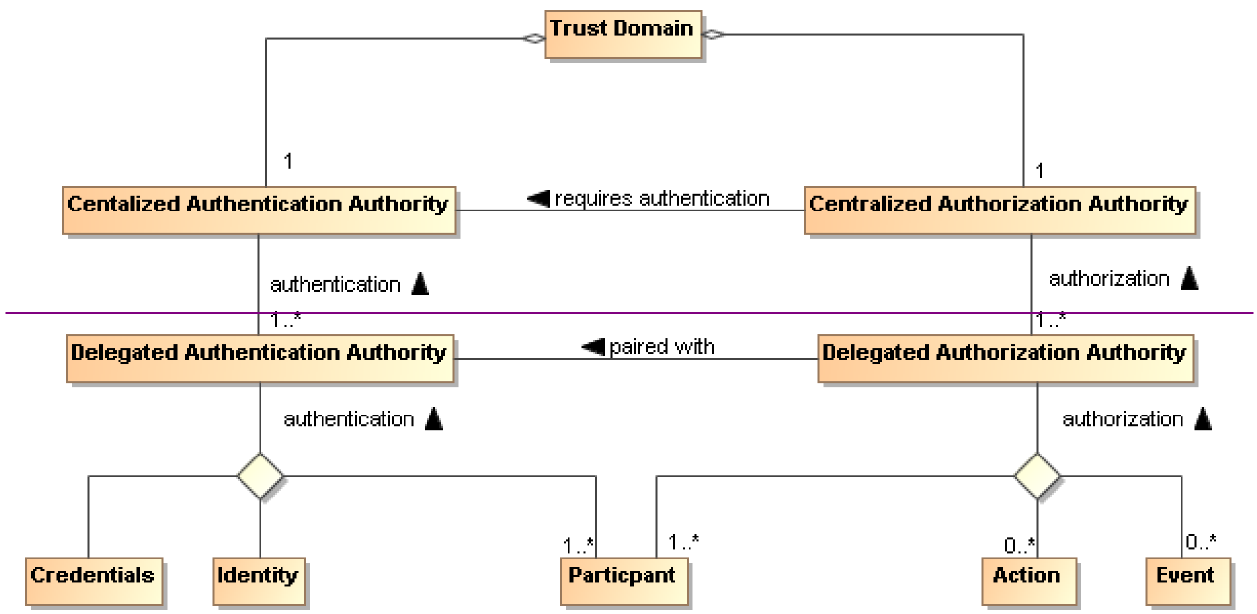

~~Figure Centralized Trust Authority~~

~~depicts a hierarchical central trust authority. A participant's credentials and identity are authenticated by a centralized authentication authority. A web browser will often use a centralized authority in establishing secure communications with a service provider such as a bank. Actions and events also have centralized authorization authorities in this model. Centralized trust authorities tend to provide stronger regulatory control and more efficient revocation of participants.~~

~~In the context of a SOA that is used by many people, there may not be a single repository for information that can justify trust. Often different aspects of trust are managed by different entities. For example, a corporate directory might be used to verify the employment of an individual, whereas a bank would be used to verify their credit worthiness and a government agency used to verify their residency.   depicts chains of trust between participants that are established by participants who introduce other participants into the chain of trust.~~

Copyright © OASIS Open 2011. All Rights Reserved.

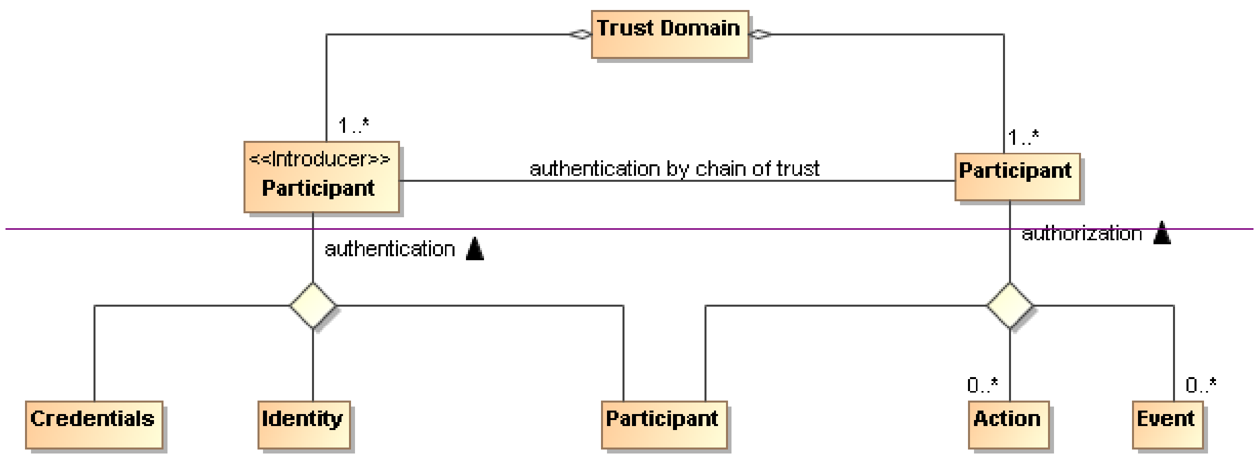

~~*Figure Decentralized Trust Authority*~~

~~Together, the various entities that provide corroboration of an individual's authenticity and trustworthiness to perform actions and raise events form a chain of trust. Chain's of trust need not be functionally organized: third parties who are known to both may also be used to facilitate trust. A long chain of trust is likely to be more fragile and less trustworthy than a simple one.~~

~~Complex trust domains are likely to be composed of a combination of centralized and decentralized trust authorities. For SOA, the level of complexity of a trust domain can achieve is dependent on the policy language's and IT mechanism's ability to express trust relationships.~~

~~5.2.4~~5.2.3 ~~Policy Mechanisms for~~ Security Threats

~~When a participant wishes to perform an action that requires access to a trust domain, depending on the policies that are in place, he/she must provide suitable identification and/or credentials before continuing the interaction.~~

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

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

~~depicts security interactions based on Section . In the context of security, the diagram has been modified with recognized policy, identity, and attribute authorities in the SOA ecosystem. Additional auditing has also been depicted.~~

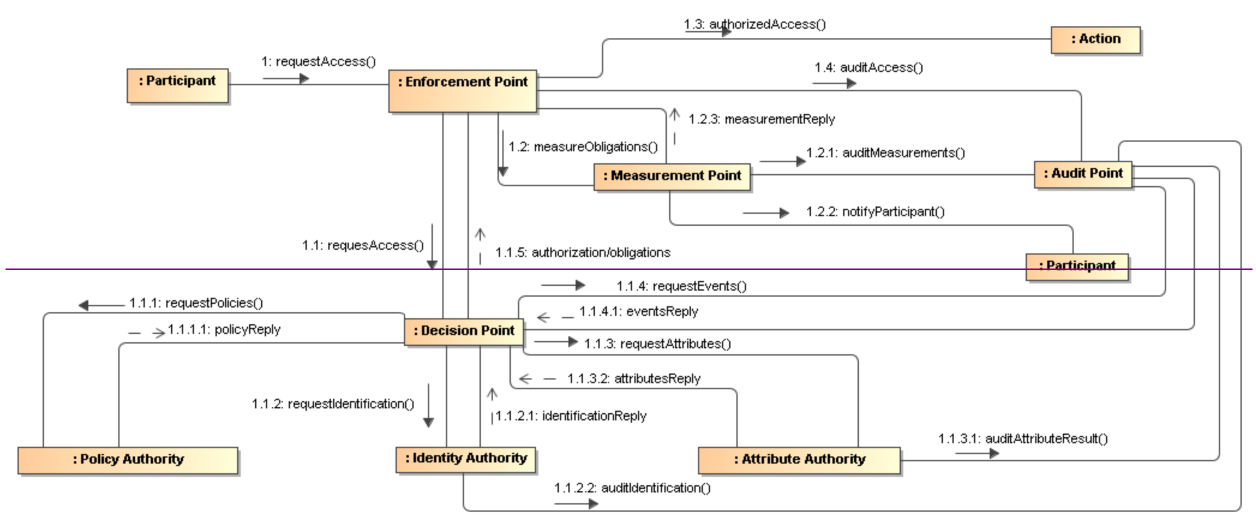

~~*Figure Policy Based Security*~~

Copyright © OASIS Open 2011. All Rights Reserved.

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

~~5.2.5 Security Layers~~

~~Security concepts can be described in terms of three primary layers when discussing the deployment of SOA-based systems. The commonly known OSI seven-layer model provides an expanded view of these three primary layers, each one of the OSI seven layers requires specific application of security. However, discussing the seven layers of the OSI seven-layer model is beyond the scope of this reference architecture.~~

~~depicts three generalized layers of security to consider and their relationship to ownership domains when deploying SOA-based systems. The lowest level of abstraction is the network layer, the next level of abstraction is the transport layer, and the third level of abstraction is the application layer.~~

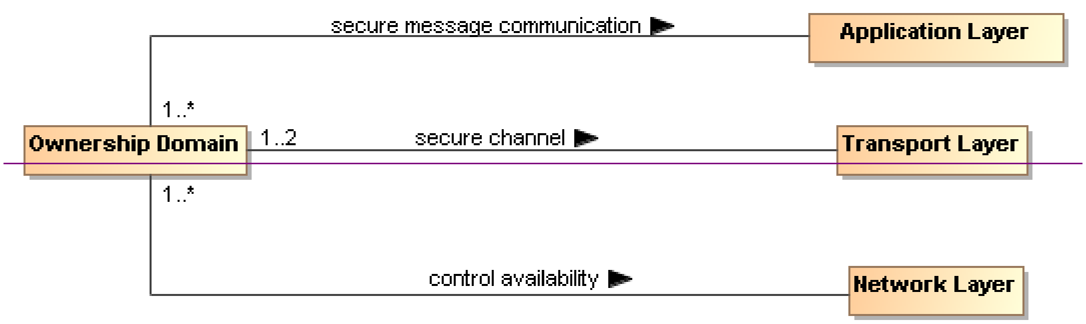

~~*Figure Security Layers*~~

~~5.2.5.1 Network Layer~~

~~At the lowest level of abstraction in the security model are the network devices and the hardware that links the network devices, referred to as the network layer. The network layer includes devices like routers and firewall appliances and it also includes protocols such as the Internet Protocol (IP), Border Gateway Protocol (BGP), Open Shortest Path First (OSPF) protocol, etc. Network devices, however, can have policy-based SOA security mechanisms built in so there is not always a clear distinction between network device and network layer.~~

~~In order for a SOA-based system to operate, the network must be available to provide network services. Control of the network layer is required in order to address the security concept of availability such as protection from Denial of Service (DoS) attacks.~~

~~The network layer may also address general availability by defining policies or service level agreements (SLAs) about the quality of service of the network layer operation and then translating hose commitments into measurable constraints carried out by the network devices for such things as guaranteed service delivery or specific bandwidth allocations.~~

~~5.2.5.2 Transport Layer~~

~~The transport layer may pass through network layers belonging to many ownership domains. The transport layer is primarily concerned with establishing a secure communications channel between sender and receiver, a good example being the interaction with a bank through a web browser. The transport layer may include protocols like HTTP over Transport Layer Security (TLS) as well as HTTP over Secure Sockets Layer (SSL).~~

~~Given the nature of SOA-based communications across multiple ownership boundaries, security provided at the transport layer cannot be relied upon for protection of message confidentiality.~~

Copyright © OASIS Open 2011. All Rights Reserved.

~~5.2.5.3 Application Layer~~

~~The application layer accounts for the security of messaging between participants within a SOA ecosystem, where participants may have policy based roles and authority to act within and across ownership domains. Web service standards like WS-Security, XML Digital Signature, XML Encryption, and SAML are all examples of standards addressing the security concepts at the application layer.~~

~~Application layer security for SOAs may be built into network devices so network devices may have network layer and application layer security built in.~~

~~In a SOA ecosystem where participants interact through many ownership domains and any number of unknown network domains, the application layer may be the only layer the basic security principles of confidentiality, integrity, authentication, authorization, and non-repudiation are assured. Assurance of availability is addressed at the network layer but may be controlled by the application layer and/or transport layer.~~

~~5.2.6 Threat Model~~

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

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

Message alteration

> If an attacker is able to modify the content (or even the order) of messages that are exchanged without the legitimate participants~~participants~~ being aware of it then the attacker has successfully compromised the security of the system. In effect, the participants~~participants~~ may unwittingly serve the needs of the attacker rather than their own.
>
> An attacker may not need to completely replace a message with his own to achieve his objective: replacing the identity of the beneficiary of a transaction may be enough.

Message interception

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

Man in the middle

> In a man-in-the-middle attack, the legitimate participants~~participants~~ believe that they are interacting with each other; but are in fact interacting with the attacker. The attacker attempts to convince each participant~~participant~~ that he is their correspondent; whereas in fact he is not.
>
> In a successful man-in-the-middle attack, legitimate participants~~participants will often~~ do not have ~~a true~~anaccurate understanding of the state of the other participants~~participants.~~. The attacker can use this to subvert the intentions of the participants~~participants.~~.

Spoofing

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

Denial of service attack

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

 Copyright © OASIS Open 2011. All Rights Reserved.

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

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

Replay attack

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

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

False ~~Repudiation~~repudiation

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

~~5.2.7~~5.2.4 Security ~~Response Model~~Responses

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

Performing threat assessments, devising mitigation strategies, and determining acceptable levels of risk are the foundation for an effective process to mitigating threats in a cost-effective way.[19] The choice in hardware and software to realize a SOA will be ~~the~~a basis for threat assessments and mitigation strategies. The stakeholder~~stakeholders~~s of a specific SOA implementation should determine acceptable levels of risk based on threat assessments and the cost of mitigating those threats. ~~Example mitigation strategies are provided for threats listed in Section .~~

~~5.2.7.1~~5.2.4.1 Privacy Enforcement

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

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

 Copyright © OASIS Open 2011. All Rights Reserved.

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

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

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

~~5.2.7.2~~5.2.4.2 Integrity Protection

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

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

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

~~5.2.7.3~~5.2.4.3 Message Replay Protection

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

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

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

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

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

~~5.2.7.4~~5.2.4.4 Auditing and Logging

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

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

 Copyright © OASIS Open 2011. All Rights Reserved.

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

~~5.2.7.5~~5.2.4.5 Graduated engagement

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

~~5.2.8~~5.2.5 Architectural Implications of SOA Security

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

- Security expressed through policies have the same architectural implications as described in Section 4.4.3 for policies and contracts~~contracts~~ architectural implications.
- Security policies require mechanisms to support security description administration, storage, and distribution.
- ~~Security policies should:~~
    - ~~be able to express trust relationships and trust domains;~~
    - ~~provide the ability to update policy trust relationships and trust domains in a way that does not require upgrades to software and hardware;~~
    - ~~be able to express standard protocols used to provide confidentiality, integrity, authentication, authorization, non-repudiation, and availability.~~
- Service descriptions supporting security policies should:
    - have a meta-structure sufficiently rich to support security policies;
    - be able to reference one or more security policy artifacts;
    - have a framework for resolving conflicts between security policies.
- The mechanisms that make-up the execution context in secure SOA-based ~~message exchanges~~systems should:
    - provide protection of the confidentiality and integrity of message exchanges;
    - be distributed so as to provide centralized or decentralized policy-based identification, authentication, and authorization;
    - ensure service availability to consumers;
    - be able to scale to support security for a growing ecosystem of services;
    - be able to support security between different communication technologies;
- Common security services include:
    - services that abstract encryption techniques;
    - services for auditing and logging interactions and security violations;
    - services for identification;
    - services for authentication;
    - services for authorization;
    - services for intrusion detection and prevention;
    - services for availability including support for quality of service specifications and metrics.

 Copyright © OASIS Open 2011. All Rights Reserved.

~~5.3 Services as Managed Entities Model~~

~~5.4~~5.3 Management Model

5.3.1 Management

Management is ~~the control~~a process of ~~the use, configuration, and availability of~~controlling resources in accordance with the policies ~~of the stakeholders involved~~and principles defined by Governance.

There are three separate but linked domains of interest within the management of SOA~~-based systems. The first and most obvious is~~ ecosystem:

1. the management and support of the resources that are involved in any complex ~~system~~ structures – of which SOA~~-based systems~~ ecosystems are excellent examples~~. The second is~~;
2. the promulgation and enforcement of the policies and service contracts agreed to by the stakeholders in the SOA~~-based systems. The third domain is~~ ecosystem;
3. the management of the relationships of the participants ~~in SOA-based systems~~ – both to each other and to the services that they use and offer.

There are many artifacts ~~in a large system that may need~~ related to management. ~~As soon as there is the possibility of more than one instance of a thing, the issue of managing those things becomes relevant.~~ Historically, systems management capabilities have been organized by the ~~following functional groups known as~~ "FCAPS" functions (based on ITU-T Rec. M.3400 (02/2000), "TMN Management Functions"): ~~Fault~~

- fault management,
- configuration management,
- account management,
- performance and security management.

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

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

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

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

Management in a SOA ecosystem is thus concerned with management taking actions that will establish the condition of trust that must be present before engaging in service interactions. These concerns should largely be handled within the governance of the ecosystem. The policies, agreements, and practices defined through governance provide the boundaries within which management operates and for which management must provide enforcement and feedback. However, governance alone cannot foresee all circumstances but must offer sufficient guidance where agreement between all stakeholders cannot be

 Copyright © OASIS Open 2011. All Rights Reserved.

reached. Management in these cases must be flexible and adaptable to handle unanticipated conditions without unnecessarily breaking trust relationships.

Service management is the process – manual, automated, or a combination – of proactively monitoring and controlling the behavior of a service or a set of services. Service management operates under constraints attributed to the business and social context. Specific policies may be used to govern cross-boundary relationships. Managing ~~systems~~solutions based on such policies (and that may be used across ownership boundaries) raises issues that are not ~~normally~~ typically present when managing a ~~system~~service within a single ownership domain. ~~For example, care~~Care is therefore required in managing a service when the owner of the service, the provider of the service, the host of the service and ~~access~~ mediators to the service may all belong to different stakeholders. ~~In addition, it may be important to allow service consumers to communicate their requirements to the service provider so that they are satisfied in a timely manner.~~

~~A given~~Cross-boundary service ~~may be provided and consumed~~management takes place in ~~more than one version. Version control~~, at least, the following situations:

- using combinations of services ~~is~~that belong to different ownership domains
- using of services that mediate between ownership domains
- sharing monitoring and reporting means and results.

These situations are particularly important ~~both for service providers and service consumers (who may need to ensure certainty~~ in ~~the version of~~ ecosystems that are highly decentralized, in which the ~~service they are interacting with).~~participants interact as peers as well as in the “master-servant” mode.

The management model shown in Figure 41~~In fact, managing a service has quite a few similarities to using a service: suggesting that we can use the service oriented model to manage SOA-based systems as well as provide them. A management service would be distinguished from a non-management service more by the nature of the capabilities involved (i.e., capabilities that relate to managing services) than by any intrinsic difference.~~

~~In this model, we show~~ conveys how the SOA ~~framework may apply~~applies to managing services ~~as well as using and offering them. There are, of course, some special considerations that apply to service~~. Services management ~~which we bring out: namely that we will be managing the life-cycle of services, managing any service level~~ operates via service metadata, such as service lifecycles and attributes~~,~~ ~~managing dependencies between services and so on.~~ associated with service use, which are typically collected in or accessed through the service description.

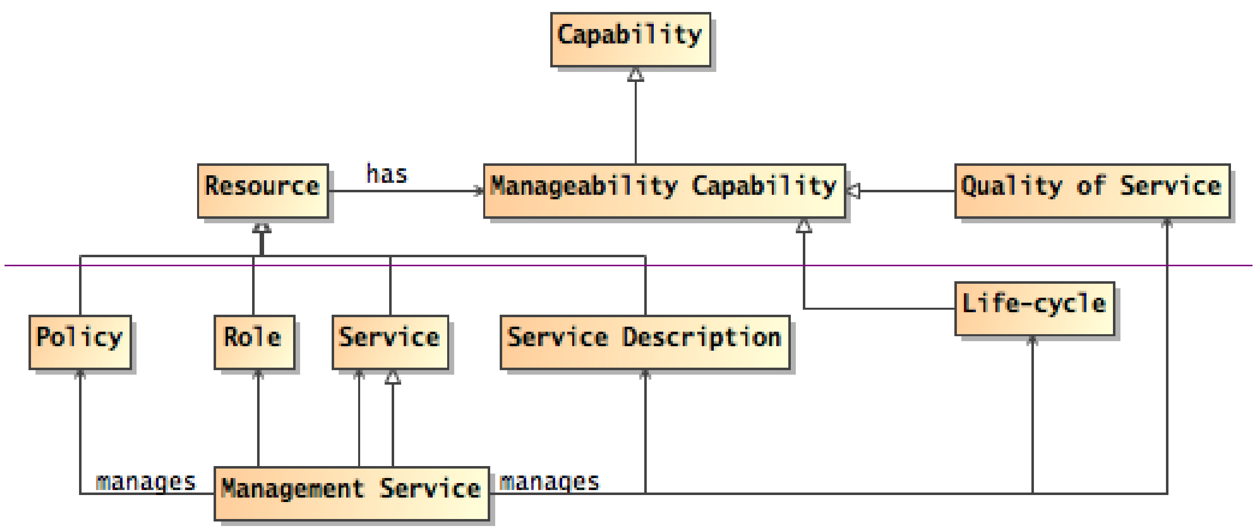

Copyright © OASIS Open 2011. All Rights Reserved.

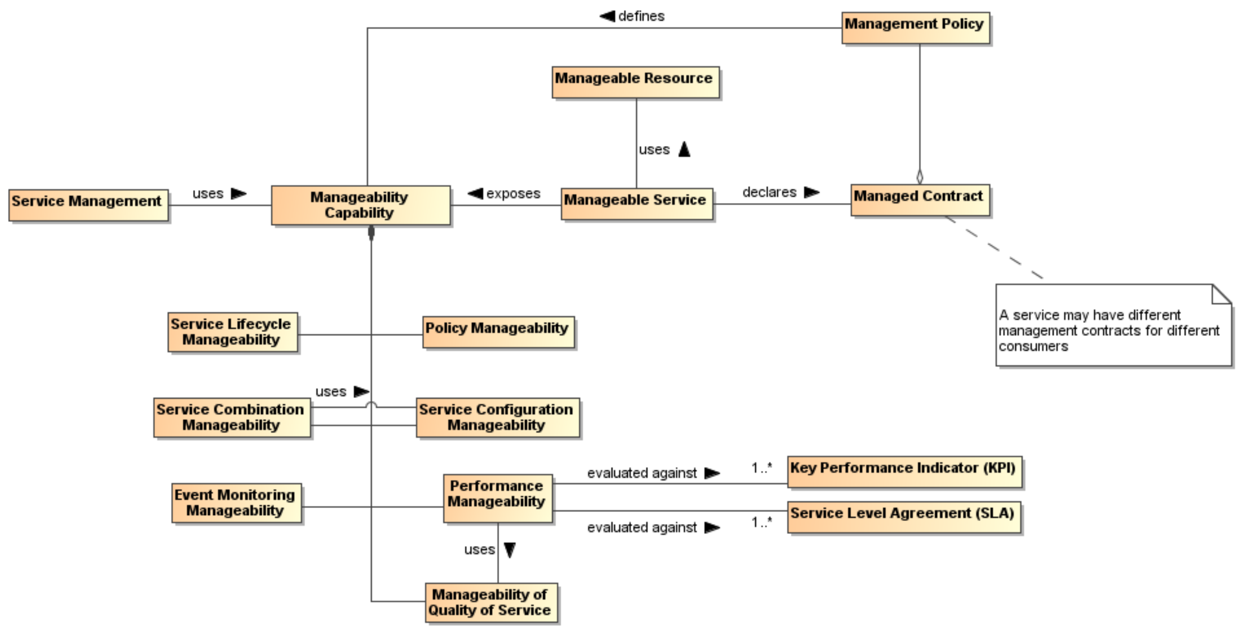

Figure 41 ~~Managing resources in a SOA~~

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

\- *Manageability capabilities in SOA ecosystem ~~Capability~~*

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

Manageability of a resource is the capability that allows it to be ~~managed~~controlled, monitored, and reported on with respect to some property. Note that manageability ~~capabilities are~~ is not necessarily a part of the managed entities themselves and are generally considered to be external to the managed entities.

> ~~Manageability capabilities are the core resources that management systems use to manage: each~~Each resource ~~that~~ may be managed ~~in some way has~~through a number of aspects ~~that~~ of management, and the resources may be ~~managed.~~grouped based on similar aspects. For example, ~~a service's life-cycle~~ resources may be ~~manageable, as may its Quality of Service parameter; a policy may also~~ grouped according to the aspect referred to as "Configuration Manageability" for the collection of services. Some resources may not be managed ~~for life-cycle but Quality of Service would not normally apply.~~

~~Life-cycle manageability~~

~~A manageability~~under a particular capability ~~associated~~if there are no manageability aspects with a ~~resource that permits the life cycle of the resource to be managed.~~clear meaning or use. As ~~noted above, the life-cycle manageability capability of a resource is unlikely to reside~~an example, all resources within ~~the resource itself (you cannot tell a system that~~a SOA ecosystem have a lifecycle that is meaningful within the ecosystem. Thus, all resources are manageable under Lifecycle Manageability. In contrast, not all resources report or handle events. Thus, Event Manageability is ~~not running to start itself).~~only concerned with those resources for which events are meaningful.

> ~~The life~~**Life-cycle** ~~management~~**Manageability** of a ~~resource~~service typically refers to how the ~~resource~~service is created, how it is destroyed and ~~what dependencies there might exist that~~ how service versions must be ~~simultaneously~~ managed.

~~Configuration~~ This manageability

Copyright © OASIS Open 2011. All Rights Reserved.

~~A capability that permits~~ is a feature of the SOA ecosystem because the service cannot manage its own life cycle.

Another important consideration is that services may have resource requirements that must be established at various points in the services' life cycles. ~~configuration of~~ However actual providers of these resources may not be known at the time of the service creation and, thus, have to be managed~~.~~ at service run-time.

Combination Manageability of a service addresses management of service characteristics that allow for creating and changing combinations in which the service participates or that the service combines itself. Known models of such combinations are aggregations and compositions. Examples of patterns of combinations are choreography and orchestration. Combination Manageability drives implementation of the Service ~~configuration, in particular, may be complex~~Composability Principle of service orientation.

Service combination manageability resonates with the methodology of process management. Combination Manageability may be applied at different phases of service creation and execution and, in some cases~~ where there are dependencies between~~, can utilize Configuration Manageability.

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

Configuration Manageability of a service allows managing the identity of and the interactions among internal elements of the service. Also, Configuration Manageability correlates with the management of service versions and configuration of the deployment of new services into the ecosystem. Configuration Management differs from the Combination Manageability in the scope and scale of manageability, and addresses lower level concerns than the architectural combination of services.

Event Monitoring Manageability allows managing the categories of events of interest related to services and reporting recognized events to the interested stakeholders. Such events may be the ones that trigger service invocations as well as execution of particular functionality provided by the service.

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

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

The performance business- and technical-level characteristics of the service should be known from the service contract. The service provider and consumer must be able to monitor and measure these characteristics or be informed about the results measured by a third party.

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

> **Manageability of Quality of Service** deals with management of service non-functional characteristics that may be of significant value to the service consumers and other ~~resources.~~

~~**Event monitoring manageability**~~

> ~~Managing the reporting of events and faults is one of the key lower-level manageability capabilities.~~

~~**Accounting manageability**~~

> ~~A capability associated with resources that allows for the use of those resources to be measured and accounted for. This implies that not only can the *use* of resources be properly measured, but also that those *using* those resources also be properly identified.~~

> ~~Accounting for the use of resources by participants in the SOA supports the proper budgeting and allocation of funding by participants.~~

 Copyright © OASIS Open 2011. All Rights Reserved.

~~**Quality of service manageability**~~

~~A manageability capability associated with a resource that permits any quality of service associated with the resource to be managed.~~stakeholders in the SOA ecosystem. Classic examples of this include bandwidth ~~requirements and~~ offerings associated with a service.

~~**Business performance manageability**~~

> ~~A manageability capability~~Manageability of quality of service assumes that ~~is~~ the properties associated with ~~services that permits the service's business performance to be~~ service qualities are monitored ~~and managed. In particular, if there are business-level service level agreements that apply to a service, being able to monitor and manage those SLAs is an important role for management systems.~~

~~Building support for arbitrary business~~during the service execution. Results of monitoring ~~is likely to be challenging. However, given a *measure* for determining a service's compliance to business service level agreements, management systems can monitor that performance~~may be challenged against SLA and even KPI, which results in ~~a way that is entirely similar to other management tasks~~the continuous validation of how the service contract is preserved by the service provider.

Policy ~~**manageability**~~

~~Where~~**Manageability** allows additions, changes and replacements of the policies associated with a resource ~~may be complex and dynamic, so those policies themselves may require management.~~in the SOA ecosystem. The ability to manage those policies (such as promulgating policies, retiring policies and ensuring that policy decision points and enforcement points are current) ~~is a management function.~~enables the ecosystem to apply policies and *evaluate* the results.

~~In the~~ The capability to manage, i.e. use a particular ~~case of~~ manageability, requires policies~~, there is a special relationship~~ from governance to be translated into detailed rules and regulations which are measured and monitored providing corresponding feedback for enforcement.

Management of SOA ecosystem recognises the manageability challenge and requires manageability properties to be considered for all aforementioned manageability cases. In the following sub-sections, we describe how thses properties are used in the management. Also we describe some relationships between management and ~~policies. Just like~~ other ~~artifacts, policies require management in a SOA. However, much~~components of SOA ecosystem.

5.3.2 Management Means and Relationships

A minimal set of management issues for the SOA ecosystem is shown on Figure 42~~is about *applying* policies also: where governance is often about what the policies regarding artifacts and services should be, a key management role is to ensure that those policies are consistently applied~~ and elaborated in the following sections.

 Copyright © OASIS Open 2011. All Rights Reserved.

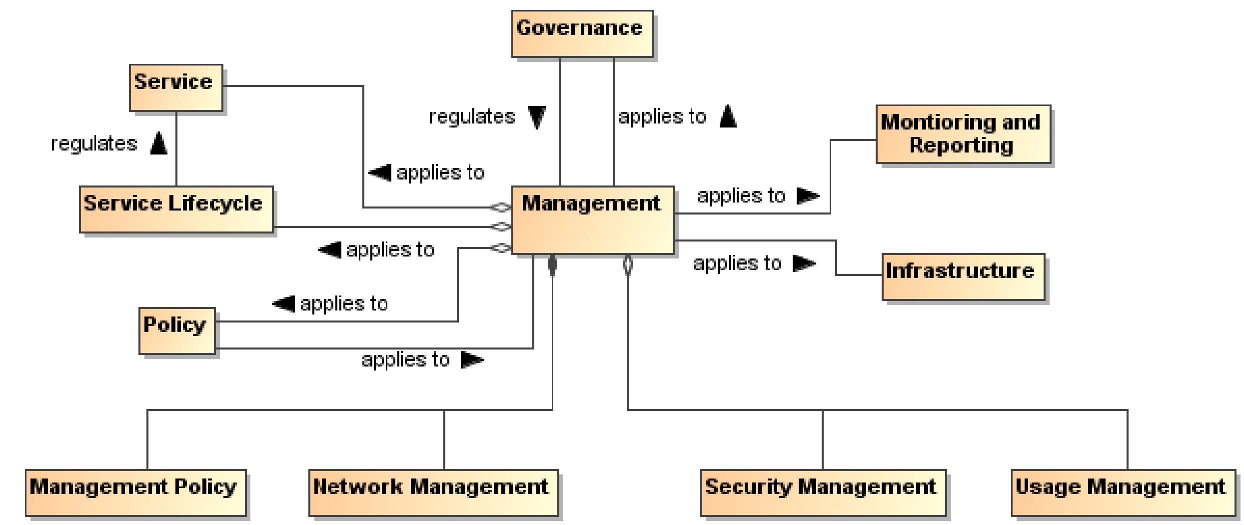

Figure 42 - Management ~~service~~

~~A management service is a service that manages other services~~Means and ~~resources.~~Relationships in SOA ecosystem

~~5.4.1.1~~5.3.2.1 Management Policy

~~A management policy is a policy whose topic is a management topic. Just as with other aspects of a SOA, the management~~The management of resources within the SOA ecosystem may be governed by management policies~~, contracts (such as SLAs).~~

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

~~There are additional~~Additional concepts of management ~~that often~~ also apply to IT management~~:~~.

~~**Systems management**~~

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

~~5.4.1.2~~5.3.2.2 Network ~~management~~ Management

Network management ~~refers to~~deals with the maintenance and administration of large ~~-~~scale physical networks such as computer networks and telecommunication networks. ~~Systems~~Specifics of the networks may affect service interactions from performance and operational perspectives.

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

5.3.2.3 ~~However, for the purposes~~Security Management

Security Management includes identification of ~~this Reference Architecture, while recognizing their importance, we do not focus on systems~~roles, permissions, access rights, and policy attributes defining security boundaries and events that may trigger a security response.

Security management ~~or network~~within a SOA ecosystem is essential to maintaining the trust relationships between participants residing in different ownership domains. Security management must

 Copyright © OASIS Open 2011. All Rights Reserved.

consider not just the internal properties related to interactions between participants but ecosystem properties that preserve the integrity of the ecosystem from external threats.

~~- the specific identifier is not prescribed by this Reference Architecture but the structure and semantics of the identifier must be indicated for the identifier value to be properly used. For example, part of identity may include version identification.~~

~~For this, the configuration management plan or similar document from which the version number is derived must be identified.~~

5.3.2.4 Usage Management

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

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

~~5.4.2~~5.3.3 Management and Governance

The primary role of governance in the context of ~~SOA is to allow~~a SOA ecosystem is to foster an atmosphere of predictability, trust, and efficiency, and it accomplishes this by allowing the stakeholders ~~in the SOA to be able~~ to negotiate and set the key policies that govern the running of the ~~system.~~SOA-based solution. Recall that in an ecosystems perspective, the goal of governance is less to have complete fine-grained control but more to enable the individual participants to work together. ~~Policies that are set at the governance of a SOA will tend to focus on the rules of engagement between participants – what kind of interacts are permissible, how to resolve disputes, and so on.~~

~~While governance may be primarily focused on setting policies, management is more focused on realization and enforcement of policies.~~

~~5.4.3 Management Contracts and Policies~~

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

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

 Copyright © OASIS Open 2011. All Rights Reserved.

5.3.4 Management and Contracts

5.3.4.1 Management for Contracts and Policies

As we noted above, management can often be viewed as the application of contracts and individual policies to ensure the smooth running of the SOA~~. Policies~~ ecosystem. Policies and service contracts specify the service characteristics that have to be monitored, analysed and managed and play an important ~~part in managing systems both as~~role as the guiding constraints for management, as well as being artifacts ~~that~~ (e.g., policy and contractual documents) that also need to be managed.

5.3.4.2 Contracts

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

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

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

Service Contract

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

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

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

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

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

 Copyright © OASIS Open 2011. All Rights Reserved.

prorated cost.

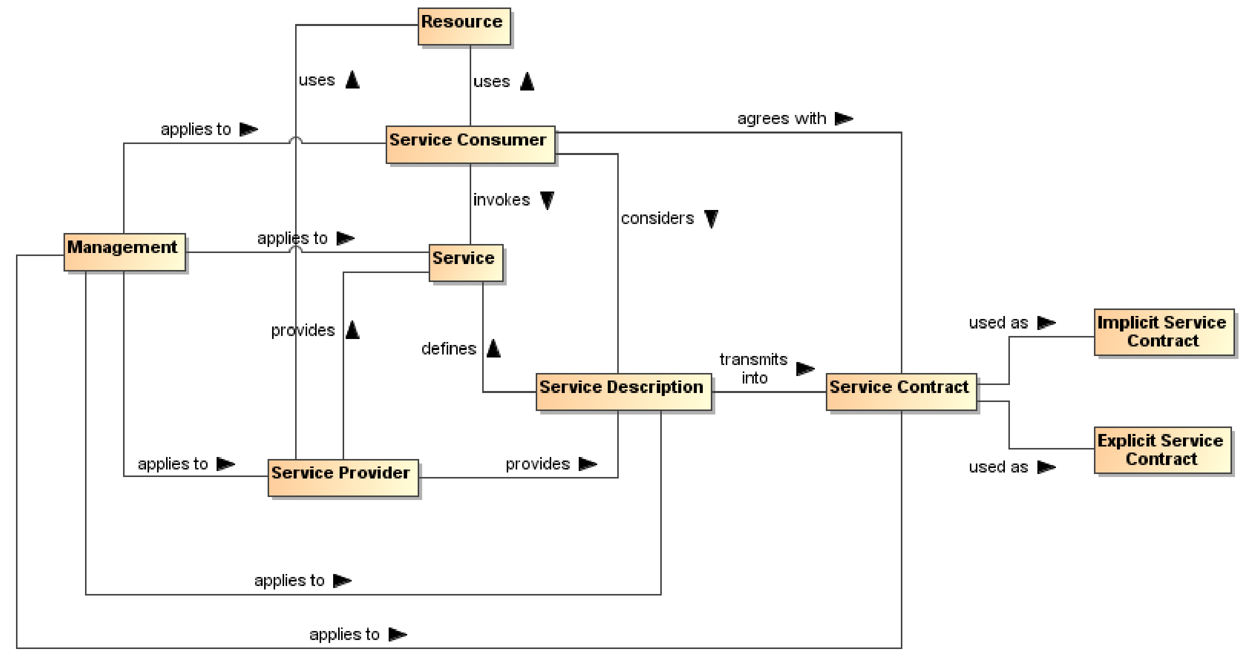

Figure 43 ~~and as the guiding constraints to determine how the SOA should be managed~~ - Management of the service interaction

Consequently, an implicit service contract is an agreement (1) on the consumer side with the terms, conditions, features and interaction means specified in the service description "as is" or (2) a selection from alternatives that are made available through mechanisms included in the service description, and neither of these require any a priori interactions between the service consumer and the service provider. An explicit service contract always requires a form of interaction between the service consumer and the service provider prior to the service invocation. This interaction may regard the choice or selection of the subset of the elements of the service description or other alternatives introduced through the formal agreement process that would be applicable to the interaction with the service and affect related joint action.

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

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

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

 Copyright © OASIS Open 2011. All Rights Reserved.

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

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

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

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

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

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

~~5.4.3.1~~5.3.4.3 Policies

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

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

~~5.4.3.2~~5.3.4.4 Service Description and Management ~~Infrastructure~~

~~In order for a service or other resource to be manageable there must be a corresponding manageability capability that can effect that management. The particulars of this capability will vary somewhat depending on the nature of the capability. For example, a service life-cycle manageability capability requires the ability to start a service, to stop the service, and potentially to pause the service. Conversely, in order to manage document-like artifacts, such as service descriptions, the capability of storing the artifacts, controlling access to those artifacts, allowing updates of the artifacts to be deployed are all important capabilities for managing them.~~

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

- ~~Integrate with existing security services~~
- ~~Monitoring~~
- ~~Heartbeat and Ping~~
- ~~Alerting~~
- ~~Pause/Restore/Restart Service Access~~
- ~~Logging, Auditing, Non-Repudiation~~

 Copyright © OASIS Open 2011. All Rights Reserved.

- ~~Runtime Version Management~~
- ~~Complement other infrastructure services (discovery, messaging, mediation)~~

~~* Message Routing and Redirection~~
~~* Failover~~
~~* Load-balancing~~

~~* QoS, Management of Service Level Objects and Agreements~~
~~* Availability~~
~~* Response Time~~
~~* Throughput~~

- ~~Fault and Exception Management~~

~~5.4.4 Service Life-cycle~~

~~Managing a service's life cycle involves managing the establishment of the service, managing its steady-state performance, and managing its termination. The most obvious feature of this is that a service cannot manage its own life cycle (imagine asking a non-functioning service to start).~~ ~~Another important consideration is that services may have resource requirements that must be established at various points in the services' life cycles.~~ The service description identifies several management objects such as a set of service interfaces and related set of SLAs. Service behavioral characteristics and performances specified in the SLA depend on the interface type and its Execution Context. In the service description, a service consumer can find references to management policies, SLA metrics, and the means of accessing measured values that together increase assurance in the service quality. At the same time, service description is an artifact that needs to be managed.

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

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

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

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

 Copyright © OASIS Open 2011. All Rights Reserved.

5.3.5 Management for Monitoring and Reporting

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

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

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

5.3.6 Management for Infrastructure

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

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

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

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

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

The combination of manageability capabilities and infrastructure elements constitutes a certain level of SOA management maturity. While several maturity models exist, this topic is out of the scope of the current document.

5.3.7 Architectural Implication of the Management Model

5.4 SOA Testing Model

Testing for SOA combines the typical challenges of software testing and certification with the additional needs of accommodating the distributed nature of the resources, the greater access of a more unbounded consumer population, and the desired flexibility to create new solutions from existing components over which the solution developer has little if any control. The purpose of testing is to demonstrate a required level of reliability, correctness, and effectiveness that enable prospective

 Copyright © OASIS Open 2011. All Rights Reserved.

consumers to have adequate confidence in using a service. Adequacy is defined by the consumer based on the consumer's needs and context of use. Absolute correctness and completeness cannot be proven by testing; however, for SOA, it is critical for the prospective consumer to know what testing has been performed, how it has been performed, and what were the results.

5.4.1 Traditional Software Testing as Basis for SOA Testing

SOA services are largely software artifacts and can leverage the body of experience that has evolved around software testing. IEEE-829 **[IEEE-829]** specifies the basic set of software test documents while allowing flexibility for tailored use. As such, the document structure can also provide guidance to SOA testing.

IEEE-829 covers test specification and test reporting through use of the following document types:

- *Test plan* documenting the scope (what is to be tested, both which entity and what features of the entity), the approach (how it is tested), and the needed resources (who does the testing, for how long), with details contained in the:
- *Test design specification*: features to be tested, test conditions (e.g. test cases, test procedures needed) and expected results (criteria for passing test); entrance and exit criteria
- *Test case specification*: test data used for input and expected output
- *Test procedure specification*: steps required to run the test, including any set-up preconditions
- *Test item transmittal* to identify the test items being transmitted for testing
- *Test log* to record what occurred during test, i.e. which tests run, who ran, what order, what happened
- *Test incident report* to capture any event that happened during test which requires further investigation
- *Test summary* as a management report summarizing test run and results, conclusions

In summary, IEEE-829 captures (1) what was tested, (2) how it was tested, e.g. the test procedure used, and (3) the results of the test.

5.4.1.1 Types of Testing

There are numerous aspects of testing that, in total, work to establish that an entity is (1) built as required per policies and related specifications prescribed by the entity's owner, and (2) delivers the functionality required by its intended users. This is often referred to as verification and validation.

Policies, as described in Section 1.1.1, that are related to testing may prescribe but are not limited to the business processes to be followed, the standards with which an implementation must comply, and the qualifications of and restrictions on the users. In addition to the functional requirements prescribing what an entity does, there may also be non-functional performance and/or quality metrics that state how well the entity does it. The relation of these policies to SOA testing is discussed further below.

The identification of policies is the purview of governance (section 5.1) and the assuring of compliance (including response to noncompliance) with policies is a matter for management (section 1.1).

5.4.1.2 Range of Test Conditions

Test conditions and expected responses are detailed in the test case specification. The test conditions should be designed to cover the areas for which the entity's response must be documented and may include:

- nominal conditions
- boundaries and extremes of expected conditions
- breaking point where the entity has degraded below a certain level or has otherwise ceased effective functioning
- random conditions to investigate unidentified dependencies among combinations of conditions
- errors conditions to test error handling

 Copyright © OASIS Open 2011. All Rights Reserved.

The specification of how each of these conditions should be tested for SOA resources, including the infrastructure elements of the SOA ecosystem, is beyond the scope of this document but is an area that evolves along with operational SOA experience.

5.4.1.3 Configuration Management of Test Artifacts

The test item transmittal provides an unambiguous identification of the entity being tested, thus REQUIRING that the configuration of the entity is appropriately tracked and documented. In addition, the test documents (such as those specified by IEEE-829) MUST also be under a documented and appropriately audited configuration management process, as should other resources used for testing. The description of each artifact would follow the general description model as discussed in section 4.1.1.1; in particular, it would include a version number for the artifact and reference to the documentation describing the versioning scheme from which the version number is derived.

5.4.2 Testing and the SOA Ecosystem

Testing of SOA artifacts for use in the SOA ecosystem differs from traditional software testing for several reasons. First, a highly touted benefit of SOA is to enable unanticipated consumers to make use of services for unanticipated purposes. Examples of this could include the consumer using a service for a result that was not considered the primary one by the provider, or the service may be used in combination with other services in a scenario that is different from the one considered when designing for the initial target consumer community. It is unlikely that a new consumer will push the services back to anything resembling the initial test phase to test the new use, and thus additional paradigms for testing are necessary. Some testing may depend on the availability of test resources made available as a service outside the initial test community, while some testing is likely to be done as part of limited use in the operational setting. The potential responsibilities related to such "consumer testing" is discussed further below.

Secondly, in addition to consumers who interact with a service to realize the described real world effects, the developer community is also intended to be a consumer. In the SOA vision of reuse, the developer composes new solutions using existing services, where the existing services provides access to some desired real world effects that are needed by the new solution. The new solution is a consumer of the existing services, enabling repeated interactions with the existing services playing the role of reusable components. Note, those components are used at the locations where they individually reside and are not typically duplicated for the new solution. The new solution may itself be offered as a SOA service, and a consumer of the service composition representing the new solution may be totally unaware of the component services being used. (See section 4.3.4 for further discussion on service compositions.)

Another difference from traditional testing is that the distributed, unbounded nature of the SOA ecosystem makes it unlikely to have an isolated test environment that duplicates the operational environment. A traditional testing approach often makes use of a test system that is identical to the eventual operational system but isolated for testing. After testing is successfully completed, the tested entity would be migrated to the operational environment, or the test environment may be delivered as part of the system to become operational. This is not feasible for the SOA ecosystem as a whole.

SOA services must be testable in the environment and under the conditions that can be encountered in the operational SOA ecosystem. As the ecosystem is in a state of constant change, so some level of testing is continuous through the lifetime of the service, leveraging utility services used by the ecosystem infrastructure to monitor its own health and respond to situations that could lead to degraded performance. This implies the test resources must incorporate aspects of the SOA paradigm, and a category of services may be created to specifically support and enable effective monitoring and continuous testing for resources participating in the SOA ecosystem.

While SOA within an enterprise may represent a more constrained and predictable operational environment, the composability and unanticipated use aspects are highly touted within the enterprise. The expanded perspective on testing may not be as demanding within an enterprise but fuller consideration of the ecosystem enables the enterprise to be more responsive should conditions change.

 Copyright © OASIS Open 2011. All Rights Reserved.

5.4.3 Elements of SOA Testing

IEEE-829 identifies fundamental aspects of testing, and many of these should carry over to SOA testing: in particular, the identification of what is to be tested, how it is to be tested, and by whom the testing is to be done. While IEEE-829 identifies a suggested document tree, the availability of these documents in the SOA ecosystem is discussed below.

5.4.3.1 What is to be Tested

The focus of this discussion is the SOA service. It is recognized that the infrastructure components of any SOA environment are likely to also be SOA services and, as such, falls under the same testing guidance. Other resources that contribute to a SOA environment may not be SOA services, but are expected to satisfy the intent if not the letter of guidance presented here. Specific differences for such resources are as yet largely undefined and further elaboration is beyond the scope of the SOA-RAF.

The following discussion often focuses on a singular SOA service but it is implicit that any service may be a composite of other services. As such, testing the functionality of a composite service may effectively be testing an end-to-end business process that is being provided by the composite service. If new versions are available for the component services, appropriate end-to-end testing of the composite may be required in order to verify that the composite functionality is still adequately provided. The level of required testing of an updated composite depends on policies of those providing the service, policies of those using the service, and mission criticality of those depending on the service results.

The SOA service to be tested MUST be unambiguously identified as specified by its applicable configuration management scheme. Specifying such a scheme is beyond the scope of the SOA-RAF other than to say the scheme should be documented and itself under configuration management.

5.4.3.1.1 Origin of Test Requirements

In the Service Description model (Figure 13), the aspects of a service that need to be described are:

- the service functionality and technical assumptions that underlie the functionality;
- the policies that describe conditions of use;
- the service interface that defines information exchange with the service;
- service reachability that identifies how and where message exchange is to occur; and
- metrics access for any participant to have information on how a service is performing.

Service testing must provide adequate assurance that each of these aspects is operational as defined.

The information in the service description comes from different sources. The functionality is defined through whatever process identifies needs and the community for which these needs are addressed. The process may be ad hoc as serves the prospective service owner or strictly governed, but defining the functionality is an essential first step in development. It is also an early and ongoing focus of testing to ensure the service accurately reflects the described functionality and the described functionality accurately addresses the consumer needs.

Policies define the conditions of development and conditions of use for a service and are typically specified as part of the governance process. Policies constraining service development, such as coding standards and best practices, require appropriate testing and auditing during development to ensure compliance. While the governance process identifies development policies, these are likely to originate from the technical community responsible for development activities. Policies that define conditions of use often define business practices that service owners and providers or those responsible for the SOA infrastructure want followed. These policies are initially tested during service development and are continuously monitored during the operational lifetime of the service.

The testing of the service interface and service reachability are often related but essentially reflect different motivations and needs. The service interface is specified as a joint product of the service owners and providers who define service functionality, the prospective consumer community, the service developer, and the governance process. The semantics of the information model must align with the semantics of those who consume the service in order for there to be meaningful exchange of information. The structure of the information is influenced by the consumer semantics and the requirements and constraints of the representation as interpreted by the service developer. The service process model that

 Copyright © OASIS Open 2011. All Rights Reserved.

defines actions which can be performed against a service and any temporal dependencies derive from the defined functionality and may be influenced by the development process. Any of these constraints may be identified and expressed as policy through the governance process.

Service reachability conditions are the purview of the service provider who identifies the service endpoint and the protocols recognized at the endpoint. These may be constrained by governance decisions on how endpoint addresses may be allocated and what protocols should be used.

While the considerations for defining the service interface derive from several sources, testing of the service interface is more straightforward and isolated in the testing process. At any point where the interface is modified or exposes a new resource, the message exchange should be monitored both to ensure the message reaches its intended destination and it is parsed correctly once received. Once an interface has been shown to function properly, it is unlikely to fail later unless something fundamental to the service changes.

The service interface is also tested when the service endpoint changes. Testing of the endpoint ensures message exchange can occur at the time of testing and the initial testing shows the interface is being processed properly at the new endpoint. Functioning of a service endpoint at one time does not guarantee it is functioning at another time, e.g. the server with the endpoint address may be down, making testing of service reachability a continual monitoring function through the life of the service's use of the endpoint. Also, while testing of the service endpoint is a necessary and most commonly noted part of the test regiment, it is not in itself sufficient to ensure the other aspects of testing discussed in this section.

Finally, governance is impossible without the collection of metrics against which service behavior can be assessed. Metrics are also a key indicator for consumers to decide if a service is adequate for their needs. For instance, the average response time or the recent availability can be determining factors even if there are no rules or regulations promulgated through the governance process against which these metrics are assessed. The available metrics are a combination of those expected by the consumer community and those mandated through the governance process. The total set of metrics will evolve over time with SOA experience. Testing of the services that gather and provide access to the metrics will follow testing as described in this section, but for an individual service, testing will ensure that the metrics access indicated in the service description is accurate.

The individual test requirements highlight aspects of the service that testing must consider but testing must establish more than isolated behavior. The emphasis is the holistic results of interacting with the service in the SOA environment. Recall that the execution context is the set of agreements between a consumer and a provider that define the conditions under which service interaction occurs. The agreements are expected to be predominantly the acceptance of the standard conditions as enumerated by the service provider, but it may include the identification of alternate conditions that will govern the interaction.

For example, the provider may prefer a policy where it can sell the contact information of its consumers but will honor the request of a consumer to keep such information private. The identification of the alternate privacy policy is part of the execution context, and it is the application of and compliance with this policy that operational monitoring will attempt to measure. The collection of metrics showing this condition is indeed met when chosen is considered part of the ongoing testing of the service.

Other variations in the execution context also require monitoring to ensure that different combinations of conditions perform together as desired. For example, if a new privacy policy takes additional resources to apply, this may affect quality of service and propagate other effects. These could not be tested during the original testing if the alternate policy did not exist at that time.

5.4.3.1.2 Testing Against Non-Functional Requirements

Testing against non-functional requirements constitutes testing of business usability of the service. In a marketplace of services, non-functional characteristics may be the primary differentiator between services that produce essentially the same real world effects.

As noted in the previous section, non-functional characteristics are often associated with policies or other terms of use and may be collected in service level contracts offered by the service providers. Non-functional requirements may also reflect the network and hardware infrastructure that support communication with the service, and changes may impact quality of service. The service consumer and

Copyright © OASIS Open 2011. All Rights Reserved.

even the service provider may not be aware of all such infrastructure changes but the changes may manifest in shared states that impact the usability of the service.

In general, a change in the non-functional requirements results in a change to the execution context, but as with any collection of information that constitutes a description, the execution context is unable to explicitly capture all non-functional requirements that may apply. A change in non-functional requirements, whether explicitly part of the execution context or an implicit contributor, may require retesting of the service even if its functionality and the implementation of the functionality has not changed. Depending on the circumstances, retesting may require a formal recertifying of end-to-end behavior or more likely will be part of the continuous monitoring that applies throughout the service lifetime.

5.4.3.1.3 Testing Content and the Interests of Consumers

As noted in section 5.4.1.1, testing may involve verification of conformance with respect to policies and technical specifications and validation with respect to sufficiency of functionality to meet some prescribed use. It may also include demonstration of performance and quality aspects. For some of these items, such as demonstrating the business processes followed in developing the service or the use of standards in implementing the service, the testing or relevant auditing is done internal to the service development process and follows traditional software testing and quality assurance. If it is believed of value to potential consumers, information about such testing could be included in the service description. However, it is not required that all test or compliance artifacts be available to consumers, as many of the details tested may be part of the opacity of the service implementation.

Some aspects of the service being tested will reflect directly on the real world effects realized through interaction with the service. In these cases, it is more likely that testing results will be directly relevant to potential consumers. For example, if the service was designed to correspond to certain elements of a business process or that a certain workflow is followed, testing should verify that the real world effects reflect that the business process or workflow were satisfactorily captured.

The testing may also need to demonstrate that specified conditions of use are satisfied. For example, policies may be asserted that require certain qualifications of or impose restrictions on the consumers who may interact with the service. The service testing must demonstrate that the service independently enforces the policies or it provides the required information exchanges with the SOA ecosystem so other resources can ensure the specified conditions.

The completeness of the testing, both in terms of the features tested and the range of parameters for which response is tested, depends on the context of expected use: the more critical the use, the more complete the testing. There are always limits on the resources available for testing, if nothing else than the service must be available for use in a finite amount of time.

This again emphasizes the need for adequate documentation to be available. If the original testing is very thorough, it may be adequate for less demanding uses in the future. If the original testing was more constrained, then well-documented test results establish the foundation on which further testing can be defined and executed.

5.4.3.2 How Testing is to be Done

Testing should follow well-defined methodologies and, if possible, should reuse test artifacts that have proven generally useful for past testing. For example, IEEE-829 notes that test cases are separated from test designs to allow for use in more than one design and to allow for reuse in other situations. In the SOA ecosystem, description of such artifacts, as with description of a service, enables awareness of the item and describes how the artifact may be accessed or used.

As with traditional testing, the specific test procedures and test case inputs are important so the tests are unambiguously defined and entities can be retested in the future. Automated testing and regression testing may be more important in the SOA ecosystem in order to re-verify a service is still acceptable when incorporated in a new use. For example, if a new use requires the services to deal with input parameters outside the range of initial testing, the tests could be rerun with the new parameters. If the testing resources are available to consumers within the SOA ecosystem, the testing as designed by test professionals could be consumed through a service accessed by consumers, and their results could augment those already in place. This is discussed further in the next section.

 Copyright © OASIS Open 2011. All Rights Reserved.

5.4.3.3 Who Performs the Testing

As with any software, the first line of testing is unit testing done by software developers. It is likely that initial testing will be done by those developing the software but may also be done independently by other developers. For SOA development, unit testing is likely confined to a development sandbox isolated from the SOA ecosystem.

SOA testing will differ from traditional software testing in that testing beyond the development sandbox must incorporate aspects of the SOA ecosystem, and those doing the testing must be familiar with both the characteristics and responses of the ecosystem and the tools, especially those available as services, to facilitate and standardize testing. Test professionals will know what level of assurance must be established as the exposure of the service to the ecosystem and ecosystem to the service increases towards operational status. These test professionals may be internal resources to an organization or may evolve as a separate discipline provided through external contracting.

As noted above, it is unlikely that a complete duplicate of the SOA ecosystem will be available for isolated testing, and thus use of ecosystem resources will manifest as a transition process rather than a step change from a test environment to an operational one. This is especially true for new composite services that incorporate existing operational services to achieve the new functionality. The test professionals will need to understand the available resources and the ramifications of this transition.

As with current software development, a stage beyond work by test professionals will make use of a select group of typical users, commonly referred to as beta testers, to report on service response during typical intended use. This establishes fitness by the consumers, providing final validation of previously verified processes, requirements, and final implementation.

In traditional software development, beta testing is the end of testing for a given version of the software. However, although the initial test phase can establish an appropriate level of confidence consistent with the designed use for the initial target consumer community, the operational service will exist in an evolving ecosystem, and later conditions of use may differ from those thought to be sufficient during the initial testing. Thus, operational monitoring becomes an extension of testing through the service lifetime. This continuous testing will attempt to ensure that a service does not consume an inordinate amount of ecosystem resources or display other behavior that degrades the ecosystem, but it will not undercover functional errors that may surface over time.

As with any software, it is the responsibility of the consumers to consider the reasonableness of solutions in order to spot errors in either the software or the way the software is being used. This is especially important for consumers with unanticipated uses that may go beyond the original test conditions. It is unlikely the consumers will initiate a new round of formal testing unless the new use requires a significantly higher level of confidence in the service. Rather the consumer becomes a new extension to the testing regiment. Obvious testing would include a sanity check of results during the new use. However, if the details of legacy testing are associated with the service through the service description and if testing resources are available through automated testing services, then the new consumers can rerun and extend previous testing to include the extended test conditions. If the test results are acceptable, these can be added to the documentation of previous results and become the extended basis for future decisions by prospective consumers on the appropriateness of the service. If the results are not acceptable or in some way questionable, the responsible party for the service or testing professionals can be brought in to decide if remedial action is necessary.

5.4.3.4 How Testing Results are Reported

For any SOA service, an accurate reporting of the testing a service has undergone and the results of the testing is vital to consumers deciding whether a service is appropriate for intended use. Appropriateness may be defined by a consumer organization and require specific test regiments culminating in a certification; appropriateness could be established by accepting testing and certifications that have been conferred by others.

The testing and certification information should be identified in the service description. Referring to the general description model of Figure 11, tests conducted by or under a request from the service owner (see ownership in section 3.1.3) would be captured under Annotations from Owners. Testing done by others, such as consumers with unanticipated uses, could be associated through Annotations from 3rd

 Copyright © OASIS Open 2011. All Rights Reserved.

Parties. The annotations should clearly indicate what was tested, how the testing was done, who did the testing, and the testing results. The clear description of each of these artifacts and of standardized testing protocols for various levels of sophistication and completeness of testing would enable a common understanding and comparison of test coverage. It will also make it more straightforward to conduct and report on future testing, facilitating the maintenance of the service description.

Consumer testing and the reporting of results raises additional issues. While stating who did the testing is mandatory, there may be formal requirements for authentication of the tester to ensure traceability of the testing claims. In some circumstances, persons or organizations would not be allowed to state testing claims unless the tester was an approved entity. In other cases, ensuring the tester had a valid email may be sufficient. In either case, it would be at the discretion of the potential consumer to decide what level of authentication was acceptable and which testers are considered authoritative in the context of their anticipated use.

Finally, in a world of openly shared information, we would see an ever-expanding set of testing information as new uses and new consumers interact with a service. In reality, these new uses may represent proprietary processes or classified use that should only be available to authorized parties. Testing information, as with other elements of description, may require special access controls to ensure appropriate access and use.

5.4.4 Testing SOA Services

Testing of SOA services should be consistent with the SOA paradigm. In particular, testing resources and artifacts should be visible in support of service interaction between providers and consumers, where here the interaction is between the testing resource and the tester. In addition, the idea of opacity of the implementation should limit the details that need to be available for effective use of the test resources. Testing that requires knowledge of the internal structure of the service or its underlying capability should be performed as part of unit testing in the development sandbox, and should represent a minimum level of confidence before the service begins its transition to further testing and eventual operation in the SOA ecosystem.

5.4.4.1 Progression of SOA Testing

Software testing is a gradual exercise going from micro inspection to testing macro effects. The first step in testing is likely the traditional code reviews. SOA considerations would account for the distributed nature of SOA, including issues of distributed security and best practices to ensure secure resources. It would also set the groundwork for opacity of implementation, hiding programming details and simplifying the use of the service.

Code review is likely followed by unit testing in a development sandbox isolated from the operational environment. The unit testing is done with full knowledge of the service internal structure and knowledge of resources representing underlying capabilities. It tests the interface to ensure exchanged messages are as specified in the service description and the messages can be parsed and interpreted as intended. Unit testing also verifies intended functionality and that the software has dealt correctly with internal dependencies, such as structure of a file system or access to other dedicated resources.

Some aspects of unit testing require external dependencies be satisfied, and this is often done using mock objects to substitute for the external resources. In particular, it will likely be necessary to include mocks of existing operational services, both those provided as part of the SOA infrastructure and services from other providers.

Service Mock

> A service mock is an entity that mimics some aspect of the performance of an operational service without committing to the real world effects that the operational service would produce.

Mocks are discussed in detail in sections 5.4.4.3 and 5.4.4.4.

After unit testing has demonstrated an adequate level of confidence in the service, the testing must transition from the tightly controlled environment of the development sandbox to an environment that more clearly resembles the operational SOA ecosystem or, at a minimum, the intended enterprise. While sandbox testing will use simple mocks of some aspects of the SOA environment, such as an interface to

 Copyright © OASIS Open 2011. All Rights Reserved.

a security service without the security service functionality, the dynamic nature of SOA makes a full simulation infeasible to create or maintain. This is especially true when a new composite service makes use of operational services provided by others. Thus, at some point before testing is complete, the service will need to demonstrate its functionality by using resources and dealing with conditions that only exist in the full ecosystem or the intended enterprise. Some of these resources may still provide test interfaces -- more on this below -- but the interfaces will be accessible using the SOA environment and not just implemented for the sandbox.

At this stage, the opacity of the service becomes important as the details of interacting with the service now rely on correct use of the service interface and not knowledge of the service internals. The workings of the service will only be observable through the real world effects realized through service interactions and external indications that conditions of use, such as user authentication, are satisfied. Monitoring the behavior of the service will depend on service interfaces that expose internal monitoring or provide required information to the SOA infrastructure monitoring function. The monitoring required to test a new service is likely to have significant overlap with the monitoring the SOA infrastructure includes to monitor its own health and to identify and isolate behavior outside of acceptable bounds. This is exactly what is needed as part of service testing, and it is reasonable to assume that the ecosystem transition includes use of operational monitoring rather than solely dedicated monitoring for each service being tested.

Use of SOA monitoring resources during the explicit testing phase sets the stage for monitoring and a level of continual testing throughout the service lifetime.

5.4.4.2 Testing Traditional Dependencies vs. Service Interactions

A SOA service is not required to make use of other operational services beyond what may be required for monitoring by the ecosystem infrastructure. The service can implement hardcoded dependencies which have been tested in the development sandbox through the use of dedicated mocks. While coordination may be required with real data sources during integration testing, the dependencies can be constrained to things that can be tested in a more traditional manner. Policies can also be set to restrict access to pre-approved users, and thus the question of unanticipated users and unanticipated uses can be eliminated. Operational readiness can be defined in terms of what can be proven in isolated testing. While all this may provide more confidence in the service for its designed purpose, such a service will not fully participate in the benefits or challenges of the ecosystem. This is akin to filling a swimming pool with sea water and having someone in the pool say they are swimming in the ocean.

In considering the testing needed for a fully participating service, consider the example of a new composite service that combines the real world effects and complies with the conditions of use of five existing operational services. The developer of the composite service does not own any of the component services and has limited, if any, ability to get the distributed owners to do any customization. The developer also is limited by the principle of opacity to information comprising the service description, and does not know internal details of the component services. The developer of the composite service must use the component services as they exist as part of the SOA environment, including what is provided to support testing by new users. This introduces requirements for what is needed in the way of service mocks.

5.4.4.3 Use of Service Mocks

Service mocks enables the tested service to respond to specific features of an operational service that is being used as a component. It allows service testing to proceed without needing access to or with only limited engagement with the component service. Mocks can also mimic difficult to create situations for which it is desired to test the new service response. For composite services using multiple component services, mocks may be used in combination to function for any number of the components. Note, when using service mocks, it is important to remember that it is not the component service that is being tested (although anomalous behavior may be uncovered during testing) but the use of the component in the new composite.

Individual service mocks can emphasize different features of the component service they represent but any given mock does not have to mimic all features. For example, a mock of the service interface can echo a sent message and demonstrate the message is reaching its intended destination. A mock could go further and parse the sent message to demonstrate the message not only reached its destination but

 Copyright © OASIS Open 2011. All Rights Reserved.

was understood. As a final step, the mock could report back what actions would have been taken by the component service and what real world effects would result. If the response mimicked the operational response, functional testing could proceed as if the real world effect actually occurred.

There are numerous ways to provide mock functionality. The service mock could be a simulation of the operational service and return simulated results in a realistic response message or event notification. It is also possible for the operational service to act as its own mock and simply not execute the commit stage of its functionality. The service mock could use a combination of simulation and service action without commit to generate a report of what would have occurred during the defined interaction with the operational service.

As the service proceeds through testing, mocks should be systematically replaced by the component resources accessed through their operational interfaces. Before beta testing begins, end-to-end testing, i.e. proceeding from the beginning of the service interaction to the resulting real world results, should be accomplished using component resources via their operational interfaces.

5.4.4.4 Providers of Service Mocks

In traditional testing, it is often the test professionals who design and develop the mocks, but in the distributed world of SOA, this may not be efficient or desirable.

In the development sandbox, it is likely the new service developer or test professionals working with the developer will create mocks adequate for unit testing. Given that most of this testing is to verify the new service is performing as designed, it is not necessary to have high fidelity models of other resources being accessed. In addition, given opacity of SOA implementation, the developer of the new service may not have sufficient detailed knowledge of a component service to build a detailed mock of the component service functionality. Sharing existing mocks at this stage may be possible but the mocks would need to be implemented in the sandbox, and for simple models it is likely easier to build the mock from scratch.

As testing begins its transition to the wider SOA environment, mocks may be available as services. For existing resources, it is possible that an Open Source model could evolve where service mocks of available functions can be catalogued and used during initial interaction of the tested service and the operational environment. Widely used functions may have numerous service mocks, some mimicking detailed conditions within the SOA infrastructure. However, the Open Source model is less likely to be sufficient for specialty services that are not widely used by a large consumer community.

The service developer is probably best qualified for also developing more detailed service mocks or for mock modes of operational services. This implies that in addition to their operational interfaces, services will routinely provide test interfaces to enable service mocks to be used as services. As noted above, a new service developer wanting to build a mock of component services is limited to the description provided by the component service developer or owner. The description typically will detail real world effects and conditions of use but will not provide implementation details, some of which may be proprietary. Just as important in the SOA ecosystem, if it becomes standard protocol for developers to create service mocks of their own services, a new service developer is only responsible for building his own mocks and can expect other mocks to be available from other developers. This reduces duplication of effort where multiple developers would be trying to build the same mocks from the same insufficient information. Finally, a service developer is probably best qualified to know when and how a service mock should be updated to reflect modified functionality or message exchange.

It is also possible that testing organizations will evolve to provide high-fidelity test harnesses for new services. The harnesses would allow new services to plug into a test environment and would facilitate accessing mocks of component services. However, it will remain a constant challenge for such organizations to capture evolving uses and characteristics of service interactions in the real SOA environment and maintain the fidelity and accuracy of the test systems.

5.4.4.5 Fundamental Questions for SOA Testing

In order for the transition to the SOA operational environment to proceed, it is necessary to answer two fundamental questions:

- Who provides what testing resources for the SOA operational environment, e.g. mocks of interfaces, mocks of functionality, monitoring tools?

 Copyright © OASIS Open 2011. All Rights Reserved.

- What testing needs to be accomplished before operational environment resources can be accessed for further testing?

The discussion in section 5.4.4.4 notes various levels of sophistication of service mocks and different communities are likely to be responsible for different levels. Section 5.4.4.4 advocates a significant role for service developers, but there needs to be community consensus that such mocks are needed and that service developers will agree to fulfilling this role. There is also a need for consensus as to what tools should be available as services from the SOA infrastructure.

As for use of the service mocks and SOA environment monitoring services, practical experience is needed upon which guidelines can be established for when a new service has been adequately tested to proceed with a greater level of exposure with the SOA environment. Malfunctioning services could cause serious problems if they cannot be identified and isolated. On the other hand, without adequate testing under SOA operational conditions, it is unlikely that problems can be uncovered and corrected before they reach an operational stage.

As noted in section 5.4.4.2~~These dependencies may take the form of other services being established; possibly even services that are not exposed by the service's own interface.~~

, some of these questions can be avoided by restricting services to more traditional use scenarios. However, such restriction will limit the effectiveness of SOA use and the result will resemble the constraints of traditional integration activities we are trying to move beyond.

5.4.5 Architectural Implications for SOA Testing

The discussion of SOA Testing indicates numerous architectural implications on the SOA ecosystem:

- The distributed, boundary-less nature of the SOA ecosystem makes it infeasible to create and maintain a single mock of the entire ecosystem to support testing activities.
- A standard suite of monitoring services needs to be defined, developed, and maintained. This should be done in a manner consistent with the evolving nature of the ecosystem.
- Services should provide interfaces that support access in a test mode.
- Testing resources must be described and their descriptions must be catalogued in a manner that enables their discovery and access.
- Guidelines for testing and ecosystem access need to be established and the ecosystem must be able to enforce those guidelines asserted as policies.
- Services should be available to support automated testing and regression testing.
- Services should be available to facilitate updating service description by anyone who has performed testing of a service.

 Copyright © OASIS Open 2011. All Rights Reserved.

6 Conformance

This Reference Architecture Framework is an abstract ~~architecture~~architectural description of Service Oriented Architecture, which means that it is especially difficult to construct ~~automated~~ tests for conformance to the architecture. In addition, conformance to an architectural specification does not, by itself, guarantee any form of interoperability between multiple implementations.

However, ~~in order to be~~ it *is* possible to decide whether or not a given architecture is conformant to an architectural description such as this one. In discussions of conformance we use the term **target architecture**~~, it should be possible~~ to identify ~~in a~~ the (typically concrete~~ implementation the key~~) architecture that may be viewable as conforming to the abstract principles outlined in this document.

Target Architecture

> A target architecture is an architectural description of a system that is intended to be viewed as conforming to the SOA-RAF.

While we cannot guarantee interoperability between target architectures (or more specifically between applications and systems residing within the ecosystems of those target architectures), interoperability between target architectures is promoted by conformance to this Reference Architecture Framework as it reduces the semantic impedance mismatch between the different ecosystems.

The primary measure of conformance is whether given concepts ~~and components of this architecture, albeit~~as described in document have corresponding concepts in ~~abstracted form~~the target architecture. Such a correspondence MUST honor the relationships identified within this document for the target architecture to be considered conforming.

For example, in Section 3.1.3.1 we identify resource as a key concept. A resource is associated with an owner and a number of identifiers. For a target architecture to conform to the SOA-RAF, it must be possible to find corresponding concepts of resource, identifier and owner within the target architecture: say *entity*, *token* and *user* . Furthermore, the relationships between *entity*, *token* and *user* MUST mirror the relationships between resource, identifier and owner appropriately.

Clearly, such correspondence is simpler if the terminology within the target architecture is identical to that in the SOA-RAF. But so long as the ‘graph’ of concepts and relationships is consistent, that is all that is required for the target architecture to conform to this Reference Architecture Framework.

[EDITOR’S NOTE: The conformance section is not complete]

 Copyright © OASIS Open 2011. All Rights Reserved.

A. Acknowledgements

The following individuals have participated in the work of the technical committee responsible for creation of this specification and are gratefully acknowledged:

Participants:

Chris Bashioum, MITRE Corporation
Rex Brooks, ~~individual member~~Individual Member
Peter Brown, ~~Pensive.eu~~Individual Member
Scott Came, Search Group Inc.
Joseph Chiusano, Booz Allen Hamilton
Robert Ellinger, ~~Northrup~~Northrop Grumman Corporation
David Ellis, Sandia National Laboratories
Jeff A. Estefan, Jet Propulsion Laboratory
Don Flinn, Individual Member
Anil John, Johns Hopkins University
Ken Laskey, MITRE Corporation
Boris Lublinsky, Nokia Corporation
Francis G. McCabe, Individual Member
Christopher McDaniels, USSTRATCOM
Tom Merkle, Lockheed ~~martin~~Martin Corporation
Jyoti Namjoshi, Patni Computer Systems Ltd.
Duane Nickull, Adobe Inc.
James Odell, Associate
Michael Poulin, Fidelity Investments
Michael Stiefel, ~~Reliable Software~~Associate
Danny Thornton, ~~Individual~~Northrop Grumman
Timothy Vibbert, Lockheed Martin Corporation
Robert Vitello, New York Dept. of Labor

The committee would particularly like to underline the significant contributions made by Rex Brooks, Jeff Estefan, Ken Laskey, Boris Lublinsky, Frank McCabe, Michael Poulin and Danny Thornton

B. Index of Defined Terms

The first page number refers to the first use of the term. The second, where necessary, refers to the page where the term is formally defined.

Copyright © OASIS Open 2011. All Rights Reserved.

Ownership Boundary
Participant
Peer
Permission
Policy
Policy Conflict
Policy Conflict Resolution
Policy Constraint
Policy Decision
Policy Enforcement
Policy Framework
Policy Object
Policy Ontology
Policy Owner
Policy Subject
Presence
Private State
Protocol
Public Semantics
Qualification
Real World Effect
Regulation
Resource
Responsibility
Right
Risk
Role
Rule
Security
Semantic Engagement
Service Action
Service Consumer
Service Level Real World Effect
Service Mediator
Service Provider
Shared State
Skill
Social Structure
Stakeholder
State
System
System Stakeholder

 Copyright © OASIS Open 2011. All Rights Reserved.

Trust
View
Viewpoint

 Copyright © OASIS Open 2011. All Rights Reserved.

C. The Unified Modeling Language, UML

Figure 44 illustrates an annotated example of a UML class diagram that is used to represent a visual model depiction of the Resources Model in the *Participation in a SOA Ecosystem* view.

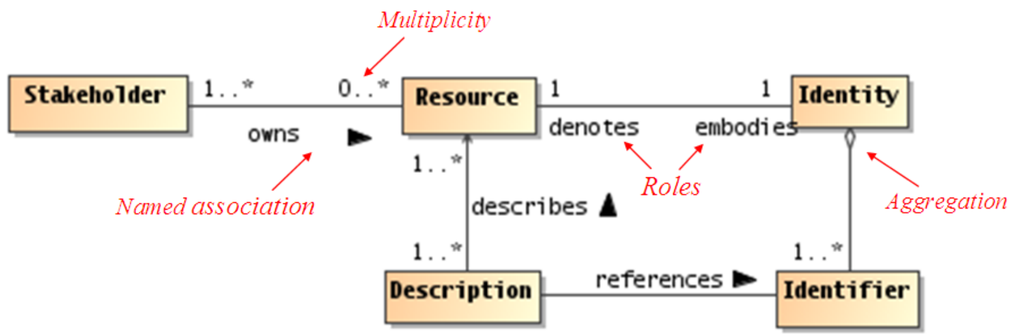

Figure 44 - Example UML class diagram—Resources

Lines connecting boxes (classifiers) represent associations between things. An association has two roles (one in each direction). A role can have cardinality, for example, one or more ("1..*") stakeholders own zero or more ("0..*) resources. The role from classifier A to B is labeled closest to B, and vice versa, for example, the role between resource to Identity can be read as resource embodies Identity, and Identity denotes a resource.

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

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

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

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

 Copyright © OASIS Open 2011. All Rights Reserved.

~~B.~~D. Critical Factors Analysis

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. CFA is particularly suitable for capturing quality attributes of a project, often referred to as “non-functional” or “other-than-functional” requirements~~ of a project~~: for example, security, scalability, wide-spread adoption, and so on. As such, CFA complements rather than attempts to replace other requirements capture techniques.

~~B.1~~D.1 Goals

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.

~~B.2~~D.2 Critical Success Factors

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.

~~B.3~~D.3 Requirements

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.

~~B.4~~D.4 CFA Diagrams

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.

The legend:

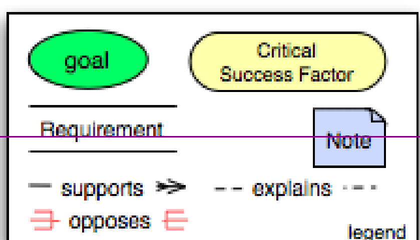

Copyright © OASIS Open 2011. All Rights Reserved.

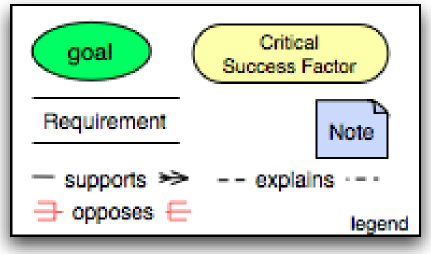

illustrates the key elements of the graphical notation. Goals are written in round ovals, critical success factors are written in round-ended rectangles and requirements are written using open-ended rectangles. The arrows show whether a CSF/goal/requirement is supported by another element or opposed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

 Copyright © OASIS Open 2011. All Rights Reserved.

 Copyright © OASIS Open 2011. All Rights Reserved.

E. Relationship to other SOA Open Standards

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

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

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

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

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

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

 Copyright © OASIS Open 2011. All Rights Reserved.

6088 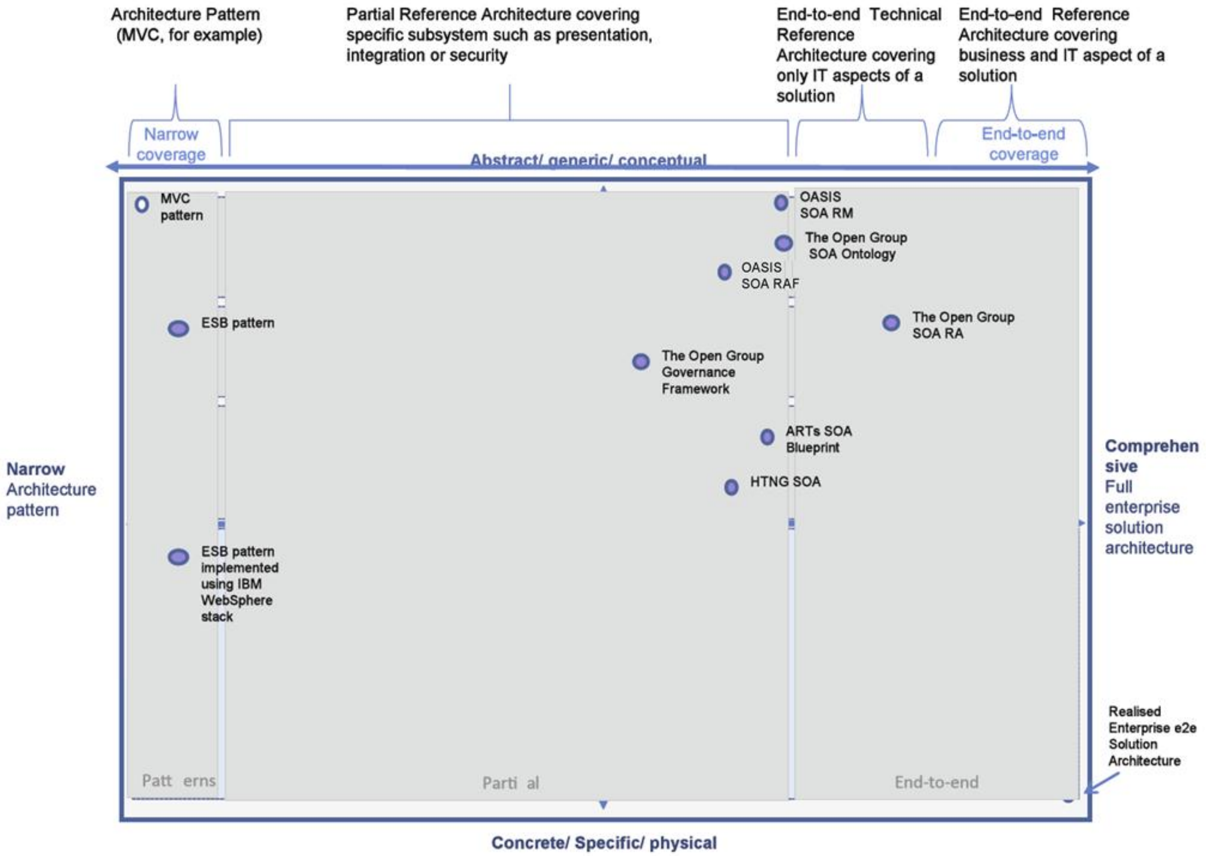

6089 *Figure 45 - SOA Reference Architecture Positioning (from "Navigating the SOA Open Standards Landscape Around*
6090 *Architecture, © OASIS, OMG, The Open Group)*